CyberPeace

CAPABILITIES WITHOUT SAFEGUARDS:

REFRAMING INDIA'S AI-IN-HEALTHCARE BOOM AS A DATA-RIGHTS PROBLEM

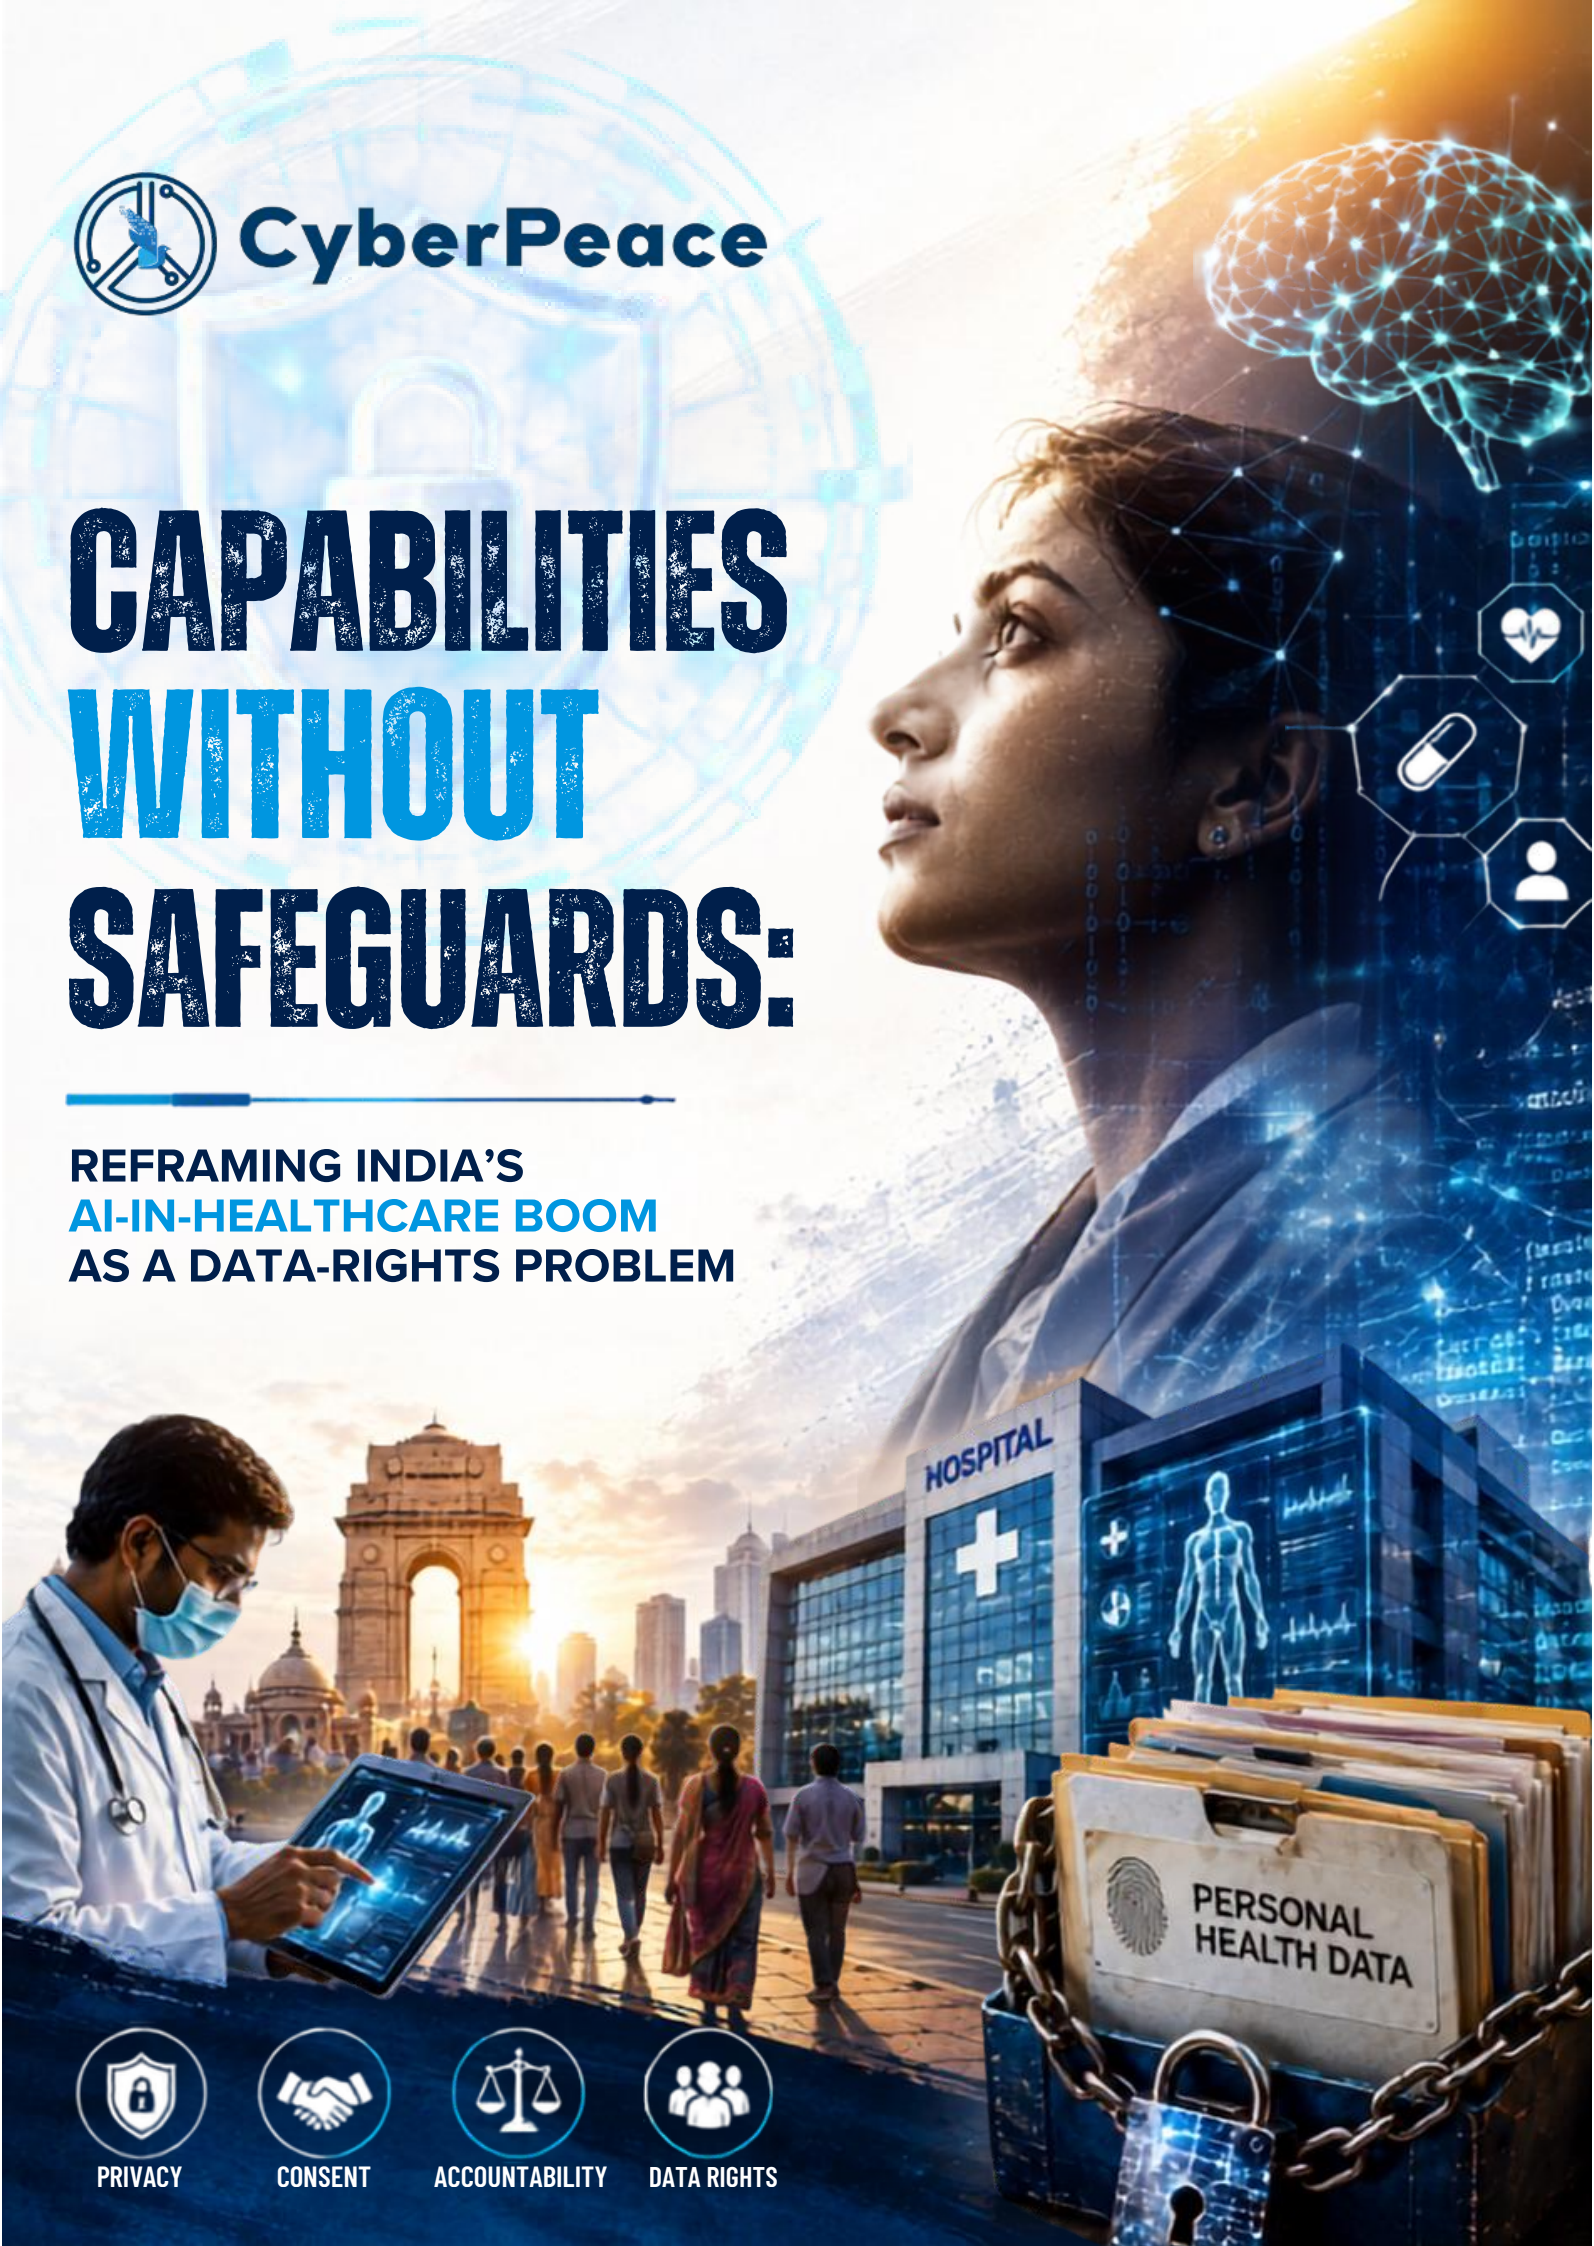

Executive Summary

India is assembling, within a few years, much of the technical scaffolding for AI-driven healthcare that other health systems took a decade to build: point-of-care diagnostic tools such as Qure.ai's chest-X-ray reader, now used across more than a thousand facilities;[1] ambient AI scribes entering hospital wards through global platforms like Epic and homegrown startups;[2] and a national medical-imaging data pipeline, MIDAS, already spanning thirteen hospitals and four disease areas.[3] The legal and technical safeguards meant to govern the patient data feeding these systems are maturing far more slowly: India's principal data-protection statute has no dedicated category for health data;[4] its enabling rules are being phased in over several years;[5] and its only health-specific ethics instrument is a voluntary guideline rather than an enforceable law.[6] This paper argues that policy discussion of AI in Indian healthcare should shift from a capabilities register to a needs-and-rights register. It closes with recommendations for policymakers, hospitals, AI developers, and civil society.

---

[1]How AI and Ayushman Bharat Digital Mission Are Powering an AI-Driven Healthcare Revolution in India, ORGANISER (Feb. 25, 2026), https://organiser.org/2026/02/25/341713/bharat/how-ai-and-ayushman-bharat-digital-mission-are-powering-an-ai-driven-healthcare-revolution-in-india/.

[2]Epic's AI Scribe Goes Live as EHR Giant Touts Strong Adoption of Built-In AI Features, FIERCE HEALTHCARE (Feb. 5, 2026), https://www.fiercehealthcare.com/ai-and-machine-learning/epic-rolls-out-ai-charting-and-more-built-automation-clinicians-and.

