

सर्वजन हिताय | सर्वजन सुखाय
WELFARE FOR ALL | HAPPINESS OF ALL



INDIA AI IMPACT SUMMIT 2026

COMPENDIUM

Real-World Impact of AI in Health

Disclaimer

The case studies included in this compendium have been evaluated based on information submitted by the respective authors and participating organizations. Responsibility for the accuracy of data, metrics, and representations rests solely with the submitting authors. The evaluation committee and partner institutions shall not be held liable for any discrepancies, omissions, or subsequent changes in the information provided.

Partners: IndiaAI Mission, Ministry of Electronics and Information Technology (MeitY), Government of India, World Health Organization (WHO), Ministry of Health and Family Welfare (MoH&FW), Government of India and the Indian Council for Medical Research (ICMR), Government of India.



Table of Content

Foreword	I
List of Abbreviations	X
Introduction	
• Scaling AI for Health: Initiatives of the Ministry of Health & Family Welfare (Government of India)	01
• Applied AI in Health: From Innovation to Responsible Implementation (World Health Organization)	03
AI for Computer Vision in Health	
1. Bridging the Neuroradiology Divide: Deploying Specialized, Indigenous AI to Augment Diagnostic Expertise in India	09
2. AI-Assisted Visual Inspection for Cervical Cancer Screening in Rwanda	21
3. Development of a Comprehensive Deep Learning Neural Network Algorithm for Analysis of Gram-Stained Sputum Smears for Adequacy	31
4. Harnessing AI for Cervical Cancer Surveillance and Clinical Decision Support in Kenya	41
5. High-Accuracy ECG Image Interpretation using Parameter-Efficient LoRA Fine-Tuning with Multimodal LLaMA 3.2	49
6. Leading with AI: Kerala's Nayanamritham 2.0 – A Government-Led AI-Assisted Chronic Eye Disease Screening Initiative	57
7. Leveraging Detection and Classification of Alzheimer's Disease and Brain Tumour Detection using Transformer Models	69
8. LIRA: An AI-Powered Multi-modality Radiology Platform Enhancing Diagnostic Equity in Tuberculosis Screening in India	83
9. Machine Learning-assisted Digital Twin for Diabetic Retinopathy	93
10. AI-Supported Neuroimaging: Artificial Intelligence in the Triage of Acute Stroke and Head Injury	105
11. Usability and Efficacy of Swaasa An AI-Platform for Rapid Respiratory Health Assessment using Cough as a marker	119



AI for Clinical Decision Support

12. Blueprint for getting AI ready data from Radiation Oncology: Insights from a large Indian tertiary care setting **129**
13. From Static Portal to Smart Assistant: Reimagining NTD Program Support with AI **139**
14. Intelligent General Practitioner (iGP), an AI-powered, CHWs-driven, People-Centered Model to Revolutionize Primary Healthcare Delivery in LMICs towards UHC **149**
15. Supporting Antibiotic Stewardship Through an AI-Driven Digital Intervention to Combat Antimicrobial Resistance **159**

AI for Predictive Analysis in Health

16. AI-Driven Environmental Health Risk Prediction in Urban Morocco: The OCT7 S1 Framework Meknès **171**
17. "AyurVAID D-RISK": Advanced, Non-Invasive, Self-Administered, H2O AutoML Model for Detection of Undiagnosed Diabetes **179**
18. DEEP-DEPRESSION: An AI-Powered EEG-Guided Brain-Computer Interface (BCI) Enabled Solution for Objective Depression Diagnosis to Strengthen Mental Health Care in Low-Resource Settings **189**
19. Spatiotemporal Probabilistic Encoder-fed Climate-aware Transfer-learning-ready Risk-forecasting Adversarial (SPECTRA) Architecture for Climate-informed Infectious Disease **199**

Ambient AI for Health

20. Data privacy essentials for data de-identification in radiology images, a practical approach in AI research **209**
21. From Text to Insight: Harnessing the power of AI to transform Unstructured Data in Brazil's Health Information Systems **221**
22. HealthVaani: The Generative AI-powered Voice of Knowledge for India's Frontline Health Workers **231**

AI and Digital Twins in Health

23. Predictive Virtual Cardiac Twin: An AI-Powered Immersive Visualization Platform for Advanced Cardiac Care **241**

24. AIMS-Artificial Intelligence Medical Simulator	249
--	-----

Snapshots

1. Pre-generated Offline AI for Community Health Workers in Africa	256
2. AI-Enabled Diabetic Retinopathy Screening in India's Primary Healthcare System	258
3. AI Diagnostic Tool for Ear Pathology (Primary Care ENT Screening)	260
4. Swasth-Derma: AI-Powered Pre-Medical Dermatology Triage for Public Health	262
5. AI-Led Multilayered Cancer and Psycho-Oncology Pathway for India's High-Burden Cancers	264
6. Precision Beyond Genes with QPOP: An AI – Guided Oncology Approach	266
7. Blockchain-Enabled, AI-Driven EMR for Health Equity Across the Global South	268
8. Advanced Data Analytics for Public Health Action and Research Venture (ADARV) – A Public Health Data Warehouse	270
9. Moskeet: Climate-Responsive AI & IoT Surveillance for Vector-Borne Disease Control	272
10. AI-driven Decision Support for Ayurveda: A Retrieval-Augmented Generation (RAG) Approach	274
11. AIIMS Delhi Ransomware Attack: CyberPeace Foundation's Insights for Forging Secure and Resilient Digital Health Ecosystems in India	276
12. Digital Health Validitron: A unified service for de-risking clinical AI solutions	278

Foreword 1

The India AI Impact Summit 2026 is being convened at a pivotal moment when global discourse on Artificial Intelligence (AI) is transitioning from principles and commitments to demonstrable outcomes and real-world impact. In this context, it is essential to highlight the role of sectoral AI solutions in advancing socio-economic development in India and across the world.

Over the past three years, the international community has come together through a series of landmark summits to shape the global AI agenda. The Bletchley Park AI Safety Summit (November 2023) focussed on identifying AI safety risks and promoting risk-based policy frameworks. The Seoul AI Summit (May 2024) expanded this agenda by emphasizing innovation and inclusivity. Subsequently, the Paris AI Action Summit (February 2025), co-chaired by France and India, brought together over 1,000 participants from more than 100 countries, broadening discussions to encompass economic competitiveness, environmental sustainability, and practical applications of AI. With India hosting the AI Impact Summit in February 2026, the global conversation advances further towards measurable impact.


India continues to advocate for the democratization of AI and equitable access to its benefits. Our approach is guided by the principles of safety, security, trustworthiness, and inclusivity across all sectors. At the national level, the IndiaAI Mission has made significant progress in enabling AI solutions in priority areas such as agriculture, healthcare, climate change, and security—domains that directly affect the



well-being of millions of citizens. In healthcare, focused efforts are underway in collaboration with leading institutions to develop high-quality datasets and deploy AI-enabled diagnostics, health-risk forecasting, and clinical decision support systems.

The Casebook on Real-World Impact of AI in Health, developed in partnership with the World Health Organization (WHO), represents a key milestone in these efforts. This casebook serves as a comprehensive reference for policymakers, innovators, and researchers, with the objective of promoting responsible and ethical adoption of AI, enabling replication and scaling of proven solutions across the Global South, and supporting digital transformation of healthcare systems. It stands as an example of how international collaboration can deliver practical and scalable outcomes for developing economies.

As India prepares to host the AI Impact Summit 2026, the Ministry of Electronics and Information Technology reaffirms its



commitment to fostering an AI ecosystem that is inclusive, ethical, and firmly grounded in real-world impact. We acknowledge with appreciation the partnership of the World Health Organization, and the leadership of the Ministry of Health and Family Welfare and the Indian Council of Medical Research in advancing responsible AI in healthcare. We also recognize the contributions of

the innovator and research communities in developing AI-based solutions for public good.

Together, these collective efforts demonstrate that artificial intelligence can serve as a powerful instrument for equitable development—transforming lives while upholding the highest standards of safety, transparency, and human dignity.

Shri S. Krishnan

Secretary
Ministry of Electronics and Information Technology
Government of India

Foreword 2


The India AI Impact Summit 2026—the fourth in the series of international AI summits and the first one in Global South—marks an important milestone in the evolution of collective efforts to translate the potential of AI into demonstrable and measurable outcomes.

The global AI agenda has steadily progressed towards real-world implementation of AI solutions across critical sectors including healthcare. AI innovations in areas such as diabetic retinopathy detection, tuberculosis screening, and maternal health monitoring demonstrate scalable and cost-effective approaches relevant to countries facing public health and resource constraints. The IndiaAI Mission's emphasis on democratised access to AI resources and collaborative innovation reflects the conviction that solutions developed in one context can be adapted and scaled across regions and nations.

The Casebook on Real-World Impact of AI in Health, developed in partnership with the World Health Organisation (WHO), represents the realisation of this collaborative vision. This joint initiative invited contributions from researchers, innovators and institutions worldwide, resulting in an enthusiastic response of more than 200 abstracts from diverse geographies and healthcare contexts. The submissions underwent a multi-stage technical evaluation process. Selected



contributors submitted full chapters detailing solution design, deployment models, impact, ethical considerations and lessons learned. Each chapter was reviewed by an expert Evaluation Committee with representatives from the Ministry of Health and Family Welfare, Government of India, the Indian Council of Medical Research and the WHO, to secure the inclusion of most impactful, evidence-based and replicable solutions. The resulting Casebook presents a curated set of 24 AI health applications, ranging from AI-enabled medical devices to innovations in early detection and classification of neurological and cardiological diseases. We extend sincere appreciation to the Ministry of Health and Family Welfare and the Indian Council of Medical Research for their leadership in advancing responsible AI in healthcare, and to the World Health Organisation for its steadfast partnership in the development of this Casebook.



Together, these collective efforts reaffirm that AI, when developed and deployed responsibly, ethically and inclusively, can serve as a powerful instrument for strengthening health systems and advancing global health equity.

Shri Abhishek Singh

CEO, IndiaAI Mission
Additional Secretary, Ministry of Electronics
and Information Technology
Government of India

Foreword 3

India has made significant gains in public health indicators over the past decade. However, delivering high-quality healthcare to vulnerable populations, particularly those in remote rural areas, urban slums, and marginalized communities, continues to require scalable, high impact, and cost-effective innovations. Artificial intelligence (AI) and digital health technologies offer promising pathways to strengthen reach, quality, and equity in healthcare delivery. Several AI-driven applications, including tools for screening cancers, tuberculosis, anemia, diabetic retinopathy, and cataracts, have been developed and are currently being piloted across the country. Yet, sustainable adoption and system-wide integration of these technologies demand careful attention to some foundational challenges.


Foremost among these is the development of India-specific, accurate AI algorithms, which depend on robust, secure, and representative health data systems that safeguard privacy while enabling responsible innovation. To address this need, the Indian Council of Medical Research (ICMR) has established the Health Research Data Repository, providing researchers with access to standardized, interoperable, and ethically governed health research data. Equally critical is rigorous evaluation and evidence generation to ensure that AI solutions demonstrate real-world effectiveness, safety, and equity. In this regard, ICMR's National Institute for



Research in Digital Health and Data Science (NIRDHDS) is collaborating with clinical, technical, and governmental partners to build a network for systematic technical, clinical, economic, and implementation evaluations of AI health technologies. ICMR has further strengthened the translational research ecosystem through initiatives such as MedTech Mitra and Patent Mitra, which support innovators from validation to regulatory and intellectual property guidance.

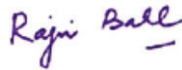
Beyond infrastructure and evidence generation, capacity building and awareness are equally essential. Healthcare providers, frontline workers, data scientists, and policymakers must be equipped to responsibly adopt and manage AI systems.

In conclusion, building a healthier India through AI demands collective action across sectors. Academia must generate



evidence, industry drive innovation, government ensure ethical governance, and communities guide equitable implementation. With its scientific expertise, strong policy vision, and initiatives such

as the IndiaAI Mission, India is well positioned to lead globally in transforming healthcare through efficient, equitable, and responsible adoption of AI and digital health technologies.



Dr. Rajiv Baahl

Director General
Indian Council of Medical Resear

Foreword 4

Artificial intelligence is rapidly reshaping how health systems synthesize knowledge, optimize services, and support decision-making. Across the world, AI is moving from promise to practice, helping detect disease earlier, extend the capacity of health workers, strengthen surveillance, and build more responsive and resilient systems. The question is no longer whether AI will influence the future of health, but how we ensure it does so responsibly, ethically, and equitably.

This compendium arrives at an important moment. Countries are actively translating AI's potential into real-world impact while protecting safety, ethics, and public trust. WHO has consistently emphasized that the value of AI in health lies not in technology alone, but in its integration into health systems, supported by strong evidence, rigorous evaluation, and alignment with public health priorities.


Through its normative mandate, WHO is working with Member States and partners to strengthen the foundations for trustworthy AI in health. This includes advancing guidance on ethics and governance, supporting regulatory readiness, strengthening evidence generation and evaluation, and promoting implementation approaches that prioritize safety, transparency, and accountability across the AI lifecycle. Central to this effort is the Global Initiative on Artificial Intelligence for Health, co-led with the International Telecommunication Union (ITU) and the World Intellectual Property Organization (WIPO), which supports countries in building the policy, regulatory, and technical ecosystems needed for responsible AI adoption.



Developed in collaboration with IndiaAI, the Ministry of Electronics and Information Technology, the Ministry of Health and Family Welfare, and the Indian Council of Medical Research, this compendium brings together real-world examples of AI designed, deployed, and evaluated across diverse health domains. It highlights the growing leadership of countries in the Global South in shaping the global AI-for-health landscape and advancing people-centred innovation.

Across clinical care, public health, workforce support, and data systems, the cases illustrate how AI is being used to augment human expertise, expand access, and address persistent system bottlenecks. Many solutions are embedded in real workflows and guided by human oversight, illustrating AI's potential as an enabler of Universal Health Coverage.

At the same time, the compendium underscores the importance of lifecycle monitoring, interoperability, regulatory



clarity, robust governance, and sustainable financing to move from pilots to scalable impact. This publication is more than a collection of case studies—it is a contribution to the global dialogue on responsible AI for health and a step

toward bridging innovation and responsible scale. On behalf of WHO, I extend my sincere appreciation to all partners and contributors whose ongoing efforts help ensure that the benefits of AI translate into better health for all.

Dr Alain Labrique

Director, Department of Data, Digital Health, Analytics and AI Health Systems and Data Division
World Health Organization





List of Abbreviations

AI	- Artificial Intelligence	EMR	- Electronic Medical Record
AIMS	- Artificial Intelligence Medical Simulator	FHIR	- Fast Healthcare Interoperability Resources
AMR	- Antimicrobial Resistance	FLW	- Frontline Health Worker
API	- Application Programming Interface	GPU	- Graphics Processing Unit
BCI	- Brain-Computer Interface	HIS	- Health Information System
BERT	- Bidirectional Encoder Representations from Transformers	HL7	- Health Level Seven (interoperability standard)
CDSS	- Clinical Decision Support System	LLM	- Large Language Model
CHW	- Community Health Worker	ML	- Machine Learning
CNN	- Convolutional Neural Network	MMR	- Maternal Mortality Ratio
CV	- Computer Vision	NLP	- Natural Language Processing
DICOM	- Digital Imaging and Communications in Medicine	NTD	- Neglected Tropical Disease
DR	- Diabetic Retinopathy	OCR	- Optical Character Recognition
DSS	- Decision Support System	PHC	- Primary Health Care
ECG	- Electrocardiogram / Electrocardiography	RAG	- Retrieval-Augmented Generation
EEG	- Electroencephalogram / Electroencephalography	SaMD	- Software as a Medical Device
EHR	- Electronic Health Record	TB	- Tuberculosis
		UHC	- Universal Health Coverage
		VIA	- Visual Inspection with Acetic Acid



Scaling AI for Health: Initiatives of the Ministry of Health and Family Welfare, Government of India*

The Ministry of Health and Family Welfare (MoH&FW), Government of India, is leveraging artificial intelligence (AI) to strengthen public health delivery and accelerate progress toward national health priorities. MoH&FW has designated AIIMS New Delhi, PGIMER Chandigarh, and AIIMS Rishikesh as Centres of Excellence (CoE) for Artificial Intelligence in Health, supporting the development, validation, and scale-up of AI-enabled solutions.

The Ministry has fostered a collaborative ecosystem across government and research institutions, working with the Ministry of Electronics and Information Technology (MeitY), the Ministry of Higher Education and national technical agencies like Central Tuberculosis Division, National Centre for Disease Control (NCDC), C-DAC Mohali, ICMR, Indian Institute of Science (IISc), and the National Health Systems Resource Centre (NHSRC). Through public-private partnerships, MoH&FW has also collaborated with Wadhvani AI to provide technical support for the development and deployment of AI-based tools.

AI portfolio under MoH&FW spans multiple stages of maturity: deployed, advanced (ready for deployment), and under development, across disease surveillance, clinical decision support, and national disease control programmes.


Deployed solutions include the AI-enabled Clinical Decision Support System (CDSS)

integrated with the national telemedicine platform eSanjeevani; Media Disease Surveillance (MDS) embedded within the Integrated Disease Surveillance Programme (IDSP); and multiple AI tools supporting tuberculosis elimination efforts, including Cough Against TB (CATB), Prediction of Adverse TB Outcomes (PATO), TB vulnerability mapping, and Line Probe Assay Automation (LPA). An AI-based diabetic retinopathy screening solution is at an advanced deployment stage, while additional AI tools in dermatology, hepatology, and oral cancer screening are under development.

The Clinical Decision Support System (CDSS) integrated into eSanjeevani enhances the quality and consistency of teleconsultations by structuring patient complaint capture and generating AI-supported differential diagnosis suggestions. Since its integration in April 2023, more than 289 million eSanjeevani consultations have benefited from standardized clinical data capture across Health and Wellness Centres.

Media Disease Surveillance (MDS) is an AI-driven event-based surveillance tool embedded in the Integrated Health Information Platform (IHIP) under IDSP at NCDC. Operational since April 2022, it scans digital news sources in 13 Indian languages and shares validated alerts with state and district authorities for early response. The system has generated more than 7,000 alerts, supporting earlier outbreak detection and mitigation.

*Based on the inputs provided by the Ministry of Health and Family Welfare, Government of India



Under the National Tuberculosis Elimination Programme, 'Cough Against TB' (CATB) supports community-based pulmonary TB screening and has demonstrated an additional detection yield of approximately 12–16% compared with conventional screening approaches, with more than 175,000 individuals screened. TB vulnerability mapping uses AI to analyze geographic and health indicators to identify high-risk areas and has been deployed nationwide through integration with the Ni-Kshay portal.

Prediction of Adverse TB Outcomes (PATO) identifies patients at elevated risk of treatment loss-to-follow-up or mortality. The tool is deployed nationwide and has analyzed over 2.2 million TB patient records, flagging approximately 750,000 patients as higher risk to enable targeted follow-up and support.

Line Probe Assay Automation (LPA) applies computer vision to automate drug-resistant TB test result interpretation, reducing human error and improving standardization. The solution is integrated with the Ni-Kshay portal and is being rolled out across culture and drug susceptibility testing laboratories.

The AI-based diabetic retinopathy screening solution MadhuNetrAI enables non-specialist health workers to conduct retinal screening using automated image analysis. The system classifies disease severity across standard grading categories and prioritizes urgent referrals. It has been implemented in 41 facilities across 11 states and has supported the screening of more than 16,000 retinal images, benefiting over 9,000 patients.

Applied Artificial Intelligence in Health: From Innovation to Responsible Implementation

The Growing Opportunity for Artificial Intelligence in Health

Public health has delivered some of humanity's most significant collective achievements. The eradication of smallpox and the near-eradication of polio demonstrate the transformative impact of sustained global cooperation, scientific innovation, and strong health systems. Yet, despite these historic successes, health systems today face a new and complex convergence of pressures, including workforce shortages, rising disease burden, increasing demand for services, and widening inequities in access to care. These challenges are particularly pronounced in low- and middle-income countries (LMICs), where resource constraints and infrastructure gaps continue to limit the reach of essential services.

Despite decades of progress, major gaps remain. More than half of the world's population still lacks access to essential health services, and around 2.1 billion people continue to experience financial hardship due to out-of-pocket health spending. At the same time, approximately 930 million people experience catastrophic health expenditure each year, and about 90 million are pushed into extreme poverty due to out-of-pocket health spending. These persistent inequities underscore the urgency of strengthening health systems to achieve universal health coverage by 2030.


Against this backdrop, recent advances in machine learning, computer vision, and

large-scale data processing have accelerated the development of AI systems capable of supporting diagnostics, public health surveillance, risk prediction, and clinical decision-making. Global evidence suggests that AI can improve screening and triage, enhance disease surveillance, support clinical workflows, and enable more efficient use of scarce resources. AI is therefore moving from a domain of experimental research toward operational implementation across health systems.

At the same time, the rapid expansion of AI in health has raised critical questions regarding safety, effectiveness, equity, governance, and accountability. WHO has emphasized that the benefits of AI in health can be fully realized when systems are developed and deployed in ways that are ethical, science-based, and aligned with public health priorities. The global health community is therefore moving toward a dual agenda: accelerating innovation while strengthening the governance, regulatory, and evaluation frameworks required for responsible adoption.

From Promise to Practice: Why This Compendium Matters

Over the past decade, AI has become one of the most prominent areas of innovation in global health. Governments, international organizations, academia, researchers, and industry actors are investing heavily in AI-driven solutions, and a growing body of literature highlights their potential to transform health



systems. Yet much of the global conversation continues to focus on future possibilities rather than real-world implementation, particularly in low- and middle-income countries (LMICs).

For health leaders and policymakers, a persistent gap remains between aspiration and action. While strategies, guidance documents, and pilot initiatives are rapidly emerging, there is still limited visibility into how AI is being successfully deployed within real health systems, under real resource constraints, and in ways that respond to local public health priorities.

This compendium responds directly to that gap.

Rather than presenting AI as a distant or experimental technology, the case studies presented here document how AI is already being applied across diverse settings to pressing health challenges. They demonstrate how AI can support health workers, strengthening surveillance, improving diagnostics, optimizing service delivery, and enabling more responsive and resilient health systems.

Importantly, the compendium centers perspectives from LMICs and the Global South. These contexts often face the greatest health system constraints, yet they also serve as powerful sites of innovation. Solutions developed in these contexts illustrate how AI can be designed for environments with limited infrastructure, constrained workforces, and diverse populations, conditions that increasingly characterize health systems worldwide.

The compendium also aims to move beyond technology alone. Successful implementation


depends not only on technical performance, but also on governance, regulation, workforce readiness, financing, and integration into existing systems. By documenting real-world experiences, this collection highlights both the opportunities and the challenges encountered when moving from pilot projects to sustained implementation. By sharing practical lessons, enabling peer learning, and amplifying examples from across LMICs, it aims to support the global effort to translate AI's promise into measurable health impact.

Rapid Global Engagement: Response to the Call for Abstracts

The global call for abstracts generated a strong and rapid response, reflecting the growing momentum of AI innovation in health. Within a two-week submission window, 213 abstracts were received from 16 countries, with the majority originating from low- and middle-income country contexts. The volume and diversity of submissions within a short timeframe underscores the accelerating pace of AI adoption and the strong interest among innovators, researchers, and public health practitioners in sharing real-world experiences.

The breadth of submissions spanned multiple disease areas, health system functions, and stages of implementation, from early pilots to scaled deployments. This diversity illustrates both the maturity and the dynamism of the emerging AI in health ecosystem across the Global South.

This response also highlights the clear demand for platforms that enable knowledge exchange, cross-country learning, and the dissemination of practical implementation highlights. By bridging



together contributions from diverse settings, this compendium provides a snapshot of a rapidly evolving field and reflects the growing global commitment to harnessing AI to address pressing public health challenges.

Evaluation and Selection Methodology

All eligible submissions underwent a structured, multi-stage evaluation process designed to ensure methodological rigor, relevance, and diversity of representation. Each abstract was independently reviewed by at least three experts from WHO's Department of Digital Health, Data, Analytics and AI, using a standardized scoring framework.

Evaluation criteria included:

- Health system relevance and impact
- Implementation maturity and real-world deployment
- Technical and methodological soundness
- Scalability potential
- Governance and safety safeguards
- Documentation transparency and clarity

Following abstract evaluation, shortlisted contributors were invited to develop full chapters. These submissions underwent a second round of review by an evaluation committee comprising experts from WHO, the Ministry of Health and Family Welfare (MoH&FW), and the Indian Council of Medical Research (ICMR).


Ultimately, 44 abstracts were shortlisted, of which 24 chapters were selected for inclusion in the final compendium. In addition, a number of promising submissions were invited to contribute to a complementary "Snapshot"

section designed to highlight emerging initiatives and early-stage innovations. A total of 24 snapshot contributions were invited, and 12 are featured in this section. This rigorous and collaborative selection process ensured a balanced portfolio of case studies spanning diverse AI modalities, health domains, geographic contexts, and levels of implementation maturity.

Structure of the Compendium and AI Domains

The selected chapters are organized into six functional AI domains that reflect major modalities currently shaping AI adoption in health systems:

- **AI for Computer Vision in Health-** applications that analyze medical images and visual data to support screening, diagnosis, and monitoring, such as radiology, pathology, and ophthalmology.
- **AI for Predictive Analytics in Health-** models that use historical and real-time data to forecast risks, detect patterns, and support population health planning and resource allocation.
- **AI for Clinical Decision Support-** tools that assist health professionals by providing evidence-based recommendations, triage support, and workflow optimization within clinical settings.
- **Ambient AI for Health-** technologies that operate in the background of care environments, including remote monitoring, speech recognition, and automated clinical documentation.
- **AI and Digital Twins in Health-** emerging approaches that use virtual representations of patients or health systems to simulate



disease progression, treatment response, and system performance.

- **AI-Enabled Robotics in Health-** robotic systems that support surgery, rehabilitation, logistics, and assistive care.

These domains provide practical framework for understanding how AI is being applied across different layers of health systems, from diagnostics and clinical workflows to population health, system optimization, and care delivery.

Across these domains, the case studies span a wide range of health priorities, including tuberculosis, cardiovascular disease, maternal and child health, antimicrobial resistance, mental health, pharmacovigilance, and public health surveillance. The breadth of applications reflects both the diversity of health needs across LMIC settings and the expanding scope of AI-enabled solutions. While no submissions were received within the robotics domain, its inclusion reflects the broader landscape of AI-enabled health innovation and the growing relevance of robotics in future health systems.

Together, this structure highlights the growing maturity of AI in health and provides readers with a clear lens through which to explore the range of emerging use cases and implementation experiences.

Cross-Cutting Insights and Governance Priorities

Across the selected cases, several common patterns emerge. AI systems are predominantly positioned to support health systems. Human oversight remains central in safety-critical applications.


The cases also highlight important areas of continued development. Many implementations remain at pilot or early deployment stages, with fewer examples of large-scale validation and long-term monitoring. This reflects a broader global challenge: the need for stronger evidence generation, standardized evaluation methodologies, and clearer regulatory pathways for AI in health.

WHO guidance emphasizes that robust validation, lifecycle monitoring, and governance frameworks are essential to ensure safety, effectiveness, and equity. These priorities align closely with WHO-ITU-WIPO led Global Initiative on AI for Health, to support the safe and effective adoption of AI across health systems. Through this initiative, WHO is developing guidance, standards, and tools to help Member States navigate the opportunities and risks associated with AI, including work on ethics and governance, regulatory considerations, and evidence generation for AI-based medical devices.

The experiences documented in this compendium reinforce the importance of sustained investment in regulatory capacity, data governance, and ethical oversight as AI moves from innovation to scaled implementation.

Supporting Responsible Adoption and Scaling

By documenting real-world use cases and highlighting cross-cutting lessons, this compendium contributes to global efforts to move from pilot projects toward responsible scaling. The experiences presented demonstrate how AI can support health workers, extend diagnostic capacity, strengthen



surveillance, and improve access to care, particularly in underserved and resource-constrained settings.

At the same time, the cases underscore the importance of sustained collaboration between

governments, researchers, and international organizations to strengthen governance, regulation, and evaluation frameworks. Such collaboration will be essential to ensure that AI delivers meaningful and equitable health impact.

Acknowledgements

The World Health Organization gratefully acknowledges the leadership of the Ministry of Electronics and Information Technology (MeitY), Government of India, for convening this initiative as part of the India AI Impact Summit. We also thank the Ministry of Health and Family Welfare (MoH&FW) and the Indian Council of Medical Research (ICMR) for their collaboration and expert contributions to the review and selection process.

WHO extends its appreciation to the members of the Evaluation Committee: Madhukar Bhagat, Joint Secretary, MoH&FW, Mona Duggal, Director, Indian Council of Medical Research – National Institute for Research in Digital Health and Data Science, (ICMR - NIRDHDS), Sameer Pujari, AI Lead, Department of Data, Digital Health, Analytics and AI, WHO, and Shikha Dahiya, Joint Director, IndiaAI Mission.

From the WHO Department of Data, Digital Health, Analytics and AI, serving as a knowledge partner, Sameer Pujari provided overall oversight

for the development and coordination of the compendium. Technical coordination was led by Rajeshwari Singh, with support from Kanika Kalra, Richelle George, Yu Zhao, Shada AlSalamah, Karthik Adapa (WHO South-East Asia Regional Office), and Grace Achungura (WHO Country Office, India).

From the Ministry of Electronics and Information Technology (MeitY) and the IndiaAI Mission, coordination was led by Kavita Bhatia, COO, IndiaAI. Mohammed Y. Safirulla K., Director, IndiaAI Mission, oversaw the development of the compendium, with technical support from Ishani Agnihotri and Simrandeep Singh, IndiaAI Mission.

We also acknowledge the support of ACCESS Health International, India and the Gates Foundation for their contributions to the design and production of this compendium.



References

Topol, E. J. (2019). High-performance medicine: The convergence of human and artificial intelligence. *Nature Medicine*, 25(1), 44–56. <https://doi.org/10.1038/s41591-018-0300-7>

World Health Organization. (1980). The global eradication of smallpox: *Final report of the Global Commission for the Certification of Smallpox Eradication*. World Health Organization. <https://iris.who.int/handle/10665/39253>

World Health Organization. (2016). *Global strategy on human resources for health: Workforce 2030*. World Health Organization.

World Health Organization. (2020). *WHO results report 2020 – Outcome 1.2: Reduced number of people suffering financial hardship*. <https://www.who.int/about/accountability/results/who-results-report-2020-mtr/outcome/2020/1.2-reduced-number-of-people-suffering-financial-hardship-results-report-achievements>


World Health Organization. (2021). *Ethics and governance of artificial intelligence for health*. World Health Organization.

World Health Organization. (2021). *Generating evidence for AI-based medical devices*. World Health Organization.

World Health Organization. (2023). *Regulatory considerations on artificial intelligence for health*. World Health Organization.

World Health Organization. (n.d.). *Poliomyelitis (polio)*. <https://www.who.int/news-room/fact-sheets/detail/poliomyelitis>

World Health Organization, & World Bank. (2025, December 6). *Most countries make progress towards universal health coverage but major challenges remain*. <https://www.who.int/news/item/06-12-2025-most-countries-make-progress-towards-universal-health-coverage-but-major-challenges-remain-who-world-bank-report-finds>



Bridging the Neuroradiology Divide: Deploying Specialized, Indigenous AI to Augment Diagnostic Expertise in India

Prassanna Langote, Mahesh Shankar,
Kumar Surender Sinwar, Pankaj Chuladhara



Bridging the Neuroradiology Divide: Deploying Specialized, Indigenous AI to Augment Diagnostic Expertise in India

Prassanna Langote, Mahesh Shankar, Kumar Surender Sinwar, Pankaj Chuladhara*

Abstract

A persistent diagnostic capacity gap exists in brain imaging services where access to subspecialty neuroradiology expertise is concentrated in major urban centres, while tier-2 and tier-3 regions often rely on general radiologists managing high and diverse workloads. This imbalance contributes to variability in reporting turnaround time and interpretive consistency for brain CT studies, particularly in emergency and high-volume settings.

This use case describes the deployment of an AI-enabled brain CT decision-support system designed to assist radiologists by providing multi-pathology detection, segmentation, and quantitative measurements across a wide range of neurological findings. The system was implemented on indigenous data-centre infrastructure to align with national data residency and digital health priorities. Multi-site validation and field deployment have been conducted across more than 30 hospitals, including government medical colleges, with over 15,000 CT studies processed under operational conditions.

Observed implementation experience indicates that AI-assisted workflows can support radiologists in prioritising critical findings, standardising measurements, and managing higher reporting volumes. Reported efficiency gains and workflow improvements should be interpreted as operational indicators rather than independently verified outcome effects. This case illustrates how sovereign, workflow-integrated AI decision support may contribute to narrowing specialist access gaps when deployed with appropriate governance and human oversight.

Background

Human Resource Challenge

Radiology workload, particularly for emergency and on-call CT imaging, has quadrupled in 15 years, driven by >500% growth in on-call CT, including head trauma and neurovascular CT.

Growth in emergency CT utilisation, including neuroimaging, has been associated with higher reporting volumes, night-time reading, and fatigue-related performance risks. Published studies indicate that off-hours interpretation is associated with moderately higher error rates compared with daytime reporting, and that fatigue, cognitive load, and workload spikes are associated with increased discrepancy rates.

*MLHealth 360 India Private Limited, Rajasthan, India


Burnout indicators among neuroradiologists are widely reported, with perceived error risk increasing under time pressure. Across routine imaging, published estimates suggest interpretation error rates in the range of 3–5%, with a subset associated with clinically significant delays and adverse outcomes.

Training pathways for independent neuroradiology reporting are lengthy, and many health systems rely on a mix of junior radiologists, generalists, and outsourced teleradiology services. While teleradiology expands coverage, it can introduce additional coordination layers and variable trust, sometimes requiring duplicate local review. These workforce dynamics create operational strain and variability in complex imaging interpretation.

Clinical Challenge

Brain CT interpretation is clinically demanding due to the need to detect subtle or early findings such as small haemorrhages, early ischemic change, ventricular abnormalities, skull fractures, and midline shift. Brain CT represents a substantial (25-30%) share of total CT volume in many hospitals, while the proportion of immediately critical findings is relatively low. This high-volume, low-yield pattern places sustained cognitive demands on reporting radiologists and increases the importance of consistent triage and measurement support.

Figure 01: Gap Analysis of Radiology Infrastructure and Performance Metrics in India Compared to International Benchmarks for Brain CT Imaging (↓ = lower is better; ↑ = higher is better) (ref: 1-6)



Indian Neuroradiology Divide

CT infrastructure has expanded across India, including in non-metropolitan regions, but subspecialty neuroradiology expertise remains concentrated in major cities. Tier-2 and tier-3 facilities frequently depend on general radiologists who must cover multiple modalities and organ systems. This distribution contributes to uneven access to subspecialty-level interpretation and variable turnaround times for brain CT reporting outside metropolitan centres.

National Alignment

India's digital health strategy and national AI initiatives emphasise scalable, locally governed digital solutions that support equitable access while maintaining data sovereignty. In diagnostic imaging, this includes interest in AI systems that can operate within existing hospital infrastructure, support radiologist decision-making, and avoid external data transfer dependencies.

Within this context, the Scaida BrainCT system was developed as a specialised AI decision-support module intended to assist radiologists with multi-pathology brain CT analysis while operating within sovereign infrastructure.

AI that supports radiologists without external data transfer

AI Use Case Description

Solution Description

AI solution — evidence context

Published evaluations of brain CT AI tools report a range of sensitivities (75.6% to 98.9%) and specificities (from 87.35% to over 99%) for detection of selected pathologies such as intracranial haemorrhage. Studies also indicate that combined AI-plus-radiologist workflows can improve detection consistency (sensitivity reaching 95.2% and overall accuracy up to 98.8%,) and case prioritisation compared with unaided review, particularly for urgent findings.

At the same time, false positives, workflow friction, and integration challenges are recognised limitations, and performance varies by dataset, setting, and pathology type.

The Scaida BrainCT system is designed as a neuroradiology decision-support tool that extends beyond binary triage by providing multi-class segmentation and quantitative outputs across multiple pathology categories. Its intended role is to support, not replace, radiologist interpretation.

Commercial system limitations

Several widely used brain CT AI systems focus primarily on binary or limited-class triage (for example, presence or absence of intracranial haemorrhage) to support worklist reprioritisation achieving sensitivities of 0.90-0.95 and specificities of

0.88-0.99. These systems typically provide alerting functions with limited spatial or volumetric detail. Multi-pathology and multi-class segmentation capabilities remain less common and are often restricted to single disease categories.

Brain CT Core Functionality — Multi-Class Segmentation Architecture

Scaida BrainCT implements multi-class (44) semantic segmentation across 40+ neurological and related findings using a labelled dataset derived from approximately 2,000 studies ((48,774 positive slices, 71,000+ radiologist-validated annotations) and tens of thousands of annotated slices validated by radiologists. Rather than binary classification, the model performs voxel-

level labelling across multiple pathology groups within a single inference workflow.

These categories are mapped to clinical priority groupings to support triage and review order decisions. Annotation distributions reflect emergency and trauma-heavy case mixes typical of Indian public hospital settings.

AI-driven voxel-level labelling for brain CT triage

Table 01: Distribution of Annotations Across Clinical Priority Categories in the SCIADA Brain CT Dataset.


Risk Stratification by Pathology Category

Category	Classes	Annotations	Clinical Priority
Hemorrhage (6 subtypes)	SDH, IPH, EDH, IVH, SAH, contusion	5,279	Emergent—surgical triage
Ischemia (4 subtypes)	Acute/chronic arterial, lacunar, venous	4,909	Urgent—thrombolysis window
Trauma (3 types)	Skull/temporal fracture, pneumocephalus	4,882	Emergent—penetrating injury
Structural (7 findings)	Mass effect, ventricular dilation, atrophy	6,715	Moderate—herniation risk
Soft tissue/PNS (10 findings)	Scalp swelling, DNS, sinusitis	>10000	Low—incidental/trauma correlation

AI Methodology — Architecture

The model uses a deep learning segmentation architecture optimised through iterative training with parameter tuning across learning rate, batch size, epochs, and loss weighting. Training included repeated refinement cycles to improve segmentation accuracy across heterogeneous image quality conditions. The architecture is designed for medical image segmentation rather than general classification, with emphasis on spatial precision.

Figure 02: Model Architecture



Patent

The multi-pathology segmentation architecture is reported to be protected under an Indian patent No. 576522 (2025), supporting local intellectual property ownership. Patent status relates to technical design and does not itself establish clinical effectiveness.

Quantitative Clinical Outputs

For detected abnormalities, the system generates structured quantitative outputs intended to support radiologist review, including:


- segmentation-derived volumes (cm^3) replacing ABC/2 estimation for ICH surgical planning,
- anatomical localization with hemisphere, lobe, and slice coordinates,
- displacement metrics (midline shift in mm at septum pellucidum),
- severity scores based on volume thresholds and eloquent region involvement,

- color-coded overlays (red: hemorrhage, blue: infarct, yellow: edema)

These outputs are designed to complement radiologist assessment and manual measurement rather than replace them.

Performance and Dataset Characteristics

Training data include substantial representation of trauma-related and chronic findings like—substantial representation of skull fractures ($n=3,892$), chronic infarcts ($n=3,250$), and PNS pathology ($n=9,316$) frequently encountered in Indian emergency settings. Data augmentation and windowing strategies were used to support model robustness across scanner types and acquisition protocols. Reported performance metrics are derived from internal and multi-site validation and should be interpreted as context-specific.



Deployment & Implementation

DICOM-Native Integration and Workflow Efficiency

The system is integrated into radiology workflows through DICOM-native outputs using DICOM SEG-compliant objects that can be visualised within standard PACS viewers. Automated routing sends brain CT studies to the AI engine in parallel with PACS ingestion. AI outputs are returned as structured objects and overlays prior to radiologist review, avoiding separate logins or proprietary viewers.

This “silent” or background integration model is intended to minimise workflow disruption. AI outputs support lesion measurement, anatomical labelling, and prioritisation flags, while radiologists retain full control over interpretation and reporting.

Cloud-Agnostic Architecture and Data Sovereignty

The platform is designed to be vendor-agnostic and portable across sovereign cloud and hybrid environments. Production inference and patient data processing are configured to remain within Indian data centres in alignment with national data residency expectations. Hybrid deployment options allow local control of PACS data with encrypted transfers to managed compute resources where required.

Hub-and-Spoke Teleradiology for Tier 2/3 Access

The system supports hub-and-spoke teleradiology models in which central

reading hubs assist peripheral facilities (Tier 2 and Tier 3). AI-supported triage and quantification can assist small neuroradiology teams in prioritising urgent cases from distributed sites. This model is intended to complement, not replace, human subspecialty reporting.

Adaptability and Standards Compliance

Integration requires standard DICOM node configuration and routing rules rather than multi-vendor marketplace integration. Use of open standards supports IT governance review and reduces vendor lock-in risk. The system is positioned as an efficiency-enhancing decision support that does not override radiologist authority.

Real-World Deployment & Impact

Scale of Deployment

Deployments have been reported across more than 30 healthcare facilities in tier-2 and tier-3 districts across selected Indian states, including government and teaching institutions. More than 15,000 brain CT studies have been processed through the system in operational settings, with several hundred clinicians exposed to AI-assisted outputs.

- Programme reports indicate large numbers of detected abnormalities across ischemic ~4000 detections (critical for stroke management), haemorrhagic (~3000 detections (identifying acute

bleeds), traumatic (~7000 detections (fractures and pneumocephalus), structural ~8000 detections (mass effect, atrophy, etc.) and soft-tissue/PNS categories (~9000 detections (sinusitis, scalp swelling). These counts reflect model detections rather than independently adjudicated clinical outcomes.

Demographic distributions across processed studies show use across age groups and both sexes, suggesting operational applicability across varied patient populations.

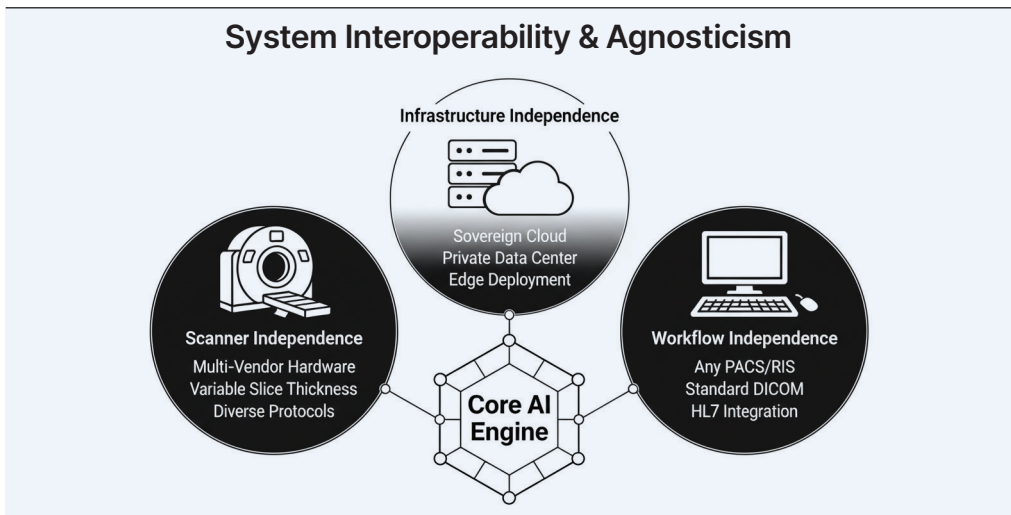
Infrastructure-neutral design has supported deployment in both higher-capacity urban hospitals and lower-resource district settings. Integration within at least MBS Medical College Kota, Rajasthan- a government medical college environment demonstrates feasibility within public sector workflows, subject to local governance approvals.


Stakeholder Benefits

Reported stakeholder-level benefits are operational and indicative rather than definitive:

- Radiologists in non-metropolitan settings report increased confidence in measurement and detection support
- Teleradiology networks report potential capacity expansion per radiologist when AI triage and quantification are used
- Neurosurgeons and neurologists receive standardised quantitative measurements that may support planning discussions
- Trainees can use structured overlays as educational aids
- Hospitals report potential reductions in outsourcing dependence
- Patients in non-metropolitan areas may receive faster preliminary prioritisation

Figure 03: System Interoperability and Agnostic Architecture of SCIADA Brain CT AI Platform





These benefits are based on programme reports and should be interpreted as context-specific operational observations.

Impact

Reported impact indicators include reduced average reporting turnaround times in AI-assisted workflows compared with selected conventional pathways, and higher per-radiologist study throughput in some deployments. Efficiency ratios such as “2–3× volume handling” or large percentage time reductions are based on internal operational measurements and may vary by setting.

Ethics & Governance

The system is designed as clinician decision support requiring radiologist sign-off on all reports, preserving human accountability and medico-legal responsibility. AI outputs are assistive and do not generate final diagnoses autonomously.

Data processing is configured to remain within national infrastructure to support data sovereignty and privacy requirements. Governance features include auditability of outputs, version control, and quality management processes typical of Software as a Medical Device frameworks.

Model design emphasises quantitative accuracy for measurements such as lesion volume and midline shift, but radiologist verification remains required. These practices align with widely recognised principles for responsible AI in health, including human oversight, transparency, and accountability.

Discussion


Implementation experience suggests that clinical value depends not only on model performance but also on workflow alignment, standards-based integration, and clear role definition. Passive, PACS-embedded delivery of AI outputs appears to support higher acceptance than tools requiring separate interfaces.

Key operational lessons include the need to train models on heterogeneous, real-world imaging conditions and to communicate clearly that AI functions as augmentation rather than replacement. Challenges include limited data for rare pathologies, adoption hesitancy among some clinicians, and the need for continuous post-deployment monitoring for performance drift.

Future plans described by the developer include structured post-deployment quality assurance, drift monitoring, and feedback-driven retraining, alongside horizontal scale-out across additional non-metropolitan networks using the same standards-based architecture.

Challenges

Data scarcity in rare pathologies presents a significant challenge. Critical yet rare conditions like venous infarcts (n=56) and intraaxial tumours (n=5) lack sufficient training data. This “long-tail” distribution is particularly difficult for deep learning models, which typically need large balanced datasets to generalise without bias. To overcome this, targeted active learning and rigorous curation were employed to mitigate edge-case presentations.



Workforce anxiety and adoption resistance also posed challenges. There was cultural resistance due to fears of professional displacement, with studies showing 39-90% of radiologists

perceive AI as a threat to job security. Training emphasised that AI tools are intended to augment radiologist capabilities rather than substitute for them.

Conclusion

This use case describes the deployment of a specialised, multi-pathology brain CT AI decision-support system designed to assist radiologists in high-volume and resource-variable settings. By combining segmentation, quantification, and standards-based workflow integration within sovereign infrastructure, the approach seeks to support more consistent neuroradiology-level analysis outside major urban centres.

References


1. Rao, P., Mathur, N., & Kalyanpur, A. (2024). Utilization of teleradiology by intensive care units: A cohort study. *Indian Journal of Critical Care Medicine*, 28(1), 20–25. <https://doi.org/10.5005/jp-journals-10071-24593>
2. Medical Buyer. (2025, June 30). *India has alarming ratio of one radiologist per lakh people*. <https://medicalbuyer.co.in/india-has-alarming-ratio-of-one-radiologist-per-lakh-people/>
3. Manipal Hospitals. (2025, June 13). *Impact of teleradiology in rural India*. <https://radiologygroup.manipalhospitals.com/blog/impact-of-teleradiology-in-rural-india>
4. Shah, A., & Diwan, A. (2023). Stroke thrombolysis: Beating the clock. *Indian Journal of Critical Care Medicine*, 27(2), 107–110. <https://doi.org/10.5005/jp-journals-10071-24405>
5. Insights10. (2024, July 31). *India radiology service market report 2023 to 2030*. <https://www.insights10.com/report/india-radiology-service-market-analysis/>
6. European Society of Radiology. (2025). Guidelines and recommendations for radiologist staffing, education and training in Europe: What do we need? *Insights into Imaging*, 16(1), Article 50. <https://doi.org/10.1186/s13244-025-01887-0>
7. Bruls, R. J. M., & Kwee, R. M. (2020). Workload for radiologists during on-call hours: Dramatic increase in the past 15 years. *Insights into Imaging*, 11(1), Article 121. <https://doi.org/10.1186/s13244-020-00925-z>
8. Patel, A. G., Pizzitola, V. J., Johnson, C. D., Zhang, N., & Patel, M. D. (2020). Radiologists make more errors interpreting off-hours body CT studies during overnight assignments as compared with daytime assignments. *Radiology*, 297(2), 374–379. <https://doi.org/10.1148/radiol.2020201558>
9. Wu, H., Li, M., Chen, J., Fan, X., He, M., Ji, Z., & Cao, H. (2025). Reflections on training and management of junior radiologists during night shifts in China. *BMC Medical Education*, 25(1), Article 879. <https://doi.org/10.1186/s12909-025-07573-1>

10. Kasalak, Ö., Alnahwi, H., Toxopeus, R., Pennings, J. P., Yakar, D., & Kwee, T. C. (2023). Work overload and diagnostic errors in radiology. *European Journal of Radiology*, 167, Article 111032. <https://doi.org/10.1016/j.ejrad.2023.111032>
11. Hanna, T. N., Zygmunt, M. E., Peterson, R., Theriot, D., Krupinski, E. A., & Berland, L. L. (2018). The effects of fatigue from overnight shifts on radiology search patterns and diagnostic performance. *Journal of the American College of Radiology*, 15(12), 1709–1716. <https://doi.org/10.1016/j.jacr.2017.12.019>
12. Chen, J., Pijnappel, R. M., Prokop, M., Schaefer-Prokop, C. M., & Kwee, T. C. (2022). Added value of double reading in diagnostic radiology, a systematic review. *European Journal of Radiology*, 147, Article 110153. <https://doi.org/10.1016/j.ejrad.2022.110153>
13. Chen, J. Y., Vedantham, S., & Lexa, F. J. (2025). 3rd comprehensive survey of the neuroradiology work environment in the United States with reported trends in clinical work, nonclinical work, errors, burnout and retirement. *American Journal of Neuroradiology*. <https://doi.org/10.3174/ajnr.a8913>
14. Brady, A. P., & Becker, C. D. (2019). Teleradiological outsourcing—compromises and hidden costs. *European Radiology*, 29(4), 1647–1648. <https://doi.org/10.1007/s00330-019-6014-5>
15. Agrawal, A. (2022). Emergency teleradiology—past, present, and, is there a future? *Frontiers in Radiology*, 2, Article 866643. <https://doi.org/10.3389/fradi.2022.866643>
16. Szymanski, K. A., Hoang, A. T., Peterson, D. M., Hardy, A., Lee, C., & Luo, Y. (2025). On-call radiology resident preliminary report major discrepancies: A meta-analysis. *Academic Radiology*, 32(4), 2342–2356. <https://doi.org/10.1016/j.acra.2025.01.015>
17. Orejuela-Zapata, J. F., Mejía-Quiñones, V., & Granados-Sánchez, A. M. (2023). Discrepancy rate and clinical impact of preliminary reports from radiology residents. *Heliyon*, 9(8), Article e18714. <https://doi.org/10.1016/j.heliyon.2023.e18714>
18. Goh, G., Kearney, A., & Bourne, R. (2019). An audit of the accuracy of after-hours trainee radiology reporting in an Australian teaching hospital. *Journal of Medical Imaging and Radiation Oncology*, 63(6), 792–798. <https://doi.org/10.1111/1754-9485.12921>
19. Weinberg, B. D., Richter, M. D., Champine, J. G., Morriss, M. C., & Browning, T. (2015). Radiology resident preliminary reporting in an independent call environment. *Journal of the American College of Radiology*, 12(1), 95–100. <https://doi.org/10.1016/j.jacr.2014.08.005>
20. Howie, A. R., Clarke, S., Nolan, E., & Bebbington, N. (2023). After-hours radiology reporting in an Australian tertiary hospital – A retrospective review. *Journal of Medical Imaging and Radiation Oncology*, 67(1), 13–19. <https://doi.org/10.1111/1754-9485.13445>
21. Gurney, J., Muir, L., Tao, L., & Swain, C. (2021). Discrepancies in preliminary radiology resident interpretations of imaging investigations in the emergency department. *Emergency Radiology*, 28(6), 1015–1022. <https://doi.org/10.1007/s10140-021-01950-w>
22. Mohan, C., Nagesh, C., Mathur, S., Ramachandra, P., & Ananthakrishnan, R. (2016). Teleradiology and its impact on patient management. *Indian Journal of Radiology and Imaging*, 26(1), 1–3. <https://doi.org/10.4103/0971-3026.178278>
23. Del Gaizo, A. J., Osborne, T. F., Shahoumian, T., & Sherrier, R. (2024). Deep learning to detect intracranial hemorrhage in a national teleradiology program and the impact on interpretation time. *Radiology: Artificial Intelligence*, 6(5), Article e240067. <https://doi.org/10.1148/ryai.240067>
24. Abramson, R. G., Su, P. F., & Shyr, Y. (2020). Integrating AI into radiology workflow: Levels of research, production, and feedback maturity. *Journal of Medical Imaging*, 7(1), Article 016502. <https://doi.org/10.1117/1.JMI.7.1.016502>
25. European Society of Radiology. (2020). The role of radiology in the diagnosis and management of COVID-19: A statement by ESR. *Insights into Imaging*, 11(1), Article 131. <https://doi.org/10.1186/s13244-020-00931-1>



AI-Assisted Visual Inspection for Cervical Cancer Screening in Rwanda

Maya Fakhfakh



AI-Assisted Visual Inspection for Cervical Cancer Screening in Rwanda

Maya Fakhfakh*

Abstract

Cervical cancer remains a leading cause of preventable cancer mortality in low- and middle-income countries, where screening coverage and diagnostic capacity are often limited. Visual inspection with acetic acid (VIA) is widely used in primary care because it is low cost and provides immediate results, but interpretation depends heavily on provider training and experience, leading to variability in screening quality.

This use case describes an AI-assisted VIA mobile application developed by MednTech to support frontline at Kicukiro Health Center in Kigali, Rwanda providers with image-based decision support during cervical cancer screening. The tool analyzes cervical images captured after acetic acid application and generates a screening classification to complement provider assessment. A pilot deployment at a public primary health centre in Kigali, Rwanda, screened 260 women under routine programme conditions. While the algorithm achieved 97% accuracy in controlled testing, this dropped to 84.5% in real-world primary care conditions with sensitivity of 33.3% and specificity of 92.2%. These results illustrate the common gap between development datasets and real-world implementation environments.

A stakeholder workshop involving providers, policymakers, and a patient representative identified enabling factors, including a supportive national policy environment and provider interest, alongside implementation considerations such as regulatory clearance, infrastructure readiness, and workflow integration. This case illustrates that successful translation of AI-supported screening tools depends not only on model performance, but also on governance, infrastructure, training, and health system alignment.

*Founder and Executive Director, MednTech, Washington DC, USA



Background

Cervical cancer is largely preventable through timely screening and early treatment; however, approximately 94% of global cervical cancer deaths occur in low- and middle-income countries, reflecting persistent gaps in access to preventive and early detection services. Many high-income countries have reduced mortality through cytology and HPV based screening programmes supported by laboratory and specialist infrastructure. In contrast, many resource-constrained settings rely on alternative screening approaches that can be delivered at primary care level.

In Rwanda, cervical cancer is the second common cancer affecting women. Half of cases are detected at advanced stages, reflecting gaps in early screening coverage and diagnostic capacity. Published assessments indicate that only a minority (17.6%) of frontline providers have received comprehensive practical training in cervical screening techniques, contributing to variability in service delivery.

Visual inspection with acetic acid has been adopted widely as a pragmatic screening method in low-resource settings because it is inexpensive, produces immediate results, and can be linked to same visit treatment in screen-and-treat programmes. However, VIA interpretation is inherently subjective and depends on provider skill, lighting, image quality, and examination technique. Inter-observer variability has been documented across multiple settings.

Digital and AI-supported image analysis tools are being explored as a way to support

more consistent VIA interpretation. Rwanda's Accelerated Plan for the Elimination of Cervical Cancer (2024-2027) explicitly prioritizes expanding access to high-precision screening, aligning with the WHO's 90-70-90 targets and other global cervical cancer elimination targets. While AI-assisted VIA approaches have shown promising results in controlled research settings, relatively few have been evaluated under routine primary care conditions. This use case describes one such implementation-focused pilot and associated system-level lessons.

MednTech is a nonprofit organisation working with African health systems and communities to co-design AI-supported preventive care tools. The AI-assisted VIA application described here was deployed within Rwanda's existing screening infrastructure to examine operational feasibility and performance under field conditions.

AI Use Case Description

Solution Description

The AI-VIA application is a clinical decision-support tool, designed to be used on mobile devices and intended to assist frontline providers conducting VIA screening. Using a smartphone camera, providers capture cervical images before and after application of 3–5% acetic acid. The embedded AI model analyzes these images and produces a binary screening output (VIA-negative or VIA-positive) to complement provider judgement.

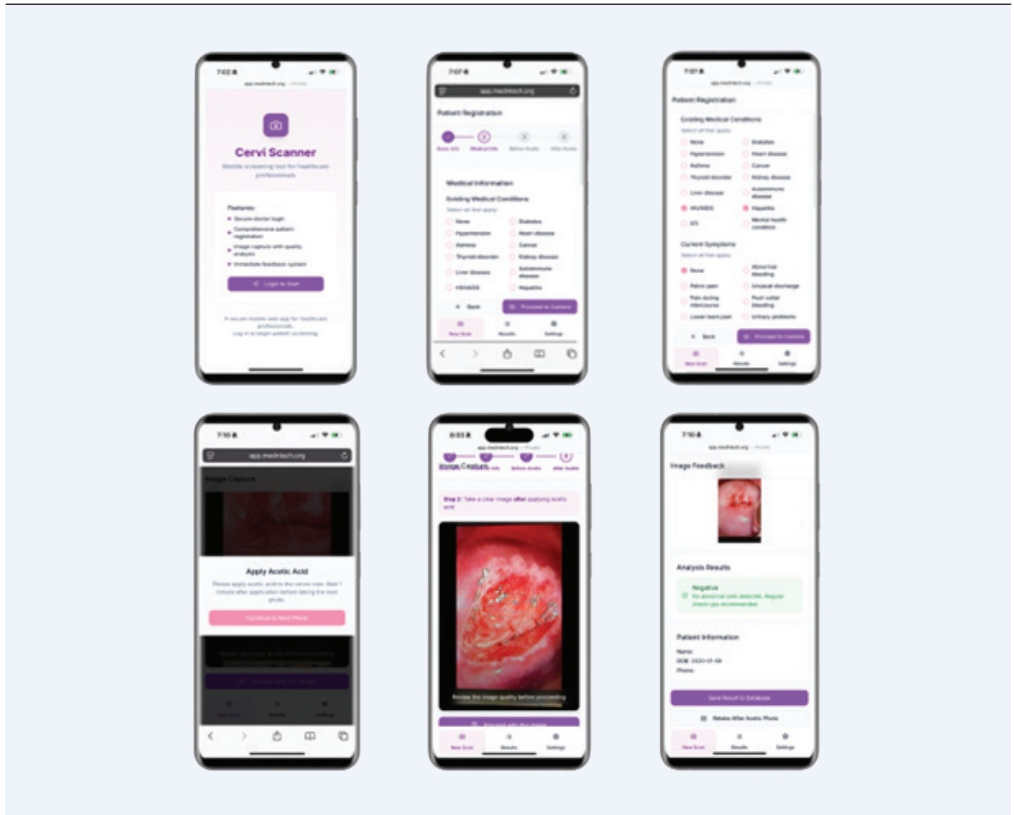
The underlying algorithm is a convolutional neural network based on an EfficientNet-B0 architecture. In addition to image analysis, the

application collects structured patient information across four domains: sociodemographic characteristics, medical history, reproductive history, and selected lifestyle factors. These data support structured documentation and may be used in future analyses, but they are not used by the AI model for screening classification,

The baseline model was trained using annotated VIA images from the International Agency for Research on Cancer (IARC) cervical image bank.

The original dataset comprised 186 labelled cases (363 images). To address class imbalance, the training set was expanded to 500 samples per class using controlled duplication and standard augmentation techniques, including rotation, cropping, and colour adjustments. In a controlled evaluation on a held-out test set of 200 images, the model achieved 97% overall accuracy, with F1-scores ranging from 0.97 to 0.98 across classes.

Figure 01: AI-VIA application workflow



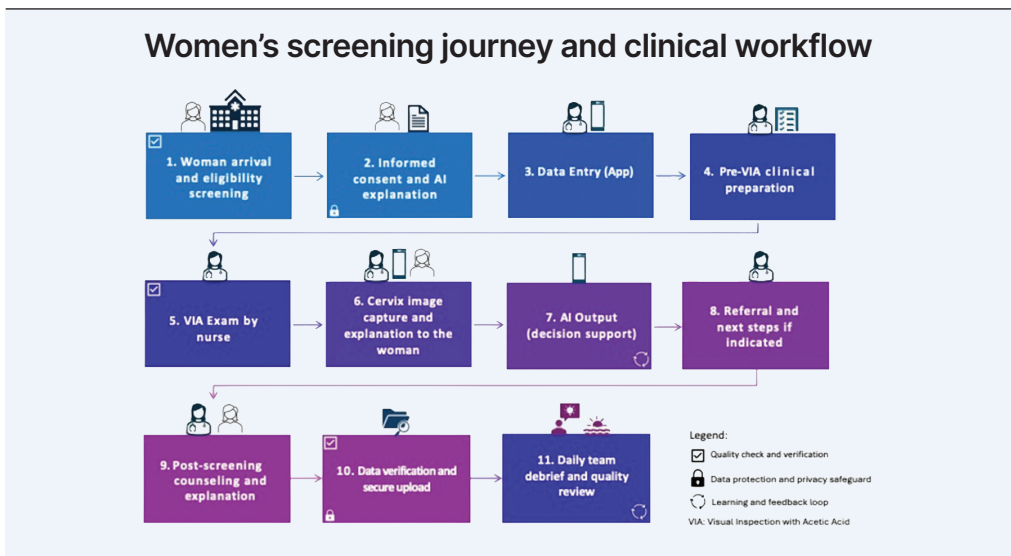
The intended role of the tool is screening decision support. It does not provide a formal diagnosis and is not designed to replace clinical examination, histopathology, or guideline-based management.


Deployment and Implementation

The pilot was conducted from 18 June to 25 July 2025 at Kicukiro Health Center, a public primary health care facility in Kicukiro District, Kigali, and an active implementation site under Rwanda's national cervical cancer screening programme. Prior to deployment, three nurses completed structured training covering image acquisition, informed consent procedures, data entry, and ethical use of the AI application.

Screening workflows followed national VIA protocols (Figure 2). Eligible women were identified, provided informed consent with an explanation of the AI-supported tool, and had relevant clinical information recorded in the application. Following clinical preparation, nurses performed VIA examination, captured cervical images, and communicated findings to the screened women. During the pilot phase, the application was used in an investigational capacity with mandatory human review. AI-generated outputs were used strictly for decision support and referral guidance, with all clinical decisions remaining under provider oversight. Each screening session concluded with post-screening counselling, data verification, and secure data upload, supplemented by daily team debriefings for quality assurance.

Figure 02: Women's screening journey and clinical workflow





After completion of the pilot, a stakeholder workshop in Kigali convened health-care providers, public health experts, policymakers, and a patient representative to review findings and inform considerations for potential scale-up.

Identified Users and Scale of Deployment

Primary users are frontline screening providers, including nurses, midwives, and community-level health workers delivering VIA screening in primary care settings. The application is designed to augment provider capacity, particularly where formal VIA training is limited, by offering decision-support outputs. End beneficiaries are women in the recommended screening age group accessing cervical cancer screening services, especially in settings with limited specialist availability.

During the pilot, 260 women were screened, and 251 complete records were available for analysis after excluding cases with inadequate images or technical data loss. The mean age was 35.6 years \pm 6.3. Most participants had primary education (65.7%), were unemployed (88.1%), and were married (32.7%) or partnered (24.3%). The majority were non-smokers (95.6%), with most reporting 3-4 pregnancies (43.4%) and 3-4 live births (42.6%).

Impact

In real-world primary care testing, the AI algorithm achieved an overall accuracy of 84.5% compared with clinical expert consensus as the reference standard, with sensitivity of 33.3% and specificity of 92.2%.


By comparison, trained nurses achieved sensitivity of 51.5% and specificity of 98.2% under the same conditions.

The VIA-positive prevalence in the study sample (13.1%) is broadly consistent with reported screening prevalence ranges in sub-Saharan Africa (approximately 11–12%). At relatively low prevalence, sensitivity estimates are more variable. The observed high negative predictive value (97.3%) indicates that the tool ruled out disease in most screened individuals; however, this should be interpreted in the context of low VIA-positive prevalence (13.1%) in the study sample.

Stakeholders noted potential positioning of AI-assisted VIA as a triage component within HPV-based screening pathways, consistent with WHO recommendations, rather than as a standalone screening replacement. Participants also noted that further performance improvement particularly in sensitivity would likely require expanded, locally representative training datasets before broader frontline use, especially in rural and underserved settings where HPV testing capacity is limited.

Ethics and Governance

Ethical approval for the pilot was granted by the Rwanda National Ethics Committee (RNEC 782/2025). The consent process comprised two distinct components: consent for participation in the screening pilot and separate consent for the use of anonymised clinical data for model improvement. Both components were explained to participants prior to screening.



Data were stored using secure systems in compliance with Rwanda's Law No. 058/2021 on Data Protection and Privacy.

Stakeholder consultations identified key governance requirements for potential scale-up, including medical device classification and regulatory review by the Rwanda FDA prior to integration into national guidelines, as well as interoperability with Electronic Medical Record (EMR) systems to ensure continuity of care.

Discussion

Successes

The pilot indicated operational feasibility of integrating an AI-assisted VIA tool into an existing primary care screening workflow. Rwanda's policy environment, which supports digital health innovation and cervical cancer elimination efforts, facilitated engagement and institutional dialogue.


Provider feedback indicated interest in decision-support tools for VIA interpretation, particularly among staff with limited prior experience. The application required limited workflow modification because VIA was already established in routine services. Some providers reported that visualising cervical images with patients supported counselling and shared understanding, representing a potential co-benefit for patient engagement.

From a technical perspective, high specificity and negative predictive value (NPV) in this pilot suggest possible value in rule-out or triage-oriented roles, subject to further validation. These observations are consistent with other studies exploring AI-supported visual triage approaches, though cross-study comparability is limited by methodological differences.

Challenges

A key implementation challenge is the gap between controlled test performance and real-world results. The model achieved 97% accuracy on held-out test data from the IARC image bank, but performance decreased to 84.5% under field conditions, with sensitivity of 33.3% in primary care settings. This pattern is not unique. A systematic review of AI-assisted VIA algorithms reported wide sensitivity ranges (22–93%), reflecting the influence of image quality, lighting, acetic acid application technique, and population characteristics on model performance. One contributing factor is limited availability of locally representative training data, which may reduce model robustness across diverse clinical presentations.

Infrastructure constraints further affect deployment feasibility. Published estimates indicate that reliable electricity and connectivity remain uneven across many facilities in sub-Saharan Africa, reinforcing the importance of offline-capable workflows to maintain service continuity. Regulatory pathways for AI-enabled medical devices



are evolving; while national digital health strategies are supportive, detailed classification and approval frameworks are still under development in many settings.

Human factors also influence adoption. Although overall provider response was positive, some concerns were expressed about potential role displacement. Experience from digital health implementation suggests that early engagement, clear scope definition, and demonstration of clinical value are important to supporting acceptance and appropriate use.

Lessons Learned

First, implementation outcomes were primarily driven by context rather than by technology alone. Effective use depended on integration with existing workflows, alignment with health information systems, and provider acceptance. Early assessment of these factors is critical for successful deployment. Second, AI tools are best positioned as clinical decision support, not replacements for professional judgement. Clear communication on intended use and limitations helps build trust and supports appropriate adoption. Third, real-world deployment requires iterative refinement. Continuous monitoring, user feedback, and phased scale-up with evaluation checkpoints are important to maintain performance and safety outside controlled settings. Fourth, impact is greater when AI is embedded within established screening pathways. AI-assisted VIA may be most effective as a triage component within

HPV-based programmes and can support point-of-care screen-and-treat approaches to reduce loss to follow-up.

Conclusion

This pilot indicates that AI-assisted VIA screening is operationally feasible and generally acceptable in primary care settings in Rwanda. The observed high negative predictive value supports potential use as a triage support tool, although sensitivity will need to improve through expansion of locally representative training data. Prospects for scale-up will depend not only on further technical optimisation, but also on contextual readiness, infrastructure capacity, governance, and sustained workforce engagement.

