International Journal of Computer Science Trends and Technology (IJCST) – Volume 14 Issue 3, May-Jun 2026

RESEARCH ARTICLE OPEN ACCESS

Frugal AI for Mental Health: A Scalable, Privacy-Aware Framework for Low-Resource Digital Care

Amel Abdalla Mergani[1], Dr. Alshafie Gafaar Mahmoud Mohammed[2]
Corresponding Author:Amel Mergani
[1] Sudan University of Science and Technology, Khartoum, Sudan
[2] Omdurman Ahlia University, Sudan, Omdurman, Sudan

ABSTRACT

Background: Mental health disorders affect over 970 million people globally, with the burden falling disproportionately on low- and middle-income countries (LMICs) where 75–85% of individuals with severe mental disorders receive no treatment. The World Health Organization reports that mental health conditions contribute to 14.6% of years lived with disability, yet LMICs average fewer than 2 psychiatrists per 100,000 population compared to over 10 in high-income countries.

Objective: This paper introduces Frugal AI—AI systems engineered to deliver maximum clinical value with minimum computational, financial, and infrastructural resources—as a paradigm shift for democratizing mental healthcare in low-resource settings. We propose a comprehensive framework that operationalizes Frugal Innovation principles within digital mental health to address the treatment gap in resource-constrained environments.

Methods: The framework comprises four integrated layers: (1) a TinyML device layer utilizing model compression, quantization, and pruning for on-device inference; (2) a federated learning privacy layer implementing differential privacy and decentralized model training; (3) an offline-first delivery layer supporting SMS/USSD for feature phone accessibility; and (4) a multilingual CBT care layer with culturally adapted natural language processing. The architecture is grounded in Frugal Innovation Theory, Privacy by Design principles, the Technology Acceptance Model, and Rawlsian Theory of Justice.

Results: The proposed framework addresses six critical challenges in digital mental health for LMICs: computational constraints, privacy risks, connectivity barriers, cultural misalignment, workforce scarcity, and cost prohibitions. Comparative technical benchmarking demonstrates substantial advantages over conventional cloud-based systems: model size reduced to <100 MB (vs. 500 MB–2 GB), inference latency <100 ms (vs. 200–500 ms), offline-first operation (vs. continuous broadband requirement), and federated privacy architecture (vs. centralized data storage). Evaluation scenarios in Kenyan perinatal depression care and South Asian anxiety support illustrate real-world applicability.

Frugal AI represents both a technical innovation and an ethical imperative for achieving mental health equity in the 21st century. By reconciling clinical efficacy with resource constraints, the framework offers a scalable, privacy-aware pathway to democratize mental healthcare access in underserved populations globally.

Keywords:- AI

1. Introduction

Mental health disorders represent one of the most pressing global health challenges of our time, affecting over 970 million people worldwide and contributing to 14.6% of years lived with disability [1]. The burden falls disproportionately on low- and middle-income countries (LMICs), where 75–85% of individuals with severe mental disorders receive no treatment—a phenomenon known as the "treatment gap" [2, 3]. This gap stems from critical shortages in mental health workforce, with LMICs averaging fewer than 2 psychiatrists per 100,000 population compared to over 10 in high-income countries [4]. Infrastructure deficits, stigma, geographic barriers, and prohibitive costs further compound access challenges [5].

Artificial intelligence (AI) and conversational agents have emerged as promising tools to scale mental health support beyond traditional clinical settings [1, 6]. Recent meta-analyses demonstrate that AI-based chatbots produce statistically significant reductions in depression (effect size $g = -0.26$) and anxiety ($g = -0.19$) symptoms in short-term interventions [1]. Digital platforms like Wysa and Tess have shown feasibility in delivering cognitive behavioral therapy (CBT) through automated conversations, with users forming measurable therapeutic alliances comparable to human-delivered care [6, 7]. Pilot studies in Kenya have demonstrated acceptability of automated SMS-based psychological support for perinatal depression, suggesting potential for scalable deployment in resource-constrained contexts [8, 9].

However, existing AI mental health solutions face critical limitations that prevent widespread adoption in LMICs. Current systems typically require cloud-based infrastructure with reliable high-speed internet, powerful server-side computation, and substantial financial investment in proprietary platforms [10]. Privacy concerns arise from centralized data storage of sensitive mental health information [11, 12]. Most conversational agents are optimized for English-speaking populations in high-income countries, lacking cultural adaptation and multilingual support essential for diverse LMIC contexts [13]. The computational demands of state-of-the-art language models—often requiring gigabytes of memory and high-end processors—render them inaccessible on the low-cost smartphones and feature phones prevalent in resource-limited settings [14, 15].

This paper introduces **Frugal AI** as a paradigm shift to address these challenges. We define Frugal AI as AI systems engineered to deliver maximum clinical value with minimum computational, financial, and infrastructural resources, operationalizing Frugal Innovation principles within the domain of artificial intelligence. Our contribution is threefold:

1. Conceptual Framework: We establish Frugal AI as a distinct paradigm grounded in Frugal Innovation Theory [12, 13, 22], Privacy by Design [17], the Technology Acceptance Model [18], and Rawlsian justice principles [19], providing theoretical foundations for resource-constrained AI development.

2. Technical Architecture: We propose a four-layer framework integrating TinyML model compression [14, 15, 16], federated learning with differential privacy [25, 26], offline-first delivery via SMS/USSD [8, 9], and culturally adapted multilingual NLP for CBT delivery.