[3]ARTPARK, MIDAS — Building India's National AI Data Infrastructure for Health, https://www.midas.iisc.ac.in/fe/about (last visited July 27, 2026).

[4]Health Data & the DPDP Act — Navigating India's New Privacy Regime, MONDAQ (Dec. 18, 2025), https://www.mondaq.com/pdf/1722484.pdf.

[5]DLA Piper, Data Protection Laws of the World: India, https://www.dlapiperdataprotection.com/?t=law&c=IN (last visited July 27, 2026).

[6]Indian Council of Medical Research, Ethical Guidelines for Application of Artificial Intelligence in Biomedical Research and Healthcare (2023), https://www.icmr.gov.in/ethical-guidelines-for-application-of-artificial-intelligence-in-biomedical-research-and-healthcare.

Introduction

Artificial intelligence is entering Indian healthcare faster than almost any other public infrastructure in the country's recent history. Diagnostic algorithms read chest X-rays for tuberculosis in minutes rather than the weeks a radiologist shortage might otherwise impose;[7] ambient listening software drafts clinical notes directly from the patient encounter;[8] and a nationally coordinated hub-and-spoke network is assembling India's first large, standardised, AI-ready medical-imaging datasets.[9] Government and industry framing of this shift understandably emphasises the capability of the paramount efficiency it can reach.

A quieter, parallel conversation about what happens to the recording of a patient's voice or the scan of their body once it leaves the examination room has struggled to keep pace. India's own policy analysts have begun flagging this asymmetry at the national level, as most AI investment has concentrated on supply-side infrastructure like compute and foundations without a dedicated strategy for the demand side.[10]

