

AI Rules in India → A Healthcare Perspective



A practical guide to the rules, laws, and frameworks shaping AI in India's healthcare ecosystem.

Foreword

Welcome to the 9th Edition of our Innovation Series: **AI Rules in India → A Healthcare Perspective.**

As artificial intelligence becomes more deeply embedded in healthcare systems, the conversation is no longer limited to what AI can do, but how it should be **built, governed, and used** in practice. In India, this shift is reflected not through a single policy or regulation, but through a set of evolving **laws, guidelines, and system-level frameworks** that together shape the direction of healthcare AI. Understanding this landscape requires looking beyond individual rules and viewing how these elements interact within real-world healthcare environments.

This guide is written for a broad audience, from medical students, clinicians, and public health practitioners to engineers, founders, product leaders, researchers, and policymakers working at the intersection of healthcare and AI. It is intentionally designed for readers who may not have deep expertise in law, policy, or technical architecture, but who want to understand how India's AI rules and frameworks influence real-world healthcare systems, decisions, and outcomes.

Unlike regulatory summaries that focus on legal interpretation or compliance detail, this guide takes a systems view. It does not attempt to **evaluate or critique individual laws** or frameworks. Instead, it aims to bring clarity to how different layers, ranging from guidance and validation to data governance, interoperability, ethics, and regulations collectively shaping how AI is **built, tested, deployed, and used** in healthcare. In this sense, the focus is not on any single rule, but on the structure that emerges when they are viewed together.

This volume (**Vol. 01**) is a high-level overview designed to simplify a complex and evolving space. It provides a structured way to understand the key components of India's healthcare AI ecosystem without going into detailed legal or technical analysis. The intention is to create a common foundation that can be used by different stakeholders to navigate this landscape with greater clarity.

At HealthInnovation Toolbox, our mission is to support responsible and meaningful innovation in healthcare through accessible knowledge, practical frameworks, and ecosystem-focused thinking. This guide is part of a broader effort to strengthen understanding at the intersection of **healthcare, data, and intelligent systems**, enabling more informed and responsible adoption of AI across diverse settings.


— **Team HealthInnovation Toolbox**

What's Inside This Guide?

INDIA'S AI MOMENT & WHY THIS MATTERS NOW?	05-07
From "Move Fast" → "Move Responsibly"	06
When One Model Doesn't Fit All Patients	07
NEW & EXISTING RULES → DECODED	08-10
Six key layers shaping how AI is built, tested, and used in healthcare in India	09
AI Rules → Healthcare Impact Matrix	10
SAHI → THE GUIDANCE LAYER	11-16
What is SAHI?	12
SAHI → Key Focus Areas	13
SAHI Stakeholders Mapping	14
SAHI Across Healthcare Stakeholders	15-16
BODH → THE VALIDATION LAYER	17-22
What is BODH?	18
BODH → Key Focus Areas	19
BODH Stakeholders Mapping	20
BODH Across Healthcare Stakeholders	21-22
DPDP → DATA PROTECTION LAYER	23-28
What is DPDP?	24
DPDP → Key Focus Areas	25
DPDP Stakeholders Mapping	26
DPDP Across Healthcare Stakeholders	27-28

What's Inside This Guide?

ABDM → THE DATA EXCHANGE LAYER	29-34
What is ABDM?	30
ABDM → Key Focus Areas	31
ABDM Stakeholders Mapping	32
ABDM Across Healthcare Stakeholders	33-34
ICMR → THE ETHICS LAYER	35-41
What is ICMR?	36
A Landmark Move Toward India-Centric Healthcare Evidence	37
ICMR → Key Focus Areas	38
ICMR Stakeholders Mapping	39
ICMR Across Healthcare Stakeholders	40-41
CDSCO → THE REGULATORY LAYER	42-47
What is CDSCO?	43
CDSCO → Key Focus Areas	44
CDSCO Stakeholders Mapping	45
CDSCO Across Healthcare Stakeholders	46-47
PUTTING IT ALL TOGETHER	48



India's AI moment – & Why this matters now?

The **AI landscape** is evolving rapidly, both globally and in **India**. Over the past few years, governments have begun introducing frameworks and guidelines to better understand and manage the growing impact of artificial intelligence. India is now taking clearer steps in this direction, signalling a move toward more structured oversight. The focus is not on restricting innovation, but on ensuring that AI systems are developed and used in ways that are **responsible, transparent, and aligned with public interest**.

This shift becomes especially important in healthcare, where the impact of AI is both promising and sensitive. From supporting diagnostics to improving clinical workflows, AI has the potential to address real gaps in access and efficiency. At the same time, it raises important questions around **data quality, decision-making, and patient awareness**. As India shapes its approach to AI governance, healthcare will be a key space where balancing innovation with safety and trust becomes essential. This guide breaks down these emerging AI directions into clear, practical insights, with a specific focus on what they mean for real-world healthcare use. Let's go....

FROM “MOVE FAST” → “MOVE RESPONSIBLY”

For a long time, AI progress was measured by how quickly something could be built and deployed. The underlying assumption was simple: **if a system works, it should scale**. But as AI began moving from controlled environments into everyday decision-making especially in areas like healthcare, where the definition of “**working**” started to change. Accuracy alone was no longer enough; questions around reliability across populations, explainability of outcomes, and accountability for errors began to surface more clearly.

This shift is not about slowing innovation, but about redefining what good innovation looks like. Moving responsibly means thinking beyond performance metrics to include **context, consequences, and trust**. In India, this transition is becoming more visible as AI systems are increasingly expected to demonstrate not just capability, but also clarity, fairness, and alignment with real-world needs.

Generalization gap in healthcare AI

AI diagnostic performance can drop by **20–40%** when applied to new populations outside training data.

Representation gap

Up to 80% of health data used in AI comes from **limited** geographic or demographic sources.

Major real-world risk

Nearly 60% of clinical AI models lack external validation before deployment.