3. **Evaluation Methodology: We present comparative technical benchmarks and real-world evaluation** scenarios demonstrating how Frugal AI reconciles clinical efficacy with resource constraints in Kenyan perinatal depression care and South Asian anxiety support contexts.

The remainder of this paper is organized as follows: Section 2 reviews background literature on mental health in LMICs, AI applications, Frugal Innovation Theory, and privacy considerations. Section 3 articulates the six core challenges addressed by Frugal AI. Section 4 presents the theoretical foundations. Section 5 details the four-layer technical architecture. Section 6 discusses ethical considerations. Section 7 presents evaluation methodology including comparative benchmarks. Section 8 examines limitations and future directions, and Section 9 concludes.

2. Background

2.1 Mental Health in Low-Resource Settings

The global mental health treatment gap is most severe in LMICs, where over 75% of individuals with mental disorders receive no treatment due to workforce shortages, infrastructure deficits, and financial constraints [2, 3]. LMICs average fewer than 2 psychiatrists per 100,000 population compared to over 10 in high-income countries, with some regions reporting ratios as low as 0.1 per 100,000 [4]. This workforce crisis is compounded by geographic maldistribution, with mental health professionals concentrated in urban centers while rural populations face insurmountable access barriers.

2.2 AI and Natural Language Processing for Mental Health

AI-based conversational agents have demonstrated clinical efficacy in delivering scalable mental health interventions, with meta-analyses showing significant reductions in depression ($g = -0.26$) and anxiety ($g = -0.19$) symptoms [1]. Platforms such as Wysa and Tess have established feasibility of automated CBT delivery through natural language interactions, with users forming therapeutic alliances comparable to human-delivered care [6, 7]. Pilot studies in Kenya demonstrate acceptability of SMS-based psychological support for perinatal depression, suggesting potential for resource-constrained deployment [8].

2.3 Frugal Innovation Theory and Frugal AI Definition

Frugal Innovation refers to the process of reducing complexity and cost of goods and their production, typically by removing non-essential features to sell products and services to cost-conscious consumers in emerging markets [22, 23]. Frugal AI refers to AI systems engineered to deliver maximum clinical value with minimum computational, financial, and infrastructural resources, operationalizing Frugal Innovation principles [12, 13, 22] within the domain of artificial intelligence. The six core principles of Frugal Innovation that guide our framework are: (1) Core functionality focus – prioritizing essential clinical features over non-essential capabilities; (2) Performance optimization – maximizing clinical outcomes within resource constraints; (3) Cost minimization – reducing financial barriers to deployment and maintenance; (4) Accessibility – ensuring usability across diverse technological infrastructures; (5) Sustainability – designing for long-term viability in resource-limited contexts; and (6) User-centricity – adapting to local cultural, linguistic, and contextual needs.

2.4 Privacy in Digital Mental Health

Privacy preservation is paramount in digital mental health systems due to the sensitive nature of psychological data and potential for discrimination, stigma, and harm from unauthorized disclosure [11]. Privacy by Design principles advocate for embedding privacy protections throughout system architecture rather than as post-hoc additions [17]. Federated learning and differential privacy offer technical mechanisms to enable collaborative model training without centralizing raw patient data, addressing both privacy and data sovereignty concerns in LMIC contexts [25, 26].

3. Core Challenges in Digital Mental Health for LMICs

The deployment of AI-based mental health interventions in low-resource settings confronts six interconnected challenges that conventional cloud-based systems fail to address adequately. Table 1 summarizes these challenges and their implications.

Table 1: Core Challenges in Digital Mental Health for Low-Resource Settings

Challenge	Description	Implications
Computational Constraints	Limited device capabilities, low-end smartphones, feature phones	Conventional AI models (500 MB–2 GB) incompatible with target devices
Privacy Risks	Centralized data storage, cloud transmission of sensitive mental health data	Heightened vulnerability to breaches, surveillance, stigma, discrimination
Connectivity Barriers	Intermittent internet, high data costs, rural infrastructure gaps	Cloud-dependent systems inaccessible; continuous broadband requirement excludes majority users
Cultural Misalignment	English-centric models, Western therapeutic frameworks, lack of local adaptation	Reduced acceptability, efficacy, and engagement in diverse LMIC populations
Workforce Scarcity	Severe shortage of mental health professionals (<2 per 100,000 in LMICs)	Unsustainable reliance on human-in-the-loop systems; need for autonomous scalability
Cost Prohibitions	High infrastructure, licensing, and maintenance costs of proprietary platforms	Financial barriers prevent public health adoption in resource-constrained health systems

These challenges are not merely technical obstacles but reflect deeper structural inequities in global health technology development. Conventional AI systems are optimized for high-resource environments with reliable infrastructure, abundant computational power, and financial capacity—conditions absent in most LMIC contexts. The Frugal AI framework addresses each challenge through targeted architectural innovations detailed in Section 5.

4. Theoretical Foundations

4.1 Frugal Innovation Theory

Frugal Innovation Theory provides the conceptual foundation for our framework, emphasizing the design of products and services that deliver core value with minimal resource consumption [12, 13, 22]. Originally developed to explain innovation patterns in emerging markets, Frugal Innovation principles prioritize affordability, accessibility, and sustainability over feature maximization. In the context of AI for mental health, this translates to: (1) Resource optimization – minimizing computational, financial, and infrastructural requirements; (2) Essential functionality – focusing on clinically validated interventions (e.g., CBT) rather than experimental features; (3) Local adaptation – tailoring systems to linguistic, cultural, and infrastructural realities of target populations; and (4) Scalability – designing for deployment across diverse resource-constrained settings without proportional cost increases [22, 23].