This paper takes that asymmetry as its starting point. It is not an anti-AI paper; genuine products like ambient scribes that let a doctor look at a patient instead of a screen or diagnostic tools that catch tuberculosis in an underserved district are a boon[11]. The argument is narrower, i.e., the safeguards around this technology are not yet proportionate to the scale of data pooling now underway, and closing that gap deserves first-order attention by applying privacy-by-design principles.

The Landscape: What's Actually Being Deployed in India

Four categories dominate current deployment:

[7] NITI Aayog Frontier Tech, AI-Driven Scans, Early Action: Qure.ai's Answer to India's TB Challenge (June 19, 2025), https://frontiertech.niti.gov.in/story/ai-driven-tuberculosis-diagnosis-a-transformative-solution-for-indias-healthcare-challenges.

[8] Doctors Increasingly See AI Scribes in a Positive Light. But Hiccups Persist., KFF HEALTH NEWS (Jan. 27, 2026), https://kffhealthnews.org/news/article/ambient-ai-scribes-doctor-appointments-note-taking-ehr-epic/.

[9] MIDAS: A Technology-Enabled Hub-and-Spoke System for the Collection and Dissemination of High-Quality Medical Datasets in India, BMC MED. INFORMATICS & DECISION MAKING (2025), https://link.springer.com/article/10.1186/s12911-025-03092-7.

[10] Observer Research Foundation, Protocol in Every Pocket: AI at India's Healthcare Frontline (2026), https://www.orfonline.org/expert-speak/protocol-in-every-pocket-ai-at-india-s-healthcare-frontline.

[11] NITI Aayog Frontier Tech, *supra* note 7.

