EkStep people + ai

# USE CASE ADOPTION FRAMEWORK

How DPI Thinking Drives
AI Adoption

**Authors:**

SHANKAR MARUWADA, SHALINI KAPOOR, TANVI LALL

# INTRODUCTION

*People+ai, an initiative of the EkStep Foundation, together with its knowledge partners, has developed the Use Case Adoption Framework to document and enable AI use cases from the Global South that aim for widespread adoption and meaningful impact.*

*This draft version is being shared for feedback and refinement. It is currently under review at Carnegie India as part of an ongoing effort to validate and improve the framework through expert input and sector-specific insights.*

# CONTENTS

# EXECUTIVE SUMMARY

Artificial intelligence has the potential to transform how citizens and societies learn, work, and grow, but its impact depends on real benefits realised by the people on the ground. Realising those real benefits is the biggest challenge and there are several views on how the same can be achieved. The challenges of taking AI from proof of concepts to production, availability of in-shore compute, local datasets, bias in current models, and unavailability of population scale testing, are for real and have emerged as AI adoption has grown over the past few years.

Digital Public Infrastructures (DPI) have seen this journey. Open digital ecosystems built on shared data, protocols, and governance have already transformed access to finance, identity, and public services. The biggest challenge here was to create a unification of perspectives from policy, technology, funding, governments, into a common thread. The exponential transformation happened not just through technology but a balanced amalgamation of ideas to create a unified way for people.

The next step is to apply similar thinking to AI with use cases as the anchor, to drive us towards that north star of a unified approach. We are still discovering AI as a science, and there is a long way to go, marked with enough confusions, fears and opinions. But usage by the people and for the people can be the way forward for us to discover the challenges and opportunities of AI.

DPI offers a model, a way of thinking for us to build a unified framework for AI Adoption- Use Case Adoption Framework. Newer use cases will unearth newer possibilities.

This paper discusses how connecting real-world AI use cases across sectors such as health, education, agriculture, and livelihoods can help policymakers, innovators, and institutions align around a shared goal: ensuring that AI works for everyone. The framework will evolve as AI adoption and usage by people evolves

# The AI Adoption Challenge:
## Why pilots don't scale

Everyone is seeking "scale," yet many AI initiatives stall after initial deployment: they work in a few sites or sandboxes but don't generalize across states or nationwide. Figure 1 shows that only 10-30% of proof of concepts reach production, and even fewer achieve adoption at scale. Moving from zero to population-scale is an organizational and infrastructural shift - not a bigger pilot. The gap isn't just technical; it's institutional, ethical, and operational - and when that gap persists, citizens and beneficiaries pay the price: services remain patchy, workflows remain broken, and trust erodes as pilots promise much but deliver little where it matters - on the ground.