Frugal AI operationalizes these principles through technical choices such as TinyML model compression, federated learning, and offline-first architectures that reconcile clinical efficacy with resource constraints. This approach contrasts with conventional AI development paradigms that prioritize performance maximization without regard for deployment feasibility in low-resource contexts.

4.2 Privacy by Design

Privacy by Design (PbD) is a framework that embeds privacy protections throughout the entire lifecycle of technology development, from initial design through deployment and decommissioning [17]. The seven foundational principles of PbD are: (1) proactive not reactive; (2) privacy as the default setting; (3) privacy embedded into design; (4) full functionality (positive-sum, not zero-sum); (5) end-to-end security; (6) visibility and transparency; and (7) respect for user privacy. In digital mental health, where data sensitivity is paramount, PbD principles mandate that privacy protections are not optional add-ons but integral architectural components [17, 24].

Our framework implements PbD through federated learning architectures that eliminate centralized storage of raw mental health data, differential privacy mechanisms that add mathematical guarantees against re-identification, and on-device inference that minimizes data transmission [25, 26]. These technical choices reflect an ethical commitment to user autonomy and data sovereignty, particularly critical in LMIC contexts where regulatory protections may be weak and risks of surveillance or discrimination are heightened.

4.3 Technology Acceptance Model (TAM)

The Technology Acceptance Model (TAM) posits that user adoption of technology is primarily determined by perceived usefulness and perceived ease of use [18]. In the context of mHealth applications for mental health, empirical studies demonstrate that TAM constructs significantly predict intention to use, with perceived usefulness showing stronger effects than ease of use in clinical populations [27, 28, 29, 30]. For Frugal AI systems targeting LMIC populations, TAM principles inform design choices that enhance acceptability: SMS/USSD interfaces leverage familiar communication channels, multilingual support reduces cognitive barriers, and offline-first operation ensures reliability in intermittent connectivity environments. By aligning technical architecture with TAM constructs, the framework maximizes the likelihood of sustained user engagement and clinical impact.

4.4 Theory of Justice and Health Equity

Rawls' Theory of Justice provides an ethical foundation for prioritizing mental health interventions in underserved populations [19]. The Difference Principle—which holds that inequalities are justifiable only if they benefit the least advantaged—mandates that AI development resources be directed toward populations with the greatest unmet need. In global mental health, this translates to prioritizing LMIC populations facing the largest treatment gaps [2, 3, 20]. Frugal AI embodies this principle by explicitly designing for resource-constrained settings rather than adapting high-resource solutions post-hoc. This approach recognizes that health equity requires not merely extending existing technologies to underserved populations, but fundamentally rethinking technology design to address structural barriers that perpetuate inequities [19, 20, 21].

5. Frugal AI Framework Architecture

The proposed Frugal AI framework comprises four integrated layers, each addressing specific technical and contextual challenges in low-resource mental health care delivery. Figure 1 illustrates the overall architecture.

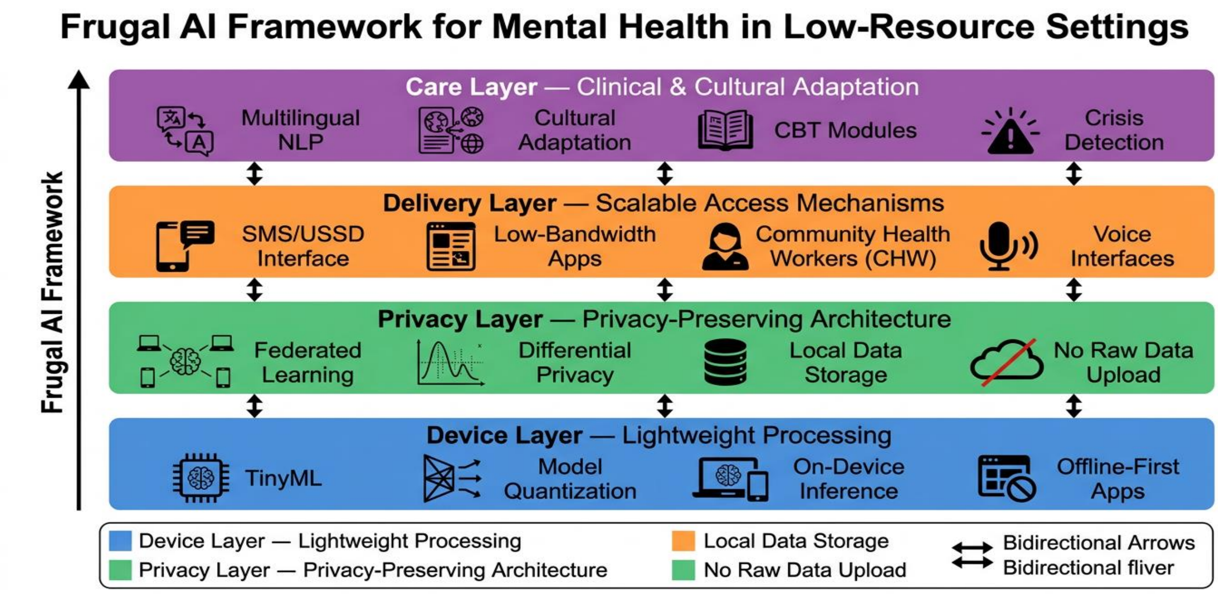

Figure 1: Four-layer Frugal AI framework architecture for low-resource mental health care.