1. **Point-of-care diagnostic AI:** Qure.ai's qXR product interprets chest X-rays for tuberculosis and related conditions and, by industry accounts, is used at more than a thousand facilities across India, alongside comparable tools from Niramai, 5C Network, DeepTek, and SigTuple.[12] These tools increasingly qualify as regulated medical devices: the Central Drugs Standard Control Organisation classifies diagnostic AI as "software as a medical device" under the Medical Device Rules, with explicit provisions added in a 2023 update.[13]

2. **Ambient and clinical documentation AI:** Ambient scribes that listen to a clinical encounter and auto-draft notes have moved from pilot to mainstream in months rather than years. Epic, which by one industry estimate holds roughly 42% of the U.S. hospital electronic health record market, launched its own built-in "AI Charting" feature in early 2026,[14] within India, developers such as the Chennai-based startup Ekko are building multilingual equivalents tuned to outpatient departments where a single consultant may see more than sixty patients in a session.[15]

3. **EHR platforms with embedded AI:** Global EHR vendors are increasingly bundling generative AI copilots directly into the record system rather than offering them as separate add-ons, collapsing the line between documentation tool and data pipeline.[16]

4. **National data-pooling infrastructure:** The clearest instance is MIDAS (Medical Imaging Datasets for India), a joint initiative of ICMR, IISc, and ARTPARK organised as a hub-and-spoke network: the ICMR-IISc Centre for Health Data serves as the root node, with designated hospitals and medical colleges as thematic hubs drawing data from further spoke about institutions.[17] MIDAS already spans thirteen hospitals across four disease verticals and more than a terabyte of multimodal data, and its builders describe it as an "ethics-compliant", open-source digital public good.[18]

[12] 5C Network, Radiology AI in India: 2026 Hospital Guide, https://www.5cnetwork.com/resources/radiology-ai-india-guide (last updated May 20, 2026).

[13] 5C Network, *supra* note 12.

[14] Fierce Healthcare, *supra* note 2.

[15] OC Academy, Ambient AI in Healthcare: Empowering Indian Doctors (May 27, 2026), https://www.ocacademy.in/blogs/ambient-ai-in-healthcare-india/.

[16] AI Listens to Patients, Physicians Face New Questions, MEDSCAPE (June 22, 2026), https://www.medscape.com/viewarticle/ai-listens-patients-physicians-face-new-questions-2026a1000kxc.

[17] BMC Med. Informatics & Decision Making, *supra* note 9; ARTPARK, *supra* note 3.

Tool type	Patient data collected	Who controls it
Diagnostic AI (Qure.ai, Niramai, 5C, DeepTek)	Radiology images, scan metadata	Private AI vendor + hospital
Ambient scribes (Epic AI Charting, Ekko)	Recorded patient–doctor conversation, transcript	EHR vendor / scribe startup + hospital
EHR-embedded AI copilots	Full longitudinal medical record	EHR vendor, hospital IT
MIDAS national pooling	De-identified imaging + demographic data	ICMR-IISc Centre for Health Data, ARTPARK, hub hospitals

[18] ARTPARK, *supra* note 3; IndiaAI, MIDAS: A Platform to Create India-Centric Datasets for AI-Powered Healthcare (last visited July 27, 2026).

THE DATA TRAIL:
WHERE PATIENT DATA ACTUALLY GOES

Tracing a single data point, be it an MRI scan or a recorded consultation, through this ecosystem surfaces four recurring gaps.

4.1 Capture and consent

India's flagship digital health platform, the Ayushman Bharat Digital Mission (ABDM), states that health data will not be shared without a patient's explicit consent. '19 In practice, investigative reporting has found that the consent form forms sign authors the National Digital Health Mission and its "ecosystem partners" to use and disclose data without describing the specific purpose of that processing closer to blanket authorisation than to purpose-limited, granular consent.'20 A 2024 survey of hospital outpatients found that while most welcomed digital health tools, nearly all wanted consent sought at every step of data use rather than once at enartment.'21 Separately, reporters found that a public-facing insurance scheme dashboard exposed treatment-level details of individual beneficiaries.'22

4.2 Storage and processing

Health AI tools typically run on vendor or cloud infrastructure rather than hospital-owned servers, which is part of why the DPP Act restricts cross-border data transfer and imposes heightened obligations on entities designated "Significant Data Fiduciares".'3

4.3 Secondary use and training

MIDAS and comparable pooling efforts describe the automatic de-identification and harmonisation of data before it is made available for AI model training, and per its builders, for benchmarking diagnostic tools, including foreign-developed ones, against Indian population baselines.'24 That is a legitimate scientific goal, but one whose safety depends entirely on the strength of the de-identification step, examined in Part 5.

4.4 Multi-institutional aggregation risk

A hub-and-spoke architecture spanning thirteen or more hospitals is, from a security standpoint, a larger and more attractive target than any single member site. A compromised spoke, vendor contract, or credential can expose records aggregated from institutions that individually might have been reasonably well-defended.

Tool type	Patient data collected	Who controls it
Diagnostic AI (Qure, ai, Niramai, SC, DeepTek)	Radiology images, scan metadata	Private AI vendor + hospital
Ambient scribes (Epic AI Charting, Ekko)	Recorded patient-doctor conversation, transcript	EHR vendor / scribe startup + hospital
EHR-embedded (Epic AI copilots)	Full longitudinal medical record	EHR vendor, hospital IT
MIDAS national pooling	De-identified imaging + demographic data	ICMR-IISC Centre for Health Data, ARTPARK, hub hospitals

1999Privacy Issues Still Persistent Despite ADBD's Promised Safeguards, MEDIANAMA https://www.medicinenama.com/2012/12/21/abm privacy-by-design dashboard exposes patient details/

Ayushman Bharat Digital Mission: Are Indian Health Records Safe? MEDIANAERA (Sept. 20, 2024)

https://www.medicinenama.com/2012/09/29/99999999999 privacy-and-integrity-ayushman-bharat-digital-mission-pmid-dashboard/

2023Medianama, supra note 19.

2023Cyril Amarchant Mangoldas, FAQs — The Digital Personal Data Protection Act, 2023 (Dec. 2025). https://www.cyriallisc.com/wp-content/uploads/2025/12/FAQs-DPDA.pdf.

2023IndiaAl, supra note 18:
The MIDAS Touch, KERNEL — Indian Institute of Science (2023), https://kernel.iisc.ac.in/the-midas-touch/

Threat Model

For practitioners, understanding AI's cybersecurity and privacy risks is as important as understanding its clinical potential. Deploying AI in healthcare requires recognising where vulnerabilities arise across data collection, model development, deployment, and institutional infrastructure, and implementing safeguards that protect patient confidentiality, system integrity, and regulatory compliance.

