AI-Driven Healthcare Delivery in Pakistan: A Framework for Systemic Improvement

Imama Zahoor
Lahore University of Management Sciences
Pakistan

Shiza Ihtsham
Lahore University of Management Sciences
Pakistan

Umar Ramzan
Lahore University of Management Sciences
Pakistan

Agha Ali Raza
Lahore University of Management Sciences
Pakistan

Basmaa Ali
Lahore University of Management Sciences
Pakistan

ABSTRACT

In Low- and Middle-Income Countries (LMICs), poor health outcomes come from a high burden of disease, a shortage of healthcare professionals, and inefficient health information exchange leading to substantial economic losses. In this paper, we highlight critical gaps in healthcare delivery in Pakistan and propose solutions to improve patient outcomes in resource-constrained environments. We have built *Darcheeni*, an AI-driven healthcare framework that leverages artificial intelligence to assist and supplement physicians, streamline healthcare processes, and prioritize patient-centered care. *Darcheeni* analyzes doctor-patient interactions in real-time, integrates lab and imaging data, generates and distributes care plans customized to the patient's needs, and sends them directly to patients' smartphones. We also discuss the challenges and limitations associated with sustainable AI integration by centering our learnings from the pilot deployment of *Darcheeni*. By focusing on Pakistan as a case study, this work offers practical insights and strategies for deploying AI-driven technologies sustainably in similar resource-constrained environments and contributes to the broader discourse on the role of AI in global health improvement.

CCS CONCEPTS

• **Information systems → Expert systems**; • **Social and professional topics → Medical technologies**; • **Human-centered computing →** *Accessibility theory, concepts and paradigms.*

KEYWORDS

healthcare process optimization, resource-constrained environments, artificial intelligence, sustainable AI integration, pilot deployment

ACM Reference Format:
Imama Zahoor, Shiza Ihtsham, Umar Ramzan, Agha Ali Raza, and Basmaa Ali. 2024. AI-Driven Healthcare Delivery in Pakistan: A Framework for Systemic Improvement. In *Proceedings of ACM SIGCAS/SIGCHI Conference on Computing and Sustainable Societies (COMPASS '24).* ACM, New York, NY, USA, 8 pages. https://doi.org/XXXXXXX.XXXXXXX

COMPASS '24, July 8 - 11, 2024, New Delhi, India

ACM ISBN 978-1-4503-XXXX-X/18/06
https://doi.org/XXXXXXX.XXXXXXX

1 INTRODUCTION

Approximately 8 million individuals die in Low- and Middle-Income Countries (LMICs) every year due to inefficient healthcare practices, with 60% of these deaths stemming from poor quality care [29]. In 2015 alone, such fatalities resulted in economic losses exceeding US$6 trillion [22]. In Pakistan, there is a higher burden of disease, ineffective health information exchange (HIE) resulting in resource wastage due to poor allocation and duplication of efforts, premature and avoidable deaths, and a cycle of poverty stemming from health expenditures [20]. Additionally, the doctor-to-patient ratio in Pakistan stands at 1:1300, notably below the WHO's suggested ratio of 1:1000 [19]. This undermines the quality of healthcare by increasing the workload for existing practitioners and resulting in overburdened and burnt-out healthcare professionals. A major aspect of this challenge lies in documentation practices in Electronic Health Records (EHRs) and the physician role has become two-fold: 1) attending to the patient's needs and 2) entering diagnoses, orders, visit notes, and additional administrative data into the EHR [12].

The proliferation of Artificial Intelligence (AI) in healthcare has emerged as a tool for augmenting physician expertise and patient care. Recent research has explored its application in LMICs in the following major areas: clinical decision support systems, treatment planning, screening, triage assistants, and health chatbots [1, 2, 8]. However, most of these applications are separate and standalone systems researched and developed in high-income countries. This presents an opportunity for a solution that is effective, scalable, seamless to integrate into the current system, and well-aligned with local norms and health patterns. In this paper, we:

- Identify the foundational gaps in healthcare facilities in Pakistan.
- Introduce *Darcheeni*, an AI-driven framework designed to fill these gaps and develop an approach to healthcare delivery in resource-constrained environments.
- Discuss the challanges from the pilot deployment of *Darcheeni*.

It is important to note that while *Darcheeni* is formulated as a comprehensive framework based on foundational research, only a portion of it has been deployed to measure real-world implications and challenges. Therefore, we refer to *Darcheeni* as a framework rather than a full working system. Through this short paper, we hope to contribute to the discourse on AI in healthcare and specifically, provide practical insights and strategies for the sustainable deployment of such technologies in resource-constrained environments.

2 BACKGROUND

AI is a rapidly evolving field which promises to construct machines that can perform tasks that typically require human intelligence. It comprises a range of techniques such as Machine Learning (ML), Deep Learning (DL), and Natural Language Processing (NLP). In particular, AI is poised to transform healthcare by improving the quality, accessibility, and affordability of care.

2.1 AI in Healthcare

AI can be utilized in several ways to improve healthcare. Traditional Machine Learning techniques provide the ability to analyze large and complex patient datasets to predict outcomes, optimize medical dosages, and create personalized treatment plans [16, 32, 36]. They can mitigate the risk of adverse drug effects by predicting them in advance [26]. Similarly, patients at risk of developing chronic diseases can be identified, leading to faster interventions and a reduced strain on the healthcare system in the future [30]. Clinicians can be assisted in making treatment decisions, particularly in predicting therapy responses [16] and in conducting triage for patients based on urgency, prioritizing high-risk cases, leading to lower waiting times and improved patient flows [11].