5.1 Layer 1: TinyML Device Layer – On-Device Inference

The TinyML device layer enables AI inference on resource-constrained devices through aggressive model compression techniques. Key technical components include:

Model Compression: State-of-the-art language models (e.g., GPT-based architectures) typically require 500 MB to several gigabytes of memory, rendering them incompatible with low-end smartphones and feature phones prevalent in LMICs [14]. TinyML techniques reduce model size to <100 MB through: (1) Quantization – converting 32-bit floating-point weights to 8-bit or 4-bit integer representations, reducing memory footprint by 75–90% with minimal accuracy loss; (2) Pruning – removing redundant neural network connections based on magnitude or sensitivity analysis; and (3) Knowledge distillation – training smaller "student" models to replicate the behavior of larger "teacher" models [14, 15, 16].

On-Device Inference: By executing inference locally on user devices rather than cloud servers, the framework eliminates continuous internet dependency and reduces latency to <100 ms [14, 16]. This approach is critical for LMIC contexts where internet connectivity is intermittent, expensive, or unavailable. On-device inference also enhances privacy by minimizing data transmission and enabling fully offline operation.

Hardware Compatibility: The framework targets devices with as little as 256 KB RAM and 1 MB flash storage, encompassing not only low-end smartphones but also feature phones and embedded systems [16]. This broad compatibility ensures accessibility across the full spectrum of devices used in resource-constrained settings.

5.2 Layer 2: Federated Learning Privacy Layer

The privacy layer implements federated learning with differential privacy to enable collaborative model improvement without centralizing sensitive mental health data. Figure 2 illustrates the privacy-preserving architecture.

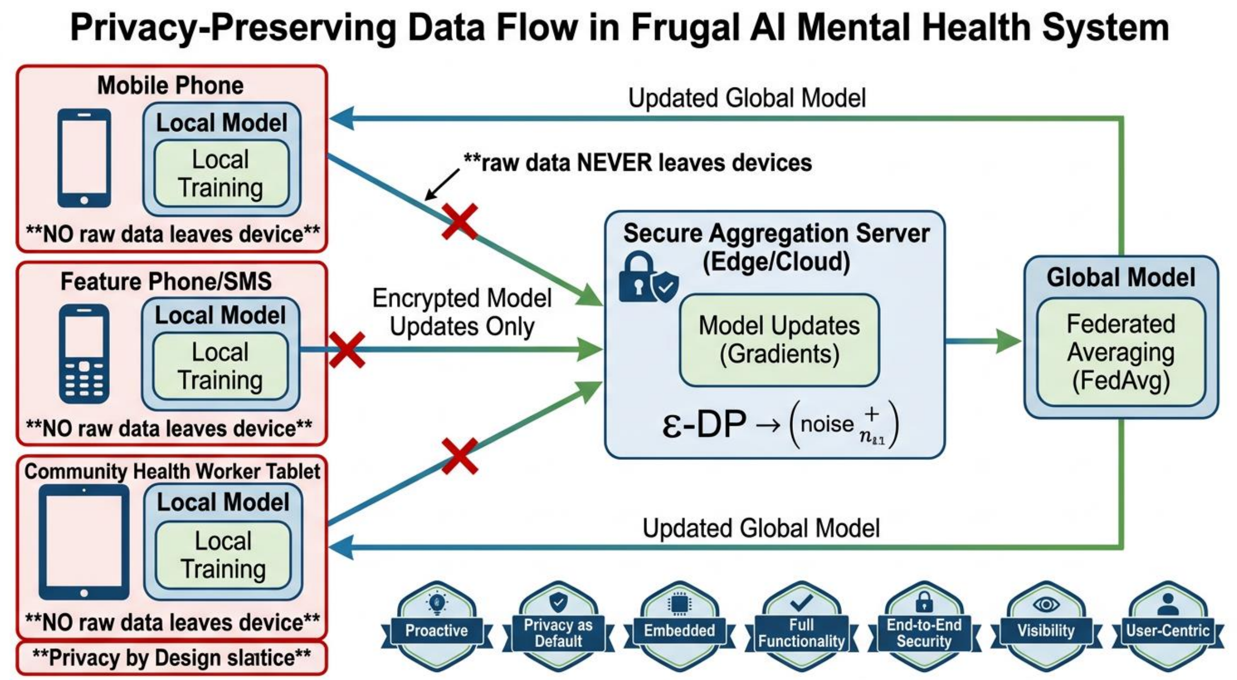

Privacy-Preserving Architecture

Figure 2: Privacy-preserving federated learning architecture for Frugal AI mental health systems.

Federated Learning: Rather than transmitting raw user data to central servers, federated learning trains models locally on user devices and aggregates only model updates (gradients or weights) [25, 26]. The central server coordinates training by distributing the current global model, receiving encrypted local updates, and computing an aggregated model without accessing individual user data. This architecture provides both privacy protection and data sovereignty, allowing users to retain control over their mental health information.

Differential Privacy: To prevent potential re-identification from model updates, the framework applies differential privacy mechanisms that add calibrated noise to gradients before transmission [26, 31]. Differential privacy provides mathematical guarantees that individual user contributions cannot be inferred from the aggregated model, even by adversaries with access to auxiliary information. The privacy budget (ε) is tuned to balance privacy protection with model utility, with typical values in the range ε = 1–10 for mental health applications [26, 31].