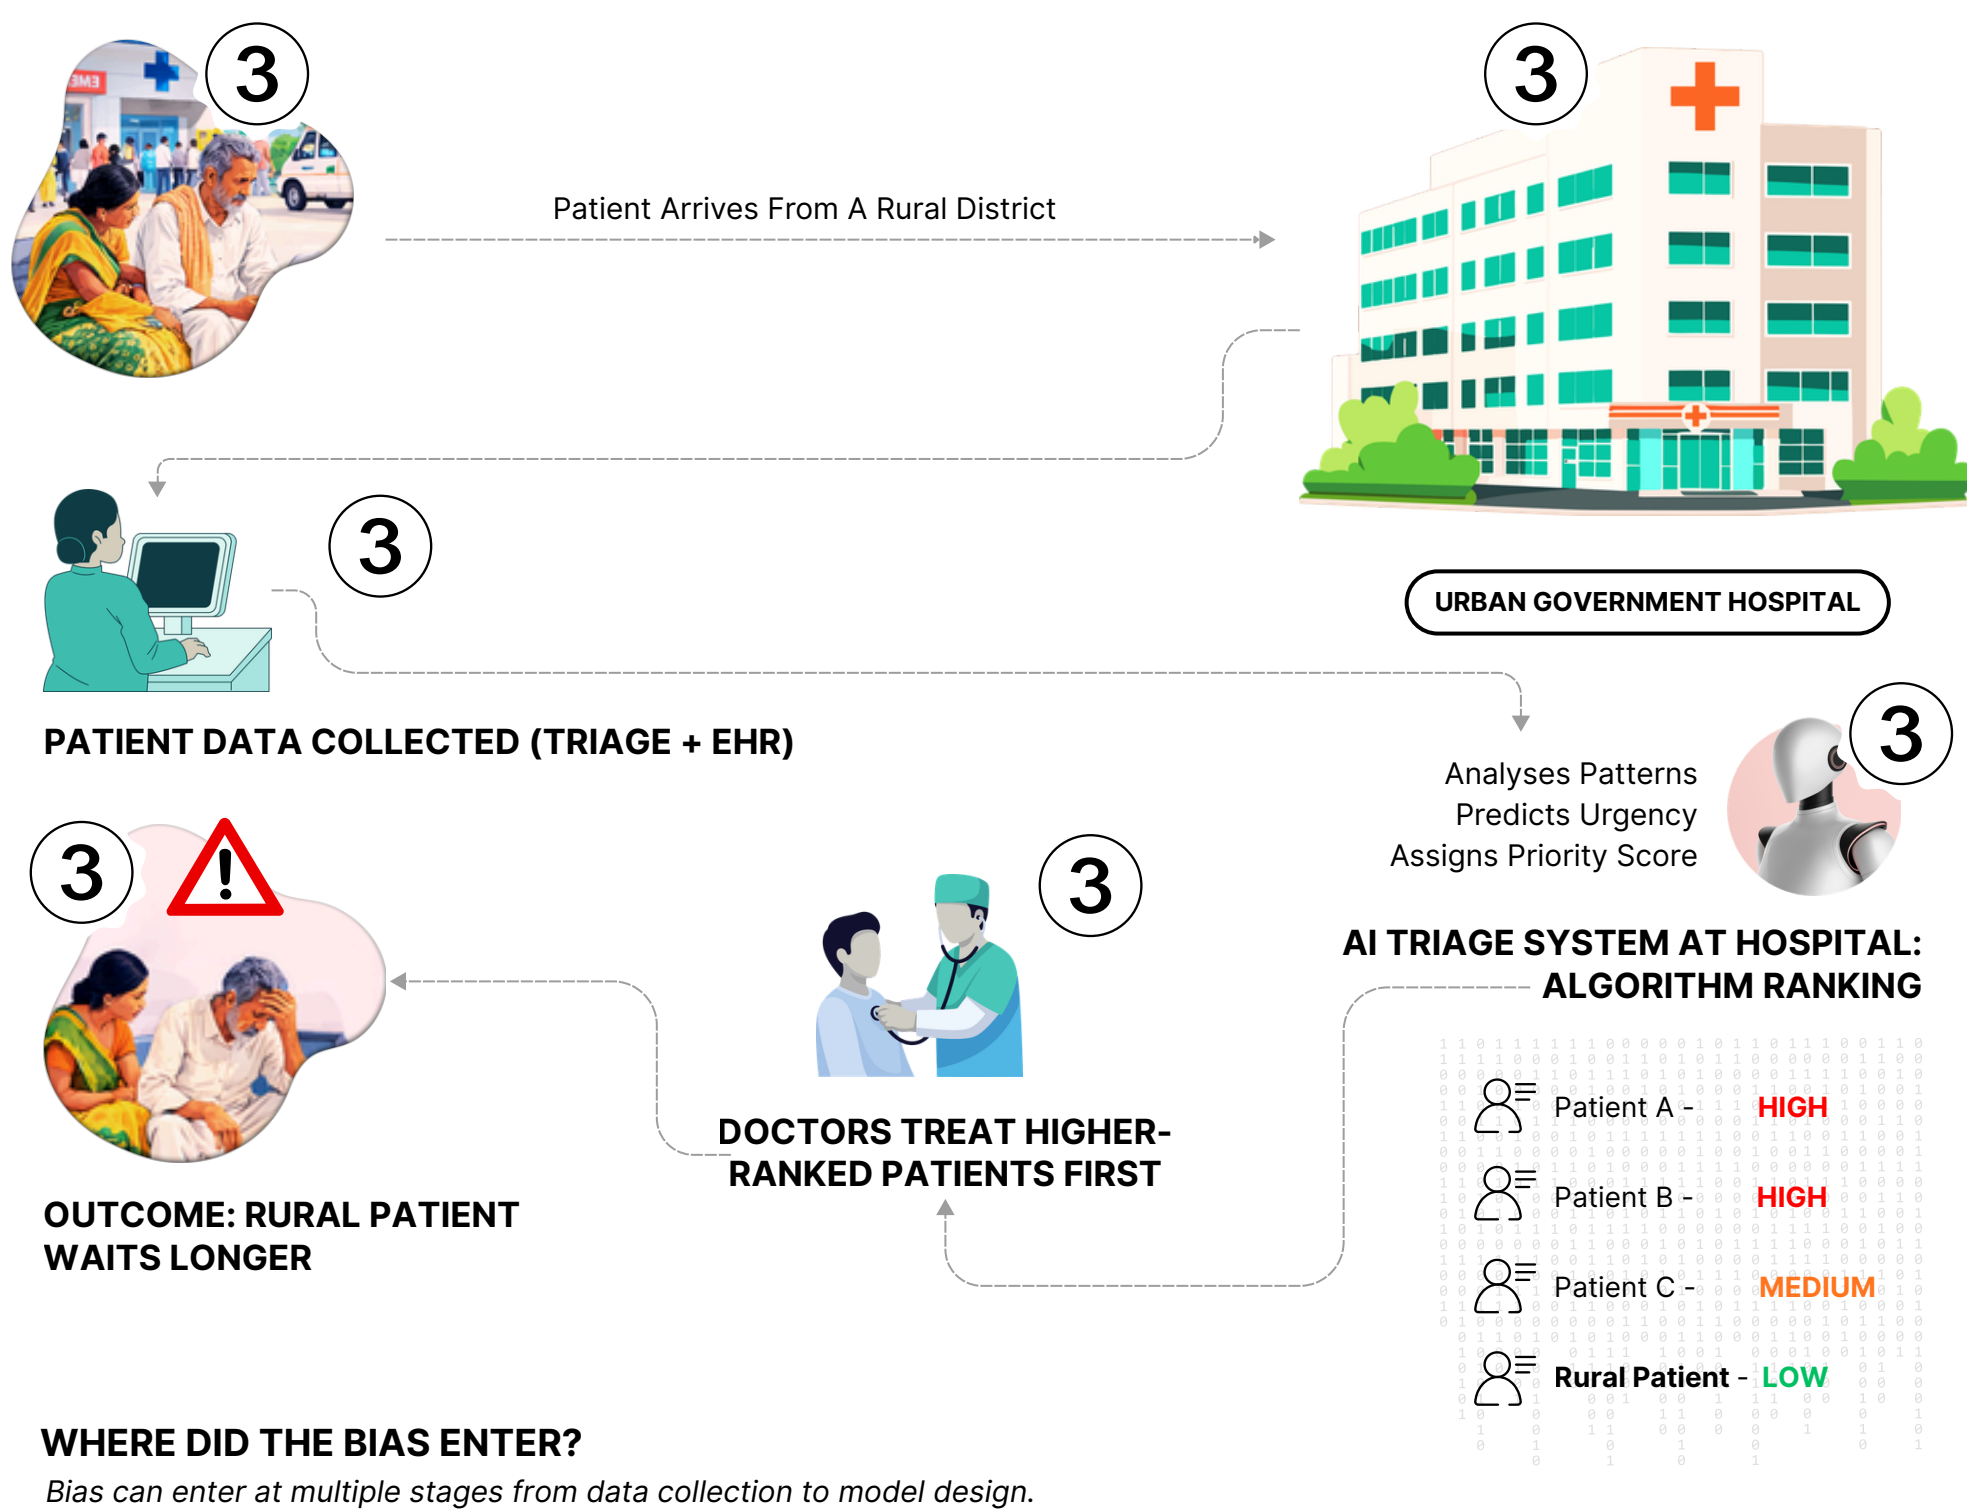
Hidden Bias at Scale

Bias in healthcare algorithms has been shown to affect millions of patients by underestimating care needs.


