Guidelines for Developing and Implementing AI-Powered Health Decision Support Technologies in Drug Outlets

Developed by Inspired Ideas Research Foundation and No Discipline LLC, in collaboration with Apotheker Health Access Initiative (AHAI) in Tanzania

Table of Contents

Primary Contributing Authors

Inspired Ideas Research Foundation & No Discipline LLC: Ally Salim Jr, Megan Allen, Dr. Kelvin Mariki, Moses Minja, Thomas Pallangyo, Rose Lawrence Maina, Fatma Ally

Executive Summary

These guidelines are meant to provide recommendations related to the integration of Artificial Intelligence (AI) into the frontline drug shop ecosystem to close the critical gap between healthcare access and quality. In many resource-limited settings, drug shops and retail pharmacies serve as the primary point of care for millions, yet they often operate with limited diagnostic tools and varying levels of formal training.

These guidelines provide a strategic roadmap for leveraging AI-powered health decision support systems to augment the expertise of dispensers, ensuring that every patient encounter is guided by evidence-based, high-quality clinical standards.

Target Audience

This document is designed for a multi-disciplinary audience of stakeholders committed to health system strengthening:

- **Government Regulators and Ministries of Health:** To provide a framework for the safe, legal, and effective oversight of digital health tools in the private sector.
- **Developers and Technologists:** To offer design principles—such as "offline-first" architecture and medication-centered workflows—that address the unique constraints of limited-resource healthcare.
- **Funders and International Agencies:** To demonstrate the cost-benefit and scalability of AI interventions that align with global health equity goals.
- **Implementing Organizations:** To provide practical lessons on dispenser training, community sensitization, and the maintenance of complex digital networks in low-infrastructure environments.

Why This Is Important

Poor quality of care in primary health systems contributes to millions of preventable deaths and unintended pregnancies annually. While drug shops provide essential proximity and affordability, they are prone to inconsistent knowledge application and diagnostic errors. Implementing AI in this setting is not merely a technological upgrade; it is a vital public health intervention. Through tasks such as automating complex pediatric dosing, identifying "red flag" symptoms for immediate referral, and providing real-time disease surveillance, AI-enabled drug shops can transform from medicine retailers into high-functioning nodes of a comprehensive primary care network.

How to Read This Document

This document serves as a comprehensive framework for governments, health organizations, and developers to implement AI-powered health decision support systems within drug outlets. It is structured to guide users through the entire lifecycle of an AI health tool—from initial ecosystem understanding to long-term scaling.

- *Part 1: Foundational Context*: Establishes the role of medicine retail as a parallel primary healthcare system and the specific quality gaps AI is designed to address.
- *Part 2: Implementation Framework:* Outlines the technical, human, and regulatory requirements needed to prepare an environment for digital health interventions.
- *Part 3: Design and Deployment:* Provides technical specifications for AI system design, focusing on offline-first architectures and clinical governance.
- *Part 4: Training and Change Management:* Details the "high-touch" training models and stakeholder engagement strategies required for successful technology adoption.
- *Part 5: Post-Deployment Operations:* Focuses on the maintenance, technical support, and rigorous monitoring and evaluation (M&E) necessary for safety and effectiveness.
- *Part 6: Challenges and Sustainability*: Addresses common implementation risks—such as user turnover and connectivity—and provides a roadmap for national institutionalization.

Blue Boxes: Implementation Insights

Throughout these guidelines, you will find *Blue Boxes* that detail real-world findings from our AI health implementation in Tanzania. These boxes offer contextual "deep dives" into:

- **What Worked:** Successful strategies for dispenser motivation and clinical accuracy.
- **Lessons Learned:** Critical adjustments made to software design and training protocols based on field feedback.
- **Real-world Data:** Metrics on how AI integration impacted drug dispenser decision making and health outcomes in a live drug outlet environment.

The Afya-Tek Project

These guidelines are directly informed by the multi-year implementation and success of the Afya-Tek Project in Tanzania.[1] Afya-Tek is a digitally-enabled continuum of care program that integrates Community Health Workers (CHWs), private drug outlets (ADDOs), and primary health facilities into a unified system. Led by Apotheker Health Access Initiative with funding from Fondation Botnar, the project demonstrated that digital health tools can improve healthcare delivery in Tanzania by addressing challenges like poor care quality and coordination.

As a part of the project, Inspired Ideas Research Foundation also piloted AI-powered decision support tools to improve the accuracy of clinical decisions at the point of care, including Elsa DOTS (see more on page 10).

This document distills the technical, operational, and regulatory learnings from the Afya-Tek consortium—into a replicable framework for global health practitioners.

8 Key Recommendations

Prioritize Decision Support Over Autonomy

AI must serve as a clinical aid that *augments* the user's expertise, ensuring the human provider retains final authority and accountability for the patient.

Design for "Offline-First" Functionality

In areas with unstable power or network, systems must function locally on the device to avoid disrupting care, with automated data synchronization once connectivity is restored.

Align with Business Values

To ensure sustained adoption by private actors, AI tools must be bundled with business features like inventory management, stock-out alerts, and automated revenue reporting.

Ensure Model Explainability

Avoid "black box" outputs - use interpretable and explainable designs and clear language to explain *why* the AI is making a recommendation, which builds dispenser trust and ensures clinical reasoning.

Build for Low Tech Literate Popultions

Design intuitive, multimodal interfaces—incorporating icons, voice prompts, and simplified workflows. Translate all clinical logic, training materials, and user interfaces in the local language to eliminate ambiguity, ensure accurate data collection, and reduce the risk of errors.

Encode National Guidelines

Map AI logic and clinical recommendations to national Standard Treatment Guidelines and ensure use of the specific list of medicines approved for the level of care.

Invest in "High-Touch" Training and Support

Initial training should be in-person and cohort-based, followed by intense on-site support to address the high turnover of staff and the steep learning curve of digital tools.

Engage Local and National Government

Engage national regulatory bodies early to align the tool with national treatment guidelines and data privacy laws, ensuring the system is a permanent part of the national health strategy.

Part 1: Foundational Context

1. Understanding the Drug Shop Ecosystem

1.1 The Medicine Retail Landscape in LMICs

Before discussing technology for drug shops, it's important to understand the broader context of medicine retail in low- and middle-income countries (LMICs), where formal healthcare infrastructure often requires improvements to meet population needs.

Public health facilities often face geographic inaccessibility, chronic understaffing, long wait times, and inadequate medication supply chains. Perhaps most critically, these facilities are financially prohibitive for many families despite being nominally "free" or low-cost; the hidden costs of transportation, time away from income-generating work, and unofficial fees create significant barriers to access, particularly for the most vulnerable populations.

The Emergence of Medicine Retail as Primary Healthcare

In this context, small private medicine shops have been successful across LMICs to fill the gap, effectively becoming a parallel primary healthcare system. These outlets have emerged organically in response to community needs and typically share several key characteristics that make them attractive to patients:

- *Accessibility and Affordability:* Located within residential neighborhoods, offering services within walking distance.
- *Convenience*: Extended operating hours and immediate service without the multi-hour waits common at public facilities.
- *Trust*: Providers are often community members who offer culturally familiar care in local languages.

Studies across LMICs consistently show that drug shops are often the first point of contact for health care, particularly for acute illnesses (fever, diarrhea, respiratory infections), minor injuries and ailments, reproductive health needs, and chronic disease medication refills.

The Quality-Access Tradeoff

Despite their reach, these outlets can lack formal training, consistent regulatory oversight, and clear boundaries for their scope of practice. Additionally, dispensing practices can vary

widely and are often driven by consumer preferences, which can affect patient health outcomes.

> **Note on Terminology**: While this guideline is built on experience with Accredited Drug Dispensing Outlets (ADDOs) in Tanzania, the principles and recommendations apply broadly to similar medicine retail outlets in resource-limited settings globally. These outlets operate under different names and regulatory frameworks across countries, but serve similar roles in their communities.
>
> Readers should interpret recommendations through the lens of their own country's drug retail systems. The core challenges, opportunities, and implementation considerations remain relevant across contexts, even where specific regulatory details differ.

1.2 The Role of Drug Shops in Primary Healthcare Delivery

Tanzania's ADDO Model: A Case Study in Formalization

In Tanzania, Accredited Drug Dispensing Outlets (ADDOs) are privately operated retail outlets that have been trained and licensed to sell essential medicines, including selected prescription drugs, particularly in rural and underserved areas.[2] The ADDO program was launched in 2003 by the Tanzanian Food and Drug Authority (TFDA) as an innovative response to a critical healthcare access gap: the widespread presence of unregulated drug shops that were selling medicines without proper training, quality assurance, or regulatory oversight.

The program takes a comprehensive approach combining owner and dispenser training, government accreditation based on standards, business incentives, and local regulatory enforcement, with efforts to increase consumer demand for quality products.[3] ADDOs have become the principal source of medicines in Tanzania and represent an important component of the country's multi-faceted healthcare system.[4] They provide accessibility, affordability, convenience, and cultural familiarity.

Why Tanzania's ADDO Model is Instructive

While each country's context differs, the ADDO program is internationally recognized as one of the most systematic efforts to formalize and regulate medicine retail. It offers:

- *Defined standards and scope of practice:* Defined training, infrastructure, and operational requirements are provided for ADDO owners and dispensers. Specific medicines list and conditions appropriate for ADDO-level care are defined at the national level.

- *Government ownership:* The system is led and controlled by the national regulatory authority.
- *Scale*: There are over 10,000 registered ADDOs in Tanzania across diverse settings.
- *Longevity*: Over two decades of implementation, learning, and evolution has built a strong base of evidence for quality improvement and evaluating effectiveness.

These characteristics make it a useful model for understanding how artificial intelligence (AI) can be integrated into formalized medicine retail systems. *Recommendations could also be generalized to other types of systems, including those with different regulatory structures, pharmacy-based systems with more highly trained providers, and community health worker programs with medicine distribution authority.*

1.3 Demographics and Health-Seeking Behaviors of Patients Receiving Services at Drug Shops

Understanding the characteristics and health-seeking behaviors of individuals who utilize drug shops is critical for designing appropriate AI interventions.

Target Population

Drug shop clients are a mix of rural and peri-urban populations, with a large proportion of women and children. They seek care for a range of health issues, including both acute and chronic disease management. Based on surveillance data, the most frequent health concerns presented include febrile illnesses (malaria, non-specific fever), respiratory tract infections, gastrointestinal complaints, skin conditions, reproductive health needs, and minor injuries. Clients hold a mix of health beliefs and cultural expectations. For example, many expect "strong" treatment (injections or antibiotics) even when not clinically indicated, putting pressure on dispensers.