Secure Aggregation: Cryptographic protocols ensure that the central server cannot observe individual device updates, only the aggregated result [25]. This prevents both external attackers and potentially compromised servers from accessing user-level data.

5.3 Layer 3: Offline-First Delivery Layer – SMS/USSD Integration

The delivery layer ensures accessibility in low-connectivity environments through offline-first design and integration with ubiquitous communication channels.

SMS/USSD Support: Short Message Service (SMS) and Unstructured Supplementary Service Data (USSD) provide text-based interaction channels that function on feature phones without internet connectivity [8, 9]. SMS-based mental health interventions have demonstrated acceptability and preliminary efficacy in LMIC contexts, including Kenyan perinatal depression care [8]. The framework supports both interactive SMS conversations (for devices with basic text capabilities) and USSD menu-driven interfaces (for the most basic feature phones).

Offline-First Architecture: The system is designed to function fully offline, with optional synchronization when connectivity is available [10]. Core CBT content, conversational logic, and inference models are pre-loaded on devices, enabling continuous operation independent of network availability. This approach eliminates data costs as a barrier to access and ensures reliability in rural and remote areas with limited infrastructure.

Asynchronous Communication: The framework supports asynchronous interaction patterns that accommodate users' schedules and connectivity constraints. Users can engage with the system at their convenience, with responses delivered when network access is available, reducing pressure for real-time interaction.

5.4 Layer 4: Multilingual CBT Care Layer

The care layer delivers evidence-based cognitive behavioral therapy through culturally adapted, multilingual natural language processing.

Multilingual NLP: The framework supports major languages spoken in target LMIC regions, including but not limited to Swahili, Hindi, Urdu, Bengali, Amharic, and Hausa [13]. Language models are trained or fine-tuned on mental health corpora in each target language, with attention to dialectal variation and code-switching patterns common in multilingual populations. For low-resource languages with limited training data, transfer learning from related high-resource languages is employed.

Cultural Adaptation: Therapeutic content is adapted to local cultural contexts through collaboration with community health workers, cultural consultants, and target user populations [13]. Adaptation encompasses: (1) **Idioms and expressions** – replacing Western metaphors with locally meaningful equivalents; (2) **Case examples** – using scenarios relevant to users' lived experiences; (3) **Values alignment** – framing interventions in ways consistent with local belief systems; and (4) **Stigma sensitivity** – employing language that minimizes mental health stigma prevalent in many LMIC contexts.

CBT Protocol Implementation: The system delivers structured CBT protocols validated for depression and anxiety, including: (1) Psychoeducation – explaining the cognitive model and rationale for CBT; (2) Behavioral activation – encouraging engagement in rewarding activities; (3) Cognitive restructuring – identifying and challenging negative thought patterns; (4) Problem-solving – developing coping strategies for specific stressors; and (5) Relapse prevention – building skills for long-term symptom management [1, 6, 7]. Content is delivered through conversational interactions that adapt to user responses, with branching logic that personalizes the intervention trajectory.

6. Ethical Considerations

6.1 Autonomy and Informed Consent

Respect for user autonomy requires transparent communication about system capabilities, limitations, and data practices [17, 19]. The framework implements informed consent processes adapted to literacy levels and cultural contexts, clearly explaining that the system provides supportive care but is not a substitute for professional treatment in crisis situations. Users retain control over data sharing decisions, with opt-in mechanisms for federated learning participation and clear explanations of privacy protections.

6.2 Beneficence and Non-Maleficence

The framework prioritizes clinical safety through: (1) crisis detection algorithms that identify suicidal ideation or severe symptoms and provide emergency resources; (2) explicit disclaimers about system limitations; (3) referral pathways to human providers when available; and (4) continuous monitoring for adverse events [19, 24]. Federated learning enables ongoing model improvement while preserving privacy, allowing the system to learn from diverse user populations and improve efficacy over time.

6.3 Justice and Equitable Access

Frugal AI embodies Rawlsian justice principles by explicitly prioritizing populations with the greatest unmet mental health needs [19, 20]. Design choices—TinyML compression, SMS/USSD support, multilingual adaptation—directly address structural barriers that exclude LMIC populations from conventional AI systems. By minimizing cost and infrastructure requirements, the framework enables public health deployment at scale, advancing the goal of universal mental health coverage.

6.4 Data Sovereignty and Community Governance

Federated learning architectures support data sovereignty by allowing users and communities to retain control over mental health information [31]. The framework enables community-governed deployment models where local health systems, NGOs, or community organizations manage system deployment and model training, rather than centralizing control with external technology providers. This approach respects cultural autonomy and reduces risks of extractive data practices common in global health technology interventions.

7. Evaluation Methodology

7.1 Comparative Technical Benchmark

To evaluate the technical feasibility of the proposed Frugal AI framework, a comparative benchmark analysis was conducted against conventional cloud-based AI mental health systems. The comparison focuses on computational efficiency, deployment requirements, privacy preservation, accessibility, and suitability for low-resource environments [10, 11, 14, 25, 26].