**AI delivers impact
when embedded
in screening
pathways**

References

Adair-Rohani, H., Zukor, K., Bonjour, S., Wilburn, S., Kuesel, A. C., Hebert, R., & Fletcher, E. R. (2013). Limited electricity access in health facilities of sub-Saharan Africa: A systematic review of data on electricity access, sources, and reliability. *Global Health: Science and Practice*, 1(2), 249–261. <https://doi.org/10.9745/GHSP-D-13-00037>

de Sanjosé, S., Perkins, R. B., Campos, N., Inturrisi, F., Egemen, D., Befano, B., et al. (2024). Design of the HPV-automated visual evaluation (PAVE) study: Validating a novel cervical screening strategy. *eLife*, 12.

Geda, Y. F., Lamiso, Y. Y., Berhe, T. M., Mohammed, S. J., Chibsa, S. E., & Endalew, D. A. (2024). Visual inspection with acetic acid positivity in screening and early detection of cervical dysplasia in Africa, 2023: A meta-analysis. *Ethiopian Journal of Health Sciences*, 34(1), 3–14.

Mulongo, M., & Chibwesa, C. J. (2022). Prevention of cervical cancer in low-resource African settings. *Obstetrics and Gynecology Clinics of North America*, 49(4), 771–781. <https://doi.org/10.1016/j.ogc.2022.07.010>

Nkurunziza, C., Ghebre, R., Magriples, U., Ntsumbumuyange, D., & Bazzett-Matabele, L. (2021). Healthcare provider challenges to early detection of cervical cancer at primary healthcare level in Rwanda. *Gynecologic Oncology Reports*, 37, 100810. <https://doi.org/10.1016/j.j.gore.2021.100810>

Republic of Rwanda, Ministry of Health. (2025). *Accelerated plan for elimination of cervical cancer in Rwanda (2024–2027)*. Ministry of Health.

Sankaranarayanan, R., Shyamalakumary, B., Wesley, R., Sreedevi Amma, N., Parkin, D. M., & Nair, M. K. (1999). Visual inspection with acetic acid in the early detection of cervical cancer and precursors. *International Journal of Cancer*, 80(1), 161–163. [https://doi.org/10.1002/\(SICI\)1097-0215\(19990105\)80:1<161::AID-IJC28>3.0.CO;2-4](https://doi.org/10.1002/(SICI)1097-0215(19990105)80:1<161::AID-IJC28>3.0.CO;2-4)

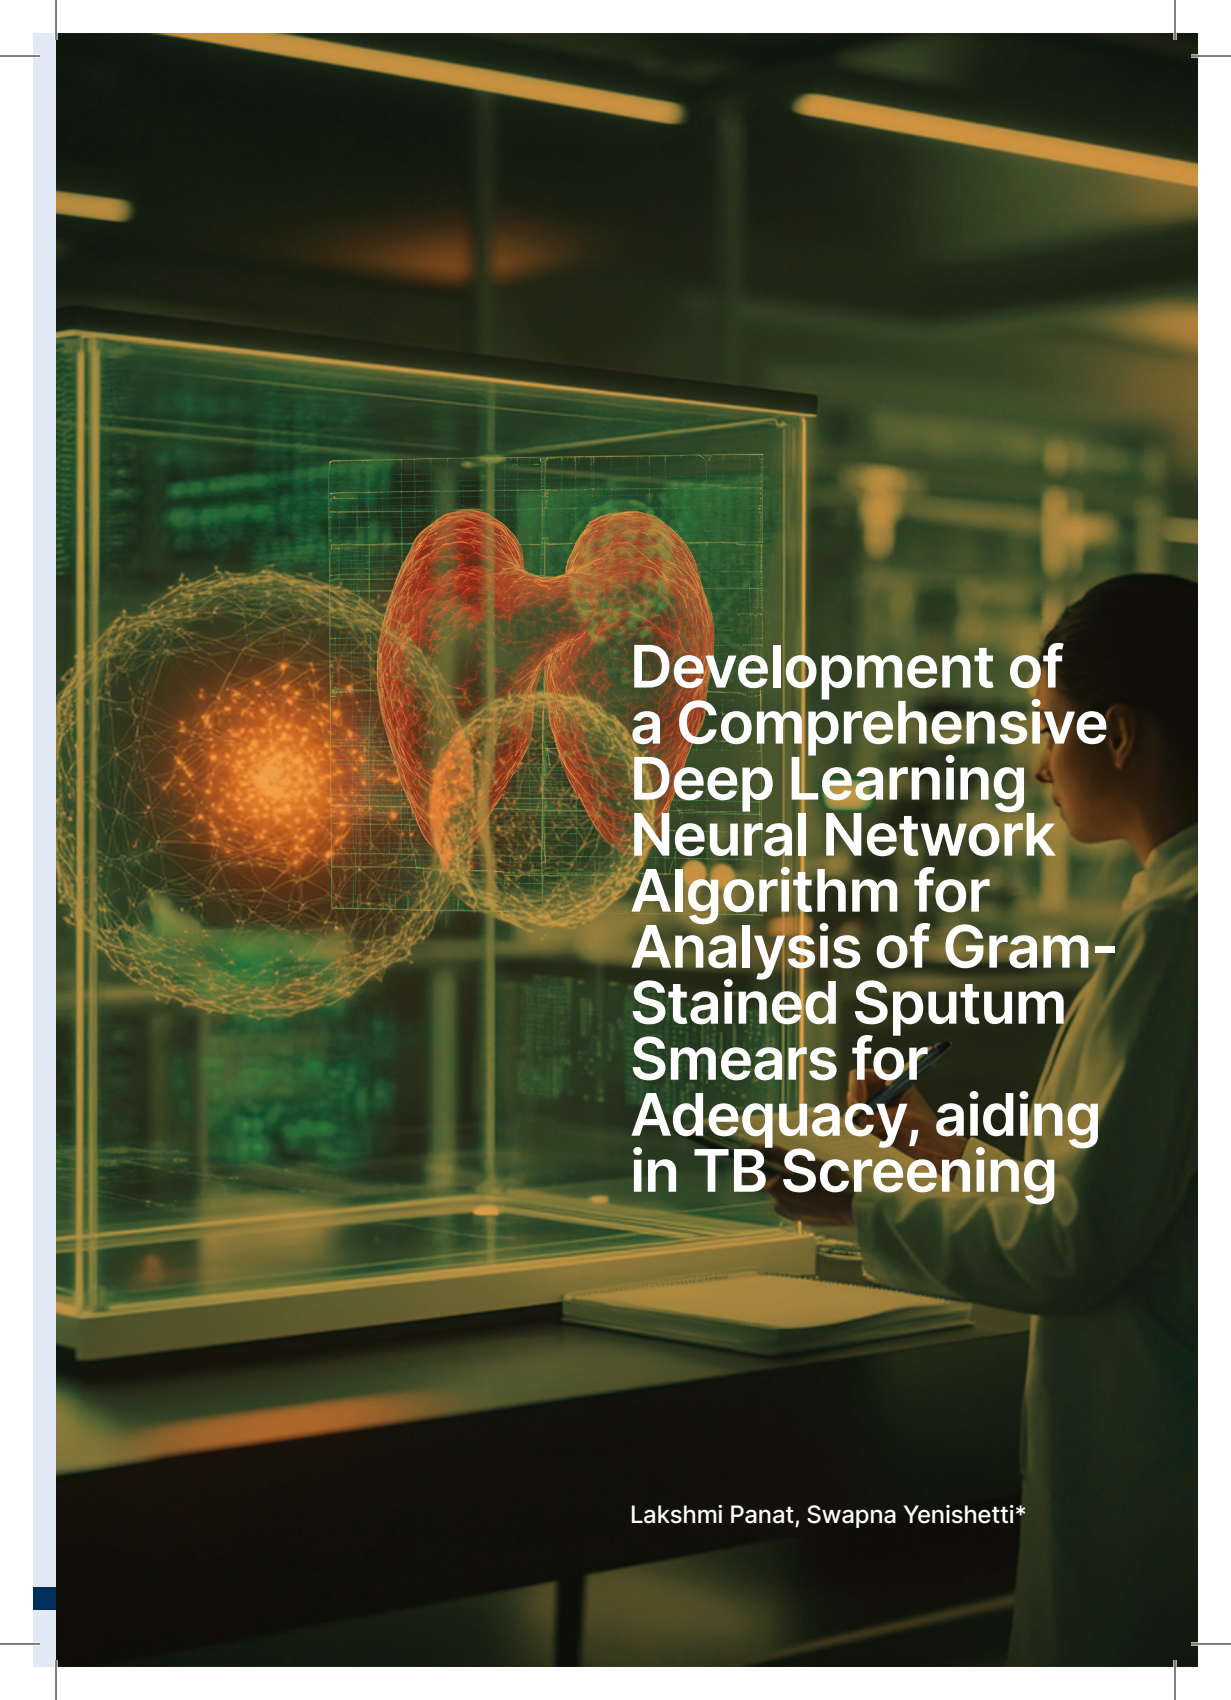
Vedantham, H., Silver, M. I., Kalpana, B., Rekha, C., Karuna, B. P., Vidyadhari, K., et al. (2010). Determinants of VIA (visual inspection of the cervix after acetic acid application) positivity in cervical cancer screening of women in a peri-urban area in Andhra Pradesh, India. *Cancer Epidemiology, Biomarkers & Prevention*, 19(5), 1373–1380. <https://doi.org/10.1158/1055-9965.EPI-09-1282>

Viñals, R., Jonnalagedda, M., Petignat, P., Thiran, J.-P., & Vassilakos, P. (2023). Artificial intelligence-based cervical cancer screening on images taken during visual inspection with acetic acid: A systematic review. *Diagnostics*, 13(5). <https://doi.org/10.3390/diagnostics13050871>

World Health Organization. (2021). *WHO guideline for screening and treatment of cervical pre-cancer lesions for cervical cancer prevention* (2nd ed.). World Health Organization.


World Health Organization. (2022). *Cervical cancer*. <https://www.who.int/news-room/fact-sheets/detail/cervical-cancer>

Zhang, X., Zeng, Q., Cai, W., & Ruan, W. (2021). Trends of cervical cancer at global, regional, and national level: Data from the Global Burden of Disease study 2019. *BMC Public Health*, 21(1), 894. <https://doi.org/10.1186/s12889-021-10907-5>

A person in a white lab coat is looking at a futuristic, glowing digital interface. The interface displays a 3D model of human lungs and a network graph. The background is a dark, modern laboratory or office setting with green and blue lighting.

**Development of
a Comprehensive
Deep Learning
Neural Network
Algorithm for
Analysis of Gram-
Stained Sputum
Smears for
Adequacy, aiding
in TB Screening**

Lakshmi Panat, Swapna Yenishetti*



Development of a Comprehensive Deep Learning Neural Network Algorithm for Analysis of Gram-Stained Sputum Smears for Adequacy, aiding in TB Screening

Lakshmi Panat, Swapna Yenishetti*

Abstract

Tuberculosis (TB) remains a major public health concern in India, which accounts for a roughly one-quarter of all new TB cases worldwide with an estimated 2.8 million incident cases in 2024, despite notable declines in incidence and mortality over the past decade under the National Tuberculosis Elimination Program (NTEP). Sputum smear microscopy continues to serve as an important first-line laboratory method for evaluating lower respiratory tract infections, including pulmonary TB, particularly in resource-constrained settings. Gram-stained sputum smear screening is commonly used to assess specimen adequacy before further diagnostic processing. Adequacy assessment using Bartlett's criteria depends on manual counting of neutrophils and squamous epithelial cells and can be labour-intensive and observer-dependent.

This use case describes the development and pilot evaluation of a deep learning-based image analysis algorithm designed to support automated assessment of sputum smear adequacy from Gram-stained slides. The system uses an object-detection neural network to identify and count neutrophils and squamous epithelial cells and applies Bartlett's criteria to generate an adequacy classification. Image acquisition is performed using mobile phone based microscopy capture, and the model is integrated into a mobile and web-supported workflow.

Pilot validation on a limited annotated dataset demonstrated moderate-to-high agreement with manual assessment under study conditions. The system is designed as a laboratory decision-support tool to assist technicians and microbiologists in adequacy screening, not as a replacement for expert judgement. Reported performance metrics reflect pilot evaluation results and should not be interpreted as definitive or generalisable across all laboratory settings. The case illustrates how AI-based microscopy support tools may complement laboratory workflows when implemented with appropriate governance, validation, and human oversight.

*Centre for Development of Advanced Computing, Pune, India



Background

India carries one of the highest national burdens of tuberculosis globally, with millions of estimated incident cases reported annually. National TB control programmes have contributed to declines in incidence and mortality over the past decade, but timely and accurate diagnosis remains a central operational priority. Laboratory diagnostics play a critical role in this pathway.

Expectorated sputum is widely used to diagnose lower respiratory tract infections, including pulmonary TB, and Gram staining provides a rapid early assessment to guide specimen quality checks and initial treatment decisions. Because sputum samples are often contaminated with oral secretions, adequacy screening is essential, commonly using Bartlett's criteria based on squamous epithelial cell and neutrophil counts. However, manual microscopy-based assessment is resource-intensive and subject to observer variability and workload effects, and poor-quality samples are linked to lower pathogen detection rates. These limitations have driven interest in computational image analysis to enable more consistent and reproducible adequacy screening.

Advances in artificial intelligence, particularly deep learning-based computer vision, have enabled automated detection and classification of objects in microscopy images. Object detection models can identify and localise multiple cell types within a field of view. This creates a potential opportunity to support automated sputum adequacy assessment as a laboratory decision-support function.

AI Use Case

Artificial intelligence methods in computer vision can be applied to microscopy images to detect and localise cellular structures. Deep learning models based on convolutional neural networks can be trained to recognise specific cell types and generate bounding boxes indicating their locations within an image. Such systems produce structured outputs, including object class and count that can support downstream rule-based decisions. In this use case, AI is applied to support automated detection and counting of squamous epithelial cells and neutrophils in Gram-stained sputum smear images. The intended role of the AI system is to assist laboratory personnel in adequacy screening using Bartlett's criteria. It does not diagnose TB or identify specific pathogens and is not intended to replace confirmatory laboratory methods or expert review.

Outputs are designed to be reviewed by trained personnel and incorporated into existing laboratory workflows as decision-support signals.

Solution

The model development process used Gram-stained sputum smear slides labelled as adequate or inadequate according to Bartlett's criteria. A balanced developmental dataset consisting of equal numbers of adequate and inadequate smears was assembled for pilot training and evaluation. All images were manually annotated by domain experts using an image-labelling tool, marking squamous epithelial cells

and neutrophils with bounding boxes. An initial image-segmentation approach was explored, followed by development of an object-detection model using a Faster Region-Based Convolutional Neural Network (Faster R-CNN) architecture. A pre-trained Faster R-CNN model with a ResNet-50 backbone was adapted and fine-tuned using TensorFlow object detection frameworks. Training and evaluation processes were monitored using TensorBoard to track loss curves and detection metrics.

Image acquisition during development was performed using mobile phone-based capture through microscope eyepieces. This approach was selected to reflect potential field deployment conditions and lower-cost digitisation pathways. However, mobile capture introduces variability in lighting, focus, and framing, which can influence model performance and requires standardisation protocols.

Figure 01: Image acquisition through mobile phone



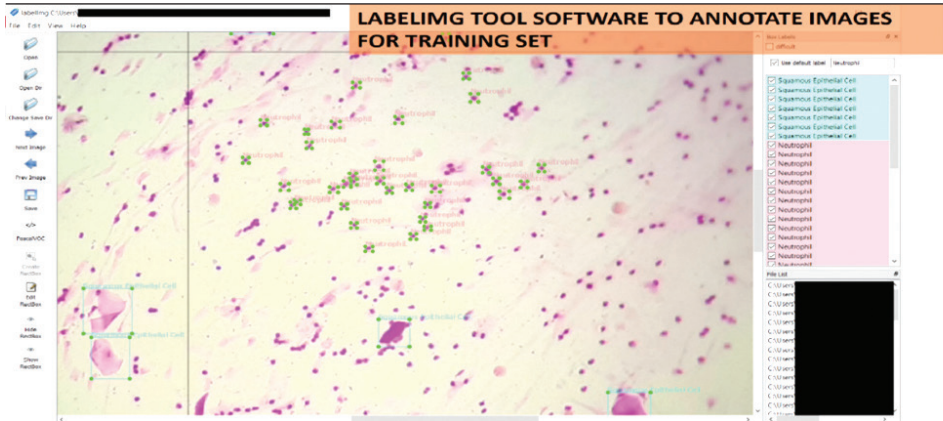
IMAGE CAPTURE USING SMARTPHONE CAMERA

Following training and iterative correction based on expert feedback, the model was tested on a separate pilot set of 100 Gram-stained sputum smear slides. Model outputs included counts of detected squamous epithelial cells and neutrophils per image. These counts were mapped to adequacy classifications

using Bartlett's criteria thresholds. Results were compared with manual microscopy-based adequacy determinations using cross-classification tables.

All the collected images were manually annotated through 'labellmg' tool as shown below.

Figure 02: Annotation through labelling

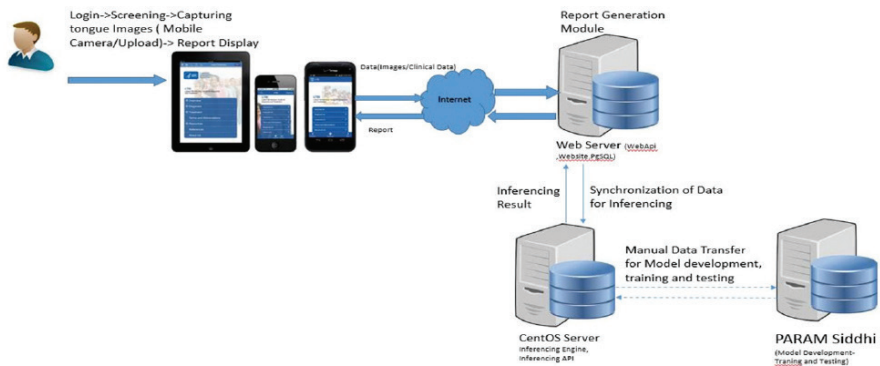


Deployment & Implementation

A deployment workflow was designed to support integration of the AI model into laboratory practice as a decision-support layer. The proposed workflow includes a mobile application

for image capture, a web-based interface for review, and backend AI infrastructure for model inference and retraining.

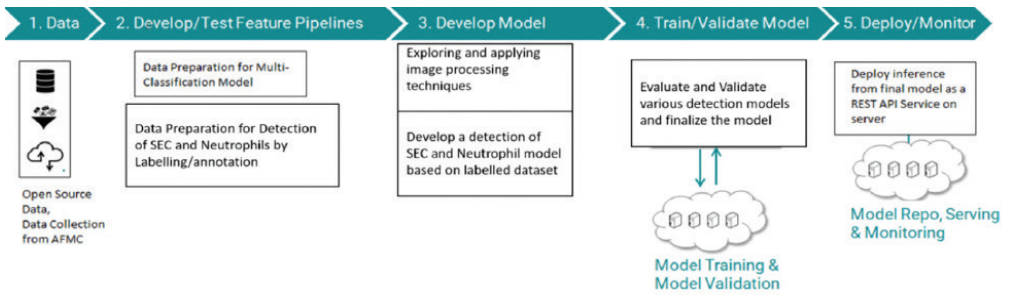
Figure 03: Comprehensive AI Workflow



In the end-to-end pipeline, authenticated users capture sputum smear images through a microscope using a mobile device at defined magnification levels. Images are transmitted to an inference service hosted on dedicated AI compute infrastructure. The inference service processes each image using the trained object-detection model and returns quantitative outputs, including detected counts of squamous epithelial cells and neutrophils.

Deep learning (Faster R-CNN) model to assess sputum smear adequacy using Gram stained microscopy images


Figure 04: End to end AI pipeline



Pilot implementation activities included user onboarding, role-based access configuration, and basic training for laboratory technicians and microbiologists on image acquisition, system use, and interpretation of AI outputs. A pilot deployment was conducted at Armed Forces Medical College (AFMC), Pune, a tertiary medical institution laboratory setting, with ongoing monitoring of performance and user feedback used to guide iterative refinements.

Scale of Deployment

The technical architecture uses a mobile capture interface combined with centralised inference services, allowing additional laboratories to be onboarded with limited configuration changes where compatible infrastructure is available. This design supports potential multi-site deployment from a technical perspective. However, broader scale-up depends on additional factors, including regulatory approvals, laboratory validation studies, governance approvals, connectivity,



device standardisation, and training capacity. Current evidence reflects pilot-scale deployment and testing rather than national-scale implementation.

Identified Users

Primary intended users are laboratory technicians, microbiologists, and authorised diagnostic personnel involved in sputum microscopy and specimen screening workflows.

Impact

Pilot deployment findings indicate that AI-assisted sputum adequacy assessment can produce automated cell counts and adequacy classifications that show moderate-to-high agreement with manual assessment under study conditions. In the pilot evaluation dataset, agreement between model-detected and manually counted cells was higher for neutrophils than for squamous epithelial cells. Reported sensitivity, specificity, and overall accuracy for adequacy classification reflect pilot test results on a limited dataset. From an operational perspective, AI-assisted counting may support more standardised and reproducible adequacy screening and may reduce manual screening time per slide in some workflows. Potential public health value lies in supporting laboratory efficiency and consistency where microscopy workload is high and expert availability is limited, subject to further validation.

Ethics & Governance

Model development used ethically approved and institutionally authorised datasets. Institutional ethical clearance and permissions were obtained before project initiation. Data used for training

and evaluation were handled under laboratory governance protocols. Governance measures include performance monitoring, audit logging, and periodic review to detect model drift or systematic error. The system is designed with human-in-the-loop oversight, with laboratory experts retaining responsibility for adequacy decisions and downstream diagnostic actions. Role-based access controls, authentication, and secure data handling are incorporated into deployment design.

Discussion

This use case describes a pilot application of deep learning-based object detection to support automated sputum smear adequacy assessment using Bartlett's criteria. The workflow spans specimen imaging, annotation, model training, inference, and laboratory integration. It illustrates how computer vision methods can be adapted to microscopy quality-control tasks.

Successes

Pilot results demonstrate technical feasibility of detecting and counting key cell types in Gram-stained sputum images using a Faster R-CNN-based model. Agreement with manual counts was higher for neutrophils and moderate for squamous epithelial cells under study conditions. The system produced adequacy classifications with moderate-to-high sensitivity and moderate specificity in the pilot dataset.

As compared to the manual method, algorithm reported a Sensitivity of 90.56% a specificity of 78.72% and an overall of 85% in determining whether a slide is adequate or not. Refer Fig-8

Challenges

Key challenges included variability in smear quality and microscopy conditions, difficulty in accurately detecting small and overlapping cells, and the need to standardize image capture across users. Additional challenges involved expert data annotation, workflow integration, system connectivity, and ensuring data privacy and governance compliance.

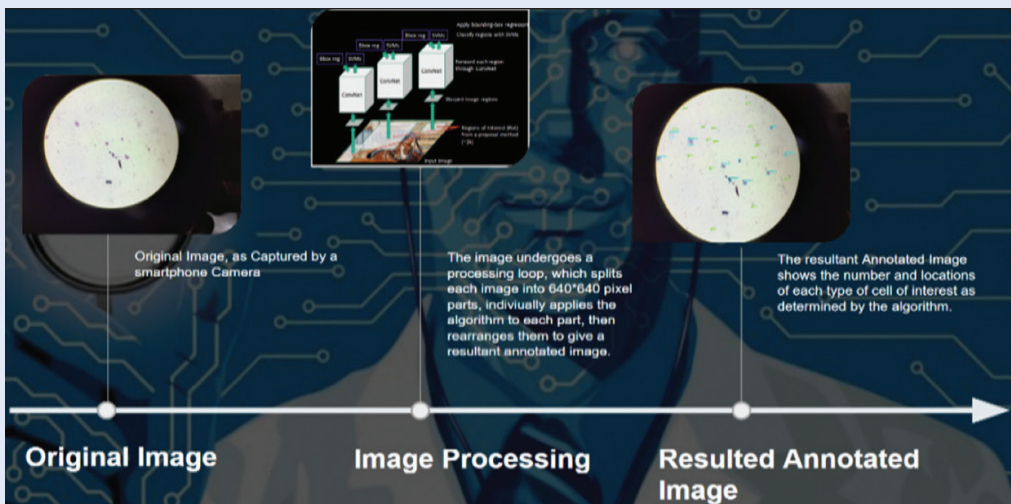
Lessons Learned

Standardised image acquisition protocols and user training are critical for consistent performance. Close domain-expert collaboration improves annotation quality and model relevance. Iterative testing in real laboratory environments helps identify usability and robustness gaps. Early attention to governance, access control, and scalability supports trust and sustainability.

Conclusion


The F-RCNN based algorithm indicated the feasibility of an approach for detection of Squamous Epithelial Cells and Neutrophils in Gram Stained Sputum Smears and thus determining their adequacy using Bartlett's Criteria. Further validation studies are needed to refine the AI model, and this approach may inform the development of similar AI-based tools to support TB diagnostic efforts. This is particularly relevant in resource-limited settings, where trained microbiologists may not be easily available and a large volume of screening of slides is required. AI-based tools such as this may support improved diagnostic accuracy, facilitate earlier treatment initiation, and complement public health surveillance efforts in such contexts.

Figure 05: Transformation of Sputum smear slide image




References

- Belciug, S., & Gorunescu, F. (2020). Era of intelligent systems in healthcare. In *Intelligent decision support systems: A journey to smarter healthcare* (pp. 1–55).
- Budayanti, N. S., Suryawan, K., Iswari, I. S., & Sukrama, D. M. (2019). The quality of sputum specimens as a predictor of isolated bacteria from patients with lower respiratory tract infections at a tertiary referral hospital, Denpasar, Bali–Indonesia. *Frontiers in Medicine, 6*, 1–6.
- Central TB Division, Directorate General of Health Services, Ministry of Health and Family Welfare, Government of India. (2017). *National strategic plan for tuberculosis elimination 2017–2025*.
- Chinnusamy, N., Arumugam, V., & Vedachalam, D. (2016). A study on analysis of the sputum Gram staining and culture in patients with lower respiratory tract infections attending a tertiary care hospital. *Indian Journal of Microbiology Research, 3*(1), 24.
- Flournoy, D. J. (1998). Interpreting the sputum. *Laboratory Medicine, 29*(12), 763–768.
- Fukuyama, H., Yamashiro, S., Kinjo, K., Tamaki, H., & Kishaba, T. (2014). Validation of sputum Gram stain for treatment of community-acquired pneumonia and healthcare-associated pneumonia: A prospective observational study. *BMC Infectious Diseases, 14*(1).
- Maciel, E. L. N., do Prado, T. N., Peres, R. L., Palaci, M., Johnson, J. L., & Dietze, R. (2009). Guided sputum sample collection and culture contamination rates in the diagnosis of pulmonary TB. *Jornal Brasileiro de Pneumologia, 35*(5), 460–463.
- Peiffer-Smadja, N., Dellière, S., Rodriguez, C., et al. (2020). Machine learning in the clinical microbiology laboratory: Has the time come for routine practice? *Clinical Microbiology and Infection, 26*(10), 1300–1309.
- Qu, K., Guo, F., Liu, X., Lin, Y., & Zou, Q. (2019). Application of machine learning in microbiology. *Frontiers in Microbiology, 10*, 827.
- Said Agrebi, A. L. (2020). Use of artificial intelligence in infectious diseases. In *Artificial intelligence in precision health* (pp. 415–438).
- Samuel, L. P., Balada-Llasat, J. M., Harrington, A., & Cavagnolo, R. (2016). Multicenter assessment of Gram stain error rates. *Journal of Clinical Microbiology, 54*(6), 1442–1447.
- World Health Organization. (2024). *Global tuberculosis report 2024*. <https://intep.in/node/352/CP-burden-tb-india>

The image features a hand holding a glowing pink awareness ribbon. The ribbon is intertwined with digital data points and network lines, symbolizing the integration of AI and technology in healthcare. The background is a gradient of purple and blue, with a starry, digital pattern at the bottom.

Harnessing AI for Cervical Cancer Surveillance and Clinical Decision Support in Kenya

Terry Mochire*



Harnessing AI for Cervical Cancer Surveillance and Clinical Decision Support in Kenya

Terry Mochire*

Abstract

Cervical cancer remains a leading cause of cancer-related mortality among women in Kenya, despite being largely preventable through timely screening, diagnosis, and treatment. Persistent challenges—including fragmented data systems, paper-based clinical workflows, and limited access to decision-support tools—have constrained the effective use of existing cancer data across clinical, programme, and policy settings.

In response, IntelliSOFT Consulting Limited, in collaboration with the National Cancer Institute of Kenya (NCI-K), has developed an AI-enabled digital platform designed to strengthen cervical cancer care across the continuum, from point-of-care clinical decision support to national surveillance and research. The platform integrates a FHIR-compliant mobile application for structured clinical data capture and guideline-aligned recommendations, an interactive public health dashboard for programme monitoring and equity analysis, and a research interface using retrieval-augmented generation to support evidence-informed enquiry.

Built in alignment with WHO SMART Guidelines and national digital health standards, the solution is designed for low-resource settings, supporting offline-first workflows, interoperability with national systems, and role-based access. Early implementation as a research-backed proof of concept indicates the feasibility of responsible AI to support standardised clinical workflows, strengthen the use of data for planning and policy, and expand access to cancer-related information and analysis. This chapter describes an early approach to embedding AI within a public health system, with potential relevance to other low- and middle-income countries, aligned with national cancer control priorities and global cervical cancer elimination goals.

*IntelliSOFT Consulting Limited, Nairobi, Kenya



Background

Cervical cancer remains a major public health challenge in low- and middle-income countries, accounting for nearly 94% of global cervical cancer deaths. In Kenya, it is among the most common cancers affecting women and a leading cause of female cancer mortality, with an estimated 5,200 deaths annually. This burden reflects gaps in prevention and care, including limited screening coverage, delayed diagnosis, and loss to follow-up before treatment.

Global and national strategies recognise cervical cancer as a priority for elimination. The World Health Organization's 2030 targets call for 90% HPV vaccination coverage, 70% screening coverage, and 90% treatment for women diagnosed with cervical disease. Kenya's National Cancer Control Strategy aligns with these goals and prioritises scaling up screening and treatment services nationwide. Despite progress, coverage remains uneven: as of 2023, approximately one-third of eligible women had ever been screened, and fewer than 40% of women who screened positive completed appropriate care.

Structural and systemic constraints contribute to these gaps. Fragmented information systems, reliance on paper-based documentation, and limited tools for translating guidelines into frontline decision-making hinder continuity of care. Weak referral tracking and limited visibility of programme performance further constrain effective planning and equity-focused interventions.

Kenya's broader digital health agenda provides an enabling context for innovation.


National investments in connectivity, the Health Information Exchange (HIE), and interoperable digital infrastructure create opportunities to strengthen data-driven service delivery. Within this landscape, IntelliSOFT initiated an AI-enabled cervical cancer platform in 2024 to support prevention, early detection, and surveillance. The initiative seeks to bridge the gap between guideline-based policy and frontline practice by translating national clinical guidance into actionable, context-aware decision support, underpinned by formal ethical approval and research authorisation from national bodies.

AI Use Case

Solution Description

The solution is an integrated, standards-based digital platform designed to support cervical cancer prevention, diagnosis, surveillance, and research across multiple levels of the Kenyan health system. Rather than functioning as a standalone application, the platform operates as a cohesive ecosystem comprising three interoperable components: a mobile clinical application, a public health analytics dashboard, and a research and knowledge discovery interface.

At the clinical level, a FHIR-compliant mobile application enables frontline healthcare workers to capture structured data related to screening, diagnosis, and follow-up using workflows aligned with national tools and guidelines. Artificial intelligence is applied to support, rather than replace, clinical judgement by translating existing clinical guidelines into context-aware recommendations that reflect patient



characteristics and facility-level resources. These recommendations are generated through a combination of rules-based logic and probabilistic reasoning informed by aggregated data patterns.

To support evidence-informed enquiry, the platform incorporates a retrieval-augmented generation (RAG) approach that integrates structured clinical and registry data with curated unstructured sources, including national guidelines and peer-reviewed literature. Large language models are used to synthesise responses to natural language queries while maintaining traceability to source material. The Med-Gemma model was selected for its suitability for biomedical text, controlled deployment, and governance within a public health context.

Explainability is a core design principle. AI-generated outputs are accompanied by references to the underlying data and guidance used, enabling users to understand, interrogate, and contextualise recommendations. The platform is localised to Kenyan clinical protocols and reporting structures, supporting relevance and alignment with national practice.

Deployment and Implementation

The platform is designed to function within routine service delivery environments. Frontline providers use the mobile application during patient encounters to capture demographic information, screening history, clinical findings, and management actions through structured forms aligned with existing paper-based workflows.

An offline-first architecture allows data to

be captured securely on the device and synchronised with a central FHIR server when connectivity is available. Synchronisation is incremental and resilient to intermittent network conditions. Each clinical interaction is recorded as a FHIR resource, supporting interoperability with national systems such as the cancer registry and routine health information platforms.


Once synchronised, data feed into an analytics layer that supports both operational monitoring and AI-assisted insights. Programme managers access aggregated, de-identified data through a dashboard that visualises trends, disparities, and performance indicators. Researchers and analysts interact with data through a dedicated interface that supports structured queries and natural language exploration.

Role-based access controls ensure appropriate data visibility across user groups, and all interactions are logged to support auditability. AI outputs are presented as advisory, with accountability for clinical and programme decisions retained by human users.

Scale of Deployment

The platform is implemented as a research-backed proof of concept intended for validation and learning rather than immediate nationwide deployment. Activities focus on demonstrating technical feasibility, workflow integration, usability, and acceptability within real clinical and programme contexts.

The system architecture and governance model are designed at national scale, enabling future expansion without structural redesign. Piloting activities are being conducted in collaboration with the National Cancer Institute of Kenya and



selected stakeholders to generate evidence and operational guidance required for broader adoption. A phased approach to scale prioritises safety, quality, and governance over rapid expansion.

Identified Users and Beneficiaries

Primary users include frontline clinicians such as nurses, clinical officers, and medical officers who use the mobile application to support guideline-aligned care during patient encounters. Public health and programme managers are secondary users, benefiting from analytics that support planning, supervision, and resource allocation at facility, county, and national levels.

Researchers and policymakers are beneficiaries of the research interface, which supports evidence generation and policy analysis through accessible interaction with structured datasets and relevant literature. Collectively, these user groups contribute to a more integrated and learning-oriented cancer control ecosystem.

Impact

Providers are guided through standardised workflows aligned with national guidelines. However, given the platform's current development stage, its impact is best observed at the process level rather than at the population level.

At the programme level, enhanced visibility of screening and follow-up data supports equity-focused analysis, enabling identification of geographic and demographic disparities. Stakeholder engagement indicates strong demand for tools that bridge the gap between data collection and actionable decision-making.

The platform's modular architecture enables progressive scaling, with additional facilities, counties, or provider networks onboarded as implementation readiness grows. This phased expansion strategy prioritizes quality, safety, and governance over rapid rollout, while preserving readiness for county-level and national deployment once defined validation milestones are achieved..

Ethics and Governance

Ethical and governance considerations are embedded throughout the platform's design and deployment. The initiative operates under formal ethical approval and research authorisation from Kenya's National Commission for Science, Technology and Innovation. Data protection practices align with the Kenya Data Protection Act, including provisions for consent, data minimisation, and secure storage.

Access to data and AI outputs is governed through role-based controls. Patient-identifiable data are restricted to clinical contexts, while aggregated or anonymised data support surveillance and research. AI outputs are transparent and traceable, reinforcing clinical accountability and informed decision-making.

Institutional oversight is provided through collaboration with the National Cancer Institute of Kenya, ensuring alignment with national priorities and standards. This governance framework reflects WHO principles for responsible AI in health, including human oversight, transparency, accountability, and equity.



Discussion

Successes

Early experience suggests that AI can be integrated within existing health system structures in a manner that appears operationally feasible and acceptable to stakeholders, though formal evaluation is ongoing. Alignment with recognised standards such as HL7 FHIR and national workflows facilitated integration without disrupting routine practice.

Strong institutional collaboration, particularly with the National Cancer Institute of Kenya, ensured alignment with national cancer control priorities and governance frameworks. The platform's multi-level design is intended to support continuity across clinical care, programme monitoring, and research within a single ecosystem.

Challenges

Despite early successes, several challenges remain. Data availability and quality are key constraints: historical registry data are often incomplete or heterogeneous, limiting immediate analytical depth and requiring substantial preprocessing. This highlights the need for incremental gains and realistic expectations in data-constrained settings. Technical challenges have also included optimizing model performance and infrastructure for low-latency, cost-predictable operation, with iterative trade-offs needed to balance explainability, accuracy, and resource use when integrating AI models into health system workflows. Operationally, sustained change management is essential. Introducing AI-supported tools requires ongoing training, clear communication of roles, and reinforcement that AI outputs are advisory rather than directive.

Lessons Learned

Key lessons include the importance of interoperability and governance as foundational design elements, the value of structured proof-of-concept implementations for generating operational knowledge, and the central role of trust and institutional relationships in responsible AI deployment. Designing AI to complement human expertise, rather than obscure accountability, remains critical to adoption and sustainability.



Conclusion


This use case describes the development and early implementation of an AI-enabled platform for cervical cancer surveillance and clinical decision support in Kenya. By embedding AI within existing health system structures and aligning with national and global standards, the initiative illustrates a pragmatic approach to strengthening decision-making across the cancer care continuum.

While still at a proof-of-concept stage, the platform provides insights into how responsible AI can support integrated, people-centred cancer control in low- and middle-income settings. As Kenya advances toward national and global cervical cancer elimination targets, this experience offers transferable lessons in governance, interoperability, and ethical deployment relevant to other contexts seeking to harness AI for public health impact.



ECG Image Interpretation using Parameter-Efficient LoRA Fine-Tuning with Multimodal LLaMA V.3.2

Nandakishor M, Anjali M*



ECG Image Interpretation using Parameter-Efficient LoRA Fine-Tuning with Multimodal LLaMA V.3.2

Nandakishor M, Anjali M*

Abstract

Cardiovascular diseases remain a leading cause of mortality globally. Electrocardiography (ECG) is a widely used diagnostic modality for assessing cardiac electrical activity, but interpretation accuracy varies with training and experience. In many resource-constrained settings, ECGs are available primarily as printed or scanned images, limiting the applicability of signal-based automated analysis tools.

This use case describes ECG-LLaMA, a multimodal AI model designed to support interpretation of ECG images and generation of structured clinical reports. The model is based on multimodal LLaMA V.3.2 and was adapted using parameter-efficient Low-Rank Adaptation (LoRA) fine-tuning on a large instruction-style ECG image dataset. Reported validation metrics indicate technical discrimination performance and structured report generation quality under test conditions. The model achieved an Area Under the Curve (AUC) of 0.98 and a report generation quality score of 85.4.

The system is intended as a clinical decision-support tool that assists clinicians in ECG interpretation workflows, particularly where cardiology expertise is limited. The case illustrates how multimodal AI models can be adapted for image-based ECG interpretation with human oversight and governance safeguards.


Background

Cardiovascular diseases account for a substantial share of global morbidity and mortality. Timely and accurate ECG interpretation supports triage and management decisions across emergency, primary care, and inpatient settings. However, ECG interpretation requires training and experience, and

variability in interpretation accuracy has been documented among non-specialist clinicians.

In many low-resource and distributed care environments, ECG machines are available but specialist interpretation is not consistently accessible. In addition, ECG outputs are often

*Convai Innovations Pvt. Ltd., Kasaragod, Kerala



stored or transmitted as printed strips or scanned images rather than raw digital signals. This limits the applicability of many automated ECG algorithms that depend on waveform signal input.

Advances in multimodal large language models (MLLMs) have created new possibilities for image-based clinical interpretation tasks. These models can jointly process visual inputs and generate structured explanatory text. However, general-purpose models are not inherently reliable for clinical use and require domain-specific adaptation and evaluation.

The ECG-LLaMA project was developed to explore whether a multimodal model can be adapted to support ECG image interpretation and structured reporting. The goal is to support, not replace clinician judgement, particularly for non-specialist providers who require rapid interpretive support.

AI Use Case

Solution Description

ECG-LLaMA is a fine-tuned multimodal AI model designed to interpret ECG images and generate structured clinical interpretation text. The system uses a multimodal vision language architecture based on LLaMA V.3.2 and applies parameter-efficient LoRA fine-tuning to adapt the base model to ECG interpretation tasks.

The system processes ECG images rather than raw signal traces. This supports workflows where only printed or scanned ECG outputs are available. The intended function is structured interpretation support and draft report generation for clinician review.

The model architecture includes:

- a vision encoder that processes ECG image inputs
- a language model component that generates descriptive and interpretive text
- instruction-style fine-tuning to align outputs with clinical reporting patterns

Training used the ECGInstruct dataset, described as a large instruction-style ECG image-report corpus (approximately one million diverse ECG samples) with expert-informed annotations and synthetic augmentation to reflect real-world image artefacts such as noise, folds, and layout variation. Dataset composition and annotation processes influence model behaviour and represent a potential source of bias or limitation.

LoRA fine-tuning was applied to adapt the base model while limiting the number of updated parameters. This approach reduces compute requirements compared with full-model fine-tuning but does not eliminate the need for domain validation.

The model produces interpretive text and suggested findings. It does not directly access patient clinical context and does not make treatment recommendations.

Deployment & Implementation

Model development followed a staged pipeline including dataset assembly, model adaptation, and validation.

Data preparation:

Public ECG datasets and curated ECG image sources were aggregated. Synthetic instruction-style pairs were generated and reviewed with cardiology input. Data were de-identified prior to model use. Dataset representativeness remains a limiting factor and may not cover all device formats or population variations.

Model training:

Fine-tuning was conducted using a parameter-efficient LoRA strategy on multimodal LLaMA V.3.2. Efficiency-oriented optimisation frameworks were used to reduce compute overhead. Training configuration affects model behaviour and reproducibility and should be documented in deployment settings.

Validation:

Performance was evaluated using discrimination metrics (including AUC and F1-score) and structured report quality scoring using automated evaluators with human expert verification. Automated evaluators are themselves imperfect, and results should be interpreted with caution.

Operational Workflow: The system is designed for deployment as either a cloud-based API or as an on-premise inference engine.

1. Input:

A healthcare worker captures an image of a printed ECG strip using a smartphone camera or scanner.

2. Processing:

The image is securely uploaded to the ECG-LLaMA system for analysis.

3. Output:

The system produces a structured clinical report describing cardiac rhythm, measured intervals, morphological abnormalities, and suggested findings for clinician review.

4. Clinical review:

A qualified clinician reviews and validates the AI-generated findings to inform the final clinical decision.

Scale of Deployment

Geographical coverage:

The solution has been developed in India, with initial pilots and validation partnerships established in Kerala, including at Dr. Moopen's Medical College. The underlying model has been open-sourced on Hugging Face to support broader global adoption, independent validation, and further research (ECG-Instruct-Llama-3.2-11B-Vision).

User base:

The primary target users are general practitioners, emergency room physicians, and medical officers at primary health centres who may not have immediate access to cardiology consultation. The open-source release has also attracted interest from the medical AI research community.

Identified Users

Primary users: Non-specialist physicians, nurses, and technicians in primary and secondary care settings who perform or review ECGs.

Beneficiaries:

Patients undergoing ECG assessment, particularly in settings with limited specialist availability. Potential benefit relates to interpretive support and triage assistance rather than automated diagnosis.

Impact

The solution demonstrated improvements in diagnostic performance and reporting quality during validation:

- **Diagnostic precision:**
Achieved an AUC of 0.98 for abnormality detection, compared with 0.51 for the baseline model.
- **Report quality:**
Reached a report quality score of 85.4/100, indicating structured alignment with clinical reporting conventions under test conditions.
- **Efficiency:**
The LoRA-based fine-tuning approach reduces computational requirements, enabling deployment on comparatively lower-cost hardware.
- **Error detection and correction:**
The model achieved over 90% accuracy in identifying common life-threatening arrhythmias, including atrial fibrillation and ventricular tachycardia.

Ethics & Governance

Model training used de-identified public datasets and curated image sources. No direct patient identifiers were included in the training corpus. Deployment environments must still implement local data protection controls for uploaded ECG images. The system is positioned as a clinical decision-support tool with mandatory human review. Outputs are intended to assist, not replace, clinician judgement. Appropriate use requires qualified oversight. Open-source release of model weights and inference code supports transparency and independent audit. Known risks for multimodal language models include hallucinated or over-confident text generation. Mitigation approaches used include instruction-grounded fine-tuning and evaluation focused on consistency with visual evidence, but residual risk cannot be eliminated. Users should treat outputs as advisory. Equity considerations include potential variation in performance across device types, print formats, and population characteristics.



Discussion

Successes

The project demonstrates that multimodal foundation models can be adapted for ECG image interpretation tasks using parameter-efficient fine-tuning approaches. Instruction-style domain datasets supported generation of structured explanatory outputs rather than simple labels.

Challenges

Data representativeness remains a central constraint. Ensuring that training data reflect real-world ECG image variability, including lighting, distortion, and print artefacts, required extensive augmentation and curation. Generative medical AI models carry hallucination and over-generalisation risks. Mitigation strategies reduce but do not eliminate this risk. Subtle waveform abnormalities were more difficult for the model to classify reliably and required iterative retraining.

Evaluation of generative clinical text requires hybrid approaches combining automated scoring and human expert review; reliance on a single metric is insufficient.

Lessons Learned

Dataset quality and instruction design strongly influence model behaviour. Domain-specific instruction tuning contributed more to usable outputs than base model scale alone. Hybrid evaluation combining automated metrics and clinician review is necessary for credible assessment. Open release can accelerate feedback and independent validation, but governance guidance remains necessary to support responsible use.

Conclusion

ECG-LLaMA illustrates how a multimodal foundation model can be adapted using parameter-efficient fine-tuning to support ECG image interpretation and structured reporting. Validation results indicate strong technical performance under test conditions, and pilot use suggests feasibility as a clinician decision-support aid. The system is not a substitute for cardiology expertise and should be used only with qualified human oversight. As part of a supervised workflow, such models may complement clinician interpretation capacity, particularly in settings with limited specialist access, when implemented with appropriate governance and safeguards.



References

Mukkunoth, N., et al. (2025). High-accuracy ECG image interpretation using parameter-efficient low-rank adaptation (LoRA) fine-tuning with multimodal LLaMA V.3.2. *BMJ Digital Health & AI*, 1, e000031. <https://doi.org/10.1136/bmjdhai-2025-000031>

Meta AI. (2024). *Llama 3*. <https://llama.meta.com/llama3/>

Hu, E. J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., & Chen, W. (2022). LoRA: Low-rank adaptation of large language models. In *Proceedings of the International Conference on Learning Representations (ICLR)*. <https://arxiv.org/abs/2106.09685>


Liu, R., Bai, Y., Yue, X., & Zhang, P. (2024). *Teach multimodal LLMs to comprehend electrocardiographic images*. arXiv. <https://arxiv.org/abs/2410.19008>

Salerno, S. M., Alguire, P. C., & Waxman, H. S. (2003). Competency in interpretation of 12-lead electrocardiograms: A summary and appraisal of published evidence. *Annals of Internal Medicine*, 138(9), 751–760.



Leading with AI: Kerala's Nayanamritham 2.0 – A Government- Led AI-Assisted Chronic Eye Disease Screening Initiative

Anusha Purushotham, Karthika Kumar,
Bipin Gopal, Anand Sivaraman, Divya Rao



Leading with AI: Kerala's Nayanamritham 2.0 – A Government-Led AI-Assisted Chronic Eye Disease Screening Initiative

Anusha Purushotham¹, Bipin Gopal², Karthika Kumar¹, Anand Sivaraman¹, Divya Rao^{1*}

Abstract

Kerala has been an early adopter of artificial intelligence (AI) within public health service delivery, integrating AI-enabled tools into government-led screening programmes to address chronic disease burden. Nayanamritham 2.0 is a state-led initiative that incorporates AI-assisted retinal image analysis into routine public-sector eye care to support the early detection of diabetic retinopathy, glaucoma, and age-related macular degeneration.

Implemented across district hospitals, community health centres, and primary health centres, the programme uses an offline, AI-enabled screening system operated by trained optometrists, with clinical oversight maintained by ophthalmologists. AI functions as a decision-support tool to assist with screening and referral prioritisation, rather than as a diagnostic replacement.

This use case examines the evolution of Nayanamritham from pilot to programme, drawing on programme data and qualitative insights from frontline health workers and administrators. Using implementation frameworks, it highlights enabling factors, challenges, and governance considerations relevant to the responsible integration of AI into public health systems. Kerala's experience offers policy-relevant lessons for governments in similar settings seeking to scale AI-supported screening while maintaining equity, accountability, and human oversight.

Background


Burden of Chronic Eye Disease in India and Kerala

India carries a disproportionate share of the global burden of visual impairment, with an estimated 8 million people who are blind and more than 60 million living with moderate to severe visual impairment, a substantial proportion of which is preventable through early detection

and timely treatment. Chronic eye conditions particularly diabetic retinopathy (DR), glaucoma, and age-related macular degeneration (AMD) are contributing an increasing share of this burden as the country undergoes rapid epidemiological and demographic transition.

Diabetic retinopathy is closely associated with India's expanding diabetes prevalence. National

*Remidio Innovative Solutions Pvt Ltd ²Department of Health Services, Government of Kerala



estimates indicate that approximately 12–17% of adults with diabetes have some degree of DR, with about 4% affected by sight-threatening disease. Kerala represents a high-risk setting, with adult diabetes prevalence estimated at 15–20%, among the highest in the country. Community- and facility-based studies suggest that more than 90% of individuals with DR are unaware of their condition at the time of diagnosis, reflecting the typically asymptomatic progression of early disease and limited routine screening coverage.

Glaucoma presents a similar public health challenge. Often asymptomatic until advanced stages, it accounts for an estimated 5–6% of blindness in India, although the true burden is likely higher due to under-detection and limited population-level screening. AMD, historically under-recognised in low- and middle-income settings, is becoming more prominent in Kerala as life expectancy rises beyond 75 years, increasing the prevalence of age-related ocular disorders. Across these conditions, delayed diagnosis and late presentation remain common, contributing to avoidable and irreversible vision loss as well as higher long-term health system costs.

Health System and Policy Context in Kerala

India's public health system has traditionally relied on specialist-led eye care concentrated at secondary and tertiary levels. Retinal screening capacity at the primary care level has been limited, and shortages of ophthalmologists constrain the feasibility of universal screening using conventional models.


Kerala's health system provides a conducive environment for innovation, characterised by strong public-sector delivery, decentralised

governance, and a policy emphasis on prevention and primary care strengthening. Alignment with national initiatives, including the National Programme for Control of Blindness and Visual Impairment (NPCB&VI) and the Ayushman Bharat Digital Mission (ABDM) has enabled experimentation with digitally supported service delivery models. Within this context, Nayanamritham was conceived as a pathway to decentralise eye screening while maintaining clinical quality and accountability.

AI Use Case Description

Overview of Nayanamritham 1.0 and Nayanamritham 2.0

Nayanamritham 2.0 represents the evolution of Kerala's earlier diabetic retinopathy screening pilot into a state-supported, AI-assisted chronic eye disease screening programme embedded within routine services. The original Nayanamritham pilot was implemented in 2019 across 16 Family Health Centres (FHCs) in Thiruvananthapuram district as a tele-ophthalmology-supported screening pathway. A total of 5,307 individuals with diabetes were screened, with 85.3% of retinal images assessed as gradable. The age- and sex-standardised prevalence of any DR was 17.4%, with sight-threatening DR identified in 3.3% of screened individuals, indicating a substantial burden of previously undiagnosed disease and gaps in routine primary-level screening. The pilot demonstrated a substantial burden of previously undiagnosed diabetic retinopathy and highlighted the feasibility of primary-care-based screening.



However, reliance on specialist image grading and connectivity-dependent workflows posed challenges for sustainability at scale.

Building on these findings, the Government of Kerala launched Nayanamritham 2.0 in 2025 to expand disease coverage beyond diabetic retinopathy and to integrate AI-assisted screening into frontline service delivery. The programme aims to decentralise screening, support early identification of referable disease, and optimise specialist workload, while retaining human oversight and referral confirmation within the public health system.

Solution Description: MediosHI 3-in-1 Edge AI

MediosHI 3-in-1 Edge AI is a software-as-a-medical-device (SaMD) that analyses retinal images to support screening for diabetic retinopathy, glaucoma, and age-related macular degeneration. The system uses convolutional neural networks trained on large datasets of retinal images graded by expert ophthalmologists using established clinical standards. The AI algorithms were developed and validated using a dataset of over 200,000 retinal images reflecting variability across ethnic groups, disease severity, clinical presentation, imaging environments, and fundus camera systems, including images sourced from established retinal datasets such as AREDS and EyePACS. All images were annotated by qualified ophthalmologists using internationally recognised grading and classification standards, providing a reliable reference standard for training and performance evaluation.

The AI application is embedded in portable, smartphone-based fundus cameras and

operates entirely offline, enabling on-device analysis at the point of care. This design addresses connectivity constraints common in public-sector facilities and reduces dependence on cloud-based data transfer.

AI outputs indicate whether findings are likely to be referable and are accompanied by visual explanations highlighting image regions contributing to the assessment. These outputs are intended to support optometrists in screening and referral decisions, with final clinical confirmation and management undertaken by ophthalmologists at higher-level facilities.

The system architecture of the MediosHI 3-in-1 AI application is illustrated in Figure 1, demonstrating how offline AI enables decentralised, high-volume screening within public health settings.

Deployment and Implementation

The programme is implemented across public health facilities, where eligible individuals are identified through outpatient services and non-communicable disease clinics. Trained optometrists or allied personnel capture retinal images and generate AI-assisted screening outputs during routine visits.

AI-generated outputs inform counselling and referral initiation but do not replace clinical decision-making. Referred cases are reviewed by ophthalmologists, ensuring a human-in-the-loop model throughout the care pathway. Qualitative feedback from providers indicates that AI-supported screening has reduced ambiguity at the primary care level and improved the quality of referrals reaching secondary facilities.

Figure 01




Scale of Deployment

By 2024–2025, more than 250 fundus cameras had been deployed across Kerala’s public health system, with AI-assisted screening operational in over 50 centres. Approximately 30 centres are currently under active monitoring through routine reporting mechanisms, with phased scale-up underway. More than 100 optometrists have completed structured training programmes to support service delivery. Strong state ownership,

supported by dedicated budgetary allocations, has enabled Nayanamritham 2.0 to transition from time-bound pilot initiatives to routine service delivery—an essential determinant of long-term sustainability within public health systems.

Identified Users and Beneficiaries

Primary users include optometrists conducting screening and ophthalmologists responsible



for confirmation and treatment. Programme managers at district and state levels use aggregated data to support monitoring and planning. The primary beneficiaries are patients accessing public health services, particularly those in rural and underserved areas who previously faced barriers to specialist eye care.

Impact

Observed outcomes from Nayanamritham 2.0 indicate improvements in service delivery processes and workforce functioning. AI-assisted screening has supported higher screening volumes within existing staffing structures and more consistent identification of referable disease compared with earlier manual workflows.

Detection patterns suggest that decentralised screening is identifying previously undiagnosed cases. AI-supported triaging has contributed to more focused use of ophthalmologist time, allowing prioritisation of higher-risk cases. Workforce feedback highlights increased confidence among optometrists and improved communication with patients through visual screening outputs.

These observations reflect operational and system-level benefits rather than definitive population-level health impact. Ongoing monitoring and longer-term evaluation are required to assess effects on disease progression, treatment uptake, and visual outcomes.

Ethics and Governance

Ethical and governance considerations are integral to Nayanamritham 2.0. The system architecture prioritises data protection through offline, on-device processing, with data stewardship retained within government systems. AI functions strictly as a screening and triage support tool, with human oversight embedded at all stages.

Transparency is supported through interpretable outputs, while accountability is maintained through government ownership, procurement oversight, and integration within existing clinical governance structures. These arrangements align with principles for responsible AI use in public health, including human oversight, transparency, and equity.

250+ fundus cameras deployed

across Kerala's public health system (2024–2025)

100+ optometrists trained

through structured capacity-building programmes

Table 01: NASSS Framework as applied to Nayanamritham 2.0

Time (IST)	Session
Condition (C)	High prevalence of diabetic retinopathy (17.4%) and sight-threatening diabetic retinopathy (3.3%) among public-sector patients, combined with baseline awareness less than 1%, establishes DR as a high-impact, under-addressed chronic condition requiring systematic screening (JCM 2021).
Technology (T)	Offline, on-device AI enables reliable screening in settings with variable connectivity, designed to support consistent performance while addressing infrastructure and data-privacy constraints in public health facilities.
Value Proposition (V)	An economic evaluation estimated an ICER of approximately INR 22,000 per QALY, which the authors consider favourable under WHO cost-effectiveness thresholds.
Adopters (A)	Optometrists reported improved diagnostic confidence and professional empowerment, while ophthalmologists experienced workload optimisation without loss of clinical oversight, supporting cross-cadre acceptance (qualitative findings).
Organisation (O)	Integration within Non-Communicable Disease (NCD) clinics and Family Health Centres leveraged existing workflows validated during the pilot phase, enabling institutional adoption within routine service delivery.
Wider System (S)	Alignment with the Aardram Mission, Ayushman Bharat Digital Mission (ABDM) principles, and WHO NCD strategies facilitated policy coherence and system-level integration.
Sustainability (S)	Transition from pilot-phase tele-grading to AI-assisted scale-up can reduce marginal screening costs and dependence on central reading centres, supporting long-term operational and financial sustainability.



Discussion

Success Factors

Key enabling factors include strong government leadership, alignment with public health priorities, and sustained collaboration between public institutions and technology developers. Designing the system for offline use proved critical in addressing real-world infrastructure constraints. Continuous training and engagement supported trust and adoption among frontline workers.

Challenges Encountered

The programme faced several implementation challenges. Initial scepticism among clinicians and e-health departments focused on AI reliability, data privacy, and accountability. Operational constraints included staff attrition, heterogeneity in facility-level infrastructure, fragmented screening and treatment data systems, and the absence of a fully centralised data governance framework—reflecting broader limitations in public-sector digital infrastructure.

Change management remained a continuous requirement. Integrating AI into routine public services necessitated sustained stakeholder engagement, workflow adaptation, and institutional reassurance, highlighting that AI adoption is fundamentally an organisational as well as a technological process.

These challenges align with broader evidence from public-sector AI deployments, where connectivity dependence, data transfer costs, and governance risks have been identified as

key constraints for cloud-based solutions in health settings. By enabling offline inference and retaining patient data within government systems, Nayanamritham 2.0 mitigated these risks while supporting continuity of care in facilities with variable digital capacity. This design approach aligns with the emerging consensus that AI for LMIC health systems must be tailored to real-world operating conditions rather than idealised digital environments.

Lessons Learned

The experience of Nayanamritham 2.0 highlights the importance of human–AI collaboration, the need to design for offline and low-resource settings, and the importance of governance and workforce training for sustained adoption. Scalability depends not only on algorithmic performance but also on governance, financing, and workforce readiness.

Evidence indicates that AI adoption falters when technologies are perceived as externally imposed or threatening to professional roles, but improves when AI enhances clinical agency. Embedding AI within existing public systems enhances sustainability but requires long-term commitment beyond pilot phases. Post-pandemic disruptions underscored these risks, highlighting the importance of long-term asset management, continuous training, and policy-backed financing. Ultimately, the programme's significance lies not in algorithmic novelty, but in the successful institutionalisation of AI within a public health system.




Conclusion

Nayanamritham 2.0 offers an example of how AI can be responsibly integrated into a public health system to support early detection of chronic disease at scale. By combining government stewardship, offline AI design, and empowered health workers, the programme strengthens

primary care and supports more equitable access to eye screening. Kerala's experience offers transferable lessons for other settings seeking to deploy AI as a decision-support tool within people-centred, accountable health systems.

References

- World Health Organization. (2019). *World report on vision*. World Health Organization.
- Ministry of Health and Family Welfare. (2020). *National blindness and visual impairment survey India 2015–2019: A summary report*. Government of India.
- Yau, J. W. Y., Rogers, S. L., Kawasaki, R., Lamoureux, E. L., Kowalski, J. W., Bek, T., Chen, S. J., & Wong, T. Y. (2012). Global prevalence and major risk factors of diabetic retinopathy. *Diabetes Care*, 35(3), 556–564. <https://doi.org/10.2337/dc11-1909>
- Raman, R., Ganesan, S., Pal, S. S., Kulothungan, V., Sharma, T., & Raman, R. (2012). Prevalence of diabetic retinopathy in India. *Indian Journal of Ophthalmology*, 60(3), 204–210. <https://doi.org/10.4103/0301-4738.95873>
- Anjana, R. M., Deepa, M., Pradeepa, R., Mahanta, J., Narain, K., Das, H. K., Adhikari, P., Rao, P. V., Saboo, B., Kumar, A., Bhansali, A., John, M., Luaia, R., Reang, T., Ningombam, S., Jampa, L., Budnah, R. O., Elangovan, N., Subashini, R., ... Mohan, V. (2017). Prevalence of diabetes and prediabetes in 15 states of India: Results from the ICMR–INDIAB population-based cross-sectional study. *The Lancet Diabetes & Endocrinology*, 5(8), 585–596. [https://doi.org/10.1016/S2213-8587\(17\)30174-2](https://doi.org/10.1016/S2213-8587(17)30174-2)
- Sivaprasad, S., Sahasranamam, V. I., George, S., et al. (2021). Burden of diabetic retinopathy amongst people with diabetes attending primary care in Kerala: The Nayanamritham Project. *Journal of Clinical Medicine*, 10(24), 5903. <https://doi.org/10.3390/jcm10245903>
- George, R., Ve, R. S., & Vijaya, L. (2010). Glaucoma in India: Estimated burden of disease. *Journal of Glaucoma*, 19(6), 391–397. <https://doi.org/10.1097/IJG.0b013e3181c4ac02>
- Wong, W. L., Su, X., Li, X., Cheung, C. M. G., Klein, R., Cheng, C. Y., & Wong, T. Y. (2014). Global prevalence of age-related macular degeneration and disease burden projection for 2020 and 2040. *The Lancet Global Health*, 2(2), e106–e116. [https://doi.org/10.1016/S2214-109X\(13\)70145-1](https://doi.org/10.1016/S2214-109X(13)70145-1)
- Rao, G. N., Khanna, R., Athota, S. M., Rajshekar, V., & Rani, P. K. (2012). Integrated model of primary and secondary eye care for underserved rural areas: The L V Prasad Eye Institute experience. *Indian Journal of Ophthalmology*, 60(5), 396–400. <https://doi.org/10.4103/0301-4738.100541>
- Blanchet, K., Gilbert, C., & de Savigny, D. (2014). Rethinking eye health systems to achieve universal coverage. *Bulletin of the World Health Organization*, 92, 386–387. <https://doi.org/10.2471/BLT.13.135343>
- Ministry of Health and Family Welfare. (2017). *National programme for control of blindness and visual impairment: Operational guidelines*. Government of India.
- National Health Authority. (2021). *Ayushman Bharat Digital Mission: Strategy overview*. Government of India.
- Conroy, D., Jeeva, I., Muralidharan, V., et al. (2021). The ORNATE India project: United Kingdom–India collaboration in retinal disease research. *Eye*. <https://doi.org/10.1038/s41433-021-01628-4>
- Express Healthcare News Bureau. (2025, February 18). Kerala government and Remidio launch Nayanamritham 2.0, an AI-assisted chronic eye disease screening programme. *Express Healthcare*.
- Wittenberg, R., et al. (2024). Costs and cost-effectiveness of the Kerala pilot screening programme for diabetic retinopathy in the public health system. *Eye*, 38(17), 3352–3356. <https://doi.org/10.1038/s41433-023-02642-1>
- Rudnicka, A. R., Shakespeare, R., Chambers, R., Bolter, L., Anderson, J., Fajtl, J., ... Willis, K. (2025). Automated retinal image analysis systems to triage for grading of diabetic retinopathy: A large-scale, open-label, national screening programme in England. *The Lancet Digital Health*, 7(11). <https://doi.org/10.1016/j.landig.2025.100914>
- Gulshan, V., Peng, L., Coram, M., et al. (2016). Development and validation of a deep learning algorithm for detection of diabetic retinopathy in retinal fundus photographs. *JAMA*, 316(22), 2402–2410. <https://doi.org/10.1001/jama.2016.17216>
- Ting, D. S. W., Cheung, C. Y., Lim, G., et al. (2017). Development and validation of a deep learning system for diabetic retinopathy and related eye diseases using retinal images from multiethnic populations. *JAMA*, 318(22), 22112223. <https://doi.org/10.1001/jama.2017.18152>
- Abràmoff, M. D., Lavin, P. T., Birch, M., Shah, N., & Folk, J. C. (2018). Pivotal trial of an autonomous AI-based diagnostic system for detection of diabetic retinopathy in primary care offices. *NPJ Digital Medicine*, 1, 39. <https://doi.org/10.1038/s41746-018-0040-6>
- Xie, Y., Nguyen, Q. D., Hamzah, H., et al. (2020). Artificial intelligence for teleophthalmology-based diabetic retinopathy screening: A systematic review. *The Lancet Digital Health*, 2(10), e527–e536. [https://doi.org/10.1016/S2589-7500\(20\)30102-4](https://doi.org/10.1016/S2589-7500(20)30102-4)



Rogers, T. W., Gonzalez-Bueno, J., Garcia, F. R., et al. (2020). Evaluation of artificial intelligence systems for diabetic retinopathy screening: Real-world considerations. *Eye*, 34(7), 1220–1230. <https://doi.org/10.1038/s41433-019-0650-4>

Chauhan, A., Vale, L., Kankaria, A., Gupta, V., Kaur, G., N, N., Sood, N., Tigari, B., Bhadada, S., & Duggal, M. (2025). Assessing adoption of human and AI-enabled diabetic retinopathy screening in primary healthcare settings: findings from a pragmatic trial. *Archives of public health = Archives belges de sante publique*, 83(1), 268. <https://doi.org/10.1186/s13690-025-01757-3>

Thomas, S. M., MacKenzie, G. E., et al. (2014). Task-shifting and task-sharing in eye care: A systematic review. *The Lancet Global Health*, 2(4), e234–e243. [https://doi.org/10.1016/S2214-109X\(14\)70040-8](https://doi.org/10.1016/S2214-109X(14)70040-8)

Wahl, B., Cossy-Gantner, A., Germann, S., & Schwalbe, N. R. (2018). Artificial intelligence and global health: How can AI contribute to health in resource-poor settings? *BMJ Global Health*, 3, e000798. <https://doi.org/10.1136/bmjgh-2018-000798>

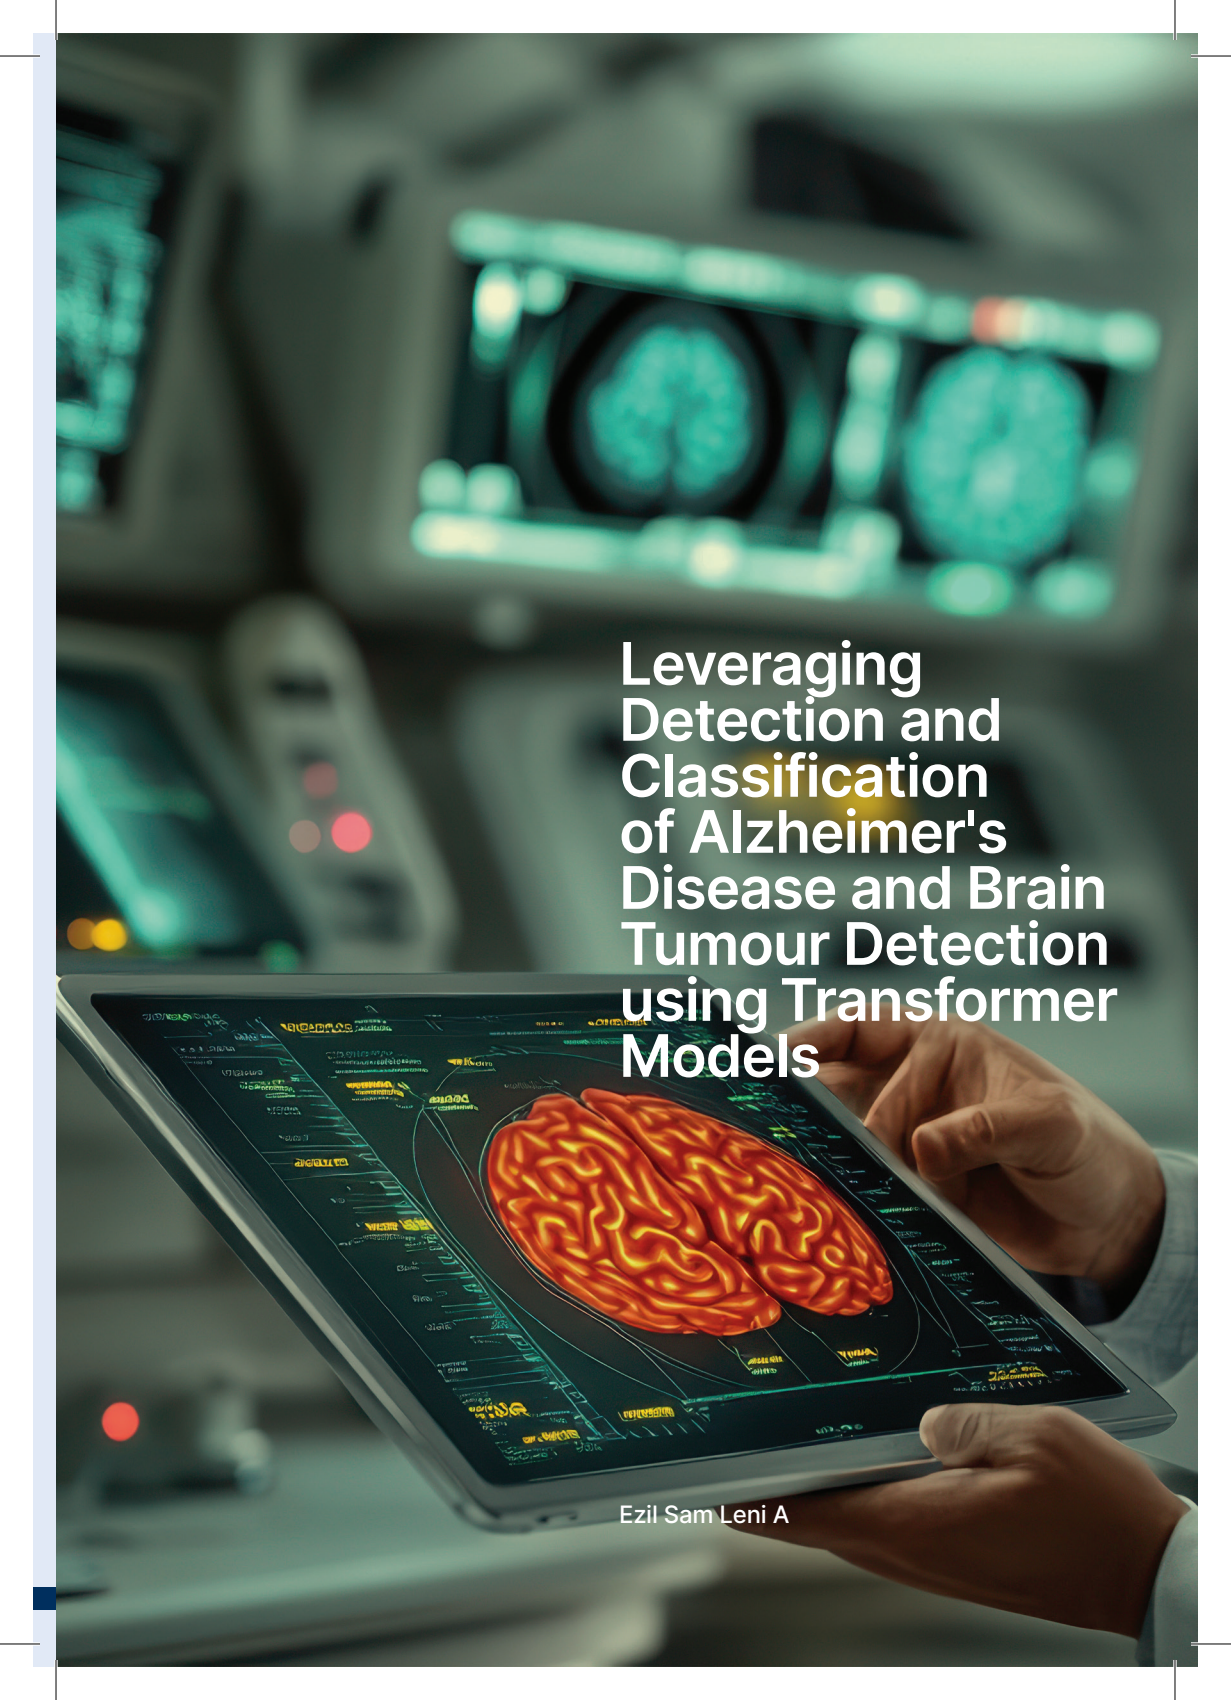
Mehta, P., Stahl, S., Germann, S., et al. (2021). Constraints to digital health adoption in low-resource settings. *Journal of Global Health*, 11, 04028. <https://doi.org/10.7189/jogh.11.04028>

Rajpurkar, P., Chen, E., Banerjee, O., et al. (2022). AI in health systems: The need for context-aware design. *NPJ Digital Medicine*, 5, 44. <https://doi.org/10.1038/s41746-022-00578-1>

Sendak, M. P., Ratliff, W., Sarro, D., et al. (2020). Real-world integration of artificial intelligence into clinical care. *NPJ Digital Medicine*, 3, 46. <https://doi.org/10.1038/s41746-020-0282-1>


Rogers, E. M. (2003). *Diffusion of innovations* (5th ed.). Free Press.