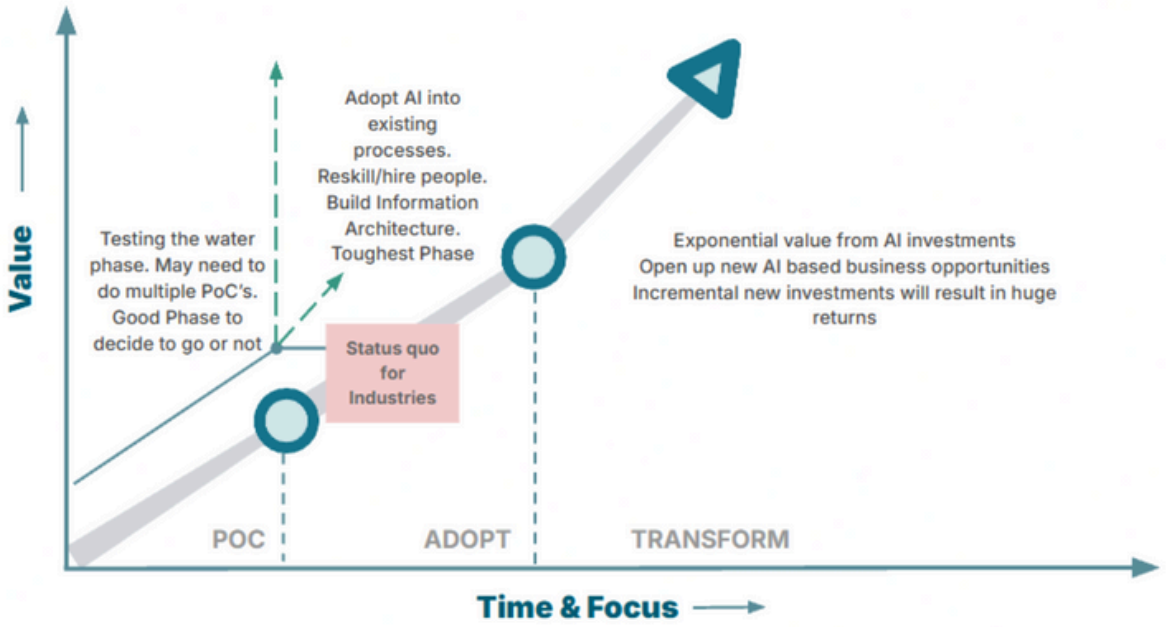


**Figure 1: Only 10–30% of AI pilots are adopted at scale (refer to Appendix 1)**

**[Source: AI for You - Shalini Kapoor & Sameep Mehta]**

## What's blocking AI adoption (beyond first deployments)

- **Islands of deployment**: Bespoke solutions work locally but don't port across volumes, geographies, or programs.

- **No shared vocabulary:** Core terms - use case, infrastructure, impact, benchmark - mean different things to different stakeholders, derailing alignment, governance, and procurement.
- **Pilot-grade plumbing:** Early rollouts depend on ad-hoc workflows that buckle under broader operational load.
- **Safety bolted on late:** Benchmarks, evaluations, and grievance pathways arrive after expansion - too late to prevent harm or build trust at scale.
- **Institutional gaps:** Ownership and accountability are unclear - creating hesitation among builders and beneficiaries.

# The AI Adoption Journey

AI use cases follow a journey where pilots are designed to achieve product-market fit and then scaled for widespread adoption to millions and billions of people on the planet. A typical progression includes two stages:

## Zero to One: Achieving product-market fit

This stage focuses on identifying a real problem, developing an AI solution that addresses it, and proving that it delivers value. It requires close collaboration between problem owners, domain experts, and technology teams. Success is defined by evidence: measurable improvement in outcomes, clear understanding of risks, and documented safety checks.

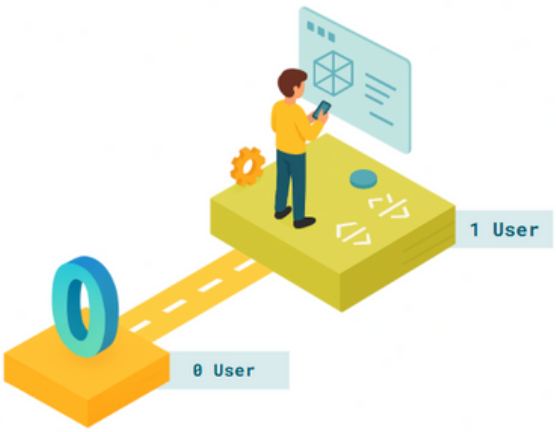


Figure 2: From 0 to 1

### From Validation to Population Scale

Once a solution demonstrates clear value in real-world conditions where it measurably improves outcomes for its intended users and can operate safely within existing systems, the challenge shifts to extending that impact across geographies, contexts, and user groups without losing reliability or trust.

This scaling stage requires different infrastructure than the proof stage. It depends on robust data systems that can handle increased load, affordable compute that doesn't create cost barriers, multilingual capabilities that ensure accessibility, institutional support from government and civil society, and continuous monitoring to catch problems early. Policies and procurement systems must also evolve to ensure that rapid growth doesn't compromise safety, fairness, or accountability.

Safe impact at scale means growth never comes at the cost of harm or exclusion. Whether reaching thousands or millions of users, three things must hold: measurable public benefit, responsible governance, and trust by users and institutions.

Progress from 0 to population-scale adoption depends on consistent focus, strong institutions, and safety built into every stage. The challenges unfold here.