From our own research and market analysis of this population, we have observed that:

- Over 50% of individuals arrive at drug outlets "already knowing their condition," often assuming they have malaria or a UTI.
- Patients are often reluctant to seek laboratory testing due to associated costs and time, preferring immediate medication.
- Dispensers face pressure to provide "strong" treatments (like antibiotics) to meet patient expectations and maintain business credibility.

Care Seeking Behaviors

Individuals choose drug shops as primary points of care for multiple reasons that reflect both practical and cultural preferences. Cost sensitivity drives many to seek affordable care without consultation fees, while convenience factors—proximity to home or work and flexible operating hours—make drug shops accessible. Prior relationships matter significantly, with patients returning to familiar dispensers they trust and feel comfortable speaking with openly. For some conditions, particularly sexually transmitted infections and family planning needs, drug shops offer privacy and reduced stigma compared to public facilities. Many patients seek care at drug shops first and escalate to clinics only if symptoms persist or worsen, making these outlets true first-line providers.

1.4 Capabilities, Limitations, and Potential for AI

Capabilities

Drug shops are typically permitted to sell medicines from an approved list that includes both over-the-counter products and specific prescription medicines for common conditions. This includes medicines such as antimalarials, antibiotics for common infections, oral rehydration solutions, family planning products, and other essential medicines deemed appropriate for community-level distribution.

Beyond selling products, dispensers are expected to provide basic health counseling and medication advice, helping patients understand how to use medicines correctly and when to seek higher-level care. They maintain records of medicine sales and patient interactions, contributing to both business management and potential disease surveillance. Dispensers are also expected to refer patients to health facilities when conditions exceed their scope of practice.

Training

Drug dispensers undergo initial training that typically covers basic pharmacology and medicine use, recognizing symptoms and appropriate treatment, dispensing practices and patient counseling, business management and record-keeping, and regulatory compliance requirements. However, studies have shown that drug shop dispensers often express poor knowledge of basic pharmacology,[5] and the depth of training varies considerably across outlets and regions.

Limitations and Quality Considerations

Despite the training and regulatory framework, inappropriate medicines use and dispensing remains a challenge.[6] This is due in large part to limited diagnostic capability

and knowledge gaps around critical clinical decisions: knowing when to refer versus treat, understanding appropriate dosing for different age groups and conditions, recognizing drug interactions and contraindications, and identifying danger signs requiring urgent referral. Guideline adherence is also challenging, particularly for conditions like malaria, pneumonia, and diarrheal diseases where treatment protocols evolve based on emerging evidence and changing resistance patterns.

2. Leveraging Artificial Intelligence to Support Drug Shop Operations

2.1 Defining Artificial Intelligence

The term "artificial intelligence" encompasses a broad range of technologies. For the purposes of these guidelines, we distinguish between AI-Powered Systems and Non-AI Digital Health tools.

AI-Powered Systems: Systems that use expert knowledge, machine learning, natural language processing, or other AI techniques to perform activities such as:

- Analyzing patient symptoms and recommending differential diagnoses
- Suggesting appropriate treatments based on patient-specific factors
- Predicting risk levels and recommending referral
- Learning from patterns in data to improve performance over time
- Adapting recommendations based on context (e.g., local disease prevalence, stock availability)

These systems exhibit characteristics such as pattern recognition, contextual reasoning, probabilistic decision-making, and continuous learning.

Non-AI Digital Health Tools are complementary but distinct. They can include but are not limited to electronic medical records (EMRs), SMS reminder systems, supply chain management software, and telemedicine platforms.

Box 1: Elsa DOTS (Drug Outlet Technology System) Application

Elsa Dots is an AI-powered application developed by Inspired Ideas Research Foundation in Tanzania and piloted as a part of the AfyaTek Program. DOTS is an intelligent health assistant that uses data and expert knowledge to support business

operations, sales and inventory management, and health decision making (symptom assessment and drug-drug interactions). The tool is delivered to drug dispenser users through a mobile application and has a simple-to-use interface for quick data collection, evaluation, and reporting.

DOTS includes robust stock management tools and business reporting. Additionally, the platform includes proprietary AI algorithms for **analyzing symptoms**, i**dentifying drug interactions**, and **providing next steps recommendations** including treatment and dosage suggestions. Elsa's algorithms support over 200 conditions for children, adolescents, and adults, ranging from pediatric illnesses like otitis media, malaria, and pneumonia to chronic illnesses such as hypertension and diabetes.

2.2 Range of Potential AI Applications for Drug Shops

AI can be applied across multiple domains within drug shop operations. Understanding this spectrum helps implementers prioritize use cases. The following indicates just a handful of examples where AI can complement drug shop operations. More detailed case examples can be found in Section #3.

For initial implementation, we recommend prioritizing health decision support, as these directly address the most critical quality gaps previously identified. Operational and public health applications can be phased in as capacity grows.

Clinical Decision Support	• Symptom Assessment and Diagnosis
	  ○ Structured symptom collection
	  ○ Generation of differential diagnoses with probabilities
	  ○ Red flag detection for immediate referral
	  ○ Patient interviews (voice or text-based)
	• Treatment Recommendation
	  ○ Evidence-based medicine selection from drug shop regulations
	  ○ Patient-specific dosing (age, weight, pregnancy status, comorbidities)
	  ○ Duration and frequency guidance
	• Drug Interaction Assessments
	  ○ Contraindication checking
	  ○ Recommendation of alternative medications based on patient presentations and potential drug interactions

	• Risk Stratification
	  ○ Assessment of severity and likelihood of complications
	  ○ Need for follow-up and/or referral
	• Patient Counseling Support
	  ○ Generating clear explanations in appropriate language that a dispenser can communicate to the client, including lifestyle and prevention advice
	  ○ When-to-return instructions

Operational and Business Support (AI approaches that build on top of tools that perform basic inventory tracking and shop management)	• Inventory/ Stock Management
	  ○ Demand forecasting based on seasonal patterns and disease trends
	  ○ Automated reordering and expiry date tracking
	  ○ Optimal stock level recommendations
	• Business Analytics
	  ○ Revenue optimization
	  ○ Identifying best-selling products and recommending new pricing strategies

Public Health and Surveillance	• Disease Surveillance
	  ○ Real-time syndromic surveillance across a region or country
	  ○ Outbreak detection and recommendations for resource allocation or prevention efforts
	  ○ Early warning systems for emerging or reemerging conditions
	• Quality Monitoring
	  ○ Detecting inappropriate prescribing patterns and evaluating antimicrobial resistance
	  ○ Identifying training needs
	  ○ Benchmarking performance across outlets

Continuing Education	• Adaptive Learning Systems
	  ○ Personalized training based on knowledge gaps
	  ○ Case-based learning with feedback
	  ○ Competency assessments

2.2 Support vs. Autonomy in Health Decision Support Systems

A fundamental principle for AI in drug shops is that the AI-powered system serves as decision support, not autonomous decision-making. This distinction has critical implications

for safety, regulation, liability, and user acceptance. It is recommended that developers and implementers of AI-powered digital health tools carefully consider the role that the technology plays and ensure that the human end-user retains control over the final health decisions.

	AI as Decision Support (Recommended)	**AI as Autonomous Decision-Maker**
Authority	Human dispenser retains final say.	AI makes the final treatment decision.
Logic	Transparent reasoning ("why").	"Black box" output.
Override	Dispenser can override based on context.	Little room for contextual judgment.
Risk	Lower regulatory and liability burden.	Higher regulatory burden and risk.

2.3 What AI is Not

While AI can significantly enhance drug shop services, it has clear boundaries. Setting realistic expectations prevents disappointment and misuse. AI does not replace:

- *Human Connection:* Empathy, cultural sensitivity, and personal trust.
- *Physical Examination*: Assessing vital signs or recognizing life-threatening visual cues.
- *Systemic Solutions*: AI cannot fix medicine stock-outs, extreme poverty, or inadequate hospital capacity for referrals.

Box 2: Risks of (Over)Reliance & Deskilling

AI in clinical settings demonstrates risks when users stop thinking critically and defer blindly to AI, lose skills they do not practice, or miss errors because they assume the AI is always correct.

To better understand this, Inspired Ideas Research Foundation conducted research to evaluate how reliant drug shop dispensers were on AI-powered technologies when determining a differential diagnosis for a presented clinical case vignette. We explored how the drug dispensers responded to technology that is framed as "always

correct" in an attempt to measure whether they begin to rely on it without any critical thought of their own.

We found that dispensers relied on the AI's decision 25% of the time, even when the AI provided no explanation for its decision. We've observed that the dispensers are more likely to change their final answer for conditions they don't encounter frequently, signaling a low confidence in their first choice.

Reference: https://arxiv.org/abs/2302.09487

3. Case Examples of AI in Drug Shops

The drug shop model has achieved considerable success in expanding access to medicines, but quality and safety gaps persist despite training and supervision efforts. AI offers unique capabilities to address problems that traditional approaches have struggled to solve. The following case examples illustrate relevant challenges related to drug shops and how AI tools might be able to offer novel solutions.

Problem 1: Inconsistent Application of Knowledge

Situation: Even well-trained dispensers apply their knowledge inconsistently. A dispenser who correctly manages pneumonia in the morning might forget to check for contraindications in the afternoon when tired or distracted.

How AI Tools Can Help: Tools can apply the same standards to every patient encounter. AI serves as a reliable cognitive aid that:

- Prompts systematic assessment rather than relying on memory
- Checks contraindications every time
- Applies guidelines consistently
- Reduces errors of omission (forgetting to ask important questions)

Problem 2: Knowledge Decay Over Time

Situation: Dispensers receive initial training, but knowledge can degrade over time without regular refreshers. With thousands of dispersed drug shop

dispensers, ongoing training is resource-intensive. Printed materials become outdated and are rarely consulted during patient interactions.

How AI Tools Can Help:

- Provide real-time, case-specific education during patient encounters
- Update automatically when guidelines change, ensuring current recommendations
- Reinforce learning through repetition and feedback
- Identify individual knowledge gaps for targeted continuing education

Problem 3: Diagnostic Complexity Beyond Current Capacity

Situation: Dispensers must differentiate between conditions based only on symptom histories, without diagnostic tests or specialist consultation. Many conditions present similarly (e.g., fever can indicate malaria, typhoid, pneumonia, UTI, or other causes), leading to presumptive treatment that may be inappropriate.

How AI Tools Can Help:

- Process complex combinations of symptoms to generate differential diagnoses
- Incorporate epidemiological data (local disease prevalence, seasonal patterns)
- Use probabilistic reasoning beyond human mental calculation
- Identify subtle patterns that distinguish similar conditions
- Recommend when testing or referral is needed rather than presumptive treatment

Problem 4: Inadequate Pediatric Dosing

Situation: Calculating weight-based doses for children is error-prone, especially under time pressure. Dispensers may round inappropriately, use adult doses, or avoid treating children out of uncertainty. Dosing charts help but require looking up information while managing impatient queues.