Sources: Nature Medicine; The Lancet Digital Health; BMJ; WHO; Science Direct

A patient from a rural village travels several hours to seek care at a large urban government hospital. At the emergency department, an AI-based triage system helps prioritize patients using patterns learned from historical clinical data. After entering his details → age, symptoms, location, and limited medical history, the system automatically ranks his case lower than others waiting for treatment. Doctors begin attending to higher-ranked patients first. But this raises an important question:

If the algorithm was trained primarily on data from urban hospital populations, could rural patients be unintentionally under-prioritized?



LONGER WAITING TIMES MAY BE THE SMALLEST COST OF BIASED TRIAGE ALGORITHMS → MISDIAGNOSIS, DELAYED TREATMENT, AND PREVENTABLE COMPLICATIONS ARE THE FAR MORE DANGEROUS CONSEQUENCES.



New & Existing Rules → Decoded

India's approach to AI is not defined by a single law, but by a set of **emerging frameworks** that together shape how **AI is built, tested, and deployed**, especially in healthcare. These new directions span data protection, clinical validation, ethical use, and regulatory oversight, each addressing a different part of the AI lifecycle.




Rather than viewing them in isolation, it is more useful to see them as interconnected layers that collectively define expectations around **safety, transparency, and accountability**. This section brings together the key initiatives that matter most for healthcare AI, and decodes how they influence real-world development and use.

SIX KEY LAYERS SHAPING HOW AI IS BUILT, TESTED, AND USED IN HEALTHCARE IN INDIA.

These **six layers** together outline how healthcare AI is shaped in India, from how data is accessed and protected, to how models are validated, regulated, and applied in real clinical settings. Each layer plays a distinct role, but their combined effect is what ensures that AI systems are not just functional, but **safe, reliable, and aligned** with healthcare needs. These six elements represent a mix of **laws, guidelines, and operational frameworks** that together shape how healthcare AI is evolving in India.

From legal instruments governing data protection and medical device regulation, to ethical guidelines and system-level frameworks enabling **consent, validation, and deployment**, each plays a distinct role across the AI lifecycle. Viewed collectively, they provide a more complete picture of how AI is being structured, not just to function effectively, but to operate safely, responsibly, and in alignment with real-world healthcare needs.






 Legal  Strategic Frameworks  Operational Frameworks

AI Rules → Healthcare Impact Matrix

This matrix provides a simplified view of how each rule and framework influences different aspects of **healthcare AI**, from how data is handled to how systems are validated and deployed. Rather than operating in isolation, each plays a role across the lifecycle of an AI solution. For example, data-focused frameworks like **DPDP** and **ABDM** shape how patient information is accessed and shared, while systems like **BODH** and **CDSCO** focus more on validation and clinical safety. Together, they highlight that building healthcare AI in India requires attention not just to performance, but to compliance, trust, and real-world accountability.

Rule / Framework	Data & Privacy	Clinical Safety	Validation	Transparency	Deployment	Accountability
SAHI	●		●		●	
BODH	—			●	●	●
DPDP Act		●	—		●	
ABDM		●	—	●		●
ICMR	●		●	●	●	
CDSCO	●			●		

-  High Impact
-  Moderate Impact
-  Limited / indirect Impact



SAHI → The Guidance Layer

Setting the direction for how AI should be built and used in healthcare.

What is SAHI?

The SAHI (Safe & Responsible AI for Healthcare) framework represents India's early structured effort to define how AI should be developed and used within the healthcare ecosystem. Emerging from policy discussions and national strategy work (including contributions from bodies like **NITI Aayog** and broader digital health initiatives), SAHI is not a single enforceable law, but a **guiding framework that outlines expectations for responsible AI adoption**. It reflects a shift toward building AI systems that are not only effective, but also safe, transparent, and aligned with **public health priorities**.

Why Its Introduced?

India's healthcare system presents unique challenges for AI adoption, with wide variations in infrastructure, access, and patient demographics across regions. From urban hospitals to resource-constrained rural settings, AI systems are expected to perform reliably in highly diverse and often unpredictable environments. In such a context, even small errors or biases can have amplified consequences on patient care.

As AI applications began expanding rapidly across healthcare, particularly in diagnostics, triage, and predictive analytics, it became clear that traditional evaluation metrics were not sufficient. Questions around **bias, explainability, and accountability** started gaining importance, especially in high-stakes clinical environments. SAHI was introduced to address this gap, providing a structured direction for how AI systems should be designed and deployed in a way that **balances innovation with patient safety and trust**.

SAHI responds to this reality by emphasizing the need for reliability across **diverse populations, transparency** in decision-making, and a clear understanding of system limitations. It also highlights the importance of aligning AI tools with existing clinical workflows and ensuring meaningful human oversight. Rather than focusing only on controlled performance metrics, SAHI encourages **evaluating how AI systems behave in real-world Indian healthcare settings**, where variability and scale are defining factors.

SAHI → Key Focus Areas


SAHI brings structure to an otherwise fragmented AI landscape by outlining the key dimensions that need attention before systems are used in real healthcare settings. Rather than focusing only on model performance, it highlights the broader considerations that influence how AI behaves in practice across different environments, users, and patient groups.

Safety-first design



Minimizing risk in high-impact use cases such as diagnosis or treatment recommendations, where errors can directly affect patient outcomes. In the Indian context, this includes ensuring systems work reliably even in low-resource settings and under variable data quality.

Explainability



Making AI outputs understandable to clinicians, so they can interpret, trust, and appropriately act on recommendations. This is especially important in India where AI is often used to support, not replace, clinical judgment.

Fairness



Reducing bias across diverse demographic and regional populations, including variations in age, gender, geography, and socio-economic background. Models must perform consistently across India's highly heterogeneous population.

Accountability



Clearly defining who is responsible when AI systems influence decisions whether developers, deployers, or healthcare providers, ensuring there are mechanisms to review and address errors.

Lifecycle thinking



Looking beyond initial deployment to continuously monitor, update, and validate AI systems as they interact with real-world data, clinical workflows, and evolving healthcare needs.

SAHI Stakeholders Mapping

SAHI does not affect the healthcare ecosystem uniformly; its influence varies depending on how different stakeholders interact with AI systems. From those who build and deploy AI to those who use it in care delivery or are impacted by its outcomes, each group experiences a different set of expectations and responsibilities. For **developers**, it shapes how systems are designed; for **institutions**, how they are evaluated and integrated; and for **care providers**, how they are used in practice. **Patients**, in turn, experience its effects through fairness, transparency, and quality of care.