Greenhalgh, T., & Papoutsi, C. (2018). Studying complexity in health services research: Desperately seeking an overdue paradigm shift. *BMC Medicine*, 16, 95. <https://doi.org/10.1186/s12916-018-1089-4>



Leveraging Detection and Classification of Alzheimer's Disease and Brain Tumour Detection using Transformer Models

Ezil Sam Leni A



Leveraging Detection and Classification of Alzheimer's Disease and Brain Tumour Detection using Transformer Models

Ezil Sam Leni A*

Abstract

Alzheimer's disease (AD) is a progressive neurodegenerative condition and a leading cause of dementia worldwide. Brain tumors also represent a significant neurological disease burden requiring timely imaging-based evaluation. Early-stage detection and structured classification remain challenging in resource-constrained settings due to specialist shortages, diagnostic delays, and variability in interpretation. Deep learning methods have shown potential in MRI-based analysis, though many earlier approaches relied primarily on convolutional neural networks (CNNs) without integrated segmentation support or interpretability layers.

This use describes two AI-based neuroimaging pipelines: one for Alzheimer's disease stage classification and one for brain tumor segmentation and classification. The Alzheimer's pipeline combines CNN-based feature extraction with the Segment Anything Model (SAM) for segmentation-informed feature representation, using late-fusion classification. The brain tumor pipeline combines SAM-based region-of-interest segmentation with a Swin Transformer architecture for hierarchical feature learning. Both pipelines were evaluated on publicly available MRI datasets under controlled experimental conditions.

Reported results indicate improved classification metrics compared with selected standalone baseline models within the study datasets. The systems are designed to support radiological review and research workflows, not replace clinical judgement. This use case illustrates how hybrid CNN-transformer architectures with segmentation support may enhance AI-assisted neuroimaging analysis under appropriate governance and oversight.

*Department of CSE, Alliance University, Karnataka



Background

With a population of approximately 1.4 billion, India faces a growing burden of age-related neurological conditions. An estimated 8.8 million older adults are affected by Alzheimer's disease, and more than 7.4% of the population is aged over 60, underscoring the need to strengthen dementia care infrastructure. Brain tumours also represent a significant and rising health concern, with roughly 30,000 new cases reported annually and an estimated incidence of 5–10 per 100,000 population.

MRI is central to both dementia evaluation and tumor detection. However, interpretation requires specialist expertise and time. In many health systems particularly in low- and middle-income settings, there are shortages of trained neuroradiologists and neurologists. This can contribute to diagnostic delays and uneven access to expert review.

AI-based image analysis methods are being explored to support radiologists by assisting with segmentation, pattern recognition, and classification tasks. Earlier work has often focused on CNN-only architectures, which can perform well in feature extraction but may provide limited anatomical localisation or segmentation interpretability. More recent transformer-based and hybrid approaches enable segmentation-aware and multi-scale feature modelling.

This investigation aligns with national digital and AI health missions by exploring how hybrid deep learning and transformer models may support MRI-based analysis for Alzheimer's disease staging and brain tumor classification. These

tools are intended as decision-support aids for imaging interpretation, not autonomous diagnostic systems.

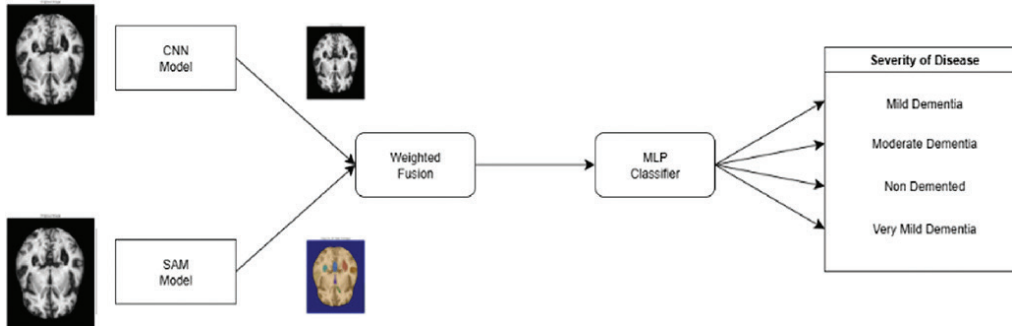
AI Use Case

Two related AI pipelines are described: an Alzheimer's disease classification pipeline and a brain tumor segmentation–classification pipeline.

Alzheimer's Pipeline: Alzheimer's disease investigation used the OASIS dataset, which contains T1-weighted brain MRI images categorized as non-demented, very mild, mild, and moderate dementia. Input images are processed through a convolutional neural network (CNN) to extract spatial features and a SAM module to derive semantic feature representations. These features are combined using a weighted averaging fusion approach and classified using a multilayer perceptron (MLP). Model performance was evaluated against baseline approaches. The overall Alzheimer's pipeline architecture is illustrated in Figure 01.

Uses hybrid pipelines combining CNN, SAM and Swin Transformers for MRI analysis

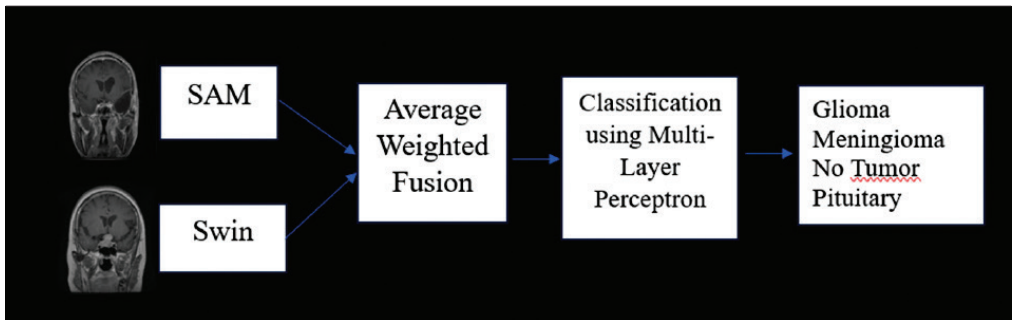
Figure 01: System Architecture-Alzheimer's Investigation



Brain Tumor Pipeline:

MRI images are first segmented using SAM to identify candidate regions of interest. These segmented regions are then processed by a Swin Transformer model to extract hierarchical multi-scale features for tumor-type classification across four categories (glioma, meningioma, pituitary tumor, no tumor).

Figure 02: System Architecture- Brain Tumor Investigation



Both pipelines are evaluated against selected baseline single-model approaches. Results are reported as comparative model performance under experimental conditions rather than clinical deployment outcomes.

Solution Description

Alzheimer's Investigation

The Alzheimer's pipeline follows a staged workflow: dataset curation, preprocessing, model training, feature fusion, and evaluation. CNNs are used for hierarchical spatial feature extraction from MRI images. SAM is used to generate segmentation masks using prompt-based or zero-shot segmentation capabilities. Segmentation outputs are converted into feature representations. CNN and SAM-derived features are fused using weighted averaging and classified via MLP. The OASIS dataset is used (Table 1), containing T1-weighted MRI images across four dementia severity categories. Sample instances of non-demented and moderately demented brain images are shown in Figures 3a and 3b, respectively. Class imbalance is present and acknowledged as a limitation.

Figure 03 a): Non-Demented Brain




Figure 03 b): Moderately Demented Brain




Table 01: Distribution of classes in Oasis Dataset

Class	Instances
Mild Dementia	5002
Moderate Dementia	488
Non-Demented	67,222
Very Mild Dementia	13,725

Brain Tumor Investigation

The brain tumor pipeline uses SAM for segmentation and a Swin Transformer for classification. Swin Transformer architectures apply shifted-window attention to capture multi-scale contextual features efficiently. The Kaggle Brain Tumor MRI dataset is used (Table-2), with four classes. SAM is used to generate region-of-interest masks with minimal supervision. The Swin model processes segmented inputs for classification. Feature fusion and classification are again performed using an MLP layer.

Table 02: Distribution of classes in Kaggle Brain MRI Dataset

Tumor Types	Instances
Glioma	1,621
Meningioma	1,645
Pituitary Tumor	1,757
No Tumor	2000

Across both pipelines, preprocessing includes resizing, normalization, and stratified dataset splitting. Results reflect performance within these curated datasets. (Table-3)

Table 03: Abstract of AI Methodology and technique used

Stages	Alzheimer's Investigation	Brain Tumour Investigation
Data Collection	OASIS Alzheimer's Datasets	Kaggle Brain Tumor MRI Dataset
	T1-weighted brain MRI images	T1-weighted brain MRI slices
Data Preprocessing	Uses PyTorch's 'transforms.Compose' pipeline.	Uses PyTorch's 'transforms.Compose' pipeline.
	Uniform resizing 128×128 pixels.	Uniform resizing 128×128 pixels.
	Training, validation, and testing ratios of 70-15-15.	Training, validation, and testing ratios of 70-15-15.



Model Training	Experimental Setup: PyTorch and TensorFlow in Jupyter Notebooks	Experimental Setup: PyTorch and TensorFlow in Jupyter Notebooks
	Models Used: CNN, SAM	Models Used: SAM, Swin
	Hyperparameters Used:	Hyperparameters Used:
	Batch Size -32	Batch Size -32
	Epochs (OASIS) -15	Epochs -15
	Dropout - 0.3 (CNN MLP), 0.1-0.3 (SAM)	Dropout - 0.3 0.1-0.3 (SAM), 0.1-0.3 (Swin)
	Fusion Weight $\alpha=0.5$ (balanced contribution)	Fusion Weight $\alpha=0.5$ (balanced contribution)
	Learning Rate - 0.001	Learning Rate - 0.001
Model Evaluation	Confusion Matrix Analysis: The CNN + SAM model demonstrates improved effectiveness over the standalone CNN model.	Confusion Matrix Analysis: The confusion matrices of SAM + Swin model demonstrate the effectiveness over the standalone models.
	Segmentation Metrics: Empirical findings indicated that CNN+SAM outperformed the standalone transformer models.	Segmentation Metrics: Empirical findings indicated that SAM+Swin outperformed the standalone transformer models.
	Classification Metrics: The CNN+SAM model showed an increase of 18% improvement in accuracy over the standalone CNN model.	Classification Metrics: The SAM +Swin model showed an increase of 20% improvement in accuracy over the standalone CNN model.

Deployment & Implementation:

Development workflows include dataset curation, preprocessing standardisation, model training, hyperparameter tuning, and evaluation using held-out test splits. Both pipelines are described as prototype-ready systems suitable for controlled deployment testing rather than validated clinical tools. (Table-4)

Designed as decision support for imaging workflows - not a stand alone diagnostic tool

Table 04: Abstract of AI Methodology and technique used

Process/Approach/ Workflow	Alzheimer's Investigation	Brain Tumour Investigation
Development process (To build and validate the AI model for creating a robust prototype)	<ul style="list-style-type: none"> i) Curate the OASIS & MRI 4-Classes Dataset as non-demented, very mild dementia, mild dementia, and moderate dementia. ii) Preprocess the images and resize them to 124 × 124. iii) Train the CNN and SAM Models and fuse using weighted averaging. iv) Classify the fused features with MLP. v) Evaluate the performance and fine tune it for improvement. 	<ul style="list-style-type: none"> i) Curate the Kaggle MRI images such as glioma, meningioma, pituitary, and no-tumor classes. ii) Preprocess the images and resize them to 124 × 124. iii) Train the SAM and Swin Models and fuse using weighted averaging. iv) Classify the fused features with MLP. v) Evaluate the performance and fine tune it for improvement.
Deployment approach:	<ul style="list-style-type: none"> i) Convert the trained CNN and SAM model to a standard REST/DICOMweb APIs. ii) Integrate the model with PACS and radiology stations. iii) Deploy as per IndiaAI guidelines in hospital data centers or cloud infrastructure. 	<ul style="list-style-type: none"> i) Convert the trained SAM and Swin model to a standard REST/DICOMweb APIs. ii) Integrate the model with PACS and radiology stations. iii) Deploy as per IndiaAI guidelines in hospital data centers or cloud infrastructure.
Operational workflow (For live clinical deployment)	<ul style="list-style-type: none"> i) Upload the MRI images either automatically or by the radiologists. ii) Execute the CNN and SAM AI pipeline and infer the results. iii) Review the inference with the radiologists. iv) Generate the final reports signed by the radiologists. Track the real-time performance for trustworthy deployment as per India AI guidelines. 	<ul style="list-style-type: none"> i) Upload the MRI images either automatically or by the radiologists. ii) Execute the SAM and Swin Alpipeline and infer the results. iii) Review the inference with the radiologists. iv) Generate the final reports signed by the radiologists. Track the real-time performance for trustworthy deployment as per India AI guidelines.

Identified Users

Primary users include radiologists, neurologists, and neuroimaging researchers. Secondary users include academic research groups and imaging centres exploring AI-assisted workflows. Beneficiaries may include health institutions seeking decision-support tools for MRI interpretation and research settings studying automated neuroimaging analysis. These tools are not intended for unsupervised or patient-facing use.

Impact

Reported performance metrics are based on experimental dataset evaluation. The

Alzheimer's disease investigation reported the following metrics on the OASIS dataset: accuracy of 0.95, F1-score of 0.94, precision of 0.96, and recall of 0.93 using the integrated CNN+SAM feature framework. Comparative results between the standalone CNN model and the combined CNN+SAM model are presented in Table 5.

The SAM-Swin brain tumour detection and classification study reported robust performance across precision, recall, accuracy, and F1-score metrics using zero-shot segmentation and hierarchical feature attention. Detailed results are provided in Table 6.


Table 05: Performance Analysis-Alzheimer's Investigation

Model	Accuracy	F1-Score	Precision	Recall
CNN	0.87	0.85	0.86	0.84
CNN+SAM	0.95	0.94	0.96	0.93

Table 06: Performance Analysis-Brain tumor Investigation

Tumor Type	Precision	Recall	F1 Score	Accuracy
Glioma	95.90	95.45	95.49	95.30
Meningioma	96.5	96.75	96.5	96.75
No Tumor	97.09	97.32	97.09	97.09
Pituitary Tumor	95.62	95.21	94.99	95.62

Figure 04: Heatmap Using Grad-CAM - Brain Tumor Investigation



For the brain tumor pipeline, class-wise precision, recall, F1-score, and accuracy values are reported across tumor categories. Grad-CAM heatmaps are used to support visual interpretability of model attention regions.

These metrics represent technical model performance under controlled dataset conditions. They do not establish clinical diagnostic accuracy, patient outcome improvement, or workflow efficiency gains in real-world settings. External validation and prospective clinical evaluation would be required for such claims.

Ethics & Governance

Ethical and governance considerations for the proposed deployments focus on informed consent, data privacy, and secure handling of MRI data. Data management is designed to support localisation requirements, role-based access controls, and DICOM-aligned de-identification. Known dataset limitations are explicitly recognized, including demographic imbalance in the Alzheimer's dataset and class imbalance with uncertain demographic representativeness in the tumour dataset. Mitigation measures include expanded and more diverse data collection and routine subgroup performance monitoring. The approach follows responsible AI practices, including fairness review, transparency, and accountability consistent with WHO guidance on AI ethics, and recommends alignment with national AI and digital health policies. Human clinical oversight remains mandatory for all diagnostic or decision-support use.



Discussion

Successes

Hybrid CNN–transformer and segmentation-informed models showed improved experimental performance compared with selected standalone baselines within the study datasets. The architectures also support explainability techniques and segmentation overlays, features that are valued by radiologists for visual validation and clinical interpretability.

Challenges

Key limitations include dataset bias, class imbalance, variable image quality, and limited generalisability across scanners and population groups. Zero-shot segmentation performance can be inconsistent in atypical or low-quality scans. Additional technical complexity arises in multimodal image registration and pediatric imaging contexts.

Lessons learned

Explainability and interpretability are central to clinician trust and uptake; performance metrics alone are not sufficient for adoption. Federated and other privacy-preserving learning approaches may help meet regulatory and data-governance requirements. Robust deployment requires diverse preprocessing strategies and continuous bias and subgroup performance monitoring.



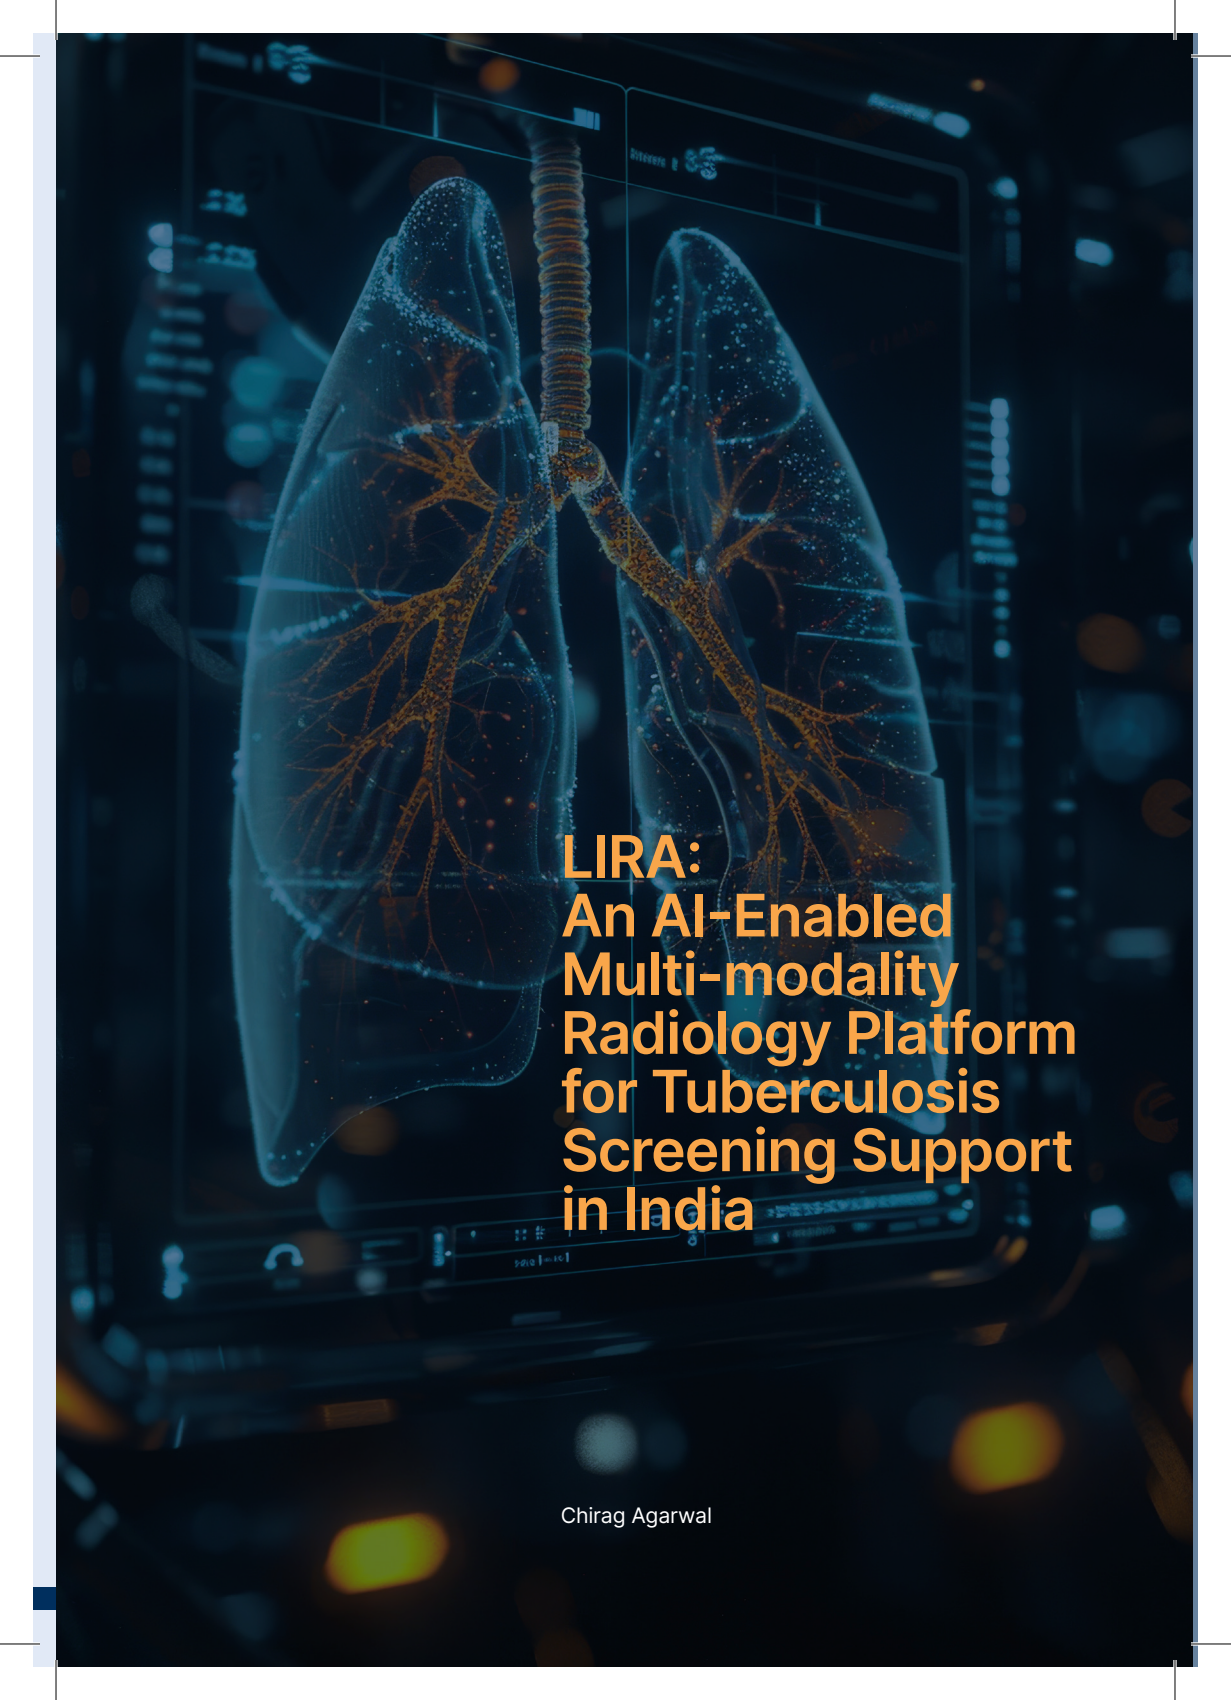
Conclusion

The Alzheimer's disease use case employs an integrated CNN-SAM framework, with experimental results indicating improved performance compared with a standalone CNN model. The brain tumour use case combines SAM with a Swin transformer to capture semantic and multi-scale features for detection and classification across glioma, meningioma, pituitary, and no-tumour categories. Across both use cases, evaluation metrics including accuracy, F1-score, precision, and recall showed encouraging results. The work documents the end-to-end AI pipeline, from model development through deployment strategy and operational workflow.

Primary beneficiaries include neurologists, radiologists, medical researchers, research laboratories, and pharmaceutical groups. Deployment and operationalization are expected to follow ethical AI and regulatory guidance, with emphasis on patient safety, data privacy, and clinical trust.



References

- Alharbi, D., Almulifi, A., Alharthi, M., Aljuaid, N., Alshammari, N., & Althobaiti, M. (2025). Brain tumor segmentation using transformer-based U-Net and deep explainable AI. *Computers in Biology and Medicine*, *173*, 108274. <https://doi.org/10.1016/j.combiomed.2025.108274>
- Almalki, H., Khadidos, A. O., Alhebaishi, N., & Senan, E. M. (2025). Early detection of Alzheimer's disease progression stages using hybrid of CNN and transformer encoder models. *Scientific Reports*, *15*(1), 16799. <https://doi.org/10.1038/s41598-025-01072-5>
- Asiri, A. A., Shaf, A., Ali, T., Pasha, M. A., Khan, A., Irfan, M., Alqahtani, S., Alghamdi, A., Alghamdi, A. H., Alshamrani, A. F. A., Alelyani, M., & Alamri, S. (2024). Advancing brain tumor detection: Harnessing the Swin Transformer's power for accurate classification and performance analysis. *PeerJ Computer Science*, *10*, e1867. <https://doi.org/10.7717/peerj-cs.1867>
- Gao, H., Wan, Y., Xu, H., Chen, L., Xiao, J., & Ran, Q. (2025). SwinBTC: Transfer learning to brain tumor classification for healthcare electronics using augmented MR images. *IEEE Transactions on Consumer Electronics*, *71*(1), 2297–2308. <https://doi.org/10.1109/TCE.2025.3527061>
- Herington, J., McCradden, M. D., Creel, K., Boellaard, R., Jones, E. C., Jha, A. K., Rahmim, A., Scott, P. J. H., Sunderland, J. J., Wahl, R. L., Zuehlsdorff, S., & Saboury, B. (2023). Ethical considerations for artificial intelligence in medical imaging: Data collection, development, and evaluation. *Journal of Nuclear Medicine*, *64*(12), 1848–1854. <https://doi.org/10.2967/jnumed.123.266080>
- Hossain, S., Chakrabarty, A., Gadekallu, T. R., Alazab, M., & Piran, M. J. (2024). Vision transformers, ensemble model, and transfer learning leveraging explainable AI for brain tumor detection and classification. *IEEE Journal of Biomedical and Health Informatics*, *28*(3), 1261–1272. <https://doi.org/10.1109/JBHI.2023.3266614>
- Hu, Z., Li, Y., Wang, Z., Zhang, S., & Hou, W. (2023). Conv-Swinformer: Integration of CNN and shift window attention for Alzheimer's disease classification. *Computers in Biology and Medicine*, *164*, 107304. <https://doi.org/10.1016/j.combiomed.2023.107304>
- Ingole, B. S., Ramineni, V., Pulipeta, N. K., Kathiriya, M. J., Krishnappa, M. S., & Jayaram, V. (2024). *The dual impact of artificial intelligence in healthcare: Balancing advancements with ethical and operational challenges*(arXiv:2411.16691). arXiv. <https://arxiv.org/abs/2411.16691>
- Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A. C., Lo, W., Dollár, P., & Girshick, R. (2023). *Segment anything* (arXiv:2304.02643). arXiv. <https://arxiv.org/abs/2304.02643>
- Li, Z., Silamu, W., Wang, Y., & Zhe, W. (2023). DenseTrans: Multimodal brain tumor segmentation using Swin Transformer. *IEEE Access*, *11*, 42895–42908.
- Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., & Guo, B. (2021). Swin Transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)* (pp. 9992–10002). <https://doi.org/10.1109/ICCV48922.2021.00986>
- Mahmud, S. M. H., Ali, M. M., Shahriar, M. F., Al-Zahrani, F. A., Ahmed, K., et al. (2023). Detection of different stages of Alzheimer's disease using CNN classifier. *Computers, Materials & Continua*, *76*(3), 3933–3948. <https://doi.org/10.32604/cmc.2023.039020>
- Mishra, S. S., Ahirwar, M. K., Nanda, S., & Rath, S. M. (2025). Epidemiology of primary intracranial tumors at a tertiary care hospital in central India: A retrospective study. *Cureus*, *17*(7), e88404. <https://doi.org/10.7759/cureus.88404>
- Nickparvar, M. (n.d.). *Brain tumor MRI dataset* [Data set]. Kaggle. <https://www.kaggle.com/datasets/masoudnickparvar/brain-tumor-mri-dataset>
- Open Access Series of Imaging Studies (OASIS). (2007). *OASIS-1: Cross-sectional MRI dataset*. <https://doi.org/10.1162/jocn.2007.19.9.1498>
- World Health Organization. (n.d.). *Dementia*. <https://www.who.int/news-room/fact-sheets/detail/dementia>



LIRA:
An AI-Enabled
Multi-modality
Radiology Platform
for Tuberculosis
Screening Support
in India

Chirag Agarwal



LIRA: An AI-Enabled Multi-modality Radiology Platform for Tuberculosis Screening Support in India

Chirag Agarwal*

Abstract

Access to diagnostic imaging in India remains unevenly distributed, with a substantial share of radiology infrastructure concentrated in urban centres while rural and underserved populations face limited availability of imaging and specialist interpretation. Workforce constraints, including a low radiologist-to-population ratio of 1:100,000 (average doctor-to-population ratio in India: 1:834) contribute to reporting delays and variability in access to timely image interpretation. These constraints affect tuberculosis (TB) screening pathways, where chest imaging plays an important role in early detection and triage.

LIRA is an AI-enabled radiology decision-support platform designed to assist with triage, interpretation support, and structured reporting for adult chest X-rays (CXR) and low-dose CT scans (LDCT). The system generates abnormality classifications, visual heatmaps, and draft structured outputs within DICOM-compliant workflows. Models have been trained on large annotated Indian datasets developed with research partners including National Institute for Research in Tuberculosis, ICMR and evaluated across multiple clinical sites.

Validation studies conducted across partner hospitals and research centres report high sensitivity and moderate-to-high specificity for abnormality detection and TB-suggestive pattern detection under study conditions. Real-world deployment across multiple facilities indicate operational feasibility in diverse care settings. Lenek Technologies partnered with hospitals and health centres including UPUMS Hospital, Ashwini Hospital, CHC Tonk, Leelamani Hospital, IIT Kanpur Health Center, and Saint John's Research Institute to validate performance using approximately 9,790 confirmed TB cases. Evaluation results showed a sensitivity of 99% and specificity of 92% for abnormality detection, and a sensitivity of 98.5% with specificity of 80% for tuberculosis detection. Together, these findings suggest that the solution may be compatible with existing imaging infrastructure in both urban and resource-constrained environments.

* Lenek Technologies, Uttar Pradesh, India



Background

LIRA was developed within the context of national efforts to strengthen indigenous medical technology capacity and reduce dependence on imported diagnostic devices. The initiative emerged from collaboration across academic, research, and startup ecosystems aligned with national med-tech innovation missions, including industry-academia engagement at IIT Kanpur and incubation support through SIIC IIT Kanpur. Lenek Technologies focused on tuberculosis (TB) screening after field assessments highlighted persistent diagnostic gaps, particularly in rural and underserved settings.

Early observations from district hospitals and community health centres showed that delays in TB diagnosis were frequently linked to slow imaging interpretation rather than lack of treatment availability. Chest X-rays were often acquired locally but interpreted remotely, with reporting delays that could extend to several days and affect follow-up testing and continuity of care.

Diagnostic imaging infrastructure in India remains unevenly distributed, with a large proportion concentrated in metropolitan areas while rural and peri-urban facilities operate with limited radiologist availability. Approximately 75% of India's diagnostic imaging infrastructure is concentrated in urban centres that serve only about 25% of the population, resulting in limited access to timely radiology services in rural and underserved areas. The radiologist-to-population ratio estimated at around 1:100,000 further intensifies this gap, with a relatively small specialist workforce managing high daily

reporting volumes, often exceeding 100 scans across multiple facilities. For the National TB Elimination Programme (NTEP), which depends significantly on chest imaging for early case detection and triage, these structural constraints contribute to delays between initial presentation and confirmed diagnosis


The guiding design question for LIRA's development was whether an AI system could be trained on locally representative datasets to support chest imaging interpretation and triage in settings without on-site radiologists, while maintaining clinician oversight and compatibility with existing radiology systems. LIRA was therefore designed as a software layer that integrates with current imaging workflows rather than replacing them.

AI Use Case Description

Solution Description

LIRA is an AI-enabled radiology decision-support system that processes DICOM-compliant chest X-rays and low-dose CT scans and produces structured assistive outputs. These include abnormality classifications, localisation heatmaps, and draft structured report elements intended to support radiologist or clinician review. The system uses deep learning models based primarily on convolutional neural network architectures trained on annotated Indian imaging datasets from government and private sector partners.

Model objectives include multi-class abnormality detection, identifying patterns such as consolidations, cavitations, nodules, and other



thoracic findings that may be associated with TB or other lung pathology. Outputs are presented with visual overlays to indicate regions contributing to model predictions, supporting interpretability and review. The system has undergone vulnerability and penetration testing by certified agencies, and the software platform holds national regulatory approval for manufacture and sale as medical device software.

For TB-focused models, training datasets were prepared with the National Institute for Research in Tuberculosis (NIRT) under protocols aligned with national research and privacy guidance. Iterative training included radiologist feedback cycles in which model outputs, including false positives and false negatives, were reviewed and used to refine performance. This collaborative annotation and validation process aimed to improve relevance to Indian imaging characteristics and disease presentation patterns.

LIRA has undergone clinical validation on approximately 100,000 chest X-rays, demonstrating 99% sensitivity and 92% specificity in differentiating normal from abnormal findings. (Any report of clinical study published) For tuberculosis detection, the system was further validated using around 9,790 confirmed TB cases, achieving a sensitivity of 98.5% and specificity of 80%. Validation studies were conducted across multiple clinical sites, including UPUMS Hospital (Saifai, Uttar Pradesh), Ashwini Hospital (Solapur, Maharashtra), CHC Tonk (Rajasthan), Leelamani Hospital (Kanpur, Uttar Pradesh), IIT Kanpur Health Center, and Saint John's Research Institute (Bengaluru). Model outputs include heatmap overlays highlighting image regions that most influenced the algorithm's assessment, supporting result

interpretability and clinician trust. These metrics reflect study performance under defined datasets and should not be interpreted as guaranteed performance in all field settings.

LIRA is also being extended to additional use cases, including lung nodule detection and cardiovascular risk-related imaging features. These modules are under development and evaluation and are not yet established as programme-standard tools.

Deployment

LIRA's deployment model was designed to minimise additional infrastructure requirements. The platform integrates with existing hospital PACS and RIS systems through application programming interfaces or connects directly with digital X-ray consoles where PACS is not present. This plug-in approach allows AI inference to occur within established radiology workflows.

TB-specific model development relied on annotated Indian population datasets prepared with national research partners, with attention to de-identification and privacy protection. Training and validation followed structured protocols with documented review steps.

Deployment architecture supports both cloud-based and edge-based inference. In tertiary centres with stable connectivity, LIRA operates through secure cloud-hosted environments linked to PACS, returning decision-support outputs shortly after image acquisition. In bandwidth-constrained settings like Tonk, lightweight edge servers support on-site inference, with asynchronous synchronisation when connectivity is available. This hybrid model is intended to support operational continuity across variable digital environments.

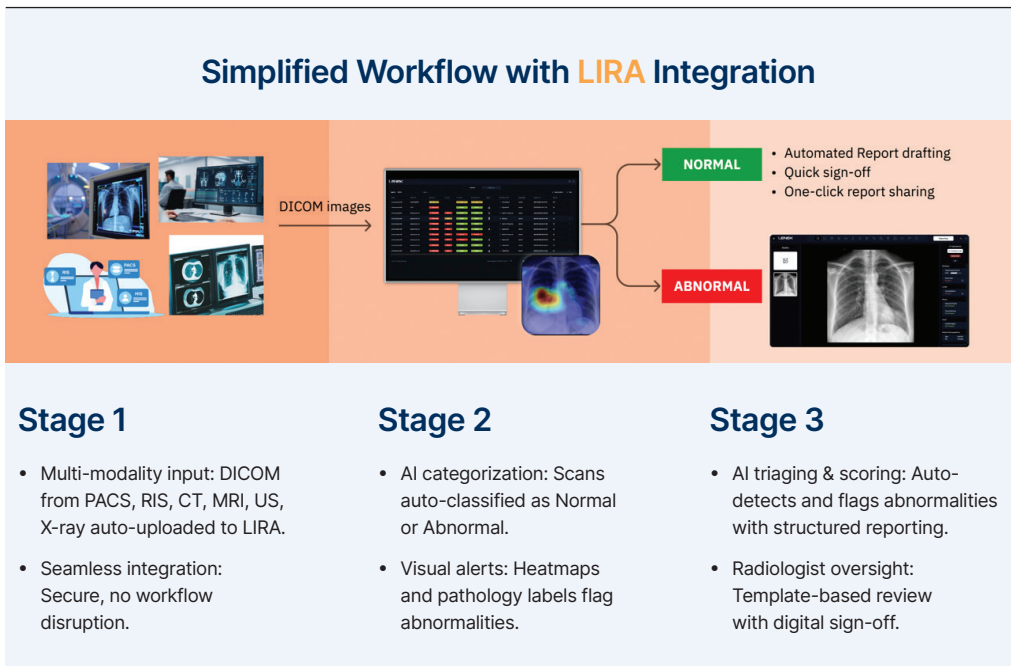
Implementation

Field implementation workflows include provider training on system use, image capture quality, and interpretation of AI outputs. Patients are registered with documented informed consent, and chest X-rays are obtained using existing or portable devices. LIRA generates a screening-support classification and highlights regions of interest. Cases flagged as TB-suggestive are referred for confirmatory testing using molecular or sputum-based diagnostics according to programme protocols. Operational

reporting is aggregated for programme monitoring.

Data security measures include de-identification for analytics, encryption, and role-based access controls aligned with national privacy requirements. Integration pathways are being developed with national digital health architecture components, including ABDM sandbox environments, FHIR-based exchange, and linkage to longitudinal digital health accounts where available. These integrations are intended to support continuity of records rather than automated decision-making.

Figure 01: Workflow



Scale of Deployment

By late 2025, LIRA had been deployed across more than ten sites in multiple Indian states, including tertiary hospitals, medical colleges, private diagnostic centres, and community health facilities. Prominent clinical sites include UPUMS Hospital (Saifai, Uttar Pradesh), Ashwini Hospital (Solapur, Maharashtra), CHC Tonk (Rajasthan), Leelamani Hospital (Kanpur, Uttar Pradesh), IIT Kanpur Health Center, and Saint John's Research Institute (Bengaluru). Validation and deployment partners span public hospitals, district health facilities, research institutes, and academic health centres, reflecting testing across diverse care delivery environments. Across early deployments, more than 50,000 chest scans were processed using AI-assisted workflows.

These deployments suggest operational scalability across both high-volume urban centres and resource-constrained facilities when infrastructure integration is feasible. Reported experience indicates that AI-assisted outputs can be generated even in sites without on-site radiologists, providing preliminary structured interpretation support pending clinical review.

Programme reports from participating sites indicate that AI-assisted triage enabled earlier identification of scans with potentially significant abnormalities for expedited review. These operational observations are context-specific and based on site-reported experience rather than controlled comparative trials.

Identified Users


LIRA supports multiple user groups across the health system. Radiologists and specialist readers in higher-level facilities use AI outputs and heatmaps to support case prioritisation and structured reporting. General physicians and clinicians in district and community facilities use simplified dashboards to identify scans likely to be normal, abnormal, or TB-suggestive to guide referral and confirmatory testing decisions.

In outreach and screening camp settings, trained technicians and frontline workers operate AI-enabled imaging workflows under protocol, capturing images and reviewing AI-assisted classifications before referring high-risk cases for further testing. Programme managers and administrators use aggregated dashboards to monitor throughput, abnormality rates, and turnaround times across facilities.

Patients generally do not interact directly with the AI system but may experience indirect effects through shorter reporting delays, more structured imaging reports, and earlier referral for confirmatory diagnostics. In digitally integrated settings, AI-assisted reports may be attached to longitudinal digital health records.

Impact

Reported impact from early deployments focuses primarily on operational and workflow indicators rather than confirmed health outcomes. Sites report reductions in reporting latency where AI-assisted



triage and structured outputs are used, although exact time savings vary and are not yet uniformly measured across settings. Developer estimates of large percentage reductions in latency should be interpreted as indicative potential rather than independently verified benchmarks.

Validation studies report high sensitivity for abnormality detection and TB-suggestive pattern detection under study conditions, with moderate specificity. In screening workflows, high sensitivity may support prioritisation of cases for confirmatory testing, while moderate specificity implies that clinical and laboratory confirmation remains essential.

Operational reports suggest that AI-assisted prioritisation may support more targeted use of downstream confirmatory tests, potentially improving efficiency of molecular testing resources. Early economic observations from selected rural programmes suggest possible per-scan cost efficiencies linked to fewer repeat visits and faster triage, though formal cost-effectiveness studies are not yet available.

Ethics and Governance

LIRA's development and deployment incorporate governance and ethical safeguards aligned with medical device and digital health expectations. AI outputs are positioned as decision-support only, with final interpretation and reporting responsibility retained by clinicians. No autonomous diagnostic decisions are made by the system.

Data governance measures include de-identification, consent-based data use, encryption, and access controls. Development collaborations with national research institutions included dataset governance and privacy protections. Alignment with national digital data protection frameworks and digital health architecture standards is part of deployment planning.

Bias and performance audits across demographic and site subgroups have been conducted by the developer, with mitigation strategies such as dataset expansion and recalibration applied where performance variation was observed. Exploratory use of privacy-preserving learning approaches, such as federated updating, has been reported in selected contexts.

Governance processes include internal review mechanisms, clinician advisory input, and regulatory compliance processes under national medical device software pathways. These measures support accountability, traceability, and post-deployment monitoring.

LIRA supports clinical decisions; final diagnosis remains with clinicians



Discussion

Successes

A key strength of LIRA's implementation approach is standards-based integration with existing PACS and RIS systems, enabling adoption without major hardware replacement. Hardware-agnostic and API-based design supports use across diverse facility types. Visual explainability through heatmaps and structured outputs supports clinician review and acceptance.

Collaborative dataset development with national research institutions improved the contextual relevance of models. Regulatory approval and national digital health alignment supported transition from pilot deployments toward repeatable programme implementations.

Operational experience suggests that AI-assisted radiology decision support can increase screening throughput and support earlier triage of abnormal scans when embedded in routine workflows, particularly in facilities without continuous radiologist presence.

Challenges

Image quality variability across facilities due to older machines, inconsistent acquisition protocols, and mixed analog-digital workflows affected specificity in some settings and required additional preprocessing and calibration. Connectivity limitations in low-bandwidth regions initially constrained cloud inference until hybrid edge approaches were introduced.

Regulatory navigation required substantial documentation and evidence generation. Adoption barriers included initial clinician skepticism and limited time for frontline staff training. Capacity-building workshops, transparent performance reporting, and feedback mechanisms were used to support acceptance.

Lessons Learned

Several lessons are broadly applicable. First, locally representative, well-annotated datasets are critical for reliable performance. Second, infrastructure-aware and hybrid deployment design supports resilience across varied digital environments. Third, governance and regulatory alignment should be built into development from early stages rather than added later.

Clear positioning of AI as clinician decision support surfacing high-risk cases, structuring reports, and visualising reasoning supports trust and appropriate use. Continuous monitoring, recalibration, and feedback loops are necessary for safe scale-up.



References

Frederick, A., Kubendiran, R., & Neelagandan, T. (2025). AI-CAD enhances pulmonary TB detection and yield in active case finding. *Journal of Global Health, 15*, 04201.

Menon, S., & Koura, K. G. (2025). Artificial intelligence for tuberculosis control: A scoping review of applications in public health. *Journal of Global Health, 15*, 04192.

Chiu, W.-C., Chang, S.-Y., Lin, C., Chen, T.-W., & Fang, W.-H. (2025). Integrating AI with PCR for tuberculosis diagnosis: Evaluating a deep learning model for chest X-rays. *Bioengineering, 12*(12), 1377.



Machine Learning- assisted Digital Twin for Diabetic Retinopathy

Lakshmi Panat, Swapna Yenishetti



Machine Learning-assisted Digital Twin for Diabetic Retinopathy

Lakshmi Panat, Swapna Yenishetti*

Abstract

Diabetic retinopathy (DR) is a chronic microvascular complication of diabetes and a leading cause of preventable vision impairment. Early detection and timely management are essential to reduce the risk of vision loss, particularly in populations with long-standing diabetes. Recent advances in digital twin (DT) technology and artificial intelligence (AI) have enabled the development of virtual representations of biological systems that can support understanding of disease progression and treatment effects.

This use case presents a prototype of a machine learning (ML)-assisted digital twin of the human eye, currently at an early development stage. The system is designed to explore how three-dimensional anatomical modelling, AI-based image segmentation and classification, and simulation capabilities could support visualisation and monitoring of retinal changes associated with DR. The system is intended as a decision-support and educational tool for clinicians and patients, complementing existing diagnostic pathways rather than replacing clinical judgment.

The chapter provides an overview of digital twin technology in healthcare, describes the design and implementation of the proposed solution, and discusses observed capabilities, limitations, and research challenges relevant to the application of AI-enabled digital twins in eye health.


Background

Diabetic retinopathy (DR) is a common and serious complication of diabetes that affects the retinal blood vessels and can lead to progressive vision impairment if not detected and managed in a timely manner. According to estimates from the International Diabetes Federation, the global population aged 20-79 years living with diabetes approximated 589 million in 2024. As the duration of diabetes increases, so does the risk

of developing DR, underscoring the importance of early detection and longitudinal monitoring.

A range of diagnostic and screening approaches have been developed to support DR identification, monitoring, and treatment. While these methods have contributed substantially to clinical practice, challenges remain related to accessibility, scalability, and timely detection,

*Purva Sharma Centre for Development of Advanced Computing (C-DAC), Pune, India



particularly in settings with limited specialist availability. These challenges highlight the need for supportive tools that can enhance clinical understanding and facilitate early intervention within existing health-care systems.

This use case introduces VisionTwin, a digital twin-based visualisation system that integrates AI and three-dimensional modelling to support understanding of diabetic retinopathy. The proposed approach explores how digital twin technology, when combined with machine learning, can contribute to improved visualisation of disease progression and support clinical decision-making. This use case aims to describe the conceptual foundations, technical design, and implementation considerations of this approach, while also highlighting challenges and future research directions.

AI Use Case Description

This section describes the proposed digital twin solution and its technical components, with a focus on its intended role as a supportive tool within eye-care workflows.

Solution Description

Digital twin technology in healthcare involves the creation of a virtual representation of a physical entity, such as a patient or organ system, that can be used to visualise, analyse, and simulate physiological processes. The construction process of a digital twin typically

includes data acquisition, preprocessing, anatomical modelling, AI-based analysis, and simulation.

In this solution, a digital twin of the human eye is developed to support visualisation and monitoring of diabetic retinopathy. The system integrates AI-based retinal image segmentation and classification with three-dimensional eye modelling to enable interactive exploration of retinal abnormalities. The solution is designed to assist clinicians and patients in understanding disease features and potential progression, rather than to provide autonomous diagnosis.

The development process begins with detailed three-dimensional modelling of ocular anatomy, including structures such as the cornea, lens, and retina. Retinal fundus images are mapped onto a three-dimensional retinal surface using UV mapping techniques to generate a detailed and interactive visual representation. AI-based segmentation models (UNet with a ResNet-34 backbone) are applied to identify retinal features associated with DR, such as haemorrhages, exudates, and microaneurysms. In parallel, classification models (EfficientNet-B4) are used to estimate DR severity across graded categories.

The outputs of these models are visualised on the digital twin, allowing users to explore segmented regions and severity indicators within the three-dimensional eye model. These visualisations are intended to support clinical interpretation, patient education, and research exploration, while final diagnostic and treatment decisions remain under clinical oversight.

Figure 01: Construction Process of DT in Healthcare

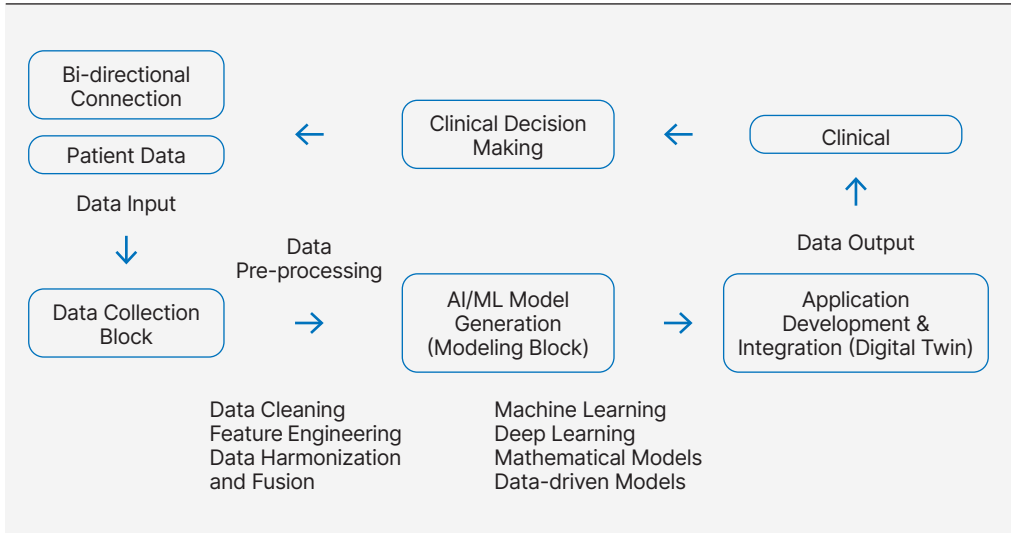


Figure 02: 3D eye modeling



Figure 03: 3D retinal sphere with UV mapping overlaying retinal/fundus images for detailed visualization of the human retina

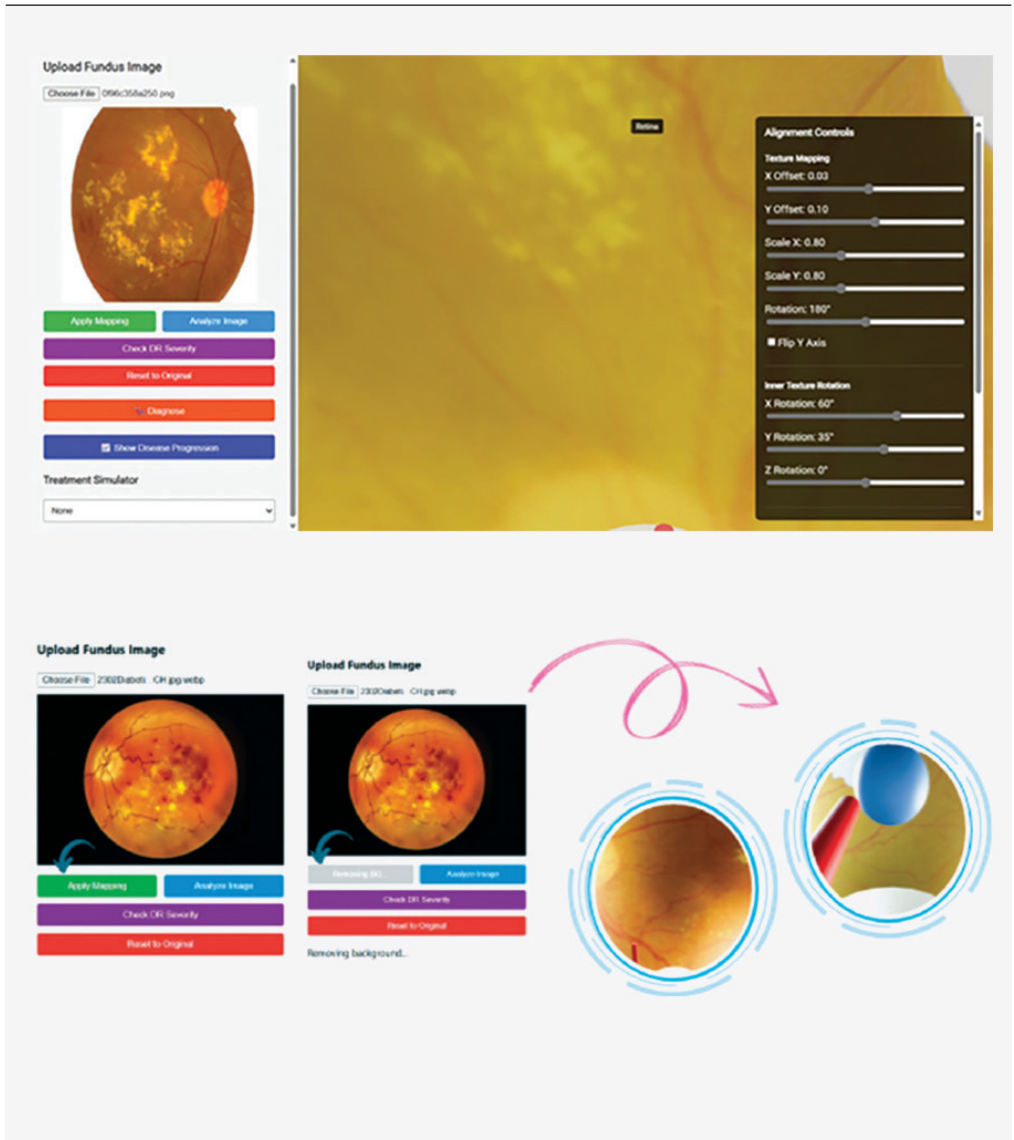


Figure 04(a): Prediction results: Detection of retinal abnormalities using UNet with ResNet34 backbone

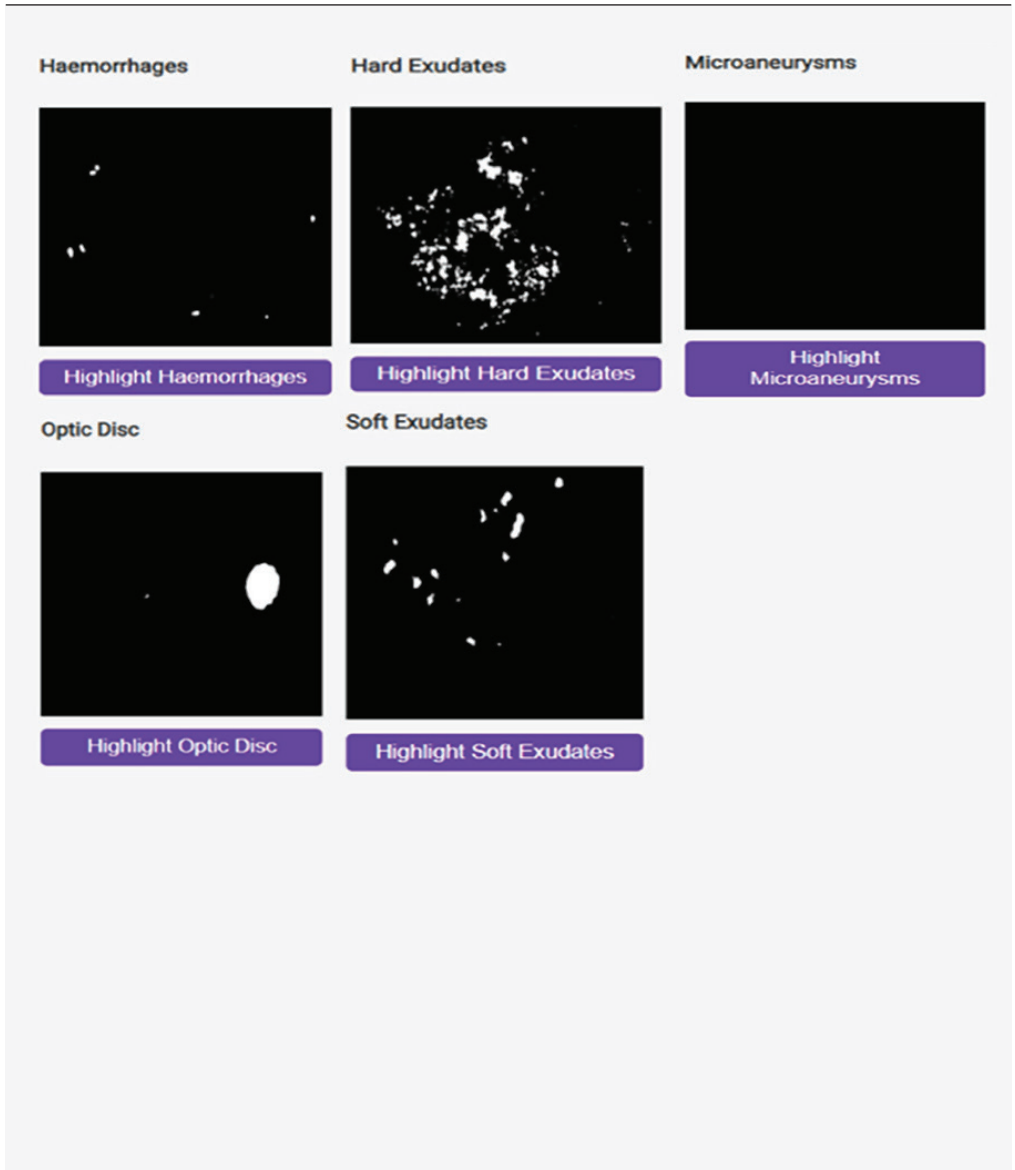



Figure 04(b): Severity Prediction of DR from grade 0 to 4 using EfficientNet-B4



Development and Implementation


The solution has been implemented as a prototype integrating AI models with interactive three-dimensional visualisation. Publicly available, annotated datasets including IDRiD and APTOS 2019 were used for model training and evaluation. AI models are deployed through a backend service using FastAPI to support near-real-time inference.

The frontend interface is developed using web-based technologies, including React, Next.js, and Three.js, to enable responsive and interactive three-dimensional visualisation. Cloud-based deployment has been explored to

support accessibility across different health-care settings, including hospitals and smaller clinics. The current implementation demonstrates technical feasibility, with further validation required prior to routine clinical use.

Scale of Deployment

At present, the AI-assisted digital twin system remains at a prototype stage. While it has been designed with adaptability to different clinical contexts in mind, it has not yet been deployed at scale within routine health-care services. Future development plans include enhanced simulation of disease progression, incorporation of treatment visualisation modules (such as



laser therapy and anti-VEGF interventions), and exploration of integration with electronic health records.

Additional features under consideration include immersive visualisation through augmented and virtual reality technologies and the use of higher-resolution datasets to improve model robustness. These developments are exploratory and would require further evaluation, governance arrangements, and resource assessment prior to broader adoption.

Identified Users

The primary intended users of the proposed system are ophthalmologists and retina specialists, who may use the platform as a supplementary tool for visualisation, education, and research related to diabetic retinopathy. Patients may also benefit from improved understanding of their condition through interactive visual representations. Researchers may use the system to explore disease characteristics and evaluate modelling approaches.

Impact

Observed outcomes from the prototype indicate that the system can generate detailed visualisations of retinal features and provide AI-assisted estimates of DR severity. These outputs may support diagnostic interpretation and patient engagement when used alongside standard clinical assessment.

Performance metrics derived from experimental evaluation (including segmentation and classification accuracy) are presented to demonstrate technical feasibility rather

than validated clinical effectiveness. The system's potential contribution lies primarily in enhancing understanding, supporting communication, and facilitating research exploration. Formal clinical evaluation would be required to assess its impact on diagnostic accuracy, care pathways, or health outcomes.

Ethics & Governance

The development and prospective deployment of AI-assisted digital twin systems raise important ethical and governance considerations. In the current prototype, publicly available datasets are used to minimise privacy risks. Any future deployment involving patient data would require compliance with applicable data protection regulations and institutional governance frameworks.

The system is explicitly designed to support, rather than replace, clinical expertise. Human oversight is essential, with clinicians retaining responsibility for interpretation and decision-making. Ongoing monitoring, model updating, and bias assessment would be necessary to ensure reliability and fairness, particularly if applied across diverse populations. These considerations align with principles articulated by the World Health Organization for responsible use of AI in health.

Human judgement remains central in every AI-assisted health decision

Discussion

This section reflects on observed capabilities, challenges, and insights gained from the development of the AI-assisted digital twin prototype.


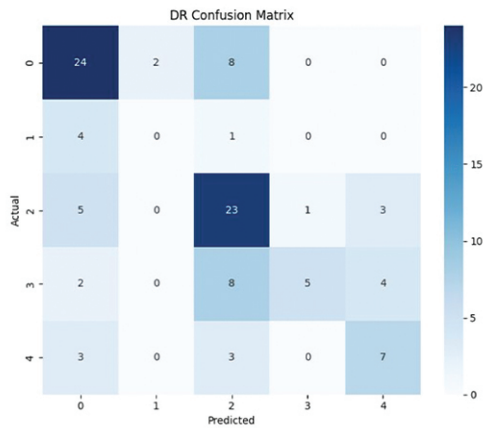
Successes


Prototype evaluation supports the feasibility of integrating AI-based retinal image analysis with three-dimensional digital twin visualisation. Experimental results indicate that the system can identify retinal features and classify disease severity under controlled conditions as depicted in Figure 5.

Figure 05: Evaluation results

=== Diabetic Retinopathy (DR) Metrics ===

	precision	recall	f1-score	support
0	0.6316	0.7059	0.6667	34
1	0.0000	0.0000	0.0000	5
2	0.5349	0.7188	0.6133	32
3	0.8333	0.2632	0.4000	19
4	0.5000	0.5385	0.5185	13
accuracy			0.5728	103
macro avg	0.5000	0.4453	0.4397	103
weighted avg	0.5915	0.5728	0.5498	103





Challenges

Several challenges were identified during development. The availability and quality of annotated datasets remain a key constraint, with variability in image quality and population representation affecting model generalisability. AI models may exhibit bias or reduced performance when applied beyond their training data.

Additional challenges include the computational complexity associated with detailed three-dimensional modelling and real-time visualisation, which may limit feasibility in low-resource environments. These challenges highlight the importance of careful system design, validation, and context-specific adaptation.

Lessons Learned

Key lessons include the importance of early and sustained engagement with ophthalmologists to ensure clinical relevance and trust. Access to diverse, high-quality datasets is critical to improving model robustness. Balancing user experience with computational requirements is essential to ensure accessibility, particularly in settings with constrained infrastructure. Emerging approaches such as edge or cloud computing may help address some of these constraints.