How AI Tools Can Help:

- Instantly calculate precise doses based on weight/age
- Account for contraindications specific to children
- Reduces underdosing (treatment failure) and overdosing (toxicity)

Problem 5: Weak Referral Decision-Making

Situation: Deciding when to treat vs. refer is challenging. Over-referral costs time and money; under-referral risks poor outcomes. Dispensers may feel pressure to treat (patient expectations, business concerns) even when referral is appropriate.

How AI Tools Can Help: Tools can perform objective risk stratification based on danger signs and provide clear criteria for when higher level care is needed. This can reduce liability concerns by providing evidence-based recommendations

Problem 6: Limited Supervision and Quality Monitoring

Situation: Supervisors can't observe real-time practice or review every encounter. Problem patterns may go undetected until serious incidents occur.

How AI Tools Can Help:

- Generate real-time quality metrics
- Flag concerning patterns for supervisor review
- Provide objective performance data for supportive supervision
- Create audit trails for accountability

Problem 7: Missed Public Health Surveillance Opportunities

Situation: Drug shop dispensers are frontline disease detectors but lack systems to aggregate and analyze patterns. Outbreaks may be recognized late. Surveillance data often isn't collected or is of poor quality.

How AI Tools Can Help: Tools can automatically capture syndromic data from encounters and detect unusual patterns suggesting outbreaks. They can also:

- Enable early warning before cases reach formal health facilities
- Provide geographic and temporal granularity
- Reduce reporting burden on dispensers

Part 2: Implementation Readiness

4. Assessing Readiness for AI Health Tools

Before system development and deployment, a comprehensive readiness assessment must be conducted to ensure the environment can support AI functionality without compromising care.

4.1 Infrastructure Requirements

At a minimum, drug shops and outlets require basic power for device charging and at least intermittent network connectivity for data synchronization. To ensure the resilience of AI-powered systems in low-resource environments, the following technical standards are recommended:

Power Access and Grid Stability

- *Stability Evaluation:* Sites must be evaluated for grid stability; in areas with high instability or frequent load shedding, primary reliance on the national grid is insufficient. Mitigation strategies include:
    - *Solar-Powered Charging:* For off-grid or highly unstable areas, solar microgrids are recommended to maintain 24-hour operation.
    - *High-Capacity Power Banks:* Outlets should be equipped with high-capacity power banks to prevent system downtime during longer outages.
    - Recommended use of rechargeable mobile devices over computer systems that must be plugged in.

Network and Connectivity

- *Offline-First Architecture:* Because network outages are frequent, it is recommended that the AI decision support engine reside locally on the device. This ensures clinical logic and symptom assessments remain available even when the internet is down.
- *Data Sync Protocols (if the deployment has a shared server that needs to be synced to):*

- Automated Queuing: Systems should automatically queue and upload data once a connection is re-established.
- Sync Indicators: Interfaces should clearly display "Sync Status" so dispensers know whether their records have been successfully reported to the central server.

- *Device-to-Device Offline Communication:* In settings where multiple devices are used (e.g., a dispenser phone and a shop manager tablet), tools should utilize Bluetooth or Wi-Fi Direct for local data sharing without requiring external internet access.

Technology Availability and Device Specifications

- *Hardware Sufficiency:* Devices must be selected with future-proofing in mind. Low-end phones often lack the RAM or NPU (Neural Processing Unit) required for computationally intensive on-device AI workloads. Minimum recommended specifications:
  - *Processor*: Efficient mid-range chipsets that balance performance with heat management.
  - *Battery*: Phone should last for 12 hours with light periodic use or 4-6 hours of extensive use.
  - *Storage*: Minimum of 64GB local storage is recommended to allow for comfortable storage of data and process algorithms without slow down.
  - *Memory*: Requires a minimum of 4GB RAM for devices dedicated to store operations, for devices used for other purposes (including personal) the minimum is 6GB.

Up-to-Date Public Health Data Infrastructure

- *Epidemiology as the Backbone:* AI models rely on prior probability—the baseline likelihood of a condition in a specific area. The effectiveness of a model improves significantly when developers have access to granular, real-time epidemiological data.
- *Syndromic Surveillance:* By digitizing drug outlets, we create a bi-directional data flow where current outbreak data (e.g., a spike in local cough cases) updates the AI's diagnostic weighting in real-time.

4.2 Human Capacity Considerations

A successful AI deployment is contingent upon a deep understanding of the "human-in-the-loop." Before introducing digital tools, implementers must evaluate the baseline competencies and digital readiness of the dispensers who will be using them.

Dispenser Education and Health Literacy

- *Baseline Clinical Knowledge:* Many ADDO dispensers have limited formal health training or may exhibit limited knowledge of basic pharmacology. Assessments should identify these gaps to determine if foundational health training is required before or alongside AI onboarding.
- *Health Literacy Levels:* The general public (patients) often hold specific health beliefs and expect "strong" treatments like antibiotics or injections. Dispensers must be assessed on their ability to navigate these cultural expectations while adhering to AI-driven rational medicine use principles.

Technology and Digital Health Literacy

- *Digital Skill Disparities:* Technology literacy varies significantly between urban and rural settings, with rural dispensers often reporting lower digital health literacy and less access to technical support.
- *Functional vs. Applied Literacy:* Assessments should distinguish between basic operational use (e.g., navigating a smartphone) and applied understanding (e.g., interpreting an AI's probabilistic recommendation for clinical decision-making).
- *Language and Localization:* Given the complexity of medical terminology, the preference for local languages (e.g., Swahili) is high. Systems that use "lay" language familiar to dispensers are more likely to be accepted than those relying on English medical jargon.

Stakeholder Readiness and Motivation

- *Owner and Operator Buy-in:* Drug shop owners are primarily driven by business outcomes, such as profit and customer retention. Assessments must ensure that owners perceive the tool as a value-add for shop efficiency or credibility rather than a business burden.
- *Dispenser Professional Identity:* Some providers view AI as an "add-on" that increases their local credibility. Others may resist tools due to change fatigue or a fear of being replaced. Understanding these psychological drivers is essential for tailoring change management strategies.

Readiness Assessment Methodology

- *Qualitative Interviews:* Conduct in-depth discussions to explore perceptions, lived experiences, and contextual challenges.
- *Observation:* Use "improv-based" scenarios or shadow dispensers during patient interactions to see how they handle real-world workflow pressures.
- *Pre-Training Competency Checks*: Use standardized clinical vignettes or basic device navigation tasks to establish a baseline for tailoring training curricula.

4.3 Regulatory and Legal Readiness

The most critical factor in regulatory readiness is early and continuous government and organizational buy-in. Moving a tool from pilot to a recognized health intervention requires that national authorities are active participants in the development process.

- *Strategic Government Alignment:* Successful implementation requires a dedicated team focused on engaging with national regulatory bodies. In Tanzania, this could include the Tanzania Medicines and Medical Devices Authority (TMDA), the Ministry of Health, the Pharmacy Council, or district/ regional-level leadership. Having government representatives on the implementation team is essential to navigate bureaucracy and ensure the tool aligns with national priorities.
- *Licensing and Accreditation:* AI-powered decision support should be integrated into existing accreditation frameworks where possible and required. This ensures that the use of the tool is legally recognized as part of the dispenser's scope of practice.
- *Liability*: Although the legal frameworks for AI health tools are still new in many LMICs, the implementing organization can ensure that the professional accountability remains with the human provider. Because the AI serves as decision support rather than an autonomous decision-maker, the dispenser retains final authority and responsibility for the patient encounter.
- *Data Protection and Privacy Compliance:* Systems must be designed to comply with national data or consumer protection laws. This includes establishing clear protocols for patient consent, data minimization, and secure storage.

5. Economic Considerations

Implementing AI in the drug shop ecosystem requires a shift from viewing technology as a one-time project cost to a long-term health system investment.

5.1 Aligning AI with Business Values

Drug shops, pharmacies, and ADDO dispensers are strongly motivated by profit and business success. AI adoption is most successful when bundled with tools that solve business pain points, such as:

- *Inventory Management:* Digital tracking of stock, expiry dates, and automated reordering.
- *Business Analytics*: Revenue optimization and identifying best-selling products.
- *Credibility Dividend:* Using modern technology increases patient faith in the dispenser's thoroughness.

5.2 Cost Benefit Analysis Framework

The economic value of AI in drug outlets is measured not only by the cost of the software but by the prevention of clinical errors and the optimization of resource use. AI interventions generate value by reducing the financial burden on public hospitals through accurate "red flag" detection and referral, and by curbing the long-term economic threat of antimicrobial resistance through appropriate antibiotic prescribing. There are two types of costs that should be considered when deploying an AI-powered tool in this context:

Upfront Investment: This includes the procurement of mobile devices, software development and licensing, research or pilot costs, potential data collection for AI development, initial server setup, and the training and (potential) certification of dispensers.

Ongoing Operational Costs: Sustainable budgets must account for monthly data connectivity (particularly in the pilot phase to support accessibility for users), cloud hosting fees, software maintenance/security updates, and a hardware replacement fund to account for the typical 2-3 year lifespan of mobile devices in the field. We also accounted for 10% device loss or damage.

5.3 Financing Models

For many governments and agencies in resource-limited settings, the initial investment for AI integration is external donor or grant funding. To ensure these partnerships are successful and sustainable, it is critical to align national health priorities with the strategic goals of international funders. Large organizations often prioritize innovation, scalability, and measurable improvements in population health. By framing AI in drug outlets as a tool for "health system strengthening" and "data-driven equity," implementers can secure long-term support while building the local infrastructure necessary for eventual independence.

Beyond initial grant/ funding alignment, a diversified financing strategy is essential for long-term operations. Examples of ongoing support include:

- *Donor/NGO Support and Sustainability Planning:* Initial phases often rely on external grants to cover high upfront R&D and hardware costs; however, every donor-funded pilot must include a clear "exit strategy" that transitions operational responsibility to local entities.
- *Public-Private Partnerships (PPPs):* Collaborations where the government provides regulatory oversight while private tech vendors or NGOs manage the technical infrastructure and cloud maintenance.
- *Government Financing:* Transitioning technology costs into national health budgets ensures that the system is viewed as a permanent public utility rather than a temporary project.
- *User Fees and Insurance Integration:* Exploring models where AI-enabled dispensers are integrated into national health insurance schemes; the high quality of care and data-backed diagnostic accuracy can justify reimbursement for services rendered at the retail level.
- *Revenue-Generating Data Insights:* While maintaining strict privacy, aggregated and anonymized epidemiological data can be valuable for pharmaceutical supply chain optimization or national disease surveillance programs, potentially creating a self-sustaining revenue stream.