Understanding these differences is important because the effectiveness of AI in healthcare depends not just on the technology itself, but on how each stakeholder engages with it. The following sections outline how SAHI translates into practical implications across the healthcare ecosystem.

Stakeholder	Design Expectations	Operational Changes	Risk & Accountability	Adoption & Trust
AI Builders (India)	✓	●	✓	●
Global AI Companies	✓	●	✓	●
Hospitals & Systems	●	✓	✓	✓
Clinical Workforce	●	●	✓	✓
Patients & Public	—	●	✓	✓



High Impact



Moderate Impact



Limited / indirect Impact

SAHI Across Healthcare Stakeholders

For Hospitals & Health Systems

SAHI shapes how hospitals evaluate and adopt AI, pushing them to look beyond performance claims and assess whether systems are **reliable, interpretable, and suited** to their specific care settings. It also emphasizes the need to ensure AI tools fit into existing workflows and infrastructure. At the same time, it highlights the importance of ongoing oversight, encouraging hospitals to monitor performance, manage risks, and take responsibility for how AI is used in real-world clinical environments.

Clinical Workforce

SAHI shapes how AI is used across the clinical workforce, including doctors, nurses, and allied staffs, by influencing how these tools fit into everyday care workflows. It emphasizes that AI should support, not replace, clinical judgment, with outputs that are clear and usable in real-world settings. At the same time, it brings greater focus on responsible use encouraging clinical teams to engage actively with AI systems, understand their limitations, and apply them thoughtfully in patient care.

For Patients & Public

SAHI shapes how patients experience AI-driven care by encouraging systems that are fair, transparent, and appropriate across diverse populations. It highlights the need to reduce bias and ensure that AI does not disadvantage certain groups, **especially in a country as diverse as India**. At the same time, it reinforces the importance of clarity around how AI is used in healthcare, supporting more informed and trustworthy interactions between patients and digital systems.

For AI Builders (India)

SAHI influences how startups and developers design AI systems from the outset, pushing them to consider real-world variability, data diversity, and clinical usability. It encourages moving beyond model performance to ensure systems are relevant and reliable in Indian healthcare settings. It also signals the need to build with responsibility in mind, incorporating explainability, documenting limitations, and preparing for how systems behave after deployment.

For Global AI Companies

SAHI signals that solutions built for other markets may not directly translate to the Indian healthcare context. **Companies entering India are encouraged to adapt their models to local data, workflows, and population diversity.** It also sets expectations around responsible deployment, requiring global players to align with India's emerging approach to transparency, fairness, and real-world performance rather than relying solely on existing benchmarks.

SAHI brings a shift in how healthcare AI is approached across the ecosystem from isolated model development to more **context-aware and responsible deployment**. It highlights that building effective AI is not just about performance, but about how systems function across diverse settings, users, and patient populations. For stakeholders, this translates into greater attention to real-world usability, clearer understanding of limitations, and more active engagement with how AI is applied in practice. Rather than being a single checkpoint, responsibility becomes an ongoing consideration across the lifecycle of AI systems. In this way, **SAHI** sets the foundation for a more structured and thoughtful approach to integrating AI into healthcare.

SAHI defines the direction → how well it is implemented will determine its real impact.



BODH → The Validation Layer

Ensuring AI systems are tested, validated, and ready for real-world healthcare use.

What is BODH?

The **BODH (Benchmarking, Evaluation & Validation of Health AI Systems) framework** represents India's effort to bring structure to how AI systems are tested before they are used in real healthcare settings. Emerging as part of India's broader push toward **responsible AI** adoption, BODH focuses on ensuring that AI models are not only technically sound but also reliable across real-world clinical scenarios. Unlike guiding frameworks such as SAHI, which define how AI should be built, BODH addresses how these systems should be evaluated and validated before deployment.

As AI tools increasingly move into areas like diagnostics, triage, and decision support, the need for **standardized testing** becomes critical. Performance measured in controlled environments often does not reflect how systems behave across different hospitals, datasets, and patient populations. BODH is introduced to address this gap by promoting structured evaluation approaches that consider variability, robustness, and real-world applicability of AI systems.






Why Its Introduced?

Healthcare AI systems can perform well in testing but fail when exposed to new populations or different clinical settings. In India, this challenge is amplified due to diversity in **data, infrastructure, and care delivery environments**. BODH is introduced to ensure that AI systems are evaluated beyond internal testing, with greater focus on consistency, reliability, and performance across varied conditions.

It reflects a shift from “build and deploy” → “test, validate, then deploy”.

BODH → Key Focus Areas

BODH brings attention to what happens between **development and deployment**, an often overlooked stage where many risks emerge. It highlights the need to systematically evaluate how AI systems perform under different conditions before they are integrated into real healthcare workflows.

Benchmarking		Comparing AI systems against standard datasets or performance baselines to understand relative effectiveness. This helps identify how a model performs in comparison to accepted standards rather than in isolation.
Robustness Testing		Ensuring models perform consistently across different populations, hospitals, and data variations. This is critical in India where variations in data quality and patient demographics can significantly impact outcomes.
Validation Beyond Training Data		Moving beyond internal validation to test systems in external, real-world environments. This ensures that models generalize well and are not limited to the conditions they were originally trained on.
Error Analysis		Identifying where and why systems fail, especially in high-risk clinical scenarios. Understanding failure patterns helps improve reliability and reduces the risk of unintended consequences in care delivery.
Deployment Readiness		Assessing whether systems are suitable for real-world use, not just experimental success. This includes evaluating integration with workflows, usability, and performance under real operational conditions.

BODH Stakeholders Mapping

Unlike SAHI, which sets the direction for how AI systems should be designed, **BODH** focuses on how confidence in these systems is established before they are introduced into real healthcare environments. It brings attention to the critical stage between development and deployment, where systems need to be **evaluated under varied and practical conditions**. This includes assessing how models perform across different datasets, settings, and use cases, and whether they can maintain consistency beyond controlled testing scenarios. As a result, BODH influences how different stakeholders approach validation, interpret evidence of performance, and make decisions about readiness for real-world adoption.

Stakeholder	Design Expectations	Operational Changes	Risk & Accountability	Adoption & Trust
AI Builders (India)		●		●
Global AI Companies		●		●
Hospitals & Systems	●			
Clinical Workforce	—	●		●
Patients & Public	—	—	●	●

-  High Impact
-  Moderate Impact
-  Limited / indirect Impact

BODH Across Healthcare Stakeholders

For Hospitals & Health Systems

Hospitals are encouraged to look for evidence of external validation and robustness before adopting AI tools. This shifts adoption decisions toward more evidence-based evaluation rather than vendor claims. It also increases the importance of selecting tools that have been tested across similar healthcare environments.

For Clinical Workforce

BODH supports more reliable AI tools by ensuring they have been tested across varied conditions, helping clinical teams engage with systems that are better validated for real-world use. This allows providers to use AI with greater confidence while still applying clinical judgment in practice.

For Patients & Public

Patients benefit from AI systems that have undergone more rigorous testing, reducing the risk of errors and improving consistency of care across different settings. It also contributes to more dependable outcomes, especially when systems are used across diverse populations and care environments.

For AI Builders (India)

BODH pushes developers to move beyond internal testing and validate their models across diverse datasets and real-world conditions. It encourages building systems that can demonstrate consistent performance rather than relying only on controlled results. This also means preparing for variability in data quality, infrastructure, and clinical settings that are typical in India.