1. **Attack surface:** The relevant perimeter now includes EHR-vendor APIs, third-party AI inference endpoints, ambient-scribe pipelines that transmit raw consultation audio to cloud processing, and the inter-hospital data links a hub-and-spoke network like MIDAS requires by design. Each additional integration point is another place where access controls, encryption, or vendor vetting can fail.

2. **Re-identification risk:** "Anonymised" carries more weight in policy documents than the underlying science supports. Researchers have shown that a generative statistical model can correctly re-identify 99.98% of Americans in a dataset using just fifteen demographic attributes, even in heavily sampled or incomplete datasets, directly undercutting the release-and-forget model of anonymisation that pooled health datasets typically rely on.[25] Work specific to biomedical signals reinforces the point: genomic data can, in some circumstances, be linked back to public face images and even electrocardiogram traces, a routine input to diagnostic AI which carries individually distinctive features that a trained neural network can use to re-identify a person in an ostensibly de-identified dataset.[26]

3. **Model-level risks:** Beyond raw data, trained models can themselves leak information about the people whose records shaped them. Membership-inference attacks, formalized by Shokri and colleagues in 2017, let an adversary determine whether a specific individual's record was used to train a model.[27] Applied to brain-MRI age-prediction models, researchers correctly identified whether a given scan had been used in training 60% to more than 80% of the time, depending on model complexity and the attacker's assumptions, which include federated-learning setups explicitly designed to avoid sharing raw data.[28] Model-inversion attacks, which attempt to reconstruct input data from a model's outputs, add a further layer of exposure for any diagnostic model exposed through an API.[29]

[25] Luc Rocher, Julien M. Hendrickx & Yves-Alexandre de Montjoye, *Estimating the Success of Re-Identifications in Incomplete Datasets Using Generative Models*, 10 Nature Commc'ns 3069 (2019).

[26] Siamese Neural Network-Enhanced Electrocardiography Can Re-Identify Anonymized Healthcare Data, PMC (2025), https://www.ncbi.nlm.nih.gov/pmc/articles/PMC12088719/.

4. Insider risk and access sprawl:

A hub-and-spoke network multiplies the number of institutions, IT departments, and individual staff with some level of legitimate access and, with it, the number of potential points of insider misuse or credential compromise.

5. Incident precedent:

India's healthcare sector already has a documented breach history at scale. A November 2022 ransomware attack on AIIMS Delhi encrypted roughly 1.3 terabytes of data across five servers and put the records of an estimated 40 million patients at risk, forcing a two-week reversion to manual operations.[30] An October 2023 leak traced to India's COVID-testing data pipeline put an estimated 815 million records, including unmasked Aadhaar numbers, up for sale on a dark-web forum.[31] In October 2024, Fortis Healthcare and Star Health Insurance each disclosed breaches affecting patient and policyholder data, including, in Fortis's case, allegedly

[27]Reza Shokri, Marco Stronati, Congzheng Song & Vitaly Shmatikov, *Membership Inference Attacks Against Machine Learning Models*, 2017 IEEE Symp. on Security & Privacy 3.

[28]Umang Gupta, Dimitris Stripelis, Pradeep K. Lam, Paul M. Thompson, José Luis Ambite & Greg Ver Steeg, *Membership Inference Attacks on Deep Regression Models for Neuroimaging*, 143 Proc. Machine Learning Rsch. 228 (2021).

[29]Georgios Kaissis et al., *End-to-End Privacy Preserving Deep Learning on Multi-Institutional Medical Imaging*, Nature Mach. Intel. (2021), https://www.nature.com/articles/s42256-021-00337-8.

[30]Cyberattack on Top Indian Hospital Highlights Security Risk, AL JAZEERA (Dec. 7, 2022), https://www.aljazeera.com/economy/2022/127/cyberattack-on-top-indian-hospital-highlights-security-risk; Threats in the Cyberspace: Critical Systems at Risk, DECCAN HERALD, https://www.deccanherald.com/amp/story/india%2Fthreats-in-the-cyberspace-critical-systems-at-risk-1192686.html.

[31]Leaked Covid Test Data Largest Breach in India's History, TECHINFORMED (Oct. 31, 2023), https://techinformed.com/icmr-data-breach-one-of-largest-in-indias-history/; ICMR India 815 Million Records Breach 2023 — How, RINGSAFE (Apr. 24, 2026), https://ringsafe.in/icmr-india-815-million-records-breach-october-2023-analysis/.

exposed medical records and financial statements.[32] A further vulnerability disclosed in 2025 exposed unauthenticated organ-donor records through an AIIMS portal.[33] None of these incidents involved an AI model directly, but each shows that institutions now expected to feed AI-training pipelines have not yet secured the basic patient-record infrastructure those pipelines will depend on.

[32]Leaked Covid Test Data Largest Breach in India's History, TECHINFORMED (Oct. 31, 2023), https://techinformed.com/icmr-data-breach-one-of-largest-in-indias-history/; ICMR India 815 Million Records Breach 2023 — How, RINGSAFE (Apr. 24, 2026), https://ringsafe.in/icmr-india-815-million-records-breach-october-2023-analysis/.

[33]Umang Gupta, Dimitris Stripelis, Pradeep K. Lam, Paul M. Thompson, José Luis Ambite & Greg Ver Steeg, *Membership Inference Attacks on Deep Regression Models for Neuroimaging*, 143 Proc. Machine Learning Rsch. 228 (2021).

