Responsible AI for resource-constrained settings

Terms of Reference

1. Introduction

The GSMA is a global organisation unifying the mobile ecosystem to discover, develop and deliver innovation foundational to positive business environments and societal change. Our vision is to unlock the full power of connectivity so that people, industry and society thrive. Representing mobile operators and organisations across the mobile ecosystem and adjacent industries, the GSMA delivers for its members across three broad pillars: Connectivity for Good, Industry Solutions and Events.

GSMA Mobile for Development Foundation

The GSMA Mobile for Development Foundation, Inc. ("GSMA") is a U.S. 501(c)(3) charitable organisation. The Mobile for Development team (M4D) operates at the intersection of the mobile ecosystem and the development sector. Our aim is to stimulate digital innovation and deliver both sustainable business and large-scale socio-economic impact. Our research and insights platform, in-market expertise and community of partners push forward digital innovations and implementations that empower underserved populations. To date, we have impacted the lives of over 220 million people.

Emerging Tech programme

The EmergingTech programme aims to accelerate impact and climate action by fostering the adoption of AI and emerging technologies in low- and middle-income countries (LMICs). The programme works with mobile network operators (MNOs) and impact-focused innovators to design and deliver AI and emerging tech powered products and services for underserved populations.

AI for Impact (AI4I)

The GSMA AI for Impact initiative, supported by major mobile operators connecting over two billion people, is driving the responsible and transformative adoption of AI across the mobile industry.

2. Background, scope and objectives

Background

AI deployment across impact focused products and services is increasing faster than governance frameworks can keep pace in low- and middle-income countries (LMICs). Innovators are leveraging AI to solve challenges across health, financial inclusion, education, agriculture, climate and digital inclusion. (The GSMA innovation fund for impactful AI received 625 applications from 40 countries)

Organisations like the ITU, IDRC, UNESCO and the GSMA included have been actively developing frameworks, policy documents and best practices for responsible and ethical AI development. For instance, the GSMA Responsible AI Maturity Roadmap was designed for mobile operators and telco-adjacent organizations. It is a robust, globally recognized tool, grounded in ISO 42001 and NIST AI RMF. Results from the pilot survey across 15 operators reveal that even at telco scale, the industry average maturity is at the Evolving Stage, with the weakest performance in the Third-party Ecosystem and Change Management Dimension of the roadmap.

Impact driven innovators building and deploying AI at the last mile, often as the first and only AI-powered service available to the end users, operate without the institutional resources, legal teams, or governance infrastructure that the mainstream responsible AI frameworks assume. This gap is further amplified with the dependency these innovators have on hyperscalers such as OpenAI, Anthropic, Google and Meta. Another element to consider here is that most AI innovation in LMICs is happening at the application layer of the AI stack i.e. innovators are building on top of foundation models they did not train and in many cases cannot audit.

The harm that can result from irresponsible AI at this stage, such as a missed maternal health diagnosis, a discriminatory credit decision, a biased educational assessment lands on individuals who have no recourse, no alternative, and quite often no awareness that an AI system was involved at all.

Scope

The GSMA is undertaking a mixed-methods study to address evidence gaps on the feasibility, economics, and enabling conditions for responsible AI development and deployments in LMICs.

The GSMA is seeking a supplier to audit existing responsible AI frameworks against an LMIC-fit rubric, develop and validate a deployment context typology based on decision stakes and recourse availability, build a harm inventory of documented AI-related harms and practitioner-reported near-misses in LMIC deployment contexts, and map the practical constraints that prevent early-stage innovators from embedding responsible AI practices, including the influence of business models and funding structures.

Framework Architecture: The Deployment Context Typology

The framework will be organised around a Deployment Context Typology rather than a sector taxonomy or a technical layer analysis. The typology classifies AI deployments by two axes that are sector-agnostic and directly relevant to LMIC contexts:

- **Decision Stakes:** the severity, reversibility, and breadth of harm when the AI is wrong. Assessed across five dimensions: physical safety, economic exclusion, legal jeopardy, dignity and discrimination, and epistemic harm.
- **Recourse Availability:** the realistic ability of the affected person to identify that AI caused them harm, challenge the decision, and obtain a remedy. Assessed across five resource types: legal/regulatory, institutional, informational, market (can the user go elsewhere?), and social/community.

These axes produce four deployment zones. Zone 2 (high stakes, no recourse) is the primary focus because it is the most dangerous zone and the one most concentrated in LMIC last-mile deployment contexts. Health AI in rural settings, informal credit scoring, and GBV legal advisory tools are all Zone 2 by default in most LMIC markets.