For Global AI Companies

BODH signals that validation done in other markets may not be sufficient for India. Companies entering India are expected to test their models against local data and healthcare settings before deployment. It reinforces the need to adapt solutions to **regional diversity** rather than relying on global benchmarks alone.

BODH strengthens a critical layer in the healthcare AI lifecycle by ensuring that systems are not only built well, but also rigorously tested before real-world use. It shifts attention from isolated performance to demonstrated reliability across diverse settings, datasets, and patient populations. For stakeholders, this brings greater focus on **evidence, validation, and readiness** before adoption, making evaluation a continuous part of how AI systems are trusted and used in practice.

If SAHI defines how to build responsibly → BODH defines how to prove it works.



DPDP → Data Protection Layer

Defining how personal health data is collected, used, and protected.

What is DPDP?

The Digital Personal Data Protection (DPDP) Act, 2023 is India's primary law governing how personal data is collected, processed, stored, and shared. It establishes a legal framework that defines the rights of individuals over their data and the responsibilities of entities handling that data. In the context of healthcare AI, where large volumes of sensitive patient information are involved, **DPDP plays a central role** in shaping how data can be accessed and used.

Unlike frameworks such as SAHI and BODH, which focus on how AI systems are built and validated, **DPDP addresses the foundation on which these systems operate data**. It introduces clear expectations around consent, purpose limitation, and data protection, making data governance a critical part of AI development and deployment.

Why It Was Introduced?





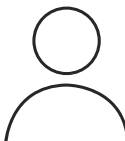
As digital systems expanded across sectors, concerns around misuse, unauthorized access, and lack of control over personal data became more prominent. In healthcare, these concerns are even more significant due to the sensitive nature of patient information. DPDP was introduced to:

- Give individuals greater control over their data
- Establish accountability for organizations handling data
- Create a structured legal environment for digital systems

DPDP defines the rules of data before AI can define outcomes.

DPDP → Key Focus Areas

DPDP shifts attention to how data is handled throughout the AI lifecycle, highlighting that access to data is not unrestricted but governed by clear rules and responsibilities. It brings structure to how personal data is **collected, used, and shared**, particularly in sensitive sectors like healthcare.

Consent Management		Ensuring that personal data is collected and used only with clear, informed, and specific consent from individuals. In healthcare AI, this means patients must understand how their data will be used, and systems must be designed to capture, manage, and respect consent throughout the data lifecycle.
Purpose Limitation		Data must be used only for the purpose it was collected, preventing misuse or overreach. For AI systems, this restricts using patient data beyond defined use cases, ensuring that models are trained and applied within clearly stated boundaries.
Data Minimization		Collecting only the data that is necessary, reducing exposure and risk. This encourages developers to avoid excessive data collection and focus on relevant inputs, which also helps reduce bias and improves system efficiency.
Data Security		Protecting data against breaches, unauthorized access, and misuse. In healthcare, this requires strong safeguards across storage, transmission, and processing to ensure sensitive patient information remains protected at all stages.
User Rights		Enabling individuals to access, correct, or request deletion of their data. This gives patients greater control over their information and requires systems to be designed with mechanisms that support these rights in practice.

DPDP Stakeholders Mapping

DPDP influences how data flows across the healthcare ecosystem by introducing clear rules around how personal information is accessed, shared, and used. It shapes not only how data is collected, but also how it is managed throughout its lifecycle, from **storage and processing to eventual use in AI systems**. In a healthcare setting, where multiple actors interact with the same data, this creates a more structured approach to handling sensitive information.

Its impact becomes most visible in how **responsibility, consent, and control** are distributed across different stakeholders. From developers designing data pipelines to hospitals managing patient records and providers accessing information during care, each participant operates within defined boundaries. At the same time, patients gain a more active role through rights and consent mechanisms, making data usage more transparent and accountable across the system.

Stakeholder	Design Expectations	Operational Changes	Risk & Accountability	Adoption & Trust
AI Builders (India)				●
Global AI Companies				●
Hospitals & Systems	●			
Clinical Workforce	—	●		●
Patients & Public		●		

-  High Impact
-  Moderate Impact
-  Limited / indirect Impact

DPDP Across Healthcare Stakeholders

For Hospitals & Health Systems

Hospitals must ensure that patient data is collected, stored, and shared in compliance with DPDP requirements. This includes managing consent, securing data systems, and maintaining accountability for how data is used across digital platforms.

For Clinical Workforce

DPDP influences how healthcare providers access and use patient data, encouraging more structured and responsible handling of information. It reinforces the importance of using data appropriately within defined purposes.

For Patients & Public

DPDP strengthens patient control over personal data, enabling individuals to understand, manage, and exercise rights over how their information is used. It supports greater transparency and trust in digital healthcare systems.

For AI Builders (India)

DPDP requires developers to design systems that are compliant with data protection principles from the outset. This includes integrating consent mechanisms, limiting data collection, and ensuring secure data handling practices throughout the lifecycle.

For Global AI Companies

DPDP sets clear expectations for how personal data must be handled within India, requiring global companies to align with local data protection laws. This may involve adapting **data storage, consent flows, and processing practices** to meet regulatory requirements.

DPDP establishes a foundational layer for healthcare AI by bringing structure and discipline to how data is handled across the system. It moves the ecosystem away from open or loosely governed data access toward **clearly defined, consent-driven processes**. In practice, this affects how data is collected, stored, shared, and integrated into AI workflows, making data governance a core consideration rather than a backend function.

For stakeholders, this translates into operational changes, such as building consent-aware systems, maintaining clear data usage boundaries, and ensuring secure data flows across platforms. It also requires organizations to rethink how data is accessed within teams and across partners, with greater emphasis on accountability and traceability. As a result, compliance becomes embedded in everyday decisions, shaping not just how AI systems are developed, but how they function and evolve in real-world healthcare environments.

If SAHI defines how to build → and BODH defines how to validate → DPDP defines how data can be used.



ABDM → The Data Exchange Layer

Enabling secure and consent-based exchange of health data across systems.

What is ABDM?

The Ayushman Bharat Digital Mission (ABDM) is India's national digital health initiative aimed at creating an interoperable health data ecosystem. It provides the infrastructure for securely linking health records across hospitals, labs, and digital platforms using standardized formats and unique health identifiers. Unlike DPDP, which governs how data should be protected, **ABDM focuses on enabling how data can be shared and accessed across the healthcare system.**

In the context of healthcare AI, ABDM plays a critical role by making structured, **longitudinal patient data** available across systems, subject to consent. It lays the foundation for building AI solutions that can operate on integrated datasets rather than isolated silos, improving both the quality and continuity of care.

Why It Was Introduced?

India's healthcare data has traditionally been fragmented across multiple **providers, formats, and systems**, with patient records often stored in silos across hospitals, clinics, labs, and digital platforms. This fragmentation makes it difficult to access complete patient histories, leading to repeated tests, gaps in clinical understanding, and limited ability to generate meaningful insights for care or analytics. Variations in data formats and lack of standardization further reduce the usability of this data for digital systems, including AI.

ABDM was introduced to address these challenges by creating a more connected and standardized health data ecosystem. It aims to:

- Enable **seamless and secure data exchange** across healthcare providers, allowing information to move across systems when needed
- **Standardize how health data is structured and shared**, improving interoperability and making data usable across platforms
- **Improve continuity and coordination of care**, ensuring that patient information is available across different points of care for better clinical decision-making

AI needs data, but healthcare needs connected data.