GOVERNANCE AND REGULATORY GAP ANALYSIS

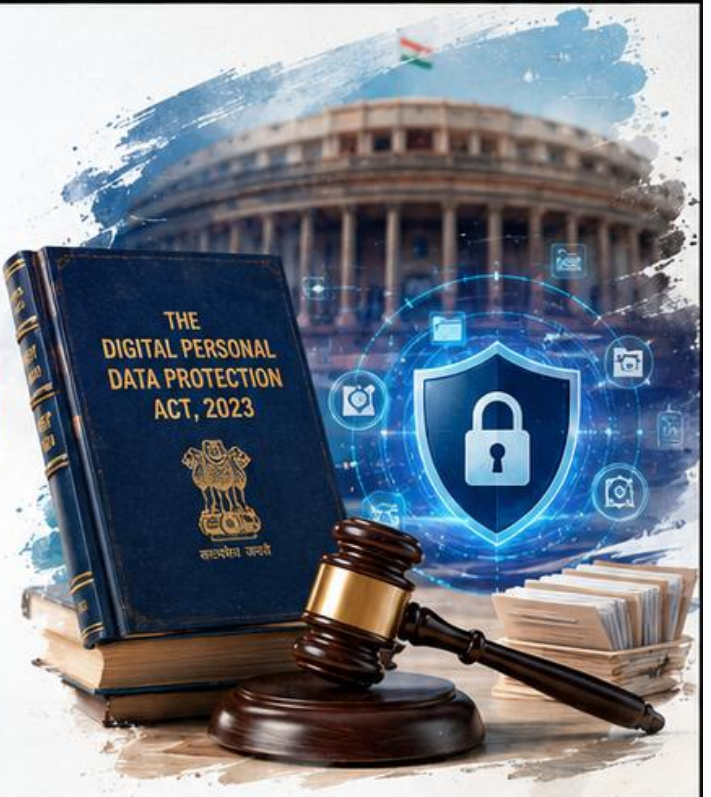

India's Digital Personal Data Protection Act, 2023, is, on paper, a significant advance: it establishes a Data Protection Board of India, mandates consent and breach notification, and restricts cross-border data flows.[34] Two gaps matter most for health data specifically. First, unlike the dedicated "sensitive personal data" category found in India's earlier 2011 rules and in comparable foreign statutes, the DPDP Act applies the same general standard to health data as to any other personal data; guidance on what counts as health data must be inferred from the Act read together with its 2025 Rules rather than stated directly.[35] Second, enforcement is not yet fully operative: the Rules were notified only in November 2025, and core provisions are being brought into force on a multi-year, phased schedule.[36] For the next several years, India's health-AI infrastructure is being built under a statute whose enforcement machinery is still being assembled.

> **A STRONG LAW ON PAPER. A LONG ROAD TO FULL ENFORCEMENT.**

The Indian Council of Medical Research's 2023 Ethical Guidelines for AI in Biomedical Research and Healthcare fill part of the gap, setting out ten principles including autonomy, data privacy, risk minimisation, and accountability and requiring institutional ethics committee review of AI research proposals.[37] But these are guidelines, not statute: they operate through the ethics committee and research-funding system rather than through courts, and legal commentators have observed that the guidelines were finalised without fully reconciling their provisions against the data-protection law that was, at the time, still in draft.

[32]When Cyberattacks Turn Deadly: The Silent Crisis in Healthcare Cybersecurity, NEUROLOGY INDIA (2025), https://journals.lww.com/neur/fulltext/2025/05000/when_cyberattacks_turn_deadly___the_silent_crisis.27.aspx.

[33]AIIMS Portal Vulnerability Exposed Sensitive Donor Data, THE CYBER EXPRESS (July 28, 2025), https://thecyberexpress.com/aiims-portal-vulnerability/.

[34]The Digital Personal Data Protection Act, 2023, No. 22 of 2023, Acts of Parliament, 2023 (India), https://www.meity.gov.in/static/uploads/2024/06/2bbf10e9f04e6fb4f8fef35e82c42aa5.pdf; PRS Legislative Research, The Digital Personal Data Protection Bill, 2023, https://prsindia.org/billtrack/digital-personal-data-protection-bill-2023 (last visited July 27, 2026).

[35]Mondaq, *supra* note 4.

[36]DLA Piper, *supra* note 5.

[37]ICMR, *supra* note 6; Use of AI & Data Privacy Implications in Healthcare Sector, AMLEGALS (June 18, 2025), https://amlegals.com/use-of-ai-data-privacy-implications-in-healthcare-sector/.

operate through the ethics committee and research-funding system rather than through courts, and legal commentators have observed that the guidelines were finalised without fully reconciling their provisions against the data-protection law that was, at the time, still in draft.

Dimension	DPDP Act, 2023 (India)	HIPAA (United States)	GDPR (European Union)
Dedicated health-data category	No, general standard applies	Yes – Protected Health Information	Yes, "special category" data
Breach notification	72 hours, to the Board and to affected individuals	60 days, to individuals and HHS	72 hours, to supervisory authority
Independent audit of AI-specific ethics claims	No dedicated body; ICMR guidance is non-statutory	Sector regulators (HHS/OCR)	Data Protection Authorities; DPIA required

**The table is compiled from the sources cited throughout this Part.*[38]

> This leaves an open institutional question that runs through this paper: nobody currently holds a clear statutory mandate to audit whether an infrastructure project's claim to be "ethics compliant" is actually verified, as opposed to aspirational.

[38] Sakrishna & Associates, The Indian Council of Medical Research Released the Ethical Guidelines for Application of Artificial Intelligence in Biomedical Research and Healthcare (Nov. 23, 2023), https://www.saikrishnassociates.com/the-indian-council-of-medical-research-released-the-ethical-guidelines-for-application-of-artificial-intelligence-in-biomedical-research-and-healthcare/.