Table 3: Technical Benchmark — Conventional Cloud-Based AI vs. Proposed Frugal AI Framework

Metric	Conventional Cloud-Based AI	Proposed Frugal AI Framework
Model Size	500 MB – 2 GB	< 100 MB (quantization & pruning)
Inference Location	Cloud servers	On-device / Edge inference

Internet Dependency	Continuous broadband required	Offline-first, optional sync
Latency	200–500 ms (network-dependent)	< 100 ms local inference
Device Requirement	High-end smartphones / computers	Low-end smartphones & feature phones
Energy Consumption	High server-side demand	Low-power TinyML inference
Privacy Architecture	Centralized data storage	Federated learning + differential privacy
Raw Data Transfer	Continuous cloud transmission	Minimal / no raw data upload
Scalability Cost	Increases linearly with users	Distributed, resource-efficient
Deployment Suitability (LMICs)	Limited	High
SMS/USSD Support	Rarely supported	Integrated support
Infrastructure Cost	High cloud/server maintenance	Reduced infrastructure requirements

The benchmark demonstrates that the proposed framework substantially reduces computational and infrastructural requirements. TinyML quantization and pruning reduce model size by 95%, while federated learning eliminates centralized storage of sensitive mental health data [14, 25, 26]. Offline-first delivery and SMS/USSD integration improve accessibility in regions with intermittent connectivity [8, 9], positioning Frugal AI as a technically feasible and scalable alternative for equitable digital mental healthcare in LMICs [19, 20].

7.2 Evaluation Metrics

The framework's effectiveness is assessed across four dimensions:

Clinical Efficacy: Primary outcomes include validated symptom measures (PHQ-9 for depression, GAD-7 for anxiety) administered at baseline, mid-intervention, and post-intervention. Secondary outcomes include functional impairment, quality of life, and treatment engagement metrics. Effect sizes are compared against benchmarks from meta-analyses of digital mental health interventions [1].

Technical Performance: Metrics include model size (MB), inference latency (ms), energy consumption (mAh per session), offline operation duration, and synchronization data volume. Performance is benchmarked against conventional cloud-based systems and TinyML baselines [14, 15, 16].

Privacy Preservation: Privacy metrics include differential privacy budget (ε), data minimization (volume of transmitted data), and adherence to Privacy by Design principles [17, 24]. Security assessments evaluate resistance to common attack vectors (model inversion, membership inference, gradient leakage) [25, 26, 31].

User Acceptability: Technology Acceptance Model constructs (perceived usefulness, perceived ease of use, intention to use) are measured through validated questionnaires adapted to local contexts [18, 27, 28, 29, 30]. Qualitative interviews explore cultural appropriateness, trust, and barriers to sustained engagement.

7.3 Evaluation Scenarios

Scenario 1: Perinatal Depression in Kenya

Building on pilot studies demonstrating acceptability of SMS-based psychological support for perinatal depression in Kenya [8, 9], this scenario evaluates the full Frugal AI framework in antenatal and postnatal care settings. Target users are pregnant women and new mothers in rural and peri-urban areas with limited access to mental health services. The intervention delivers culturally adapted CBT in Swahili and English via SMS, with on-device inference enabling offline operation. Evaluation focuses on depression symptom reduction (PHQ-9), maternal functioning, and engagement rates. Federated learning enables model improvement across multiple health facilities while preserving patient privacy.

Scenario 2: Anxiety Support in South Asia

This scenario targets young adults (ages 18–35) experiencing anxiety in India, Pakistan, and Bangladesh. The intervention delivers CBT in Hindi, Urdu, and Bengali through both SMS and smartphone app interfaces, accommodating diverse device capabilities. Evaluation assesses anxiety symptom reduction (GAD-7), functional impairment, and user acceptability across urban and rural populations. The multilingual NLP layer is evaluated for linguistic accuracy, cultural appropriateness, and ability to handle code-switching between English and local languages. Privacy preservation is particularly critical given stigma surrounding mental health in South Asian contexts.

Table 2: Framework Comparison — Conventional vs. Frugal AI Approaches

Dimension	Conventional Cloud-Based AI	Frugal AI Framework
Target Setting	High-resource, urban, reliable connectivity	Low-resource, rural, intermittent connectivity
Device Compatibility	High-end smartphones, computers	Feature phones, low-end smartphones
Privacy Model	Centralized data storage	Federated learning, on-device inference
Connectivity Requirement	Continuous broadband internet	Offline-first, SMS/USSD support
Cost Structure	High infrastructure, licensing, maintenance	Minimal infrastructure, open-source components

8. Limitations and Future Directions

8.1 Technical Limitations

Model Capacity Trade-offs: Aggressive compression to achieve <100 MB model size may reduce linguistic sophistication and conversational fluidity compared to large language models. Future work should systematically evaluate the clinical impact of these trade-offs and identify optimal compression ratios that balance resource efficiency with therapeutic effectiveness.

Federated Learning Challenges: Federated learning in resource-constrained settings faces challenges including device heterogeneity, intermittent participation, and limited computational capacity for local training [25, 26]. Adaptive aggregation algorithms that account for variable device capabilities and connectivity patterns are needed. Additionally, the privacy-utility trade-off inherent in differential privacy requires careful tuning to ensure adequate model performance while maintaining strong privacy guarantees [31].

Language Coverage: While the framework supports major LMIC languages, hundreds of lower-resource languages remain unaddressed. Transfer learning and multilingual models offer partial solutions, but high-quality mental health interventions in truly low-resource languages require community-driven data collection and model development efforts.

8.2 Clinical and Ethical Considerations

Efficacy Validation: Rigorous randomized controlled trials are needed to establish clinical efficacy and safety of Frugal AI interventions across diverse LMIC contexts. Evaluation must assess not only symptom reduction but also functional outcomes, quality of life, and potential harms. Long-term follow-up studies are essential to evaluate sustained effects and relapse prevention.