Similarly, researchers and practitioners alike have utilized computer vision techniques to analyze medical imagery and extract relevant diagnostic information. Robust models for disease classification, region segmentation, and nodule detection are developed for conditions where the relevant data can be collected [38]. Radiology, pathology, ophthalmology, and dermatology have, in particular, received substantial attention. The use of such methods in radiology has been so significant that it has quickly grown into its own field of research [7, 27, 33]. These developments have allowed for a rapid response in times of crisis - for instance in the development and deployment of detection models for COVID-19 [41].

Previous research has indicated that physicians can dedicate up to 35% of their time towards documentation [17]. Large Language Models (LLMs) can be used to automatically generate such documentation. This can save valuable time, enabling healthcare professionals to focus on patient care, making the entire process more satisfying for both physicians and patients. LLMs can also be used to generate a concise summary of a patient's complete medical history, providing the physician with relevant information and facilitating increased efficiency [23, 35, 39]. They can also be incorporated into traditional computer-aided diagnosis systems, allowing physicians to ask open-ended questions regarding specific diagnoses so as to better understand the system.

2.2 AI in LMICs

AI holds the potential for tremendous opportunities for LMICs which are lacking in medical resources, expertise, and infrastructure [13, 34]. Rural and isolated areas, in particular, may benefit from the introduction of AI-based health applications. By utilizing software solutions, such applications could dramatically reduce the costs associated with screening and treatment plan selection for diseases that require specialized expertise or expensive equipment [5, 9, 18, 31]. Other potential benefits include the reduction of waiting times and provision of a private channel for those suffering from stigmatizing diseases such as psychiatric pathologies [21, 24]. Maternal and child health, which are major issues in LMICs, may also find such systems to be particularly useful [14, 37]. Automated translation services can tailor such solutions for the local language and culture, improving the accessibility and compliance of such services in areas where cultural beliefs may pose a barrier to healthcare [25].

However, health applications utilizing AI have been developed predominantly within high-income countries and have only recently begun to be used within LMICs [37]. As such, obtaining an evaluation for these systems within a local context is a difficult process, often resulting in uninformed decision making within that context [14]. Therefore, it is pertinent to consider the associated risks and challenges that are unique to LMICs. For instance, the training of AI systems often requires large quantities of high-quality data which is difficult to obtain in LMICs [34]. Simply relying on data collected from other countries can lead to biased or defective solutions [14]. Applications developed in the context of high-income countries may recommend treatments that are either inaccessible or prohibitively expensive in low-income countries [37]. There also exists a lack of the expertise and infrastructure required to create governance models used to guide such technologies, resulting in an adverse effect on the quality and safety of the system [14]. Such systems may also require a large initial investment to develop, leading to a lack of deployment in low-income countries. Furthermore, the widespread practice of informal medicine in certain societies may instead lead to the spread of non-compliant AI applications.

Research regarding the challenges faced in the practical implementation of healthcare AI systems in LMICs is a new and burgeoning field. We seek to add to this discourse by providing a comprehensive overview of the implementation of such a system in Pakistan, where relevant research is particularly scarce. The insights gained from this implementation may be used to more effectively guide the design and creation of AI-based health applications in Pakistan and other LMICs.

3 FOUNDATIONAL RESEARCH

3.1 Methodology

The development of *Darcheeni* involved an investigation of healthcare delivery across a set of 8 outpatient facilities in Lahore, Pakistan. Our observational study was conducted over a year, from June 2022 to June 2023, and aimed to identify the gaps in the current healthcare delivery systems that AI could reimagine. The facilities included:

- Purpose-built Outpatient Departments (OPD) in large tertiary care hospitals,
- Charity-funded clinics,
- Private clinics in upper-middle-class and middle-class, neighborhoods, and
- Purpose-built OPD in large secondary care hospitals.

The observational team consisted of the first, second, and fifth authors. The first author is a computer science research assistant while the second and fifth authors are medical doctors with experience at the intersection of technology and healthcare. During the study period, the team accumulated over 200 hours of observation across the aforementioned facilities. Each observation session lasted anywhere from 2 to 5 hours during clinic hours and we observed over 100 clinical encounters in total with a diverse demographic of patients in terms of age, gender, and socio-economic status. Each facility was visited twice and each visit was conducted by at least two researchers from the team.

Given the sensitivity of the healthcare environment, particularly patient consultations, we did not audio or video record to respect privacy and confidentiality. Instead, our data collection relied on detailed handwritten field notes taken by researchers during and immediately after observation sessions. These notes covered aspects such as patient flow, documentation practices, patient-physician interactions, and the use of medical devices and technology within the clinical setting. We also conducted informal interviews with healthcare providers, including physicians, nurses, and administrative staff, alongside patients and their families when possible. These conversations were instrumental in providing deeper insights into the practical challenges and expectations of the healthcare delivery ecosystem from those who experience it firsthand. All conversations were paraphrased in our field notes, with identifiers removed to maintain confidentiality. We then employed axial qualitative coding to identify broad themes and patterns. By incorporating these personal anecdotes and testimonies into our analysis, we aim to present a more comprehensive understanding of the healthcare landscape.