[39] The Digital Personal Data Protection Act, 2023, *supra* note 35; Cyril Amarchand Mangaldas, *supra* note 23; HIPAA vs DPDP Act: Complete Comparison Comparison for Healthcare AI in India (2026), WKDRS (Apr. 22, 2026), https://www.wekhors.com/blog/hipaa-vs-dpdp-act/; GDPR vs HIPAA: 7 Differences You Must Look Out For, CAPTAIN COMPLIANCE (Jan. 20, 2025), https://captaincompliance.com/education/gdpr-vs-hipaa/.

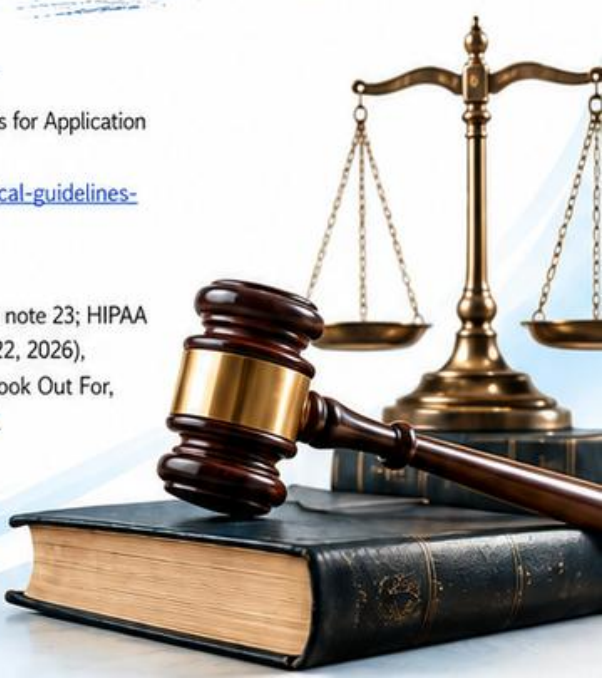

Stakeholder Incentives and Accountability

A useful test for any data-pooling project is to ask who can walk away from it. Hospitals, AI vendors, and research institutions that build or license these systems have commercial, reputational, or funding incentives to participate, along with comparatively strong bargaining power over the terms; government bodies set the rules and often co-own the resulting infrastructure. Patients supply the data with the least practical ability to negotiate its terms, a gap that is widest for patients in the rural and low-resource districts prioritised by tuberculosis- and cancer-screening AI programmes.[40]

The ABDM consent form discussed in Part 4 is one concrete instance of this asymmetry: it formally satisfies a consent requirement while functionally shifting the burden of understanding what will happen to the data onto the patient signing it, typically mid-treatment and without an easily accessible grievance channel.[41] Meaningful consent would require, at minimum, purpose-specific disclosure understandable at the point of care and a genuine ability to decline without losing access to treatment; current practice, on the evidence gathered here, falls short of that standard.

CyberPeace Recommendations

8.1 For policymakers

1. Accelerate and fully resource the phased DPDP Rules implementation rather than let the multi-year timeline become a de facto grace period;[42]
2. Issue sectoral rules treating health data as a heightened category, closing the gap identified in Governance and Regulatory Gap Analysis mentioned above under this paper.
3. Create or designate an independent body with statutory authority to audit "ethics compliant" and "privacy by design" claims made by national infrastructure projects, rather than relying on self-certification;
4. Reconcile CDSCO's medical-device approval process with DPDP obligations into a single compliance pathway for diagnostic AI vendors.

[40] NTI Aayog Frontier Tech, *supra* note 7.

[41] Medianama, *supra* note 19; Internet Freedom Foundation, AB PM-JAY Counts Patients but Discounts Patient Privacy (Jan. 8, 2024), https://internetfreedom.in/ab-pm-jay-patient-privacy/.

[42] DLA Piper, *supra* note 5.

8.2 For hospitals and institutions

1. Replace blanket, enrolment-time consent with granular, purpose-specific consent captured at each new use of data;
2. Apply data minimization by default to ambient-scribe audio and EHR-embedded AI tools, retaining raw recordings only as long as clinically necessary;
3. Build breach-response capability to the DPDP's notification clock, treating the two-week AIIMS outage as a cautionary baseline rather than an acceptable one.[43]
4. Insert enforceable security and data-use requirements into every third-party AI vendor contract.

8.3 For AI developers

1. Treat membership-inference and model-inversion resistance as a pre-deployment testing requirement for any model trained on pooled hospital data, not an academic afterthought;[44]
2. Adopt differential privacy and federated learning wherever pooling is unavoidable, building on MIDAS's own stated move toward differential-privacy tooling through a collaboration involving Harvard's Dataverse as an early template worth generalizing;
3. Publish security and privacy audits of deployed models, not only accuracy benchmarks.

8.4 For civil society and NGOs

1. Monitor the phased DPDP Rules rollout and flag gaps between the commencement schedule and real-world deployment;
2. File Right to Information requests with ICMR-IISc, ARTPARK, and hub hospitals on the data-sharing agreements underlying MIDAS and comparable projects; run data-rights literacy campaigns targeted specifically at the rural and low-resource populations disproportiontebly captured by point-of-care screening programs, who are least likely to encounter an accessible consent process or a functioning grievance channel;[45]
3. Track the actual, not merely promised, resolution of known vulnerabilities, following the model of the independent researcher who disclosed the AIIMS donor-portal exposure through responsible channels.[46]

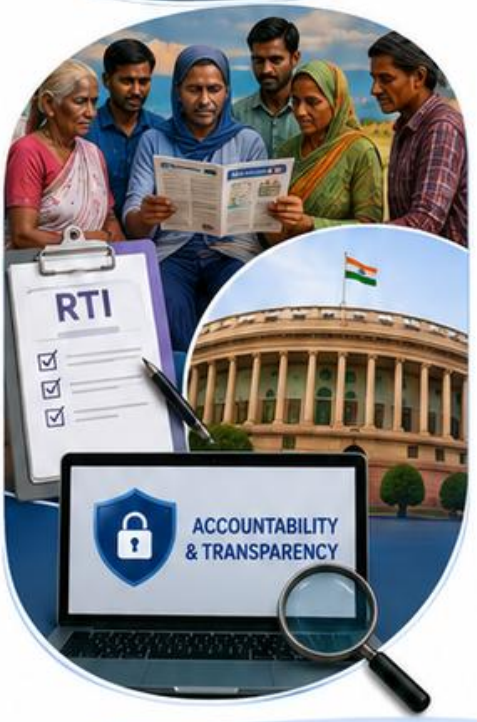

[43]Al Jazeera; Deccan Herald, *supra* note 31.