ABDM → Key Focus Areas

ABDM shifts attention to how data moves across the healthcare system, emphasizing that value is created not just by collecting data, but by enabling it to flow securely and meaningfully between different actors. In a system where patient information is often spread across multiple providers, the ability to access connected and longitudinal data becomes critical for both care delivery and digital innovation.

By creating a structured environment for data exchange, ABDM highlights the importance of **interoperability, standardization, and controlled access**. It enables different parts of the healthcare ecosystem, hospitals, labs, digital platforms, and insurers to operate on shared information while maintaining clear boundaries around consent and usage. This not only improves continuity of care but also allows technologies like AI to work on more complete and context-rich datasets.

Interoperability



Ensuring different healthcare systems can exchange and interpret data seamlessly using common standards. This allows information from one provider to be understood and used by another without loss of meaning or context.

Consent-Based Data Sharing



Allowing patients to control when and how their data is shared across providers and platforms. This ensures that data flow is not automatic but governed by explicit permissions, maintaining trust while enabling access.

Health Data Standards



Structuring data in standardized formats to enable consistency and usability across systems. Standardization reduces fragmentation and ensures that data can be reliably used for clinical decisions and digital applications.

Unique Health Identity (ABHA)



Creating a unified identifier to link patient records across multiple touchpoints. This enables different healthcare interactions to be connected into a single, longitudinal view of the patient.

Ecosystem Integration



Connecting hospitals, labs, insurers, and digital platforms into a unified health data network. This enables coordinated care, reduces duplication, and supports more efficient healthcare delivery at scale.

ABDM Stakeholders Mapping

ABDM influences how data is accessed and shared by creating a **structured pathway for information** to move across different parts of the healthcare ecosystem. Instead of data being confined within individual providers, it enables controlled access across systems, allowing stakeholders to interact with patient information in a more coordinated manner. This changes not just availability of data, but how it is used across care journeys.

Its impact becomes most visible in how different stakeholders connect through shared data whether it is hospitals accessing records from other providers, clinicians viewing more complete patient histories, or digital platforms integrating information from multiple sources. By enabling this continuity, ABDM supports more **informed decision-making, reduces duplication**, and allows healthcare delivery to become more connected and efficient across settings.

Stakeholder	Design Expectations	Operational Changes	Risk & Accountability	Adoption & Trust
AI Builders (India)	●		●	
Global AI Companies	●		●	
Hospitals & Systems	●			
Clinical Workforce	—		●	
Patients & Public		●		



High Impact



Moderate Impact



Limited / indirect Impact

ABDM Across Healthcare Stakeholders

For Hospitals & Health Systems

Hospitals benefit from improved data accessibility and continuity, enabling better coordination of care across providers. ABDM also requires them to adopt standardized systems and integrate with national health data infrastructure.

For Clinical Workforce

ABDM allows providers to access more complete patient histories across different care settings, supporting more informed decision-making. It also changes how data is accessed during care delivery.

For Patients & Public

ABDM gives patients greater control over their health data while enabling easier sharing across providers. This improves continuity of care and reduces the need to repeatedly provide the same information.

For AI Builders (India)

ABDM enables access to more structured and interconnected health data, allowing developers to build AI systems on richer datasets. This supports better model performance and more context-aware solutions across the care continuum.

For Global AI Companies

ABDM creates a standardized environment for accessing and integrating health data, making it easier for global companies to plug into India's healthcare ecosystem. However, it also requires alignment with local data standards and consent mechanisms.

ABDM transforms how healthcare data is used by enabling it to move across systems in a structured and controlled manner, rather than remaining confined within individual providers. It shifts the focus from isolated **data silos to connected data** ecosystems, where information can be accessed across different points of care when needed. This allows for more complete patient views, reduces duplication of tests and records, and supports continuity across care journeys.

For stakeholders, this means adapting to standardized systems, integrating with shared infrastructure, and building workflows that can operate on connected data. It also requires aligning with consent-driven data access and ensuring that data is used appropriately across interactions. Over time, this creates a more coordinated and efficient healthcare system, where decisions are informed by richer, more complete information rather than fragmented inputs.

If DPDP defines how data is protected → ABDM defines how data flows securely and meaningfully across the healthcare ecosystem.



ICMR → The Ethics Layer

Guiding how AI is used responsibly in patient care.

What is ICMR?

The Indian Council of Medical Research (ICMR) Ethical Guidelines for AI in Healthcare provide a structured framework to ensure that AI systems are developed and used in ways that align with clinical ethics and patient well-being. Unlike regulatory or data-focused frameworks, ICMR focuses on how AI behaves in real care settings, where decisions affect **diagnosis, treatment, and patient outcomes**.

In practice, these guidelines direct on, how AI tools should be introduced into clinical workflows, how their outputs should be interpreted, and how risks should be managed when systems influence care decisions. They address key concerns such as whether AI recommendations are understandable, whether they perform fairly across patient groups, and who remains accountable when outcomes are affected. The emphasis is not just on what AI can do, but on how it should be used responsibly alongside human judgment.

Why It Was Introduced?

As AI systems began to move closer to clinical decision-making, concerns emerged around bias, lack of transparency, and unclear responsibility for outcomes. In real-world settings, clinicians may rely on AI outputs without fully understanding how they are generated, and patients may be affected by decisions shaped by algorithms they are unaware of. ICMR introduced these guidelines to address these gaps by:

- Ensuring AI supports clinical decision-making without replacing professional judgment
- Reducing **risks of bias** and **unintended harm** in patient care
- Bringing clarity on responsibility when AI influences outcomes
- Providing practical direction on how AI should be used within care settings

AI can assist care, but ethics must guide it.

A Landmark Move Toward India-Centric Healthcare Evidence

Recent direction from the Indian Council of Medical Research (ICMR) highlights a long-standing gap in how clinical evidence has been generated and applied in India. For decades, a significant portion of global medical research has been based on Western populations, with differences in genetics, diet, environment, and disease patterns often underrepresented. As a result, applying such evidence directly to Indian populations may not always lead to optimal or accurate clinical outcomes.

In its latest guidance, ICMR has more clearly emphasized the need for India-centric research, underscoring that medical evidence and systems must be grounded in data that reflects Indian populations and healthcare realities. This has important implications for how clinical guidelines, treatment protocols, and AI models are developed and validated. For healthcare AI in particular, it reinforces the need for training and testing systems on data that reflects Indian patient populations, rather than relying primarily on imported datasets or assumptions.


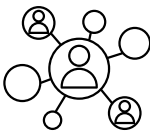
This direction aligns closely with broader themes of fairness, reliability, and real-world applicability. It highlights that responsible AI in healthcare is not just about technical performance, but about ensuring that systems are relevant, accurate, and appropriate for the populations they serve.

ONE POPULATION'S DATA CANNOT DEFINE ANOTHER'S CARE.

ICMR → Key Focus Areas

ICMR shapes how ethical considerations are embedded across the healthcare ecosystem by guiding how AI systems are designed, introduced, and used in real clinical environments. It brings attention to the human and decision-making aspects of AI, ensuring that technology aligns with established principles of patient care rather than operating in isolation.