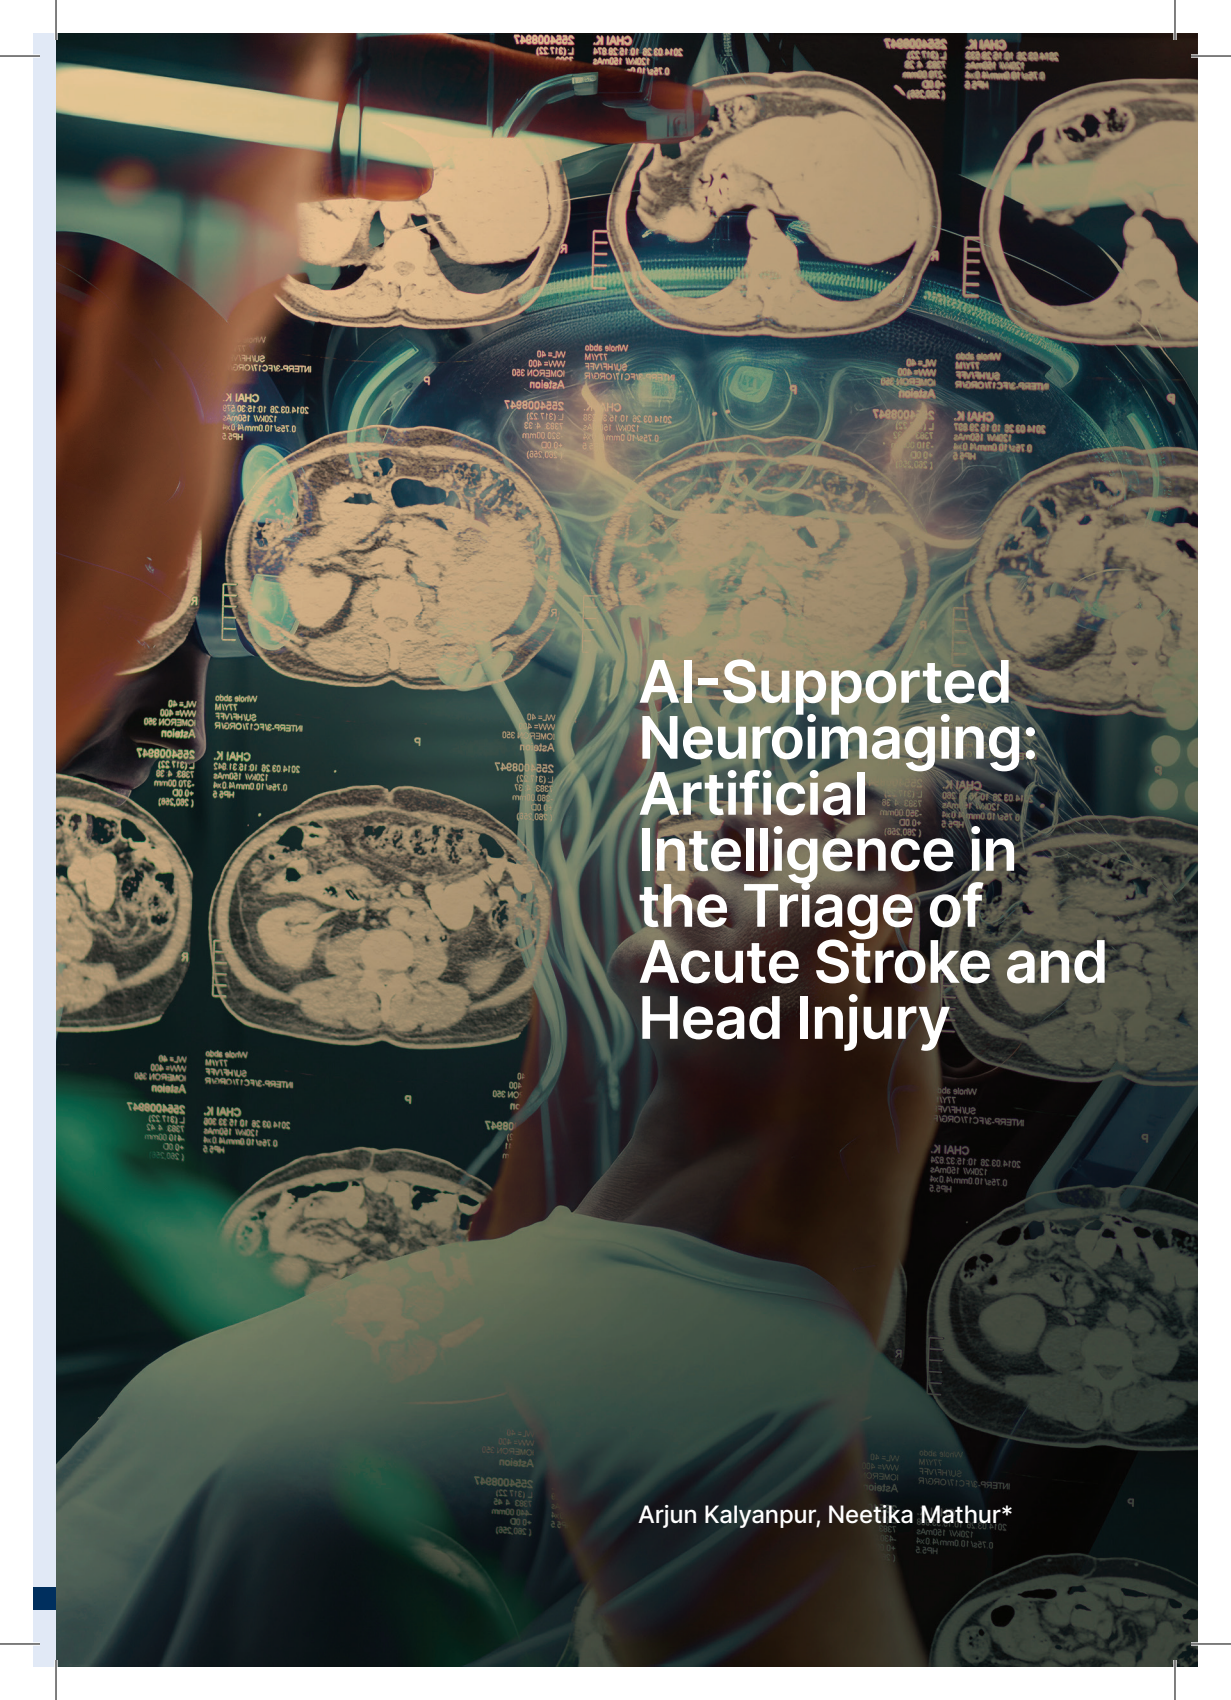
Conclusion

This use case describes an exploratory AI-assisted digital twin approach for visualising and supporting understanding of diabetic retinopathy. By integrating machine learning, three-dimensional modelling, and simulation, the proposed system illustrates how digital twins may complement existing diagnostic and educational tools in eye care.

While the prototype demonstrates technical feasibility, further evaluation, governance planning, and contextual adaptation are required before such approaches can be integrated into routine clinical practice. Continued research and clinical validation will be essential to determine the potential of AI-enabled digital twins in supporting equitable and people-centred eye-health systems.

References

- Alavee, K. A., et al. (2024). Enhancing early detection of diabetic retinopathy through the integration of deep learning models and explainable artificial intelligence. *IEEE Access*, 12, 73950–73969. <https://doi.org/10.1109/ACCESS.2024.3405570>
- Alyoubi, W. L., Shalash, W. M., & Abulkhair, M. F. (2020). Diabetic retinopathy detection through deep learning techniques: A review. *Informatics in Medicine Unlocked*, 20, 100377. <https://doi.org/10.1016/j.imu.2020.100377>
- APTOS 2019 blindness detection dataset. (2019). Kaggle. <https://www.kaggle.com/c/aptos2019-blindness-detection/data>
- Asia, A.-O., Zhu, C.-Z., Althubiti, S. A., Al-Alimi, D., Xiao, Y.-L., Ouyang, P.-B., & Al-Qaness, M. A. A. (2022). Detection of diabetic retinopathy in retinal fundus images using CNN classification models. *Electronics*, 11(17), 2740. <https://doi.org/10.3390/electronics11172740>
- Dai, L., Wu, L., Li, H., et al. (2021). A deep learning system for detecting diabetic retinopathy across the disease spectrum. *Nature Communications*, 12, 3242. <https://doi.org/10.1038/s41467-021-23458-5>
- Gargeya, R., & Theodore, L. (2017). Automated identification of diabetic retinopathy using deep learning. *Ophthalmology*, 124(7), 962–969. <https://doi.org/10.1016/j.ophtha.2017.02.008>
- Giroti, I., Das, J. K. A., Harshith, N. M., & Thahniyath, G. (2023). Diabetic retinopathy detection and classification using EfficientNet model. In *Proceedings of the 2023 International Conference on Artificial Intelligence and Applications (ICAIA)* (pp. 1–6). IEEE. <https://doi.org/10.1109/ICAIA57370.2023.10169756>
- Hill, L., & Makaroff, L. E. (2016). Early detection and timely treatment can prevent or delay diabetic retinopathy. *Diabetes Research and Clinical Practice*, 120, 241–243. <https://doi.org/10.1016/j.diabres.2016.09.004>
- International Diabetes Federation. (2025). *IDF diabetes atlas* (11th ed.). https://diabetesatlas.org/media/uploads/sites/3/2025/04/IDF_Atlas_11th_Edition_2025.pdf
- International Diabetes Federation. (2025). *IDF diabetes atlas* (11th ed.). International Diabetes Federation. <https://diabetesatlas.org>
- Iliuță, M.-E., Moiescu, M.-A., Caramihai, S.-I., Cernian, A., Pop, E., Chiș, D.-I., & Mitulescu, T.-C. (2024). Digital twin models for personalised and predictive medicine in ophthalmology. *Technologies*, 12(4), 55. <https://doi.org/10.3390/technologies12040055>
- Li, T., Gao, Y., Wang, K., Guo, S., Liu, H., & Kang, H. (2019). Diagnostic assessment of deep learning algorithms for diabetic retinopathy screening. *Information Sciences*, 501, 511–522. <https://doi.org/10.1016/j.ins.2019.06.011>
- Li, Y., Song, Z., Kang, S., Jung, S., & Kang, W. (2021). Semi-supervised autoencoder graph network for diabetic retinopathy grading. *IEEE Access*, 9, 140759–140767. <https://doi.org/10.1109/ACCESS.2021.3119434>
- Lim, J. I., Regillo, C. D., Sadda, S. R., Ipp, E., Bhaskaranand, M., et al. (2023). Artificial intelligence detection of diabetic retinopathy: Subgroup comparison of the EyeArt system with ophthalmologists' dilated examinations. *Ophthalmology Science*, 3(1), 100228. <https://doi.org/10.1016/j.xops.2022.100228>
- Mishra, S., Hanchate, S., & Saquib, Z. (2020). Diabetic retinopathy detection using deep learning. In *Proceedings of the 2020 International Conference on Smart Technologies in Computing, Electrical and Electronics (ICSTCEE)* (pp. 515–520). IEEE. <https://doi.org/10.1109/ICSTCEE49637.2020.9277506>
- Porwal, P., Pachade, S., Kokare, M., Deshmukh, G., et al. (2020). IDRid: Diabetic retinopathy—Segmentation and grading challenge. *Medical Image Analysis*, 59, 101561. <https://doi.org/10.1016/j.media.2019.101561>
- Senapati, A., Tripathy, H. K., Sharma, V., & Gandomi, A. H. (2024). Artificial intelligence for diabetic retinopathy detection: A systematic review. *Informatics in Medicine Unlocked*, 45, 101445. <https://doi.org/10.1016/j.imu.2024.101445>
- Zhu, S., Xiong, C., Zhong, Q., & Yao, Y. (2024). Diabetic retinopathy classification with deep learning via fundus images: A short survey. *IEEE Access*, 12, 20540–20558. <https://doi.org/10.1109/ACCESS.2024.3361944>
- International Diabetes Federation. (2025). *IDF diabetes atlas* (11th ed.). International Diabetes Federation. <https://diabetesatlas.org>



AI-Supported: Artificial Intelligence in the Triage of Acute Stroke and Head Injury

Arjun Kalyanpur, Neetika Mathur*



AI-Supported Neuroimaging: Artificial Intelligence in the Triage of Acute Stroke and Head Injury

Arjun Kalyanpur, Neetika Mathur*

Abstract

Stroke and head injury remain major causes of mortality and disability in India and globally. In acute stroke care, timely diagnosis is closely linked to clinical decision pathways, while early identification of intracranial bleeding in head injury informs urgent management decisions. This use case describes the operational use of an artificial intelligence (AI) enabled tool, Neural Assist, integrated within a teleradiology workflow to support radiologists in the triage and interpretation of non-contrast head CT scans for suspected stroke and head injury.

Non-contrast digital head CT scans from hospitals across the country are uploaded to an AI-enabled, cloud-based telereporting workflow platform, RADspa, which integrates the radiology information system (RIS) and picture archiving and communication system (PACS). The embedded AI algorithm automatically analyses images to identify features requiring urgent attention, including intracranial hemorrhage, hemorrhage type, midline shift (brain swelling), skull fracture, and vessel-density patterns associated with ischaemia. Suspected positive cases are flagged within the radiologist's worklist to support prioritisation and faster triage. In internal validation, NeuralAssist was reported to detect haemorrhage with approximately 89% sensitivity and 91% specificity, classify haemorrhage type with approximately 90% accuracy, identify midline shift with approximately 87% sensitivity and 88% specificity, dense MCA with approximately 87% sensitivity and 92% specificity. and detect skull fractures with approximately 90% sensitivity and 93% specificity. AI outputs are provided as decision support, while final interpretation and reporting remain the responsibility of qualified radiologists, and validated reports are distributed to hospitals through the same platform.

Reported performance metrics from internal validation datasets indicate sensitivity and specificity ranges across multiple detection tasks. These figures should be interpreted as technical performance indicators under defined conditions rather than definitive clinical outcome measures. This use case focuses on workflow integration, operational use, governance safeguards, and implementation lessons from AI-supported neuroimaging triage within a teleradiology model.

*dAlgnostiX / Teleradiology Solutions, Bangalore, India



Background

Stroke and head injury are leading contributors to death and long-term disability worldwide. In stroke care, the principle that delays in diagnosis may reduce the range of available interventions is widely recognised. Similarly, early detection of intracranial bleeding following trauma informs urgent referral and treatment pathways. Imaging, particularly non-contrast CT of the head, plays a central role in early assessment.

Artificial intelligence tools are increasingly being explored to support radiology workflows, particularly for triage, prioritisation, and structured detection support. This use case describes NeuralAssist, an AI-based image analysis tool integrated into a teleradiology reporting environment. In this implementation, AI is positioned as a decision-support and triage aid, not as an autonomous diagnostic system.

NeuralAssist Model Description

NeuralAssist is an AI-enabled clinical decision-support system designed to support radiologists in reviewing non-contrast CT (NCCT) head scans. The system combines multiple machine learning and deep learning components to identify imaging patterns associated with urgent neurological conditions. The model focuses on four categories of imaging indicators: Intracranial hemorrhage detection and quantification; shift detection; Dense MCA sign detection; and Skull fracture detection. The system combines neural network-based segmentation and detection models with rule-based post-processing steps intended to reduce false

positives and support output consistency. Training datasets were derived from large, curated teleradiology image archives. Use of these datasets was conducted under institutional ethics committee approval.


AI outputs are presented as overlays, measurements, and probability indicators to support, not replace radiologist judgment.

Hemorrhage Detection and Quantification

A deep learning semantic segmentation framework is used to identify hemorrhagic regions at pixel-level resolution and supports multi-class detection of intracranial hemorrhage, including intraparenchymal (IPH), intraventricular (IVH), subarachnoid (SAH), subdural (SDH), and epidural (EDH) types. The processing pipeline includes artifact exclusion, scan series selection, Hounsfield Unit windowing, image normalization, and format standardization to ensure consistent input quality. Following segmentation, the system generates derived measurements such as dimensional estimates and categorical severity indicators. Internal validation testing has reported approximately 89% sensitivity and 91% specificity, with subtype classification accuracy above 90%; however, these figures are based on defined evaluation datasets and may vary across populations, scanner types, and acquisition protocols, and should therefore be interpreted as technical validation metrics rather than direct measures of clinical effectiveness.

Midline Shift Detection

A segmentation-based model is used to estimate midline position and measure



displacement relative to a geometric reference derived from ventricular anatomy. Outputs include both visual overlays and numeric measurements in millimeters. Preprocessing includes exclusion of scans with severe artifacts or post-operative changes and standardized intensity normalization. A predefined displacement threshold is used to flag potential mass effect. Reported internal validation metrics include sensitivity of approximately 87% and specificity of approximately 88%. These results indicate supportive detection capability under evaluation conditions and require local validation before broader clinical reliance.

Dense MCA Detection

Detection of dense middle cerebral artery (MCA) sign, an early imaging marker sometimes associated with acute ischemic stroke is performed using a classical machine learning classifier based on intensity and asymmetry features. A Support Vector Machine model is used to support interpretability and stability. The system provides a binary present/absent flag with a confidence score to assist radiologist review. Reported internal validation metrics include sensitivity of approximately 87% and specificity of approximately 92%. This output is intended as a supportive indicator and not as a standalone diagnostic determination.

Skull Fracture Detection

A multi-stage detection pipeline supports skull fracture identification. Components include deep learning-based detection, multi-class differentiation between fractures and sutures, ensemble slice-level decision logic, and post-processing for false-positive reduction. Preprocessing includes exclusion of pediatric

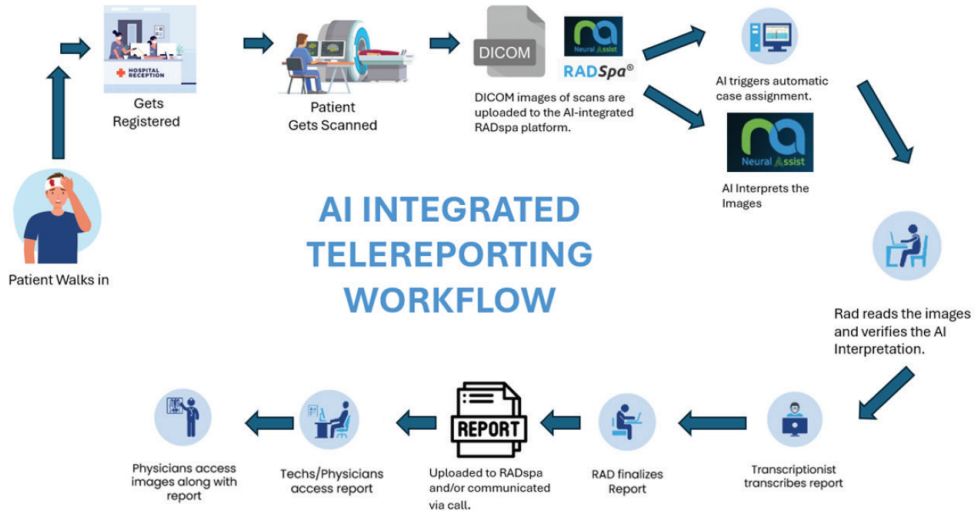
scans and cases with major artifacts or post-operative alterations. Internal validation reports sensitivity near 90% and specificity near 93%, with fracture localization agreement around 90%. These measures describe algorithm performance under test conditions and may not generalize across all settings.

AI Integrated Telerreporting Workflow

Neural Assist is integrated into a cloud-based teleradiology workflow platform (RADspa) that includes PACS and RIS functionality. DICOM images transmitted from participating hospitals are processed by the AI model prior to radiologist review.

AI outputs are used to support triage and prioritisation by flagging cases with imaging features that may indicate urgency. The system may generate structured preliminary annotations and draft text elements; however, final reports are issued only after radiologist validation. Radiologists review AI outputs alongside source images and retain full authority to accept, modify, or reject AI-generated findings. Reports are then transmitted back to referring hospitals through standard workflow channels, with critical findings communicated according to established clinical protocols. This workflow model positions AI as a prioritisation and workload-support layer, while preserving human clinical accountability.

Figure 01: AI-integrated teleradiology workflow.



Case Examples

The following cases illustrate operational examples where AI outputs were compared with radiologist interpretations. These examples demonstrate concordance in selected cases but should not be interpreted as proof of universal performance.

Case 1. Acute Multicompartmental Intracranial Hemorrhage with Significant Midline Shift

The patient had worsening confusion and a known history of subdural and subarachnoid hemorrhage. A non-contrast head CT was

performed using helical axial acquisition with sagittal and coronal reformats, and iterative reconstruction techniques were applied to reduce radiation exposure (total exam DLP: 811 mGy/cm). The radiologist identified a 5.2 × 7.1 × 6.1 cm acute intraparenchymal hemorrhage in the right frontal lobe, a 0.5 cm subdural hemorrhage along the right frontoparietal convexity, and subarachnoid hemorrhages involving the bilateral frontal and temporal lobes, with associated intraventricular hemorrhage and a 1.8 cm leftward midline shift. Overall, the AI findings were concordant with the radiologist's report (Figure 2a-e).

Figure 02: (a) AI report and images showing (b and c) intraventricular IVH and subdural hemorrhage SDH, (d)intraparenchymal and intraventricular hemorrhage (IVH) (slice 18 and 34) and (e) subarachnoid hemorrhage (SAH) (slice 32)



DRAFT REPORT BY NEURAL ASSIST

Study UID: .

Case Type	Positive
-----------	----------

Hemorrhage	Midline Shift	Calvarial Fracture	Dense MCA
Positive	Positive	Negative	Negative

Hemorrhage Findings

Type of Hemorrhage	Hemorrhage Location	Quantification
Intraventricular hemorrhage (IVH)	Right Lateral Ventricle (Slice - 34). Fourth Ventricle (Slice - 18)	Extensive


(a)




(b)



(c)



(e)



(f)

Case 2. An acute nondisplaced fracture of the left occipital calvarium

The patient had a clinical history of a small post-traumatic subarachnoid hemorrhage in the right frontal lobe. A non-contrast CT (NCCT) scan of the head was performed with axial acquisition and sagittal and coronal reformats. The radiologist confirmed the presence of the small subarachnoid hemorrhage (not shown). The AI algorithm additionally flagged notice of an acute, nondisplaced fracture of the left occipital calvarium, which was subsequently confirmed on review (Figure 3a,b).

AI tool is integrated into a cloud teleradiology workflow to triage urgent head CT scans

Embedded directly in PACS/RIS enabled RADSpa platform for seamless workflow use

Figure 03: (a) AI report and (b) image showing left occipital bone fracture



DRAFT REPORT BY NEURAL ASSIST

Study UID:


Case Type	Positive
-----------	----------

Hemorrhage	Midline Shift	Calvarial Fracture	Dense MCA
Suspicious	Negative	Positive	Positive

Calvarial Fracture Findings

Location	Slice No.
Left Occipital Bone Fracture	8

(a)



(b)

Provides visual overlays measurements and probability scores for radiologist review

Case 3. Stable multicompartmental acute Subdural Hemorrhage (SDH) with Right Parietal skull Bone Fracture

The patient had a history of head trauma and subdural hematoma. A non-contrast head CT (NCCT) was performed with axial acquisition and sagittal and coronal reformats. The radiologist

identified an acute subdural hematoma along the right cerebral convexity and the anterior falx, along with a stable, acute nondisplaced fracture of the right frontoparietotemporal calvarium extending into the right mastoid region. The AI algorithm flagged the case as positive, with findings concordant with the radiologist’s report (Figure 4a–f).

Figure 04: (a) AI report and (b, c) image showing subdural haematoma (SDH) along the right cerebral convexity and (d, e) acute nondisplaced fracture of the right parietal bone. Image (d) also shows a fracture mimic in the form of a skull suture (site of fusion).



DRAFT REPORT BY NEURAL ASSIST

Study UID:

Case Type	Positive
-----------	----------

Hemorrhage	Midline Shift	Calvarial Fracture	Dense MCA
Positive	Negative	Positive	Negative


Hemorrhage Findings

Type of Hemorrhage	Hemorrhage Location	Quantification
Epidural Hematoma (EDH)/Subdural Hematoma (SDH)	Right Convexity (Slice - 38)	Thickness: 2 mm
Subdural Hematoma (SDH)	Right Convexity (Slice - 42)	Thickness: 5 mm
Subarachnoid Hemorrhage (SAH)	Right Convexity Sulci (Slice - 20)	Mild


Calvarial Fracture Findings

Location	Slice No.
Right Parietal Bone Fracture	39
Right Temporal Bone Fracture	25


(a)




(b)



(c)



(d)



(e)



Ethics and Governance

Integration of Neural Assist into the reporting workflow was reviewed by an institutional ethics committee, with governance safeguards in place including mandatory human radiologist oversight for all final reports, use of AI strictly for triage and decision support rather than autonomous diagnosis or treatment decisions, maintenance of audit trails for AI outputs and report modifications, and implementation of data security and access controls within the PACS/RIS infrastructure. The cases used in the study were completely de-identified.

Challenges and Future Directions

Key implementation challenges include generalizability across imaging protocols and populations, interoperability with heterogeneous hospital IT systems, and the need for diverse multicenter training datasets. Broader validation across demographic and device variation remains important. Explainable AI approaches, clearer visualization tools, and user-centered interface design may improve interpretability and adoption. Integration with longitudinal imaging and clinical data systems may further support workflow value, subject to governance and privacy safeguards.



Conclusion

AI-supported neuroimaging triage tools can complement radiology workflows by highlighting potentially urgent findings and supporting prioritization within high-volume environments. NeuralAssist illustrates how combined deep learning and rule-based approaches can be embedded into teleradiology systems to support but not replace radiologist decision-making.

Observed technical performance metrics indicate supportive detection capability under defined validation conditions. Real-world clinical impact depends on workflow integration, human oversight, governance safeguards, and local validation. This model may inform future development of AI-supported emergency neuroimaging networks, particularly in settings seeking to extend specialist access through teleradiology infrastructure.


References

- Akbarzadeh, M. A., Sanaie, S., Kuchaki Rafsanjani, M., & Hosseini, M.-S. (2021). Role of imaging in early diagnosis of acute ischemic stroke: A literature review. *The Egyptian Journal of Neurology, Psychiatry and Neurosurgery*, 57(1), 175. <https://doi.org/10.1186/s41983-021-00432-y>
- Kalyanpur, A., & Mathur, N. (2023). A Teleradiology System for Early Ischemic and Hemorrhagic Stroke Evaluation and Management. *Journal of Clinical Interventional Radiology ISVIR*, 07(03), 183–189. <https://doi.org/10.1055/s-0043-1771379>
- Al-Janabi, O. M., El Refaei, A., Elgazzar, T., Mahmood, Y. M., Bakir, D., Gajjar, A., Alateya, A., Jha, S. K., Ghozy, S., Kallmes, D. F., & Brinjikij, W. (2024). Current Stroke Solutions Using Artificial Intelligence: A Review of the Literature. *Brain Sciences*, 14(12), 1182. <https://doi.org/10.3390/brainsci14121182>
- Liu, Y., Wen, Z., Wang, Y., Zhong, Y., Wang, J., Hu, Y., Zhou, P., & Guo, S. (2024). Artificial intelligence in ischemic stroke images: Current applications and future directions. *Frontiers in Neurology*, 15, 1418060. <https://doi.org/10.3389/fneur.2024.1418060>
- Kalyanpur, A., & Mathur, N. (2025). Applications of artificial intelligence in acute stroke imaging. *Artificial Intelligence in Health*, 2(4), 1. <https://doi.org/10.36922/AIH025140025>
- Gilotra, K., Swarna, S., Mani, R., Basem, J., & Dashti, R. (2023). Role of artificial intelligence and machine learning in the diagnosis of cerebrovascular disease. *Frontiers in Human Neuroscience*, 17, 1254417. <https://doi.org/10.3389/fnhum.2023.1254417>
- Seyam, M., Weikert, T., Sauter, A., Brehm, A., Psychogios, M.-N., & Blackham, K. A. (2022). Utilization of Artificial Intelligence–based Intracranial Hemorrhage Detection on Emergent Noncontrast CT Images in Clinical Workflow. *Radiology: Artificial Intelligence*, 4(2), e210168. <https://doi.org/10.1148/ryai.210168>
- Banayan, N. E., Shalu, H., Hatzoglou, V., Swinburne, N., Holodny, A., Zhang, Z., & Stember, J. (2025). Automated Midline Shift Detection in Head CT Using Localization and Symmetry Techniques Based on User-Selected Slice. *American Journal of Neuroradiology*, 46(9), 1859–1867. <https://doi.org/10.3174/ajnr.A8767>
- Lu, C.-Y., Wang, Y.-H., Chen, H.-L., Goh, Y.-X., Chiu, I.-M., Hou, Y.-Y., Kuo, K.-H., & Lin, W.-C. (2024). Artificial Intelligence Application in Skull Bone Fracture with Segmentation Approach. *Journal of Imaging Informatics in Medicine*, 38(1), 31–46. <https://doi.org/10.1007/s10278-024-01156-0>
- Kalyanpur, A., & Mathur, N. (2024). Integration of Teleradiology and Artificial Intelligence: Opportunities and Challenges. *Medical Research Archives*, 12(10). <https://doi.org/10.18103/mra.v12i10.5904>



Usability and Efficacy of Swaasa An AI-Platform for Rapid Respiratory Health Assessment using Cough as a marker

Narayana Rao Sripada, Dr Sujith Thomas Chandy,
Dr Harshal Ramesh Salve



Usability and Efficacy of Swaasa An AI-Platform for Rapid Respiratory Health Assessment using Cough as a marker

Narayana Rao Sripada¹, Dr Sujith Thomas Chandu², Dr Harshal Ramesh Salve³

Abstract

Respiratory diseases such as chronic obstructive pulmonary disease (COPD), asthma, tuberculosis, acute lower respiratory infections, and lung cancer remain major contributors to preventable morbidity and mortality. Early screening and monitoring are often limited in primary care settings where pulmonary diagnostics such as spirometry and imaging are not consistently available.

Swaasa® is an AI-enabled respiratory screening platform that analyses short, solicited cough recordings to support rapid risk assessment for selected respiratory conditions. Using machine learning and deep learning models applied to cough acoustics, combined with structured symptom inputs, the system produces a Lung Health Index indicating screening-level risk. It is designed as a screening and clinical decision-support tool that complements, but does not replace, established diagnostic pathways.


Clinical and operational studies conducted across tertiary, primary, and community settings report performance characteristics consistent with screening applications and indicate feasibility in low-resource environments. Programme deployments through public and non-governmental partners indicate operational scalability. This chapter presents Swaasa® as a real-world implementation case of AI-supported respiratory screening, highlighting feasibility, governance safeguards, and implementation lessons.

Background

Respiratory diseases account for a substantial share of the global disease burden and frequently require early detection and follow-up to reduce avoidable complications. However, access to diagnostic tools remains uneven, particularly in primary care and community settings where spirometry, imaging, and specialist consultation may be limited.

Conventional respiratory assessment relies on clinical examination, pulmonary function testing, imaging, and laboratory investigations. While effective, these approaches are resource-intensive and not always accessible at first contact points. Screening gaps contribute to delayed diagnosis, especially for chronic and infectious respiratory diseases.

1 Salcit Technologies Private Limited, Telangana, India, 2 Christian Medical College, Vellore, Tamil Nadu, India
3 AIIMS (All India Institute of Medical Sciences, New Delhi)



Cough is a common symptom across many respiratory conditions and reflects airway and lung mechanics. Advances in digital signal processing and AI have enabled analysis of cough sounds as a potential non-invasive screening signal. Cough-based analysis does not provide a diagnosis but may support early risk identification and referral decisions.

Swaasa® is a HIPAA-, ISO 27001-, and ISO 13485-certified Software as a Medical Device (SaMD) that uses AI-based cough sound analysis and selected respiratory symptoms for respiratory assessment. It is licensed by CDSCO as a Class B medical device (Drugs Control Administration, Telangana). The platform is built on over eight years of R&D supported by Government of India agencies (BIRAC, MeitY, DCIS), the UK FCDO, and USAID, and is protected by multiple Indian and international patents. Published findings report sensitivity and specificity ranges consistent with screening use. Independent usability evaluations have reported high user acceptance, while noting operational dependencies such as internet connectivity.

The system combines cough analysis with a structured respiratory questionnaire to generate a Lung Health Index (LHI) categorised into low, medium, or high risk. Cough signals contribute predictive features, while questionnaire responses support contextual risk stratification.

AI Use Case Description

Swaasa® uses short cough recordings captured through a smartphone and analysed by AI models trained to recognise disease-associated acoustic patterns. Signal features are extracted across time and frequency domains and processed using supervised learning models. Structured symptom inputs are incorporated to support contextual interpretation.

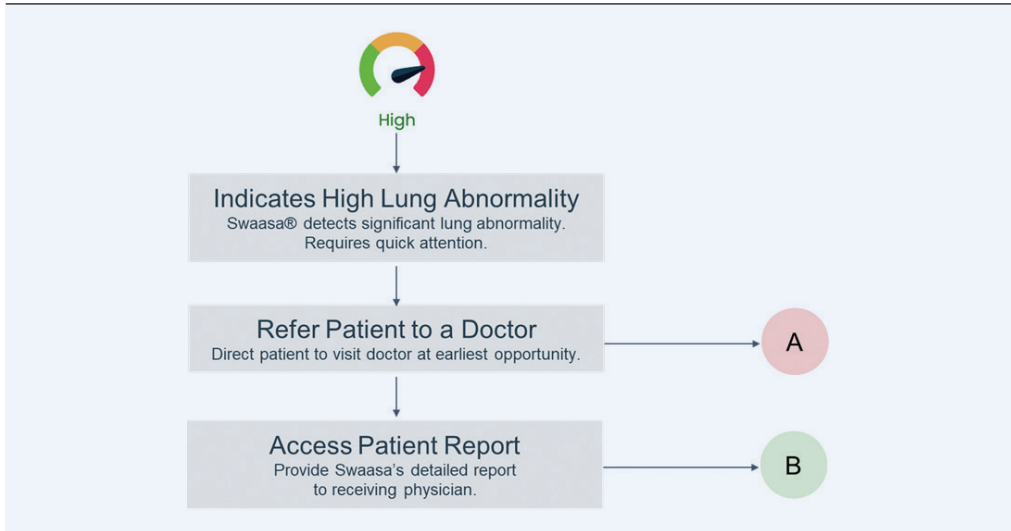
The system generates a risk score indicating the likelihood of respiratory abnormality and presents results through a clinician-facing interface in near real time. Outputs are designed to support screening, triage, and referral decisions. The platform is not intended to replace confirmatory tests such as spirometry, imaging, or laboratory diagnostics.

Low hardware requirements and short training needs support use in community and primary care programmes, particularly where conventional diagnostics are limited.

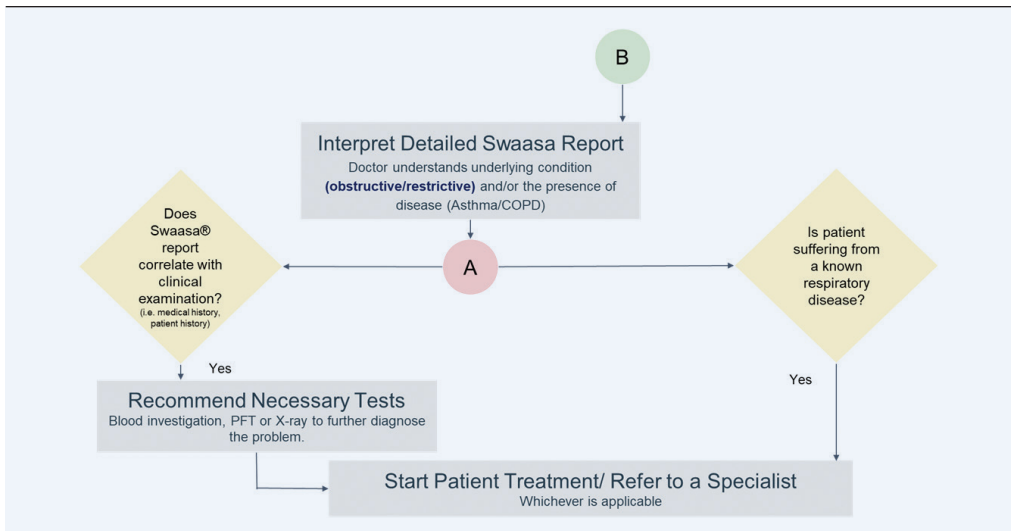
An illustrative workflow outlining the use of Swaasa is depicted below

AI-analysed cough sounds support community respiratory triage

Step-by-Step Actions Based on LHI Score (by Health Worker)



Steps to follow when a patient is referred to a Doctor.





Solution Description

Swaasa® is an AI-driven respiratory screening solution based on cough sound analysis. The technical approach combines signal processing, feature engineering, and deep learning to support rapid, non-invasive screening.

Cough audio is recorded via smartphone and pre-processed to reduce background noise and isolate valid cough events. Acoustic features are derived across temporal, spectral, and frequency dimensions associated with airway behaviour. These features are converted into spectrogram representations and processed using Advanced Audio Spectrogram Transformer-based deep learning models designed to capture time-frequency patterns.

Model training uses labelled cough datasets obtained under clinical protocols. Supervised classification models produce probabilistic outputs indicating risk of abnormal respiratory patterns. Outputs are translated into interpretable screening categories for clinical use.

The platform architecture supports scalable deployment with cloud-based processing. Performance monitoring and periodic retraining are conducted to improve robustness across populations and recording environments. Model outputs are presented as screening indicators rather than diagnoses.

Deployment & Implementation

Swaasa® is deployed as a software-only SaMD platform to enable rapid implementation without specialised hardware. A partner-led deployment model has been used across clinics, community programmes, non-governmental organisations, and digital health platforms. Short onboarding modules support training of frontline workers.

Operational workflows include cough capture using a smartphone, AI-based signal analysis, and risk-stratified outputs (low, moderate, high) that inform next steps, such as referral for spirometry, chest imaging, or specialist evaluation. This workflow supports scalable community-based screening and triage.

The approach has enabled large-scale field deployment, including nearly 450,000 assessments under CureBay's Know Your Health programme in Odisha; tuberculosis screening initiatives implemented with civil society partners such as Vasavya Mahila Mandali; and integration within women-led community health networks supported by the Healing Fields Foundation across Bihar and Uttar Pradesh. Additional deployments have been undertaken with pharmaceutical and insurance sector partners, as well as pilot implementations with state-level public authorities, including Uttarakhand and the Delhi Police.

The implementation model emphasises scalability, low operational complexity, and suitability for resource-constrained and community-based settings.

Scale of Deployment

Swaasa® has been deployed across multiple healthcare and outreach settings in eleven Indian states, including primary, secondary, tertiary, and community contexts. As a cloud-based SaMD platform, it supports high transaction volumes and has undergone controlled load testing at large daily assessment levels.

Programme reports indicate that several hundred thousand cough-based assessments have been conducted across combined deployments. These figures represent operational activity levels and should be interpreted as programme-reported reach rather than independently verified coverage metrics.

Identified Users

The table below presents the identified user groups, their primary end users, and the corresponding outcomes.

User Group / Setting	Primary End-Users	Usability Outcomes	Intended Outcomes
Community & Public Health Programs	Community health workers, NGOs, government health staff	Minimal training required; smartphone-based use; rapid test completion	May support earlier case detection and referral
Health Camps (Mass Screening)	Camp organizers, paramedical staff	High throughput; quick turnaround time; portable deployment	Intended to support identification of previously undetected respiratory conditions
Primary Care Clinics / Doctors	General physicians, clinic staff	Easy integration into routine OPD workflow	Intended to complement clinical assessment and support diagnostic prioritization
Occupational Health	Industrial health teams, on-site medical staff	Non-invasive, repeatable screening; operational convenience	May support identification of occupational respiratory risk
Corporate Wellness Programs	Employees, wellness vendors	High user acceptance; scalable via API integration	Intended to support preventive screening and risk stratification
Insurance Customer Engagement	Policyholders, insurer wellness teams	Remote usability; app-based access	May support longitudinal risk monitoring
Remote Patient Monitoring (RPM)	COPD and chronic respiratory patients; hospital RPM teams	Home-based monitoring; low patient burden	Intended to support early exacerbation detection



Impact

The observed results from the Swaasa® AI platform indicate operational value in enabling rapid, non-invasive respiratory screening in high-burden settings with limited diagnostic capacity. Deployed across primary and community health facilities under the SAMRIDH (USAID-IPE Global) initiative, the platform supported screening of 15,877 individuals and showed strong concordance with spirometry-based risk classification. Field implementation was associated with improved triage and referral processes. A threefold increase in chest X-ray referrals was observed at rural centres during the implementation period. These findings are broadly consistent with WHO digital health principles related to accessibility, efficiency, equity, and scalability in LMIC contexts.

Ethics & Governance

Ethical and governance considerations are integrated into system design and deployment. The platform follows a human-in-the-loop model in which AI outputs support, but do not replace, clinician judgement. Intended use is limited to screening and decision support.

Data protection measures include minimum necessary data collection, encryption, role-based access, and audit logging. The platform reports alignment with recognised information security and SaMD quality standards. Regulatory licensing and quality management processes support version control, traceability, and post-market performance monitoring.

Responsible AI practices include use of clinically validated training datasets, periodic model

review, and monitoring for bias and performance drift. Outputs are presented as interpretable risk categories to support transparency. Deployment guidance emphasises intended use and limitations to reduce misuse and over-reliance, inspired by WHO principles for responsible AI in health.

Discussion

Successes

Implementation experience indicates that cough-based AI screening is operationally feasible as a rapid, non-invasive approach using widely available smartphones. The platform supports usability across clinic and community settings and can be incorporated into screening and triage workflows. Reported performance metrics are consistent with screening support use.

Challenges

Challenges include variable clinician awareness, cautious adoption patterns, workflow integration needs, and internet connectivity dependence. Early deployments also identified the need for clearer communication that the tool is a screening aid rather than a diagnostic replacement.

Lessons Learned

Targeted training, clear positioning as a decision-support tool, and simple workflow integration improve adoption. Use in pre-consult screening and referral support roles appears operationally suitable. Engagement of clinical champions and programme partners supports trust and sustained use.



Conclusion

Swaasa® illustrates how AI-enabled cough sound analysis can support rapid respiratory risk screening using standard mobile devices. By combining acoustic modelling with structured symptom inputs, the platform produces interpretable risk indicators that support triage and referral decisions. Clinical and feasibility studies indicate performance and usability consistent with screening applications, and deployments show feasibility across community and primary care contexts.

References

Google Research. (2024, August 19). *This AI model is helping researchers detect disease based on coughs*. Google Blog. <https://blog.google/technology/health/ai-model-cough-disease-detection/>


Lotheti, S. K., Vissamsetti, H., Pamarthi, K., Bhimarasetty, D. M., Rudraraju, G., Sripada, N. R., Gottipulla, C., Firmal, P., Mamidgi, B., Palreddy, S. D., Bhoge, N. K. R., Narreddy, H. V. R., Yechuri, V., Jain, M., Peddireddi, V. S., Joshi, N., Vijayan, S., Turaga, S., & Avasarala, V. (2025). A cross-sectional feasibility study to evaluate the usability and efficacy of Swaasa AI platform for rapid respiratory health assessment. *Scientific Reports*, 15, Article 38379. <https://doi.org/10.1038/s41598-025-22215-8>

Pentakota, P., Rudraraju, G., Sripada, N. R., Mamidgi, B., Gottipulla, C., Jalukuru, C., Palreddy, S. D., Bhoge, N. K. R., Firmal, P., Yechuri, V., Jain, M., Peddireddi, V. S., Bhimarasetty, D. M., Sreenivas, S., Prasad, K. L., Joshi, N., Vijayan, S., Turaga, S., & Avasarala, V. (2023). Screening COVID-19 by Swaasa® AI platform using cough sounds: A cross-sectional study. *Scientific Reports*, 13, Article 18284. <https://doi.org/10.1038/s41598-023-45104-4>

SAMRIDH Healthcare Blended Finance Facility. (2024, June 11). *Revolutionizing respiratory healthcare with Swaasa® AI*. <https://samridhhealth.org/revolutionizing-respiratory-healthcare-with-swaasa-ai/>



Thangakunam, B., Chandy, S. T., Rudraraju, G., Sripada, N. R., Jayanthi, G., Gottipulla, C., Mamidgi, B., Palreddy, S. D., Bhoge, N. K. R., Narreddy, H. V. R., Samuel, P. P., Christopher, D. J., Jain, M., & Yechuri, V. (2025). Clinical validation of Swaasa® artificial intelligence platform using cough sounds for screening and diagnosis of respiratory diseases. *medRxiv*. Advance online publication. <https://doi.org/10.1101/2025.01.27.25321216>

Yellapu, G. D., Rudraraju, G., Sripada, N. R., Mamidgi, B., Jalukuru, C., Firmal, P., Yechuri, V., Varanasi, S., Peddireddi, V. S., Bhimarasetty, D. M., Kanisetty, S., Joshi, N., Mohapatra, P., & Pamarthi, K. (2023). Development and clinical validation of Swaasa AI platform for screening and prioritization of pulmonary TB. *Scientific Reports*, 13, Article 4740. <https://doi.org/10.1038/s41598-023-31772-9>



Blueprint for getting AI ready data from Radiation Oncology: Insights from a large Indian tertiary care setting

Hannah Mary Thomas T^{1,2}, Balu Krishna S¹,
Rajendra Benny Kuchipidi¹, Joy Mammen²,
Arun Zechariah³, Thomas Samuel Ram¹,
Simon Pradeep Pavamani^{1*}



Blueprint for getting AI ready data from Radiation Oncology: Insights from a large Indian tertiary care setting

Hannah Mary Thomas T^{1,2}, Balu Krishna S¹, Rajendra Benny Kuchipidi¹, Joy Mammen², Arun Zechariah³, Thomas Samuel Ram¹, Simon Pradeep Pavamani^{1*}

Abstract

Artificial intelligence is increasingly explored in medical imaging and oncology applications. However, in routine hospital environments particularly high-volume, technology-intensive departments such as radiation oncology, the primary constraint is often not algorithm availability but the presence of reliable, well-governed, and reusable clinical data. Without appropriate digital infrastructure, data stewardship processes, and workforce capacity, AI tools are difficult to develop, validate, and deploy responsibly.

This use case describes the experience of an Indian tertiary cancer centre in developing an AI-ready clinical and imaging data pipeline within a radiation oncology department. The work originated from a radiomics-based research question on recurrence prediction and evolved into a broader digital and governance transformation effort. The focus is on infrastructure, workflow redesign, and data stewardship rather than model performance.

Observed results relate primarily to infrastructure readiness, data completeness, governance maturity, and project enablement. The approach illustrates how AI readiness in clinical settings depends on interoperable systems, structured data capture, quality control, de-identification, secure storage, and human oversight. The framework is context-specific and is presented as an implementation blueprint that may inform similar efforts, with adaptation to local regulatory, governance, and infrastructure conditions.


Background

Radiation oncology departments rely on multiple digital and device-based systems for simulation, treatment planning, delivery, and follow-up. In settings with high patient volumes and constrained resources, including many centres in India, operational pressures coexist with

increasing interest in AI applications such as automated segmentation, planning support, and outcome prediction.

A prospective observational study was initiated in 2020 to explore a radiomics-based model

*¹Department of Radiation Oncology, ²Biomedical Informatics Unit, ³Education Technology Unit, Christian Medical College Vellore, Tamil Nadu, India



for predicting locoregional recurrence in head and neck cancer. The key question was asked, "Is this AI question relevant to our clinicians and patients, and if so, why should we pursue it?". An early feasibility assessment identified data fragmentation, missingness, and follow-up variability as limiting factors. This led to a reframing of priorities: rather than focusing first on model optimisation, the department prioritised building an AI-ready digital backbone and structured data pipeline embedded within routine care.

An interdisciplinary team including radiation oncologists, physicists, radiologists, data scientists, and IT specialists defined a patient-centred objective and recognised that sustainable AI efforts would depend on a redesigned data systems and governance processes rather than isolated model development.

AI ready use case

Solution

The initial modelling objective required access to planning CT images, clinical variables, treatment history, and longitudinal follow-up outcomes. Supervised models were first constructed using radiomic features describing image intensity, texture, and shape, linked with recurrence outcomes. Substantial data sparsity and missingness across key variables highlighted the need for an AI-ready data infrastructure capable of generating comprehensive, well-curated, and de-identified datasets aligned with

data-centric best practices for medical imaging and clinical data.

An end-to-end pipeline was therefore established with linked stages (1) patient identification and consent, (2) clinical data capture in the electronic medical record (EMR), (3) imaging acquisition and routing, (4) contouring and treatment planning, (5) automated and manual quality control (QC), (6) de identification and pseudonymisation, (7) secure storage and indexing, and (8) controlled access for analysis and model development (9) mitigation of longitudinal data loss during follow-up. Each stage was implemented through multidisciplinary coordination. The architecture leveraged existing enterprise clinical systems including the hospital information system (HIS), oncology information system (ARIA, Varian Medical Systems), treatment planning system (Eclipse, Varian Medical Systems), and Picture Archival and Communications Systems (PACS) platforms such as the open source Orthanc and XNAT to ensure AI readiness was embedded within routine clinical workflows (ARIA Oncology Information System | Varian, n.d.; Gutman et al., 2014; Jodogne, 2018; Marcus et al., 2007). Teleconsultation services aligned with national guidance (e-sanjeevani guidelines) were incorporated to support follow-up data capture (eSanjeevani, n.d.). The pipeline supports AI model development and evaluation but does not itself perform clinical decision-making.

Deployment and implementation

Digital transformation of radiation oncology records

Prior to 2019, the department used a hybrid

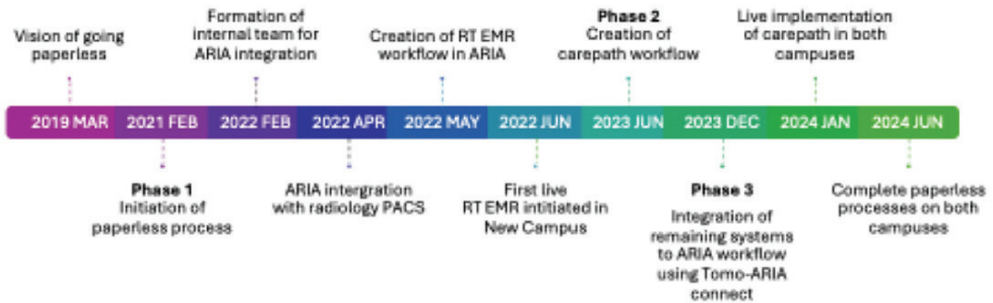
information model combining institutional digital systems with radiation-oncology-specific paper charts. While functional in a single-campus setting, this model limited the data reuse and scalability and produced non-machine-readable records.

Beginning in 2021, the department transitioned to an integrated radiation oncology EMR linked to the hospital information system. A multidisciplinary mapping exercise translated all major paper-based touchpoints into structured electronic equivalents. The ARIA Oncology Information System was selected based on

existing licences, domain functionality, and remote access capability

Implementation followed a phased approach aligned with clinical operations (Figure 1). Phase 1 converted new-campus workflows to structured electronic capture. Phase 2 extended structured capture to the original campus with a temporary dual-system period until completeness and reliability were verified. Phase 3 connected treatment platforms not natively integrated with the treatment planning system through available connectivity modules.

Figure 01: Timeline of digital transformation of radiation oncology records



This transformation supported structured and interoperable data capture. It functions as enabling infrastructure rather than an AI model or decision system.

AI-ready data pipeline in routine care

Pipeline operationalisation was aligned with routine workflows so that data curation emerged as a by-product of standard clinical processes (Figure 2).

- Consent and registration:** Consent processes were revised to include secondary AI-related data use, long-term storage, and controlled sharing, using plain-language explanations. This supports autonomy, transparency, and informed participation.
- Standardised image acquisition:** Imaging protocols were harmonised for slice thickness, reconstruction kernels and

contrast timing, and radiology where feasible across scanners. Site-specific checklists and technologist training were introduced to reduce avoidable variability. Residual heterogeneity remains and is documented.

- **Annotation and peer review:**

Disease-site specific contouring templates and labelling conventions for tumour and organs at risk were standardised. A subset of cases underwent double reading with consensus to characterise inter-observer variability.

- **Automated quality control:**

Automated scripts screened DICOM data for missing series, inconsistent parameters, and incomplete structure sets, flagging cases for human review prior to dataset inclusion.

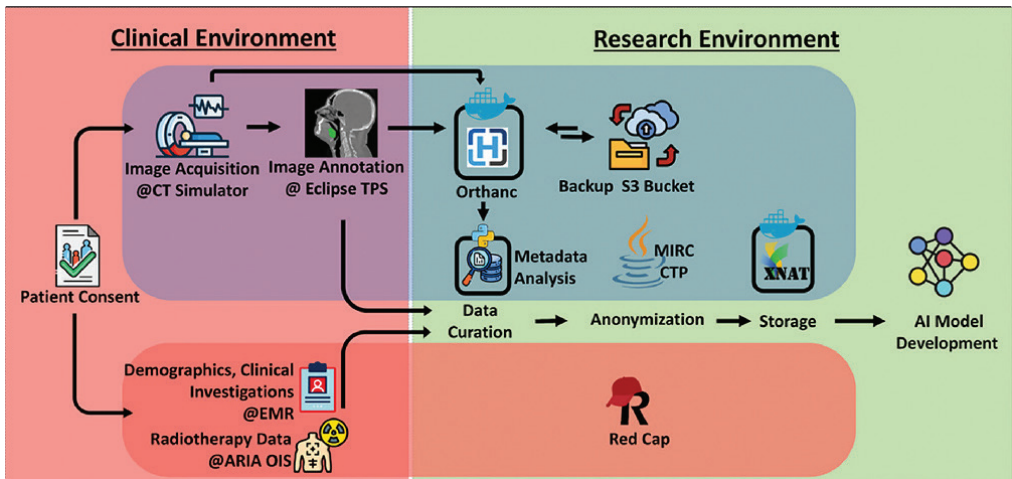
- **De-identification and storage:**

Vendor-neutral Orthanc DICOM servers applied configurable de-identification profiles before forwarding images to curated archives with role-based access. Structured clinical datasets were stored in controlled databases on Recap with audit trails.

- **Follow-up data:**

Teleconsultation appointments were scheduled at discharge to support follow-up capture. Follow-up completeness increased over time; however, this change occurred alongside multiple workflow interventions and should not be attributed to a single component.

Figure 02: AI-ready data pipeline



Scale of deployment

Since implementation of the integrated EMR and imaging pipeline, more than 2,500 radiation oncology electronic records have been captured across two campuses. The infrastructure has supported multiple imaging-based and clinical data projects across disease sites.

Orthanc and XNAT deployments have supported increasing data volumes under current institutional conditions. Extension to multi-institutional networks is technically feasible but would require additional governance agreements, interoperability alignment, and regulatory approvals. Current evidence reflects single-institution, multi-campus scale.

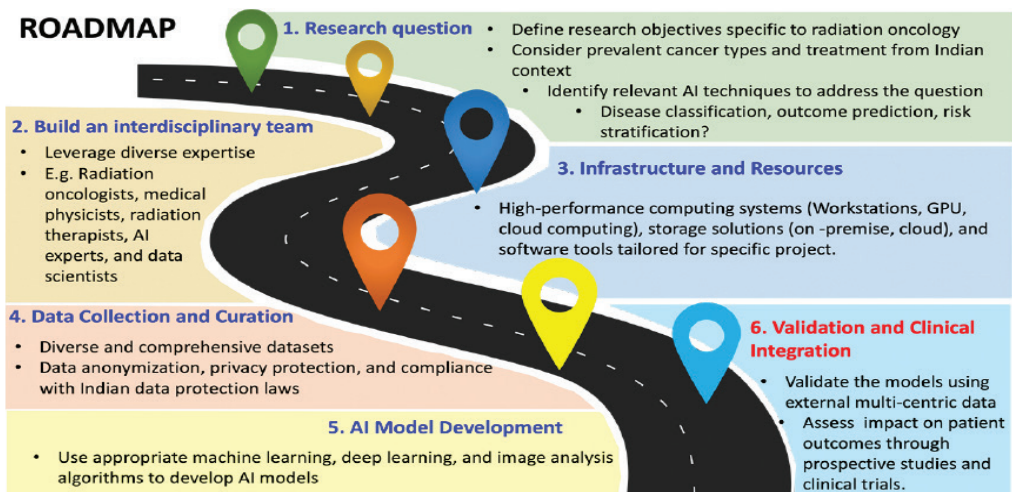
Identified users and beneficiaries

Three principal user groups are involved:

- **Data producers:**
Clinicians, physicists, and therapists generating structured clinical and imaging data within operational systems.
- **Data stewards:**
IT teams, physicists, and data managers responsible for storage, de-identification, security, and access control.
- **Data consumers:**
Approved clinicians, researchers, trainees, and data scientists using curated datasets for analysis, model development, education, and quality improvement.

Patients are indirect beneficiaries through improved data quality and governance supporting future decision-support tools.

Figure 03: Roadmap highlighting components that are required for implementation of AI-projects





Impact

The primary impacts have been infrastructural and organizational rather than algorithmic. Once established, the pipeline (Figure 3) was observed to be reusable and scalable, supporting additional disease sites and AI use cases with minimal reconfiguration, thereby potentially reducing marginal time and cost for new research efforts (Inampudi et al., 2024). Augmented follow-up data capture through teleconsultation workflows was associated with an increase in outcome data completeness from approximately 30% to 80% over four years. Governance and auditability were strengthened through standardized de-identification profiles, role-based access controls, and comprehensive logging, supporting transparency and accountability principles reflected in global AI ethics guidance. As individual models mature, their reporting can align with frameworks such as TRIPOD+AI, while benefiting from datasets generated through a pipeline already designed to meet reproducibility and data stewardship requirements

Ethics and governance

Ethics and governance were treated as core design requirements. Institutional ethics committees reviewed protocols addressing secondary use, cross-border collaboration, and long-term retention. Governance measures address human oversight, accountability, transparency, data protection, and inclusiveness. Technical safeguards include de-identification, project-level permissions, secure infrastructure, and audit trails. Oversight is provided by a cross-disciplinary steering group. Consent materials clarify that AI tools developed from these data are intended to support, not replace, clinician decision-making.

Discussion

Successes

Prioritising data pipelines before model development yielded a reusable institutional asset supporting multiple projects. Embedding structured capture within routine workflows reduced reliance on parallel research systems. Standards-based and open platforms supported interoperability and cost control. Stakeholder engagement and training supported adoption and data stewardship culture.

Challenges

Implementation required workflow negotiation to integrate structured consent, imaging checks, and documentation standards without disrupting high-volume services. Infrastructure variability and resource constraints complicated harmonisation. Infrastructure work was less visible than model development, making funding and recognition more challenging. Balancing research access with privacy and regulatory limits required ongoing governance alignment.

Lessons learned

Key lessons include treating data stewardship as a primary ethical safeguard, building modular pipelines, preferring interoperable standards, investing in local champions, and aligning digital capture with clinical workflows. AI readiness should be considered a staged infrastructure and governance journey rather than a single technical project.



Conclusion

This experience indicates that AI readiness in radiation oncology depends primarily on structured data pipelines, interoperable systems, governance safeguards, and human capacity rather than algorithm development alone. The framework described here was implemented incrementally within routine workflows and supports production of curated datasets for AI research and evaluation.

The approach is context-specific and requires adaptation to local governance, infrastructure, and regulatory environments. Sustainable AI in clinical settings is supported by deliberate investment in data infrastructure and stewardship, enabling future decision-support tools to be developed and assessed responsibly.

References

- Akinrinmade, A. O., Adebile, T. M., Ezuma-Ebong, C., Bolaji, K., Ajufo, A., & Adigun, A. O. (2023). Artificial intelligence in healthcare: Perception and reality. *Cureus*, 15. <https://doi.org/10.7759/cureus.45594>
- Cheng, A. C., Duda, S. N., Taylor, R., Delacqua, F., Lewis, A. A., & Bosler, T. (2021). REDCap on FHIR: Clinical data interoperability services. *Journal of Biomedical Informatics*, 121, 103871. <https://doi.org/10.1016/j.jbi.2021.103871>
- Collins, G. S., Moons, K. G. M., Dhiman, P., Riley, R. D., Beam, A. L., Van Calster, B., Ghassemi, M., Liu, X., Reitsma, J. B., Van Smeden, M., Boulesteix, A.-L., Camaradou, J. C., Celi, L. A., Denaxas, S., Denniston, A. K., Glocker, B., Golub, R. M., Harvey, H., Heinze, G., ... Logullo, P. (2024). TRIPOD+AI statement: Updated guidance for reporting clinical prediction models that use regression or machine learning methods. *BMJ*, 385, e078378. <https://doi.org/10.1136/bmj-2023-078378>
- Directorate-General for Parliamentary Research Services (European Parliament), Lekadir, K., Quaglio, G., Tselioudis Garmendia, A., & Gallin, C. (2022). *Artificial intelligence in healthcare: Applications, risks, and ethical and societal impacts*. Publications Office of the European Union. <https://doi.org/10.2861/568473>
- eSanjeevani. (n.d.). National telemedicine service. <https://esanjeevani.mohfw.gov.in/#/>
- Gutman, D. A., Dunn, W. D., Cobb, J., Stoner, R. M., Kalpathy-Cramer, J., & Erickson, B. (2014). Web-based tools for visualizing imaging data and development of XNATView, a zero footprint image viewer. *Frontiers in Neuroinformatics*, 8, 53. <https://doi.org/10.3389/fninf.2014.00053>
- Hannah Mary Thomas, T., Varghese, A. J., B. K. S., Mathew, M., Devakumar, D., Irodi, A., Joel, T., Peace, T., Godson, H. F., Isiah, R., & Pavamani, S. (2024). Data preparation pipeline for artificial intelligence in radiation oncology: Lessons from a large prospective imaging trial. In *Proceedings of the XXth International Conference on the Use of Computers in Radiation Therapy* (Lyon, France).
- Inampudi, S., Rajkumar, E., Gopi, A., Mol, K. S. V., & Sruthi, K. S. (2024). Barriers to implementation of digital transformation in the Indian health sector: A systematic review. *Humanities and Social Sciences Communications*, 11, 632. <https://doi.org/10.1057/s41599-024-03081-7>
- Jodogne, S. (2018). The Orthanc ecosystem for medical imaging. *Journal of Digital Imaging*, 31(3), 341–352. <https://doi.org/10.1007/s10278-018-0082-y>
- Kondylakis, H., Kalokyri, V., Sfakianakis, S., Marias, K., Tsiknakis, M., Jimenez-Pastor, A., Camacho-Ramos,

E., Blanquer, I., Segrelles, J. D., López-Huguet, S., Barelle, C., Kogut-Czarkowska, M., Tsakou, G., Siopis, N., Sakellariou, Z., Bizopoulos, P., Drossou, V., Lalas, A., Votis, K., & Lekadir, K. (2023). Data infrastructures for AI in medical imaging: A report on the experiences of five EU projects. *European Radiology Experimental*, 7(1), 20. <https://doi.org/10.1186/s41747-023-00336-x>

Kumar, A. (n.d.). *National strategy for artificial intelligence*.

Lin, D., Lapen, K., Sherer, M. V., Kantor, J., Zhang, Z., Boyce, L. M., Bosch, W., Korenstein, D., & Gillespie, E. F. (2020). A systematic review of contouring guidelines in radiation oncology: Analysis of frequency, methodology, and delivery of consensus recommendations. *International Journal of Radiation Oncology, Biology, Physics*, 107(4), 827–835. <https://doi.org/10.1016/j.ijrobp.2020.04.011>

Marcus, D. S., Olsen, T. R., Ramaratnam, M., & Buckner, R. L. (2007). The extensible neuroimaging archive toolkit: An informatics platform for managing, exploring, and sharing neuroimaging data. *Neuroinformatics*, 5(1), 11–34. <https://doi.org/10.1385/ni.5:1:11>

Moons, K. G. M., Altman, D. G., Reitsma, J. B., Ioannidis, J. P. A., Macaskill, P., Steyerberg, E. W., Vickers, A. J., Ransohoff, D. F., & Collins, G. S. (2015). Transparent reporting of a multivariable prediction model for individual prognosis or diagnosis (TRIPOD): Explanation and elaboration. *Annals of Internal Medicine*, 162(1), W1–W73. <https://doi.org/10.7326/M14-0698>

NASSCOM. (n.d.). *How AI is transforming the future of healthcare in India*. <https://nasscom.in/knowledge-center/publications/how-ai-transforming-future-healthcare-india>

Sankarapillai, J., Ramamoorthy, T., Sarveswaran, G., Venugopal, G., & Mathur, P. (2025). Epidemiological analysis of radiation therapy utilization and its implications for cancer care in India—Insights from National Cancer Registry Programme. *BMC Cancer*, 25, 1062. <https://doi.org/10.1186/s12885-025-14440-1>

van Kessel, R., Seghers, L.-E., Anderson, M., Schutte, N. M., Monti, G., Haig, M., Schmidt, J., Wharton, G., Roman-Urrestarazu, A., Larrain, B., Sapanel, Y., Stüwe, L., Bourbonneux, A., Yoon, J., Lee, M., Paccoud, I., Borga, L., Ndili, N., Sutherland, E., & Mossialos, E. (2025). A scoping review and expert consensus on digital determinants of health. *Bulletin of the World Health Organization*, 103(2). <https://doi.org/10.2471/BLT.24.292057>

Varian Medical Systems. (n.d.). *ARIA oncology information system*. <https://www.varian.com/products/software/digital-oncology/oncology-management-systems/aria-oncology-information-system>

Wiens, J., Saria, S., Sendak, M., Ghassemi, M., Liu, V. X., Doshi-Velez, F., Jung, K., Heller, K., Kale, D., Saeed, M., Ossorio, P. N., Thadane-Israni, S., & Goldenberg, A. (2019). Do no harm: A roadmap for responsible machine learning for health care. *Nature Medicine*, 25(9), 1337–1340. <https://doi.org/10.1038/s41591-019-0548-6>

World Health Organization. (n.d.). *Ethics and governance of artificial intelligence for health*. <https://www.who.int/publications/i/item/9789240029200>


REDCap Consortium. (n.d.). *REDCap software*. <https://projectredcap.org/software/>

DHR. (n.d.). *HMSC guidelines*. <https://hmsc.dhr.gov.in/hmscguidelines>

From Static Portal to Smart Assistant: Reimagining NTD Program Support with AI



Namrata Tomar, Jorge Cano, Lungi Okoko



From Static Portal to Smart Assistant: Reimagining NTD Program Support with AI

Namrata Tomar¹, Jorge Cano², Lungi Okoko³

Abstract

The ESPEN Portal, developed by the WHO Regional Office for Africa (AFRO) through the Expanded Special Project for the Elimination of Neglected Tropical Diseases (ESPEN), serves as a central data and knowledge platform supporting national NTD programmes across 48 endemic countries. The portal hosts one of the world's largest open-access repositories of subnational NTD data, geospatial maps, technical guidelines, and operational resources, forming a critical foundation for programme planning, implementation, and monitoring.

As the volume and diversity of content expanded, users increasingly faced challenges navigating the portal's extensive resources, contributing to underutilisation and a high demand for direct technical support from WHO AFRO/ESPEN staff. To address this operational bottleneck, WHO AFRO/ESPEN, Dimagi, and the Gates Foundation collaboratively developed a Generative AI-enabled assistant designed to complement the existing portal.

The assistant applies a Retrieval-Augmented Generation (RAG) approach to provide multilingual, source-cited summaries and data responses drawn exclusively from verified ESPEN datasets and WHO technical documents. Early deployment experience indicates that the tool supports faster navigation of guidance and data, enabling users to independently access information that previously required manual search or staff mediation. The assistant complements, rather than replaces, human technical assistance and illustrates how responsible, context-aware AI can strengthen public health information infrastructure and equitable access to knowledge in support of NTD elimination goals.

¹Dimagi, Dublin, Ireland ²Expanded Special Project for the Elimination of NTDs (ESPEN), WHO-AFRO, Brazzaville, Congo ³Gates Foundation, Virginia, United States of America



Background

The WHO NTD Road Map 2021–2030 emphasises a shift from fragmented, disease-specific activities toward integrated, country-led programmes that rely on timely, granular data for planning and accountability. In the African Region, which bears a substantial share of the global NTD burden, effective coordination of mass drug administration and surveillance across 48 endemic countries depends on robust data systems.

ESPEN was established by WHO AFRO to support this need, including the development of a web-based portal that functions as a central repository for subnational datasets, maps, national reports, and technical guidance. According to internal analytics, the portal contains millions of rows of programme data and thousands of geospatial and policy resources.

Despite its role as a “single source of truth,” users increasingly encountered challenges translating data availability into actionable information. Much of the content was embedded in static documents, requiring time-intensive manual navigation. Baseline user research conducted during the design phase indicated that routine programme questions could require up to 15 minutes to answer through manual search, contributing to low engagement and increasing reliance on WHO staff for direct support.

Within this context, WHO AFRO/ESPEN and partners explored whether AI-supported interfaces could unlock the value of existing public health data without introducing new data collection burdens. The resulting

initiative focused on augmenting, rather than replacing, the ESPEN Portal by transforming static content into an interactive, query-based support tool grounded in verified WHO information.

AI Use Case

Solution

The ESPEN GenAI Assistant is a conversational, agent-based system built on Open Chat Studio and designed to operate as an interactive, query-based interface to the ESPEN Portal. It employs a closed-domain Retrieval-Augmented Generation (RAG) architecture to ensure that all responses are derived exclusively from verified ESPEN datasets, WHO technical guidelines, and approved documentation.

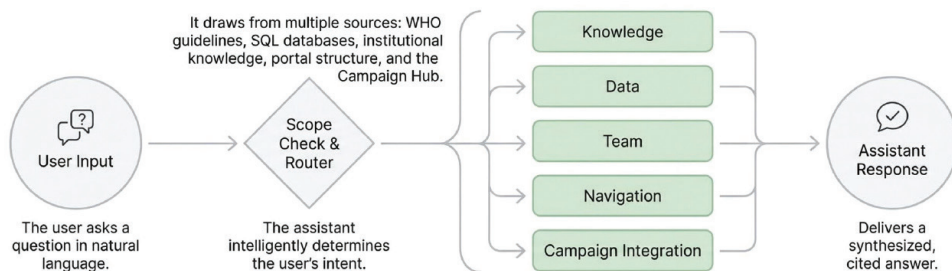
A Scope Check and Router mechanism evaluates user intent against five domains: data retrieval, summarization, campaign intelligence, navigation, and team connectivity, directing queries to specialized tools:

- **Data Requests:**

For queries requiring quantitative precision (e.g., subnational prevalence or MDA statistics), the system triggers custom SQL agents that query WHO data tables directly. This ensures answers are mathematically accurate rather than probabilistically generated. This module supports dynamic visualization, allowing users to generate graphs and charts from retrieved data.