Part 3: Design and Deployment

6. Technology System Design

6.1 Engaging Users in Human Centered Design

Human Centered Design (HCD) activities can provide a wealth of information that allows for the co-development of technology systems with end users. The purpose of these activities are to understand the unique challenges that users (i.e. drug dispensers at ADDOs and pharmacies) face when doing their jobs, as well as to explore potential solutions together. HCD activities should be done in an open, honest, and friendly setting – one where they are not being corrected but instead free to share the reality of their situation.

HCD activities come in many types, however *Improv-Based Scenarios* offer a unique opportunity to understand the experience of the user and to find out how they handle specific situations in their drug shop.

Participants take on different roles (dispenser, customer, etc.), come up with a scenario, and act it out – just like they would if they were in a drug shop. To increase the complexity and observe edge-case scenarios, a facilitator can contribute to the scenario by giving additional directions (for example: "the client wants medication that you know they should not have"). The scenarios can be both realistic and unrealistic.

Facilitators document the process and outcomes of each scenario. They look for information on:

- How dispensers handle clients who are in a hurry
- How dispensers handle clients who refuse to take a certain medicine
- Whether or not dispensers ask a client questions to see if their drugs are going to have a bad interaction
- How dispensers handle stock reordering
- How dispensers find out what medicine to give a client

These insights are used to develop workflows and technology system features.

Box 3: Example Learnings from Human Centered Design Activities during the AfyaTek Program

The Inspired Ideas Research Foundation deployment of Elsa DOTS in Tanzania utilized Human-Centered Design (HCD) to move beyond theoretical models and align the

technology with the daily realities of drug dispensers. By conducting improv-based scenario observations and open feedback sessions, we identified critical behaviors that directly reshaped our tool's design and implementation strategy.

Interactions Related to Dispensing Medicines

- *Visual Identification of Drugs:* Clients often refer to drugs by color, shape, or local names rather than brand names or ingredients. This can impact communications between the client and dispenser; if a medication does not match the description provided by the patient, they might be dissatisfied with the service provided.
- *Partial Dosing*: Patients often buy only one or two days' worth of medicine due to limited funds or a belief that treatment ends when symptoms fade.
    - *Business Logic:* DOTS was built to support "per-tablet" sales and inventory tracking, ensuring dispensers can maintain accurate stock records while accommodating the "partial dose" economic model.
- *Medication Misuse*: Some clients use medications incorrectly, such as applying oral amoxicillin powder directly to wounds.
    - *Decision Support*: DOTS provided patient counseling prompts in the local language (Swahili) to help dispensers effectively explain the correct administration of medications.

Stock and Stock Management

- *Counterfeit & Error Risks*: Fake medications and manufacturing errors (e.g., crystals instead of powder) threaten shop credibility.
    - *Quality Monitoring*: Future iterations of DOTS or other digital health tools could allow dispensers to flag specific batches or suppliers, creating a digital audit trail that helps identify and mitigate risks from unreliable wholesalers.
- *Capital Constraints:* Small drug shops / ADDOs often have limited capital and cannot buy in bulk, leading to higher costs and frequent sourcing from alternative suppliers.
    - *Operational Support:* DOTS included "low stock" alerts to optimize small-scale purchasing. Future iterations of the tool could include connecting dispensers with wholesale suppliers directly through the

application to automatically send medication requests and purchase orders, helping to maintain essential medicine availability.

Navigating Peer Influence and Client Preferences

A significant challenge for ADDO dispensers is that many clients prioritize the advice of peers—who share "what worked" for them—over the clinical recommendations of the health provider. Consequently, patients often arrive with their minds already made up, making it difficult for the dispenser to suggest a more appropriate or effective medication. This behavior presents a strategic opportunity to utilize a digital platform to educate the client; by visually presenting AI-validated recommendations, the dispenser can leverage the "objective" authority of the technology to counter anecdotal advice and justify the correct course of treatment.

Strengthening Safety: Allergies and Drug Interactions

Prescribing medication safely requires a rigorous assessment of potential side effects, such as drug allergies and drug-drug interactions. Currently, dispensers rely heavily on the client's memory and their own pharmacological knowledge, typically asking if a patient has experienced past reactions (most commonly to sulfur) or what other medications or foods they are consuming. HCD feedback emphasized that dispensers find it incredibly valuable for the system to provide real-time alerts regarding the severity of potential interactions during the dispensing process. By automating this check, the technology moves beyond simple record-keeping to provide a critical safety net that compensates for human memory limitations.

6.2 Technology System Design

After engaging users in human centered design activities, developers can design AI tools with the user's environment and constraints in mind. The following provide important considerations during the technology system design phase:

- *User-Centered Design for Low-Literacy Contexts:* Interfaces should be intuitive, minimizing complex text in favor of recognizable icons and logical workflows that mirror a standard patient consultation. It can be helpful to use current physical forms or commonly-used mobile applications as a design reference. See Box 4 for more details.
- *Multimodal Interfaces:* To accommodate varying literacy levels, systems should support voice-guided prompts, visual aids for identifying symptoms, and text in local languages.

- *Offline-First Architecture:* Because network connectivity is often intermittent, the AI decision engine must reside locally on the mobile device, allowing the dispenser to complete full assessments and receive recommendations without an active internet connection.
- *Cultural and Contextual Appropriateness:* The tool must respect local health beliefs and common community practices, ensuring that the AI's communication style is empathetic and culturally familiar. It must also take into account drug dispenser knowledge and context. For example, a health care provider in Dar es Salaam might consider malaria as the cause of an individual's illness more frequently than a provider in Arusha, given the local prevalence in each location.

Box 4: Recommendations for Building in Low-Tech and Low-Resource Contexts

To ensure technology is an asset rather than a burden in drug outlet settings, systems must be designed for the reality of the user's environment. The following recommendations are made based on experience in the field working with low-tech-literate users in limited-resource contexts.

For the Digital System:

- *User-Centered Design:* Prioritize intuitive interfaces that require minimal training and accommodate users with varying education levels.
- *Multimodal Interfaces:* Use a mix of icons, voice prompts, and simple text to ensure the tool is accessible to low-literacy populations.
- *Offline-First Architecture:* Clinical logic and decision support must reside on the device so the tool remains functional during frequent network outages.
- *Workflow Integration*: The digital tool should mirror the natural flow of a patient consultation to avoid increasing the time burden on the provider.
- *Cultural and Contextual Appropriateness:* Use local languages (e.g., Swahili) and respect community health beliefs to foster trust between the dispenser and the patient.

For AI and Decision Support:

- *Transparency and Reasoning:* Avoid "black box" outputs; show the logic behind a recommendation (e.g., "Malaria testing suggested due to fever in a high-transmission area").

- *Confidence Levels:* Present recommendations with visual indicators of strength or confidence to help the user weigh the advice.
- *Optimize Interpretability:* Use familiar visual metaphors to represent probabilities and uncertainty for those unfamiliar with how these are discussed.
- *Empower Human Authority:* Allow easy overrides of AI suggestions with optional documentation to respect the dispenser's contextual knowledge.
- *Flexibility in Treatment:* Provide alternative options or dosages when first-line medications are out of stock.
- *Strategic Escalation:* Clearly indicate when a case is too ambiguous or severe for the outlet level and must be escalated to a supervisor or hospital.

6.3 Hardware Specifications

Technical specifications for hardware are important to determine - or constrain - from the beginning. It is helpful to identify the following prior to development, as these factors will affect the software and AI models that can be used:

- Minimum device requirements
- Network and connectivity standards
- Security protocols
- Interoperability standards

Example minimum hardware requirements for a mobile phone-based device include::

- RAM: 4GB, or higher
- Network: HSPA (3G)
- CPU: Quad-core 1.2 GHz
- Internal Memory: 16GB
- Dimensions: 5.5” diagonal & 2.5” wide
- Android Version: 11 (Red Velvet Cake)
- GPS: Present
- Connectivity: 3G/4G capability for periodic sync

Box 5: System Design and Data Flow for the Elsa DOTS Application

The Elsa DOTS platform is offline first, where all computation and operations occur directly on the device. The remote servers are available through a cloud API endpoint that allows for user authentication, data synchronization, updating decision support models, and reporting use cases (Figure 1).

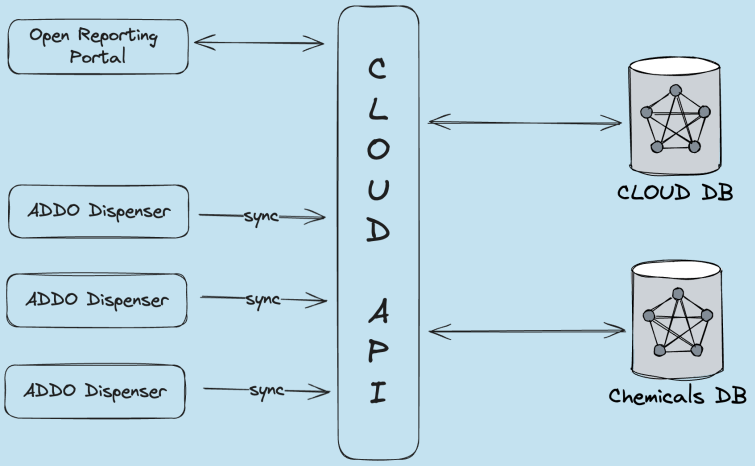

Figure 1: Components and data flow, representing how ADDO dispensers can upload their data to the cloud.

Drug dispenser users only need access to the internet when they are attempting to synchronize data, or to update their decision support algorithms and applications in general. Otherwise, the mobile applications are built as entire stand alone applications that can exist without ever connecting to the remote servers after initial authentication.