[44]Gupta et al., *supra* note 29; Kaissis et al., *supra* note 30.

CONCLUSION

The sources surveyed in this paper repeatedly frame India's AI health rollout as a race between capabilities: more hospitals joining MIDAS, more diagnostic centres running point-of-care AI, and more clinicians adopting ambient scribes. The more useful frame, and the one this paper has tried to sustain throughout, is a race between capabilities and rights between how fast data can be pooled and how fast consent architecture, breach-notification machinery, and independent audit capacity can mature alongside it. On the present timeline, rights are behind. India's core data-protection rules will not be fully operative for several more years; its principal AI-health ethics instrument remains voluntary; and its hospitals have already suffered breaches at a scale of tens of millions of records at AIIMS alone that should have made safeguard-first design the default rather than the exception. Responsible AI health infrastructure is achievable, but only if the sequencing changes: safeguards are built alongside capability, not retrofitted after the next breach forces the question.

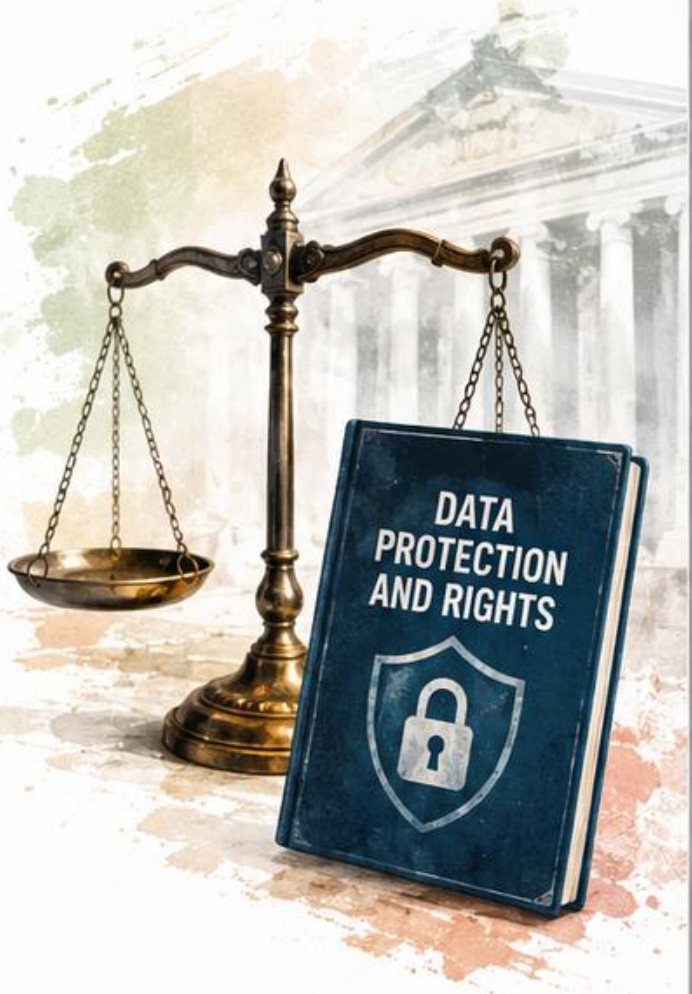

AUTHORS

1. Neeraj Soni, Sr. Researcher, Policy & Advocacy, CyberPeace
2. Isharth Kumar, NLIU Bhopal, Intern, Policy & Advocacy, CyberPeace

[45]PMC, *supra* note 21; Internet Freedom Foundation, *supra* note 41.

[46]The Cyber Express, *supra* note 34.

GLOSSARY

TERMS	DEFINITIONS
ABDM	Ayushman Bharat Digital Mission, India's national digital health platform and consent management framework.
Ambient scribe	AI software that listens to a clinical visit and automatically drafts the medical note from the conversation.
CDSCO	Central Drugs Standard Control Organisation, India's regulator for diagnostic and other medical-device software.
Data Protection Board of India	The adjudicatory body created under the DPDP Act to hear and rule on non-compliance.
De-identification	The process of stripping direct identifiers (names, ID numbers) from a dataset before it is shared or pooled.
Differential privacy	A technique that adds calibrated statistical noise to data or model outputs to limit what can be learned about any one individual.
DPDP Act	India's Digital Personal Data Protection Act, 2023, is the country's core data-protection statute.
DPDP Rules	The Digital Personal Data Protection Rules, 2025, which operationalise the Act on a phased, multi-year timeline.
ICMR	Indian Council of Medical Research, author of India's 2023 ethical guidelines for AI in biomedical research and healthcare.
Membership inference attack	A technique that lets an adversary determine whether a specific individual's record was used to train a model.
MIDAS	Medical Imaging Datasets for India, a national hub-and-spoke project pooling AI-ready medical-imaging data from hospitals.
Model inversion attack	A technique for reconstructing a model's original training or input data from its outputs.
PHI	Protected Health Information; the HIPAA term for individually identifiable health data.

CyberPeace

CAPABILITIES WITHOUT SAFEGUARDS:

REFRAMING INDIA'S AI-IN-HEALTHCARE BOOM AS A DATA-RIGHTS PROBLEM