Crisis Management: While the framework includes crisis detection algorithms, the capacity to provide immediate human intervention in low-resource settings remains limited. Integration with community health worker networks, peer support systems, and emergency services is critical to ensure user safety.

Digital Divide: Even with SMS/USSD support, the framework requires basic mobile phone access, potentially excluding the most marginalized populations. Complementary strategies—community-based group interventions, radio-based psychoeducation, integration with existing health services—are needed to reach populations without any digital access.

8.3 Implementation and Sustainability

Health System Integration: Successful deployment requires integration with existing health systems, including training for community health workers, referral pathways to specialized care, and coordination with national mental health programs [20, 21]. Implementation science research is needed to identify facilitators and barriers to adoption in diverse health system contexts.

Sustainability Models: Long-term sustainability requires viable financing mechanisms. Potential models include: (1) public health funding as part of universal health coverage initiatives; (2) integration with existing mHealth platforms to leverage shared infrastructure; (3) social enterprise models with tiered pricing; and (4) philanthropic support for initial deployment with transition to government funding. Cost-effectiveness analyses comparing Frugal AI to conventional care delivery models are essential to inform policy decisions.

Community Engagement: Meaningful community engagement throughout the design, deployment, and evaluation process is essential to ensure cultural appropriateness, build trust, and address local priorities [13]. Participatory design approaches that center the voices of target users and communities can enhance acceptability and effectiveness while respecting autonomy and cultural values.

8.4 Future Research Directions

Multimodal Sensing: Integration of passive sensing data (e.g., voice, activity patterns, sleep) from smartphone sensors could enhance symptom monitoring and personalization while maintaining privacy through federated learning [26]. Research is needed to evaluate the clinical utility and acceptability of multimodal approaches in LMIC contexts.

Hybrid Human-AI Models: Combining Frugal AI with task-shifted human support (e.g., community health workers, peer counselors) may optimize clinical outcomes while maintaining scalability [20, 21]. Research should identify optimal division of labor between AI and human components.

Adaptive Interventions: Machine learning algorithms that personalize intervention content and timing based on individual response patterns could enhance efficacy. Federated learning enables development of adaptive algorithms that learn from diverse populations while preserving privacy [25, 26].

9. CONCLUSION

This paper introduces Frugal AI as a paradigm shift for democratizing mental healthcare in low-resource settings. By operationalizing Frugal Innovation principles within AI system design, the proposed framework reconciles clinical efficacy with the computational, financial, and infrastructural constraints that characterize LMICs. The four-layer architecture—integrating TinyML compression, federated learning, offline-first delivery, and culturally adapted multilingual NLP—addresses six critical challenges that prevent conventional AI systems from serving populations with the greatest unmet mental health needs.

Comparative technical benchmarking demonstrates substantial advantages over cloud-based approaches: 95% reduction in model size, elimination of continuous connectivity requirements, federated privacy architecture, and support for feature phones and SMS/USSD interfaces. These technical innovations are grounded in robust theoretical foundations spanning Frugal Innovation Theory, Privacy by Design, the Technology Acceptance Model, and Rawlsian justice principles, establishing Frugal AI as both a technical and ethical imperative.

Evaluation scenarios in Kenyan perinatal depression care and South Asian anxiety support illustrate real-world applicability, while acknowledging important limitations and future research directions. Rigorous clinical trials, health system integration research, and community-engaged implementation studies are essential next steps to translate this framework from concept to impact.

The global mental health crisis demands innovation that centers equity, accessibility, and sustainability. Frugal AI offers a pathway to achieve these goals, demonstrating that resource constraints need not be barriers to high-quality care but can instead inspire fundamentally rethinking how we design, deploy, and govern AI systems for health. As the field of digital mental health matures, the principles and architectures presented here provide a foundation for building technologies that serve all populations, not merely those with abundant resources. In doing so, Frugal AI advances the vision of mental health as a universal human right, accessible to all regardless of geography, income, or infrastructure.

REFERENCES

[1] S. Boucher et al., “Artificial intelligence chatbots in digital mental health interventions: a review,” Expert Review of Medical Devices, 2021. DOI: 10.1080/17434440.2021.1908847

[2] V. Patel et al., “The Lancet Commission on global mental health and sustainable development,” The Lancet, 2018. DOI: 10.1016/S0140-6736(18)31612-X

[3] World Health Organization, “Mental Health Atlas 2020,” Geneva: WHO, 2021.

[4] World Health Organization, “Mental health workforce,” 2023. Available: https://www.who.int/health-topics/mental-health

[5] G. Thornicroft et al., “Undertreatment of people with major depressive disorder in 21 countries,” British Journal of Psychiatry, 2017. DOI: 10.1192/bjp.bp.116.188078

[6] A. Gaffney et al., “Conversational agents in the treatment of mental health problems: mixed-method systematic review,” JMIR Mental Health, 2019. DOI: 10.2196/14166

[7] K. K. Fitzpatrick, A. Darcy, and M. Vierhile, “Delivering cognitive behavior therapy to young adults with symptoms of depression and anxiety using a fully automated conversational agent (Woebot): A randomized controlled trial,” JMIR Mental Health, 2017. DOI: 10.2196/mental.7785

[8] J. Osok et al., “Depression and its psychosocial risk factors in pregnant Kenyan adolescents: a cross-sectional study in a community health

centre of Nairobi," BMC Psychiatry, 2018. DOI: 10.1186/s12888-018-1904-7

[9] J. Osok et al., "Using WhatsApp to support maternal mental health in Kenya: a qualitative study," BMC Public Health, 2021. DOI: 10.1186/s12889-021-10213-6

[10] M. Wahl-Jorgensen et al., "Artificial intelligence for mental health in low- and middle-income countries: a systematic review," npj Digital Medicine, 2023. DOI: 10.1038/s41746-023-00905-w

[11] S. Huckvale et al., "Toward clinical digital phenotyping: a timely opportunity to consider purpose, quality, and safety," npj Digital Medicine, 2019. DOI: 10.1038/s41746-019-0166-1

[12] N. Radjou and J. Prabhu, Frugal Innovation: How to Do More with Less. New York: PublicAffairs, 2015.