Its impact is most visible in how responsibility is shared, how fairness is maintained across diverse patient groups, and how **patient interests remain central to care decisions**. This influences how stakeholders evaluate AI systems, how they are used within workflows, and how outcomes are interpreted, ensuring that ethical considerations remain an active part of AI use, not an afterthought.


Patient Safety		Ensuring AI systems do not introduce harm and are reliable in clinical use. This includes validating performance in real-world settings, monitoring outcomes over time, and ensuring that errors or limitations do not negatively impact patient care.
Fairness & Equity		Reducing bias and ensuring equitable outcomes across populations. In a diverse country like India, this means ensuring models perform consistently across different regions, socio-economic groups, and demographic profiles.
Transparency		Making AI decisions understandable to clinicians and patients where possible. This helps users interpret recommendations, build confidence in the system, and make informed decisions rather than relying on outputs blindly.
Accountability		Defining responsibility when AI influences clinical decisions. This ensures that there is clarity on who is responsible for outcomes, whether it is developers, deploying institutions, or healthcare providers.
Human Oversight		Maintaining clinician involvement in decision-making processes. AI is positioned as a support tool, with final decisions resting with healthcare professionals who can interpret context and exercise judgment.

ICMR Stakeholders Mapping

ICMR shapes how ethical considerations are applied across the healthcare ecosystem by embedding them into the way AI systems are designed, evaluated, and used in real clinical settings. It shifts attention to the human impact of AI, ensuring that decisions influenced by these systems remain aligned with patient well-being and established standards of care. This becomes particularly important in healthcare environments where AI is integrated into diagnosis, treatment planning, and care delivery.

Its impact is most visible in how different stakeholders interpret and act on AI outputs, how responsibility is managed when systems influence outcomes, and how fairness is maintained across diverse patient populations. It also reinforces the need to keep **patient interests central**, ensuring that technology supports care rather than overriding clinical judgment. In this way, ethical considerations become part of everyday practice, guiding how AI is trusted and used across the system.

Stakeholder	Design Expectations	Operational Changes	Risk & Accountability	Adoption & Trust
AI Builders (India)		●		
Global AI Companies		●		
Hospitals & Systems	●			
Clinical Workforce	—			
Patients & Public		●		

-  High Impact
-  Moderate Impact
-  Limited / indirect Impact

ICMR Across Healthcare Stakeholders

For Hospitals & Health Systems

Hospitals are encouraged to evaluate AI systems not only for performance but also for ethical implications, including fairness and patient safety. This adds an additional layer of consideration in adoption decisions.

For Clinical Workforce

ICMR reinforces the role of clinicians in overseeing AI use, ensuring that technology supports rather than replaces clinical judgment. It emphasizes the need to interpret AI outputs responsibly.

For Patients & Public

ICMR promotes patient-centered AI, ensuring systems are fair, transparent, and do not disadvantage specific groups. It strengthens trust by emphasizing ethical use of technology in care.

For AI Builders (India)

ICMR encourages developers to integrate ethical considerations into system design, including fairness, transparency, and safety. It requires thinking beyond functionality to consider how AI impacts patient outcomes and clinical decisions.

For Global AI Companies

ICMR signals that ethical expectations in India must be addressed alongside technical performance. Companies entering India are expected to ensure their systems align with local ethical standards and healthcare priorities.

ICMR ensures that healthcare AI operates within the boundaries of clinical ethics by grounding its use in real-world patient care and decision-making. It complements technical and regulatory frameworks by focusing on how **AI is applied in practice**, where outcomes directly affect patient safety, treatment choices, and trust in the system. This shifts attention from what AI can do to how it should be used, interpreted, and managed within clinical workflows.

For stakeholders, this means actively **incorporating ethical considerations** at every stage, from how systems are designed and evaluated to how they are used during care delivery. It involves being aware of potential biases, understanding system limitations, and ensuring that patient interests remain central to all decisions. Over time, this creates a more responsible and balanced approach to AI adoption, where technology supports safe, fair, and patient-centered healthcare rather than operating independently of it.

If SAHI defines how to build → BODH defines how to validate → DPDP defines how to protect → ABDM defines how data flows → ICMR defines how AI should be used responsibly.



CDSCO → The Regulatory Layer

Defining what AI systems are approved for clinical use.

What is CDSCO?

The Central Drugs Standard Control Organization (CDSCO) is India's national regulatory authority responsible for approving and overseeing medical devices, including certain categories of AI-based healthcare solutions. As AI systems increasingly move into areas like diagnostics, imaging, and clinical decision support, many of these tools fall under the definition of software as a medical device (SaMD), bringing them within CDSCO's regulatory scope.

In practice, CDSCO defines the conditions under which healthcare AI systems can be deployed in clinical settings. It establishes requirements around **safety, performance, and validation**, ensuring that systems meet minimum standards before they are used on patients. Unlike frameworks that guide design or ethics, CDSCO acts as a **gatekeeper**, determining whether an AI solution is ready for real-world clinical use.

Why It Was Introduced?

As AI systems began to move closer to clinical decision-making, supporting diagnosis, risk assessment, and treatment planning, the need for formal oversight became increasingly important. In real healthcare settings, these tools are not just experimental; they directly influence patient outcomes, clinical workflows, and trust in care delivery. Without clear **regulatory checks**, there is a risk of deploying systems that perform well in controlled environments but fail under real-world conditions. In this context, **innovation alone is not enough**. AI systems must demonstrate safety, reliability, and consistency before they are used in patient care. This is where CDSCO plays a critical role by introducing a structured process to evaluate and approve such systems, ensuring that only those meeting defined standards are allowed into clinical use. CDSCO's role becomes important to:






- Ensure AI systems meet **clinical safety and performance standards**, particularly in high-impact use cases like diagnostics and treatment support
- Prevent deployment of unvalidated or unreliable tools, reducing the risk of harm from systems that have not been adequately tested
- Define **accountability** for products used in patient care, clarifying responsibility when AI systems influence outcomes
- Create a **structured pathway** for approval and market entry, enabling developers to navigate regulatory requirements before deployment

From innovation to approval → CDSCO makes AI clinically usable.

CDSCO → Key Focus Areas

CDSCO brings attention to the final stage of the AI lifecycle → where systems transition from development and testing into **approved clinical use**. This is the point where experimental models become real-world tools, interacting directly with patients and influencing care decisions. At this stage, the focus shifts from technical performance to demonstrated **safety, reliability, and suitability for clinical environments**.

It highlights that before AI can be used in real healthcare settings, it must pass through defined **regulatory checks** that assess not only how well the system performs, but how consistently it behaves across different conditions and use cases. These checks ensure that systems are appropriately validated, risks are understood, and safeguards are in place for patient-facing deployment. In doing so, **CDSCO acts as a critical gatekeeper**, ensuring that only those AI systems that meet required standards are introduced into clinical practice.

Regulatory Classification		Determining whether an AI system qualifies as a medical device based on its function and risk level. This defines the level of scrutiny and approval required. It also influences the documentation, testing, and compliance steps needed before the system can be deployed.
Clinical Validation Requirements		Ensuring that AI systems demonstrate safety and effectiveness through appropriate validation, including clinical evidence where required. This helps confirm that performance claims translate into real-world clinical reliability.
Approval & Certification		Establishing formal processes for reviewing and approving AI-based medical devices before they enter the market or clinical use. This ensures that only systems meeting defined regulatory standards are allowed for patient-facing deployment.
Post-Market Surveillance		Monitoring how systems perform after deployment, including tracking adverse events and ensuring continued compliance. This allows regulators and organizations to identify issues early and take corrective action if needed.
Risk-Based Oversight		Applying stricter regulatory requirements for higher-risk AI systems, particularly those involved in diagnosis or treatment decisions. This ensures that systems with greater clinical impact are subject to more rigorous evaluation and control.