We acknowledge some of the limitations of our study. Firstly, it includes a select number of healthcare facilities in Lahore - a metropolitan city in Pakistan which potentially limits the generalizability of our findings to other, more rural, regions. Secondly, the reliance on handwritten notes without audiovisual recording may omit non-verbal cues important for a complete understanding of the healthcare delivery process. Nonetheless, due to the diversity within the facilities visited, we maintain that our findings present a microcosm of the broader healthcare challenges faced nationwide and form a foundation for proposing the *Darcheeni* framework.

3.2 Findings

3.2.1 Patient Identification. Patient identification practices include tracking and managing patient information across visits, which is critical for ensuring continuity of care and accurate medical history records. At the tertiary-care hospital, each patient was assigned a new medical record number (MRN) at every visit, lacking retention of demographic or medical history details. This systemic amnesia led to incomplete patient histories, disjointed care, and increased morbidity and mortality. A conversation with a family member of a patient revealed that due to the absence of a consistent medical history, the patient was repeatedly misdiagnosed and treated for the wrong ailment which led to a deterioration of their actual condition.

While some of the larger institutions have begun integrating the Computerized National Identity Card (CNIC) into their EHR systems as a means of ensuring consistency in patient identity, smaller healthcare facilities, which constitute a majority of the accessible healthcare facilities, often have rudimentary practices due to resource constraints. In these settings, patient identification and record-keeping are frequently managed through paper records, which are prone to being misplaced or damaged. Moreover, in some instances, patients are only provided with a slip containing their prescribed medications without any formal documentation of their visit, diagnosis, or treatment plan. A patient expressed their frustration, stating,

"Every time I visit, it's like starting from zero. Last time, I just got a slip for my meds—no history, no follow-up."

3.2.2 Overcrowding and Wait Times. Overcrowding was ubiquitous across all sites. The absence of triage and patient scheduling results in the patients enduring lengthy wait times and little to no distinction between first-time and follow-up encounters. Patients either self-direct themselves to departments, rely on nurses' guesses, or are sent to general practitioners. In one instance, a patient with a broken thigh bone was sent to the emergency department (ER) after waiting for over an hour in the outpatient department (OPD). Additionally, facilities have not done a capacity assessment of their infrastructure, personnel, and technical facilities (labs and imaging). Hence, they cannot plan for timely, humane delivery of care where supply and demand match each other. This mismanagement not only strains the physical infrastructure but also overburdens the medical and administrative staff, impacting the overall quality of care delivered.

3.2.3 Substandard Care. Poor delivery of care may be a result:
1. *High work pressure on physicians.* High work pressure on physicians often leads to rushed consultations, with doctors seeing 50-150 patients per day and allocating approximately 3 minutes per patient. This prevents thorough reviews and physical examinations, crucial for accurate diagnosis and treatment planning. For example, in one of the charity-funded clinics, doctors prioritized speed over depth in patient evaluations, causing them to overlook critical symptoms or health indicators. Less than half of these oversights were identified after the patient had left.

2. *Lack of accountability.* Health outcomes are rarely tracked or directly linked to the interventions suggested by doctors, creating a gap in quality assurance and improvement mechanisms. There are also little to no mechanisms in place to hold physicians accountable for their mistakes, leading to carelessness and lapses in diligence on the part of physicians. For instance, an administrative staff at an OPD revealed that when a patient receives incorrect medication due to a misdiagnosis, there's often no system in place to address or learn from these mistakes.

3. *Poor medical record-keeping.* The majority of public setups have no medical records while some use paper records, resulting in inconsistent and often illegible notes and prescriptions. Most of the current setups delegate the responsibility of maintaining medical records and histories to the patients. Other setups rely on outdated electronic systems often containing overlooked or inaccurate information. Furthermore, there is a lack of communication of care plans to patients and little to no use of exit interviews to explain the treatment or gather feedback.

3.2.4 Sub-optimal Implementation of Patient Rights. This is categorized by inadequate protection of privacy, insufficient informed consent processes, and minimal patient engagement in treatment decisions. For instance, in some clinics, conversations about sensitive health information occur in semi-public areas, compromising patient privacy. Additionally, patients received treatment without a thorough explanation of the procedures, risks, and alternatives due to time constraints or staff shortages. For example, most patients in overcrowded facilities were administered medication without being informed about its side effects or without their explicit consent, as healthcare providers rushed to see the next in line. This adversely impacts the quality of care, patient trust, and treatment outcomes. A doctor in one of these settings commented:

"It gets so busy that sometimes we can't provide them [the patients] with all the information. We're doing the best we can under the circumstances."

3.2.5 Narrow Use of Technology. The application of technology in healthcare is predominantly geared towards administrative functions such as billing, inventory management, and scheduling appointments. This focus often neglects the potential technology holds for enhancing patient care, medical record-keeping, and health outcome analysis. In recent years, various initiatives to incorporate technology in healthcare practices, such as implementing EHR systems, have been launched with the intention of streamlining operations and improving patient care. However, these efforts frequently encounter obstacles due to a lack of infrastructure support, inadequate training for staff, absence of user research, and resistance to evolve from traditional paper-based methods. Particularly, adherence to technology-based mechanisms is difficult, with many healthcare providers reverting to manual processes due to the perceived complexity or unreliability of digital systems. A member of the IT team in one of the hospitals commented:

"We've tried integrating new software, have so many new ideas, but it's no use if they [the doctors] are not ready to use it [...] it's more trouble than it's worth."

There is also a reliance on outdated EHR systems and legacy technologies that are not user-friendly and do not meet the dynamic needs of modern healthcare delivery. These systems often lack interoperability, making it difficult to share patient information across different healthcare providers and leading to fragmented care.