Zone	High Recourse Available	Low / No Recourse
HIGH STAKES Severe, often irreversible individual harm	**ZONE 1 — Regulated Risk** High decision impact; meaningful challenge pathways exist. Example: Licensed MFI credit scoring with consumer protection framework and appeals process.	**ZONE 2 — Critical Gap** HIGH PRIORITY. Most dangerous. Harm is severe and the affected person has no realistic way to contest it. Examples: AI health triage in rural clinic; mobile credit scoring for informal traders; GBV legal advisory in customary law contexts.
LOWER STAKES Diffuse or reversible harm; alternatives exist	**ZONE 3 — Standard Risk** AI error is correctable; user has options and channels. Example: Edtech content recommendation with teacher override and user feedback mechanism.	**ZONE 4 — Diffuse Risk** Individual harm is small, but aggregate impact across user base can be significant and invisible. Example: Agricultural price information tool with embedded inaccuracies affecting 200,000 smallholder farmers.

The supplier will also develop a concise and practical responsible AI framework and self-assessment toolkit that can be used by innovators at the point of design and deployment, and ensure that the final output complements and extends GSMA's existing Responsible AI Maturity Roadmap downstream to startups, social enterprises, and ecosystem partners working at the last mile. The work should be grounded in qualitative primary and secondary research, including literature review and semi-structured interviews with relevant practitioners, and should translate the evidence into an actionable, practitioner-facing resource that is proportionate, usable, and tailored to resource-constrained LMIC settings.

The primary audience for the output includes early-stage AI startups and social enterprises building and deploying AI in LMIC contexts. Secondary audiences include innovation funds, incubators, accelerators, impact investors, mobile network operators, and technology partners supporting last-mile digital services. The framework should also be useful to donor organisations, NGOs, and policy actors seeking practical guidance on proportionate responsible AI practice.

Objectives

The objectives of this research, led by GSMA's EmergingTech programme in collaboration with the GSMA AI for Impact initiative, are

- **Audit existing frameworks:** Map the landscape of existing responsible AI frameworks against an LMIC-fit rubric, identifying structural gaps, rather than duplicating the work already done.
- **Introduce the Deployment Context Typology:** Develop and validate a sector-agnostic classification system that organises AI deployments by decision stakes and recourse availability, making responsible AI guidance portable and proportionate.

- **Build a harm inventory:** Synthesise documented cased of AI related harm in LMIC deployment contexts across three to four key impact sectors supplemented by practitioner accounts of near-misses and observed failures.
- **Map practical constraints:** Understand what actually prevents early-stage LMIC innovators from embedding responsible AI practice, grounded in evidence from Innovation Fund grantees, alumni, and ecosystem partners. This includes understanding how business model (for-profit, not-for-profit, NGO) and funding structure influence responsible AI decisions.
- **Produce an actionable guide:** Develop a concise and practical responsible AI framework and self-assessment toolkit that innovators can use at the point of design and deployment, and self-assesses they grow.
- **Extend GSMA's responsible AI work downstream:** Establish this framework as a complement to the GSMA RAI Maturity Roadmap, extending GSMA's existing responsible AI work downstream to the innovators and ecosystem partners that MNOs partner with to reach last-mile communities.

Research questions

It will seek to answer the following research questions:

Phase 1: Understanding the problem	• RQ1: What responsible AI failures and near-misses have occurred in LMIC AI deployments in key impact sectors, and what caused them? • RQ2: What practical constraints prevent early-stage LMIC innovators from embedding responsible AI practice, and how do these constraints differ by context and sector? o RQ2a: how does the business model play a role? o RQ2b: how does the funding structure play a role? • RQ3: How do cultural, linguistic and community-level factors influence what responsible AI looks like in practice across different LMIC contexts, and how should these factors inform the design and governance of AI systems built for underserved populations? • RQ4: Where do existing responsible AI frameworks fall short for resource-constrained deployers operating at the LMIC last mile, and what gap does no existing framework address?
Phase 2: Building and testing the framework	• RQ5: Does the Deployment context Typology, classifying AI deployments by decision stakes and recourse availability resonate with LMIC innovators as a useful self-placement tool, and does it lead to actual responsible AI actions across zones? • RQ6: What minimum viable responsible AI framework and practical self-assessment toolkit can support innovators operating across different deployment zones in resource-constrained LMIC contexts to deliver responsible AI solutions? • RQ7: What role can mobile network operators play in creating accountability and oversight for the AI-powered services deployed through their infrastructure and ecosystems, particularly in markets where regulatory frameworks are absent or not yet enforced?