A high-level representation of the components of the mobile application is illustrated below:

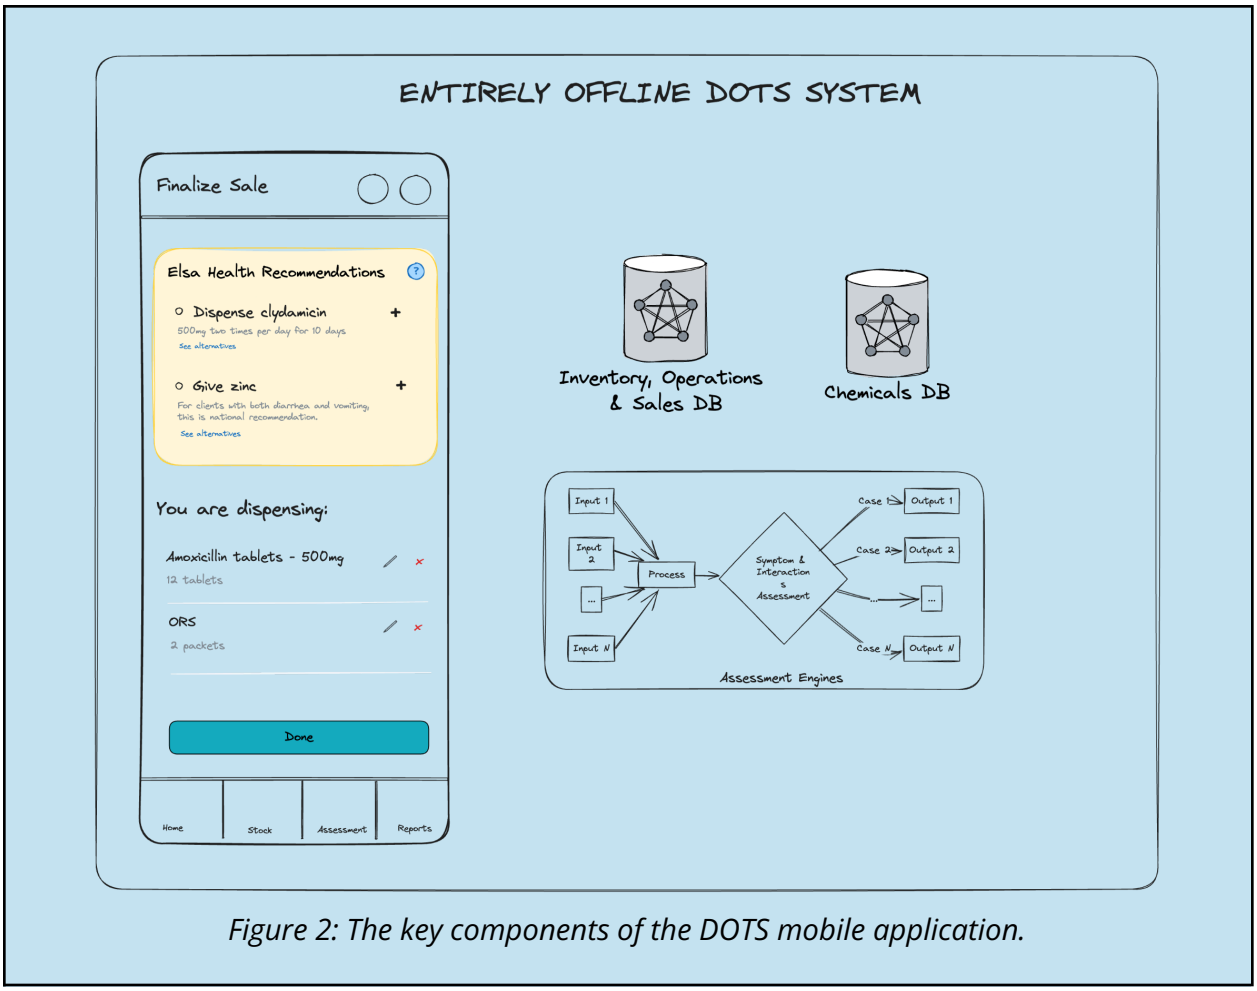

Figure 2: The key components of the DOTS mobile application.

7. AI Algorithm Design

This section provides a technical framework for AI models within drug outlets. Design choices must balance diagnostic accuracy with the unique constraints of resource-limited settings, such as intermittent connectivity and varying levels of provider expertise.

7.1 Model Archetypes: Generative ("Probabilistic") vs Deterministic ("Curve-fitting")

Depending on the goals of the digital health tool (see section #3), two primary archetypes dominate current health decision support in drug outlets: **Generative** and **Deterministic Models**

Generative Models (“Probabilistic”)

Generative models learn the underlying probability distribution of the data, allowing them to reason about uncertainty, handle incomplete information, and explain why a recommendation was made. Instead of simply mapping inputs to outputs, these models build an internal representation of how symptoms, conditions, and patient factors relate to each other.

Bayesian Networks (and others with the same operating principles, SEMs, HMB, etc)

Bayesian Networks represent clinical knowledge as a network of random variables (symptoms, signs, diseases) connected by probabilistic relationships. The probability of a given node is calculated as the posterior probability of a disease given the observed values of the other nodes.

Example: A diagnosis is calculated as the posterior probability of a disease given observed symptoms.

- *Best Use Case:* Differential diagnosis for conditions (e.g., malaria vs. dengue, or pneumonia vs gastroenteritis) where reliable, transparent probability estimates are required and evidence can be presented for the reasoning. These types of models also work really well with established protocols such as national/international guidelines for operating procedures. It is a good mix of what we know and what we can learn from data.
- *Strengths:* Highly interpretable, can function entirely offline, and provide calibrated confidence levels for each diagnosis. The network structure itself serves as documentation of clinical reasoning, and could easily be developed and adjusted by trained experts. Bespoke models like this also offer higher degrees of resistance to bad training data (missing conditions, misdiagnoses, inappropriate treatments, etc).
- *Requirements:* Requires a curated clinical knowledge base or high-quality structured datasets, typically developed in collaboration with medical experts.

These types of models, can also be referred to as “bespoke” models where experts create the foundational structures and data can be used.

Large Language Models (LLMs)

At the time of this writing (2025-2026), LLMs are the current “trending” technology being explored for health care and decision support. LLMs leverage vast amounts of medical literature and can process natural language inputs directly. In clinical settings, they are most safely deployed using Retrieval-Augmented Generation (RAG), which "grounds" the model's responses in a trusted medical corpus such as national treatment guidelines.

- *Best Use Case:* Complex medical queries, natural language based inputs and outputs, text based explanations, generating patient counseling scripts in local languages, and summarizing patient histories.
- *Strengths:* Excellent at processing unstructured text, handling conversational inputs, and generating human-readable explanations.
- *Limitations:* Significant risk of "hallucinations" (generating plausible but incorrect medical facts), high computational requirements, and typically require persistent internet connectivity for cloud-based inference. Even with RAG, outputs require careful validation before clinical use.

Hallucinations Note: Authors do understand that technically all output generated by LLMs is a “hallucination”, here we say “hallucination” to mean information that is not included in the training material and could be inaccurate

RAG Implementation Note: When deploying LLMs with RAG for drug outlets, the retrieval corpus must be strictly limited to validated sources such as the Tanzania Standard Treatment Guidelines or WHO recommendations. The model should never generate medical advice from open web content. Additionally, implementers should build in safeguards that flag when the model's confidence is low or when the query falls outside the scope of the reference corpus.

Deterministic Models ("Curve-fitting")

Deterministic models learn a direct mapping from inputs to outputs, essentially fitting a function to training data without explicitly modeling the underlying probability distribution. Given the same inputs, they will (almost) always produce the same output.

Decision Trees, Random Forests, and Gradient Boosting

Tree based models partition the input space based on feature thresholds, making sequential decisions that lead to a classification or prediction (regression tasks).

- *Best Use Case:* Well-defined classification tasks with complete input data, such as risk stratification or triage decisions where all required data points are collected upfront.
- *Strengths:* Fast inference suitable for resource-constrained devices, lower computational requirements, predictable outputs, and relatively easy to validate exhaustively against test cases. These models also offer support for non-linearities and resistance to outliers.
- *Limitations:* Do not naturally provide calibrated confidence estimates, struggle when input data is incomplete or ambiguous, and the reasoning path may be difficult to communicate to end users.

Rule-Based Expert Systems

These systems encode clinical logic as explicit if-then rules, often derived directly from treatment guidelines.

- *Best Use Case:* Enforcing protocol compliance, checking contraindications, and implementing mandatory referral triggers for "red flag" symptoms.
- *Strengths*: Completely transparent and auditable, easy to update when guidelines change, and guaranteed to behave consistently.
- *Limitations:* Cannot generalize beyond encoded rules, require significant manual effort to build and maintain, and may become unwieldy as clinical complexity increases.

Often underestimated, rule-based systems can produce extremely desirable conditions where tradeoffs are made consciously and compliance is baked in. They can also become extremely complex as in the case of our drug interactions model that encodes rules as mathematical structures and relationships using OLOGs (Ontology Logs, founded in Category Theory) and produce extremely reliable performance.

Choosing Between Archetypes

The choice between generative and deterministic approaches depends on several factors:

- *Data completeness:* If dispensers will frequently have incomplete symptom information, generative models handle this more gracefully.
- *Explainability requirements*: If dispensers need to explain recommendations to clients or supervisors, models with transparent reasoning (Bayesian Networks, rule-based systems) are preferable.
- *Connectivity constraints*: For fully offline operation, on-device Bayesian Networks or deterministic classifiers are more practical than cloud-dependent LLMs.
- *Task specificity:* For narrow, well-defined tasks (e.g., pediatric dosing calculations), deterministic approaches may be simpler and equally effective. For broader diagnostic support, generative models offer more flexibility.

In practice, many implementations benefit from combining approaches; this could look like using deterministic rules for safety-critical checks while employing probabilistic models for diagnostic reasoning.

The key rule of thumb is to use the lowest complexity (simplest) models for the current task at hand. A strategy we employ is purposely using the simplest and easiest implementation as our baseline and are often surprised by how often it beats state of the art models.

7.2 Technical Requirements for Model Building

Data Selection and Pre-processing

The quality of AI decision support is directly proportional to its training data. Different model architectures have distinct data requirements, and these vary wildly so it is important to pay attention to each:

Probabilistic/Generative Models (Bayesian Networks et al):

Requires structured symptom-to-disease mappings based on local epidemiological data and expert physician consensus. The main considerations are:

- *Expert knowledge:* To form the foundational structure of the network, or when historical data is sparse, structured interviews with clinicians can provide the necessary probability distributions. In an ideal world (never the case), this alone is enough to build the best model possible for most problems.
- *Local clinical and prevalence data*: Disease probabilities should reflect location/regional realities. A model trained on global data may overweight conditions rare in the deployment context while underweighting locally endemic diseases. Partnerships with local clinics and ministries of health can help solve this problem.

Probabilistic models often allow for simple and continuous updating of beliefs, probabilities should be updated as surveillance data accumulates from deployed systems.

For Generative Models (LLMs with RAG):

This approach requires a carefully curated and "locked" medical corpus. Models should never generate medical advice from open web content. To understand why the open web is full of dangerous medical information, simply refer to any popular public forum to witness the horrors.

Here our considerations are:

- *Corpus curation:* Limit retrieval sources to validated documents such as national Standard Treatment Guidelines, WHO recommendations, and peer-reviewed literature.
- *Version control:* Track which version of guidelines the model was trained or grounded on, and establish protocols for updating when guidelines change.
- *Local language considerations:* If the corpus is primarily in English but deployment is in Swahili, translation quality and medical terminology consistency must be validated.