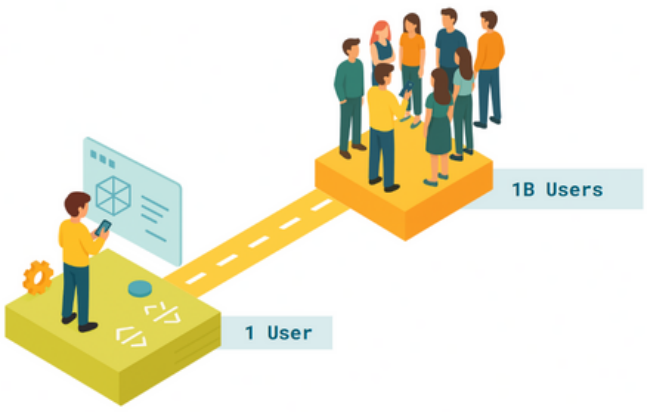


Figure 3: From Validation to Population Scale

## The Silo Syndrome of AI

AI efforts splinter across teams, vendors, and programs - each solving a local problem with its own data, infrastructure, and metrics. The result is islands of progress that don’t add up to scale.

- Scale issues: Pilots tune to one dataset, one workflow, one geography. When volume or context shifts, performance drops and costs spike. Nothing composes.
- Ethical issues: Without shared guardrails, each silo invents its own consent flows, redress paths, and evaluation thresholds creating uneven protections and real risk of harm.

Multiple stakeholders, different truths: Policy, product, procurement, and operations use incompatible definitions of use case, impact, benchmark, and even production. Decisions stall; procurement cycles reset, users and beneficiaries don't reap the benefit of AI.

Every team rebuilds the stack, and momentum dies before anything composes at scale.

AI value shows up when people actually use it - and that means moving from pilots to widespread, everyday adoption. Today, most efforts stall after initial deployment: we don't have a technology gap - we have a coordination and governance gap. Digital Public Infrastructures had seen the same challenges.

## DPI Thinking - Why AI needs it now?

Digital Public Infrastructures (DPI) have already made this journey: open ecosystems built on shared data, protocols, and governance have transformed access to finance, identity, and public services - not by technology alone, but by unifying perspectives across policy, technology, funding, and government into a single, people-centered approach. The next step is to bring the same discipline to AI, with use cases as the anchor and a clear north star of a unified, public-purpose architecture. AI is still being discovered; it carries confusion, fear, and competing opinions. Yet real-world use - by the people and for the people - remains the best way to surface challenges, refine safeguards, and realize benefits. DPI offers the model: a way of thinking to build a unified framework for AI adoption - the Use Case Adoption Framework - through which each new use case reveals new possibilities and strengthens the common rails for those that follow.

# What is a Use Case

A use case refers to a real-world application of AI that delivers measurable societal value. It is not a pilot, prototype, or research experiment, it is a repeatable model that connects a defined user need to a practical solution that demonstrates how AI contributes to better outcomes for people, systems, or communities, and how that benefit can be sustained and scaled responsibly.

## Components of a Use Case

Each use case can be understood through these simple lenses.

- **Persona**: Who is the primary user or beneficiary? This could be a farmer seeking crop advice, a teacher using adaptive assessments, or a health worker diagnosing patients.
- **Purpose**: What problem is being solved or value created? This defines the desired change, whether it's higher yields, improved learning outcomes, faster service delivery, safer streets. Technology needs to be designed to address a clear purpose and needs.
- **Application**: How is AI being used? This includes the specific solution or approach, how it fits within existing systems, and what human processes surround it. The goal is to show where AI adds value rather than replaces judgement or expertise.
- **Infrastructure**: What systems make it possible? Data, compute, language resources, connectivity, benchmarks, safety standards and institutional support all form the backbone of scalable AI use. This helps clarify what needs to be strengthened for replication.
- **Scale**: Can the solution reach more people or contexts without losing reliability or trust? Scale also includes diversity of users, accessibility, affordability, and long-term maintainability.