Geography

The research will not be limited to a fixed set of low- and middle-income countries. Instead, it should follow the strongest available evidence, drawing on documented harms, practitioner interviews, and relevant case studies wherever they are found.

Given the focus of the Emerging Tech programme, the research should prioritise LMIC contexts across Africa, South Asia, and Southeast Asia, while remaining open to evidence from other regions where it is directly relevant to the research objectives.

Limitations and Out of Scope

- A risk quantification study estimating the aggregate scale of AI harm in LMICs.
- On-the-ground validation of surrounding infrastructure, affordability, or social norms in deployment contexts. The framework will include a prompt asking innovators to consider whether end users can realistically act on AI outputs, but auditing the infrastructure itself falls outside the remit of a responsible AI framework
- AI systems built and owned by governments or large-scale enterprises in LMICs. This project focuses specifically on the startup, social enterprise, and social innovation layer
- Production of a comprehensive or exhaustive catalogue of all AI-related harms in LMICs. The harm inventory will be an honest synthesis of available evidence and practitioner accounts, with its limitations clearly named

3. Anticipated approach and delivery timeline

The below is a proposed approach but we welcome further ideas and proposed methodology from the supplier.

The research will use a mixed-methods approach structured in three phases. Phase 1 runs three parallel tracks to build the evidence base. Phase 2 translates that evidence into a draft framework and practitioner toolkit and Phase 3 provides a pilot-tested and validated responsible AI framework and practitioner toolkit.

Phase	Activities and deliverables	Timeline (indicative)
Phase 1: Inception **Framework landscape audit**	**Key activities** - Kick off call to discuss methodology, scope, timelines and methodological approach - **Framework landscape audit** A structured review of existing responsible AI frameworks against a consistent LMIC-fit rubric. - Stakeholder mapping and outreach plan for key-informant interviews	August-September 2026
	Deliverables - Inception report/deck with literature review highlights, a gap map to ensure that the project builds only in spaces that are genuinely unoccupied. - KII guide, a shortlist of stakeholders and advisory group member suggestions	

<table>
<tr><td rowspan="2">Phase 1:
Research

Harm taxonomy

&

Practitioner constraint mapping</td><td>Key activities
- Building Harm Taxonomy
Developing a bottom-up inventory of documented AI-related harms in LMIC deployment contexts, organised by sector, harm type (physical safety, economic exclusion, legal risk, dignity and discrimination, epistemic harm), and cause (bias in training data, absent recourse, confident hallucination in low-resource contexts, value misalignment).
○ Desk-based evidence gathering
○ Practitioner-reported evidence
- Practitioner Constraint Mapping
25 – xs30 semi-structured key informant interviews (KIIs) with LMIC innovators, including GSMA Innovation Fund grantees, alumni integrating AI into their products and services, and AI4I ecosystem partners.</td><td rowspan="2">September-November 2026</td></tr>
<tr><td>Deliverables
- Interview tracker and detailed summary of interview notes.
- Interim report containing practitioner constraint insights and a documented set of harm cases based on direct field experience from KIIs.</td></tr>
<tr><td rowspan="2">Phase 2:

Building the framework and practitioner self-assessment toolkit</td><td>Key activities
- Synthesising the insights from phase 1 into a draft responsible AI framework and practitioner self-assessment toolkit and report drafting
- Framework architecture and standards
Translating the deployment context typology from a conceptual model into an operating tool by
○ Developing a zone placement guide
○ Defining minimum viable responsible AI standards for each zone
○ Developing sector-specific guidance for priority impact sectors
○ Identifying specific points in the AI development cycle where responsible AI practice matters most for resource-constrained deployers
- Practitioner self-assessment toolkit
Developing the practitioner-facing product that innovators can use. The toolkit should include
○ a zone placement tool and a pre-deployment checklist tailored to each zone
○ a data governance minimum standards guide, a community accountability protocol template, and a harm identification guide for the highest-risk decision points.</td><td rowspan="2">November – December 2026</td></tr>
<tr><td>Deliverables
- Draft framework document
- Draft practitioner toolkit (modular, sector-aware, standalone)
- Full draft report</td></tr>
</table>