4 DESIGN STRATEGIES

Based on the findings from the foundational research, we outlined six goals to address the challenges.

4.1 Development Goals

The landscape of health systems operating at multiple tiers necessitates system-wide action and foundational reforms, particularly in Pakistan's context where healthcare challenges are amplified by resource constraints and a high disease burden. Recognizing that fixes at the micro-level (e.g., healthcare provider or clinic) alone are unlikely to alter the underlying performance of the entire system. It is also essential to note that AI systems are not a cure-all for the structural problems faced by healthcare facilities and require a nuanced consolidation of policy and technology. Keeping this in mind, our proposed framework rests on a multifaceted approach encompassing the following strategic components:

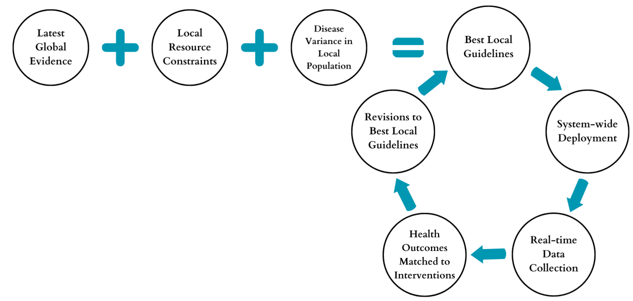

Figure 1: Development of Local Guidelines in *Darcheeni*.

4.1.1 Developing Local Standards of Care. Establishing care standards to regional specificities to ensure that quality care is universally accessible (Figure 1). In regions where healthcare disparities are stark, developing localized care guidelines can significantly reduce preventable fatalities at a cost sustainable for households and communities.

4.1.2 Empowering Primary Care. Enhancing the skill set of primary care practitioners in Pakistan not only optimizes healthcare delivery but also attenuates the long-term costs associated with untreated or poorly managed conditions.

4.1.3 Expanding Access through Skill Development. Addressing the acute shortage of healthcare professionals in Pakistan, particularly in rural and under-served areas, necessitates an expansive approach to skill development. Empowering a broader range of healthcare workers, including nurses, midwives, lady health visitors (LHVs), pharmacists, and dispensers, to provide care traditionally reserved for general practitioners can significantly widen the healthcare net, ensuring more individuals receive timely and appropriate care.

4.1.4 Real-Time Health Outcome Tracking. Greater focus on health outcomes by collecting and analyzing real-time health data. This enables the identification effective interventions and optimization of patient care pathways, ensuring healthcare delivery is efficient and outcome-focused.

4.1.5 Cultivating Accountability. Implementing real-time data collection to foster a culture of accountability within Pakistan's healthcare ecosystem not only aids in identifying care delivery lapses and delays but also ensures healthcare practitioners are cognizant of their roles and responsibilities.

4.1.6 Patient Education for Empowerment. Educating patients about their conditions and treatment options allows them to actively participate in their healthcare journey, make informed decisions, and advocate for quality care.

5 ARCHITECTURE OVERVIEW

To achieve these development goals, we introduce *Darcheeni*, a framework that leverages AI to streamline healthcare delivery processes, enhance physician support, and prioritize patient-centered care to address the challenges of healthcare delivery in Pakistan.

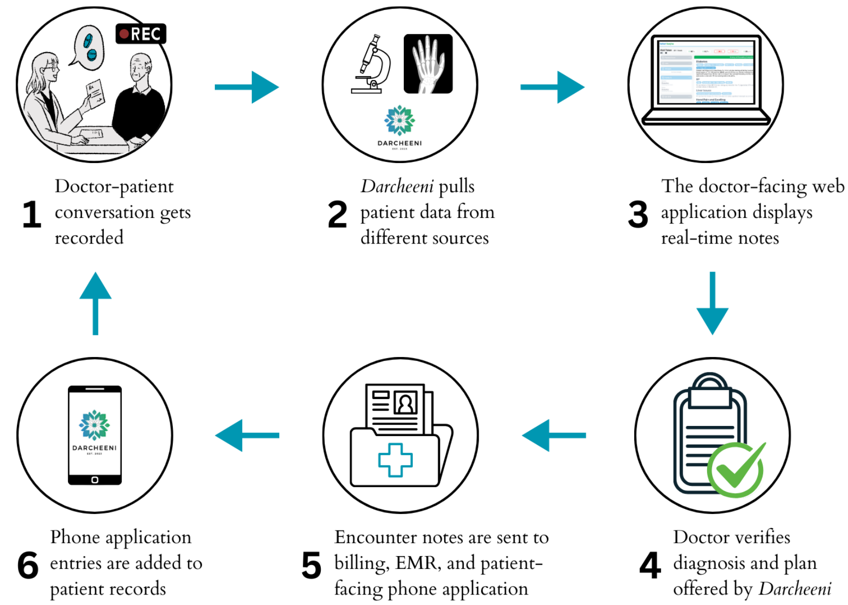

Figure 2: How *Darcheeni* works.

The framework consists of five primary modules: holistic data integration, efficient order processing and care plans, digital profiles, periodic analytics, and streamlined data transfers.

5.1 Holistic Data Integration

Darcheeni begins by actively listening to doctor-patient interactions, capturing nuances that otherwise might be overlooked in fast-paced environments. This includes a speech-to-text module that transcribes and translates local languages to English. Alongside the current interaction, it also aggregates data from diverse sources, including labs, imaging, and specialist consultations. This ensures the model has a 360-degree view of patient health to offer holistic and comprehensive advice to physicians by writing a doctor's note and a personalized assessment and plan for each encounter.