For Deterministic Models (Curve fitters):

The gold standard is labeled training datasets where inputs (symptoms, patient factors) are mapped to verified outputs (diagnoses, appropriate treatments).

We consider:

- *Label quality:* Training labels should come from confirmed diagnoses (ideally with lab confirmation), not presumptive treatment records. This is absolutely essential, and is a deal breaker.
- *Feature completeness:* Many deterministic models perform poorly with missing inputs, so training data should reflect realistic data collection scenarios.
- *Class balance:* Oversample rare but critical conditions or use your judgement to pick appropriate weighting to prevent the model from defaulting to common diagnoses.

Explainability and Auditability

Drug shop dispensers, like many healthcare personnel, are often not familiar with numerical probabilities or statistical concepts. Models must be paired with interface designs that make reasoning transparent and actionable.

- *Representing Uncertainty:* Rather than displaying raw percentages, represent confidence through familiar metaphors:
  - Traffic light systems (green/yellow/red) for risk levels
  - Ranked lists with visual weight indicators
  - Natural language hedging ("most likely," "also consider," "unlikely but serious")
- *Explaining Reasoning:* Dispensers are more likely to trust and appropriately use AI recommendations when they understand why a suggestion was made:
  - Show which symptoms contributed most to a diagnosis
  - Highlight missing information
- *Audit Trails:* Every recommendation should be logged with sufficient detail to reconstruct the decision:
  - Input data (symptoms reported, patient demographics)
  - Model version and parameters used
  - Output produced (differential diagnoses, confidence levels, treatment suggestions)
  - Dispenser action (accepted, modified, or overrode recommendation)
  - Timestamp and user identifier

These logs serve both quality improvement and accountability purposes, enabling retrospective review of cases where outcomes were poor.

Handling Uncertainty and Bias

Clinical Ambiguity:

When symptoms point to multiple conditions with similar probabilities, the system should not force a definitive diagnosis. Appropriate responses include:

- Recommending laboratory testing to differentiate between possibilities
- Suggesting referral for clinical examination beyond the dispenser's scope
- Presenting multiple possibilities with clear guidance on distinguishing features to monitor

The goal is to support appropriate clinical humility and acknowledging the limits of symptom-based assessment without diagnostic tools.

Class Imbalance:

In drug shop settings, some conditions (malaria, upper respiratory infections, skin rashes) are far more common than others (meningitis, severe pneumonia). Without proper constraints, models will over-predict common conditions and miss rare but serious ones.

- Use stratified sampling or class weighting during training
- Leverage existing clinical knowledge to supplement the data
- Set decision thresholds that prioritize sensitivity for dangerous conditions over overall accuracy
- Implement "red flag" overrides that trigger regardless of model output when specific danger signs are present - these should be provided by each country’s ministry of health.

Demographic and Geographic Bias:

Models trained on data from one population may perform poorly on another. Model validation must be performed across the following, among other metrics you may have internally:

- Age groups (pediatric dosing and disease presentation differ significantly)
- Geographic regions (urban vs. rural, highland vs. coastal) - often even a 20 minute drive in any direction results in a completely different epidemiology
- Seasonal variations (malaria transmission patterns, respiratory illness peaks)

Benchmarking

Rigorous benchmarking establishes baseline performance and enables ongoing quality monitoring. Models often experience either a context drift or a model drift, benchmarking can help spot these and should occur at multiple stages:

Pre-Deployment Validation:

Before field deployment, models should be tested against:

- *Gold standard clinical vignettes:* Standardized case descriptions with expert-consensus diagnoses, covering common conditions, rare-but-serious conditions, and ambiguous presentations.
- *Clinician comparison:* Model accuracy compared to decisions made by qualified healthcare providers on the same case set.
- *Sensitivity analysis:* Performance measured across different levels of input completeness (what happens when key symptoms are missing?).

Deployment Monitoring:

Once live, ongoing benchmarking should track:

- *Recommendation acceptance rate:* How often do dispensers follow AI suggestions? Large deviations may indicate trust issues or model inaccuracy.
- *Override patterns:* Which recommendations are most frequently overridden? This may reveal systematic model weaknesses or areas needing retraining.
- *Outcome correlation:* Where follow-up data is available, track whether AI-recommended treatments led to patient improvement or return visits.

Benchmark Metrics:

Select metrics appropriate to the clinical context:

- *Sensitivity/Specificity:* Particularly for danger sign detection, high sensitivity is critical even at the cost of specificity.
- *Calibration*: When the model says 80% confidence, is it correct 80% of the time? Poorly calibrated models undermine appropriate clinical decision-making.
- *Top-k accuracy:* For differential diagnosis, measure whether the correct condition appears in the top 3 or top 5 suggestions, not just the top 1.

7.3 The Hybrid Archetype (Recommended for Scale)

In practice, the most robust systems for health centers and facilities combine multiple model types, leveraging the strengths of each while mitigating their limitations.

A recommended architecture pairs:

- *Core Decision Support Engine (Generative/Probabilistic):* A *bespoke* probabilistic model (Bayesian Network or similar probabilistic) handles differential diagnosis and risk stratification. This component runs entirely offline, provides calibrated confidence levels, and offers transparent reasoning that dispensers can communicate to patients.
- *Safety & Compliance Layer (Deterministic/Rule-Based):* Hard-coded rules enforce critical safety checks, mandatory referral triggers for danger signs, contraindication alerts, and dosing limits. These rules override probabilistic outputs when patient safety is at stake.
- *Optional Unstructured Layer (LLM-based):* Where connectivity allows, an LLM agent can supplement the core system for tasks like generating patient counseling scripts in local languages, explaining complex drug interactions, or answering dispenser questions about unfamiliar presentations. This layer enhances usability but is not required for core clinical function. **This is the least studied approach as of this writing so care must be taken.**

Box 6: Elsa DOTS Decision Support Algorithms

Drug Interactions

Drug interactions are a very common occurrence at the point of sale. This is because identifying and understanding how different drugs interact takes years of training and a constantly high level of vigilance. These shortcomings are exactly what algorithms and computers can excel at. Inspired Ideas Research Foundation developed a graph database that describes a growing number of common medications, their chemical ingredients, and the hierarchical relationships between them. We then use this graph to compute possible interactions between any 2 (or more) drugs. This representation allows us to include any findings from research, no matter how the categorizations are described.

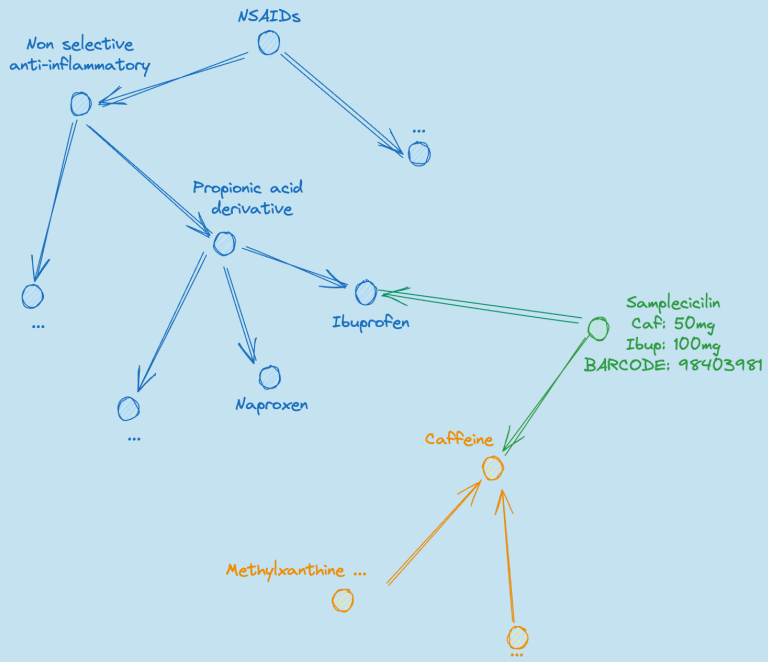

Figure 3: A visualization of the two different hierarchies of chemicals (blue and orange) and an example drug called Samplecicilin (green).

Symptom Assessment

The DOTS application includes decision support models to support symptom assessment and next steps recommendations. The models are built as an *ensemble of Bayesian methods (Bayesian Networks) and custom similarity measures* that compare the incoming data to known and documented encounters for all covered conditions. These methods allow for interactivity between the dispensers and the algorithms by responding to new information as it becomes available.

The inputs for symptom assessment are demographic information (sex, age, location, etc), signs and symptoms, laboratory investigation results where available, and other relevant patient history. The models then produce an ordered list of possible conditions based on the symptom and patient assessment, along with the uncertainties for each condition. The output also includes next steps recommendation, including referrals as needed.

The models are not designed as black boxes that take inputs and output recommendations; rather, they are built to be explainable, interpretable, and even adaptive and responsive to physician updates and thought processes. The platform relies heavily on probabilistic models, like the example shown below. These models

make probabilistic inference a more tractable goal and allow us to build a dynamic “Question Workflow” that asks the most relevant next question.

This is a part of the model that shows the relationship between multiple nodes representing different factors. From this network, it's easy to see the direction of influence that could allow a drug dispenser to work backwards from “chest pain” to “lung disease”.

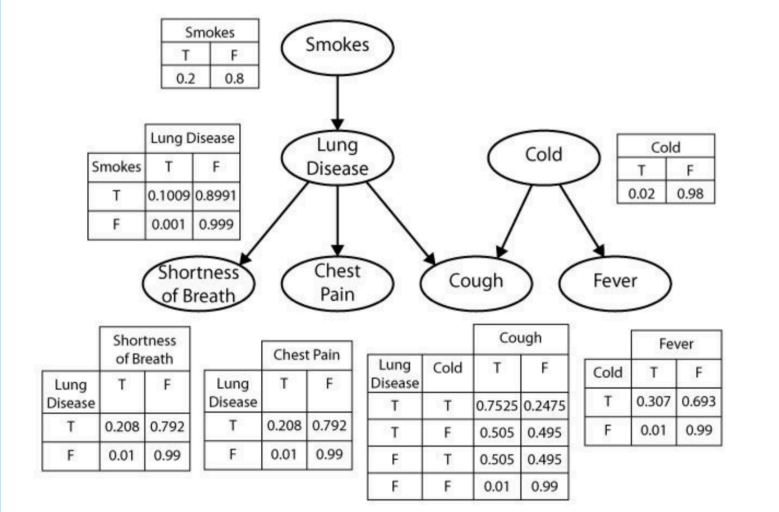

8. Clinical Governance

Clinical governance provides the framework to ensure that AI-powered health decision support systems operate safely, ethically, and in strict alignment with national health standards. This is essential in the drug outlet context, where providers have limited formal clinical training and must rely on technology as a reliable extension of official health policy.