- Knowledge Synthesis:**
 For qualitative questions (e.g., guidelines or protocols), the system utilizes a Retrieval-Augmented Generation (RAG) pipeline to search and synthesize text from trusted WHO files and technical guidelines.
- Campaign Intelligence:**
 Queries regarding operational planning are routed to the Health Campaigns Intelligence Hub, retrieving real-time data on planned public health campaigns for immunization, malaria, and NTDs and their respective population targets.
- Team Directory:**
 To facilitate direct operational support, specific queries regarding ESPEN personnel are routed to a directory module. This utilizes public information from the ESPEN website to provide users with relevant team contact details.
- The assistant utilizes a closed-domain Retrieval-Augmented Generation (RAG) architecture. This ensures the model only synthesizes answers from the verified ESPEN library, a critical design choice to mitigate the 'hallucinations' often associated with open-ended LLMs.

Figure 01: Agentic Architecture of the ESPEN GenAI Assistant





Deployment & Implementation

Development followed a user-centred, iterative design approach. National NTD programme managers participating in Kikundi, a Community of Practice convened by WHO AFRO/ESPEN, served as primary testers throughout the prototyping process. Structured feedback from these users informed refinements to language, navigation, and multilingual functionality to better reflect operational needs. The CoP conducted pre and post-testing assessments to validate the tool's impact on task efficiency before broader release. This iterative process was fueled by structured feedback from 16 representative users through dedicated assessments, complemented by qualitative insights gathered from numerous stakeholders during international conferences and partner consultations. Between the proof-of-concept (v0) and live deployment (v2), the team conducted 22+ discrete iterations driven by these feedback points.

The assistant was embedded directly into the ESPEN Portal using a lightweight widget, allowing deployment without major changes to the underlying infrastructure. Implementation proceeded in phases, with initial versions focusing on stabilising core functionality, followed by incremental expansion to include more advanced analytics and campaign intelligence features.

This approach prioritised continuity of access and minimised disruption to existing workflows, reinforcing the assistant's role as an enhancement to established systems rather than a parallel platform.

Scale of Deployment

The ESPEN GenAI Assistant is accessible through the ESPEN Portal and available to stakeholders across all 48 endemic countries. Adoption has been voluntary and organic, with usage patterns reflecting programme cycles such as reporting periods and regional meetings.


Implementation was into two distinct phases based on evolving evaluation methodology:

- **Phase I (April–July 2025):** This pilot period focused on establishing core functionalities. During this time, data collection was largely manual, focusing on total session counts and query volumes to validate the initial proof-of-concept. Initial manual tracking recorded approximately 217 sessions during this launch window.
- **Phase II (August–ongoing):** Following the stabilization of the tool, we deployed a rigorous ETL (Extract, Transform, Load) pipeline to capture high-fidelity metrics. It utilized a rigorous pipeline to filter system 'noise' and blank interactions.

According to this Phase II pipeline, the system processed 498 queries over 215 sessions through mid-December. Beyond simple volume, our analytics highlight deeper engagement: users averaged 12.9 minutes per session, indicating that they are using the assistant for substantive technical navigation rather than just quick testing.

The total operational cost for Phase II was \$11.74, averaging approximately \$0.02 per query, a supporting factor for the tool's economic sustainability

Figure 02: Phase II Operational Metrics (August – December 2025)



Identified Users

Primary users include national NTD programme managers and monitoring and evaluation officers within ministries of health. The assistant also supports data managers, implementing partners, and regional coordinators who require timely access to subnational data, campaign schedules, and technical guidance. By lowering technical barriers to information access, the tool supports more equitable participation across central and subnational levels.

Impact

The impact of the ESPN GenAI Assistant has been measured through a structured evaluation plan assessing task success, latency, and user satisfaction. Baseline testing conducted on the web portal interface during the pilot phase demonstrated efficiency gains compared to manual workflows. Specifically, the tool enabled

users to answer knowledge-related queries approximately 17 times faster and data-related queries 2 times faster than traditional navigation methods. These time savings address the critical "data friction" bottleneck, freeing program managers to focus on strategic planning rather than information retrieval.

Standardized assessments with 6 users yielded a system usability rating of 70/100 and an average rating of 4.2 out of 5 for likelihood to refer to a colleague. Qualitative feedback supports the assistant's perceived value; during the NTD Programme Managers' Meeting in Togo, a user described the tool as a "game changer for ESPN Portal data usage," highlighting its ability to generate tables and summaries that enable streamlined data reporting. Beyond individual efficiency, the tool has been utilized to create materials that effectively showcase ESPN's outcomes, supporting broader resource mobilization efforts.



Ethics & Governance

Ethical and governance considerations were integral to system design. The assistant operates within a strict scope-control framework, refusing to respond to queries outside the NTD domain. The closed-domain RAG architecture limits responses to verified ESPEN and WHO sources, and the system is programmed to acknowledge uncertainty when information cannot be retrieved.

The interface does not require user login or collection of personal data, supporting privacy and equitable access. Oversight is provided through an advisory structure involving WHO AFRO/ESPEN and partner organisations, with subject-matter experts validating all indexed content. These arrangements align with WHO principles for responsible AI in health, including transparency, accountability, and human oversight.

Discussion


Successes

A key success factor has been anchoring the assistant exclusively in trusted WHO and ESPEN data sources, which supported user trust and acceptability. Direct integration into the existing portal reduced barriers to adoption, while continuous engagement with national programme managers ensured alignment with real-world workflows.

Challenges

Despite positive implementation outcomes, deployment encountered several operational challenges. Managing the probabilistic behaviour of large language models (LLMs) required sustained technical oversight and resource investment to ensure reliability for clinical and programmatic guidance. A structured source hierarchy was therefore established to prioritise normative and officially endorsed guidance over secondary materials, ensuring that model outputs were preferentially grounded in authoritative WHO-AFRO documents. In parallel, model response parameters were calibrated to balance usability and conversational clarity with factual consistency and safety.

Risk mitigation measures included systematic source-data cleaning, development of a robust data curation pipeline to convert legacy documents into machine-readable formats, and iterative expert review (“red-teaming”) to identify and correct potential failure modes. Query-routing and scope-control mechanisms were introduced to filter



out-of-scope requests, although ongoing tuning has been required to minimise inappropriate exclusions of valid queries.

User uptake also highlighted an implementation consideration: while guided demonstrations showed strong acceptance, routine, unguided use remained more limited. This indicates that effective adoption of AI-enabled tools depends not only on technical performance but also on workflow integration, user orientation, and sustained organisational enablement.

Lessons Learned

The implementation experience highlights that technical capability alone does not ensure adoption. Guided prompts, transparent sourcing, and clear governance structures were essential to building confidence. Institutional ownership by ESPEN and the use of open-source components contributed to sustainability and local control.

Conclusion

This use case demonstrates how a Generative AI assistant can augment an existing public health data platform by transforming static repositories into interactive decision-support tools. By improving access to verified information and reducing operational bottlenecks, the ESPEN GenAI Assistant supports more efficient programme management across diverse country contexts.

This case study also highlights that responsible AI should be understood not as a static product but as a continuously managed system requiring ongoing curation, calibration, and stakeholder engagement. As the tool evolves, its grounding in open-source public goods and strong institutional ownership provides a potentially scalable approach for strengthening last-mile access to reliable information resources. With appropriate governance, quality assurance, and lifecycle oversight, such models may support more timely and equitable decision-making and contribute to progress toward 2030 disease elimination targets.

References

World Health Organization. (2021). *Ending the neglect to attain the Sustainable Development Goals: A road map for neglected tropical diseases 2021–2030*. <https://www.who.int/publications/item/9789240010352>

World Health Organization Regional Office for Africa, Expanded Special Project for Elimination of Neglected Tropical Diseases (ESPEN). (n.d.). *ESPEN | Expanded Special Project for Elimination of Neglected Tropical Diseases*. <https://espen.afro.who.int/>

World Health Organization Regional Office for Africa, Expanded Special Project for Elimination of Neglected Tropical Diseases (ESPEN). (2016). *Who we are | ESPEN*. <https://espen.afro.who.int/about-espen/who-we-are>

Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W.-T., Rocktäschel, T., Riedel, S., & Kiela, D. (2020). Retrieval-augmented generation for knowledge-intensive NLP tasks. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, & H. Lin (Eds.), *Advances in Neural Information Processing Systems* (Vol. 33). Curran Associates, Inc. <https://proceedings.neurips.cc/paper/2020/hash/6b493230205f780e1bc26945df7481e5-Abstract.html>

Kibirige, K. S., & Wandabwa, J. (2025). Enhancing access to service delivery through information transparency: A RAG-based AI-powered conversational chatbot for algorithmic transparency and regulatory compliance in digital governance. *International Journal of Science and Engineering Applications*. <https://doi.org/10.7753/ijsea1409.1008>


Lippeveld, T., Sauerborn, R., & Bodart, C. (2000). *Design and implementation of health information systems*. World Health Organization. <https://iris.who.int/handle/10665/42289>

Om, L. (2024, August 16). *Data rich and information poor*. USDM. <https://usdm.com/insights/white-papers/data-rich-and-information-poor>



Intelligent General Practitioner (iGP), an AI-powered, CHWs-driven, People-Centered Model to Support Primary Healthcare Delivery in LMICs towards UHC

Nabila Binte Ayub¹, Ananya Deb², Marzia Zaman^{1,2},
Farhana Sarkar^{1,3,4}, Khondaker A. Mamun^{1,2,5}



Intelligent General Practitioner (iGP), an AI-powered, CHWs-driven, People-Centered Model to Revolutionize Primary Healthcare Delivery in LMICs towards UHC (Bangladesh)

Nabila Binte Ayub¹, Ananya Deb², Marzia Zaman^{1,2}, Farhana Sarkar^{1,3,4}, Khondaker A. Mamun^{1,2,5}

Abstract

Achieving Universal Health Coverage (UHC) remains a central global health priority, particularly in low- and middle-income countries (LMICs), where fragmented service delivery, workforce shortages, and high out-of-pocket expenditure limit equitable access to primary healthcare. Community Health Workers (CHWs) play a critical role in extending services to underserved populations, yet their ability to deliver comprehensive, family-centred care is often constrained by limited diagnostic capacity and clinical support.

The Intelligent General Practitioner (iGP) model is a community-embedded, AI-enabled approach designed to support preventive and primary healthcare delivery through CHWs, complemented by point-of-care diagnostics, teleconsultation, and community-based clinics. AI functions as a decision-support tool to assist structured screening, triage, referral, and follow-up, while clinical accountability remains with licensed health professionals.

Implemented across rural and urban settings, the iGP model has supported the creation of digital health accounts and delivery of people-centred services at scale. Observed implementation experience indicates high service utilisation and acceptability, alongside strengthened frontline workflows. This chapter presents the iGP model as a real-world implementation case, highlighting feasibility, governance considerations, and lessons relevant to strengthening primary healthcare systems and advancing UHC in LMIC contexts.

¹CMED Health Limited, Bangladesh ²Institute of Research, Innovation, Incubation & Commercialization (IRIIC), United International University (UIU), Bangladesh ³Washington University of Science and Technology (WUST), United States ⁴Department of Computer Science and Engineering, Independent University Bangladesh (IUB), Bangladesh ⁵Department of Computer Science and Engineering, United International University (UIU), Bangladesh



Background

Universal Health Coverage aims to ensure access to essential health services without financial hardship. Despite sustained global commitment, many LMICs continue to face persistent challenges, including fragmented care pathways, shortages of trained health professionals, weak referral systems, and limited access to diagnostics at the primary level. This challenge is particularly evident in countries like: Bangladesh, India, Pakistan, Nepal, Sri Lanka, Afghanistan, where despite efforts to improve healthcare access, systemic barriers persist.

In South Asia, these challenges are particularly pronounced. Countries such as Bangladesh and India face a growing burden of non-communicable diseases (NCDs), alongside maternal, child, and adolescent health needs. Preventive services remain underutilised, and conditions such as hypertension, diabetes, and chronic respiratory disease are often detected late. High out-of-pocket expenditure continues to deter timely care-seeking, contributing to avoidable complications and financial hardship.

CHWs have demonstrated effectiveness in extending access to essential services, particularly for maternal and child health. However, their roles are often limited to vertical programmes, with restricted capacity to address broader family health needs across the life course. Limited diagnostic tools, narrow mandates, and insufficient clinical supervision constrain their contribution to comprehensive primary healthcare.

Digital health and AI-enabled tools offer potential to strengthen CHW capacity by supporting

structured decision-making, improving access to diagnostics, and facilitating referral and follow-up. Within this context, the iGP model was developed to support integrated, people-centred primary healthcare delivery, while maintaining human oversight and alignment with national health systems.

AI Use Case

Solution

The Intelligent General Practitioner (iGP) model is a CHW-driven, people-centred primary healthcare approach that integrates AI-supported decision tools, point-of-care diagnostics, teleconsultation, and community-based clinical services. The model is designed to support prevention, early detection, and continuity of care at the household and community level, while facilitating timely referral to higher levels of care when required.

Key components include:

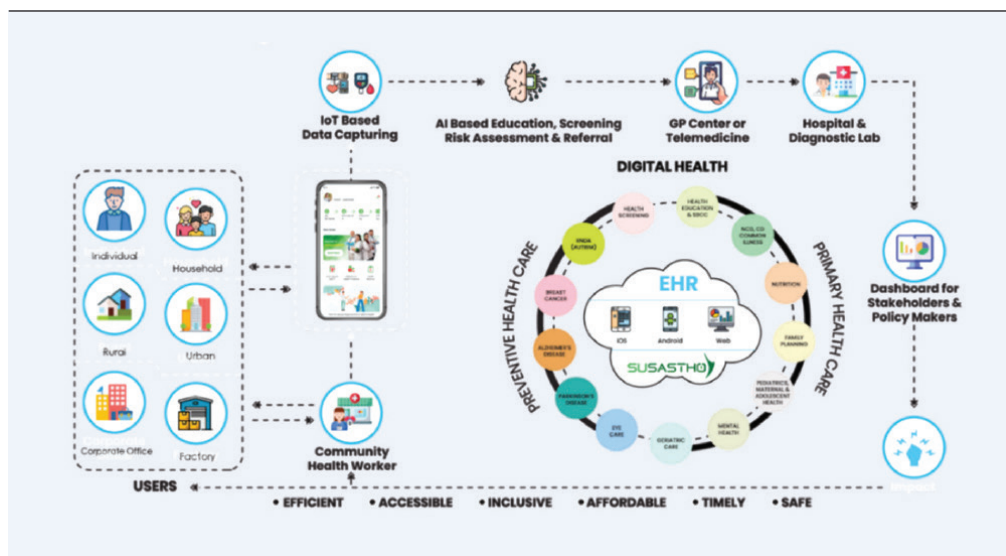
- **Unified Digital Health Accounts**-Digital health accounts are created for individuals and families to support longitudinal tracking, periodic screening, and follow-up. These accounts enable continuity of care and complement existing health system records.
- **AI-supported Decision Support and Diagnostics**- CHWs use mobile applications with AI-supported decision-support tools and IoT-enabled diagnostic devices to conduct structured screening, risk assessment, and triage. AI outputs are advisory and support, rather than replace, clinical judgement.

- **Multilingual Generative AI Co-pilot-** A multilingual generative AI assistant supports CHWs with real-time guidance, health education prompts, and structured triage support, enhancing confidence and consistency of service delivery.
- **Teleconsultation Services-** Individuals are connected to licensed clinicians via audio or video teleconsultation when indicated, enabling timely assessment and digital prescriptions where appropriate.
- **Community-based GP Centres-** Physical primary care clinics embedded within communities provide consultations, essential diagnostics, and preventive services, strengthening the continuum of care.

- **Referral, Follow-up, and Monitoring-** Structured referral pathways and digital follow-up mechanisms support continuity of care. Dashboards enable monitoring of service utilisation, follow-up completion, and user feedback to strengthen accountability.

Across all components, AI functions strictly as decision support under human oversight and within predefined clinical and governance framework.

Figure 01: Intelligent General Practitioner (iGP) Model: IoT-enabled AI-driven Comprehensive Preventive & Primary Healthcare Ecosystem



The iGP model leverages AI and IoT technologies to strengthen preventive and primary health care delivery for last-mile populations through a people-centred, hybrid (offline-online) care model. By equipping community health workers with an AI-enabled mobile application, IoT-based point-of-care diagnostic devices, and a clinical decision support system (CDSS), the platform supports early identification, screening, risk stratification, triage, referral, and follow-up within community settings, enabling the delivery of essential health services close to where people live and work.

The system integrates clinically validated screening protocols for priority conditions and applies machine-learning-based risk assessment to support early intervention and proactive care. A multilingual generative AI co-pilot, Susastho.AI, provides context-appropriate decision support, structured triage guidance, and tailored health education for frontline health workers. By reducing service fragmentation and enabling real-time data capture, the model supports population-level insights for public health planning and contributes to equitable, resilient primary health care systems in resource-constrained settings.

Figure 02: iGP platform

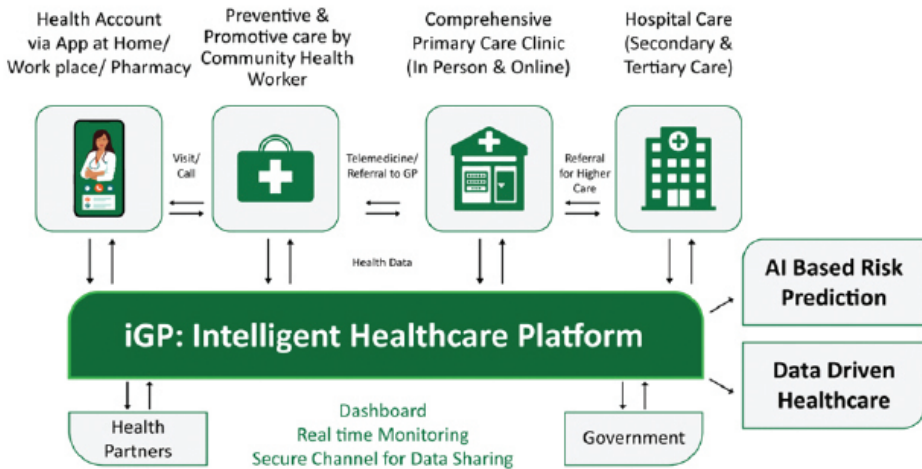
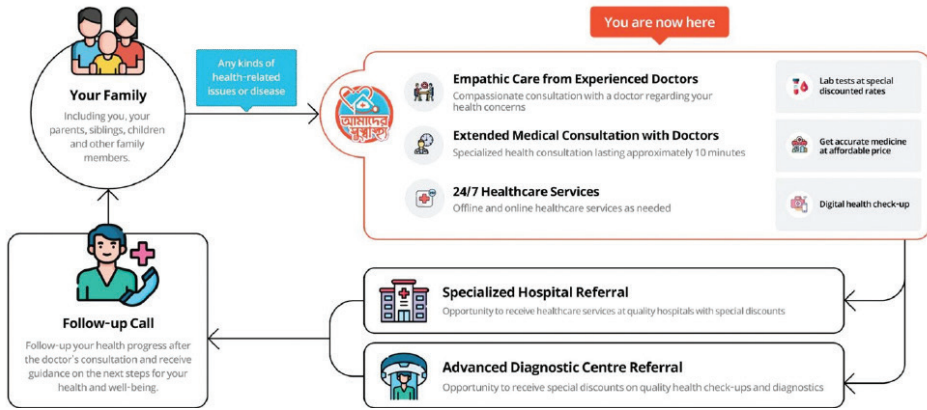


Figure 03: Preventive and Primary Healthcare Platform to Support Healthcare Service Delivery with Referral and Follow-up in the Urban Community



Deployment & Implementation


iGP model is deployed through a community-embedded approach, co-designed with CHWs and adapted iteratively based on operational feedback. Earlier implementations in Bangladesh were delivered through partnerships with non-governmental organisations, focusing on preventive care, household screening, and referral support.

The model was subsequently adapted for urban settings to address distinct challenges related to access, continuity, and fragmentation of care. This urban adaptation, implemented as Aalo Clinic, has been operationalised in six locations in Dhaka, with implementation led by UNICEF in partnership with government and civil society actors, and supported through development

cooperation funding. Building on these experiences, a self-sustaining version of the iGP model Amader Susastho is currently operational in Uttara, Dhaka, delivering an integrated, technology-enabled primary health care workflow for individuals and families. Building on this experience, the model has begun adaptation for implementation in Bangalore, India, reflecting differences in governance, service delivery, and health system organisation.

Scale of Deployment

To date, the model has been implemented across rural and urban settings in Bangladesh, with early-stage adaptation underway in Bangalore, India. In rural Bangladesh, the model has been deployed across 65 rural catchment areas, engaging over 3,000 health workers—the



majority of whom are women—and reaching approximately 2.5 million individuals from over 450,000 households. The urban implementation in Dhaka, under the Aalo Clinic initiative, has recorded over 976,000 patient interactions across 6 locations. The Amader Susastho model is currently operational in Uttara, Dhaka. Rollout in India is ongoing in an effort to scale beyond Bangladesh.

Programme monitoring data indicate that 64.5% of service users were female, with high reported satisfaction levels; however, these figures should be interpreted as indicative of operational reach rather than independently validated outcomes. Expansion across multiple geographies is supportive of the model's adaptability, notwithstanding the importance of contextual tailoring and sustained system support.

Identified Users

Primary users and beneficiaries include individuals and families receiving preventive and primary healthcare; community health workers delivering household-level services; general practitioners providing teleconsultation and clinic-based care; pharmacists supporting medication access; employers and workers participating in workplace health programmes; and health administrators and policymakers using aggregated data for planning and monitoring.

Impact

Observed implementation experience suggests improved access to preventive and primary healthcare services, particularly in underserved communities. Programme data indicate increased screening, referrals, and follow-up, alongside reported use of diagnostic and decision-support tools by CHWs.

Changes in indicators related to NCD risk detection and care-seeking behaviour are consistent with strengthened primary care processes. These findings should be interpreted as indicative of operational utility rather than definitive evidence of population-level health impact. Formal evaluations are required to assess effectiveness, cost-effectiveness, and long-term outcomes.

Ethics & Governance

Ethical and governance considerations are embedded within the iGP model. Informed consent procedures are integrated into service delivery, and AI-enabled tools operate under human oversight, with clinical accountability retained by licensed professionals. Data protection measures include role-based access controls, secure data storage, and safeguards to minimise bias and misuse. Transparency and explainability are prioritised to ensure appropriate use of AI-supported recommendations.



Discussion

Successes

The reported outcomes support the feasibility of integrating doorstep services, teleconsultation, and community-based clinics within a single continuum of care. AI-supported tools can enable CHWs to extend their scope in a structured and supervised manner, with an aim to support early detection and referral while maintaining safety.

High utilisation and reported acceptability suggest strong community engagement. The predominance of female CHWs may contribute to equity and trust at the community level. Programme monitoring data indicate changes in several operational indicators during the implementation period, including an increase in referral rates from 0.46% to 9.9%, and a reported increase in community health awareness from 11.46% to 73.36%. Female users accounted for 64.5% of service interactions. These observed changes are consistent with improved access and service delivery processes. Challenges

Key challenges include infrastructure constraints, sustaining data quality at scale, integration with fragmented health systems, and long-term financial sustainability. Continuous training, supervision, and regulatory alignment are required to support responsible scale-up.

Lessons Learned

The implementation experience suggests that AI-enabled primary healthcare may be most effective when it strengthens human relationships and frontline capacity. Simplicity, explainability, offline functionality, and strong governance are critical. Trust, community engagement, and institutional alignment underpin sustainability.

Conclusion

The iGP model illustrates how AI-supported, CHW-driven primary healthcare can be operationalised in LMIC contexts to support progress toward UHC. By integrating decision support, diagnostics, telemedicine, and structured referral pathways within community-based care, the model addresses key access and continuity gaps.


References

- A call to action and a lifecourse strategy to address the global burden of raised blood pressure on current and future generations: The Lancet Commission on hypertension. (n.d.). *The Lancet*. Retrieved January 5, 2026, from <https://www.thelancet.com/commissions-do/hypertension>
- Bangladesh infant mortality rate (1950–2025). (n.d.). *MacroTrends*. Retrieved January 5, 2026, from <https://www.macrotrends.net/global-metrics/countries/bgd/bangladesh/infant-mortality-rate>
- Bangladesh maternal mortality rate. (n.d.). *MacroTrends*. Retrieved January 5, 2026, from <https://www.macrotrends.net/global-metrics/countries/bgd/bangladesh/maternal-mortality-rate>
- Chowdhury, M. H., Ripan, R. C., Islam, A. K. M. N., Hridhee, R. A., Sarker, F., Islam, S. M. S., & Mamun, K. A. (2025). Digital health inclusion towards achieving universal health coverage for Bangladesh utilizing general practitioner model. *Health Policy and Technology*, 14(4), 100731. <https://doi.org/10.1016/j.hlpt.2023.100731>
- Ensuring universal health coverage through primary healthcare services in Bangladesh. (n.d.). *Centre for Policy Dialogue (CPD)*. Retrieved January 5, 2026, from <https://www.cpd.org.bd>
- Faruque, M., Mia, M. B., Chowdhury, M. H., Sarker, F., & Mamun, K. A. (2019). Feasibility of digital health services for educating the community people regarding lifestyle modification combating noncommunicable diseases. In *Lecture Notes in Computer Science* (Vol. 11587, pp. 333–345). Springer. https://doi.org/10.1007/978-3-030-21935-2_25
- Florence, M., Jarawan, E., Boiangiu, M., & El Yamani, F. E. K. (2025). Moving toward universal health coverage with a national health insurance program: A scoping review and narrative synthesis of experiences in eleven low- and lower-middle income countries. *PLOS Global Public Health*, 5(1), e0003651. <https://doi.org/10.1371/journal.pgph.0003651>
- Health insurance in Bangladesh: Challenges, prospects and action plan. (n.d.). *P4H Network*. Retrieved January 5, 2026, from <https://p4h.world/en/news/health-insurance-in-bangladesh-challenges-prospects-and-action-plan/>
- Health spending pushed 55 million Indians into poverty in a year. (n.d.). *Pallium India*. Retrieved January 5, 2026, from <https://palliumindia.org/2018/06/health-spending-pushed-55-million-indians-into-poverty-in-a-year>
- IDF diabetes atlas. (n.d.). *International Diabetes Federation*. Retrieved January 5, 2026, from <https://diabetesatlas.org/>
- Improving access to noncommunicable disease services for Rohingya refugees and immediate host communities in Cox's Bazar, Bangladesh. (n.d.). *World Health Organization*. Retrieved January 5, 2026, from <https://www.who.int/about/accountability/results/who-results-report-2020-mtr/country-story/2021/bangladesh>
- Joarder, T., Chaudhury, T. Z., & Mannan, I. (2019). Universal health coverage in Bangladesh: Activities, challenges, and suggestions. *Advances in Public Health*, 2019, Article 4954095. <https://doi.org/10.1155/2019/4954095>
- Noncommunicable diseases – India. (n.d.). *World Health Organization*. Retrieved January 5, 2026, from <https://www.who.int/india/health-topics/noncommunicable-diseases>
- Out-of-pocket health expenditure 74%. (n.d.). *The Business Standard*. Retrieved January 5, 2026, from <https://www.tbsnews.net/bangladesh/health/out-pocket-health-expenditure-74-89572>
- Out-of-pocket spending for healthcare push 61 lakh into poverty in 2022: BIDS. (n.d.). *The Business Standard*. Retrieved January 5, 2026, from <https://www.tbsnews.net/bangladesh/health/out-pocket-spending-healthcare-push-61-lakh-poverty-2022-bids-900326>
- Rabbani, M. G., Hossain, Z., Tisha, K. I., Nahar, K., Sakin, Z. R., Nasrin, T., Ahmed, M. W., Hasan, M. Z., Solomon, C., Jahangir, M. A., Vandenant, M., & Mahmood, S. S. (2025). Influence of a pilot urban primary healthcare model on the use of medically trained healthcare providers among the low-income slum populations in Bangladesh: Findings from an implementation research study. *International Journal for Equity in Health*, 24(1), Article 2622. <https://doi.org/10.1186/s12939-025-02622-z>
- UHC service coverage index – Bangladesh. (n.d.). *World Bank*. Retrieved January 5, 2026, from <https://datos.bancomundial.org/indicador/SH.UHC.SCI?locations=BD>
- Universal health coverage (UHC), SDG target 3.8. (n.d.). *World Health Organization*. Retrieved January 5, 2026, from <https://www.who.int/data/gho/data/themes/theme-details/GHO/universal-health-coverage>
- Zaman, M., Hridhee, R. A., Bhuiyan, R. A., Gomes, C. A., Rahman, M. M., Islam, S. M. S., Sarker, F., & Mamun, K. A. (2025). Efficacy of using a digital health intervention model using community health workers for primary health services in Bangladesh: A repeated cross-sectional observational study. *BMC Public Health*, 25(1), Article 1833. <https://doi.org/10.1186/s12889-025-22770-9>



Supporting Antibiotic Stewardship Through an AI-Driven Digital Intervention to Combat Antimicrobial Resistance

Emon Ghosh¹, Md. Jobayer Rahman¹,
Marzia Zaman^{1,2}, Marium Salwa³, Md. Atiqul Haque³,
Farhana Sarkar^{4,5}, Khondaker A. Mamun^{1,2,6}



Supporting Antibiotic Stewardship Through an AI-Driven Digital Intervention to Combat Antimicrobial Resistance

Emon Ghosh¹, Md. Jobayer Rahman¹, Marzia Zaman^{1,2}, Marium Salwa³, Md. Atiqul Haque³, Farhana Sarkar^{4,5}, Khondaker A. Mamun^{1,2,6}

Abstract

Antimicrobial resistance (AMR) represents a growing public health challenge, driven in part by irrational antibiotic use, self-medication, and constrained access to affordable, quality healthcare. In Bangladesh, rising resistance levels and widespread non-prescription antibiotic dispensing highlight the need for community-level stewardship interventions that operate within existing care-seeking behaviours.

This use case describes a pilot AI-supported digital intervention implemented across 28 community pharmacies in Dhaka, involving 2,654 participants. The intervention aimed to support rational antibiotic use by integrating pharmacists, patients, and physicians within a connected digital ecosystem. Core components included an AI-enabled symptom assessment tool for triage, an automated audit mechanism for real-time monitoring of antibiotic dispensing, and a generative AI assistant providing medication guidance via text and voice interfaces.

Implemented as a proof of concept between October 2023 and January 2024, the intervention illustrated changes in purchasing patterns consistent with improved adherence to prescription-based antibiotic use. Observed reductions in symptom-based antibiotic purchases and increases in prescription-guided transactions suggest that AI-supported stewardship tools, when embedded in pharmacy workflows and coupled with telemedicine referral pathways, can support responsible antibiotic use. The chapter presents a context-specific implementation experience, highlighting feasibility, governance considerations, and lessons relevant to antimicrobial stewardship efforts in low- and middle-income country settings.

¹ Advanced Intelligent Multidisciplinary Systems (AIMS) Lab, Institute of Research, Innovation, Incubation & Commercialization (IRIIC), United International University (UIU), Bangladesh ² CMED Health Limited, Bangladesh

³ Bangladesh Medical University (BMU), Bangladesh ⁴ Washington University of Science and Technology (WUST), USA

⁵ Center for Computational and Data Sciences, Independent University, Bangladesh (IUB), Bangladesh

⁶ Department of Computer Science and Engineering, United International University (UIU), Bangladesh



Background

Antimicrobial resistance (AMR) is widely recognized as a critical global health threat, undermining decades of progress in infectious disease control. While antimicrobials have historically contributed substantially to improved life expectancy, increasing resistance is eroding these gains. Globally, AMR was associated with an estimated 4.95 million deaths in 2019, with projections indicating a further rise if current trends persist.

In low- and middle-income countries, including Bangladesh, the AMR challenge is compounded by systemic constraints. These include limited access to affordable healthcare, inconsistent regulation of antibiotic sales, circulation of substandard medicines, and heavy reliance on community pharmacies as the first point of care. For many individuals, the cost and accessibility barriers associated with physician consultations result in widespread self-medication, often involving antibiotics obtained without prescriptions.

In Bangladesh, reported resistance rates have increased in recent years, with some commonly used antibiotics showing markedly reduced effectiveness. Antibiotic consumption has also risen substantially, reflecting broader global trends driven by LMICs. The COVID-19 pandemic further intensified antibiotic use, particularly in community settings, where symptoms resembling bacterial infections prompted empiric and often inappropriate antibiotic consumption.

A key operational gap contributing to irrational antibiotic use is the limited digitization of primary care and pharmacy workflows. Most community pharmacies lack tools to verify prescriptions,


track patient history, or audit dispensing practices. At the same time, digital health interventions including telemedicine, electronic records, and AI-supported decision tools have been identified as promising approaches to strengthen stewardship by linking pharmacies with clinicians and standardised guidance. This chapter examines one such intervention designed to support rational antibiotic use at the community level.

AI Use Case Description

Solution Description

The intervention described in this chapter is a multi-component digital ecosystem designed to support antimicrobial stewardship within community pharmacies. Implemented by CMED Health Limited in collaboration with Bangladesh Medical University, the platform integrates three interoperable components into a unified workflow.

First, an AI-driven symptom assessment module within the Health Agent App analyzes patient symptoms against standard treatment protocols to identify high-risk cases requiring physician referral. Second, a Generative AI Assistant acts as a decision support tool, providing real time medication guidance to pharmacists through the provision of dosage instructions. Third, an Intelligent Audit System functions as an automated rule-based engine that audits purchase records in real time. This system flags



irrational transactions such as antibiotics sold without a prescription or with incomplete courses, for immediate review via the admin dashboard.

To strengthen triage at the point of care, the system incorporates an AI-powered health assessment engine that distinguishes between higher-risk cases requiring physician consultation and lower-risk cases suitable for over-the-counter management. This hybrid approach combines a fine-tuned language model to structure unstructured symptom descriptions with an ensemble machine-learning model optimised for conservative, safety-oriented decision support. Importantly, AI outputs are advisory and operate within predefined rules that prioritise referral to human clinicians where risk thresholds are met.

Deployment & Implementation

The intervention was deployed using a human-in-the-loop workflow tailored to community pharmacy settings. When a patient visits a participating pharmacy, the pharmacist records the request using the Health Agent App (Figure 1) following steps (a) through (g). Requests are categorized as prescription-based, medicine-name-based, or symptom-based. In cases where antibiotics are requested without a valid prescription, the system restricts standard dispensing and initiates a telemedicine referral pathway, offering the patient access to a physician consultation for clinical assessment and digital prescription where appropriate.

Following each interaction, a digital health account is created for the patient, enabling longitudinal tracking of antibiotic use and automated reminders to support completion of prescribed courses. The system was piloted across 28 community pharmacies in two urban areas of Dhaka, Bangladesh, engaging 2,654 unique individuals over a four-month period from October 2023 to January 2024. The pilot was designed to assess feasibility, workflow integration, and behavioural signals rather than long-term clinical outcomes.

Community pharmacists (health agents) served as the primary system users, using the application to validate sales, manage inventory, and digitize their operations. While patients were the principal beneficiaries through improved access to safety checks, digital records, and telemedicine services.

Impact

Given the pilot nature of the intervention, impact is best interpreted through observed changes in dispensing patterns rather than definitive health outcomes. Analysis of transaction data suggests a shift toward more prescription-guided antibiotic use following implementation.

Safety first AI triage engines enabling rational antibiotic pathways

Figure 01: The Health Agent App interface showing the service workflow, user registration, and medicine entry screens.

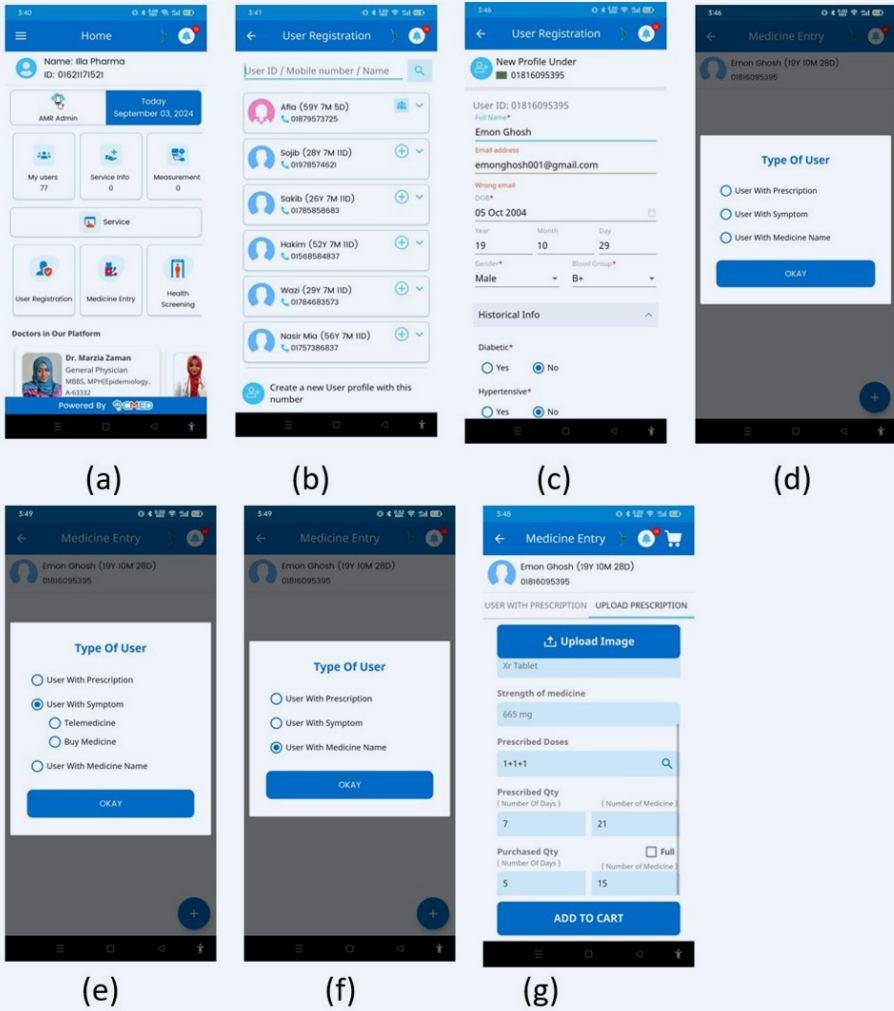
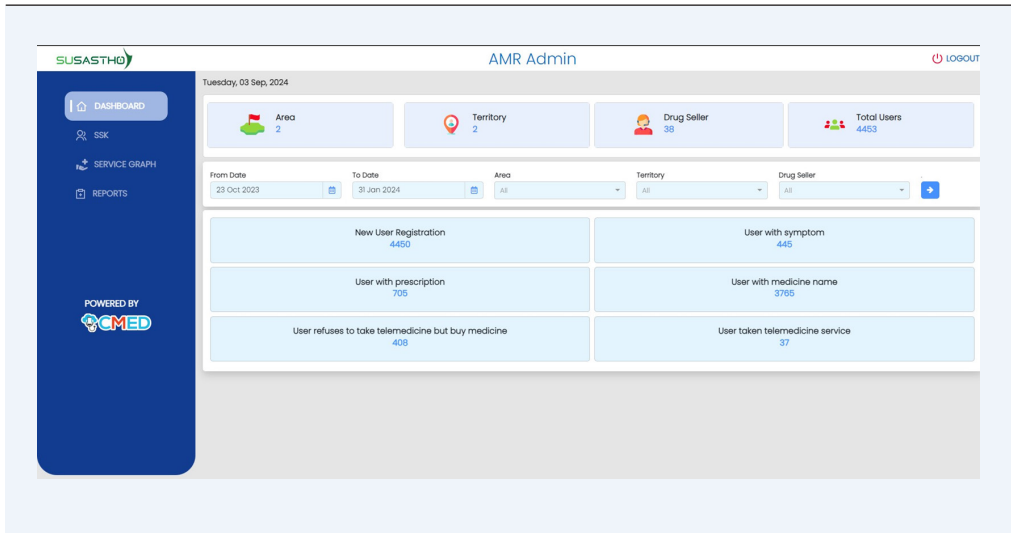


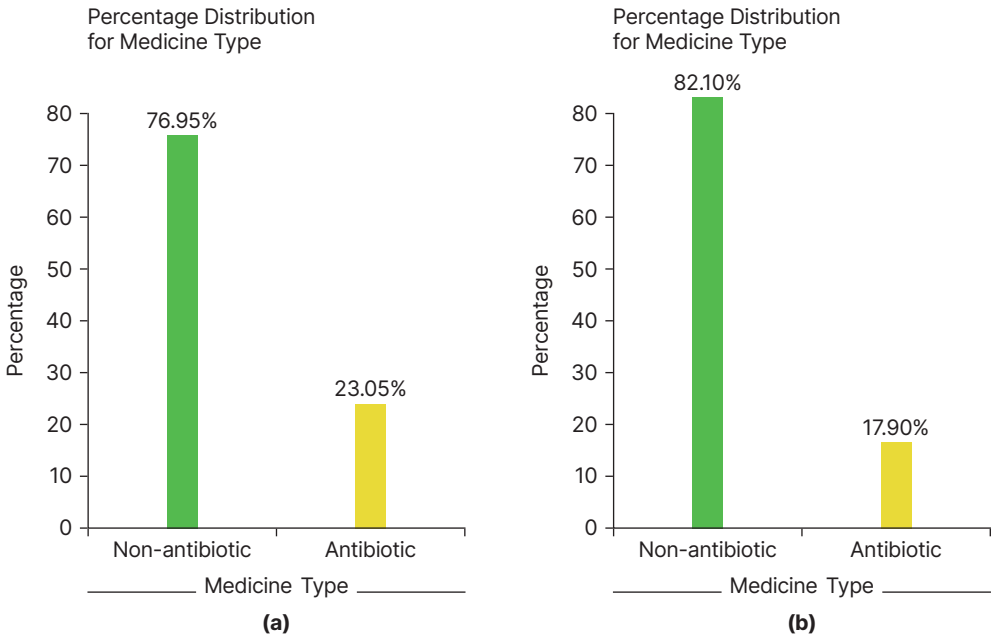
Figure 02: The AMR Admin Dashboard is used for real-time monitoring of pharmacy sales and customer statistics.



The proportion of purchases supported by valid prescriptions increased, while symptom-based antibiotic purchases declined. Analysis of the 2,654 participants reveals that the purchase of medicines with valid prescriptions increased from 27.59% to 43.49%. Conversely, irrational purchases based solely on symptoms dropped from 18.84% to 6.42%.

Protocol-aligned AI triage to curb non-prescription antibiotic use

Figure 03: Comparison of medicine type distribution showing the reduction in antibiotic purchases from 23.05% (Pre) to 17.90% (Post).



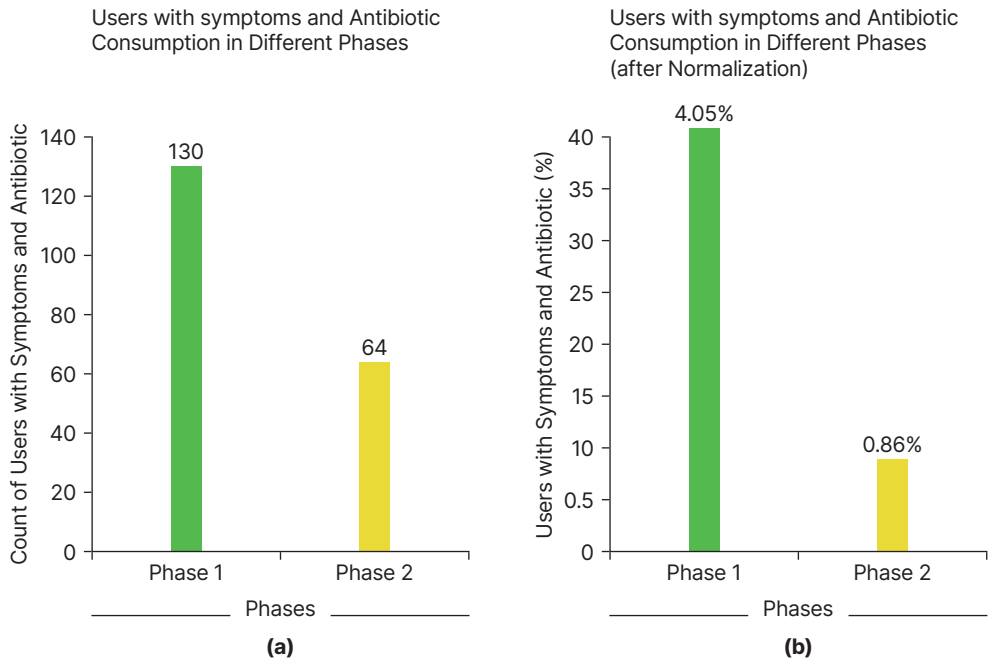
Overall, antibiotic purchases as a share of total medicine volume decreased during the post-intervention period, indicating a potential reduction in unnecessary dispensing. Notably, non-prescription antibiotic purchases based solely on symptoms were substantially reduced. These observed changes are consistent with improved adherence to stewardship protocols, though they cannot be attributed solely to the AI system without further controlled evaluation.

The observed reduction in antibiotic purchases relative to total medicine volume is consistent

with reduced unnecessary dispensing. As illustrated in Figure 3, the percentage of purchased medicines classified as antibiotics dropped from 23.05% in the pre-intervention phase to 17.90% in the post intervention phase.

Most notably, the specific purchase of antimicrobial drugs without a prescription (based purely on symptoms) was substantially reduced. As shown in Figure 4, this high risk behavior dropped from 4.05% to just 0.86% after normalization.

Figure 04: Impact on irrational use: Symptom based antibiotic purchases dropped from 4.05% to 0.86%.




The findings indicate that AI-supported workflows, when embedded within pharmacy operations and coupled with telemedicine referral mechanisms, may support more rational antibiotic use in community settings. Longer-term studies are required to assess sustainability, scalability, and downstream effects on resistance patterns.

Ethics & Governance

Ethical and governance considerations were integral to the design and implementation of

the intervention. The study operated under approval from the Ethics Review Board of United International University, with informed consent obtained from all participants. Consent processes were implemented through mobile number verification, balancing auditability with accessibility.

Data protection measures included anonymisation of personal identifiers prior to analysis and secure handling of patient information in line with standard data protection principles. The system was explicitly designed



as a decision-support tool and is technically constrained from autonomous prescribing. All high-risk triage outcomes trigger mandatory referral to a licensed physician via telemedicine, ensuring clinical accountability remains with human professionals.

Clinical governance was maintained through a human-in-the-loop workflow designed to mitigate liability risks associated with AI in healthcare. The system functions strictly as a decision-support tool and is technically restricted from autonomously prescribing medication; instead, all high-risk triage outcomes trigger a mandatory referral to a human physician via a telemedicine loop for validation. This structure, supported by clinical oversight from Bangladesh Medical University, supports the maintenance of established standards of care. Additionally, the system addresses the black box challenge of AI governance by utilizing an interpretable Ensemble ML model for risk assessment, ensuring that the logic driving triage decisions remains transparent and auditable by clinical stakeholders.

Discussion

The pilot integrated three interoperable digital platforms for pharmacists, patients, and administrators. By linking patients to doctors via telemedicine and supporting real-time digital prescriptions, the system was associated with an observed reduction in unauthorized dispensing and a reduction in the share of antibiotics in total medicine purchases from 23.05% to 17.90%. Digital health accounts further supported longitudinal patient tracking and adherence reminders.

Implementation challenges included resistance from pharmacy owners concerned about revenue loss, incomplete transaction recording, and limited patient willingness to share personal data. As the pilot was confined to two areas in Dhaka, findings are not yet generalizable nationally.

The pilot underscores the importance of aligning pharmacist incentives, strengthening patient trust through simplified and privacy-centred onboarding, and sustained engagement to support scale-up. It also illustrates the model's potential as a digital entry point to primary care, supporting antimicrobial stewardship by reducing empirical self-medication and informing data-driven public health action.

Conclusion


This use case describes a pilot AI-supported digital intervention designed to strengthen antibiotic stewardship in community pharmacy settings in Bangladesh. By integrating AI-assisted triage, digital health accounts, telemedicine referral pathways, and real-time audit mechanisms, the intervention illustrates a feasible approach to addressing irrational antibiotic use at the community level.

While the pilot indicates promising behavioural shifts consistent with improved stewardship,

findings should be interpreted within the context of a proof-of-concept implementation. Future work will focus on formal evaluation, cost-effectiveness analysis, and adaptation for broader deployment. As countries work toward national AMR action plans and global commitments, such as Universal Health Coverage and the Sustainable Development Goals, this experience offers contextually relevant insights into how AI can support responsible, people-centred antimicrobial stewardship in low- and middle-income settings.

References

- Ahmed, R., Bashir, A., Brown, J. E., Cox, J. A., Hilton, A. C., & Hilton, C. E. (2020). The drugs don't work: Evaluation of educational theatre to gauge and influence public opinion on antimicrobial resistance. *Journal of Hospital Infection*, *104*(2), 193–197. <https://doi.org/10.1016/j.jhin.2019.11.012>
- Bangladesh Medical University. (n.d.). *Bangladesh Medical University*. <https://bmu.ac.bd/>
- Chokshi, A., Sifri, Z., Cennimo, D., & Horng, H. (2019). Global contributors to antibiotic resistance. *Journal of Global Infectious Diseases*, *11*(1), 36–42. https://doi.org/10.4103/jgid.jgid_110_18
- CMED Health Limited. (n.d.). *CMED Health Limited*. <https://cmed.com.bd/>
- de Kraker, M. E. A., Stewardson, A. J., & Harbarth, S. (2016). Will 10 million people die a year due to antimicrobial resistance by 2050? *PLoS Medicine*, *13*(11), e1002184. <https://doi.org/10.1371/journal.pmed.1002184>
- Iyawa, G. E., Herselman, M., & Botha, A. (2016). Digital health innovation ecosystems: From systematic literature review to conceptual framework. *Procedia Computer Science*, *100*, 244–252. <https://doi.org/10.1016/j.procs.2016.09.149>
- Klein, E. Y., Van Boeckel, T. P., Martinez, E. M., Pant, S., Gandra, S., Levin, S. A., & Laxminarayan, R. (2018). Global increase and geographic convergence in antibiotic consumption between 2000 and 2015. *Proceedings of the National Academy of Sciences*, *115*(15), E3463–E3470. <https://doi.org/10.1073/pnas.1717295115>
- Kunin, C. M. (1993). Resistance to antimicrobial drugs—A worldwide calamity. *Annals of Internal Medicine*, *118*(7), 557–561. <https://doi.org/10.7326/0003-4819-118-7-199304010-00011>
- McNulty, C. A. M., Syeda, R. B., Brown, C. L., Bennett, C. V., Schofield, B., Allison, D. G., & Francis, N. (2020). Peer education as a tool to educate on antibiotics, resistance and use in 16–18-year-olds: A feasibility study. *Antibiotics*, *9*(4), 146. <https://doi.org/10.3390/antibiotics9040146>
- Murray, C. J. L., Ikuta, K. S., Sharara, F., Swetschinski, L., Aguilar, G. R., Gray, A., ... Naghavi, M. (2022). Global burden of bacterial antimicrobial resistance in 2019: A systematic analysis. *The Lancet*, *399*(10325), 629–655. [https://doi.org/10.1016/S0140-6736\(21\)02724-0](https://doi.org/10.1016/S0140-6736(21)02724-0)
- National Institute for Health and Care Excellence. (2017). *Antimicrobial stewardship: Changing risk-related behaviours in the general population* (NICE guideline). <https://www.nice.org.uk>




Parveen, M., Molla, M. M. A., Yeasmin, M., Nafisa, T., Barna, A. A., & Ghosh, A. K. (2020). Evidence on irrational antimicrobial prescribing and consumption among COVID-19 positive patients and possible mitigation strategies: A descriptive cross-sectional study. *Bangladesh Journal of Infectious Diseases*, 7, S3–S7.

Sultana, R., Rahman, M. M., Nayema, J., Bethe, M. I., Kabbo, S. A., Shornil, M. A. A. B., ... & Acharjee, M. (2025). Antibacterial Potential of Stevia rebaudiana Extract (Methanol and Ethanol) against Clinically Important Pathogens with Time-dependent Bactericidal Potency. *Biomedical and Biotechnology Research Journal (BBRJ)*, 9(3), 315-321.

The Business Standard. (2023, November 22). Antibiotic resistance in Bangladesh rises 11% in 5 years: IEDCR. *The Business Standard*. <https://www.tbsnews.net/bangladesh/health/antibiotic-resistance-bangladesh-rises-11-5-years-iedcr-744522>



World Health Organization. (2014). *Antimicrobial resistance: Global report on surveillance*. World Health Organization. <https://www.who.int/publications/item/9789241564748>

Zujbe, Z., Tabassum, N., Kibria, G., Ahmed, A., Chowdhury, M. M. H., & Hossain, M. S. (2025). The burden of hospital-acquired infections and antibiotic resistance in Bangladesh: Current trends and policy implications. *Discover Public Health*, 22(1), Article 290. <https://doi.org/10.1007/s44197-025-00290-x>



AI-Driven Environmental Health Risk Prediction in Urban Morocco: The OCT7 SI Framework in Meknès

Hajja Youssef



AI-Driven Environmental Health Risk Prediction in Urban Morocco: The OCT7 SI Framework in Meknès

Hajja Youssef*

Abstract

Urban environments in low- and middle-income countries are increasingly affected by environmental health risks associated with air pollution, extreme temperatures, and rapid urbanisation. In Morocco, mid-sized cities such as Meknès face growing public health challenges related to respiratory illnesses and heat-related conditions. This chapter presents OCT7 SI (Operational Climate and Health Threat Intelligence System), an AI-supported framework developed to enhance environmental health risk surveillance and preparedness in Meknès.

The system integrates multi-source datasets, including meteorological variables, air pollution indicators, and anonymised hospital admission records, and applies machine learning methods to identify patterns associated with elevated health risk periods. Rather than replacing existing surveillance mechanisms, OCT7 SI is designed to complement conventional public health reporting by providing earlier situational awareness to support preparedness and response planning.

This use case describes the system architecture, implementation process, observed operational outcomes, and governance considerations. It concludes with lessons learned and future directions relevant to cities in Morocco and comparable urban contexts in the Global South.


Background

Urbanisation is a defining demographic and environmental trend across many low- and middle-income countries, reshaping patterns of exposure and vulnerability. In Morocco, expanding urban areas are accompanied by increased traffic density, industrial activity, and land-use changes that contribute to air pollution and urban heat island effects. Mid-sized cities such as Meknès, characterised by

mixed residential, industrial, and agricultural zones, are particularly exposed to compounded environmental health risks.

A substantial body of evidence links environmental factors, including particulate matter (PM_{2.5} and PM₁₀), nitrogen dioxide (NO₂), ozone (O₃), and extreme temperatures to increased respiratory and cardiovascular

*Bio-Actives, Health & Environment Laboratory, Faculty of Sciences, Moulay Ismail University of Meknes, Morocco



morbidity. Heatwaves are becoming more frequent and intense in North Africa, amplifying risks of heat stress and exacerbation of chronic disease. Public health surveillance systems in Morocco, as in many settings, rely primarily on retrospective reporting of healthcare utilisation, which can limit the ability to anticipate periods of elevated risk and implement timely preventive measures.

Advances in artificial intelligence (AI) and machine learning offer opportunities to enhance environmental health surveillance by integrating heterogeneous datasets and identifying complex, non-linear relationships between environmental exposures and health outcomes. While such approaches have been explored extensively in high-income settings, fewer documented examples exist from African urban contexts. This use case contributes to addressing this gap by describing the design and early operational experience of an AI-supported environmental health surveillance framework developed for Meknès.

OCT7 SI system uses AI to predict environmental health risks in urban Meknès, Morocco

AI Use Case Description

Solution Description

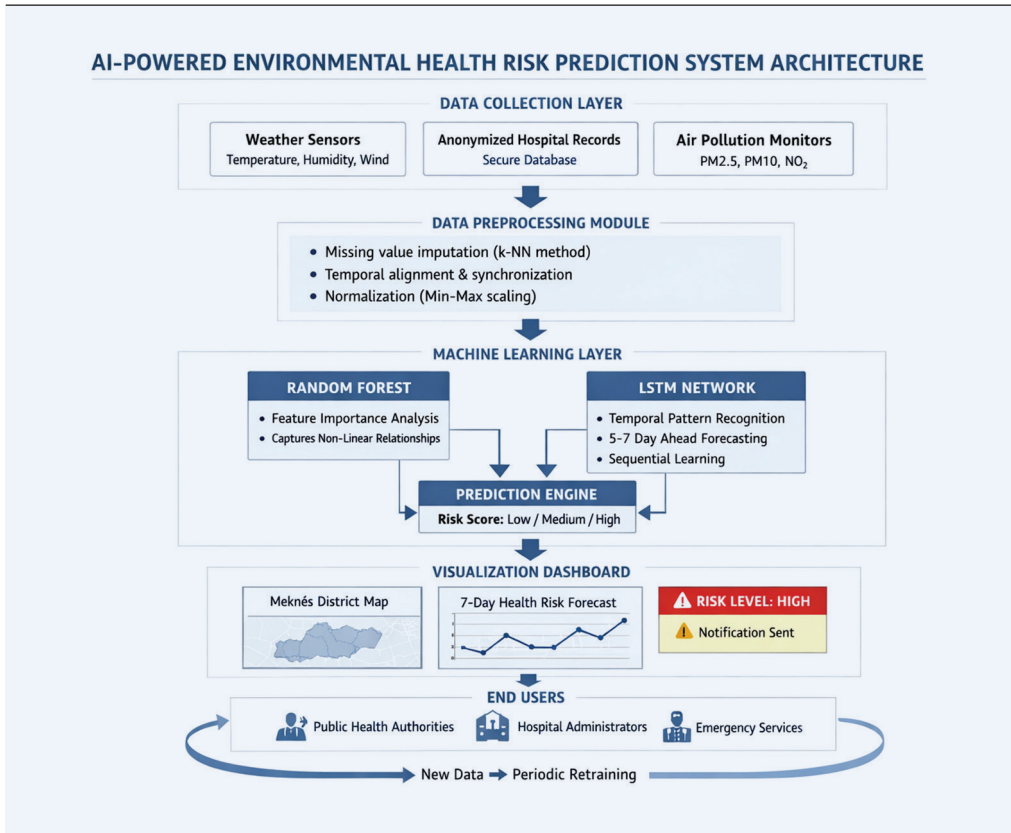
OCT7 SI (Operational Climate and Health Threat Intelligence System) is a modular framework designed to support environmental health risk surveillance at the population level. The system integrates three primary categories of data:

1. Meteorological variables, including temperature, humidity, and wind speed
2. Air pollution indicators, including PM2.5, PM10, and nitrogen dioxide
3. Anonymised hospital admission records related to respiratory and heat-related conditions

Data are aggregated at a daily temporal resolution and analysed at the district level to support sub-city risk assessment.

Two Machine learning methods were selected to address different analytical objectives. Random Forest models were used to explore non-linear relationships between environmental variables and health outcomes and to support feature relevance analysis. Long Short-Term Memory (LSTM) models were applied to capture temporal patterns and delayed associations between environmental conditions and observed health system utilisation. Model outputs are translated into categorical risk signals intended to support preparedness planning rather than precise case forecasting.

Figure 01: Presents the complete architecture of the OCT7 SI system, illustrating the workflow from data collection through preprocessing, machine learning modeling, prediction, and visualization




Deployment and Implementation

The OCT7 SI framework was developed through an iterative process encompassing data acquisition, preprocessing, model training, validation, and deployment. Meteorological and air pollution data were obtained from local monitoring stations, while hospital admission

data were provided by regional health authorities following anonymisation procedures.

Data preprocessing included handling missing values, normalisation, and temporal alignment across data streams. Models were trained on historical data covering a multi-year period and evaluated using time-aware validation



approaches appropriate for sequential data. Predictive outputs were operationalised through a secure web-based dashboard accessible to authorised public health stakeholders. The system is updated on a regular schedule to incorporate newly available data and to maintain relevance to evolving environmental conditions.

Scale of Deployment

The initial deployment of OCT7 SI focused on the urban area of Meknès, serving a population of approximately 600,000 residents. Risk signals are generated for multiple administrative districts within the city to support localised situational awareness. While implemented at a municipal scale, the system architecture is designed to support adaptation to additional cities, subject to data availability and institutional arrangements.

Identified Users

Primary users of OCT7 SI include regional public health authorities, municipal decision-makers, and hospital administrators. Secondary users include emergency response services and urban planners who may benefit from early indications of elevated environmental health risk. Community-facing outputs are designed to support targeted risk communication through existing public health channels.

Impact

The primary contribution of OCT7 SI lies in enhancing the timeliness and preparedness dimensions of environmental health surveillance. During pilot implementation, AI-supported analysis identified periods in which

combinations of air pollution and temperature conditions were associated with subsequent increases in hospital admissions for respiratory and heat-related conditions.


These risk signals supported earlier coordination among public health authorities and healthcare facilities during selected high-risk periods, including extreme heat events. While operational observations suggest that earlier situational awareness can facilitate proactive measures—such as public advisories and resource mobilisation—the system's contribution should be understood as supportive and complementary to existing public health functions.

Quantitative performance metrics were assessed to evaluate technical feasibility within the pilot context rather than to establish definitive benchmarks. Longer-term evaluation is required to assess population-level health outcomes and to better understand the conditions under which AI-supported surveillance adds the greatest value.

Ethics and Governance

Ethical and governance considerations were incorporated throughout the OCT7 SI lifecycle in line with internationally recognised principles for responsible AI in health, including those articulated by the World Health Organization. Health data were anonymised prior to analysis and processed in accordance with national data protection legislation (Morocco's Data Protection Law (09-08) and GDPR-inspired principles)

Human oversight remains central to system use. AI-generated outputs are reviewed and



interpreted by public health professionals, and the system is designed to inform population-level preparedness rather than individual-level decisions. Transparency is promoted through clear communication of system scope, assumptions, and limitations to end users. Governance arrangements emphasise accountability, inclusiveness, and alignment with public health priorities.

Discussion

Successes

A key strength of OCT7 SI is its grounding in locally available data and its alignment with operational public health needs. Collaboration between environmental monitoring entities and health authorities facilitated system development and user engagement. The combination of machine learning approaches enabled exploration of both environmental-health associations and temporal patterns relevant to preparedness planning.

Fuses weather, pollution, and hospital data streams

Challenges

Implementation challenges included data gaps, variability in data quality across sources, and limited interoperability between environmental and health information systems. Capacity constraints related to AI literacy among stakeholders necessitated targeted training and the development of accessible visualisation tools to support interpretation and trust.

Limitations

The system was developed and evaluated in a single urban context, which limits generalisability. Hospital admission data may underrepresent less severe cases managed outside formal healthcare settings. Environmental monitoring coverage may not fully capture intra-urban variability. Additionally, part of the training period coincided with the COVID-19 pandemic, which may have influenced healthcare utilisation patterns.

Lessons Learned

The OCT7 SI experience underscores the importance of context-aware system design, early stakeholder engagement, and transparent communication of system capabilities and limitations. AI-supported surveillance is most effective when embedded within existing public health processes and supported by appropriate governance and capacity-building measures.



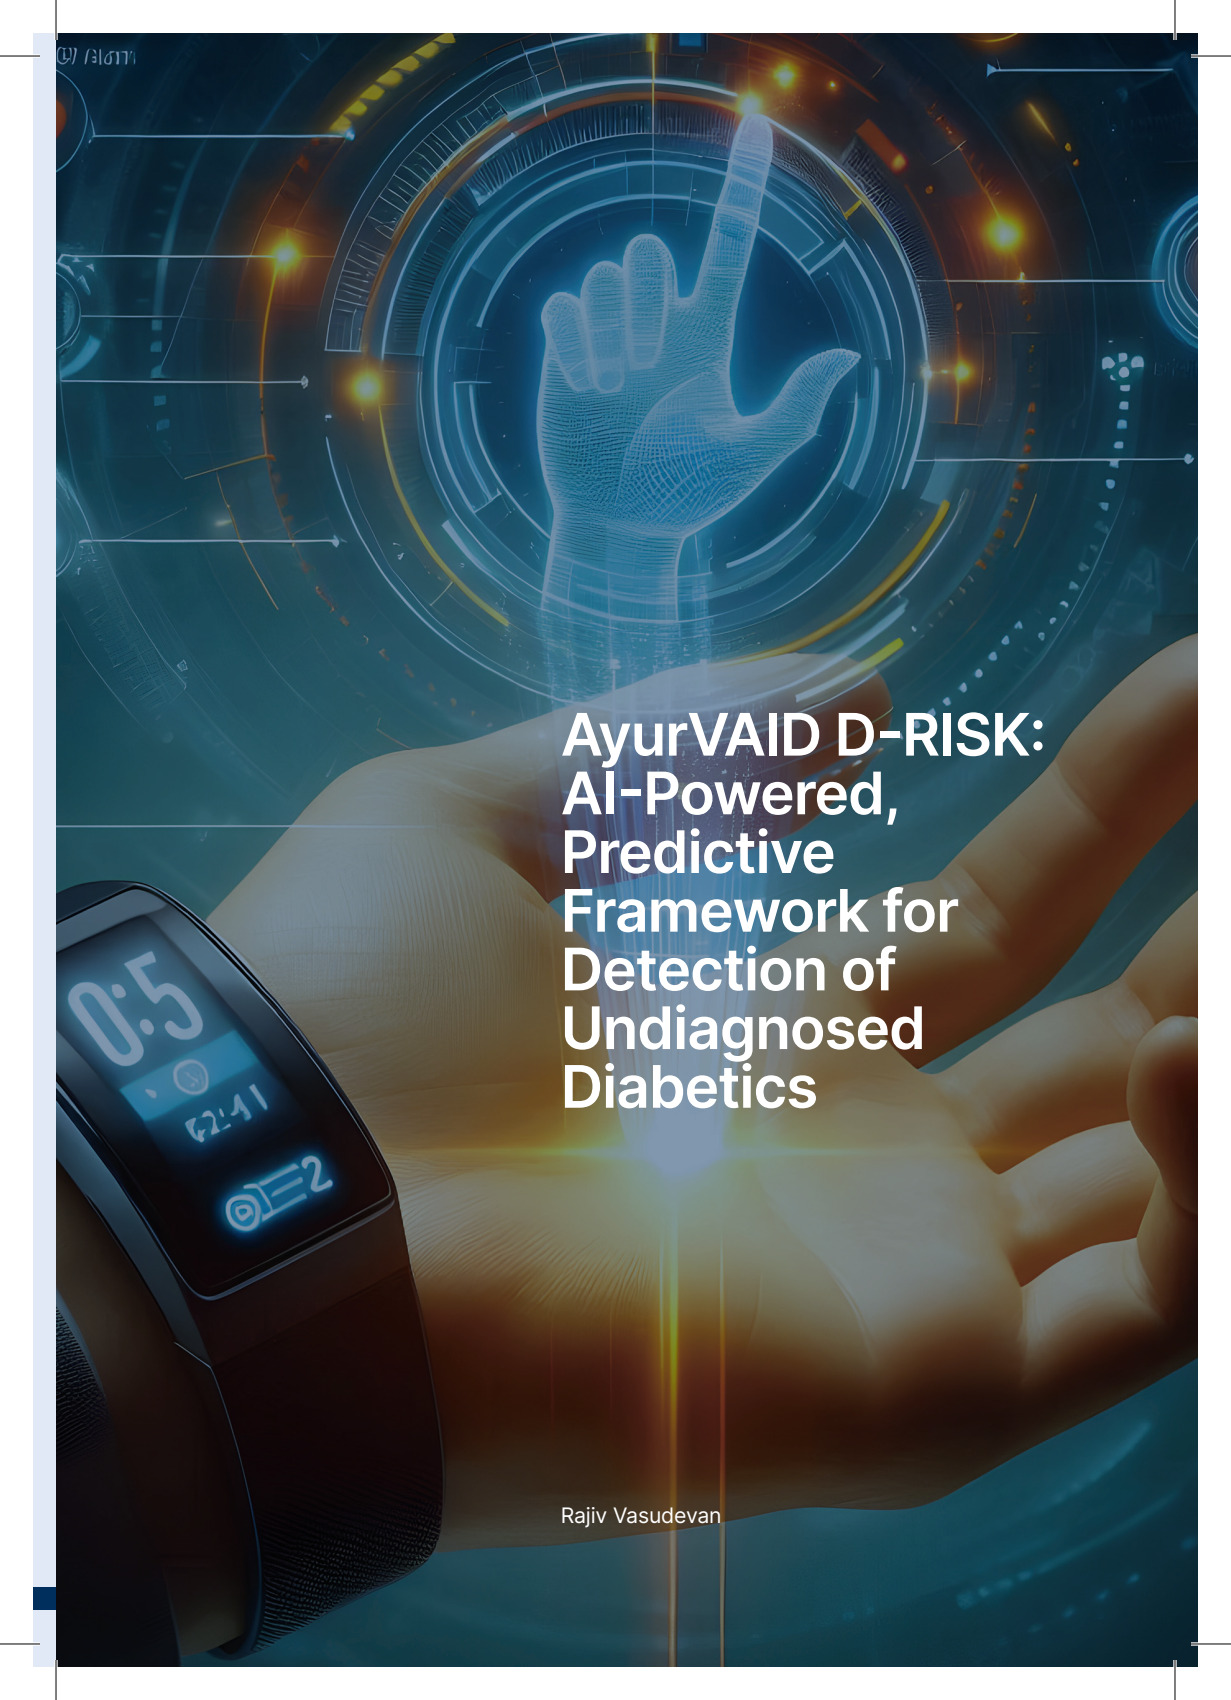
Conclusion

This case study illustrates how AI-supported analysis can be operationalised to strengthen environmental health surveillance in a mid-sized urban setting in Morocco. By integrating multi-source environmental and health data, OCT7 SI provides earlier situational awareness that can support preparedness and response planning.