5.2 Efficient Order Processing and Patient-Centric Care Plans

Once the doctor signs off orders, *Darcheeni* automates the transmission of medication, lab, and imaging orders, as well as referrals to specialists to streamline the process and lifting responsibilities from the patient. It also sends personalized care plans to patients' smartphones to encourage real-time implementation and improve patient compliance with medications and lifestyle modifications.

5.3 Smart Data Analysis and Digital Profiles

Data from smartphone apps are analyzed by *Darcheeni*, providing insights for doctors to review during subsequent appointments. Over time, it constructs digital profiles for each patient, facilitating personalized and granular care delivery. Similarly, it builds profiles for healthcare providers, including strengths, weaknesses, and error patterns. These profiles will not only help the model work better alongside the provider by personalizing each interaction but also help educate and empower both doctors and their patients.

5.4 Periodic Analytics for Continuous Improvement

Darcheeni conducts periodic analyses to create a data-driven, evidence-based healthcare system that is dynamically responsive to patient needs and healthcare delivery outcomes. It does this by collecting health data from various sources, including patient interactions, treatment responses, health outcomes, continuous monitoring, evaluation and adaptive learning. This allows for the identification of trends, patterns, and areas of improvement in healthcare delivery. By understanding which interventions lead to positive health outcomes and which areas require enhancement, healthcare providers can implement targeted strategies to improve care quality and efficiency.

5.5 Efficient Data Transfer

In many LMICs, including Pakistan, the referral process from primary or secondary care to tertiary care often lacks coordination, efficiency, complete patient information, and accessibility. Through integrated health records and automated referral pathways, *Darcheeni* allows for a smoother transition between care levels and anonymized data transfer between specialists and sub-specialists. It does this by identifying when a patient needs specialist care and promptly initiating the referral. Additionally, patients' access to their health records allows them to follow up on referrals, understand their health conditions better, and make informed choices about their care.

Figure 2 shows how *Darcheeni* works by completing the loop in the following ways:

1. The doctor and patient conversation is recorded.

2. The AI Model, also titled *Darcheeni*, pulls patient data from various sources, including lab and imaging reports, as well as clinical notes from previous encounters from a HIPAA-compliant [4] database.

3. While the encounter is ongoing, the doctor-facing web application displays real-time processing of the conversation into discrete problems, pertinent questions to prompt the physician in case of any overlooked aspects, and the distribution of important sections into bins and packets. The latter is a categorizing methodology, where bins refer to the general topics of discussion, such as medications, past medical history, and current problems, while packets refer to each instance of a bin.

4. Once the conversation is over, the doctor stops the recording, and a personalized and editable diagnosis, assessment, and plan is created. The doctor verifies and signs off on the plan, leading to the creation of three pathways: billing, storage in the EHR via an API, and transmission to the patient-facing mobile application. The personalized care plan on the patient-facing application is presented in the form of a smart to-do list, where each actionable item is sorted by priority.

5. At home, the patient records and updates the patient-facing mobile application. Most phone app entries go to the medical records and will be reviewed at the next visit. However, some critical entries notify the doctor or clinic administration for timely intervention.

6 DISCUSSION AND FUTURE WORK

6.1 Pilot Deployment

For the pilot deployment of our system, we collaborated with a tertiary care hospital, established under the approval of a joint Institutional Review Board (IRB). Within this framework, we constructed a Clinical Innovation Lab (CIL) in the OPD of the hospital to facilitate a controlled environment for system deployment. Starting in November 2023, the system's foundational capabilities, note-taking and storing patient records, were implemented and tested in a real-world clinical setting. This initial phase focused on integrating these core functionalities into the daily workflows of two doctors operating within the CIL, who collectively attended to an average of 40 patients per day.

Throughout the deployment, our research team collected over 50 hours of observational data, documenting the system's performance, user interaction, and overall impact on clinical operations. To ensure ethical compliance and respect for patient privacy, explicit verbal consent was obtained from all participating patients before recording their encounters. Any recordings made without consent were immediately discarded, and no identifying patient information was stored in our records. The setup utilized the existing computers in the OPD and wired microphones to minimize disruption to the existing clinical infrastructure. This approach not only provided valuable insights into the system's current capabilities and areas for improvement but also laid the groundwork for the next phase of deployment. The forthcoming step will involve enhancing the system with advanced functionalities, including automated diagnosis, assessment, and the creation of personalized care plans.

6.2 Challenges from Pilot Deployment

The initial deployment phase revealed several challenges that we believe to be relevant for any future implementation of a similar system in an LMIC. These include both technical and usability hurdles that must be overcome for such a system to be effective:

6.2.1 Transcription Accuracy. Accurate transcription and speaker diarization of the conversation was one such fundamental hurdle. Any audio transcription system must work well for a combination of local languages and dialects. In our case, this was a combination of English, Urdu, and Punjabi. Given that Urdu and Punjabi are low-resource languages, we designed a pipeline to collect and transcribe audio data from multiple clinics to be used to fine-tune our speech model and improve transcription accuracy. Using such context-relevant training data also allows the model to become resistant to background noise, which is a prevalent problem in many local clinics. Transcribing audio files in this way proves to be a slow and expensive process, presenting itself as the main bottleneck in the fine-tuning pipeline. Languages used within LMICs are often low-resource and so any such system would require similar fine-tuning before being deployed.