8.1 Encoding National Treatment Guidelines

The logic within any AI health tool must be a digital distillation of national standards, such as the Tanzania Standard Treatment Guidelines or WHO SMART Guidelines. Guidelines are converted into "computable artifacts" through a multi-step process:

- *Parsing*: Recommendation statements are extracted from official text and restructured into logical "if-then" decision paths.
- *Standardized Value Sets:* Clinical concepts (e.g., "fast breathing") are mapped to international medical codes to ensure interoperability and consistent interpretation.
- *Workflow Integration*: Algorithms are designed to follow the natural patient consultation flow, identifying "triggers" that match specific clinical criteria.

- *Version Control and Updates*: Health protocols are dynamic. The system must support "over-the-air" updates to instantly push new protocols (e.g., a change in first-line antimalarials) across the entire drug shop network, preventing the use of outdated medical advice.
- *Ministry of Health Validation:* Clinical algorithms should be validated by the relevant national health authority, where possible. This could include a side-by-side comparison of AI outputs against official paper-based guidelines to ensure 100% adherence.

8.2 Safety and Quality Assurance (QA)

A robust QA process protects patients and maintains the credibility of the drug shop within the community. The following are recommendations for how the AI system can incorporate safety and quality throughout the development and deployment process.

- *Red Flag and Mandatory Referral*: The system should have hard-coded "red flags" for danger signs—such as convulsions, lethargy, or severe respiratory distress—that trigger a mandatory referral recommendation to a higher-level facility. It is important to work closely with the users of the digital tool to highlight the importance of referral at this stage, as business priorities can take precedence over referral recommendations.
- **Accuracy Evaluation Protocols:**
  - *"Gold Standard" Review:* Periodically, a panel of expert clinicians should review anonymized patient cases to compare the AI's recommendations against expert clinical judgment.
  - *Real-World Monitoring:* Systems should track the frequency and rationale for human overrides. High override rates for a specific condition may indicate a logic error in the algorithm or a significant local medicine stock-out.
- *Bias and Equity Audits:* Algorithms must be tested across diverse patient groups (age, gender, rural vs. urban) to ensure that the AI does not perpetuate existing health disparities due to biased training data.

8.3 Defining the Scope of Practice

AI allows drug outlets to handle complex cases more safely, but it must not be used to expand the dispenser's scope of practice beyond what is legally permitted. The tool must clearly state what it cannot advise on, such as surgical interventions or complex chronic disease management. Additionally, the system should only offer treatment

recommendations for medicines on the approved National Essential Medicines List for that specific facility level.

As evidence from projects like AfyaTek continue to show that AI can improve clinical performance, governments may choose to re-evaluate and selectively expand the permitted scope for tech-enabled drug shops (e.g., allowing specific antibiotics only when the AI confirms diagnostic criteria).

---

9. Data Governance and Privacy

Protecting patient data is a legal requirement and a fundamental prerequisite for building community trust. In the ADDO context in Tanzania, governance must navigate the balance between digital record-keeping and the stringent protection of sensitive health information.

9.1 Data Privacy Framework

A robust framework for drug outlets focuses on the rights of the patient as the data subject. If collecting Personally Identifying Information, explicit, informed consent must be obtained from patients (or guardians for minors) before any data entry. Consent should be documented digitally within the tool and include a clear explanation of what data is collected and how it will be used. Only data essential for clinical decision-making or mandated reporting should be collected. Excessive data collection increases security risks and burdens the dispenser. Additionally, data intended for secondary use—such as training AI models or public health surveillance—must be strictly de-identified to remove PII. At rest, all stored data must be encrypted using industry standards. Data transmission between the mobile device and cloud servers must use secure protocols like TLS 1.3 to prevent interception.

9.2 Tanzania-Specific Compliance

Implementers must align with the evolving regulatory landscape in Tanzania, specifically the Personal Data Protection Act (PDPA) of 2022 and the Electronic and Postal Communications Act (EPCA).

- *PDPA Compliance (2024 Directives):* As of December 31, 2024, all private and public institutions processing personal data must be registered with the Personal Data Protection Commission (PDPC). The Act enforces the "right to be forgotten" and requires fresh consent if data is repurposed beyond its original clinical intent.

- *Online Content and AI Regulations:* Under the Electronic and Postal Communications (Online Content) (Amendment) Regulations, 2025, providers must ensure that AI-generated clinical content is not unethical, fabricated, or misleading.
- *Ministry of Health Data Policies*: All digital health interventions must align with the Ministry of Health's Digital Health Strategy, which emphasizes national data sovereignty and the hosting of sensitive health data on local servers.

9.3 Data Use Beyond Care Delivery

Aggregated, anonymized data from drug outlets can transform individual patient interactions into systemic public health insights. For example, real-time syndromic data (e.g., clusters of fever or cough) can serve as an early-warning system for outbreaks before they reach formal health facilities. Frameworks should prioritize the principle that health data ultimately belongs to the patient, with the drug shop and the implementing organization serving as "data controllers" responsible for its safe stewardship.

Part 4: Training and Change Management

10. High-Touch Training Program

A successful training program must bridge the gap between traditional medicine dispensing and technology-enabled decision support. It is important to remember that this is the starting point for an entirely new way of doing business, which can often be a difficult transition for dispensers.

10.1 Suggested Training Methodology and Format

Initial training is most effective when conducted in person using a cohort model. This allows for peer learning and immediate troubleshooting of technical hurdles. The training should be located in a central location that everyone can easily access. Training stipends can be provided to users for transportation and / or food.

- *Intensive Short-Duration Curriculum*: A 1-2 day intensive session should cover everything from basic smartphone operation to the nuances of interpreting AI recommendations. It's important to not assume that basic technical skills are present; ensure that all users know how to do things like log in to their phone, access the internet, download an app, and use their camera.
- *Structured Practice:* Training should progress from group practice and role-playing to individual "solo" assessments to build confidence before the dispenser returns to their drug shop.
- *Post-Training Field Support:* The 2-3 days immediately following training are critical. Implementation teams should conduct on-site follow-up visits at each drug shop to address real-world challenges, such as physical setup or initial patient interactions.
- *Hands-on Operational Support:* For complex features like stock management, dispensers often require "hand-holding" to perform initial data entry, such as inputting current medicine inventory into the system.

10.2 Digital Literacy and Language Considerations for Training

Training must be adapted for varying levels of technology literacy, ensuring that even those unfamiliar with smartphones can navigate the tool. To reinforce learning, develop and provide video tutorials that dispensers can review independently after the initial training to

refresh their memory on specific functions. All training materials and the AI interface itself should be in the local language (e.g., Swahili) to avoid any ambiguity in clinical care.

11. Stakeholder Engagement

For AI to be accepted and integrated into the health system, engagement must occur at every level—from high-level government officials to the patients receiving care. This ensures that the technology is understood not as a replacement for human judgment, but as a collaborative tool for better health outcomes.

Government and Regulatory Bodies

- *Early Alignment:* Early engagement with national bodies, such as the Tanzania Medicines and Medical Devices Authority (TMDA) or the Ministry of Health, is essential to ensure the program remains legally compliant and aligned with national digital health strategies.
- *Building Trust:* Demonstrating AI accuracy through pilot data helps build the necessary trust for long-term institutionalization and policy integration.

Drug Shop Owners

Owners are primarily motivated by business success. They must see tangible value, such as improved inventory tracking, reduced stock-outs, and increased patient foot traffic, to justify the time and resources spent on dispenser training. Clearly communicating how the tool aids in profit maximization and shop credibility helps reduce initial resistance from owners.

Drug Shop Dispensers

Dispensers are the primary users and the most critical link in the AI implementation chain. Their active engagement determines whether the tool is used effectively at the point of care. Box 7 describes an incentive program designed to maintain dispenser engagement.

- *Role Emphasization:* It is vital to frame the AI as a "clinical aid" or that enhances their expertise rather than threatening their role.
- *Supportive Environment:* Engagement is highest when dispensers feel supported through continuous feedback loops and accessible technical assistance.
- *Communication Training:* A key part of engagement is teaching dispensers how to describe the tool to clients. They must be able to present it as a sophisticated assistant that helps them provide more thorough and targeted care.

Communities and Patients

The ultimate success of any digital health program depends on patient acceptance. It is recommended to conduct community outreach to explain that the AI tool is a "clinical assistant" designed to improve the quality of care. Campaigns should proactively answer questions about data privacy and reassure the community that the human provider remains the final authority. In our Tanzania experience, we observed that patients often feel the assessment is more "targeted specifically for their needs" when technology is involved, which can actually increase their faith in the provider.

Box 7: Example Incentives Program for High Engagement

This is an example of an incentives program for users of a new technology system that can be implemented to increase usage and engagement. Incentives can be used at the start of a new technology deployment to:

1. Improve Research Quality
2. Increase Technology Adoption
3. Ensure Sustainable Implementation

Example Performance Metrics

Proposed Primary Metrics	
Area	**Metric(s) - What will we measure?**
System Improvement Contributions	• Number of feedback submissions • Number of bug reports or system improvement suggestions • Quality of feedback provided (rated by our team, scale 1-5) • Participation in user surveys or interviews
Consistent App Engagement	• Regular daily app opens (e.g., at least once per business day) • Number of products registered
Proposed Secondary Metrics	
App/ Decision Support	• Number of completed symptom assessments

Usage	● Number of sales made ● Number of changes to the prescribed medications
Peer Support	● Mentoring/ engaging with other users ○ Sharing best practices, troubleshooting

Example Incentive Plan and Timeline

Frequency	We Review:	Measured by (see above):	If users are engaged, we offer:
One Time	Product Registration	● Number of products registered ● % of total products registered	● Equipment (blood pressure machines, etc.)
Weekly	Basic Engagement	● System Improvement Contributions (feedback)+30 or more products registration	● Bundles for personal phones ● Small cash transfer
Monthly	Consistent Quality Engagement	● System Improvement Contributions ● Consistent App Engagement ● Quality of Clinical Decision Support Usage	● Larger airtime or data packages ● T-shirt, other merch ● Engagement certificate

Quarterly (3 months)	Excellence and Leadership	• System Improvement Contributions • Consistent App Engagement • Quality of Clinical Decision Support Usage • Peer Support	• Medical equipment (BP machines, thermometers) • Shop improvement grants • Professional development opportunities • Laptop/desktop computer

Part 5: Post-Deployment Operations

Deployment is not the end of the implementation journey, but the beginning of a continuous cycle of support and refinement. This phase ensures the AI tool remains functional, accurate, and financially viable over time.