Future development priorities include expanding the range of environmental health outcomes considered, integrating additional data sources such as satellite observations, and exploring collaborative learning approaches that preserve data privacy while improving robustness across settings. With continued evaluation, governance, and capacity building, AI-supported systems such as OCT7 SI may contribute to more resilient and climate-responsive public health surveillance in urban contexts across the Global South.


References

- Breiman, L. (2001). Random forests. *Machine Learning*, 45(1), 5–32. <https://doi.org/10.1023/A:1010933404324>
- Chen, X., et al. (2020). Machine learning for environmental health surveillance. *Environmental Health Perspectives*, 128(9), 097001. <https://doi.org/10.1289/EHP6446>
- Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), 1735–1780. <https://doi.org/10.1162/neco.1997.9.8.1735>
- Intergovernmental Panel on *Climate Change*. (2022). *Climate change 2022: Impacts, adaptation and vulnerability* (Sixth Assessment Report). Cambridge University Press. <https://doi.org/10.1017/9781009325844>
- Prüss-Ustün, A., Wolf, J., Corvalán, C., Neville, T., Bos, R., & Neira, M. (2019). *Preventing disease through healthy environments: A global assessment of the burden of disease from environmental risks*. World Health Organization. <https://www.who.int/publications/i/item/9789241565196>
- World Health Organization. (2016). *Ambient air pollution: A global assessment of exposure and burden of disease*. World Health Organization. <https://www.who.int/publications/i/item/WHO-HTM-PHE-2016.01>
- World Health Organization. (2021). *Ethics and governance of artificial intelligence for health*. World Health Organization. <https://www.who.int/publications/i/item/9789240029200>



AyurVAID D-RISK: AI-Powered, Predictive Framework for Detection of Undiagnosed Diabetics

Rajiv Vasudevan



AyurVAID D-RISK: AI-Powered, Predictive Framework for Detection of Undiagnosed Diabetics

Rajiv Vasudevan*

Abstract

Undiagnosed diabetes remains a significant public health concern globally and in many low- and middle-income settings. Conventional laboratory-based glucometric diagnostics typically identify diabetes after sustained metabolic dysregulation is established. Population-level risk scores based on limited variables may have constrained sensitivity and specificity across diverse populations.

AyurVAID D-RISK is an AI-assisted, non-invasive diabetes risk screening framework designed to support early identification of individuals who may have undiagnosed diabetes or elevated metabolic risk. The model integrates demographic, anthropometric, lifestyle, and symptom-based indicators, including features derived from classical Ayurveda (traditional medicine) descriptions of early metabolic imbalance, together with machine learning methods. In model-development, datasets of approximately 12,000 individuals, the AutoML-based framework demonstrated moderate-to-high discrimination metrics under cross-validation conditions.

D-RISK is intended as a screening and triage support tool to help prioritise individuals for confirmatory laboratory testing. Reported metrics reflect model-development validation results and should not be interpreted as definitive clinical diagnostic performance. The framework illustrates how culturally contextualised, non-invasive risk indicators combined with AI methods may support earlier risk stratification when implemented with clinical oversight and governance safeguards.


Background

A. Problem Overview

Diabetes is a chronic metabolic disorder marked by persistently elevated blood glucose due to inadequate insulin production or reduced insulin effectiveness. Its progression from normal glucose levels to overt disease is gradual and

strongly shaped by lifestyle factors such as obesity, physical inactivity, and poor diet. Insulin resistance typically appears first, prompting compensatory overactivity of pancreatic beta cells; over time this compensation fails, leading

*Apollo AyurVAID Hospitals, India



to sustained hyperglycaemia often after significant physiological decline has already occurred.

In Ayurveda, the diabetes spectrum (Prameha) describes 22 prodromal features that may appear 2–5 years before measurable hyperglycaemia, including metabolic, digestive, and systemic symptoms. A retrospective cohort study of 141 newly diagnosed patients found that 11 of these prodromal features were significantly more prevalent than in healthy controls ($p < .001$)[1]. Incorporating such indicators into the D-RISK model is intended to support earlier risk stratification, complementing conventional biomarker-based screening approaches.

B. Clinical Relevance

Early Risk Identification

Evidence indicates that insulin resistance and metabolic imbalance precede elevated/ abnormal glucose biomarkers by several years. This asymptomatic phase remains undetected by conventional screening, delaying intervention and increasing healthcare burden. Proactive, non-invasive, population-level screening tools are therefore of significant clinical and public health importance.

Limitations of Current Diagnostics

Definitive diagnosis of diabetes relies on laboratory measures such as HbA1c, fasting plasma glucose, and oral glucose tolerance testing. These tests are clinically established

but require infrastructure, logistics, and cost. They are generally applied after risk suspicion arises. Risk scores such as IDRS support pre-screening but may show variable performance across populations and contexts.

Role of AI/ML in Addressing the Gap

The D-RISK approach applies machine learning to combine more than 40 non-invasive features to estimate the probability that an individual may have undiagnosed diabetes or elevated metabolic risk. Many of these features have weak predictive value individually but may contribute useful signal when analysed jointly through multivariable models. Machine learning methods support modelling of non-linear feature interactions and subgroup variation across demographic and regional strata. They also enable post hoc interpretability analyses to estimate relative feature contributions at the model level. This approach aligns with broader WHO priorities encouraging responsible use of AI-enabled digital tools, including in traditional, complementary, and integrative medicine contexts, provided that safety, transparency, and validation requirements are met. AI outputs in this framework are intended to support screening decisions, not to replace diagnostic testing.



Development Process

Scope


The project defines a structured, self-reported questionnaire designed to support a prediction model for identifying undiagnosed diabetes risk. Gestational diabetes is excluded because it is a pregnancy-specific, short-term condition. Data governance relies on patient self-reported questionnaires covering demographics, symptoms, lifestyle factors, and medication use, with audit and provenance controls to record who entered the data, when it was entered, and whether inputs were lab-verified or self-reported. Data cleaning and standardization procedures include validation of externally captured inputs to detect and flag incorrect values for clinician review or exclusion, normalization of free-text responses into predefined categories, and a minimal, clinically conservative imputation strategy that prioritizes clinically verified replacements wherever possible.

Model Development and Validation

Model development and validation combined culturally contextualized feature engineering with automated ensemble learning. The model incorporates Ayurveda diagnostic indicators alongside conventional clinical and lifestyle variables to better reflect India's pluralistic healthcare ecosystem

and capture early metabolic imbalances not typically included in standard Western screening tools. Feature sets include demographics (age, gender, family history), anthropometric measures (waist and hip circumference), lifestyle factors (physical activity, diet patterns, sleep quality), and Ayurveda clinical indicators such as Agni Vaishamyam and Saamatva lakshanas, which correspond to metabolic dysregulation and inflammatory states associated with insulin resistance. Model training used the H2O AutoML framework to train and cross-validate multiple candidate algorithms and automatically select the top-performing ensemble (H2O Leader) based on validation results. Comparative evaluation showed the ensemble consistently outperformed individual models such as Logistic Regression, CatBoost, XGBoost, and deep learning approaches, capturing approximately 95–98% of positive cases within the top 30% of predicted risk scores—supporting efficient prioritization for confirmatory testing in resource-constrained settings. The model was trained on a dataset of about 12,000 individuals using five-fold cross-validation, with stratified sampling and threshold tuning to address class imbalance. Performance metrics include 78% accuracy, 75% sensitivity, 86% specificity, and an AUC-ROC of 0.87, indicating moderate-to-high discrimination under cross-validation conditions. SHAP (SHapley Additive exPlanations) values provide feature importance rankings and individual prediction explanations, enabling clinicians to understand and validate model decisions. These artefacts are intended to support clinical review and governance requirements.

Figure 01: Lift Chart: Comparison of H2O AutoML Ensemble and Individual Algorithms

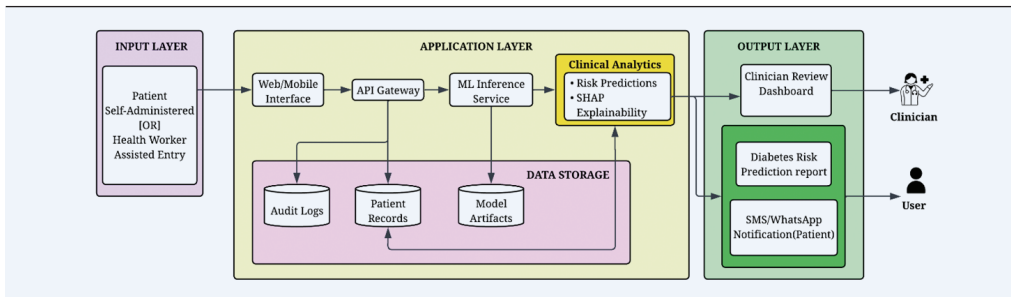


These metrics are derived from internal cross-validation and indicate discrimination performance under model-development conditions. They should be interpreted as illustrative performance indicators, not guaranteed field performance.

Clinical Integration and Deployment

A. Clinical Workflow and System Integration

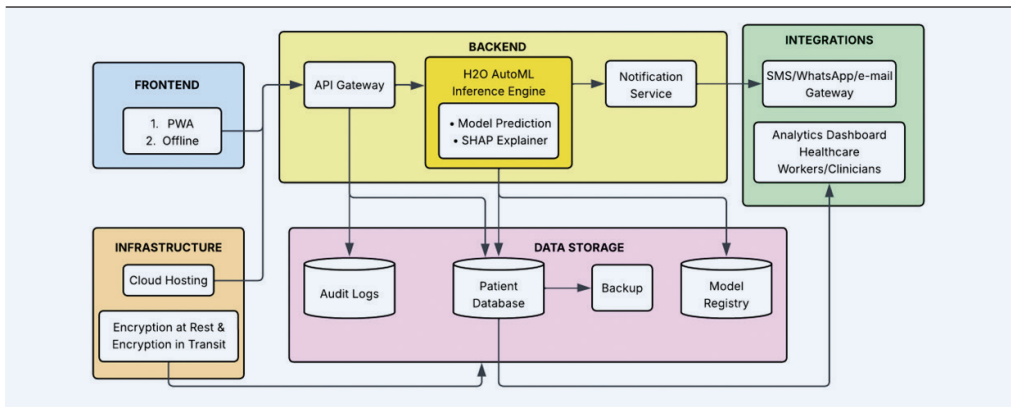
Figure 02: D-RISK Clinical Workflow and System Integration



The workflow (Figure 2) illustrates end-to-end patient flow and clinical decision support integration. Users access the platform through web/mobile interfaces with role-specific functionalities: patients complete self-assessments, health workers provide assisted data entry, and clinicians review the predictions via dashboard.

B. System Architecture and Infrastructure

Figure 03: D-RISK Multi-Tier System Architecture and Infrastructure



The D-RISK technical architecture (Figure 3) comprises five integrated layers:

- Frontend:** Progressive Web App (PWA) with offline capability, supporting patient self-assessment, health worker-assisted entry, and clinician review dashboards
- Backend:** RESTful API Gateway orchestrating requests to H2O AutoML inference engine, model result and SHAP explainer (feature importance visualization)
- Data Storage:** Encrypted databases for patient records, versioned model registry with rollback capability, immutable audit logs, and automated backup with 30-day retention
- Integrations:** Bidirectional EMR connectivity, SMS/WhatsApp notifications, and analytics dashboards for program-level monitoring
- Infrastructure:** Cloud hosting, encryption in transit, encryption at rest, DPDP and HIPAA compliant. PII & sensitive data are encrypted both at rest & in transit adhering to Data privacy regulations as per ISO/IEC 27018 – PII Protection Standard

This modular design enables phased deployment from pilot sites to national scale while maintaining security, interoperability, and clinical safety standards.

C. Roles and Responsibilities

- **Clinical Lead:** Approves thresholds, pathways, and model updates in production
- **Data Steward (IT):** Enforces validation rules and manages flowchart-based data entry constraints
- **ML Engineer:** Maintains model services, calibration, and CI/CD pipelines
- **Data Scientist:** Performance monitoring, drift analysis, retraining, and documentation (model cards)
- **Privacy Officer:** Compliance audits, consent tracking, and data-sharing controls

Implementation Scale and Geographical Coverage

A. Phased Rollout Plan

Preparation and pilot phases have been completed in selected urban and semi-urban

centres, with calibration and user-interface refinement. Regional and multi-centre scale-up phases are planned, subject to governance approvals and operational readiness. Federated learning approaches are under consideration where data-sharing constraints exist.

B. Identified Users


- **Primary:** Individuals using self-assessment tools for early risk awareness and referrals
- **Secondary:** Community health workers and clinicians using D-RISK to triage a group and prioritize confirmatory testing
- **Tertiary:** Program managers and policymakers using aggregate outputs (prevalence estimates, NNS) to plan screening and resource allocation

Impact and Value

Note: The models compared in Table 1 were developed using different populations, feature sets, data sources, and validation methodologies. These differences may substantially influence performance metrics such as sensitivity,

Table 01: Performance metrics for the US AutoML model and IDRS for diabetes detection are reported as published in the respective studies[6][7] and are presented for comparative, benchmarking purposes

Model	Sensitivity (Recall)	Specificity	PPV(Positive Predictive Value)	NPV(Negative Predictive Value)	AUC	Accuracy
US AutoML Model (JMIR AI, 2023)	70.26%	90.46%	64.10%	92.61%	90.90%	-
Indian Diabetes Risk Score – IDRS (MDRF)	72.50%	60.10%	17.00%	95.10%	-	61.30%
Apollo AyurVAID - D-RISK	75.00%	86.00%	62.00%	94.00%	87.00%	78.00%



specificity, AUC, and predictive values. Therefore, the results should not be interpreted as direct head-to-head evaluations, and caution should be exercised when comparing performance across models.

Compared to IDRS, the D-RISK model shows higher specificity, PPV, and accuracy values in development validation, though direct comparison is limited by differences in study populations and methods. When compared with a published US AutoML model trained on an older population (average age 59 years, with multimorbidity), D-RISK shows a different sensitivity–specificity balance appropriate to its target context. These cross-study comparisons are illustrative and should be interpreted with caution given differences in populations, feature sets, and validation approaches. D-RISK’s metrics reflect internal cross-validation in an Indian general population (average age 45 years) and have not yet been validated prospectively or in external datasets.

If validated prospectively, a non-invasive, low-cost screening tool of this kind could support more efficient use of confirmatory testing and help prioritise clinical follow-up in resource-constrained settings.

Ethics, Governance, and Responsible Deployment

Privacy-by-design aligned with DPDP and HIPAA principles including consent, data minimisation, pseudonymisation, and role-based access controls, aligned with applicable data protection laws. Consent for use of patient data on redacted/anonymised basis was obtained from each individual at the time of registration. Human oversight is required, and the model is deployed as a clinician-reviewed triage support tool. Explainability artefacts (model cards, SHAP summaries) document intended use, limitations, and performance. Subgroup monitoring by age and gender is planned to identify disparities and support recalibration. Responsible use depends on continuous monitoring and governance.

Discussion: Successes and Challenges

Successes

The framework demonstrates feasibility of combining culturally contextualised indicators (Ayurveda descriptors), with AutoML methods for non-invasive diabetes risk screening. Ensemble modelling showed stable discrimination in development datasets, and explainability tools supported clinician review.

Challenges

Self-reported data introduce measurement error and label noise. Missing data handling and workflow redesign were required. Class imbalance and subgroup disparities required threshold tuning and monitoring. External generalisability remains to be established.

Lessons Learned and Future Enhancements

Key lessons include the importance of domain knowledge integration, explainability for clinician acceptance, and strong data governance. Planned enhancements include prospective blinded validation against

gold-standard biomarkers, subgroup-aware recalibration, improved data capture quality, and multi-region pilots to assess operational effectiveness beyond initial settings.

Conclusion

AyurVAID D-RISK illustrates that a culturally contextualised, non-invasive, AutoML-based risk model can function as a screening and triage support tool for identifying individuals who may have undiagnosed diabetes and should receive confirmatory testing. Development-phase validation indicates useful discrimination under internal evaluation conditions. Responsible scale-up will require prospective external validation, fairness assessment, input standardisation, and operational pilots assessing clinical and economic value.

References

Sanjeev Rastogi, Neelendra Singh, Manish Gutch, Arindam Bhattacharya, Prameha purvaroopas as diabetes risk predictor - trends from a retrospective cohort study of newly diagnosed type 2 diabetes patients, *Journal of Ayurveda and Integrative Medicine*, Volume 14, Issue 1, 2023, 100671, ISSN 0975-9476, <https://doi.org/10.1016/j.jaim.2022.100671>. (<https://www.sciencedirect.com/science/article/pii/S0975947622001309>)

International Diabetes Federation. (2024). *IDF diabetes atlas* (11th ed.). <https://diabetesatlas.org/data-by-location/country/india/>

Kurian, J., Bilagi, A., & Deo, G. (2024). Etiopathogenesis and management of Prameha Poorvaroopas (prediabetes), *Prameha and Madhumeha (diabetes mellitus): An ancient perspective. Advances in Mind-Body Medicine*, 38(3), 47–51. <https://pubmed.ncbi.nlm.nih.gov/39213491> (pubmed.ncbi.nlm.nih.gov in Bing)

Rastogi, S., Chaudhari, P., & Singh, R. H. (2018). Linking Prameha etiology with diabetes mellitus: Inferences from a matched case-control study. *AYU*, 39(3), 139–145.

American Diabetes Association. (2025). Classification

and diagnosis of diabetes. *Diabetes Care*, 48(Suppl. 1), S20–S42. <https://doi.org/10.2337/dc25-S002>

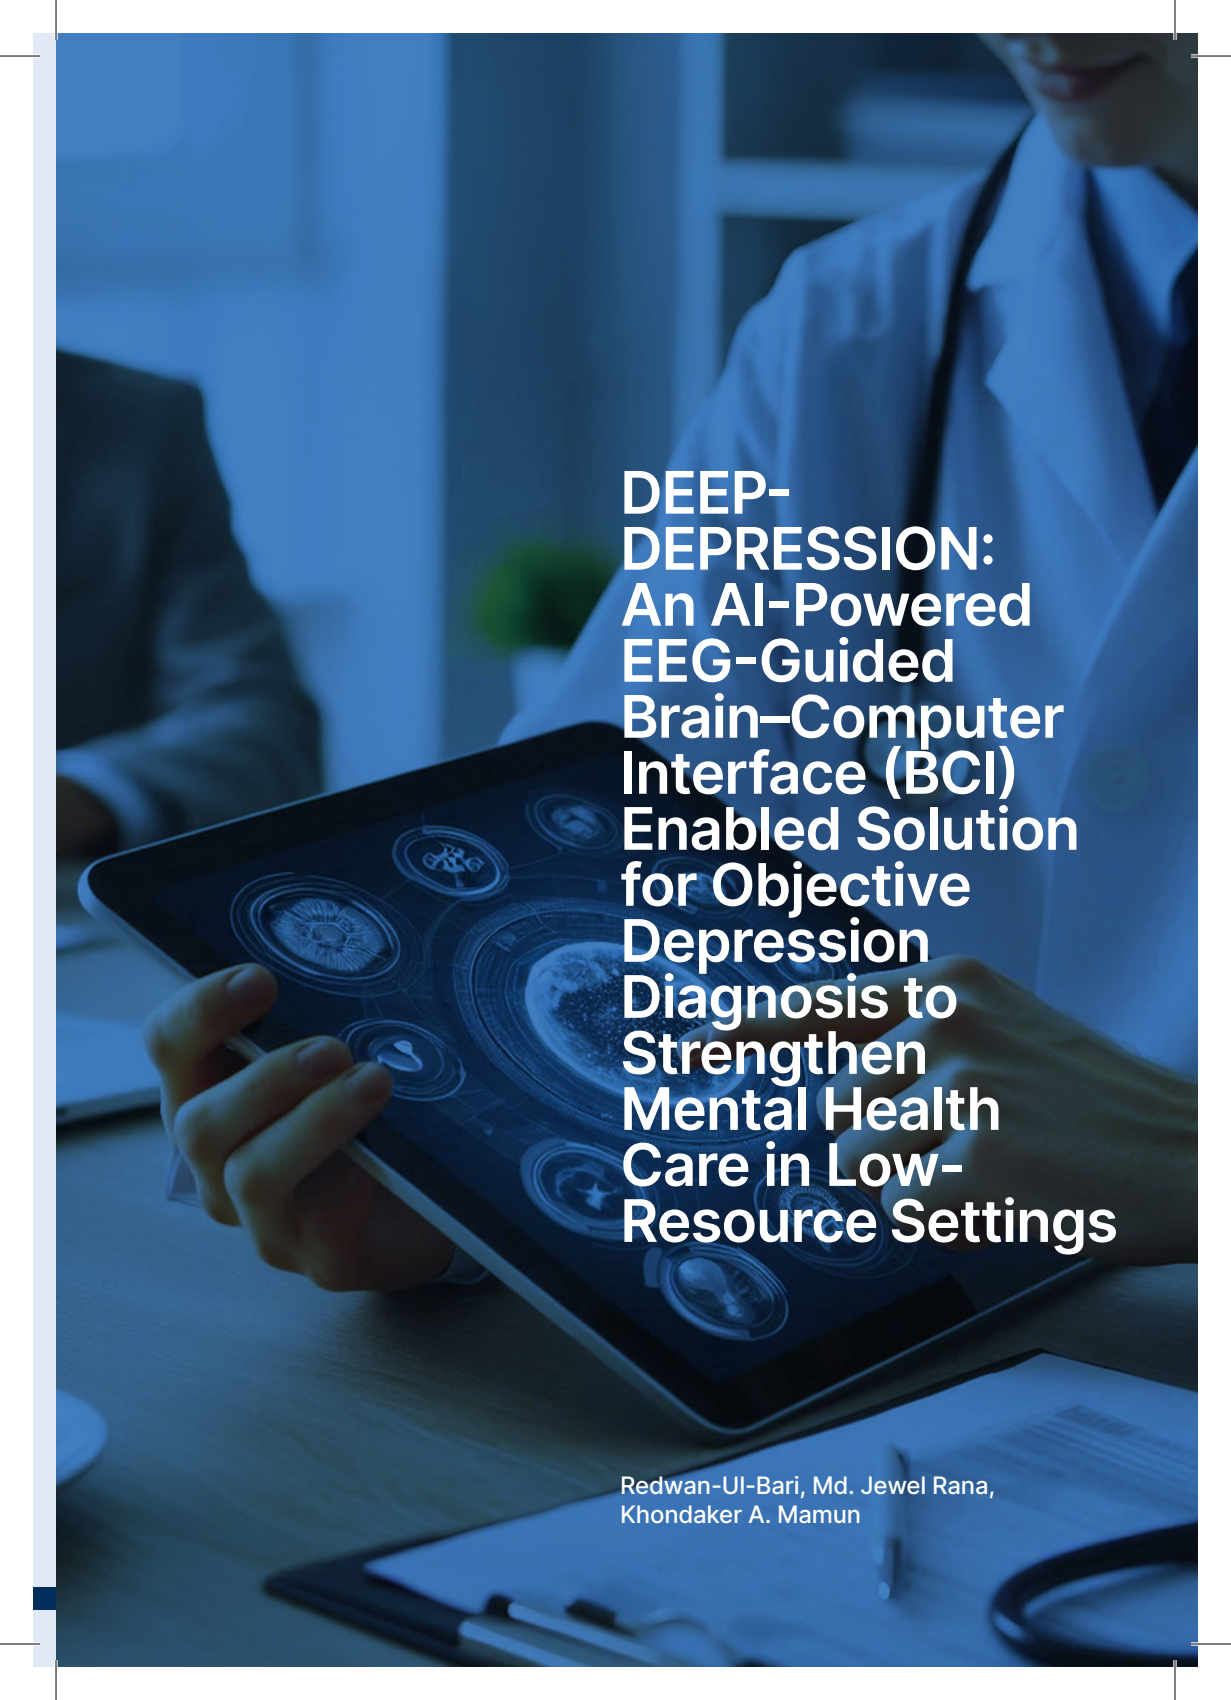
World Health Organization. (2025). *Global traditional medicine strategy 2025–2034* (Health Systems Performance Assessment toolkit for universal health coverage). Geneva: World Health Organization.

Mohan, V., & Anbalagan, V. (2013). Expanding role of the Madras Diabetes Research Foundation–Indian Diabetes Risk Score in clinical practice. *Indian Journal of Endocrinology and Metabolism*, 17(1), 31–36. <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3659903/>

Liu, J., Ssewamala, F., An, R., & Ji, M. (2025). Use of automated machine learning to detect undiagnosed diabetes in US adults: Development and validation study. *JMIR AI*, 4, e68260. <https://ai.jmir.org/2025/1/e68260/>


Government of India. (2023). *Digital Personal Data Protection Act*.

U.S. Department of Health & Human Services. (n.d.). *HIPAA security rule*.



**DEEP-
DEPRESSION:
An AI-Powered
EEG-Guided
Brain-Computer
Interface (BCI)
Enabled Solution
for Objective
Depression
Diagnosis to
Strengthen
Mental Health
Care in Low-
Resource Settings**

Redwan-UI-Bari, Md. Jewel Rana,
Khondaker A. Mamun



DEEP-DEPRESSION: An AI-Powered EEG-Guided Brain-Computer Interface (BCI) Enabled Solution for Objective Depression Diagnosis to Strengthen Mental Health Care in Low-Resource Settings

Redwan-Ul-Bari, Md. Jewel Rana, Khondaker A. Mamun*

Abstract

Mental health disorders represent a major public health burden across the Global South. Depression contributes substantially to disability in low- and middle-income countries, including Bangladesh, where diagnostic capacity is uneven and access to specialist assessment is limited. Current diagnostic approaches rely primarily on clinical interviews and self-report measures, which are clinically established but may be affected by reporting bias, time constraints, and variability in provider experience. Electroencephalography (EEG) is a non-invasive and comparatively affordable neurophysiological modality; however, its use in psychiatric assessment has been limited by signal complexity and interpretability challenges.

This use case describes a research-grade brain-computer interface (BCI) prototype that integrates EEG acquisition with artificial intelligence-based analysis of neural connectivity patterns. The system is positioned as a research-to-translation pilot designed to support validation readiness rather than immediate clinical deployment. Under controlled research conditions, the pipeline processes structured EEG recordings to characterize neurophysiological feature patterns associated with depressive cohorts. Internal validation experiments demonstrated classification performance exceeding 90% accuracy; these results are interpreted cautiously given the dataset size, sampling structure, and validation scope.

The system is designed as a clinician-facing decision-support prototype intended for use within supervised psychiatric workflows following formal clinical validation. This chapter documents technical feasibility, pilot-stage evaluation, operational design, governance safeguards, and a staged pathway toward supervised clinical validation and responsible health system integration.

*Advanced Intelligent Multidisciplinary Systems Lab (AIMS Lab), Institute of Research, Innovation, Incubation and Commercialization (IRIIC), United International University (UIU), Dhaka 1212, Bangladesh



Background

Depression presents a sustained challenge for health systems across many Global South settings due to high prevalence, delayed detection, fragmented referral pathways, and limited access to specialist care. These factors are associated with delayed diagnosis and prolonged untreated illness, contributing to avoidable disability and social and economic burden. In Bangladesh and similar contexts, mental health literacy gaps, stigma, and treatment cost barriers reduce timely help-seeking. Specialist services are concentrated in urban tertiary centers, while many individuals rely on informal or non-specialist providers. As a result, depressive disorders may be under-recognized or diagnosed late. Current diagnostic practice relies on structured interviews and standardized scales. These approaches are clinically accepted but depend on consultation time, patient disclosure, and provider experience. Underreporting, somatic presentations, and brief consultations may reduce diagnostic confidence in some settings. From a systems perspective, this creates a gap in objective complementary measures that can support but not replace clinical assessment.

EEG offers a potential complementary modality because it is non-invasive, comparatively affordable, and already used in neurology services. However, depression-related EEG signatures are not visually apparent, and raw EEG signals are high-dimensional and difficult to interpret clinically. This has limited its routine psychiatric use despite operational feasibility advantages.

Recent advances in AI enable the extraction of multivariate patterns from complex EEG


signals. However, international experience indicates that laboratory performance does not automatically translate into clinical utility. Robust validation, workflow alignment, interpretability, and governance safeguards are required. Current evidence suggests EEG-based depression indicators are best understood as multivariate feature patterns rather than standalone diagnostic biomarkers. Within this context, the present work is framed as a decision-support research pilot intended to support clinician judgment and structured evaluation, not to replace diagnostic processes. It aims to establish technical and governance readiness for supervised clinical validation in alignment with international guidance on staged AI evaluation in health.

AI Use Case Definition

Core Use Case

The core use case is an EEG-based AI decision-support system intended to assist clinicians and clinical researchers by generating quantitative neurophysiological indicators associated with Major Depressive Disorder (MDD). The system analyzes structured EEG recordings and produces pattern-based outputs designed to support supervised clinical assessment and translational research.

The system is explicitly designed to complement, not replace, clinical interviews, standardized scales, and clinician judgment. It targets diagnostic support in settings where subjectivity, time constraints, stigma,



or limited specialist availability may reduce assessment confidence. It is not designed for autonomous diagnosis or unsupervised use.

Solution Description

Implemented System (Current State)

The implemented system functions as a stable, research-grade EEG-AI pipeline with clearly defined stages and human oversight. It begins with standardized 64-channel EEG acquisition using Brain Vision hardware under controlled protocols, including resting-state (eyes open/closed) and Emotional Stroop task recordings. Data are processed through an auditable preprocessing and quality-control workflow with bandpass filtering, artifact handling, and rereferencing, with separate configurations for resting-state and task paradigms to ensure reproducibility. Feature engineering covers time-domain, frequency-domain, event-related potential, and connectivity measures, followed by wrapper-based recursive feature elimination to produce a compact, computationally efficient feature set. Supervised classification is performed using a Support Vector Machine selected for robustness and interpretability in pilot-scale datasets. Interpretability is further supported through group-level connectivity analysis using multivariate Granger causality for neurophysiological context rather than standalone diagnosis. Model performance is evaluated using 10-fold cross-validation and leave-one-subject-out validation, with structured outputs generated for clinician and researcher review within a human-in-the-loop framework. At this stage, the system is positioned as a pre-validation, clinician-facing decision-support prototype,

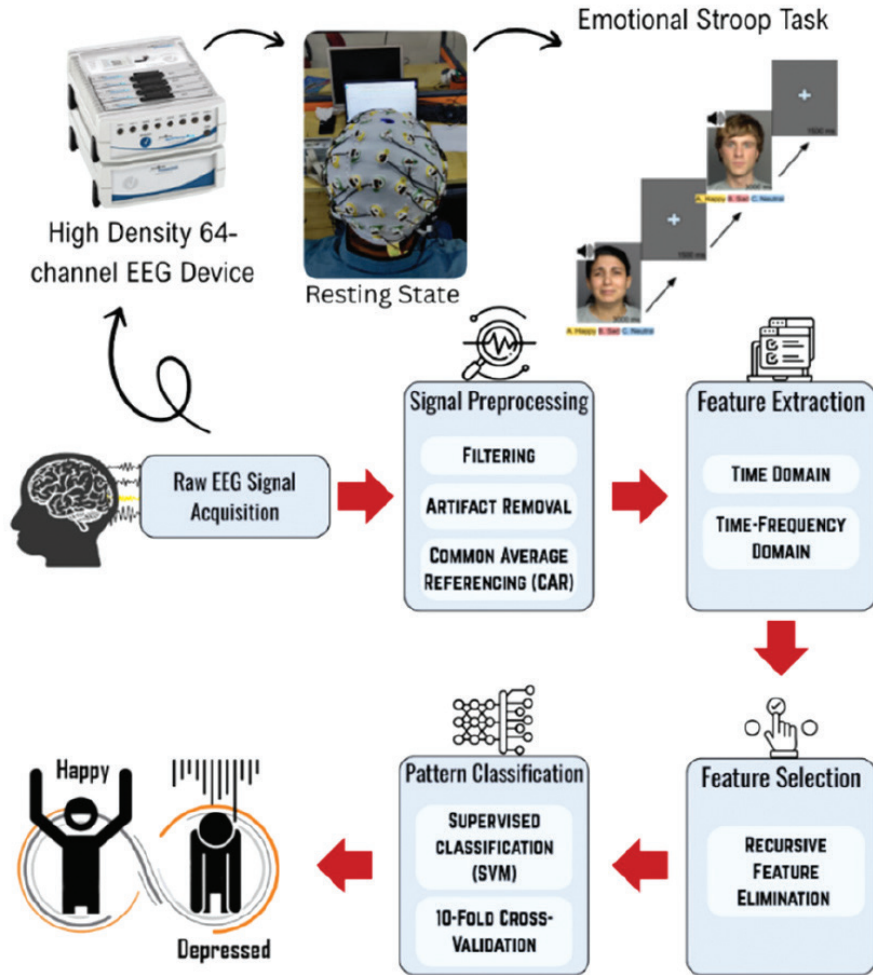
not an autonomous diagnostic tool.

Current model development draws on a combined dataset of locally collected EEG recordings (11 participants: 4 with MDD and 7 healthy controls) and a publicly available dataset of 64 participants (34 with MDD, 30 healthy controls) to support early training and stability assessment; public data supplement limited local samples, while local data are collected under ethics approval, and evaluation emphasizes internal consistency rather than cross-population generalizability.

Implementation followed a phased approach aligned with clinical operations (Figure 1). Phase 1 converted new-campus workflows to structured electronic capture. Phase 2 extended structured capture to the original campus with a temporary dual-system period until completeness and reliability were verified. Phase 3 connected treatment platforms not natively integrated with the treatment planning system through available connectivity modules.

Applies human-in-loop workflow with Psychiatrist oversight at every stage

Figure 01: Research-grade EEG-AI pipeline





Planned Enhancements and Translational Extensions

Planned enhancements prioritize clinical validation readiness, robustness, and operational feasibility rather than increasing model complexity. Key translational activities include comparative evaluation of conventional versus minimal preprocessing pipelines to establish a standardized, reproducible workflow suitable for multi-site validation; expansion of the local dataset through collection of EEG recordings from 50 clinically assessed participants (25 healthy controls and 25 with MDD) under harmonized hospital protocols to enable region-specific validation and bias assessment; and application of explainable AI methods to identify consistently informative channels and feature contributions, supporting interpretability and evidence-based simplification of acquisition protocols. Additional exploratory whole-brain feature mapping will be conducted to better characterize distributed neurophysiological patterns in depressive cohorts for interpretive and validation purposes only. Consideration of alternative model architectures will follow only after robustness and external validation milestones are met, with any progression toward broader clinical use subject to multi-site validation, regulatory review, and continued human-in-the-loop oversight. Together, these steps define a validation-first, governance-aligned pathway from research prototype to supervised clinical deployment.

Deployment & Implementation

The deployment is institution-led and government-aligned, operating as an ethics-approved multi-institutional pilot in Dhaka with AI processing at AIMS Lab and clinical recruitment through partner hospitals and psychiatric services. It follows a clinician-led, human-in-the-loop model, with supervised clinical validation planned in tertiary mental health centers. Broader rollout is contingent on validation, regulatory and institutional approval, and excludes direct-to-consumer, autonomous, or unsupervised use.

Scale of Deployment

The system is operating as an ethics-approved controlled pilot in Bangladesh, with centralized AI processing and hospital-based clinical recruitment supported by a multidisciplinary team. It follows a validation-first approach, with planned usability refinements before limited supervised clinical validation in tertiary mental health centers. Future scale-up will depend on validation results, regulatory and institutional approvals, and practical readiness factors such as hardware, workforce, workflow fit, and governance capacity.

Identified Users

Primary users are mental health professionals and clinical/AI researchers, who use the system for decision support, analysis, and validation, while retaining full clinical authority. Secondary users include trained EEG technicians for data acquisition and quality control, and health system planners who may access only aggregated, de-identified outputs for planning purposes.



Impact

Nature of Impact

The primary impact of this work is the successful implementation of a functioning EEG–AI decision-support pipeline under real-world constraints typical of low- and middle-income countries. The project delivered a research-grade prototype integrating EEG acquisition, secure data handling, machine learning–based analysis, and clinician-facing outputs within a government-approved, multi-institutional setting in Bangladesh.

Ethics & Governance

Development and evaluation of the EEG-based AI system follow established responsible AI principles for healthcare, including patient safety, fairness, transparency, accountability, and human oversight.

Discussion

Key achievements include establishing a functioning EEG–AI decision-support prototype under real-world LMIC constraints, supported by institutional infrastructure and clinician collaboration. Early governance and clinical alignment supported translational readiness.

Challenges and Constraints

Constraints include resource demands of high-density EEG, recruitment challenges linked to stigma and trust, limited dataset size, and the need for strong interpretability and workflow alignment. These are treated as design parameters rather than exceptions.

Lessons Learned

Three cross-cutting lessons emerge: governance, trust, and workflow alignment are as critical as model performance; locally developed systems are better positioned to reflect real operational conditions and constraints; and staged validation with conservative claims helps build institutional trust while supporting safety and responsible adoption.

Conclusion

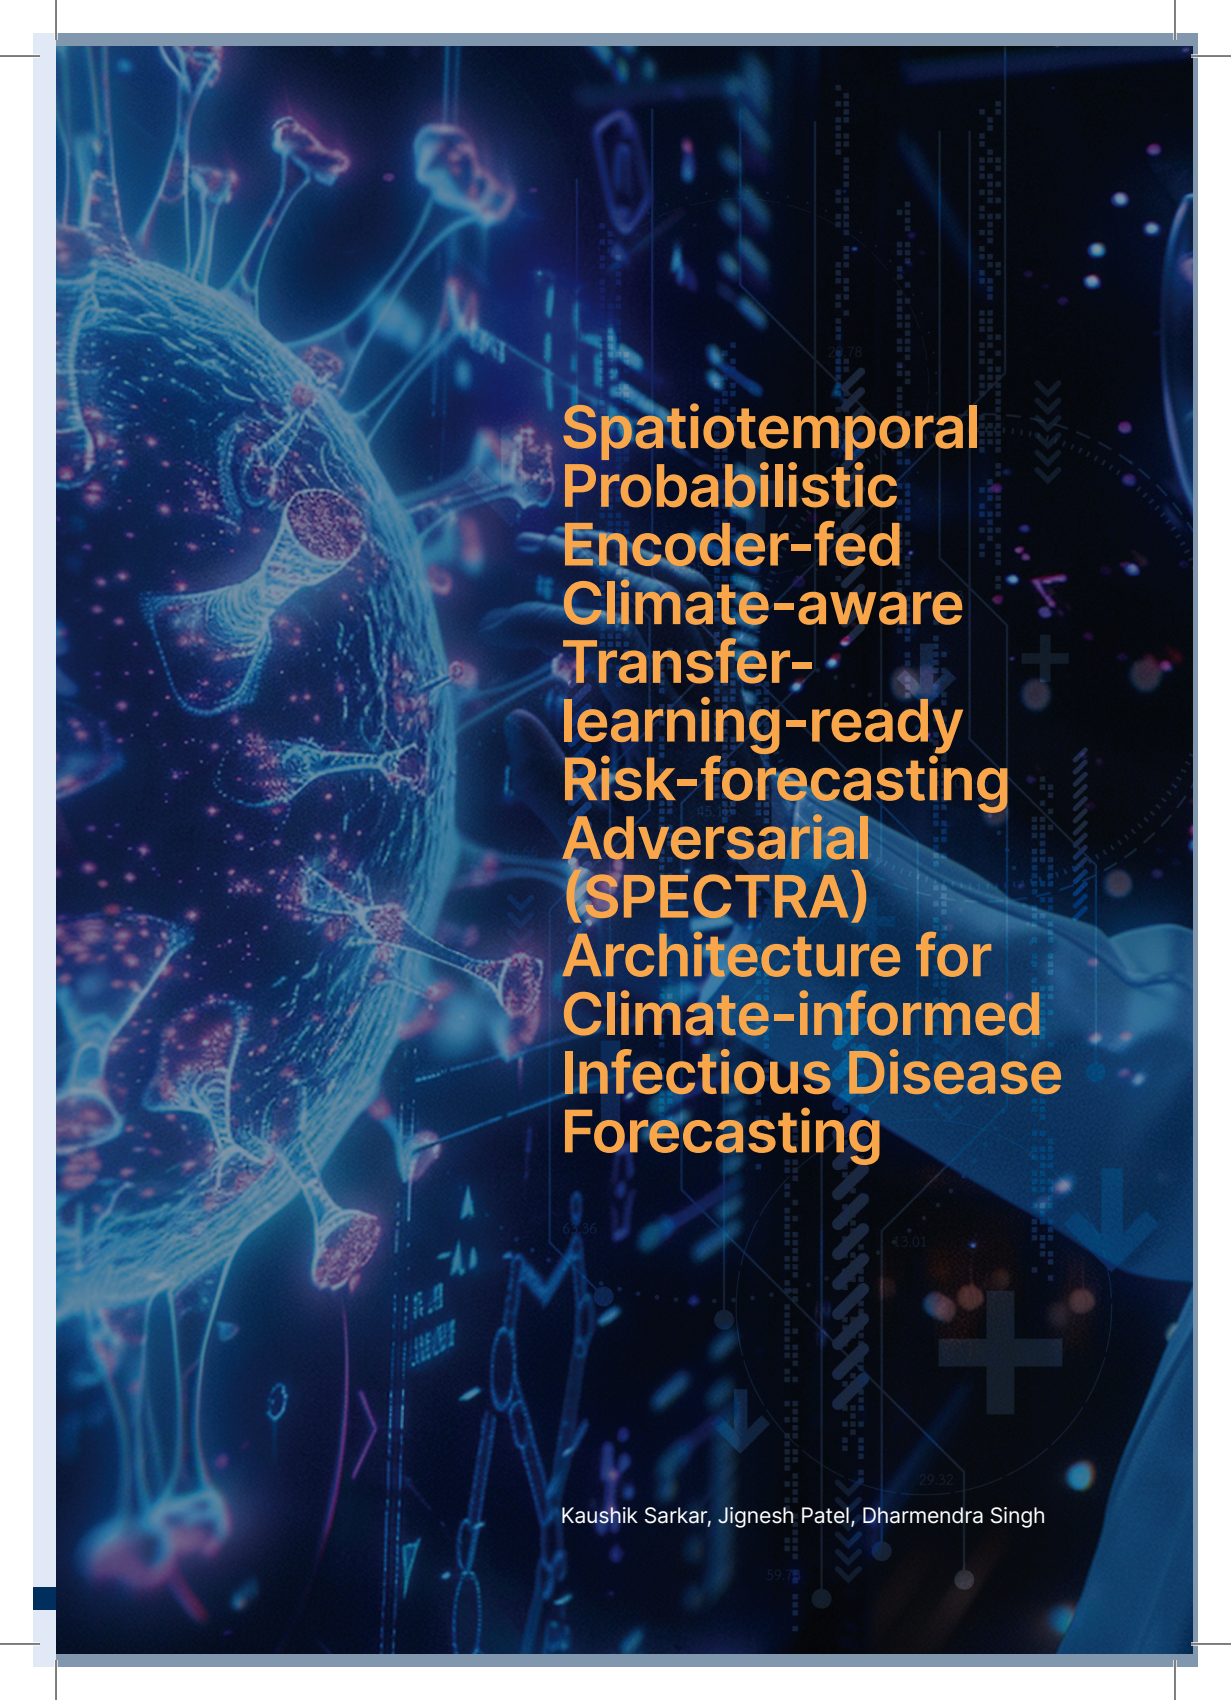
This chapter describes an EEG-based AI system for depression as a research-to-translation decision-support pilot developed under technical, ethical, and governance safeguards in a low-resource setting. The system is positioned as clinician-support infrastructure rather than a diagnostic replacement.

Its contribution lies in establishing validation-ready operational foundations including workflow design, oversight mechanisms, and interpretability that enable supervised clinical evaluation. The experience indicates that trustworthy AI in mental health depends on staged validation, transparency, and multidisciplinary collaboration, particularly in Global South contexts.

References


- Abdulnour, R.-E. E. (2025). Clinical supervision of artificial intelligence in health care. *New England Journal of Medicine*, 393(8), 786–797. <https://doi.org/10.1056/NEJMr2503232>
- Ahmed, T., Luthfa, S. T., Chaahat, A. H., Kanchon, M. R., & Sharif, A. B. (2025). Determinants of early mental health help-seeking among women in Bangladesh: A nationally representative bootstrapped regression analysis. *PLOS Mental Health*, 2(9), e0000420. <https://doi.org/10.1371/journal.pmen.0000420>
- Al-Haj, M., et al. (2025). Systematic review of EEG for depression detection. *NeuroImage: Reports*, 5(1), 100234. <https://doi.org/10.1016/j.nicl.2024.100234>
- Alloh, F. T., Regmi, P., Onche, I., van Teijlingen, E., & Trenoweth, S. (2021). Mental health in low- and middle-income countries (LMICs): Going beyond the need for funding. *Health Psychology and Behavioral Medicine*, 9(1), 316–339. <https://doi.org/10.1080/21642850.2021.1918341>
- Chandra, S. R., et al. (2017). Role of electroencephalography in the diagnosis and treatment of neuropsychiatric border zone syndromes. *Journal of Clinical and Diagnostic Research*, 11(6), OC08–OC11. <https://doi.org/10.7860/JCDR/2017/25514.9961>
- Cleveland Clinic. (2021, April 20). *Invasive EEG monitoring*. <https://my.clevelandclinic.org/health/diagnostics/17144-invasive-eeeg-monitoring>
- Cruz-Gonzalez, P., et al. (2025). Artificial intelligence in mental health care: A systematic review of diagnosis, monitoring, and intervention applications. *Current Psychology*. <https://doi.org/10.1007/s12144-024-06834-5>
- Czeisler, M. É., et al. (2022). Impact of the COVID-19 pandemic on mental health in the Americas. *The Lancet Regional Health – Americas*, 13, 100318. <https://doi.org/10.1016/j.lana.2022.100318>
- Dawson, A. (1956). Case history. *McGill Medical Journal*, 25(1), 45.
- De Aguiar Neto, F. S., & Rosa, J. L. (2019). Depression biomarkers using non-invasive EEG: A review. *Neuroscience & Biobehavioral Reviews*, 105, 83–93. <https://doi.org/10.1016/j.neubiorev.2019.07.021>
- Desai, T., Ritchie, F., & Welpton, R. (2016). Five Safes: Designing data access for research. *Journal of Privacy and Confidentiality*, 7(3), 1–23. <https://doi.org/10.29012/jpc.v7i3.405>
- Dev, A., Roy, N., Islam, M. K., Biswas, C., Ahmed, H. U., Amin, M. A., et al. (2022). Exploration of EEG-based depression biomarkers identification techniques and their applications: A systematic review. *IEEE Access*, 10, 17456–17482. <https://doi.org/10.1109/ACCESS.2022.3149544>
- Eack, S. M., Greeno, C. G., & Lee, B. J. (2006).

- Limitations of the Patient Health Questionnaire in identifying anxiety and depression: Many cases are undetected. *Research on Social Work Practice*, 16(6), 625–631. <https://doi.org/10.1177/1049731506290150>
- Fekadu, A., Misganaw, E., & Tefera, S. (2022). Under detection of depression in primary care settings in low- and middle-income countries: A systematic review and meta-analysis. *Systematic Reviews*, 11(1), 18. <https://doi.org/10.1186/s13643-022-01890-w>
- Fogel, A., Nazir, S., Hirapara, K., et al. (2025). Cultural assessment and treatment of psychiatric patients. In *StatPearls*. StatPearls Publishing. <https://www.ncbi.nlm.nih.gov/books/NBK482311/>
- Goyal, N., et al. (2025). Addressing mental health research capacity in low- and middle-income countries. *Asian Journal of Psychiatry*, 101, 104278. <https://doi.org/10.1016/j.ajp.2025.104278>
- Haque, M. A., et al. (2022). EEG-based depression detection: A systematic review. *Internet of Things and Cyber-Physical Systems*, 2, 172–187. <https://doi.org/10.1016/j.iotcps.2022.09.001>
- Hollenbenders, Y., Maier, C., & Reichenbach, A. (2024). Reproducibility of electroencephalography alpha band biomarkers for diagnosis of major depressive disorder [Preprint]. *medRxiv*. <https://doi.org/10.1101/2024.11.11.24317109>
- Hossain, M. M., Munaya, R., Ahmed, S., Tasnim, S., Edwards, J., Deeba, F., & Ma, P. (2020). Depression and suicidal behavior in South Asia: A systematic review and meta-analysis. *Global Mental Health*, 7, e31. <https://doi.org/10.1017/gmh.2020.25>
- Huque, R., Azad, A. K., Islam, K., Ahmed, H. U., & Amin, M. R. (2024). Mental health care seeking behavior in Bangladesh: Determinants and treatment gaps. *PLOS ONE*, 19(11), e0312521. <https://doi.org/10.1371/journal.pone.0312521>
- Indian Council of Medical Research. (2023). *Ethical guidelines for application of artificial intelligence in biomedical research and healthcare*. Department of Health Research, Government of India.
- Kaiser, B. N., Ticao, C., Anoj, C., Minto, J., Boglosa, J., & Kohrt, B. A. (2019). Adapting culturally appropriate mental health screening tools for use among conflict-affected and other vulnerable adolescents in Nigeria. *Global Mental Health*, 6, e10. <https://doi.org/10.1017/gmh.2019.8>
- Limenih, G., MacDougall, A., Wedlake, M., & Nouvet, E. (2023). Depression and global mental health in the Global South: A critical analysis of policy and discourse. *International Journal of Mental Health Systems*, 17(1), 40. <https://doi.org/10.1186/s13033-023-00610-8>
- Liu, P., Zhang, J., Chen, S., & Chen, S. (2024). Human-AI teaming in healthcare: 1 + 1 > 2? *Nature Medicine*. <https://doi.org/10.1038/s41591-024-03108-4>
- Liu, X., Zhang, H., Cui, Y., Zhao, T., Wang, B., Xie, X., & Zhang, L. (2024). EEG-based major depressive disorder recognition by neural oscillation and asymmetry. *Frontiers in Neuroscience*, 18, 1362111. <https://doi.org/10.3389/fnins.2024.1362111>
- Luo, Y., Tang, M., & Fan, X. (2025). Meta-analysis of resting frontal alpha asymmetry as a biomarker of depression. *npj Mental Health Research*, 4(1), 2. <https://doi.org/10.1038/s44184-025-00117-x>
- National Health Service England. (2023). *Secure data environment policy framework*. <https://digital.nhs.uk/services/secure-data-environment-service>
- Pasek, M., et al. (2024). Electroencephalographic markers in major depressive disorder: Insights from absolute, relative power, and asymmetry analyses. *Frontiers in Psychiatry*, 15, 1480228. <https://doi.org/10.3389/fpsy.2024.1480228>
- Rathod, S., Pinninti, N., Irfan, M., Gorczynski, P., Rathod, P., Gega, L., & Naeem, F. (2017). Mental health service provision in low- and middle-income countries. *Health Services Insights*, 10. <https://doi.org/10.1177/1178632917694350>
- Schmaal, L., et al. (2013). Frontal alpha asymmetry and depressive symptoms in a large cohort of patients with major depressive disorder. *Biological Psychology*, 94(1), 126–132. <https://doi.org/10.1016/j.biopsycho.2013.05.011>
- Simmatis, L., Russo, E. E., Geraci, J., Harmsen, I. E., & Samuel, N. (2023). Technical and clinical considerations for electroencephalography-based biomarkers for major depressive disorder. *npj Mental Health Research*. <https://doi.org/10.1038/s44184-023-00038-7>
- Su, C., et al. (2024). Artificial intelligence for the diagnosis and management of depression: A systematic review. *Journal of Affective Disorders*. <https://pubmed.ncbi.nlm.nih.gov/39464570/>
- Weiner, E. B., Dankwa-Mullan, I., Nelson, W. A., & Hassanpour, S. (2025). Ethical challenges and evolving strategies in the integration of artificial intelligence into clinical practice. *PLOS Digital Health*, 4(4), e0000810. <https://doi.org/10.1371/journal.pdig.0000810>
- World Health Organization. (2021). *Ethics and governance of artificial intelligence for health*. World Health Organization. <https://www.who.int/publications/i/item/97892440029200>



**Spatiotemporal
Probabilistic
Encoder-fed
Climate-aware
Transfer-
learning-ready
Risk-forecasting
Adversarial
(SPECTRA)
Architecture for
Climate-informed
Infectious Disease
Forecasting**

Kaushik Sarkar, Jignesh Patel, Dharmendra Singh



Spatiotemporal Probabilistic Encoder-fed Climate-aware Transfer-learning-ready Risk-forecasting Adversarial (SPECTRA) Architecture for Climate-informed Infectious Disease Forecasting

¹*Kaushik Sarkar, ²Jignesh Patel, ³Dharmendra Singh

Abstract

Climate-sensitive infectious diseases such as malaria, dengue, acute diarrhoeal diseases, and leptospirosis continue to contribute a substantial public health burden across many low- and middle-income countries. Variability in temperature, rainfall, humidity, and extreme weather patterns is associated with changing transmission dynamics, complicating traditional season-based planning approaches. While machine learning models are increasingly used for outbreak forecasting, many implementations remain disease-specific and country-specific, limiting reuse and scale.

This use case describes the design and implementation of a flexible artificial intelligence architecture—the Spatiotemporal Probabilistic Encoder-fed Climate-aware Transfer-learning-ready Risk-forecasting Adversarial (SPECTRA) system developed to support climate-informed infectious disease forecasting across multiple diseases and geographies. The architecture was developed using dengue surveillance data from 8 countries and 24 global meteorological datasets, with model components designed to capture climate signals and disease dynamics while quantifying uncertainty.

Adapted versions of SPECTRA have been incorporated into national multi-disease early warning and response platforms in Indonesia and Bangladesh through transfer learning and local retraining. Validation studies indicate improved forecasting accuracy relative to selected baseline models under study conditions. Operational deployments demonstrate feasibility and configurability across dashboard and scenario-modelling environments. Observed benefits relate primarily to preparedness support and earlier risk signalling rather than confirmed outcome impact. This use case illustrates how transferable, climate-aware AI architectures can support the development of integrated early warning decision-support systems when implemented with governance, local ownership, and human oversight.

*1 Institute for Health Modeling and Climate Solutions, Malaria No More, 1301 Connecticut Ave, NW, Suite 502, Washington, DC, 20036, US. ; 2Indian Institute of Information Technology Vadodara, Sector-28, Gandhinagar, 382028, Gujarat, India. 3Indian Institute of Information Technology Vadodara, Sector-28, Gandhinagar, 382028, Gujarat, India.



Background

Recurrent outbreaks of climate-sensitive infectious diseases continue to place pressure on public health systems, particularly in resource-constrained settings. A large proportion of epidemic-prone diseases are influenced by climatic and environmental factors. Changes in temperature, rainfall, and extreme weather events are associated with altered vector ecology, pathogen survival, and human exposure patterns. As climate variability increases, fixed seasonal calendars and historically derived thresholds may become less reliable for preparedness planning.

Early recognition of elevated outbreak risk supports more timely and targeted public health action, including vector control, community engagement, stockpiling, and service readiness. Surveillance and early warning systems play a central role in this process. There is increasing interest in strengthening these systems with climate-informed forecasting capabilities that can complement routine epidemiological monitoring.

However, climate–disease relationships are complex. Disease incidence is rarely driven by a single meteorological variable, and similar climatic conditions can produce different epidemiological outcomes across locations due to ecological, social, and programmatic differences. Climate drivers also operate across multiple time scales, from short-term weather fluctuations to

multi-seasonal and long-term trends. Conventional statistical and machine learning models that rely on fixed lag structures or static driver sets may be difficult to maintain when underlying relationships shift.

Operational surveillance data introduce additional complexity, including reporting delays, denominator changes, catchment shifts, and variable completeness. Many forecasting tools are built as disease-specific pilots with bespoke assumptions and thresholds, leading to fragmented toolsets that are difficult to integrate into unified national decision-support environments.

The SPECTRA architecture was developed to address these technical and operational challenges by providing a transferable, climate-aware forecasting backbone that can be adapted across diseases and settings while explicitly representing uncertainty and supporting local retraining.

Early warning systems empower targeted outbreak interventions



AI Use Case Description

Solution Description

SPECTRA is a general-purpose model for forecasting climate-sensitive infectious diseases, designed to support decision-makers in estimating disease magnitude while explicitly quantifying predictive uncertainty under varying climate conditions and historical disease trends. The model learns the forward mapping between the climate manifold and disease signals as a primary objective, enabling simultaneous prediction of disease outcomes and their uncertainty distributions.

SPECTRA consists of two core components: a climate encoder and a disease forecaster. The climate encoder is implemented as a conditional β -variational autoencoder (c β -VAE), which captures the latent structure of climate forcing from multiple meteorological variables. The disease forecaster is built as a hybrid Generative Adversarial Network (GAN) to address multiscale, non-Markovian, and non-stationary disease incidence dynamics. Its generator integrates a Causal Temporal Convolutional Network (TCN), a Long Short-Term Memory (LSTM) network, and a Multi-Head Self-Attention (MHSA) layer, while the critic employs a similarly configured LSTM followed by a fully connected layer.

Deployment and Implementation

The initial SPECTRA model was developed using dengue surveillance data from eight countries—Bangladesh, Brazil, Colombia, Indonesia, Malaysia, Nicaragua, the Philippines, and Singapore—sourced from the OpenDengue

repository (Clarke et al., 2023), together with 26 meteorological variables derived from the European Centre for Medium-Range Weather Forecasts (ECMWF) ERA5 reanalysis dataset, accessed via the Copernicus Data Store API (Hersbach et al., 2023). Model development was guided by comprehensive exploratory time-series diagnostics to characterize cross-country signal patterns and variability, complemented by a structured ablation framework and a suite of custom explainable AI methods to support interpretability. In benchmarking evaluations, SPECTRA showed improved predictive accuracy and probabilistic calibration compared with commonly used supervised learning approaches, including Linear and Ridge Regression, Random Forest, XGBoost, LightGBM, GRU, and LSTM models, under the test conditions evaluated.

SPECTRA is designed as a data-intensive forecasting system. Production-level deployment typically requires harmonized climate and health surveillance data spanning at least 36 months, with longer series improving robustness where epidemiological signals are noisy or subject to regime shifts. Model performance is dependent on data completeness and stability. Updating can be performed through scheduled batch retraining aligned with routine monthly data updates or through event-triggered retraining based on predefined data volume or performance-drift thresholds. Initial model training for a single disease generally requires approximately 6–12 A100 GPU hours, with computational demand increasing proportionally with subnational spatial granularity and meteorological feature volume, subject to memory and I/O constraints.

The SPECTRA framework has been productionized by the Institute for Health Modeling and Climate Solutions (IMACS) as a



backend forecasting and early warning layer within national decision-support systems. Through technical assistance programmes, implementations have been supported in Indonesia and Bangladesh under climate-informed disease early warning initiatives. In Indonesia, multi-disease early warning dashboards and scenario-based digital twin systems have been developed for dengue, malaria, and leptospirosis under the ClimateSmart Indonesia programme, with capacity-building and AI innovation support from multiple funding partners. Technical modules and implementation support were delivered to the in-country partner KORIKA and Ministry of Health technical teams. In Bangladesh, SPECTRA components were integrated into a climate-informed early warning surveillance platform developed with the Ministry of Health, the Institute of Epidemiology, Disease Control and Research (IEDCR), and partner organizations under a multi-country response mechanism programme.

Scale of deployment

SPECTRA serves as the analytical engine for the ClimateSmart Indonesia Early Warning System and the Climate-informed Early Warning Surveillance System in Bangladesh, both implemented at national scale. In Indonesia, the platform supports provincial-level forecasting, with targeted configurations for malaria endemic (Central Papua), malaria elimination (Lampung), and high-burden district contexts such as Mimika, as well as for dengue and leptospirosis in Central Java. In Bangladesh, the system integrates national surveillance

data for malaria, dengue, and acute watery diarrhoea reported by the Ministry of Health and the Institute of Epidemiology, Disease Control and Research (IEDCR).

Identified Users


Primary users of SPECTRA-based systems are institutional and programme-level actors rather than individual clinicians. These include national and subnational public health institutes, disease control programmes, early warning units, and meteorological-health coordination teams.

In Indonesia, in-country technology partners configured and maintained the system for use by Ministry of Health programmes in coordination with meteorological agencies. In Bangladesh, the national epidemiology institute and ministry disease programmes are the principal operational users.

Outputs are designed to support programme managers, surveillance analysts, and emergency operations centres in interpreting risk windows and planning preparedness actions. The system is not intended for direct public or patient-facing use.

Impact

Formal impact of SPECTRA-supported early warning systems on outbreak frequency, size, or health outcomes has not yet been established and would require retrospective and prospective evaluation over multiple seasons. Current evidence relates primarily to technical validation and operational utility.



SPECTRA generates calibrated disease forecasts that identify elevated-risk windows in advance of conventional early warning approaches that rely primarily on contemporaneous case counts. The model's uncertainty bounds enable risk-tiered escalation levels (for example, watch and alert thresholds), supporting more proportionate and consistent preparedness triggers within public health emergency operations workflows, including at IEDCR. This approach is intended to help reduce unnecessary alarms while supporting anticipatory action.

In Bangladesh, SPECTRA's downstream transfer-learning components were developed to address a persistent challenge in forecasting dengue surges that had proven difficult to model using conventional EWS methods, with initial results indicating improved predictive performance while preserving data sovereignty by avoiding cross-border movement of surveillance data. The Bangladesh climate-informed EWS implementation was presented at the 2025 ATACH meeting in Belém convened by WHO.

Ethics & Governance

IMACS transferred SPECTRA to Bangladesh and Indonesia as modular Python components to enable an open, extensible architecture for early warning system models while minimizing the risk of cross-border data movement. Deployment was led by in-country institutions, with IMACS providing institutional and technical enablement support.

In both countries, IMACS first configured data pipelines using mock datasets structured to mirror national health information system outputs. Model training was then conducted locally by national implementers on in-country servers using disease surveillance data obtained directly from the respective Ministries of Health. Data-sharing arrangements were established through partnerships between local organizations and ministries under non-disclosure agreements.

Climate data integration was addressed as a national coordination and capability-building task. Initial pilot use cases were developed for selected geographies and diseases using ERA5 datasets, which were subsequently replaced with official meteorological data from national weather agencies.

In Indonesia, SPECTRA implementation followed the pre-established Indonesia Blueprint for Integrating Climate and Health Data, jointly published by IMACS and KORIKA and formally endorsed by the Ministry of Health, the Badan Meteorologi, Klimatologi, dan Geofisika, and the Kementerian Komunikasi dan Digital. In Bangladesh, system oversight is currently provided through a technical working group, with a formal coordination and governance framework under development.



Discussion

Successes

SPECTRA illustrates a modular, transferable approach to climate-informed infectious disease forecasting that differs from single-model, single-disease tools. By separating climate representation and disease forecasting components, the architecture supports reuse and adaptation across diseases and geographies through transfer learning and local retraining.

Validation studies indicate that multi-component architectures can capture multi-scale temporal dependencies more effectively than some commonly used baseline models under test conditions. Integrated uncertainty estimation and explainability features support more cautious and interpretable use of forecasts in operational settings.

An additional contribution is the structured transfer-learning protocol, which supports adaptation of pretrained models to local epidemiological contexts. This approach may be useful for settings with shorter local time series, although minimum data requirements remain substantial.

Challenges and lessons learned

Implementation highlighted persistent coordination challenges in cross-ministry data sharing, particularly between health and meteorological sectors. Establishing formal data pipelines and governance

agreements required sustained institutional engagement. Capacity constraints in advanced AI and data engineering required repeated training, workshops, and technical accompaniment. While in-country teams have increasingly assumed operational responsibility, continued skills development remains important for sustainability. Disease-specific calibration and definition of actionable thresholds remain necessary even with transferable architectures. Consensus on alert definitions and validation protocols is essential to avoid over- or under-response. Sustainability depends on continued funding, institutional ownership, and routine use of forecast outputs in planning processes.

Conclusion

SPECTRA represents a transferable, climate-aware AI architecture designed to support infectious disease forecasting and early warning decision support. Its design emphasises latent climate representation, multi-scale temporal modelling, uncertainty estimation, and transfer learning to support adaptation across diseases and countries. More broadly, this use case suggests that modular, uncertainty-aware AI forecasting backends can complement surveillance and preparedness systems when implemented with strong governance, local capacity-building, and clear positioning as decision-support tools rather than autonomous predictors.