6.2.2 Speed and Scalability. For such a system to be effective, it must function in real time and scale with an increasing number of users. To achieve the required speed, we utilize a dynamic number of cloud containers that scale up to match the current number of users. During a conversation, we only transcribe short audio segments, prioritizing speed. Once a visit is complete, the entire conversation is transcribed once more so as to increase the final accuracy of the transcription. The full system must ultimately be completely integrated with each hospitals EHR system. This would, however, be a time-consuming process as each institution has a unique system and would require a custom solution.

6.2.3 System Usability. The manner in which the system is used also reveals some flaws. As the tool is based primarily on the audio collected from a conversation, it must rely on the physician to be quite descriptive. However, we observe that physicians often do not go into the required detail, especially when conducting physical examinations. Clinics are often overcrowded and noisy, with patients being accompanied by friends or family, who frequently interject into the conversation. The patients themselves are occasionally too quiet, making their voice indistinguishable from background noise. To combat this, the team eventually had to implement a "one patient at a time" policy to ensure high-quality audio recordings. Additionally, early on during the initial deployment, we realized that physicians must have the ability to alter any part of the output from the system before it is saved. This would not only increase the accuracy of the output but would also allow us to identify and correct mistakes within the system in an iterative fashion.

6.3 Potential Impact

Pakistan spends $38 per person per year on healthcare [20]. While increasing this allocation might be challenging, optimizing its utilization is crucial. In this paper, we outline how *Darcheeni* has the potential to deliver quality care, through active listening of doctor-patient interactions, generation of electronic medical records instantly, and seamless capture of vital information without any additional effort from physicians. This can enable nationwide accessibility of health records, continuity of care, delivery of quality care, preventive measures, early intervention, and the effective management of chronic conditions. By doing so, it can significantly increase the Human Development Index (HDI) and reduce healthcare costs associated with treating advanced and severe cases. Moreover, automatic record creation for every patient visit establishes a real-time accountability mechanism for healthcare providers. This holds the potential to reduce healthcare provider absenteeism and improve patient outcomes.

Similarly, primary healthcare (PHC) facilities in Pakistan are either not accessible or have limited resources [15]. 71% of primary care needs are met by general practitioners (GPs), who are not required to complete structured training and gaps exist in knowledge and skills in clinical practice [10]. Our framework involves the delivery of a local standard of care so that more patients can get treatment early in the disease, potentially cure or halt its progression, and consequently take the pressure off tertiary centers. It is also pertinent to note that it takes 17 years on average in the US for new evidence to become medical practice [28]. This measure is infinite in Pakistan because there is a lack of continuation of medical education for physicians. Upgrading *Darcheeni's* care guidelines once every 6 months will ensure that all physicians are treating their patients according to the best guidelines dictated by international evidence and resource capacity. The availability of reliable and up-to-date health data will also allow the development of highly targeted and effective health policies.

6.4 Limitations and Future Work

While the *Darcheeni* framework presents a promising approach to address critical challenges in healthcare delivery, it is essential to acknowledge several inherent limitations and constraints.

The success of this framework relies heavily on a robust technological infrastructure, requiring electricity, an internet connection, a web-connected desktop or laptop for the health caregiver, and an internet-connected smartphone for the patient. While most facilities in the larger cities possess computers and internet connectivity, the speed of the internet, software versions, and the quality of microphones vary greatly among setups, requiring significant additional investment at times. As we travel farther away from the metropolitans, the resource constraints only add up. 62.27% of Pakistanis live in rural regions with limited access to reliable internet connectivity and advanced healthcare facilities [3]. At the moment, we are being purposeful in seeking socioeconomic and ethnic diversity as we set up clinical collaborations for the acquisition of our training dataset.

Furthermore, although Pakistan has high smartphone penetration, the gender gap in smartphone (mobile phone with internet connection) ownership was 43% in 2021, with a gender gap of overall 15% in LMICs [6]. Given privacy concerns for healthcare data, future research may explore solutions so that women not only have access to their health data but also have control over access to it. Similarly, since only 22% of the total physicians in Pakistan serve in the rural areas, where the majority of the population resides, future versions of *Darcheeni* could involve allied healthcare workers (AHWs) like nurses, midwives, and compounders who make up the bulk of primary care providers in the bottom, mostly rural Pakistan and other LMICs [40].

The current deployment of *Darcheeni* is in its pilot phase, limited to a smaller scale. While this allows for iterative testing and optimization, extrapolating the findings to a larger, more diverse healthcare ecosystem may pose challenges such as scalability, affordability, and maintenance costs, which must be addressed to ensure sustainable integration within the socio-economic landscape of the deployment areas. The research team aims to validate and refine the model for nationwide deployment by the end of this year. As a future endeavor, scaling the *Darcheeni* framework to other LMICs presents an avenue. This would require language adaptation, local health needs adaptation, affordability, and regulatory compliance. Conducting a longitudinal impact assessment and soliciting feedback from users, including both healthcare providers and patients, is imperative to evaluate the sustained effectiveness of *Darcheeni*.

Currently, the only account of a doctor-patient interaction is recorded by the doctor. Using audio recordings of encounters shifts the focus from a physician's recall to the patient's narrative. Future work can incorporate the voice metadata of a patient's initial history to diagnose disease more quickly and help provide better-fitting therapeutic choices for the patient. It will receive data from patients' smartphones, analyze it, and share the insights with the care team. It will also organize the doctor's orders and instructions into a smart list that will remind the patient via timely notifications to implement the care plan with more compliance. However, we also concede that any solution which has access to private data must be carefully evaluated according to local regulatory laws and strict guardrails must be implemented accordingly.