<table>
<tr><td rowspan="2">Phase 3

Validation and final report</td><td>Key activities

- Co-designing and stress testing
○ Two structured co-design session with the advisory group and selected AI innovators to test, challenge and refine the outputs from framework architecture and standards, and practitioner self-assessment toolkit.
▪ First session will focus on reviewing the framework architecture and zone typology
▪ Second session will focus on reviewing the draft toolkit.
- Validation workshop post co-design sessions
- Incorporating feedback from validation workshop and reviews</td><td rowspan="2">January – February 2027</td></tr>
<tr><td>Deliverables

- Final report (Word document and PowerPoint) with insights, validated framework, practitioner toolkit and recommendations.</td></tr>
</table>

Firm and proposal requirements

Looking for: *Consultancy team or research organisation*

Required experience

Essential:

- *Demonstrable track record of completing similar assignments in responsible AI, digital development, or technology for development.*
- *Thematic expertise in responsible AI, ethics, and AI governance.*
- *Expertise in technology for development, mobile for development, and AI deployment in LMICs.*
- *Direct experience working with or researching early-stage startups, social enterprises, and impact-focused innovators in LMICs.*
- *Familiarity with the LMIC AI ecosystem, including awareness of foundation model dependencies, data governance challenges, and regulatory environments across Africa, South Asia, and Southeast Asia.*
- *Full working proficiency in English, writing to publication quality.*
- *A strong network of potential interviewees and advisors across the LMIC AI innovation ecosystem.*

Desirable:

- *Experience working with or for mobile network operators.*
- *Familiarity with the GSMA Responsible AI Maturity Roadmap and other AI governance frameworks.*
- *Experience conducting research in health, financial inclusion, and/or education contexts in LMICs.*
- *Presence or partnerships in the regions of primary interest: LMICs across Africa, South Asia, and Southeast Asia.*

Proposals should include a separate technical and financial proposal:

Technical proposal

- A short (1 page) statement of suitability, highlighting recent relevant experience.
- A short discussion of the proposed approach (two to four pages), including the analytical frameworks to be used, identified data sources, proposed KII sampling strategy, and initial thoughts on how the co-design sessions should be structured.
- Any proposed changes to the Terms of Reference with justification.
- Details of relevant firm or team project experience, with references where available.
- A Gantt chart outlining major project stages and timelines.
- CVs and location of all team members.

Financial proposal

1. Level of effort (person-days) by activity and team member.
2. Fee rates (per day in GBP).
3. Total project cost (GBP), without VAT[1].
4. The respondent's total price is inclusive of all costs, insurances, fees, costs, expenses, liabilities, obligations, risks, and all financial requirements for the performance of Services and provision of Deliverables.

[1] GSMA Foundation Inc is a US registered 501(c)(3) charity. Please state include separately where VAT is applicable.

5. Any charge not stated in this Proposal, which extends above to the Total Price, is not permitted.

Due to GSMA compliance requirements, exact project budgets cannot be provided at this stage. You are, however, able to provide a few implementation/budget options that can help us assess value for money and we can align our project scope to the relevant budget after a consultant has been selected.

Proposal and all deliverables are required to include a disclosure statement regarding the use of AI in the production of the work. This disclosure statement should identify which tools were used and how, and indicate which parts of the work are based on these results. The GSMA's standard terms and conditions include clauses requiring that suppliers warrant that the deliverables are the original work of this supplier; breach of this, including through the use of AI, is considered grounds for termination.

Proposal assessment and selection Process

The proposal will be scored on the following set of criteria:

Criteria	Importance	Weighting
Cost	Proposal's value for money	20%
Quality	Quality of the research approach outlined in the proposal, including degree to which it addresses the outlined research questions and proposal elements	40%
Bidder's capacity to manage the project on time and on budget	Selection of experienced high-quality research partner(s) and ability to manage the project on time and on budget	25%
Relevant experience	Bidder's experience in successfully conducting similar projects	15%

- Proposals are to be submitted no later than 4pm BST on Friday 31st July 2026 to Ibrahim Sajid (isajid@gsma.com) and Sheila Moa (smoa@gsma.com). Clarification questions should be submitted by Friday 24th July 2026; with responses to be provided by Tuesday 28th July 2026.

- The estimated proposal evaluation period is Monday 3rd – Wednesday 5th July. Shortlisted consultants may be contacted for an interview between Thursday 6th and Wednesday 12th August 2026, with final selection and invitation to contract no later than Friday 14 August 2026

GSMA Mobile for Development
Foundation Inc.
Suite 910
3445 Peachtree Rd
Atlanta
GA 30326.