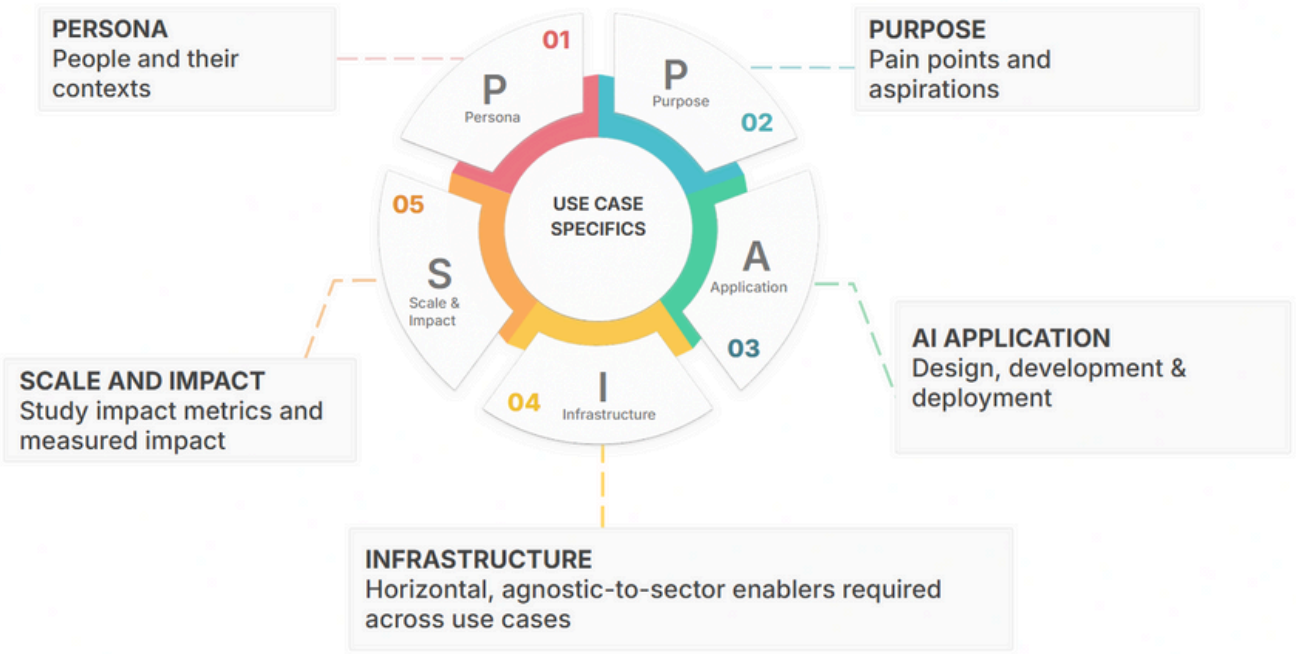


Figure 4: AI Use Cases sit at the intersection of Applications and Infrastructure

## Use Cases can become Anchors

Use cases aren't one-off projects; they are anchors. Once defined, a use case provides a stable reference point that aligns policy, technology, infrastructure and operations around a clear outcome. It clarifies who is served, what success looks like, and which rails - for example data, safety, voice technologies, and language - must be in place for reliable delivery.

Anchoring also creates discipline and portability. It establishes clear ownership and thresholds for readiness, and ties shared resources to demonstrable results rather than perpetual pilots. If the approach is standardized, progress becomes reusable instead of bespoke so learning compounds and impact scales with confidence.

## Use Cases combine vertical sectors and horizontal unlocks

- **Vertical Sectors:** AI creates real-world impact by addressing concrete problems and user needs in diverse sectors, thereby creating value for communities. The value of an AI use case lives in the sector it serves. Some of the sectors where population-scale impact can be imagined are included below. The list is illustrative, not exhaustive; we welcome feedback and suggestions for additions:

    - Agriculture: individuals and institutions such as farmers, Farmers Producers Organisations, agribusinesses and research institutes, who are the users and beneficiaries of the sector engaged in the cultivation, processing, and distribution of food and raw materials.
    - Climate: policymakers, researchers, and communities engaged in understanding, managing, and responding to climate variability and change, leveraging data-driven tools to enhance resilience and inform decision-making. (1)
    - Education: students, teachers and educational institutions engaged in acquiring, imparting, and facilitating knowledge and skills, with a focus on leveraging technology to enhance learning experiences addressing diverse needs and creating value for all participants. (2)
    - Government Services: citizens, public officials, and government agencies involved in delivering and accessing public services that ensure efficient governance, transparency, and value for communities. (3)
    - Healthcare: patients, healthcare providers, and research institutions involved in delivering and accessing medical services that aim to improve population health, patient experience, provider well-being and cost-effectiveness through integrated care, diagnostics, and therapies.(4)

(1) Artificial Intelligence in healthcare - Public Health - European Commission
(2) Artificial Intelligence and the Future of Teaching and Learning (PDF)
(3) AI in Government: Top Use Cases in the Public Sector | Salesforce.
(4) Artificial Intelligence in healthcare - Public Health - European Commission

- Justice: individuals seeking legal remedies, Judges, defenders, prosecutors, courts, legal aid institutions operating within legal systems dedicated to ensuring fairness, access to representation, and protection of[5] rights for all.[6]
- Law: Legislators, legal professionals, law firms, and regulatory bodies developing, applying and interpreting legal frameworks and policies that govern the behavior, rights, and obligations of individuals and institutions.
- Livelihoods: workers, employers, and training institutions who depend on employment, income generation, and skills development opportunities.
- Transportation: commuters, logistics providers, and transport authorities who are the users and beneficiaries of mobility and transport systems.

- **Horizontal Unlocks:** Cross-cutting, agnostic-to-sector infrastructure components digital or otherwise that enable safe adoption of AI at-scale. They encompass:

  - Technology Resources: Tools, infrastructural assets, and technical capabilities that enable AI systems to be built, deployed and scaled efficiently and effectively
  - Practical Frameworks and Policies: Best practices, structured approaches, policies and regulation that help in development and implementation of AI in real-world scenarios

(5) Redefining agriculture through artificial intelligence: Predicting the unpredictable
(6) Communicating with Justice System Stakeholders about the Right to Counsel

A few emerging horizontal unlocks are listed below:

1. AI Safety & Alignment: practical framework that ensure AI technologies are designed and used to benefit humanity while minimizing potential harm or negative outcomes, aligned with the principles of robustness, interpretability, controllability, and ethicality guiding AI models by human values and goals to be as helpful, safe, and reliable as possible. (7)
2. Accessible and Affordable Compute: technology resources that enable machines to perform complex calculations and process large amounts of data including hardware like GPUs, data centers, and software frameworks that work together to support advanced tasks such as language processing and image recognition with greater compute power. (8)
3. Democratise Data: technology resources that make organizational data widely usable for analysis and AI by ensuring it is discoverable, accessible, and reusable, well-governed and trustworthy, and available through self-serve tooling for domain teams - so product owners, model builders, and application developers can reliably discover, understand, and apply data to build, scale, and deploy inclusive AI systems.
4. Human Talent: practical framework to build adaptive, skilled, and inclusive human capacity for an AI-driven future, focusing on AI literacy, continuous reskilling and upskilling, and equitable access to future-ready education while ensuring smooth workforce transitions, minimizing displacement, and turning technological progress into a driver of productivity, innovation, and shared prosperity.
5. Inclusive: governance mechanism for the design, development, and deployment of AI systems that are accessible, equitable, and beneficial for all individuals, regardless of their gender, economic background, language, or geographic location, guided by diverse data and human-centered principles to promote fairness, sustainability, and positive social and economic outcomes. (9)

---

(7) What Is AI Alignment? | IBM , https://www.anthropic.com/news/core-views-on-ai-safety

(8) Measuring compute capacity: a critical step to capturing AI's full economic potential - OECD.AI India's Compute Conundrum | Carnegie Endowment for International Peace

(9) https://iccwbo.org/wp-content/uploads/sites/3/2025/07/2025-ICC-Achieving-inclusive-AI.pdf, OECD AI Principles overview, AI/ML for Bharat: Designing Inclusive AI for India's Next Billion Users | nasscom | The Official Community of Indian IT Industry, Designing more inclusive AI starts with data architecture | World Economic Forum