7 CONCLUSION

In this article, we introduced the architectural design of *Darcheeni* to the sustainable health research community, its potential impact on healthcare delivery in LMICs, and the foreseen challenges stemming from resource constraints. As the field of AI-driven healthcare evolves, we hope to see more technological solutions catering to an under-represented and diverse user base.

REFERENCES

[1] [n. d.]. https://www.madiro.org/

[2] [n. d.]. https://ada.com/

[3] [n. d.]. https://www.theglobaleconomy.com/Pakistan/rural_population_percent/#:~:text=Rural%20population%2C%20percent%20of%20total%20population&text=The%20latest%20value%20from%202022%20is%2062.27%20percent

[4] 1996. Health Insurance Portability and Accountability Act of 1996 (HIPAA). https://www.hhs.gov/hipaa/index.html. Accessed: 2024-03-10.

[5] Robert Caprara, Keith L Obstein, Gabriel Scozzarro, Christian Di Natali, Marco Beccani, Douglas R Morgan, and Pietro Valdastri. 2014. A platform for gastric cancer screening in low-and middle-income countries. *IEEE Transactions on Biomedical Engineering* 62, 5 (2014), 1324–1332.

[6] Isabelle Carboni, Nadia Jeffrie, Dominica Lindsey, Matthew Shanahan, Claire Sibthorpe, C Butler, et al. 2021. Connected Women-The Mobile Gender Gap Report 2021. *GSMA Intelligence: UK* 7 (2021).

[7] Garry Choy, Omid Khalilzadeh, Mark Michalski, Synho Do, Anthony E Samir, Oleg S Pianykh, J Raymond Geis, Pari V Pandharipande, James A Brink, and Keith J Dreyer. 2018. Current applications and future impact of machine learning in radiology. *Radiology* 288, 2 (2018), 318–328.

[8] Tadeusz Ciecierski-Holmes, Ritvij Singh, Miriam Axt, Stephan Brenner, and Sandra Barteit. 2022. Artificial intelligence for strengthening healthcare systems in low-and middle-income countries: a systematic scoping review. *npj Digital Medicine* 5, 1 (2022), 162.

[9] Hugo Jair Escalante, Manuel Montes-y Gómez, Jesús A González, Pilar Gómez-Gil, Leopoldo Altamirano, Carlos A Reyes, Carolina Reta, and Alejandro Rosales. 2012. Acute leukemia classification by ensemble particle swarm model selection. *Artificial intelligence in medicine* 55, 3 (2012), 163–175.

[10] Hamida Farazdaq, Jaleed A Gilani, Asra Qureshi, and Unab I Khan. 2022. Needs assessment of general practitioners in Pakistan: A descriptive cross-sectional survey. *Journal of Family Medicine and Primary Care* 11, 12 (2022), 7664–7670.

[11] Sabina Ohri Gandhi, Sabina Gandhi, and L Sabik. 2014. Emergency department visit classification using the NYU algorithm. (2014).

[12] Rebekah L Gardner, Emily Cooper, Jacqueline Haskell, Daniel A Harris, Sara Poplau, Philip J Kroth, and Mark Linzer. 2019. Physician stress and burnout: the impact of health information technology. *Journal of the American Medical Informatics Association* 26, 2 (2019), 106–114.

[13] Jonathan Guo and Bin Li. 2018. The application of medical artificial intelligence technology in rural areas of developing countries. *Health equity* 2, 1 (2018), 174–181.

[14] Ahmed Hosny and Hugo JWL Aerts. 2019. Artificial intelligence for global health. *Science* 366, 6468 (2019), 955–956.

[15] Hina Jawaid and Abdul Jalil Khan. [n. d.]. Making Pakistani Primary Care More Resilient. ([n. d.]).

[16] Kevin B Johnson, Wei-Qi Wei, Dilhan Weeraratne, Mark E Frisse, Karl Misulis, Kyu Rhee, Juan Zhao, and Jane L Snowdon. 2021. Precision medicine, AI, and the future of personalized health care. *Clinical and translational science* 14, 1 (2021), 86–93.

[17] Erik Joukes, Ameen Abu-Hanna, Ronald Cornet, and Nicolette F de Keizer. 2018. Time spent on dedicated patient care and documentation tasks before and after the introduction of a structured and standardized electronic health record. *Applied clinical informatics* 9, 01 (2018), 046–053.

[18] Shivaram Kalyanakrishnan, Rahul Alex Panicker, Sarayu Natarajan, and Shreya Rao. 2018. Opportunities and challenges for artificial intelligence in India. In *Proceedings of the 2018 AAAI/ACM conference on AI, Ethics, and Society*. 164–170.

[19] SA Khan. 2019. Situation analysis of health care system of Pakistan: post 18 amendments. *Health Care Current Reviews* 7, 3 (2019), 244.

[20] Salman J Khan, Muhammad Asif, Sadia Aslam, Wahab J Khan, and Syed A Hamza. 2023. Pakistan's healthcare system: A review of major challenges and the first comprehensive universal health coverage initiative. *Cureus* 15, 9 (2023).

[21] Ilona Kickbusch. 2019. Health promotion 4.0. , 179–181 pages.

[22] Margaret E Kruk, Anna D Gage, Naima T Joseph, Goodarz Danaei, Sebastián García-Saisó, and Joshua A Salomon. 2018. Mortality due to low-quality health systems in the universal health coverage era: a systematic analysis of amenable deaths in 137 countries. *The Lancet* 392, 10160 (2018), 2203–2212.