CDSCO Stakeholders Mapping

CDSCO influences how AI systems are brought into real clinical use by establishing **clear regulatory expectations** across the healthcare ecosystem. Its impact is most visible at the point where systems transition from being developed and tested to being approved for patient-facing deployment. This creates a structured environment in which stakeholders must align their actions with defined standards for **safety, validation, and compliance**.

Unlike earlier layers that guide design or enable data flow, ***CDSCO directly shapes what can and cannot be used in practice***. It affects how developers prepare their products for approval, how hospitals select and deploy AI tools, and how confidence is built among users and patients. By defining approval pathways and accountability, CDSCO ensures that AI adoption is not only innovative but also controlled, reliable, and aligned with clinical safety requirements. Its impact varies across stakeholders based on their role in developing, approving, deploying, or using AI systems.

Stakeholder	Design Expectations	Operational Changes	Risk & Accountability	Adoption & Trust
AI Builders (India)				
Global AI Companies				
Hospitals & Systems	●			
Clinical Workforce	—	●		
Patients & Public	—	—		

-  High Impact
-  Moderate Impact
-  Limited / indirect Impact

CDSCO Across Healthcare Stakeholders

For Hospitals & Health Systems

Hospitals must ensure that AI tools used in clinical settings are approved or compliant with regulatory requirements. This shifts adoption decisions toward verified and certified systems rather than experimental or unregulated tools.

For Clinical Workforce

CDSCO supports the use of AI systems that have been formally reviewed and approved, providing greater assurance of safety and reliability. This helps clinicians engage with AI tools that meet established standards.

For Patients & Public

Patients benefit from safeguards that ensure AI systems used in healthcare have undergone regulatory scrutiny. This reduces the risk of harm and builds confidence in the safety of AI-driven care.

For AI Builders (India)

CDSCO requires developers to design AI systems that meet regulatory standards from the outset, including **validation, documentation, and risk assessment**. It introduces the need to align product development with approval pathways, ensuring that systems are not only functional but also compliant.

For Global AI Companies

CDSCO sets clear expectations for market entry in India, requiring global companies to meet local regulatory standards. This may involve additional validation, documentation, and alignment with India-specific requirements before deployment.

CDSCO introduces a critical layer of oversight by ensuring that healthcare AI systems are formally reviewed and approved before they are used in clinical environments. This shifts the emphasis from building innovative solutions to demonstrating that those solutions are **safe, reliable, and suitable** for real-world patient care. It requires systems to meet defined standards not just in performance, but also in how they behave under different clinical conditions.

For stakeholders, this translates into integrating regulatory considerations throughout the lifecycle → from early design and validation to deployment and ongoing use. Developers must prepare for approval processes, hospitals must adopt compliant and certified tools, and providers must rely on systems that meet established standards. This makes regulatory alignment a built-in part of how AI systems are brought into practice, rather than a final step.

Over time, this creates a more structured and predictable environment for healthcare AI adoption, where innovation is supported by clear safeguards. It enables stakeholders to operate with greater confidence, knowing that systems in use have been assessed against consistent criteria for safety, performance, and clinical relevance.

If SAHI defines how to build → BODH defines how to validate → DPDP defines how to protect → ABDM defines how data flows → ICMR defines how to use responsibly → CDSCO defines what can be used.

Putting It All Together...

India's approach to AI in healthcare is not defined by a **single regulation**, but by a set of interconnected layers that together shape how AI moves from concept to clinical use. Each layer addresses a different part of the lifecycle → how systems are designed and evaluated, how data is accessed and exchanged, how ethical considerations are applied, and how final approval and oversight are ensured. When viewed together, these layers create a more complete picture of how AI is expected to function within the healthcare system.

Rather than operating in isolation, these elements **reinforce one another**. Data protection enables responsible access, data exchange enables continuity, validation ensures reliability, ethics guide usage, and regulation ensures readiness for real-world deployment. *This interconnected structure reduces fragmentation and brings greater clarity to how AI systems should be built and used in practice.*

What emerges is not a restrictive environment, but a more structured and predictable one, where innovation is expected to align with real-world conditions, patient needs, and system-level accountability. For stakeholders, this means moving beyond a purely technology-driven approach and engaging with AI as part of a larger healthcare ecosystem that requires coordination, standardization, and shared responsibility.

What This Means Going Forward?

- Healthcare AI in India is moving from experimentation to implementation.
- From isolated solutions to connected systems.
- From capability to accountability.

Understanding these layers is essential, not just to stay compliant, but to build, adopt, and use AI effectively in a rapidly evolving landscape. Lastly, while these rules and frameworks set the structure, their implementation in real-world healthcare settings remains uneven and continues to evolve.

This guide is a high-level overview of the key rules and frameworks shaping healthcare AI in India. It is intended as a starting point → a way to bring structure and clarity to a complex and evolving space.

Authors:



Dr Monika Sonu

Co-Founder Healthinnovation Toolbox



Kingshuk Chakraborty

Co-Founder Healthinnovation Toolbox



Citation

Sonu, M., & Chakraborty, K. (2026). AI Rules in India → A Healthcare Perspective (Version 01). Zenodo.

About

Healthinnovation Toolbox is a **Product Engineering Company**. With deep expertise in product engineering, we guide our partners through every phase of their digital journey - from assessing core processes, validating ideas, and engineering products to streamlining development, scaling solutions, and expediting market delivery. We help health systems design scalable AI operations, governance, and impact frameworks.

Healthinnovation Toolbox is a **Digital Machina Company**

Copyright © 2025 Health Innovation Toolbox (Digital Machina). This work is licensed under the Creative Commons Attribution 4.0 International (CC BY 4.0) license.

Citation

Sonu, M., & Chakraborty, K. (2026). *AI Rules in India → A Healthcare Perspective (Version 01)*. Zenodo.

Healthinnovation Toolbox & Digital Machina, empowering knowledge sharing and open innovation.



Disclaimer

This guide is intended solely for educational and informational purposes. The content presented does not constitute legal advice, regulatory guidance, clinical direction, or professional consultation of any kind. While efforts have been made to ensure clarity and relevance, the information may not reflect the most current legal interpretations, regulatory updates, or implementation practices across different healthcare settings.

The frameworks, laws, and initiatives described in this guide are presented at a high level to provide structural understanding of India's healthcare AI landscape. This guide does not evaluate, interpret, or provide compliance recommendations for any specific regulation, nor should it be used as a substitute for formal legal, regulatory, or policy advice.

Readers are strongly advised to consult qualified professionals in legal, regulatory, clinical, and data governance domains before applying any concepts discussed in this guide to real-world systems, products, or healthcare operations. The authors and affiliated organizations assume no responsibility or liability for any outcomes resulting from the use or interpretation of this material.

Any references to policies, institutions, or frameworks are for informational context only and do not imply endorsement or official representation. This guide is intended to support independent understanding, critical thinking, and responsible engagement with the evolving AI ecosystem in healthcare.

Use of this document implies acceptance of this disclaimer.