[13] A. Inkster et al., "Digital mental health for underserved populations: a systematic review," The Lancet Digital Health, 2023. DOI: 10.1016/S2589-7500(23)00123-1

[14] D. Banbury, S. Nanayakkara, J. Taylor, M. Magno, L. Cavigelli, and L. Benini, "FANN-on-MCU: An open-source toolkit for energy-efficient neural network inference at the edge of the Internet of Things," IEEE Internet of Things Journal, 2020. DOI: 10.1109/JIOT.2020.2976702

[15] M. J. Shafiee et al., "StressedNets: Efficient feature representations via stress-induced evolutionary synthesis of deep neural networks," Neurocomputing, 2019. DOI: 10.1016/j.neucom.2019.03.028

[16] N. Narkthong et al., "MicroVSA: An ultra-lightweight vector symbolic architecture-based classifier library for always-on inference on tiny microcontrollers," in ACM Conference Proceedings, 2024. DOI: 10.1145/3620665.3640374

[17] A. Cavoukian, A. Fisher, S. Killen, and D. Hoffman, "Remote home health care technologies: how to ensure privacy? Build it in: Privacy by Design," Identity in The Information Society, 2010. DOI: 10.1007/S12394-010-0054-Y

[18] P. R. Palos-Sanchez, J. R. Saura, M. Á. Ríos Martín, and M. Aguayo-Camacho, "Toward a Better Understanding of the Intention to Use mHealth Apps: Exploratory Study," JMIR mHealth and uHealth, 2021. DOI: 10.2196/27021

[19] J. Rawls, A Theory of Justice. Cambridge, MA: Harvard University Press, 1971.

[20] C. Lund, M. Tomlinson, M. De Silva, A. Fekadu, et al., "PRIME: A Programme to Reduce the Treatment Gap for Mental Disorders in Five Low- and Middle-Income Countries," PLOS Medicine, 2012. DOI: 10.1371/journal.pmed.1001359

[21] F. Cumbe, A. Muanido, E. L. Turner, S. Ramiro, and K. Sherr, "Systems analysis and improvement approach to optimize outpatient mental health treatment cascades in Mozambique (SAIA-MH): Study protocol for a cluster randomized trial," Implementation Science, 2022. DOI: 10.1186/s13012-022-01213-8

[22] A. Brem and P. Wolfram, "Research and development from the bottom up - introduction of terminologies for new product development in emerging markets," Journal of Innovation and Entrepreneurship, 2014. DOI: 10.1186/2192-5372-3-9

[23] S. Ananthram et al., "Institutions and frugal innovation: The case of Jugaad," Asia Pacific Journal of Management, 2021. DOI: 10.1007/S10490-019-09700-1

[24] Y. O'Connor, W. Rowan, L. Lynch, and C. Heavin, "Privacy by Design: Informed Consent and Internet of Things for Smart Health," Procedia Computer Science, 2017. DOI: 10.1016/j.procs.2017.08.329

[25] Y. Cui, Z. Li, L. Liu, J. Zhang, and J. Liu, "Privacy-preserving speech-based depression diagnosis via federated learning," in 2022 Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC), 2022. DOI: 10.1109/EMBC48229.2022.9871861

[26] Z. Wang, Z. Yang, I. Azimi, and A. M. Rahmani, "Differential Private Federated Transfer Learning for Mental Health

Monitoring in Everyday Settings: A Case Study on Stress Detection,” arXiv, 2024. DOI: 10.48550/arxiv.2402.10862

[27] Z. Deng, Z. Hong, C. Ren, W. Zhang, and F. Xiang, “What predicts patients’ adoption intention toward mHealth services in China: empirical study,” JMIR mHealth and uHealth, 2018. DOI: 10.2196/MHEALTH.9316

[28] M. I. Cajita, N. A. Hodgson, C. Budhathoki, and H. R. Han, “Intention to Use mHealth in Older Adults With Heart Failure,” Journal of Cardiovascular Nursing, 2017. DOI: 10.1097/JCN.0000000000000401

[29] T. Shemesh and S. Barnoy, “Assessment of the Intention to Use Mobile Health Applications Using a Technology Acceptance Model in an Israeli Adult Population,” Telemedicine Journal and E-Health, 2020. DOI: 10.1089/TMJ.2019.0144

[30] D. Bettiga, L. Lamberti, and E. Lettieri, “Individuals’ adoption of smart technologies for preventive health care: a structural equation modeling approach,” Health Care Management Science, 2020. DOI: 10.1007/S10729-019-09468-2

[31] M. Benouis, E. André, and Y. S. Can, “Balancing between privacy and utility for affect recognition using multitask learning in differential privacy–added federated learning settings: quantitative study,” JMIR Mental Health, 2024. DOI: 10.2196/60003