References

- Alahmari, A. A., Almuzaini, Y., Alamri, F., Alenzi, R., & Khan, A. A. (2024). Strengthening global health security through health early warning systems: A literature review and case study. *Journal of Infection and Public Health*, *17*, 85–95. <https://doi.org/10.1016/j.jiph.2024.01.019>
- Asia Pacific Malaria Elimination Network (APMEN). (2025). *APMEN case study: Adapting strategies to mitigate the impact of climate and environmental changes on vector-borne diseases in the Asia-Pacific region*. https://apmen.org/sites/default/files/all_resources/APMEN%20Case%20Study_Adapting%20Strategies_Climate%20%26%20Environmental%20Changes_%20VBDs%20in%20Asia-Pacific%20Region.pdf
- Banasiak, J., & Tchoumi, S. Y. (2024). Multiscale malaria models and their uniform in-time asymptotic analysis. *Mathematics and Computers in Simulation*, *221*, 1–18. <https://doi.org/10.1016/j.matcom.2024.02.015>
- Bangladesh EWARS: Early Warning Alert and Response System. (2025). Institute of Epidemiology, Disease Control and Research. <https://ewars.iedcr.gov.bd/>
- Beesley, L. J., Osthus, D., & Del Valle, S. Y. (2022). Addressing delayed case reporting in infectious disease forecast modeling. *PLoS Computational Biology*, *18*(6), e1010115. <https://doi.org/10.1371/journal.pcbi.1010115>
- Brook, T. (2017). Differential effects of global and local climate data in assessing environmental drivers of epidemic outbreaks. *Proceedings of the National Academy of Sciences*, *114*(49), 12845–12847. <https://doi.org/10.1073/pnas.1717723114>
- Cepella, G. M. A., Borbor-Cordova, M. J., van Elteren, M., & Petrova, D. (2025). Assessing the local context for implementing a climate-based early warning system for dengue fever outbreaks in Ecuador. *Climate Services*, *38*, 100571. <https://doi.org/10.1016/j.cliser.2025.100571>
- ClimateSmart Indonesia. (2025). *Dashboard monitoring penyakit berdasarkan perubahan iklim*. <https://climatesmart.id/>
- Ekundayo, F. (2024). Using machine learning to predict disease outbreaks and enhance public health surveillance. *World Journal of Advanced Research and Reviews*, *24*(3), 794–811. <https://doi.org/10.30574/wjarr.2024.24.3.3732>
- Faldy Anggita, F., Putra, Y. C., Afira, N., Hullyiyatus Suadaa, L., Sungkono, D. A., Riza, H., & Sarkar, K. (2025). ClimateSmart Indonesia: An integrated AI-driven platform for climate-health surveillance and early warning systems. In *Proceedings of COSITE 2025*.
- Grassly, N. C., & Fraser, C. (2006). Seasonal infectious disease epidemiology. *Proceedings of the Royal Society B: Biological Sciences*, *273*(1600), 2541–2550. <https://doi.org/10.1098/rspb.2006.3604>
- Han, L., Lin, Z., Yin, Q., Tang, M., Guan, S., & Boguñá, M. (2023). Non-Markovian epidemic spreading on temporal networks. *Chaos, Solitons & Fractals*, *173*, 113664. <https://doi.org/10.1016/j.chaos.2023.113664>
- Hersbach, H., Bell, B., Berrisford, P., Biavati, G., Horányi, A., Muñoz Sabater, J., Nicolas, J., Peubey, C., Radu, R., Rozum, I., Schepers, D., Simmons, A., Soci, C., Dee, D., & Thépaut, J.-N. (2023). *ERA5 monthly averaged data on single levels from 1940 to present*. Copernicus Climate Change Service Climate Data Store. <https://doi.org/10.24381/cds.f17050d7>
- Hu, W.-H., Sun, H.-M., Wei, Y.-Y., & Hao, Y.-T. (2025). Global infectious disease early warning models: An updated review and lessons from the COVID-19 pandemic. *Infectious Disease Modelling*, *10*(2), 410–422. <https://doi.org/10.1016/j.idm.2024.12.001>
- Clarke, J., Lim, A., Gupte, P., Pigott, D. M., van Panhuis, W. G., & Brady, O. J. (2023). *OpenDengue: Data from the OpenDengue database (Version 1.2)* [Dataset]. Figshare. <https://doi.org/10.6084/m9.figshare.24259573>
- Lawrence, A. (2024). Evaluating the effectiveness of public health measures during infectious disease outbreaks: A systematic review. *Cureus*. <https://doi.org/10.7759/cureus.55893>

- Liao, H., Lyon, C. J., Ying, B., & Hu, T. (2024). Climate change, its impact on emerging infectious diseases and new technologies to combat the challenge. *Emerging Microbes & Infections*, 13(1), 2356143. <https://doi.org/10.1080/22221751.2024.2356143>
- Mora, C., McKenzie, T., Gaw, I. M., Dean, J. M., von Hammerstein, H., Knudson, T. A., Setter, R. O., Smith, C. Z., Webster, K. M., Patz, J. A., & Franklin, E. C. (2022). Over half of known human pathogenic diseases can be aggravated by climate change. *Nature Climate Change*, 12(9), 869–875. <https://doi.org/10.1038/s41558-022-01426-1>
- Murillo, L. N., Murillo, M. S., & Perelson, A. S. (2013). Towards multiscale modeling of influenza infection. *Journal of Theoretical Biology*, 332, 267–290. <https://doi.org/10.1016/j.jtbi.2013.03.024>
- Nyawanda, B. O., Ochomo, E., Otieno, J. D., Keitany, K., Machini, B. K., & Vounatsou, P. (2025). Climate, interventions, and malaria outcomes in a warming world: Towards climate-smart malaria control in Kenya. *Tropical Medicine and Infectious Disease*, 10(12), 335. <https://doi.org/10.3390/tropicalmed10120335>
- Qasrawi, R., Issa, G., Thwib, S., AbuGhoush, R., Amro, M., Ayyad, R., Vicuna, S., Badran, E., Khader, Y., Al Qutob, R., Al Bakri, F., Trigui, H., Sokhn, E., Musa, E., & Kong, J. D. (2025). The role of machine learning in infectious disease early detection and prediction in the MENA region: A systematic review. *Informatics in Medicine Unlocked*, 56, 101651. <https://doi.org/10.1016/j.imu.2025.101651>
- Redding, D. W., Gibb, R., Dan-Nwafor, C. C., Illori, E. A., Yashe, R. U., Oladele, S. H., Amedu, M. O., Iniobong, A., Attfield, L. A., Donnelly, C. A., Abubakar, I., Jones, K. E., & Ihekweazu, C. (2021). Geographical drivers and climate-linked dynamics of Lassa fever in Nigeria. *Nature Communications*, 12, 5759. <https://doi.org/10.1038/s41467-021-25910-y>
- Rodó, X., Pascual, M., Doblas-Reyes, F. J., Gershunov, A., Stone, D. A., Giorgi, F., Hudson, P. J., Kinter, J., Rodríguez-Arias, M.-Á., Stenseth, N. C., Alonso, D., García-Serrano, J., & Dobson, A. P. (2013). Climate change and infectious diseases: Can we meet the needs for better prediction? *Climatic Change*, 118(3–4), 625–640. <https://doi.org/10.1007/s10584-013-0744-1>
- Sahai, A. K., Mandal, R., Joseph, S., Saha, S., Awate, P., Dutta, S., Dey, A., Chattopadhyay, R., Phani, R., Pattanaik, D. R., & Deshpande, S. (2020). Development of a probabilistic early health warning system based on meteorological parameters. *Scientific Reports*, 10, 14741. <https://doi.org/10.1038/s41598-020-71668-6>
- Santangelo, O. E., Gentile, V., Pizzo, S., Giordano, D., & Cedrone, F. (2023). Machine learning and prediction of infectious diseases: A systematic review. *Machine Learning and Knowledge Extraction*, 5(1), 175–198. <https://doi.org/10.3390/make5010013>
- Tissera, H. A., Jayamanne, B. D. W., Raut, R., Janaki, S. M. D., Tozan, Y., Samaraweera, P. C., Liyanage, P., Ghouse, A., Rodrigo, C., de Silva, A. M., & Fernando, S. D. (2020). Severe dengue epidemic, Sri Lanka, 2017. *Emerging Infectious Diseases*, 26(4), 682–691. <https://doi.org/10.3201/eid2604.190435>
- Towfek, S. K., & Elkanzi, M. (2024). A review on the role of machine learning in predicting the spread of infectious diseases. *Metaheuristic Optimization Review*, 2(1), 14–27. <https://doi.org/10.54216/MOR.020102>
- Tran, B.-L., Tseng, W.-C., Chen, C.-C., & Liao, S.-Y. (2020). Estimating the threshold effects of climate on dengue: A case study of Taiwan. *International Journal of Environmental Research and Public Health*, 17(4), 1392. <https://doi.org/10.3390/ijerph17041392>
- Vora, N. M., Hannah, L., Walzer, C., Vale, M. M., Lieberman, S., Emerson, A., Jennings, J., Alders, R., Bonds, M. H., Evans, J., Chilukuri, B., Cook, S., Sizer, N. C., & Epstein, J. H. (2023). Interventions to reduce risk for pathogen spillover and early disease spread to prevent outbreaks, epidemics, and pandemics. *Emerging Infectious Diseases*, 29(3), 1–9. <https://doi.org/10.3201/eid2903.221079>
- Wu, X., Lu, Y., Zhou, S., Chen, L., & Xu, B. (2016). Impact of climate change on human infectious diseases: Empirical evidence and human adaptation. *Environment International*, 86, 14–23. <https://doi.org/10.1016/j.envint.2015.09.007>

Data privacy essentials for data de-identification in radiology images, a practical approach in AI research



Manjunath K N, Abhishek Gupta, Ganesh P,
Muralikrishna S N, Sohan, Sudarshan Rawat,
Neelima Bayyapu, Ajitha Shenoy, Rajagopal KV,
Prakashini K



Data privacy essentials for data de-identification in radiology images, a practical approach in AI research

¹*Manjunath K N, ²Abhishek Gupta, ³Ganesh P, ⁴Muralikrishna S N, ⁵Sohan, ⁶Sudarshan Rawat, ⁷Neelima Bayyapu, ⁸Ajitha Shenoy, ⁹Rajagopal KV, ¹⁰Prakashini K*

Abstract

The use of medical imaging datasets for artificial intelligence (AI) research requires robust data protection and de-identification practices. This use case describes a structured approach to de-identifying Computed Tomography Coronary Angiography (CTCA) image datasets prior to their use in AI model development. A dataset of 700 CTCA volumes was prepared using a multi-step de-identification workflow aligned with recognized regulatory frameworks, including HIPAA Safe Harbor provisions, GDPR provisions on health data, and India's Digital Personal Data Protection (DPDP) Act.

De-identification was implemented through rule-based modification and removal of selected DICOM attributes while preserving the technical image integrity required for analysis. Custom Python-based tools were developed to support reproducible de-identification and report processing. A sample-based expert review by radiologists assessed residual re-identification risk using predefined criteria.

The use case illustrates an operational model for privacy-preserving dataset preparation that supports AI research while maintaining regulatory and ethical safeguards. Findings are indicative of feasibility in controlled research settings; broader validation across institutions and datasets remains necessary.

Background

This use case is multidisciplinary, integrating radiology, cardiology, artificial intelligence, and medical image analysis. Coronary artery disease (CAD) remains a leading cause of mortality globally, with India contributing a substantial share of the burden.


There is a growing need for AI-enabled automated systems to support diagnosis and prognosis. The research focuses on developing deep learning models for CAD detection using cardiac CT images in combination with radiology reports

*1,4, 5, 7, 8 School of Computer Engineering, Manipal Institute of Technology, 2, 3 Department of Cardiology, Kasturba Medical College 6 Department of Radiodiagnosis and Imaging, Manipal Hospitals Bengaluru, Bengaluru, 9, 10 Department of Radiology, Kasturba Medical College, Manipal

(Figure 1). This approach is designed to substantially reduce image interpretation time and support more efficient clinical prioritization and treatment decision-making.

Data privacy and dataset completeness were treated as core prerequisites for AI model development. Datasets were compiled from both secondary sources (standard public repositories) and primary clinical sources (local hospitals). Primary data collection received ethics approval from two South Indian hospitals a teaching hospital in Manipal (clinical site 1) and a corporate hospital in Bengaluru (clinical site 2). The study placed particular emphasis on structured dataset preparation and de-identification to ensure compliance with applicable data protection and research governance standards.

Figure 01: The application of AI in the cardiology image analysis workflow using Computed Tomography imaging modality



AI Use Case Description

This use case concerns AI research data preparation infrastructure rather than a clinical decision system. The AI-related component supports safe dataset preparation for downstream model development and evaluation. The approach emphasises responsible data handling, auditability, and reproducibility.

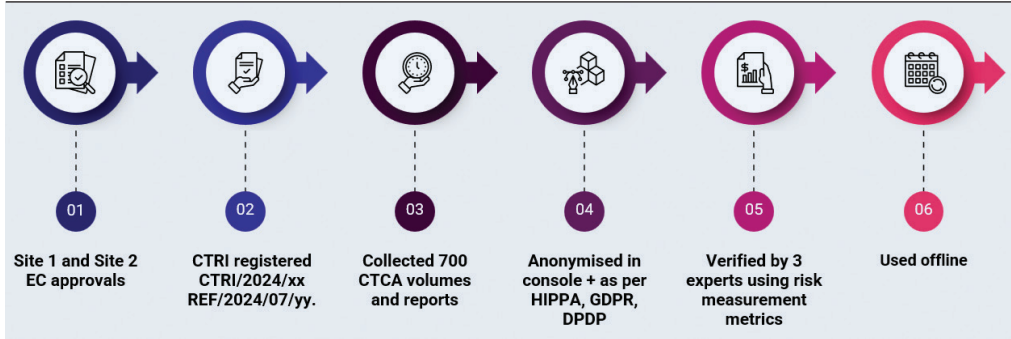
Solution Description

A structured data management and de-identification workflow was implemented for CTCA datasets and associated radiology reports. The workflow included

ethics approval, registry registration, dataset specification, layered de-identification, and expert validation.

The study proceeded through three structured preparatory stages. First, institutional ethics approvals were obtained from both participating sites (site 1 and site 2), authorizing one year of access to anonymized CT coronary angiography (CTCA) datasets and associated radiology reports. Second, the study was prospectively registered with the Clinical Trials Registry of India (CTRI) for the period 01 March 2024 to 31

Figure 02: The different steps in the healthcare dataset management in AI



March 2025. Third, sample size estimation and dataset assembly were completed, with a target of 700 diagnostic-quality CTCA volumes collected across the two centres (500 from site 1 and 200 from site 2). The dataset covered adults aged 18–88 years, with 43% female participants. Imaging acquisition parameters included multislice helical CT with contrast-enhanced 3D volume acquisition, slice thickness of 0.9–1.0 mm, matrix size 512 × 512, pixel spacing 0.27–0.67 mm, tube current 230–537 mA, tube voltage settings of 80/100/120 kVp, and 250–800 slices per study. All images were validated against DICOM PS3.1 (2022b) standards and stored in DICOM format, while corresponding radiology reports were maintained in PDF format (Figure 3). A two-level de-identification process was implemented: first, PACS-level anonymisation prior to download, and second, detailed tag-level de-identification aligned with HIPAA 45 CFR §164.514(b), GDPR Article 9, and India’s DPDP Act requirements. Sensitive

identifiers and institutional fields were removed or set to null for Type 2 and Type 3 DICOM attributes, while Type 1 attributes were replaced with default values; selected identifiers (patient ID, study ID, accession number, protocol name) were replaced with standardized pseudonymous codes to maintain internal linkage without exposing identity. Tag reviews were applied consistently across axial, curved MPR, and straightened MPR DICOM files. Re-identification risk was independently assessed by three radiologists on a random subset of 60 studies using replicability, data availability, and distinguishability criteria, with no identifiable patient information detected. The de-identified dataset was then used for ground-truth annotation of coronary arteries, model training, and evaluation.

Figure 03: DICOM data structure showing the values for each tag (DICOM element)




Table 01: Patient sensitive information from HIPPA, EDPR and DPDP act

Personal identifiers	Personal names, date of birth
Direct identifiers	Patient ID
Indirect identifiers	Height, Ethnicity, Hair color, Dates
Safe harbour methods	<ol style="list-style-type: none"> 1. Names 2. Geographic subdivisions smaller than a state 3. Admission date, discharge date, birthdate, death date, ages over 89 years 4. Telephone, cell phone, and fax numbers 5. Email addresses 6. IP addresses 7. Social Security numbers 8. Medical record numbers 9. Health plan beneficiary numbers or member ID numbers 10. Device identifiers and serial numbers 11. Certificate/license numbers 12. Credit card and bank account numbers 13. Vehicle identifiers and serial numbers including license plates 14. Website URLs 15. Full face photos and comparable images 16. Biometric identifiers (including finger and voice prints) 17. Any other unique identifying numbers, characteristics or codes.



Table 02: GDPR data attributes defined for cardiology domain

Health incidents

Psychological condition

Psychological counselling data

Prescriptions

Ergonomics data

Other health/medical data

Health reports

Identified illnesses

Drug tests

Health tests

Allergies Medical certificate

Rehabilitation data, what

rehabilitation actions are made at

the workplace, e.g. adjustments

of working hours, temporary

having other tasks, work aid etc.

This should not contain any medical data, Psycho-social related information Work accidents, injuries, etc

Through systematic review of DICOM tag types, identifying information was removed from axial CT (native slices) as well as curved MPR (CMPR) and straightened MPR (SMPR) reformatted image files. To maintain compliance with DICOM data definitions while ensuring privacy, values for Type 2 and Type 3 attributes were removed where present, and Type 1 attributes were replaced with default placeholder values (e.g.,


zero) where required. Because full de-identification would eliminate the ability to link records across files, each original patient identifier was replaced with a standardized pseudonymous code. This mapped identifier was consistently inserted as the patient ID, study ID, and accession number across the CT, CMPR, and SMPR DICOM files, preserving internal dataset linkage while preventing re-identification. (Table 3).

Table 03: List of DICOM tags whose values are changed to desired values

Tag	Name	Type	Value modification
(0018, 1030)	Protocol name	3	Changed from "KMC Coronary CTA" to "Coronary CTA"
(0010,0020)	Patient ID	2	Changed to CVD0001
(0008, 0050)	Accession number	2	Changed to CVD0001
(0020, 0010)	Study ID	2	Changed to CVD0001

Table 04: List of DICOM tags whose values are set to null and default values

Tag	Name	Type	Value
(0008, 0080)	Institution Name	3	"" if exists
(0008, 0081)	Institution address	3	"" if exists
(0008, 0090)	Referring Physician's name	3	"" if exists
(0008,1040)	Institutional Department Name	3	"" if exists
(0008, 1050)	Performing Physician name	2	"" if exists
(0008,1060)	Name of Physician(s) Reading Study	3	"" if exists
(0008,1070)	Operators' Name	3	"" if exists
(0010,0010)	Patient's Name	2	"" if exists
(0010,0030)	Birth date of the patient.	2	"" if exists
(0020,0013)	Instance Number	2	"" if exists
(0018, 1000)	Device serial number	1	0
(0008,0020)	Study Date	2	"" if exists
(0008,0030)	Study Time	2	"" if exists
(0008,0012)	Instance Creation Date	3	"" if exists
(0008,0013)	Instance creation time	3	"" if exists
(0008,0021)	Series Date	3	"" if exists
(0008,0022)	Acquisition Date	3	"" if exists
(0008,0033)	Content Time	2C	"" if exists
(0008,0031)	Series Time	3	"" if exists
(0008,0032)	Acquisition time	3	"" if exists
(0008, 0023)	Content Date	2	"" if exists
(0008, 1010)	Station name	3	"" if exists



A structured sample review was conducted by three radiologists on a subset of 60 cases out of 700 samples. Re-identification risk was assessed using criteria including replicability, distinguishability, and external data availability. The processed images were then used for annotation and ground truth preparation in downstream AI model development workflows.

Deployment & Implementation

Custom Python-based utilities were developed to enable reproducible data processing and preparation, including DICOM de-identification scripts for metadata transformation, report-parsing tools for structured extraction from PDF radiology reports, and preprocessing modules that export selected non-identifying technical parameters into analysis-ready tables. These scripts were iteratively tested on datasets from three clinical centres and refined through multiple validation cycles, and are structured as reusable research utilities, with broader release planned following additional multicentre testing. The tooling is intended solely to support research data preparation workflows and does not perform clinical interpretation or diagnostic decision-making.

Scale of Deployment

At the time of reporting, the approach has been applied to datasets from a limited number of centres. Broader multi-centre

validation across diverse imaging systems and institutional data practices is planned. Generalisability across sites, vendors, and jurisdictions requires further testing.

Identified Users

Primary users include AI and medical imaging researchers, academic institutions preparing imaging datasets, hospital-based research units, and medical imaging companies strengthening internal de-identification workflows. The tools and accompanying workflow documentation can also support ethics committee submissions by demonstrating structured privacy and data-governance safeguards, with ongoing updates planned as regulatory and standards guidance evolves.

Impact

Ethics and Governance

The code is planned for evaluation across at least 20 hospitals — spanning both public and private sectors — and approximately 5,000 CT volumes over the next five years. A key performance indicator will be successful post-de-identification reloading of images into hospital imaging software without errors or data corruption. The code is designed to be compatible with broader health data management ecosystems and may be integrated with national platforms such as ABDM to support secure contribution and reuse of radiology images.



Discussion

Successes

Early ethics engagement and advance approval planning supported timely dataset access. Cross-disciplinary collaboration between clinicians and engineers improved understanding of DICOM structures and metadata risks. Structured standards mapping helped align technical processing with regulatory expectations. The resulting workflow provides a documented and reproducible approach that other research groups may adapt and evaluate.

Challenges Faced

Ethics approvals required detailed clarification of privacy safeguards. Multi-centre dataset collection presented format and consistency challenges. CTCA datasets contained multiple DICOM series types, increasing processing complexity. Initial aggressive tag removal reduced system compatibility when images were reloaded into clinical software, requiring iterative refinement to balance de-identification with technical interoperability. These operational challenges highlight the need for domain expertise and iterative validation in privacy-preserving imaging workflows.

Lessons Learned

Companies developing medical imaging applications for radiology should periodically review applicable data-privacy and data-protection guidelines and incorporate updated requirements into their software designs. This supports researchers in accessing properly anonymised, standards-compliant datasets for legitimate research use. At the policy level, the Ministry of Information Technology could further strengthen the ecosystem by promoting greater harmonisation and clarity in data-privacy guidance for the Indian context, similar to established frameworks in the European Union and the United States.



Conclusion

This work describes a structured, standards-informed approach to de-identifying radiology image datasets for AI research use. The implemented workflow combines ethics approvals, regulatory mapping, rule-based DICOM tag processing, and expert review. Results indicate feasibility in controlled research settings and support privacy-aware dataset preparation.

Funding Source

This work was supported by the Ministry of Education, Government of India, managed by CPMU, IIT Jammu, under the AI Center of Excellence in Health project.

References

Hurkmans, C., Bibault, J., Brock, K. K., Van Elmpt, W., Feng, M., Fuller, C. D., Jereczek-Fossa, B. A., Korreman, S., Landry, G., Madesta, F., Mayo, C., McWilliam, A., Moura, F., Muren, L. P., El Naqa, I., Seuntjens, J., Valentini, V., & Velec, M. (2024). A joint ESTRO and AAPM guideline for development, clinical validation and reporting of artificial intelligence models in radiation therapy. *Radiotherapy and Oncology*, 197, 110345. <https://doi.org/10.1016/j.radonc.2024.110345>

Kocak, B., Baessler, B., Bakas, S., Cuocolo, R., Fedorov, A., Maier-Hein, L., Mercaido, N., Müller, H., Orhac, F., Santos, D. P. D., Stanzone, A., Ugga, L., & Zwanenburg, A. (2023). Checklist for evaluation of radiomics research (CLEAR): A step-by-step reporting guideline for authors and reviewers endorsed by ESR and EuSoMI. *Insights into Imaging*, 14(1), 75. <https://doi.org/10.1186/s13244-023-01415-8>

Kumar, A. S., & Sinha, N. (2020). Cardiovascular disease in India: A 360 degree overview. *Medical Journal Armed Forces India*, 76(1), 1–3. <https://doi.org/10.1016/j.mjafi.2019.12.005>

Mongan, J., Moy, L., & Kahn, C. E. (2020). Checklist for artificial intelligence in medical imaging (CLAIM): A guide for authors and reviewers. *Radiology: Artificial Intelligence*, 2(2), e200029. <https://doi.org/10.1148/ryai.2020200029>

National Electrical Manufacturers Association. (n.d.). *Digital Imaging and Communications in Medicine (DICOM) standard*, PS3.1. <https://dicom.nema.org/medical/dicom/current/output/html/part01.html>

Office of the Federal Register. (n.d.). 45 CFR § 164.514 — *Other requirements relating to uses and disclosures of protected health information*. <https://www.ecfr.gov/current/title-45/subtitle-A/subchapter-C/part-164/subpart-E/section-164.514>


PRS Legislative Research. (2023). *The Digital Personal Data Protection Act, 2023*. https://prsindia.org/files/bills_acts/bills_parliament/2023/Digital_Personal_Data_Protection_Act_2023.pdf

European Union. (2024, April 22). *General data protection regulation (GDPR): Legal text*. <https://gdpr-info.eu/>



From Text to Insight: Harnessing the power of AI to transform Unstructured Data in Brazil's Health Information Systems

Olivia Landi Corrales Guaranha



From Text to Insight: Harnessing the power of AI to transform Unstructured Data in Brazil's Health Information Systems

Olívia Landi Corrales Guaranha*

Abstract

Brazil's public health system operates multiple health information systems that generate large volumes of unstructured textual data, particularly in primary care clinical notes. Although these narratives contain clinically and socially relevant signals, they are not routinely incorporated into structured surveillance and policy analytics. This use case describes an open-source, resource-efficient artificial intelligence (AI) and natural language processing (NLP) pipeline developed to convert unstructured clinical text into structured, semantically interpretable data within Brazil's Unified Health System (SUS).

The approach combines a multilingual NLP toolkit with a semantic parser based on FrameNet, adapted to Brazilian Portuguese and extended with healthcare and violence-related semantic frames. The system is designed to operate within public infrastructure and under constrained computational conditions, avoiding dependence on proprietary large language models or external APIs. A pilot deployment on e-SUS APS data from Recife processed approximately 39 million sentences. Semantic pattern detection enabled earlier identification of underreported public health concerns, identifying an estimated 17.6% increase in detected cases of gender-based violence and suggesting substantial underreporting of self-harm and suicidal behavior.


These results indicate that semantically grounded AI methods can support expanded surveillance signal detection and analytic capacity when implemented with governance safeguards and human oversight. Findings should be interpreted as indicative surveillance support outputs, not confirmed case determinations. The chapter focuses on methodology, governance, operational feasibility, and public health decision-support value.

Background

Between 2010 and 2018, Brazil's public health system (Sistema Único de Saúde — SUS) operated dozens of national health

information systems (HIS) supporting surveillance, financing, patient registration, and planning. Most routine analyses rely on

*Vital Strategies



structured and parameterized fields, leaving free-text clinical notes underutilized for population health intelligence.


Since 2013, SUS has progressively digitized primary care records through e-SUS Atenção Primária à Saúde (e-SUS APS), enabling clinicians in Basic Health Units (UBS) to enter records electronically. Recent census data indicate that most primary care units now use digitized records. This transition has increased the volume of unstructured textual data, including clinical narratives, observations, and contextual notes.

Unstructured text can contain relevant indicators of social and health risks such as violence exposure and suicidal ideation that may not be consistently captured in structured reporting fields. However, extracting such signals requires language-aware analytic tools and governance safeguards due to the presence of sensitive personal data.

This use case describes an AI-supported semantic analysis pipeline designed to support public health surveillance and policy analytics by structuring clinical narrative data. Initial pilots focused on two public health priorities: gender-based violence and suicide risk signals. The system is positioned as an analytic support tool for surveillance teams and policymakers rather than an automated case detection system. Tiago Torrent is a Research Productivity grantee of the National Research Council for Scientific and Technological Development (CNPq grant # 311241/2025-5) for this use case.

AI Use Case Description

To address the underuse of textual data in health information systems (HIS), the authors describe a Natural Language Processing (NLP) pipeline designed for policymakers, surveillance teams, and researchers. The pipeline converts unstructured clinical text into structured, analysable data, enabling more informed policy decisions, improved detection of health-related events, and new evidence generation. Technically, the Python-based pipeline combines an established NLP toolkit (Trankit) with a custom semantic parser built on small language models (SLMs), balancing performance with low computational demand and remaining agnostic to downstream analytic tasks. The workflow reads tabular HIS datasets with textual fields, preprocesses and tokenizes text, performs linguistic annotation (sentence splitting, tokenization, lemmatization, POS and morphological tagging), and then applies semantic parsing based on a multilingual FrameNet-aligned architecture (LOME). Adapted and retrained with healthcare- and violence-specific frames from FrameNet Brasil, the parser extracts scene-based semantic structures and participant roles from Brazilian Portuguese clinical notes, producing structured representations from open text. Despite domain complexity, evaluation indicates adequate performance ($F1 \approx 62.5$ for span matching and ≈ 50.7 for exact matching), supporting practical use for large-scale, low-cost health text structuring and analysis.



To parse clinical notes written in Brazilian Portuguese, an adapted version of LOME was trained using an extended FrameNet Brasil database that includes frames specifically modeled for the healthcare and violence domains. The full database comprises 1,409 frames and 12,851 frame elements across multiple domains and abstraction levels, including 35 frames dedicated to healthcare and 43 to violence. The healthcare frames correspond to 2,776 lexical units commonly found in clinical notes, while the violence frames are evoked by 1,768 lexical items. These domain-specific frames were annotated on 2,352 sentences extracted from Brazilian Portuguese clinical records. Because frames are triggered by individual words, a single sentence may contain multiple frame annotations. In total, 14,622 frames and their associated participants were annotated. The resulting semantic parser was trained on these 14,622 domain annotations along with an additional 132,200 general-domain annotations in English, Spanish, and Portuguese, leveraging LOME's inherently multilingual architecture. The pipeline automatically annotates open-text fields with frames and frame elements. Performance is slightly below the baseline model, with F1 scores of 62.53 (baseline: 66.41) for span matching and 50.68 (baseline: 56.34) for exact matching, but remains robust given the technical variability and domain specificity of electronic medical record text.

Impact

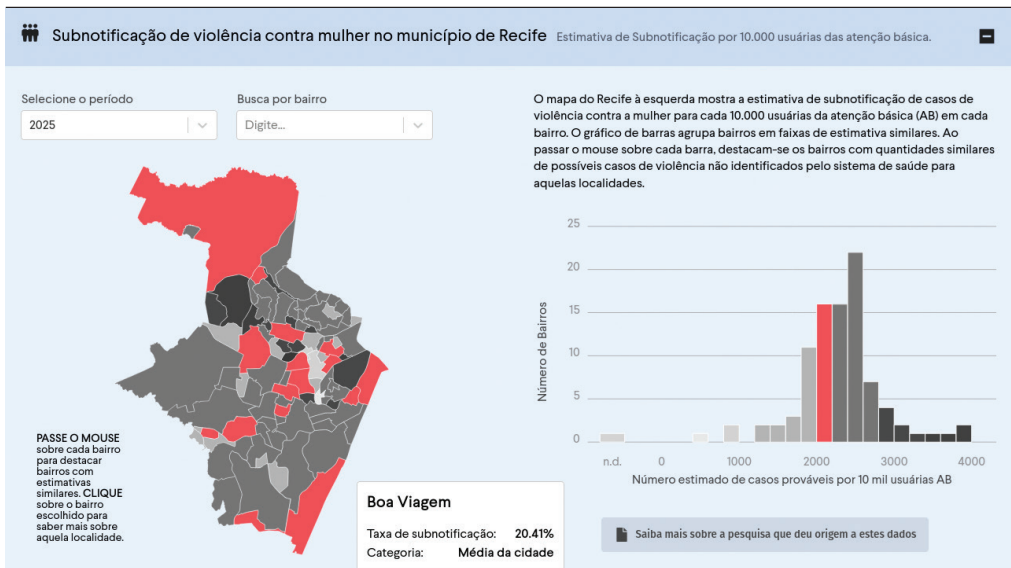
The pipeline was piloted using the e-SUS APS database of Recife, Brazil, covering

records entered between 2016 and 2025. Automatic semantic annotation of approximately 39 million sentences produced about 138 million annotation sets, suggesting the scalability of the approach. Once clinical text is semantically annotated, it can be queried for fine-grained semantic patterns using the transparent and explainable FrameNet framework, which supports definition of search criteria. The methodology has been applied to detect complex and often underreported public health issues, particularly gender-based violence (GBV) and suicide risk. GBV remains substantially underreported in healthcare systems, with estimates suggesting underreporting of physical violence near 75.9%, and most cases documented only at hospital or emergency service level. Eight semantic patterns were defined for GBV detection, including abuse, injury, sexual violence, and neglect involving family members or related persons; the pattern-matching approach achieved an average precision of 0.726. Semantic pattern detection in primary care records identified 5,142 additional GBV cases, representing a 17.6% increase over officially reported records. A similar semantic strategy was applied to suicide risk, where routine HIS data typically capture only hospitalizations, attempts, or deaths. Semantic analysis of clinical notes identified 1,962 individuals reporting suicidal ideation between 2016 and 2022 and 2,174 individuals mentioning suicide attempts. Among 2,763 individuals with a history of self-harm recorded across primary care and mortality data, 2,400 had not been officially reported, indicating 88.4% underreporting in primary healthcare services. Results

were delivered to surveillance units and health managers in Recife through interactive dashboards that integrate HIS data with semantic annotations, enabling stratification by age group, territory, and time period, longitudinal patient timelines, and visualization of lexical and semantic patterns related to violence and self-harm.

This integrated analytic layer supports earlier detection, targeted intervention, and more informed public health decision-making, with potential future extensions to support alert mechanisms for at-risk patients, subject to further validation and governance safeguards.

Figure 01: Dashboard presenting estimated underreporting of GBV across Recife. The left panel maps neighborhoods shaded by their underreporting estimates (“taxa de subnotificação”, normalized by 10 thousand PHC female users), while the right panel displays the distribution of neighborhoods according to the estimated number of potential violence cases in primary care (the x-axis indicates the number of underreporting estimates, while the y-axis indicates the number of neighborhoods that concentrate one similar underreporting interval).





Discussion

Early identification of interpersonal violence and self-harm risk is a recognized public health priority. Primary care narratives contain contextual information that is often not captured in structured fields. AI-supported semantic analysis can help make such information more accessible for surveillance analytics when implemented with safeguards. This implementation demonstrated that large-scale semantic text processing is operationally feasible within SUS infrastructure using open-source and small-model approaches. The resulting analytics supported new types of surveillance queries and managerial insights. These outputs are decision-support inputs, not determinations.

Implementation Challenges


Key challenges included strict privacy and access controls for sensitive data, heterogeneous HIS formats with limited interoperability, the absence of standardized data extraction protocols, constrained computational infrastructure, and the need for close collaboration between health and IT teams; deployment under these conditions was supported through the use of smaller models and open-source components that reduced resource requirements while maintaining functionality.

Ethics and Governance

Ethical and governance requirements are central to the design, given that clinical

text often contains personally identifiable information (PII). While structured fields can be anonymised by excluding direct identifiers, free text requires dedicated de-identification steps when human review is needed; importantly, the final structured outputs do not retain names, dates, or document numbers. To mitigate data protection and infrastructure constraints, particularly in decentralised public health systems with limited computing capacity, the solution is open-source, resource-efficient, and deployable within public infrastructure, avoiding reliance on external AI APIs and high-cost hardware.

The work has been conducted in close partnership with local health surveillance teams and has been developed under a Technical Cooperation Agreement that details data protection measures and processes. The work has been conducted in compliance with the Brazilian Data Protection Legislation (LGPD) and personal data has been accessed only by authorized personnel in secure environments. At this stage, the analysis is conducted at an aggregated and population level (e.g., neighbourhood or municipality), without producing outputs intended for individual-level clinical decision-making. Given the secondary use of existing surveillance data for public health purposes, and the absence of direct patient interaction, individual informed consent was not required. Where applicable, the work follows institutional and legal provisions that allow a waiver of consent for public health surveillance activities, consistent with Brazilian regulations and ethical frameworks.



The current implementation supports public health surveillance activities using routinely collected health data and generates outputs at aggregated levels (e.g., neighborhood or municipality). At this stage, the system does not produce automated individual-level recommendations or trigger direct patient interventions. Given the secondary use of existing surveillance data for public health purposes, and the absence of direct interaction with patients, individual informed consent was not required. Future extensions may include mechanisms to support identification of potentially at-risk individuals. Any such implementation would be subject to additional validation studies, ethical review as required by applicable regulations, and the establishment of specific governance safeguards, including clearly defined protocols for patient engagement and communication.

The AI system is designed as a decision-support tool to assist local surveillance teams in identifying trends and potential signals from large volumes of clinical text. All outputs require interpretation and validation by trained public health professionals before any operational response is initiated. The tool does not trigger automated interventions and does not replace human judgement; rather, it helps prioritise information for review within existing surveillance workflows.

Lessons Learned

Operational lessons indicate that resource-efficient models significantly improve deployability in decentralized systems, explainable semantic structures increase governance and stakeholder acceptance, domain-specific frame modeling enhances analytic relevance, cross-disciplinary collaboration between health and technical teams is essential for safe implementation, and all surveillance outputs should remain subject to human review and oversight.




Conclusion

This use case describes an open-source, resource-efficient AI and NLP pipeline designed to convert unstructured clinical text into structured semantic data for public health analytics within Brazil's SUS. A pilot deployment in Recife demonstrated technical feasibility and operational scalability and produced additional candidate surveillance signals related to gender-based violence and self-harm.

Results indicate that semantically grounded AI methods can support surveillance capacity and analytic depth when deployed within public infrastructure and governance frameworks. Outputs should be interpreted as indicative analytic signals rather than confirmed case counts. Continued validation, governance oversight, and interdisciplinary collaboration are necessary for responsible scale-up.



References

- Neto, G.C.C. & Chioro, A. (2021). Afinal, quantos Sistemas de Informação em Saúde de base nacional existem no Brasil? *Cad. Saúde Pública*, 37 (7), <https://doi.org/10.1590/0102-311X00182119>
- Cielo, AC; Raioli, T; Silva, EN; Barreto, JOM. (2022) Implementation of the e-SUS Primary Care Strategy: an analysis based on official data. *Revista de Saúde Pública*, 56:5. <https://doi.org/10.11606/s1518-8787.2022056003405>
- Brasil. Ministério da Saúde. Secretaria de Atenção Primária à Saúde. Coordenação-Geral de Programação de Financiamento da Atenção Primária. (2025) Censo das Unidades Básicas de Saúde: Sumário Executivo. Ministério da Saúde, Secretaria de Atenção Primária à Saúde. Brasília: Ministério da Saúde. Available from: <https://www.gov.br/saude/pt-br/composicao/saps/censo-das-ubs/publicacoes/censo-nacional-ubs-2024.pdf>
- Van Nguyen, M., Lai, V. D., Veyseh, A. P. B., & Nguyen, T. H. (2021). Trankit: A light-weight transformer-based toolkit for multilingual natural language processing. *arXiv preprint arXiv:2101.03289*.
- Conneau, A., Khandelwal, K., Goyal, N., Chaudhary, V., Wenzek, G., Guzmán, F., ... & Stoyanov, V. (2020, July). Unsupervised cross-lingual representation learning at scale. In *Proceedings of the 58th annual meeting of the association for computational linguistics* (pp. 8440-8451).
- Xia, P., Qin, G., Vashishtha, S., Chen, Y., Chen, T., May, C., Harman, C., Rawlins, K., White, A. S., & Van Durme, B. (2021). LOME: Large ontology multilingual extraction. In D. Gkatzia & D. Seddah (Eds.), *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: System Demonstrations* (pp. 149–159). Association for Computational Linguistics. <https://doi.org/10.18653/v1/2021.eacl-demos.19>
- Torrent, T. T., Matos, E. E. D. S., Belcavello, F., Viridiano, M., Gamonal, M. A., Costa, A. D. D., & Marim, M. C. (2022). Representing context in FrameNet: A multidimensional, multimodal approach. *Frontiers in Psychology*, 13, 838441.
- Fillmore, C. J. (1982). Frame Semantics. In *The Linguistic Society of Korea (org.). Linguistics in The Morning Calm*, (pp. 111-137). Seoul: Hanshin Publishing Co.
- Dutra, L., Lorenzi, A., Larré, L., Belcavello, F., Matos, E., Pestana, A., Brown, K., Gonçalves, M., Herbst, V., Reinach, S., Teixeira, R., de Paula, P., Pellini, A., Sequeira, C., Sabino, E., Leal, F., Conde, M., Grespan, R., & Torrent, T. (2023). Building a Frame-Semantic Model of the Healthcare Domain: Towards the identification of gender-based violence in public health data. In *Anais do XIV Simpósio Brasileiro de Tecnologia da Informação e da Linguagem Humana*, (pp. 338-346). Porto Alegre: SBC. doi:10.5753/stil.2023.234145
- Larré, L., & Torrent, T. (2024). Modelagem baseada em frames para identificação do léxico da Violência de Gênero. In *Anais do XV Simpósio Brasileiro de Tecnologia da Informação e da Linguagem Humana*, (pp. 403-412). Porto Alegre: SBC. doi:10.5753/stil.2024.245452
- Dutra, L., Lorenzi, A., Berno, L., Campos, F., Biscardi, K., Brown, K., Viridiano, M., Belcavello, F., Matos, E., Guaranha, O., Santos, E., Reinach, S. & Torrent, T. T. (2025). Frame Semantic Patterns for Identifying Underreporting of Notifiable Events in Healthcare: The Case of Gender-Based Violence. *arXiv preprint arXiv:2510.26969*.
- Prefeitura do Recife (2025). USF+ Jordão Alto é entregue e amplia acesso aos serviços de saúde, Notícias, 29 Dec. 2025. Available from: <https://www2.recife.pe.gov.br/noticias/29/12/2025/usf-jordao-alto-e-entregue-e-amplia-acesso-aos-servicos-de-saude#:~:text=AMPLIA%C3%87%C3%83O%20DE%20ACESSO%20-%20Com%20as,80%25%20em%20todo%20o%20territ%C3%B3rio>, access on 08 Jan. 2026.
- G1 (2025). Brasil tem mais de mil casos de feminicídio registrados em 2025, *Jornal Nacional*, 01 Dec. 2025. Available from: <https://g1.globo.com/jornal-nacional/noticia/2025/12/01/brasil-tem-mais-de-mil-casos-de-feminicidio-registrados-em-2025.ghtml>, access on 08 Jan. 2026.
- Vasconcelos, NM, Bernal, RTI, Souza, J., Bordoni, P.H.C, Stein, C., Coll, C. V. N., Murray, J, Malta, DC. (2024) Underreporting of Violence Against Women: an analysis of two data sources. *Ciência & Saúde Coletiva*, 29:e07732023, <https://doi.org/10.1590/1413-812320242910.07732023EN>



HealthVaani: The Generative AI-powered Voice of Knowledge for India's Frontline Health Workers

Prekshaa Arunachalam, Anika Murarka,
Vaibhav Patwardhan, Ashok Jaybhaye, Daksha Dixit,
Sambhav Solanki, Shreya Thakur, Prachi Karkhanis,
Shivi Kapil, Parag Govil, Chintan Donda, Vipin Samaria,
Alugubelli Dinesh Reddy, Saransh Khanna,
Nasir Hussain, Arpita Srivastava, Atharv Dadhich,
Ranjana N Rao, Isha Shah, Divay Khatri,
Sai Nikhilesh Reddy, Tumul Rai



HealthVaani: The Generative AI-powered Voice of Knowledge for India's Frontline Health Workers

Prekshaa Arunachalam, Anika Murarka, Vaibhav Patwardhan, Ashok Jaybhaye, Daksha Dixit, Sambhav Solanki, Shreya Thakur, Prachi Karkhanis, Shivi Kapil, Parag Govil, Chintan Donda, Vipin Samaria, Alugubelli Dinesh Reddy, Saransh Khanna, Nasir Hussain, Arpita Srivastava, Atharv Dadhich, Ranjana N Rao, Isha Shah, Divay Khatri, Sai Nikhilesh Reddy, Tumul Rai*

Abstract

India's public health and Integrated Child Development Services systems rely on over 3 million frontline health workers (FLWs), including Accredited Social Health Activists (ASHAs), Anganwadi Workers (AWWs), and Auxiliary Nurse Midwives. These workers operate in diverse and often resource-constrained settings and are expected to apply evolving programme guidance across maternal, child health, nutrition, and preventive services. Access to timely, verified, and comprehensible reference information remains a persistent operational challenge.

HealthVaani is a generative AI-enabled, multilingual knowledge assistant designed to support FLWs by providing on-demand, guideline-grounded responses through voice and text interaction. The system uses a retrieval-augmented generation architecture anchored in government-validated training materials and programme guidelines. It is designed as an in-service decision-support and knowledge-access tool, not as a diagnostic or clinical decision system.


Pilot deployments with 628 Anganwadi Workers across six Indian states examined usability, adoption, and perceived usefulness over a two-month period. Usage logs and survey responses indicate that workers engaged with the system across multiple programme domains and interaction modes. Reported benefits relate primarily to knowledge access, convenience, and user confidence. Findings should be interpreted as feasibility and adoption indicators rather than evidence of service delivery or health outcome impact. This use case illustrates how generative AI knowledge assistants may support frontline capacity when grounded in validated sources, multilingual access, and human oversight.

Background

Frontline health workers (FLWs) form the first point of contact between communities and public health and nutrition systems in India. Accredited

Social Health Activists (ASHA) support household outreach, registration of pregnant women and newborns, referral mobilisation, and promotion of

*Wadhvani AI, New Delhi, India



preventive practices. Anganwadi Workers provide supplementary nutrition, growth monitoring, and early childhood support. Auxiliary Nurse Midwives deliver clinical and preventive services at sub-centre level. Together, these cadres support maternal, child health, nutrition, and disease prevention under national programme frameworks.

Despite their central role, multiple studies have documented that FLWs often face constraints in accessing timely, reliable, and easy-to-interpret programme guidance. Information sources may be fragmented across printed manuals, registers, circulars, and supervisory channels. Knowledge recall from initial training may decline over time, while programme responsibilities continue to expand. Published assessments have identified knowledge gaps in recognising danger signs, differentiating conditions, and applying protocols consistently.

Digital access among FLWs has increased through government and partner initiatives, creating opportunities for mobile-enabled decision-support and knowledge-access tools. However, general web search and video platforms may not reliably provide programme-aligned or validated guidance and may expose workers to inconsistent or misleading content. Global public health guidance has highlighted misinformation risks in digital environments.

HealthVaani was developed to address these gaps by providing a multilingual, AI-enabled knowledge assistant grounded in validated programme materials. Its intended purpose is to support comprehension, recall, and access to approved guidance during routine fieldwork.

Solution Overview

System Overview

HealthVaani is a task-oriented AI knowledge assistant that answers work-related queries from FLWs using a retrieval-augmented generation (RAG) pipeline. The design emphasises grounding in validated sources, multilingual accessibility, and guardrails to reduce unsupported responses.

- a. **Speech-to-Text (STT):** Spoken queries in supported languages are transcribed using production speech-recognition models. This supports workers who prefer voice interaction or face typing constraints.
- b. **Translation Layer:** A language model-based translation module converts Indic-language queries into English for downstream processing and translates generated answers back into the user's selected language. Translation quality is monitored through expert review samples.
- c. **Knowledge Base Construction:** The knowledge base is derived from government-verified and expert-reviewed training materials, guidelines, and standard operating procedures across domains such as maternal health, child health, nutrition, immunisation, and growth monitoring. English documents are used as canonical references for consistency and expert validation. Documents are parsed into structured chunks and reviewed by public health experts before indexing.

d. Chunking and Embedding: Page-level chunking with domain annotations, acronyms, and metadata is applied. Chunks are embedded using sentence-level embedding models and stored in a vector database for semantic retrieval.

e. Ranking and Retrieval: User queries are embedded and matched to top-ranked chunks using a hybrid semantic and lexical retrieval approach. This supports robustness for short, noisy, or domain-specific queries.

f. Answer Generation: The language model receives the translated query and retrieved knowledge chunks with metadata and is instructed to generate answers strictly grounded in provided context. If supporting information is insufficient, the system is designed to return a non-answer response rather than speculate.

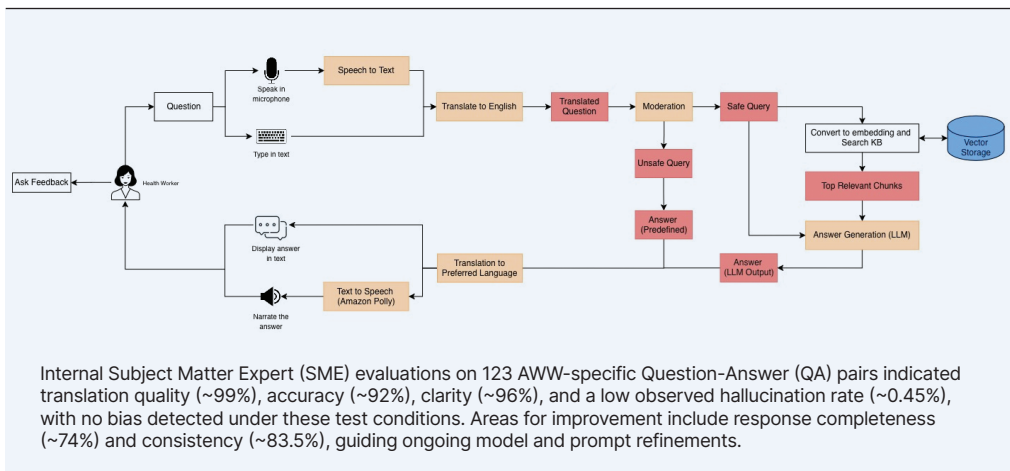
g. Moderation and Guardrails: Intent classification, personally identifiable information detection, and safety filters are applied to detect out-of-scope or unsafe queries. Queries beyond scope are declined or redirected.


h. Text-to-Speech (TTS): Optional voice output supports users with limited literacy or preference for audio responses.

Deployment & Implementation

HealthVaani was developed through an iterative, user-centred design process grounded in frontline workflows and aligned with national health and ICDS programme guidelines. Capacity building was delivered through hands-on, classroom-based demonstrations supported by Information, Education, and Communication (IEC) materials. A structured

Figure 01: HealthVaani System Overview





half-day training format was used, focusing on live demonstrations, guided installation, feature walkthroughs, and practical troubleshooting. The training approach was intentionally light-touch to enable rapid adoption and scalability across diverse implementation settings. More than 95% of AWWs reported high satisfaction with the training experience. For many participants, this represented their first interaction with an AI assistant, and the practical, demonstration-led format was reported to support user confidence and familiarisation. Following training, AWWs accessed HealthVaani via personal or government-issued smartphones to support routine service delivery tasks. Continued engagement was reinforced through simple operational supports, including WhatsApp reminders, periodic follow-ups, and field-level coordination.

Scale of Deployment

HealthVaani pilots were conducted over approximately two months across six Indian states, with one district per state and phased rollout to allow iterative learning. A total of 628 Anganwadi Workers participated across Uttar Pradesh (88), Haryana (84), Meghalaya (98), Arunachal Pradesh (110), Rajasthan (132) and Madhya Pradesh (116).

The phased model was designed to test usability, onboarding processes, and operational feasibility before larger-scale expansion. Pilot scale should be interpreted as an early implementation cohort rather than representative national coverage.

User-level Outcomes

Pilot evaluation focused on adoption, usage patterns, and perceived usefulness rather than

clinical or service delivery outcomes. Multiple data sources were used, including backend usage logs, qualitative interviews, and post-pilot surveys. Survey responses were obtained from 545 of 628 participants; results should be interpreted with this response rate in mind.

Usage Patterns: Approximately 10,800 interactions were recorded, including about 8,000 unique queries. The system handled roughly 400 queries per week, delivering responses with an average latency of 5 seconds. Interaction modes were distributed across voice, text, and selectable prompt options, consistent with the intended multimodal design. Most queries were in Hindi, with the remainder in English. Overall, 70% of queries received AI-generated responses. In-app feedback indicated that 70% of responses were rated positively, 1% negatively, and 29% received no feedback. The remaining 30% queries were beyond the scope of existing FLW guidelines or due to minor application limitations. Moreover, 74% of users reported using the AI at least once weekly. Usage was highest in Rajasthan (91%) and Madhya Pradesh (80%), and lower in Arunachal Pradesh (41%), likely reflecting connectivity and infrastructure differences. Average session duration was ~12 minutes.

User Personas: Analysis of usage logs and interviews suggested variation in digital confidence and interaction patterns. Four indicative personas- Influencers, Believers, Observers, and Onlookers were identified to guide design and onboarding strategies. Higher-engagement users were more comfortable rephrasing queries and exploring features, while lower-engagement users experienced higher cognitive and usability barriers. Personas are heuristic tools and may evolve with further study.

Table 01: Identified User Personas and Interaction Patterns

Persona	Interaction & Cognitive Load	Key Challenges / Needs
<i>Influencers</i>	Low cognitive load and frustration; are digitally fluent and comfortable phrasing and rephrasing queries using voice and text; need ideas for more themes to discover.	Need exploratory prompts (eg, “what else can I ask?”); moderate cognitive load in reaching optimal answers.
<i>Believers</i>	Low cognitive load and time pressure felt; limited rephrasing of queries and task-based use.	Friction if answers fall short or no answers are received. Friction when transcription errors occur; need guidance in rephrasing and reassurance.
<i>Observers</i>	High cognitive load, effort, and frustration; struggle with typing, using voice for input, and navigation. Needs support in articulating questions.	Easily overwhelmed by usability barriers; need simplified flows and literacy-sensitive support
<i>Onlookers</i>	Low mental demand but high frustration and difficulty in comprehending response by digital tools; minimal independent use	Depend on others to operate the app: need simplified flows and literacy-sensitive support


Perceived Usefulness: Reported satisfaction with HealthVaani responses to user queries was highest in Haryana (97%) followed by Uttar Pradesh (96%) and Rajasthan (96%). Meghalaya (80%), Madhya Pradesh (91%), and Arunachal Pradesh (92%) reported comparatively lower than the other pilot geographies. The most prominent reasons included poor internet connectivity and low digital literacy in these pilot regions.

Preferred Features and Accessibility: Voice interaction was frequently cited as helpful,

particularly where typing was difficult. Quick responses and continuous availability were also valued. Interestingly, states like Meghalaya and Arunachal Pradesh preferred typing or selecting questions instead of voice-based interaction. Other valued features included quick responses (58%), round-the-clock access (53%), ease of application navigation (48%), and language accessibility, though the latter varied substantially by region.

Efficiency and Confidence: Nearly 80% of surveyed AWW users reported reduced

Figure 02: Perceived usefulness of HealthVaani by pilot participants




time spent seeking information compared to traditional channels (supervisors and colleagues). Among respondents, 91% reported confidence in using the tool, and 75% expressed intent to continue using and recommend HealthVaani to peers, suggesting perceived practical value among pilot participants.

Observations from the pilots suggest that HealthVaani may be scalable with limited additional infrastructure, though further assessment would be needed to confirm feasibility for expansion to a larger cohort of frontline health workers.

Ethics & Governance

HealthVaani operates in a sensitive, high-stakes domain covering maternal

and child health, chronic disease management, and programme protocols, and is designed in alignment with global guidance on responsible AI for health (WHO, 2021). The system incorporates layered safety guardrails, including pre- and post-generation moderation checks to detect unsafe, out-of-scope, personally identifiable, or potentially biased content, with conservative fallback responses or refusal where reliable answers cannot be generated. All outputs are grounded exclusively in a curated knowledge base derived from validated official guidelines, and responses are provided only when supported by retrieved source content; unanswered queries are logged to identify knowledge gaps. Transparency and accountability are built in through source referencing, explicit communication



of system limitations, and end-to-end traceability via monitoring tools. Expert oversight is maintained throughout the lifecycle, including knowledge base curation, dataset preparation, and periodic expert review of system outputs. Inclusive design principles are applied through multilingual support, voice interaction, and text-to-speech features co-designed with frontline workers to reflect field realities. All pilot implementations were conducted with prior institutional permissions, with in-app opt-in consent for service data and written informed consent for survey participants, alongside strict safeguards for anonymity and confidentiality in data handling and reporting.

Discussion

Pilot deployments across six states suggest that a domain-grounded, multilingual AI knowledge assistant can be integrated into frontline workflows with moderate training and support. Adoption appears linked to relevance of content, multimodal access, and clarity about system scope and limits. Grounding answers in validated sources and communicating constraints supports user trust, consistent with human-centred AI guidance.

At the same time, several challenges were observed. The pilots were designed to assess adoption, usability and engagement, rather than clinical or service delivery outcomes. The short pilot duration, absence of control groups, and perceived outcome indicators limit causal inference. Endline

surveys were completed by 545 of 628 users via Google Forms, and findings should be interpreted in light of this partial response rate. At the solution level, limitations emerged in handling poorly phrased, ambiguous, or conversational queries, highlighting the need to improve both the AI-Assistant's robustness to linguistic variability and user support for effective query formulation


Cost and latency trade-offs influenced model and infrastructure choices, reflecting public-sector resource constraints. Persona analysis suggests that differentiated onboarding and interface scaffolding may improve inclusion for lower digital-confidence users. Integration with existing digital workflows may reduce adoption friction, as suggested in digital health implementation research.

Conclusion

HealthVaani demonstrates the feasibility of deploying a generative AI-enabled, multilingual knowledge assistant to support frontline health workers with on-demand access to validated programme guidance. Pilot results indicate user acceptance, multimodal engagement, and perceived usefulness under supported conditions. Findings relate to feasibility, usability, and adoption rather than verified service or health outcomes.


References

- Agarwal, S., Perry, H. B., Long, L.-A., & Labrique, A. B. (2015). Evidence on feasibility and effective use of mHealth strategies by frontline health workers in developing countries: A systematic review. *Tropical Medicine & International Health*, 20(8), 1003–1014.
- Ansari, M. Z., Saha, S., & Rosso, P. (2021). *Language identification of Hindi-English tweets using code-mixed pre-training* (arXiv:2107.01202). arXiv. <https://arxiv.org/abs/2107.01202>
- BHASHINI. (n.d.). *National language translation mission and ULCA platform*. <https://bhashini.gov.in/>
- Census of India. (2011). *C-16: Population by mother tongue*. Office of the Registrar General and Census Commissioner, India.
- Dong, F., Guo, W., Liu, J., Patterson, T. A., & Hong, H. (2024). BERT-based language model for accurate drug adverse event extraction from social media: Implementation, evaluation, and contributions to pharmacovigilance practices. *Frontiers in Public Health*, 12, 1392180. <https://doi.org/10.3389/fpubh.2024.1392180>
- Dudala, S. R., Ponna, S. N., Upadrasta, V. P., Bathina, H., Sadasivuni, R., Geddam, J. J. B., & Kapu, A. K. R. (2021). Assessment of gaps of knowledge and practices of frontline community workers in Chandragiri Mandal, Chittoor district, Andhra Pradesh: Maternal and child health services. *International Journal of Community Medicine and Public Health*, 8(3), 1299. <https://doi.org/10.18203/2394-6040.ijcmph20210817>
- Eysenbach, G. (2020). How to fight an infodemic: The four pillars of infodemic management. *Journal of Medical Internet Research*, 22(6), e21820. <https://doi.org/10.2196/21820>
- Frontline health workers in India. (2019, May 16). *India Development Review*. <https://idronline.org/idr-explains-frontline-health-workers/>
- Greenhalgh, T., Wherton, J., Papoutsis, C., Lynch, J., Hughes, G., A'Court, C., Hinder, S., Fahy, N., Procter, R., & Shaw, S. (2017). Beyond adoption: A new framework for theorizing and evaluating nonadoption, abandonment, and challenges to the scale-up, spread, and sustainability of health and care technologies. *Journal of Medical Internet Research*, 19(11), e8775. <https://doi.org/10.2196/jmir.8775>
- Kapadia-Kundu, N., Sullivan, T. M., Safi, B., Trivedi, G., & Velu, S. (2012). Understanding health information needs and gaps in the health care system in Uttar Pradesh, India. *Journal of Health Communication*, 17(sup2), 30–45. <https://doi.org/10.1080/10810730.2012.666625>
- Labrique, A. B., Vasudevan, L., Kochi, E., Fabricant, R., & Mehl, G. (2013). mHealth innovations as health system strengthening tools: 12 common applications and a visual framework. *Global Health: Science and Practice*, 1(2), 160–171.
- Maurya, V. K., Upadhyay, V., Dubey, P., Shukla, S., & Chaturvedi, A. (2021). Assessment of frontline healthcare workers' knowledge, attitude and practice after several months of COVID-19 pandemic. *Journal of Healthcare Quality Research*. <https://doi.org/10.1016/j.jhq.2021.07.004>
- Medhi, I., Sagar, A., & Toyama, K. (2006). Text-free user interfaces for illiterate and semi-literate users. In *Proceedings of the International Conference on Information and Communication Technologies and Development* (pp. 72–82). IEEE.
- Scott, K., George, A. S., & Ved, R. R. (2019). Taking stock of 10 years of published research on the ASHA programme: Examining India's national community health worker programme from a health systems perspective. *Health Research Policy and Systems*, 17(1). <https://doi.org/10.1186/s12961-019-0427-0>
- Sidiropoulos, G., Voskarides, N., Vakulenko, S., & Kanoulas, E. (2021). Combining lexical and dense retrieval for computationally efficient multi-hop question answering. In *Proceedings of the Second Workshop on Simple and Efficient Natural Language Processing*. <https://doi.org/10.18653/v1/2021.sustainlp-1.7>
- Sondankar, P. D., Kotnis, S. D., & Kumavat, A. P. (2015). Profile of Anganwadi workers and their knowledge regarding maternal and child health services in an urban area. *International Journal of Medical Science and Public Health*, 4, 502–507.



**Predictive Virtual
Cardiac Twin:
An AI-Powered
Immersive
Visualization
Platform for
Advanced
Cardiac Care**

S. Premlin Dass



Predictive Virtual Cardiac Twin: An AI-Powered Immersive Visualization Platform for Advanced Cardiac Care

S. Premlin Dass*

Abstract

Healthcare systems increasingly rely on advanced medical imaging to support diagnosis, intervention, and follow-up care. Although imaging resolution and acquisition capabilities have improved substantially, clinical interpretation is still largely performed through two-dimensional slice review of inherently three-dimensional anatomy. This can increase cognitive load and contribute to variability in interpretation, particularly in complex anatomical and procedural contexts.

This use case describes an AI-enabled Immersive Medical Visualization (IMV) platform developed to transform standard CT and MRI datasets into interactive three-dimensional visual representations for clinical use. The platform combines automated segmentation, cloud-based processing, and immersive visualisation interfaces to support spatial understanding, procedural planning discussions, clinical training, multidisciplinary collaboration, and patient communication.


The system reorganises and renders existing imaging data rather than generating new diagnostic information. AI components are limited to segmentation, data organisation, and spatial reconstruction. Reported benefits relate to workflow support and user-reported operational value in early deployments and modelling analyses, rather than proven clinical outcome improvement. The use case illustrates how AI-enabled visualisation platforms can complement clinical workflows when implemented with human oversight, governance safeguards, and integration into existing imaging pathways.

Background

Medical imaging is a foundational component of modern healthcare delivery. Global and national initiatives have emphasised strengthening imaging capacity and access. Over recent decades, investments in CT and MRI

technologies have increased image resolution, acquisition speed, and availability. Despite these advances, interpretation workflows continue to rely predominantly on two-dimensional slice review displayed on flat screens.

*Cresendos, Tamil Nadu, India



For many routine diagnostic scenarios, two-dimensional review is sufficient. However, in clinical situations that depend heavily on spatial relationships, including complex surgery, structural interventions, and anatomically variable conditions, clinicians must mentally reconstruct three-dimensional anatomy from sequential image slices. This process is well established but cognitively demanding and may vary by experience and training level. In high-volume and resource-constrained environments, time available for detailed mental reconstruction may be limited.

Advanced multimodal imaging is sometimes underutilised in procedural planning and training contexts due to workflow and interpretation constraints. Concurrently, hospital systems in several regions, including India, are experiencing increased patient volumes and procedural complexity. This creates demand for digital tools that enhance interpretation and communication using already-available imaging data, without requiring new acquisition protocols.

Immersive Medical Visualization platforms address this gap by converting conventional imaging datasets into interactive spatial models aligned with human visual cognition. These platforms aim to support comprehension, planning discussions, and education. They do not replace imaging acquisition or radiological interpretation and are intended as adjunct tools within existing clinical governance structures.

AI Use Case Description

Solution Description

The term “Predictive” reflects the platform’s planned future capabilities and long-term roadmap. The current system does not generate predictions, risk scores, or treatment recommendations. Any predictive functionality is prospective, subject to further development, validation, and regulatory approval. This clarification is provided to avoid confusion regarding present capabilities.

Predictive Virtual Cardiac Twin (PVCT) is described as an immersive visualisation application that converts routine CT and MRI datasets into patient-specific three-dimensional cardiac models for interactive review. The primary use case is support for spatial understanding of complex cardiac anatomy during procedural planning and multidisciplinary review. While cardiology is the reference domain, the platform architecture is designed to be adaptable to other organ systems.

In current practice, clinicians integrate multiple image series and mentally reconstruct spatial relationships when planning procedures such as coronary artery bypass grafting, valve interventions, or structural repairs. PVCT aims to support this process by presenting reconstructed spatial models derived directly from source images.

AI components are used for automated segmentation, anatomical structure identification, data validation, and spatial reconstruction. The system is intentionally designed with bounded AI scope. It does not generate diagnoses, risk scores, or treatment recommendations. It does not introduce synthetic anatomical elements or predictive overlays. Visualised structures correspond to segmented source imaging data.

A proprietary orchestration layer coordinates segmentation and reconstruction steps across variable image quality and acquisition protocols. Although model internals are proprietary, system behaviour is constrained so that outputs remain traceable to source images. Clinical interpretation and decision-making remain the responsibility of clinicians.

Future Development Pathway: Anomaly Prediction

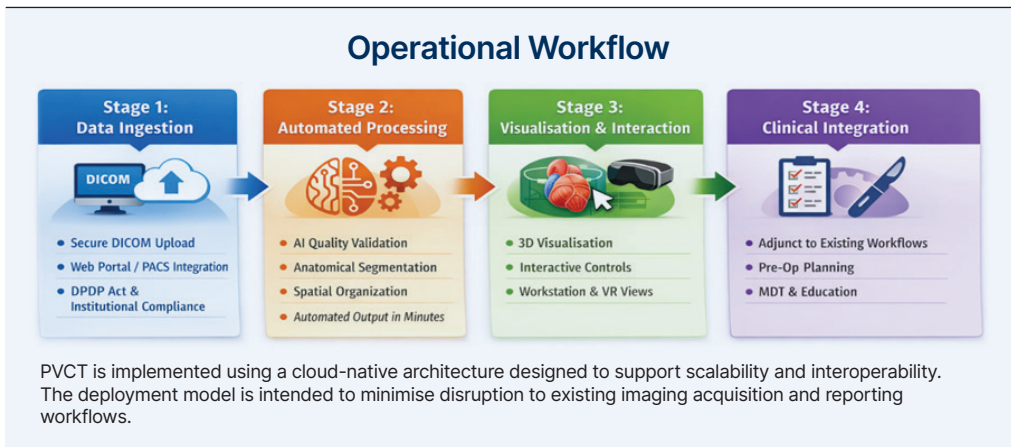
Future development pathways described for the platform include potential predictive or longitudinal anomaly visualisation capabilities. These are presented as planned or exploratory features rather than currently validated functions. Proposed capabilities would aim to support illustration of possible anatomical progression patterns for clinician discussion and patient communication.


Deployment and Implementation

The operational workflow includes four stages:

Data ingestion: Standard DICOM datasets are securely uploaded via web portals or integrated connections with PACS and

Figure 01: Operational workflow





enterprise repositories. Data handling is described as aligned with applicable data protection regulations and institutional governance requirements.

Automated processing: An AI-enabled pipeline performs image validation, anatomical segmentation, and spatial reconstruction. Tasks that are traditionally manual or semi-manual are automated. Processing time is reported as minutes under typical conditions, but may vary with dataset size and quality.

Visualisation and interaction: Outputs are delivered through configurable interfaces, including standard workstations and immersive environments. Users can rotate, isolate, and explore anatomical structures. Use cases include planning discussions, multidisciplinary meetings, training, and patient education.

Clinical integration: The platform is positioned as an adjunct to radiology reports and clinical planning processes. Typical uses include pre-operative reviews, team briefings, and educational sessions. It does not replace radiological interpretation or formal reporting.

Deployment models support both centralised cloud infrastructure and institution-controlled environments, enabling flexibility across hospitals with varying digital maturity.


Scale of Deployment

The platform is designed for use in tertiary and quaternary care centres managing

complex cardiac cases. Initial deployments have focused on selected institutions where procedural planning complexity is high. Architecture supports multi-site deployment across hospital networks. Scalability features include organ-agnostic modelling capability, automated workflows, and limited additional training requirements. Broader expansion is described as feasible but dependent on regulatory approvals, institutional readiness, and governance alignment. Reported experience reflects early and targeted deployments rather than population-level scale. The product scales seamlessly across small, mid-size, and high-volume hospitals, with flexible provisioning of computing resources as needed.

Identified Users

Primary users are clinicians involved in cardiac and structural care, including cardiac surgeons, interventional cardiologists, and radiologists. These users employ the platform to support spatial review and multidisciplinary communication. Additional clinical team members, such as anaesthesiologists, perfusionists, and operating room staff may use shared models for coordinated planning. Secondary users include trainees and educators who use models for teaching and simulation support. Hospital administrators are described as indirect beneficiaries through potential workflow efficiency and planning support. Patients may benefit indirectly through improved explanation and shared understanding. The platform is not intended for independent patient self-use.



Impact and Economic Considerations

From a health systems perspective, the intervention should be considered an efficiency-enabling investment that supports safer, more effective, and more resource-efficient care delivery. Evidence-informed modelling indicates that immersive, patient-specific three-dimensional procedural planning can reduce operating room utilisation time through improved pre-operative strategy, early identification of anatomical variability, accurate implant sizing, strengthened team coordination, and fewer unanticipated intra-operative events, thereby contributing to lower complication risk and better clinical workflow. In high-value interventions, this approach functions as a financial risk-mitigation measure by supporting correct device selection and reducing avoidable wastage and repeat procedures at a relatively small proportion of total procedural cost. At the institutional level, improved planning and efficiency can enable greater service capacity, reduce downstream resource burden, and strengthen quality, training, and trust outcomes.