[23] Chengtai Li, Yiming Zhang, Ying Weng, Boding Wang, and Zhenzhu Li. 2023. Natural language processing applications for computer-aided diagnosis in oncology. *Diagnostics* 13, 2 (2023), 286.

[24] David D Luxton. 2014. Artificial intelligence in psychological practice: Current and future applications and implications. *Professional Psychology: Research and Practice* 45, 5 (2014), 332.

[25] David D Luxton. 2014. Recommendations for the ethical use and design of artificial intelligent care providers. *Artificial intelligence in medicine* 62, 1 (2014), 1–10.

[26] Guillaume L Martin, Julien Jouganous, Romain Savidan, Axel Bellec, Clément Goehrs, Mehdi Benkebil, Ghada Miremont, Joëlle Micallef, Francesco Salvo, Antoine Pariente, et al. 2022. Validation of artificial intelligence to support the automatic coding of patient adverse drug reaction reports, using nationwide pharmacovigilance data. *Drug Safety* 45, 5 (2022), 535–548.

[27] Maciej A Mazurowski, Mateusz Buda, Ashirbani Saha, and Mustafa R Bashir. 2019. Deep learning in radiology: An overview of the concepts and a survey of the state of the art with focus on MRI. *Journal of magnetic resonance imaging* 49, 4 (2019), 939–954.

[28] Zoë Slote Morris, Steven Wooding, and Jonathan Grant. 2011. The answer is 17 years, what is the question: understanding time lags in translational research. *Journal of the royal society of medicine* 104, 12 (2011), 510–520.

[29] Engineering National Academies of Sciences and Medicine. 2018. *Crossing the Global Quality Chasm: Improving Health Care Worldwide.* The National Academies Press, Washington, DC. https://doi.org/10.17226/25152

[30] Karin M Nelson, Evelyn T Chang, Donna M Zulman, Lisa V Rubenstein, Freddy D Kirkland, and Stephan D Fihn. 2019. Using predictive analytics to guide patient care and research in a national health system. *Journal of general internal medicine* 34 (2019), 1379–1380.

[31] Allisson Dantas Oliveira, Clara Prats, Mateu Espasa, Francesc Zarzuela Serrat, Cristina Montañola Sales, Aroa Silgado, Daniel Lopez Codina, Mercia Eliane Arruda, Jordi Gomez i Prat, and Jones Albuquerque. 2017. The malaria system microapp: a new, mobile device-based tool for malaria diagnosis. *JMIR research protocols* 6, 4 (2017), e6758.

[32] Sameer Quazi. 2022. Artificial intelligence and machine learning in precision and genomic medicine. *Medical Oncology* 39, 8 (2022), 120.

[33] Luca Saba, Mainak Biswas, Venkatanareshbabu Kuppili, Elisa Cuadrado Godia, Harman S Suri, Damodar Reddy Edla, Tomaž Omerzu, John R Laird, Narendra N Khanna, Sophie Mavrogeni, et al. 2019. The present and future of deep learning in radiology. *European journal of radiology* 114 (2019), 14–24.

[34] Laura Sallstrom, Olive Morris, and Halak Mehta. 2019. Artificial intelligence in Africa's healthcare: ethical considerations. *ORF Issue Brief* 312 (2019), 1–11.

[35] Yiqiu Shen, Laura Heacock, Jonathan Elias, Keith D Hentel, Beatriu Reig, George Shih, and Linda Moy. 2023. ChatGPT and other large language models are double-edged swords. , e230163 pages.

[36] Murugan Subramanian, Anne Wojtusciszyn, Lucie Favre, Sabri Boughorbel, Jingxuan Shan, Khaled B Letaief, Nelly Pitteloud, and Lotfi Chouchane. 2020. Precision medicine in the era of artificial intelligence: implications in chronic disease management. *Journal of translational medicine* 18 (2020), 1–12.

[37] Brian Wahl, Aline Cossy-Gantner, Stefan Germann, and Nina R Schwalbe. 2018. Artificial intelligence (AI) and global health: how can AI contribute to health in resource-poor settings? *BMJ global health* 3, 4 (2018), e000798.

[38] Hongtao Xie, Dongbao Yang, Nannan Sun, Zhineng Chen, and Yongdong Zhang. 2019. Automated pulmonary nodule detection in CT images using deep convolutional neural networks. *Pattern recognition* 85 (2019), 109–119.

[39] Xi Yang, Aokun Chen, Nima PourNejatian, Hoo Chang Shin, Kaleb E Smith, Christopher Parisien, Colin Compas, Cheryl Martin, Anthony B Costa, Mona G Flores, et al. 2022. A large language model for electronic health records. *NPJ digital medicine* 5, 1 (2022), 194.

[40] Zeshan Zahid, Suleman Atique, Muhammad Hammad Saghir, Iftikhar Ali, Amna Shahid, and Rehan Ali Malik. 2017. A commentary on telerehabilitation services in Pakistan: current trends and future possibilities. *International journal of telerehabilitation* 9, 1 (2017), 71.

[41] Jianpeng Zhang, Yutong Xie, Guansong Pang, Zhibin Liao, Johan Verjans, Wenxing Li, Zongji Sun, Jian He, Yi Li, Chunhua Shen, et al. 2020. Viral pneumonia screening on chest X-rays using confidence-aware anomaly detection. *IEEE transactions on medical imaging* 40, 3 (2020), 879–890.