12. Maintenance and Technical Support

12.1 Supportive Supervision

Post-deployment, it is recommended to shift from intensive training to a model of supportive supervision where dispensers have a direct line to developers or implementers.

It can be beneficial to establish helpdesks accessible via common tools like WhatsApp, phone calls, or SMS to provide immediate technical and clinical assistance.

In pilot or research settings, feedback loops supervision touchpoints should be used to collect qualitative data on tool performance and user challenges. Use the feedback collected during supervision to continuously refine the tool's interface or logic, ensuring it remains responsive to the dispenser's real-world needs.

12.2 Software and Hardware Maintenance

In order to push changes to all users at the same time, and without as little friction as possible, implement "over-the-air" update protocols to deploy bug fixes or clinical guideline changes without requiring physical visits to every drug shop. Additionally, it is recommended to plan for hardware maintenance, including repair services and a replacement cycle for mobile devices, which typically have a lifespan of 2-3 years in these environments.

13. Monitoring, Evaluation, and Learning

For AI interventions, Monitoring and Evaluation (M&E) is more than just reporting; it is a clinical safety requirement. Because many implementation teams may lack deep technical expertise in AI auditing, we recommend a multi-disciplinary approach. This involves

forming a Quality Assurance Committee that includes software developers, medical doctors, government regulators, and frontline dispensers to review data from multiple angles.

13.1 Key Performance Indicators (KPIs)

To understand if the AI is truly improving care, it is recommended to track data across four distinct domains:

Clinical Outcomes	Monitor diagnostic accuracy (how often the AI's suggestion matches the final diagnosis), treatment appropriateness (adherence to the drug list), and referral rates (identifying if "red flags" are being caught).
Operational Metrics	Track uptake (the percentage of total shop visitors who are entered into the AI tool), user satisfaction among dispensers, and system uptime (to measure the impact of power or network outages).
Economic Metrics	Calculate the cost per patient encounter and monitor the revenue impacts for drug shop owners to ensure the tool doesn't harm their business viability.
Health System Metrics (if available or relevant)	Evaluate real-time disease surveillance patterns and the overall rate of national guideline adherence across the network.

13.2 Evaluation Methodologies: How to Measure Impact

Overall Technology Deployment

Before deploying the AI, baseline data of all metrics can be collected or curated in order to have a point of comparison. This might include data such as current dispensing practices (e.g., how often antibiotics are currently misprescribed) or diagnostic accuracy. During evaluation, it is common to use a control group, where possible, of drug dispensers alongside the AI-enabled dispensers to see if the technology is the actual cause of improved health outcomes. When possible, it is recommended to conduct interviews with patients and dispensers to understand the "why" behind the data—such as why a dispenser might choose to override an AI recommendation.

AI Accuracy and Validity

The validation of AI models varies depending on the underlying technology and the purpose of the system. Regardless, it is important to establish a continuous validation loop that includes gold standard review and real-world performance monitoring. See section 7.2 for additional information on model benchmarking and evaluation. To reiterate:

- *Gold Standard Review:* A multi-disciplinary panel of expert clinicians should regularly review a random sample of anonymized cases. They compare the AI's recommendation against what a senior doctor would have done.
- *Real-World Performance Monitoring:* Systems should flag cases where the dispenser overrode the AI. These "overrides" often highlight where the algorithm is failing to account for local context.

Box 8: Example of Evaluation Methods for Elsa DOTS Application

Qualitative User Feedback

We solicited qualitative feedback from users throughout the implementation period of the project. This happened both through the application (a form to provide direct and quick feedback), through a dedicated Whatsapp group, and through monthly meetings/ interviews with the dispensers. These feedback sessions helped inform the technical and user requirements, as well as the performance of the application.

Our M&E framework guided our discussions and data collection. The following are some of the metrics we analyzed through user feedback:

- Usability metrics (friendliness, ease of use, etc)
- Time it takes to utilize the tool for one client
- Usage differences between rural and urban users (ie user and client behaviors)
- User perceptions of validity and reliability
- User perceptions on helpfulness in decision making
- Performance of the models with and without connectivity

Evaluation of Application

Throughout development and continuous iteration of the DOTS application, we conducted:

- **Quality Assurance (QA) Testing:** After each major development update, we conducted QA testing with our team and a set of testers internally. The QA testers provide feedback on the usability of the application, the appropriateness of the content, and the accuracy of the translations. They also provide a sanity check for the models to make sure that nothing is glaringly incorrect.
- **Unit Tests:** We maintain >80% test coverage of all the technical components.
- **Integration Tests:** We support automated integration tests for all major user stories.

Evaluation of Models

Decision support models were built in collaboration with clinicians / pharmacists and were tested by a Quality Assurance Team of researchers and specialists before they went into production.

One of the ways we evaluated the models was through agreeability with expert healthcare professionals. For this process, we recruited physicians to look at a set of generated patient data and provide their expert advice on the top conditions the patient is likely to have or the drug interactions likely to take place. We then averaged those decisions across physicians and compared with the output generated by our models. This provides us with an agreeability score.

For conditions where we had testing data available, we also evaluated the models using typical machine learning techniques where we measure accuracy, MCC, F1 score, sensitivity (evidence, parameters, and structure) and Area Under the Curve.

Part 6: Scaling and Sustainability

Scaling an AI intervention from a pilot to a national program requires addressing structural, behavioral, and technical risks. This phase focuses on institutionalizing the tool within the national health framework while maintaining the quality of care.

14. Common Implementation Challenges & Mitigations

The following table outlines potential real-world challenges and strategies that can be used to address them.

Challenge Category	Description of Challenge	Mitigation and Adaptation Strategy
Client Perceptions	Patients may view digital tools as benefiting the provider or implementer rather than themselves, leading to skepticism.	Conduct community sensitization campaigns to explain the tool's value as a clinical assistant for better care.
Staff Turnover	High turnover in drug shops creates a constant need for new training.	Develop self-guided video tutorials and "training-of-trainers" models.
Model Explainability	"Black box" AI makes it difficult for dispensers to understand or explain recommendations to patients.	Use HCD to create interpretable models that show "why" a suggestion was made in a collaborative way.
Over-reliance	Risk of dispensers deferring blindly to AI without critical thinking.	Include case reviews of dispensers not using AI to maintain clinical reasoning skills.
Language Barriers	Discrepancies between formal medical terminology (English/Swahili) and "lay" language used by dispensers.	Continuously update translations based on user feedback to align with local dialects and dispenser vocabulary.
Workflow Disruption	High-intensity digital tools can disrupt the patient-provider	Build communication structures into the tool (visual aids, shared

	relationship or make consent processes feel burdensome.	screens) to facilitate rather than hinder history-taking.
Private Sector Motivation	Private providers are often driven by profit, which may conflict with rational medicine use guidelines.	Highlight how technology increases provider credibility and foot traffic, aligning business goals with public health impact.
Offline Sync & Connectivity	Users often struggle to sync data regularly, and internet bundles are frequently depleted for non-project use.	Explore "zero-rating" apps with local telcos and automate data sync protocols to protect data from being lost on damaged devices.
High-Touch Onboarding	High staff turnover in drug shops creates a constant, resource-intensive need for new training.	Move toward a "training-of-trainers" model and provide offline video tutorials for self-guided onboarding.
Hardware Risks	Lost, stolen, or damaged devices are difficult to replace on limited program budgets.	Budget for device insurance or replacement funds and work with owners to establish clear security protocols for devices.

15. Government Stewardship and Institutionalization

As the project expands, the focus must shift from technical functionality to ethical responsibility and long-term ownership. Sustainability is only achieved when the project stops being "external". Recommendations to achieve this include:

- *Integration into HIS:* The AI tool should feed data directly into the national Health Information System (HIS) to support disease surveillance.
- *Local Technical Capacity:* Shift technical maintenance from international developers to local Ministry of Health IT teams or domestic technology partners.
- *Policy Alignment:* Work with regulators to include AI-assisted dispensing in the formal scope of practice for drug outlets.

Appendix A: Example Regulatory and Compliance Checklist

This checklist can be used by developers or organizations who are engaging with national regulatory bodies to ensure all legal and safety guardrails are addressed early. Additional items should be added as needed.

- ☐ **Administrative & Legal Buy-in**
  - ☐ *Government Liaison:* Has a dedicated team member been assigned to manage relationships with the Ministry of Health and regulatory authorities?
  - ☐ *Memorandum of Understanding (MOU):* Is there a signed agreement defining the roles of the implementer vs. the government?
  - ☐ *Professional Liability:* Does the policy clearly state that the human dispenser remains the final decision-maker and holds legal accountability?

- ☐ **Clinical Safety & Validation**
  - ☐ *Guideline Alignment:* Have all AI algorithms been audited against the most recent National Standard Treatment Guidelines?
  - ☐ *"Red Flag" Logic:* Does the system have hard-coded, mandatory referral triggers for life-threatening symptoms?
  - ☐ *Expert Review Panel:* Is there a multi-disciplinary committee (doctors, pharmacists, tech leads) established to review AI accuracy?

- ☐ **Data Privacy & Security**
  - ☐ *Patient Consent*: Is the informed consent process integrated into the digital workflow?
  - ☐ *Compliance Audit:* Does the data storage and transmission architecture meet national laws?
  - ☐ *Data Minimization:* Have unnecessary data fields been removed to protect patient anonymity?

Endnotes

1. https://apotheker.or.tz/afya-tek/
2. Accredited Drug Dispensing Outlets (ADDO) | The Center for Health Market Innovations https://wayback.archive-it.org/13606/20231107020433/https://healthmarketinnovations.org/program/accredited-drug-dispensing-outlets-addo
3. Understanding the Role of Accredited Drug Dispensing Outlets in Tanzania's Health System | PLOS One https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0164332
4. Embrey M, Vialle-Valentin C, Dillip A, Kihiyo B, Mbwasi R, Semali IA, et al. (2016) Understanding the Role of Accredited Drug Dispensing Outlets in Tanzania's Health System. PLoS ONE 11(11): e0164332. https://doi.org/10.1371/journal.pone.0164332
5. Minzi, O., & Manyilizu, V. (2013). Application of basic pharmacology and dispensing practice of antibiotics in accredited drug-dispensing outlets in Tanzania. Drug, healthcare and patient safety, 5, 5–11. https://doi.org/10.2147/DHPS.S36409
6. Embrey M, Vialle-Valentin C, Dillip A, Kihiyo B, Mbwasi R, Semali IA, et al. (2016) Understanding the Role of Accredited Drug Dispensing Outlets in Tanzania's Health System. PLoS ONE 11(11): e0164332. https://doi.org/10.1371/journal.pone.0164332