Ethics and Governance

Ethical and governance considerations are integral to PVCT. Clinical authority remains with licensed clinicians. The platform provides visualisation support only and does not generate automated diagnostic or therapeutic recommendations. Transparency is maintained by clearly defining AI's bounded role in segmentation and reconstruction. Data protection measures are described as aligned with national regulations and institutional policies, including secure handling and controlled access. Accountability frameworks distinguish

responsibilities between technology providers and healthcare institutions. Design choices emphasise scalability and workflow compatibility to support broader access, including in resource-constrained settings, subject to infrastructure and governance readiness.

Discussion

Successes

Reported strengths include automation of segmentation tasks, clinician-centred interface design, and compatibility with existing imaging workflows. Bounded AI scope and traceability to source images support governance acceptability. Deployment without changes to acquisition protocols supports feasibility.

Challenges

Challenges include variability in image quality, the need to build clinician trust in AI-assisted visualisation outputs, and change management within busy clinical settings. Infrastructure and connectivity constraints may affect immersive deployment modes.

Lessons Learned

Key lessons include maintaining constrained AI scope, involving clinicians early in design, prioritising workflow compatibility, and embedding governance safeguards from the outset. These lessons are relevant for implementers and policymakers considering similar visualisation technologies.



Conclusion

Immersive Medical Visualization platforms such as PVCT demonstrate how AI-enabled segmentation and reconstruction can transform existing imaging datasets into interactive spatial models that support clinical understanding, training, and communication. The platform reorganises rather than augments diagnostic data and functions as a decision-support visualisation layer. Early deployments and modelling analyses indicate operational feasibility and perceived workflow value in complex planning contexts. Clinical interpretation and responsibility remain fully with clinicians. Further independent evaluation is needed to quantify outcome and efficiency effects across settings.

References

Philips. (2025, May 28). Strengthening medical imaging capacity globally to provide better care for more people. Philips. <https://www.philips.com/a-w/about/news/archive/standard/news/articles/2025/strengthening-medical-imaging-capacity-globally-to-provide-better-care-for-more-people.html>


Wang, J. (2025). The need for evidence-based, outcome-focused medical imaging practice. National Center for Biotechnology Information (PMC). <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC12118759/>

ANI. (2026, January 1). India's hospital sector to sustain 11–12% growth, driven by insurance penetration and medical tourism: CareEdge. The Economic Times. <https://economictimes.indiatimes.com/industry/healthcare/biotech/healthcare/indias-hospital-sector-to-sustain-11-12-growth-driven-by-insurance-penetration-and-medical-tourism-careedge/articleshow/126284241.cms>

Healthcare Asia Magazine. (2025, December 12). India's diagnostic imaging market to grow 7% through 2035. Healthcare Asia Magazine. <https://healthcareasiamagazine.com/news/indias-diagnostic-imaging-market-grow-7-through-2035>

Fortune Business Insights. (2025). Medical imaging equipment market size, share & industry analysis, by type, by application & regional forecast, 2025–2032. Fortune Business Insights. <https://www.fortunebusinessinsights.com/industry-reports/medical-imaging-equipment-market-100382>

World Health Organization. (2025). Development of medical devices policies (2nd ed.). World Health Organization. Development of medical device policies

The image depicts three medical professionals, two men and one woman, wearing VR headsets and blue lab coats. They are engaged in a simulated surgical procedure on a large, glowing digital heart model. The scene is set in a futuristic, dimly lit environment with blue and red lighting. In the background, there are screens displaying anatomical diagrams and other medical data. The overall atmosphere is high-tech and focused on advanced medical training.

**AI Medical
Simulator (AIMS):
A Cost-Optimized
AI-Enabled Platform
for Strengthening
Clinical Skills
Training in Resource-
Constrained Settings**

Alagappan Karthikeyan



AI Medical Simulator (AIMS): A Cost-Optimized AI-Enabled Platform for Strengthening Clinical Skills Training in Resource-Constrained Settings

Alagappan Karthikeyan*

Abstract

The Artificial Intelligence Medical Simulator (AIMS – AI MeduSim) is a technology-enabled training platform designed to expand access to simulation-based clinical education in resource-constrained settings. Developed in India, the platform integrates multi-sensor simulation manikins, open-source simulation software, and real-time middleware to support adaptive, scenario-based training under instructor supervision. The system is designed to reduce capital and operating costs relative to conventional high-fidelity simulators while maintaining functional realism and instructional flexibility. Early institutional deployments in South India indicate improvements in training throughput, learner performance indicators, and faculty-led scenario customization. The model demonstrates potential for scalable adoption across medical, nursing, and allied health training institutions where cost and infrastructure barriers currently limit simulation-based education.


Background

Health workforce strengthening requires safe, repeatable, and competency-oriented training environments. Simulation-based education is widely recognized as an effective method for improving procedural skills, clinical reasoning, and team-based response capacity. However, access to advanced simulation infrastructure remains uneven, particularly across low- and middle-income settings, due to high equipment costs, proprietary software restrictions, and limited instructor control over training scenarios.

Reported barriers include high acquisition and maintenance costs of commercial simulators, limited adaptability of pre-programmed scenarios, and concentration of simulation resources in well-funded urban institutions. These constraints contribute to inequities in practical skills training across the health workforce pipeline.

National competency-based medical education frameworks increasingly emphasize applied clinical skills, structured assessment, and patient safety. However,

*Cobra Softwares Private Limited, Tamil Nadu, India



implementation capacity is often limited by infrastructure and affordability constraints. Technology-enabled, modular, and lower-cost simulation platforms may help close this gap when aligned with educational standards and supported by appropriate governance and quality assurance mechanisms.

Solution Overview

AIMS is a modular hardware–software simulation platform developed to support competency-based clinical skills training. The system combines sensor-enabled manikins, an open-source simulation engine, and a real-time communication layer that enables instructor-directed scenario control and adaptive physiological responses.

The platform is designed to support structured training across multiple domains, including emergency care, critical care, anesthesia, cardiology, and obstetrics. It emphasizes instructor-led simulation, scenario adaptability, and learner performance analytics.

The platform comprises four integrated core functional components: sensor-enabled simulation manikins capable of reproducing selected physiological parameters for realistic clinical training; scenario-based simulation software with instructor control interfaces that allow real-time adjustment of case progression and learning conditions; a real-time middleware layer that supports bidirectional communication between hardware

and software to ensure responsive and synchronized simulation behavior; and performance logging and analytics modules that capture learner actions and decision pathways to support structured debriefing, competency assessment, and continuous training improvement.

Technical Architecture

The platform uses a three-layer architecture consisting of a hardware, software, and integration layer working in coordinated operation. The hardware layer includes simulation manikins equipped with respiratory, cardiovascular, and movement sensors and actuators that reproduce selected physiological signals and responses, with a modular design that allows component replacement and configuration adjustments. The software layer is built on a customized open-source simulation toolkit that supports clinical scenario execution and instructor dashboards, enabling faculty to adjust physiological parameters, trigger scenario events, pause sessions for facilitated debriefing, and expand or locally configure scenario libraries. The integration layer comprises a real-time middleware component that enables low-latency bidirectional communication between sensors, actuators, and simulation software, while also supporting structured session data logging for post-session review and assessment.



AI and Analytics Functions

Machine learning components are applied in bounded, decision-support roles within the training environment, including scenario difficulty adjustment suggestions based on learner performance trends, natural language processing support for structured review of learner notes and decision logs, predictive modeling of physiological parameter changes following simulated interventions, and pattern recognition methods that flag potential learner knowledge gaps for instructor review. All AI-generated outputs are advisory in nature and remain subject to instructor validation, oversight, and final control.

Deployment and Implementation Model

Development and deployment follow a phased model that includes hardware prototyping, software customization, integration testing, faculty validation, and pilot implementation, followed by structured operational rollout. Typical deployment includes on-site installation of simulation manikins and instructor workstations, local or hybrid data processing based on connectivity conditions, scenario library provisioning with periodic updates, faculty training and certification in simulation facilitation, and implementation of structured learner assessment and debrief workflows. Offline-capable configurations are also supported to ensure functionality in low-bandwidth environments and resource-constrained settings.

Scale of Deployment

Pilot deployments have been conducted across multiple medical training institutions in South India, involving undergraduate and postgraduate learners and faculty facilitators. Reported usage includes regular scenario-based training sessions and curriculum-linked simulation exercises.


Planned scale-up targets include broader national deployment and regional expansion through institutional partnerships and training-of-trainers models.

Target Users and Beneficiaries

Primary users include medical and nursing students, postgraduate trainees, faculty instructors and simulation facilitators, and allied health and paramedical training programmes, all of whom use the platform to support competency-based clinical skills development. Indirect beneficiaries include patients and health systems through improved provider preparedness, procedural safety, and adherence to clinical protocols. The platform is designed to promote equitable access by enabling cost reduction, modular deployment, and local scenario customization aligned with regional disease patterns and training needs.

Impact and Measurable Outcomes

Preliminary pilot data from participating institutions indicate improvements in structured clinical assessment scores,



increased simulation training throughput per institution, reduced per-learner simulation training costs, higher reported learner confidence in managing acute care scenarios, and greater integration of simulation sessions into formal curricula. These findings are derived from pilot evaluations and require continued multi-site validation using standardized outcome and performance measurement frameworks.

Ethics, Data Protection and Governance

The platform incorporates governance and ethical safeguards consistent with responsible AI and educational data use principles, including instructor-controlled simulation with human oversight of all scenario decisions, separation of personally identifiable learner data from performance analytics where feasible, role-based access controls for training data, consent-based use of learner performance data for research, alignment with applicable national data protection regulations, and documented algorithm functions to support explainability in performance feedback. Scenario content undergoes clinical expert review prior to deployment, and governance structures include clinical advisory oversight, data governance review mechanisms, and quality assurance processes to monitor scenario accuracy and algorithm performance.

Implementation Lessons

Key implementation lessons indicate that early and continuous faculty co-design improves adoption and scenario relevance,

while structured instructor training is necessary for sustained utilization. Initial deployment is more effective when it prioritizes high-frequency, high-impact clinical scenarios, and when hardware, software, and faculty development are delivered as an integrated package rather than as standalone components. Establishing outcome measurement frameworks at the time of deployment supports ongoing evaluation and improvement, and developing local manufacturing and technical support capacity enhances scalability and long-term sustainability.

Conclusion

AIMS represents a cost-optimized, AI-enabled simulation platform designed to expand access to competency-based clinical skills training in resource-constrained settings. Early deployment experience suggests that modular architecture, open-source foundations, and instructor-centered design can improve affordability, adaptability, and educational relevance.

Further multi-site evaluation, standardized outcome measurement, and alignment with regulatory and accreditation frameworks will be important to support broader scale-up. The model illustrates how AI-enabled simulation tools, when implemented with appropriate governance and human oversight, can contribute to health workforce capacity strengthening and safer clinical practice.



References

- [1] World Health Organization. (2021). Global strategic directions for nursing and midwifery 2021–2025. WHO. <https://www.who.int/publications/i/item/978-92-4-002096-0>
- [2] International Society for Simulation in Healthcare. (2023). Standards of best practice: Simulation glossary. *Clinical Simulation in Nursing*, 75, 1-24.
- [3] Issenberg, S. B., McGaghie, W. C., Petrusa, E. R., Lee Gordon, D., & Scalese, R. J. (2023). Features and uses of high-fidelity medical simulations that lead to effective learning: A BEME systematic review. *Medical Teacher*, 37(1), 10-28.
- [4] Fink, L. D. (2013). *Creating significant learning experiences: An integrated approach to designing college courses* (2nd ed.). Jossey-Bass.
- [5] Ministry of Health and Family Welfare, Government of India. (2023). National Medical Commission guidelines on competency-based medical education. <https://www.nmc.org.in>
- [6] Berliner, D. C. (2022). Learning to work with data systems in teaching. *Journal of Teacher Education*, 68(1), 87-96.
- [7] Patel, R., & Kumar, A. (2024). Machine learning approaches for adaptive clinical simulation scenarios. *Indian Journal of Medical Education*, 15(2), 112-128.
- [8] Kohn, L. T., Corrigan, J. M., & Donaldson, M. S. (2000). *To err is human: Building a safer health system*. National Academy Press.
- [9] Burt, C. W., & Sisk, T. E. (2023). Adverse events in hospitals: A literature review. *Journal of Patient Safety*, 12(3), 143-156.
- [10] Cobra Softwares Private Limited. (2025). AIMS pilot program evaluation report: Learning outcomes assessment. Internal report, Coimbatore, India.
- [11] Association for Standardized Patient Educators. (2023). Simulation-based education standards and competency framework. <https://www.aspeducators.org>
- [12] Society for Simulation in Healthcare. (2023). *Best practices in healthcare simulation: A comprehensive review*. SSH Publications.
- [13] Kumar, V., Singh, P., & Desai, S. (2024). Cost-effectiveness analysis of simulation-based medical education in resource-limited settings. *Global Health Action*, 17(1), 8234567.
- [14] European Union. (2021). Ethics guidelines for trustworthy AI. <https://digital-strategy.ec.europa.eu/en/library/ethics-guidelines-trustworthy-ai>

Pre-generated Offline AI for Community Health Workers in Africa

Jeremy Lachal*

In West Africa, Community Health Workers (CHWs) serve as the primary point of entry to the health system for millions of people living far from formal health facilities. They play a central role in disease prevention, maternal and child health, immunization, and the management of common illnesses. However, CHW training remains insufficient due to structural constraints, including geographic dispersion, high workforce turnover, limited supervisory capacity, and the high logistical costs associated with in-person training. Although access to accurate and up-to-date medical knowledge is a key determinant of care quality, most CHWs operate without reliable reference materials. Digital health tools could help bridge this gap, but limited or absent connectivity makes internet-dependent solutions difficult to deploy at scale.

kHEALTH addresses this challenge through a frugal, sovereign, offline-first AI approach designed specifically for community health systems. Co-developed with the Ministries of Health of Senegal and Côte d'Ivoire and deployed through national health networks, the platform integrates three components: (1) a mobile e-learning system that operates entirely offline; (2) an offline AI assistant trained exclusively on nationally validated medical guidelines; and (3) a structured data collection module that supports supervision and public health decision-making.

The core innovation of kHEALTH is its pre-generation methodology, which differs from conventional generative AI models. Rather than producing answers dynamically through large cloud-based systems, the approach begins by identifying the most frequent and critical questions encountered by CHWs in routine practice. These questions are systematically mapped against official national protocols to pre-generate tens of thousands of validated question-answer pairs that cover the majority of frontline use cases. During use, a lightweight natural language classifier maps each query to the closest pre-validated response.

This architecture achieves approximately 92% response relevance while supporting clinical safety (answers grounded exclusively in validated national guidance), frugality (fully offline operation without repeated model calls), data sovereignty (local server deployment), localization (adaptation to language and literacy levels), and operational efficiency (very low marginal cost per additional user), alongside an open-science orientation.

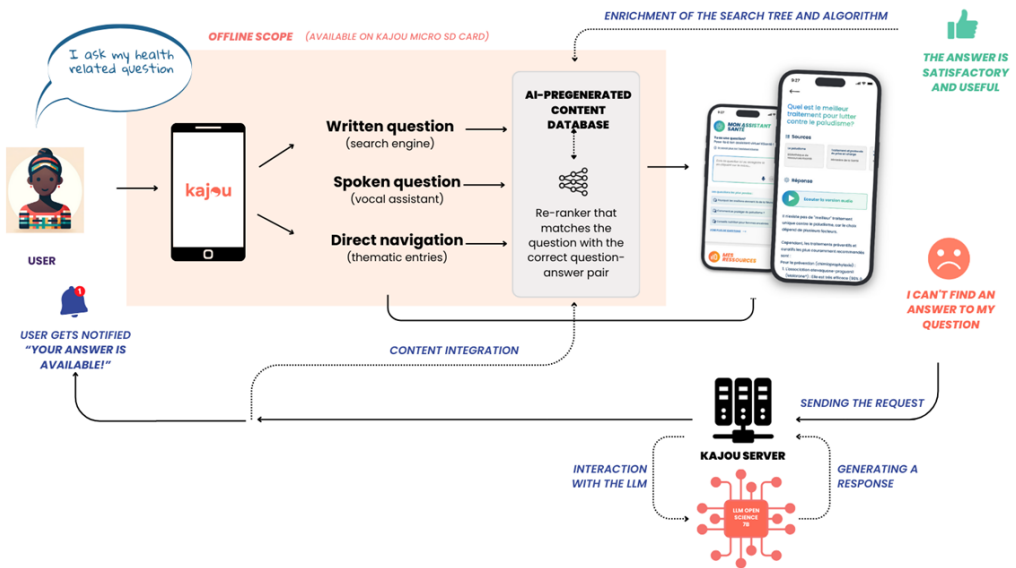
The solution is distributed via microSD card and runs directly on CHWs' personal mobile phones. It captures more than 70 anonymized usage indicators, which are synchronized when connectivity becomes available. These data enable Ministries of Health to monitor training progress, identify knowledge gaps, issue certifications, and communicate with CHWs in low-connectivity environments.

*Kajou, France

To date, more than 7,000 health professionals across Senegal, Côte d'Ivoire, and Togo are equipped with kHEALTH, with high adoption rates and early results showing an average 10% improvement in knowledge scores after several months of use. The next phase aims to scale deployment to 15,000 CHWs nationwide in Senegal and expand

the pre-generation framework to include structured diagnostic support pathways and broader local language coverage. Combined with rigorous impact evaluation, this scale-up seeks to demonstrate that pre-generated, offline AI can serve as a safe, equitable, and scalable model for knowledge delivery in primary health care systems.

Figure 01: Workflow of KHealth offline AI system



AI-Enabled Diabetic Retinopathy Screening in India's Primary Healthcare System

Anshul Chauhan¹, Luke Vale², Mona Duggal³

Diabetic retinopathy (DR) is a leading cause of preventable vision loss, and timely screening with appropriate referral is critical to avoid irreversible damage. In India, DR affects an estimated 18% of people with diabetes, compared with a global average of 9.3%, indicating a substantial and growing screening burden. This increasing demand places pressure on public health systems already constrained by limited specialist availability, cost barriers, and service capacity. Despite its importance, DR screening (DRS) remains insufficiently integrated into primary health care, particularly in rural settings. Current screening pathways are largely specialist-led and concentrated in secondary and tertiary facilities, resulting in limited population coverage. Non-AI screening approaches that depend on human graders are resource-intensive and difficult to scale in public-sector contexts, while digitization alone does not reduce specialist dependency or enable point-of-care screening at scale.

The intended use case is AI-enabled DRS as a decision-support tool within primary care to support early detection of referable DR and timely treatment. The system is designed to assist trained non-ophthalmologist health workers in conducting screening under routine public health conditions, while final diagnosis and treatment decisions remain with specialist clinicians. The deployment context spans both

clinical and public health workflows within existing primary care systems.

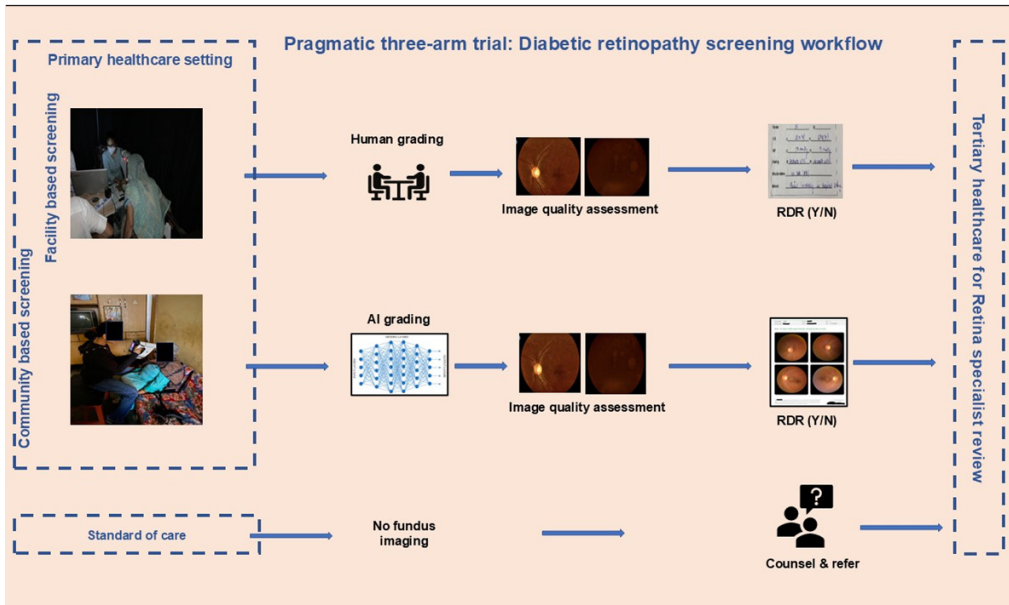
This case describes a phased, real-world deployment of AI-enabled DRS in rural Punjab, India (2023–2024), using a three-arm pragmatic trial design (n=600). The study compared facility-based human-led screening, offline AI-assisted community-based screening, and standard referral care. Retinal fundus images were captured by trained non-ophthalmologist health workers. Facility-based screening used a tabletop fundus camera, whereas community-based screening used a portable, non-mydriatic camera with an embedded offline AI system (Medios). The AI model applies convolutional neural network-based deep learning to analyze non-mydriatic retinal images, incorporating automated image-quality assessment and probabilistic classification of referable DR (yes/no) at the point of care. The algorithm was developed and validated using clinically annotated real-world fundus image datasets with built-in quality control checks to support reliable automated analysis. The system was integrated into routine screening and referral workflows and performs on-device inference without cloud dependence. Outputs are recorded within standard screening and referral documentation to maintain traceability and clinical accountability.

¹Postgraduate Institute of Medical Education and Research, India ²London School of Hygiene and Tropical Medicine UK ³Indian Council of Medical Research – National Institute for Research in Digital Health and Data Science (ICMR- NIRDHDS)

Diagnostic performance for referable DR was strongly influenced by image quality. Excluding poor-quality and misclassified images improved specificity from 85.05% to 91.0% and overall accuracy from 86.07% to 90.83%, with increased agreement (κ 0.55 to 0.57; $p < 0.001$). Ungradable images primarily associated with cataract, advanced age, or suboptimal image capture were a major source of misclassification and unnecessary referrals. These findings highlight the importance of systematic image-quality assessment and human review in ensuring appropriate referral decisions in real-world AI-assisted screening programmes.

Sustainable integration of AI-enabled DRS requires supportive policy and regulatory frameworks, including standardized workforce training, post-deployment performance monitoring, and procurement models that account for AI-related costs within public screening programmes. Scalability depends on alignment with existing clinical workflows and national health programmes. Key implementation gaps include variability in image acquisition quality, workforce training requirements, and the need for structured post-deployment monitoring and evaluation frameworks.

Figure 02: Three arm diabetic retinopathy screening workflow using human and AI grading



AI Diagnostic Tool for Ear Pathology (Primary Care ENT Screening)

Rakesh Srivastava, Shashidhar TB*

India faces a critical shortage of ENT specialists, most of whom are concentrated in urban centers, leaving rural populations with limited access to basic ear care. This imbalance contributes to delayed diagnosis and treatment of conditions such as chronic suppurative otitis media (CSOM), acute otitis media, and other ear pathologies. Children are particularly affected, as untreated infections can lead to hearing loss and developmental delays. Primary care providers including nurses, auxiliary nurse midwives (ANMs), and general physicians frequently encounter these cases but often lack the specialized training required for accurate otoscopic diagnosis.



Current solutions are insufficient to address this gap. Telemedicine models often introduce delays, training programs have limited reach, and many international AI tools are not optimized for Indian population characteristics, disease patterns, or low-resource imaging devices. OtoSensAI is designed to address this need as a large-scale AI diagnostic support tool trained specifically on Indian ear pathology data. It supports multiple use cases, including primary care decision support, remote screening in underserved areas, telemedicine consultations, medical education, and pediatric screening. The system classifies twelve categories, including normal ear, multiple infection

types, wax impaction, foreign bodies, and “not ear image” classifications for quality control. It is intended for deployment across primary health centres, school screening programs, and guided patient self-assessment contexts.

The model is trained on 12,000 curated images selected from a repository of 80,000 images collected over twelve years. Poor-quality images and diagnostically ambiguous cases requiring specialist judgment were excluded. The dataset was split into training, validation, and testing sets using an 80–10–10 ratio. The system is optimized to process low-resolution images (640×480 pixels), compatible with smartphone-based otoscope attachments commonly used in field settings. Multiple architectures including ResNet50, Vision Transformer, MobileNetV3, and Inception V3 were evaluated before selecting a Squeeze-Excitation Inception Network (SEIN). This architecture analyzes images at multiple scales to capture both fine structural details and broader anatomical patterns, while its squeeze-excitation mechanism emphasizes diagnostically relevant features and reduces background noise.

The platform is deployed on cloud infrastructure and delivers results within 2–3 seconds. Outputs include a primary classification with confidence level (high, medium, or low), ranked alternative

*Voxmedai Healthtech private Ltd, Uttar Pradesh, India



predictions, brief clinical descriptors, and a clear disclaimer that the system provides decision support and does not replace professional diagnosis. Phase 2 validation is underway using 200 ear examinations, comparing three pathways: AI-assisted diagnosis by nurses, gold-standard ENT examination, and blinded image interpretation by a second ENT specialist. This evaluation assesses both algorithmic accuracy and the feasibility of image capture by non-specialists.

Image quality remains the principal operational challenge, influenced by lighting conditions, operator technique, and patient cooperation. Standardized capture protocols and real-time quality feedback mechanisms are being developed to address this constraint, supported by pilot testing demonstrating that non-specialists can obtain usable images with appropriate guidance. Ethical safeguards focus on preventing over-reliance on AI outputs, particularly in settings where specialist confirmation is limited. Data protection measures include cybersecurity controls, informed consent

processes, and regulatory compliance. Regulatory uncertainty remains a barrier as India's AI medical device framework continues to evolve, highlighting the need for early engagement with national regulatory and research bodies. Enablers for scale include clear regulatory pathways, reimbursement alignment under public insurance schemes, integration with national ear and hearing care programmes, and clarified liability frameworks.

Sustainable deployment depends on multi-stakeholder collaboration across government agencies, academic validation partners, private technology providers, and community organizations. Although the system supports continuous model improvement, this capability requires strong governance and quality control mechanisms. Long-term success will depend on positioning the tool as a rigorously validated response to a defined service gap, implemented with clinical oversight and refined through stakeholder feedback, with the primary objective of expanding access to quality ear care in underserved populations.

Swasth-Derma: AI-Powered Pre-Medical Dermatology Triage for Public Health

Vatshal Pravinbhai Patel, Nishi Seth, Ananaya Pal*

Skin conditions are among the most common reasons for primary care visits worldwide, yet access to dermatology specialists remains limited across much of the Global South. In many communities, individuals rely on informal advice, pharmacies, or delayed consultations, particularly for rashes and inflammatory skin disorders. This creates two major gaps: delayed clinical care for potentially infectious conditions and missed opportunities to reduce disease transmission through early referral and basic hygiene measures. Conventional non-AI approaches such as static information leaflets or symptom checklists are insufficient because they cannot account for visual variability, multi-angle presentation, and real-world image capture through mobile phones. Meanwhile, many dermatology AI systems focus on multi-class diagnostic prediction, which introduces safety, regulatory, and deployment challenges, especially in low-resource settings.

Swasth-Derma addresses this gap by positioning dermatology AI as a pre-medical triage and public health guidance tool rather than a diagnostic system. Its purpose is to help patients and frontline health workers decide when to seek clinical care and how to reduce potential spread, particularly for infectious skin conditions, using one to three smartphone images per case.

The system follows a simple, clinically aligned workflow in which a patient or health worker captures one to three images—close-up, angled, and contextual views using a mobile device. Images are processed through a lightweight deep-learning pipeline that performs case-level binary risk assessment, classifying presentations as likely infectious or non-infectious. The modeling strategy uses a two-stage design. First, each image is transformed into a compact feature representation using a pretrained dermatology foundation model that remains frozen to ensure computational efficiency and model stability. Second, an attention-based Multiple Instance Learning classifier aggregates features across images from the same case to generate a single case-level prediction. This approach reflects real-world clinical viewing patterns while remaining suitable for mobile and low-resource deployment.

Data are split at the case level using a 60/20/20 train, validation, and test structure, comprising 1,401 training cases (3,026 images), 467 validation cases (996 images), and 468 held-out test cases (1,019 images). On the test set, the system achieved 80.2% case-level accuracy for infectious versus non-infectious triage. Outputs are expressed as calibrated risk probabilities rather than diagnoses. Based on configurable thresholds, users receive one of two action-oriented responses: a prompt recommending clinical evaluation when infection risk is likely, or hygiene and

*HawkFranklin Research, West Bengal, India

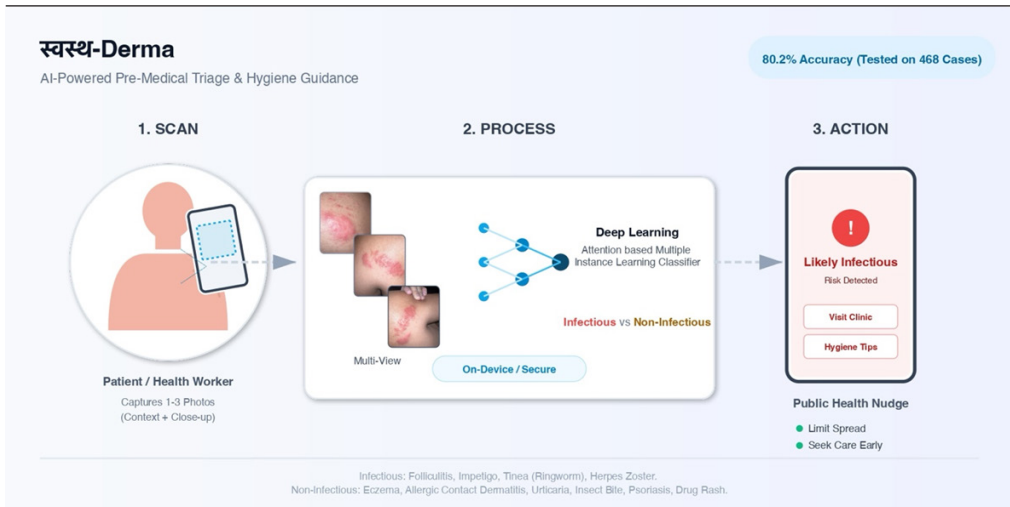
prevention guidance such as covering rashes, avoiding close skin contact, and not sharing personal items when risk appears lower.

The system supports on-device or secure inference, minimal data retention, and explicit user disclaimers indicating that it is not intended for diagnosis or emergency use. These safeguards support responsible deployment in telehealth, community health, and primary care workflows.

Key challenges include ensuring performance across diverse skin tones and population groups, maintaining user trust through

transparency, and navigating regulatory frameworks applicable to non-diagnostic AI tools. Ethical priorities include privacy protection, prevention of over-reliance on automated outputs, and clear communication of system limitations. Enablers for scale include policy recognition of pre-medical triage as decision support rather than diagnosis, supportive regulatory pathways, and integration into national digital health strategies. With appropriate governance, this approach illustrates how lightweight, privacy-preserving AI tools can strengthen early care-seeking behavior and public health outcomes in low-resource settings.

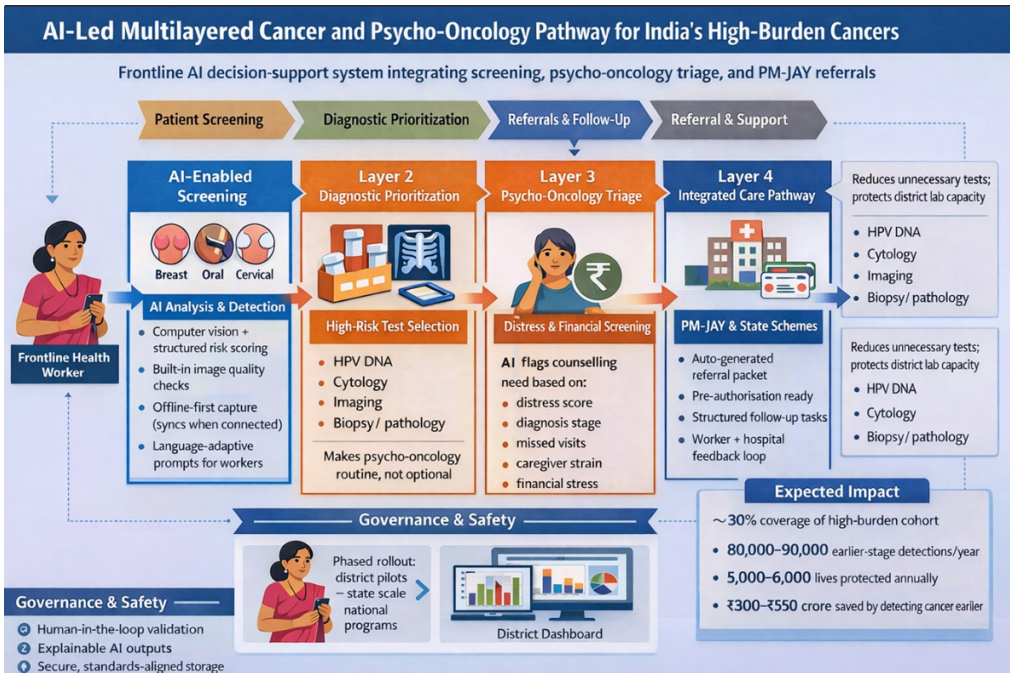
Figure 03: Swasth - Derma Technical Pipeline



AI-Led Multilayered Cancer and Psycho-Oncology Pathway for India's High-Burden Cancers

Saswata Banerjee*

Figure 04





Breast, oral and head-and-neck, and cervical cancers together account for approximately 35%–40% of India's estimated 1.5 million new cancer cases each year. Although these cancers are largely detectable at the primary-care level, nearly 70% are still diagnosed at advanced stages. This results in avoidable morbidity and mortality, places significant strain on district and

tertiary hospitals, and substantially increases public expenditure under national insurance schemes, with late-stage treatment costs several times higher than those for early-stage care.

At the same time, psycho-oncology support remains limited. An estimated 40%–50% of

*Aurigene Pharmaceuticals



patients and caregivers experience significant psychological or financial distress, which is associated with higher treatment default rates by nearly 30% in some settings. Distress screening is rarely conducted routinely and is often addressed only after crises occur. Frontline health workers operate under wide variations in literacy, language, and connectivity conditions and typically have limited access to structured decision support. Current non-AI workflows depend on manual screening methods such as visual inspection with acetic acid (VIA) and oral examination, paper-based referrals, and fragmented digital tools. These approaches do not scale effectively: referral leakage may reach about 50%, diagnostic services are burdened by low-yield testing, and connectivity constraints affect a large proportion of rural facilities.

The proposed solution is a unified, AI-enabled frontline application that integrates cancer screening, diagnostic prioritization, psycho-oncology triage, and insurance-linked referral management within a single structured workflow. The intended deployment context is public-sector screening and referral under national and state cancer programmes and insurance schemes, spanning community settings, district facilities, and medical college hospitals.

The system uses a layered design. First, AI-assisted screening applies computer vision models to analyze oral cavity images, nurse-captured VIA cervical images, and structured breast or axillary symptom inputs. Automated image-quality checks filter out non-diagnostic images, helping ensure referral-grade outputs despite variable lighting and device conditions, and enabling non-specialist workers to generate clinically useful findings

at the point of care. Second, diagnostic prioritization uses machine-learning risk scores that combine image outputs, symptoms, demographic factors, and risk indicators to prioritize confirmatory tests such as HPV DNA testing, cytology, imaging, or biopsy, thereby preserving limited diagnostic capacity by focusing on high-probability cases. Third, psycho-oncology triage is embedded directly into the workflow through brief, language-adapted distress and financial stress screening tools. AI-assisted scoring flags counseling and support needs based on distress levels, cancer stage, missed visits, caregiver burden, and social risk factors, enabling earlier and more proactive intervention. Fourth, referral and follow-up functions generate physician-ready referral packets aligned with public insurance workflows and programme requirements, while structured follow-up tasks and SMS-based feedback loops keep both frontline workers and oncology centers informed, reducing delays and loss to follow-up. Key implementation challenges include variability in image quality in field conditions, large-scale workforce training requirements, consent and data privacy management, potential bias arising from low-diversity datasets, and the need to ensure that AI outputs remain advisory rather than prescriptive. Important enabling factors include alignment with national screening priorities, integration with public insurance and digital health workflows, linkage with district cancer programmes and registries, and increasing acceptance of AI-assisted decision support. With appropriate policy support, training, and governance frameworks, this approach can make more psycho-oncology screening and referral more systematic and data-driven, support earlier detection at scale, and reduce avoidable clinical and financial burden in public oncology care.

Precision Beyond Genes with QPOP: An AI – Guided Oncology Approach

Anagha Girish*

Cancer treatment generally follows standardized protocols, yet patients with the same cancer type often show markedly different responses to therapy. In practice, oncologists frequently rely on trial-and-error drug selection, which can be time-consuming, costly, and associated with increased toxicity and adverse effects. Genomics-based precision medicine has improved targeting but does not always reliably predict actual drug response at the individual patient level. QPOP (Quadratic Phenotypic Optimization Platform) is an AI-enabled precision oncology platform designed to address this gap by predicting patient-specific drug combinations using phenotypic response data. The approach aims to shorten experimental cycles and reduce toxicity risk while operating alongside existing clinical workflows. It is intended for use in oncology settings, including hospitals, cancer centres, and academic medical institutions, and is applicable in both high-resource and resource-constrained environments. The system supports clinical decision-making and does not replace oncologist judgment.

QPOP is based primarily on phenotypic response data derived from patient tumor cells, including measurements of cell viability and drug sensitivity. Secondary inputs include basic clinical metadata and curated libraries of approved oncology drugs. Because the platform relies on direct measurement of cellular response, the input data typically demonstrate a high signal-to-noise ratio. The analytical core

is a quadratic optimization-based machine learning model that combines predictive modeling and optimization techniques. This approach is selected for its ability to learn effectively from limited datasets and to model drug-drug interaction effects.

The workflow begins with acquisition of patient-derived tumor cells, which are prepared for ex vivo testing in a manner that preserves clinically relevant phenotypes. A targeted panel of approved oncology drugs and dose combinations is then tested to generate quantitative phenotypic response profiles. These data undergo normalization, validation, and quality control before AI analysis. The quadratic optimization engine models drug effects and interactions and predicts high-performing, previously untested drug combinations. The system then produces ranked and interpretable recommendations on drug combinations and dose ranges, which are reviewed by clinicians and incorporated into treatment planning.

Outputs include ranked combination options, suggested dose ranges, comparative performance indicators, and structured decision-support insights. The platform follows a human-in-the-loop model in which clinicians review and validate all recommendations. Reliability is supported through experimental validation processes. The modular architecture includes a data ingestion layer, an AI optimization engine, and

*MES College of Arts, Commerce and Science, Karnataka, India

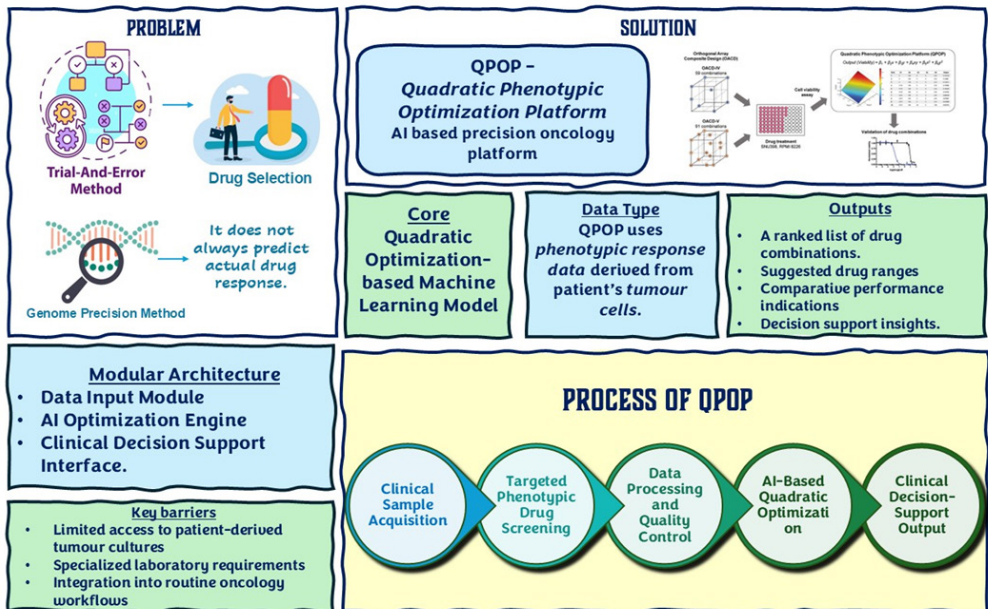
a clinical decision-support interface designed to fit hospital oncology workflows and tumor board deliberations. The platform is used in conjunction with standard diagnostic testing and clinical expertise and has undergone validation through clinical pilot studies.

Key barriers to broader adoption include limited availability of patient-derived tumor cultures, the need for specialized laboratory infrastructure, and challenges in integrating phenotypic

testing into routine oncology workflows. Wider uptake also depends on multi-centre clinical validation and clinician familiarity with AI-assisted decision-support tools. Ethical and governance considerations emphasize patient safety, transparency, de-identification of data, and strict clinician oversight. Sustainable deployment will require clear regulatory guidance, structured pilot support, and alignment with responsible AI and national digital health frameworks.

Figure 05: QPOP AI platform for phenotypic drug optimization

Precision Beyond Genes with QPOP: An AI – Guided Oncology Approach



Blockchain-Enabled, AI-Driven EMR for Health Equity Across the Global South

Chun Phillip, Atika Kumar, Lohith Deshpande*

Across much of the Global South, healthcare delivery is constrained by fragmented, paper-based, or institution-bound medical record systems. These limitations disrupt continuity of care, lead to repeated or unnecessary testing, reduce patient engagement, and disproportionately affect underserved populations and frontline providers. Many existing digital record solutions are server-intensive, costly, and poorly suited to low-resource or low-connectivity environments. Apollo Health is designed to address this gap through a secure, interoperable, and patient-controlled electronic medical record (EMR) platform that can operate across facilities, regions, and national boundaries. Its primary use case is a patient-centered EMR system that supports frontline clinics, community health workers, and small hospitals in emerging health systems.

Apollo Health is a blockchain-enabled, AI-assisted EMR platform built for equitable and distributed deployment. Patient records are secured on a permissioned blockchain ledger that creates immutable audit trails while enabling granular, patient-controlled consent for data access and modification. This architecture supports privacy, transparency, and cross-institution interoperability without relying on centralized data ownership. The platform integrates structured clinical data such as

demographics, diagnoses, medications, and laboratory results—with unstructured clinical notes and optional patient-generated data, including home vitals and wearable device inputs.

AI components include natural language processing to support clinical documentation, rule-based triage suggestions, and risk flagging for care coordination. All AI outputs are advisory and supported by explainability features to strengthen clinical trust and safe use. A cloud-light, offline-first design allows facilities with intermittent connectivity to continue operations locally and synchronize records once connectivity is restored. The system provides standards-based APIs to enable integration with existing health information systems, laboratories, pharmacies, and public health registries. Deployment follows a phased rollout model, beginning with pilot implementations in India, Nepal, and the Philippines, with planned expansion across clinics, districts, and national programmes.

Key challenges include regulatory variability across jurisdictions, differences in digital literacy among users, and early-stage provider hesitancy toward blockchain-based infrastructure. Ethical and governance priorities include protection of data privacy, meaningful informed consent, and prevention of over-reliance or misuse of

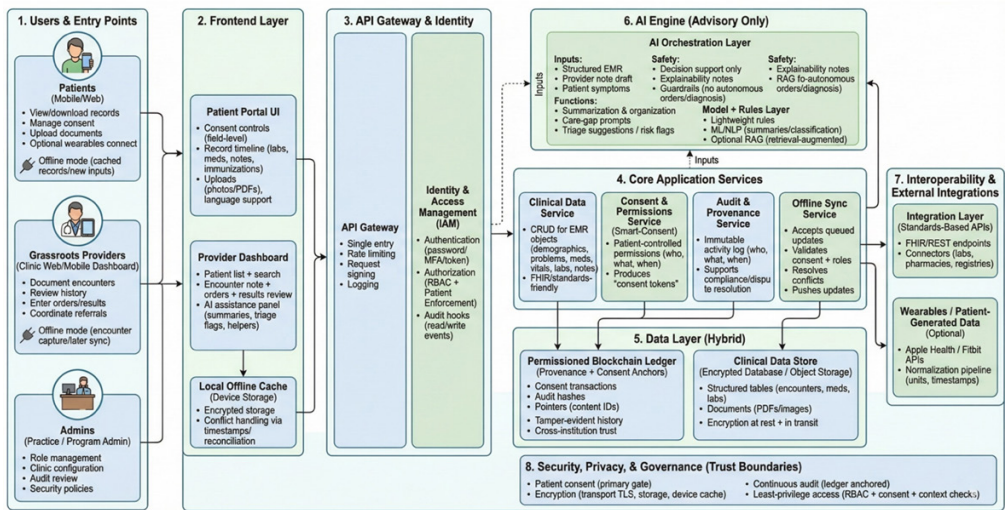
*Apollo Health, Illinois, United States

AI-generated suggestions. These risks are mitigated through patient-controlled consent frameworks, encryption, advisory-only AI functions, and alignment with applicable data protection regulations. Enabling factors include growing policy emphasis on interoperability and patient data access, increased acceptance of

AI-supported clinical workflows, and demand for affordable digital health infrastructure. With supportive regulation, public-private partnerships, and continued clinician co-design, the platform represents a scalable and responsible model for inclusive, AI-enabled digital health systems in the Global South.

Figure 06

Apollo Health - System Architecture (Patient-Owned Consent, AI-Assisted, Offline-Capable)



Advanced Data Analytics for Public Health Action and Research Venture (ADARV) – A Public Health Data Warehouse

Yogita Chaudhary¹, Sharan Murali¹, Rakshith Ramesh^{2*}

Infectious disease outbreaks remain a persistent public health challenge in India, where surveillance data systems are fragmented across programmes, formats, and jurisdictions. Outbreak datasets frequently use heterogeneous variable names and structures, requiring extensive manual cleaning and harmonization before analysis. This leads to delays in epidemiological assessment and increases the workload of field health personnel. Limited standardization also reduces comparability across outbreaks and weakens the ability to integrate findings with other public health datasets. Traditional non-AI approaches, such as rule-based mappings and manual harmonization, are error-prone and difficult to scale, while AI-based approaches raise legitimate concerns regarding data security and accuracy if not properly governed.

The Advanced Data Analytics for Public Health Action and Research Venture (ADARV) is a public health data warehouse developed in India that incorporates AI-enabled tools to standardize, analyze, and generate reports from outbreak datasets. The platform supports rapid analysis of field data to enable timely, evidence-based decision-making and also allows controlled exchange of anonymized datasets to support research and secondary analysis. ADARV is designed to reduce the time between data collection and public health action while maintaining governance and traceability.

Users upload datasets in Excel or comma-separated value formats. ADARV processes these files through an AI-assisted entity-mapping pipeline that cleans text, separates concatenated terms, and applies terminology-aware fuzzy correction aligned with SNOMED standards. It then performs entity extraction using biomedical natural language processing tools and encodes terms with domain-specific language models. Similarity search is conducted using a high-speed vector index and database-side vector search, after which results are deduplicated, merged, and ranked to return standardized SNOMED CT concepts and linked FHIR resources for each dataset.

The platform further applies statistical and epidemiological analytics using established Python and statistical libraries to generate summary tables and regression outputs. It produces geospatial spot maps and time-based epidemic curves using geospatial and visualization toolkits. Outputs are compiled into structured PDF reports that integrate person, place, and time analyses. All mappings and transformations are logged and auditable to ensure explainability and traceability. Datasets are anonymized before optional sharing through the platform's governed exchange layer. User access is verified through Aadhaar-based authentication, and data owners retain full control over whether and how their datasets are shared. Through standardized structuring and governed reuse, the platform operationalizes

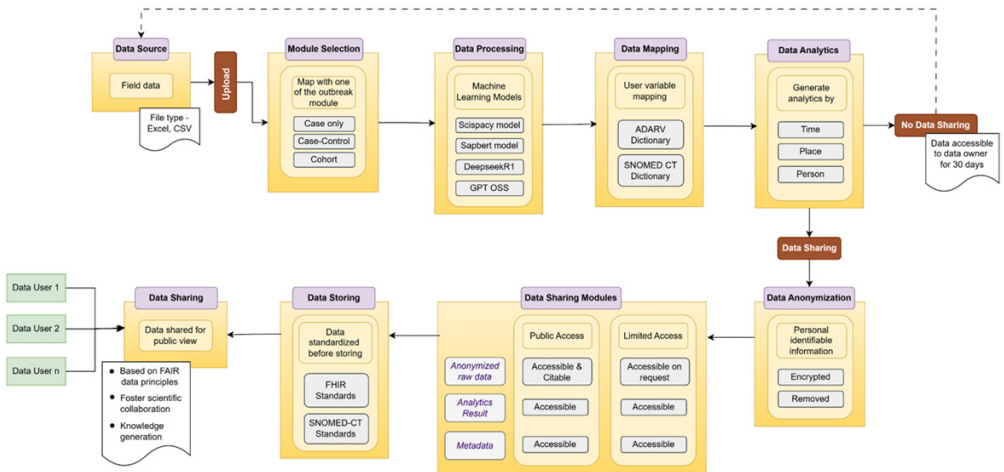
*¹ICMR National Institute of Epidemiology; ²Centre for Data for Public Good

FAIR (Findable, Accessible, Interoperable, Reusable) data principles within a public health surveillance context.

The minimum viable product is currently undergoing security audit prior to phased user onboarding, beginning with field epidemiologists working on outbreak investigations. Iterative user feedback will guide functional refinement and support scaling across geographies and public health programmes. The platform is designed for interoperability and future integration with existing public health information systems.

Key challenges include handling the wide variability of user-uploaded datasets and maintaining robustness across diverse outbreak contexts. Model tuning has been informed by practitioner experience in outbreak investigation, and legal oversight supports compliance with applicable Indian data protection requirements. With appropriate governance and scale-up support, the platform provides a practical pathway toward FAIR-aligned data practices that reduce epidemiological workload, accelerate the data-to-action cycle, and strengthen outbreak preparedness and response capacity across resource-constrained settings.

Figure 07: System architecture of ADARV illustrating the end to end workflow from outbreak data ingestion to AI driven standardisation, analytics generation, anonymisation and governed data sharing.



Moskeet: Climate-Responsive AI & IoT Surveillance for Vector-Borne Disease Control

Satish Cherukumalli*

Moskeet addresses the growing burden of mosquito-borne diseases in the Global South, where climate change, urbanization, and human mobility are reshaping vector ecology faster than many public health systems can respond. Government control programmes continue to rely largely on periodic manual surveys and blanket spraying, with limited neighbourhood-level intelligence on vector species, breeding sites, and climate-driven risk. This often results in delayed and inefficient interventions, overuse of insecticides, and accelerating resistance—particularly in dense informal settlements and climate-vulnerable areas.

Moskeet is an AI- and IoT-enabled vector surveillance and decision-support platform that generates real-time, hyperlocal risk intelligence for local government authorities. Smart electronic traps deployed across communities continuously capture mosquito wing-beat signatures using infrared sensors. These signals are processed in the cloud using machine-learning models trained on large, labeled wing-beat datasets to classify species, sex, time, and location. Outputs are delivered through role-based dashboards, field-team mobile applications, and secure APIs that integrate with municipal command-and-control systems and existing GIS platforms.

To translate raw detections into climate-aware risk estimates, Moskeet integrates trap data with multiple environmental and demographic data layers. Satellite imagery is analyzed

to derive breeding-site proxies, vegetation indices, rainfall patterns, ambient temperature, humidity, and land-surface temperature. These variables, combined with population density, are aggregated into fine-grained spatial grids (approximately 100 m × 100 m). Weighted combinations generate baseline risk scores for each grid cell. As trap data are continuously updated, these baseline scores are recalibrated using observed vector abundance by species and sex, producing dynamic entomological risk indices. Machine-learning models trained on historical climate, population, and trap datasets are then used to forecast short-term changes in vector risk.

Results are presented as intuitive risk maps, hotspot rankings, and prioritized intervention recommendations, such as targeted larval source management, selective fogging, or focused community outreach. Basic explainability indicators identify the main drivers behind predicted risk levels. Species-level intelligence further supports rotation of insecticides and more selective use of adulticides and larvicides, helping reduce resistance pressure and avoid unnecessary spraying.

Pilot deployments across communities in India covering approximately 325,000 people have demonstrated operational and epidemiological benefits. Compared with traditional larval surveys and manual trapping methods, surveillance operating costs decreased by

*Trakitnow Technologies PVT Ltd, Hyderabad, India

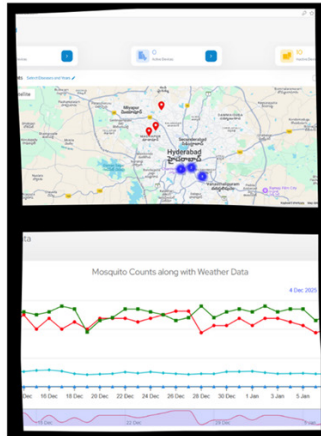
roughly 85% while enabling daily, geo-tagged species monitoring. Improved alignment of insecticide choice and timing with local vector patterns reduced fumigation expenditure by about 20%. Reported outcomes also include up to 60% reductions in adult mosquito density, along with associated declines in dengue and malaria case reports. The platform is now being expanded through phased rollouts with local health departments.

Key challenges include maintaining trap performance and data quality, minimizing bias in data-sparse areas, and safeguarding the privacy of geolocated data. Responsible scale-up will depend on supportive digital health and data-governance policies, along with integration into national vector-borne disease control programmes.

Figure 08: Moskeet AI enabled mosquito surveillance system



Moskeet Device



Dashboard



Risk Analysis

AI-driven Decision Support for Ayurveda: A Retrieval-Augmented Generation (RAG) Approach

Naman Goyal, Rinki Kumari, Sahil Singh*

Traditional medicine systems such as Ayurveda play a significant role in health care delivery in India. Clinical decision-making in Ayurveda is grounded in practitioner expertise and classical texts; however, translating this extensive and largely unstructured knowledge base into actionable, context-aware support at the point of care remains a persistent challenge. Most existing digital systems in this domain focus on administrative reporting and documentation, with limited capability to provide real-time clinical decision support aligned with the multidimensional reasoning that characterizes Ayurvedic practice. This highlights the need for structured, digitally integrated tools that complement—rather than replace—clinical expertise.

These constraints are especially pronounced across many Global South settings, where high patient volumes, workforce limitations, and uneven access to specialized knowledge reduce the effectiveness of non-AI approaches such as manual reference use, static protocols, or basic digitization. The limited availability of AI-enabled tools tailored to traditional medicine further restricts consistent and context-sensitive decision support, despite increasing policy-level efforts to integrate traditional systems into mainstream health services.

To address this gap, an AI-enabled clinical decision-support tool is proposed that combines curated Ayurvedic knowledge with de-identified clinical inputs drawn from the Ayush Hospital

Management Information System (AHMIS) of the Ministry of Ayush, Government of India. Its primary use case is point-of-care decision support for Ayurvedic practitioners, assisting with diagnostic reasoning, management approaches, formulations, lifestyle guidance, and follow-up planning. Secondary use cases include improving documentation quality and consistency, aligning records with standardized classifications, and supporting administrative and public health analytics, while preserving practitioner autonomy in all clinical decisions.

The system is built on an expert-validated knowledge base derived from Ministry of Ayush standard treatment guidelines, classical texts such as the Charaka Saṃhitā, Suśruta Saṃhitā, and Aṣṭāṅga Hṛdaya, and disease-specific protocols mapped to WHO ICD-11 TM2 classifications. These sources are digitized, normalized, and structured to support clinical accuracy and interoperability. De-identified AHMIS patient data are incorporated to reflect real-world clinical scenarios while maintaining privacy and regulatory compliance.

The technical architecture combines natural language processing with a lightweight generative model and a Retrieval-Augmented Generation (RAG) framework to produce context-aware, source-linked recommendations. The RAG approach grounds outputs in validated textual sources, reducing risks associated with standalone large language models, such as hallucination, weak explainability, and lack of

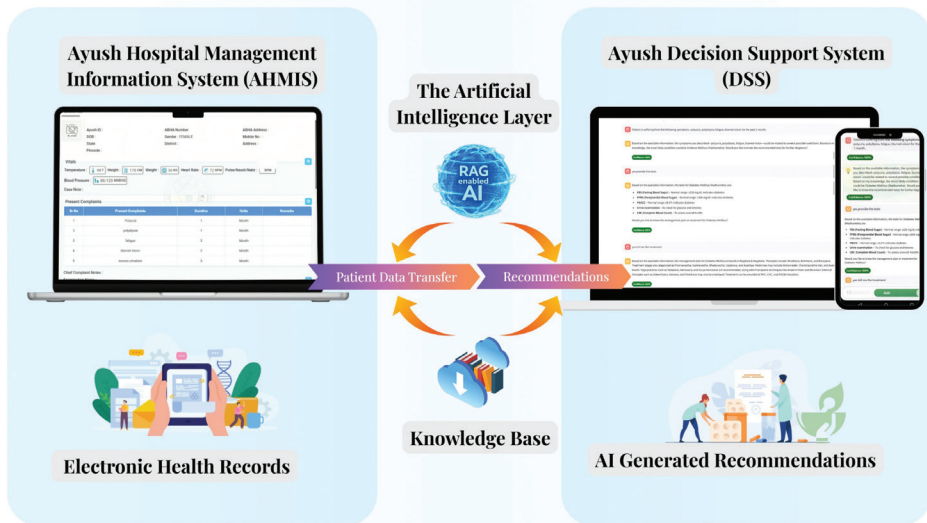
*Ministry of Ayush, Government of India

provenance. Clinical inputs trigger semantic retrieval from the curated knowledge base, followed by structured generation of diagnostic considerations, management options, dietary and lifestyle advice, and follow-up suggestions, each linked to source references.

Explainability and reliability are supported through explicit source attribution, deterministic generation parameters, rule-based constraints, and a human-in-the-loop design that enables practitioners to review, modify, or override system suggestions. The tool is designed for non-intrusive integration into existing AHMIS workflows and is being developed through a proof-of-concept and pilot-based pathway, with phased deployment supported by lightweight models and modular, cloud-compatible architecture.

Key challenges include variability in data quality, incomplete digitization of classical literature, and inconsistent clinical documentation practices. Broader adoption depends on robust clinical validation, practitioner trust, ethical safeguards, and strong data privacy protections. Important enabling factors include supportive regulatory frameworks, standardized evaluation approaches, targeted capacity building, and alignment with national digital health initiatives such as the Ayushman Bharat Digital Mission. Together, these elements can support scalable, ethical, and clinically meaningful adoption of AI-enabled decision support in traditional medicine settings across the Global South.

Figure 09: From Ayush EMR to AI driven decision support - A RAG enabled system for generating personalized clinical recommendations



AIIMS Delhi Ransomware Attack: CyberPeace Foundation's Insights for Forging Secure and Resilient Digital Health Ecosystems in India

Vineet Kumar¹, Nilakshi Jain¹, Shubham Kolaskar^{2*}

The November 2022 ransomware attack on AIIMS Delhi, India's apex public health-care institution, exposed a critical vulnerability in national digital health infrastructure: the limited ability of reactive, legacy security controls to withstand advanced and state-linked cyber threats. The breach reportedly encrypted approximately 1.3 terabytes of data and forced a prolonged reversion to manual workflows, disrupting services and delaying patient care for weeks. Conventional non-AI security tools—such as static firewalls and signature-based detection systems—are increasingly inadequate because they cannot process high-volume, real-time telemetry or reliably detect previously unseen (“zero-day”) attack patterns. The intended use case is the development of resilient digital public health infrastructure that protects continuity of care during cyber incidents. The relevant system scope includes hospital information systems, endpoint devices, and cloud backup environments within high-risk clinical settings. The objective is to close the gap between rapid digital health adoption under national programmes and comparatively weak security postures, thereby safeguarding operational continuity and patient safety during active attacks.

The proposed approach is a proactive, AI-enabled security architecture that combines forensic intelligence with real-time defensive capabilities. The system ingests telemetry from hospital information systems, network

traffic, and endpoint devices. Machine learning models are applied for behavioral anomaly detection to identify suspicious patterns and potential zero-day threats that bypass rule-based controls. Natural language processing is used to automatically parse and operationalize global threat intelligence feeds. Together, these methods aim to reduce mean time to detect and respond beyond what manual monitoring alone can achieve at scale.

Architecturally, the model shifts from perimeter-focused defence toward a Zero-Trust framework in which all users, devices, and transactions are continuously verified. Core components include AI-enabled next-generation firewalls for deep packet inspection, an AI-augmented security operations centre to coordinate detection and incident response, and encrypted, air-gapped backups to ensure recoverability and data sovereignty specifically addressing failure points observed in major ransomware incidents. The deployment model follows an infrastructure-as-a-service approach. Pilot implementations using these protocols at tertiary and affiliated facilities demonstrated substantial reductions in breach-simulation success rates within six months. The design supports scaling across national health networks through shared cyber-risk pooling models that extend protection to smaller and resource-constrained facilities.

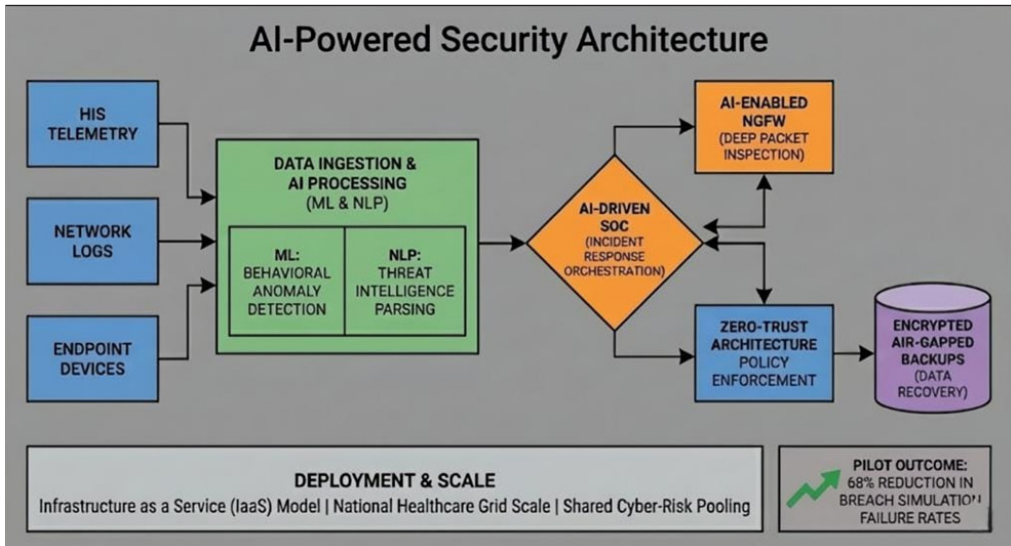
Key barriers include the cost and operational burden of upgrading legacy infrastructure,

*¹CyberPeace Foundation Delhi, ²Kutchhi Engineering College

remediating known vulnerabilities, and improving basic cyber hygiene among non-technical staff. Governance and policy alignment are also essential to ensure that enhanced monitoring does not compromise patient privacy or data protection obligations. Implementation should align with national data protection law and require secure-by-default

standards, supported by dedicated public funding for health-sector cybersecurity tooling. Long-term resilience depends on embedding a “cyberpeace” approach to digital health treating health information infrastructure as critical national assets that require continuous, layered defence comparable to physical public health systems.

Figure 10: AI Powered Security Architecture



Digital Health Validitron: A unified service for de-risking clinical AI solutions

Mahima Kalla, Kit Huckvale*

Despite growing global interest and investment in artificial intelligence (AI) for health care, adoption of AI-based tools in routine clinical practice remains limited. A systematic review by van de Sande and colleagues found that fewer than 2% of AI models developed for intensive care settings progress beyond early prototyping and testing stages. Many AI-enabled clinical decision-support tools fail to translate into practice because model performance is often evaluated in isolation, without sufficient attention to workflow integration, human-factor failure points, and real-world usability constraints.

A persistent challenge for clinical AI innovators is that conventional research and evaluation methods frequently fail to identify these socio-technical risks early in the innovation lifecycle. Approaches such as randomized controlled trials and late-stage pilot studies are often too slow, costly, and operationally complex to surface workflow, safety, and adoption barriers before deployment—particularly for tools affecting already underserved or vulnerable populations. The Digital Health Validitron, based at the Centre for Digital Transformation of Health at the University of Melbourne, was established to address this gap by providing an integrated environment for early-stage, real-world-oriented evaluation of AI and digital health technologies.

Established in 2023, the Digital Health Validitron functions as both a translational research studio and a specialist academic group supporting the design, development, evaluation, and implementation of AI-enabled and other digital health solutions. Its core methodology centers on immersive, simulation-based research to identify and mitigate factors that commonly lead to failure during real-world clinical implementation.

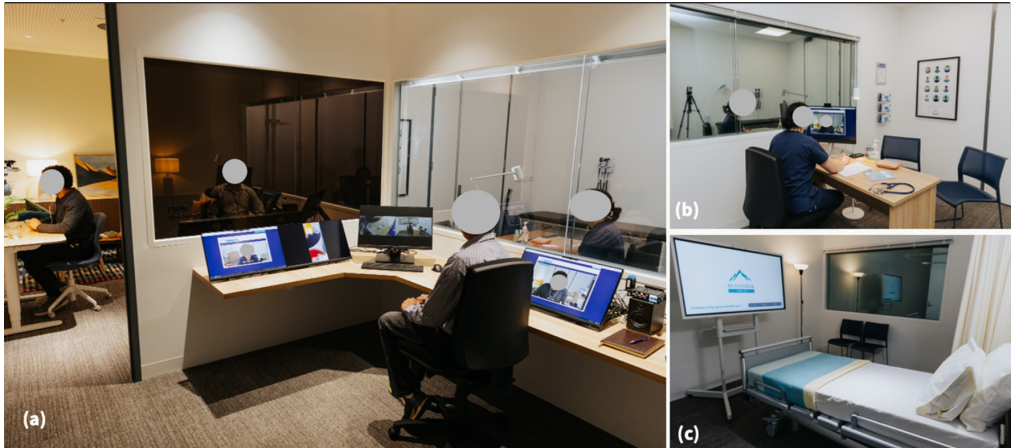
This approach is enabled through two complementary environments. The Validitron SimLab is a physical simulation facility comprising configurable home, primary care, and secondary care settings in which clinicians and patients can co-design and test technologies in realistic but controlled conditions. Simulation rooms are equipped with discreet observation areas and comprehensive video, audio, and screen-capture systems to support detailed human-technology interaction analysis. The Validitron Sandbox is a digital replica of a health information ecosystem built on established and emerging interoperability standards, including FHIR. It incorporates production, open-source, and white-label health information systems such as primary and secondary care electronic medical records, telehealth platforms, and synthetic data generators—that can be assembled modularly to rapidly construct test environments for workflow, interoperability, and technical compatibility assessment.

Centre for Digital Transformation of Health, Faculty of Medicine, Dentistry and Health Sciences, The University of Melbourne, Australia

A defining strength of the Validitron model is its multidisciplinary team, bringing together expertise in implementation science, participatory design, health equity, machine learning, health economics, software engineering, and behavioral science. This breadth supports comprehensive socio-technical evaluation rather than model-centric testing alone. The framework aligns with the RISE criteria for clinical AI adoption—regulatory

readiness, interpretability, interoperability, and structured data and evidence. Through simulation-based evidence generation, the Validitron approach supports systematic assessment of explainability, workflow fit, interoperability, and data readiness in realistic clinical scenarios, producing standards-aligned evidence to inform regulatory review, organizational decision-making, and responsible real-world deployment.

Figure 11: (a) left side – simulated patient situated in the simulated patient home space having a telehealth consultation with clinician; right side – human-computer interaction researcher situated in the control room, observing the proceedings both via one-way glass and screen mirroring; (b) clinician situated in the simulated primary care room conducting a telehealth consultation with patient next door; (c) example of another tertiary care simulated space in the laboratory (not utilised in this particular simulation study)





सर्वजन शिवाय | सर्वजन सुखाय
WELFARE FOR ALL | HAPPINESS OF ALL