6. Interpretability: practical framework to understand and explain the internal workings of AI systems, particularly how large language models represent concepts and make decisions to ensure transparency, controllability, and alignment with human values.(10)

7. Multilingual AI & Voice: technology resources and practical frameworks that enables devices and voice assistants to understand, process, and respond to multiple languages and dialects seamlessly, using a combination of advanced technologies such as speech recognition (ASR), natural language processing (NLP), and machine translation (MT) and speech synthesis, built on deep learning models enabling natural, expressive, and human-like communication across linguistic contexts, while integrating with workflows to deliver scalable, context-aware, and accessible voice experiences tailored to diverse users and use cases. (11)

## Introducing the Use Case Adoption Framework

The Use Case Adoption Framework is an actionable framework that connects vertical sectors (where value is created) with horizontal enablers (that deliver scalability and sustainability), all anchored on real-world use cases. Designed for global applicability, it establishes shared definitions to harmonize vocabulary across stakeholders and geographies, and it maps the path from ideas to large-scale impact.

Its purpose is alignment and execution: to align multiple actors around a common language, surface shared needs and gaps, and enable concerted action on the bottlenecks that block scale.

(10) The engineering challenges of scaling interpretability \ Anthropic
(11) Teaching the Google Assistant to be Multilingual

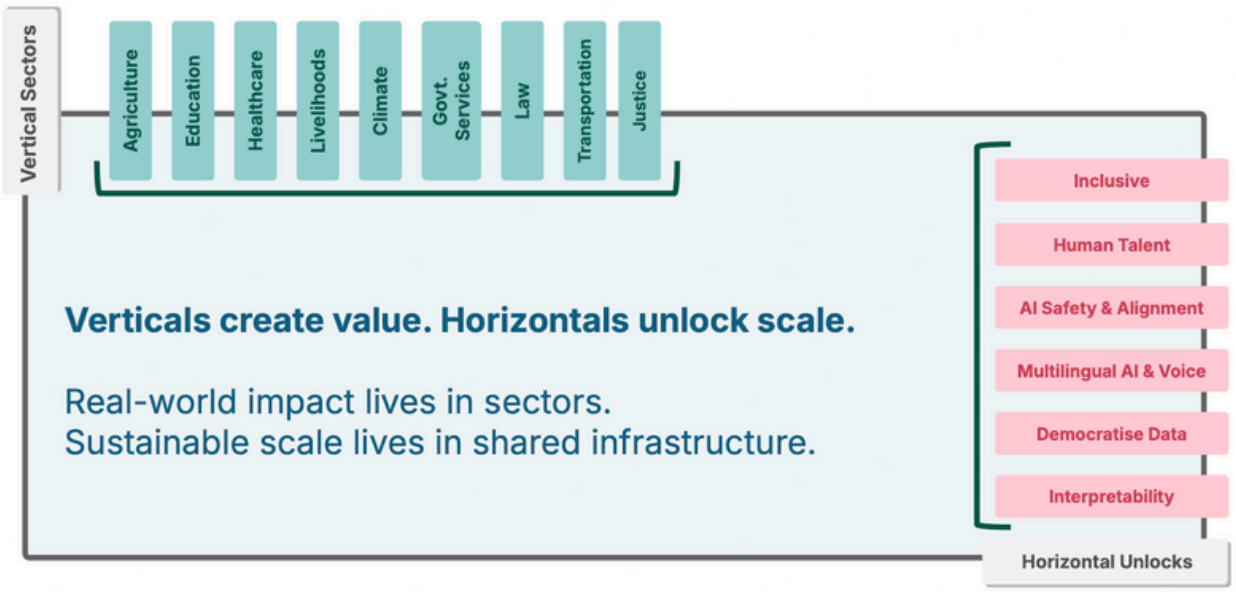


Figure 5: Connecting Sectors and Unlocks in AI Adoption

The framework is designed for use by multiple stakeholders:

- Public institutions and governance bodies can use it to track where new AI use cases are emerging, identify common bottlenecks, and plan coordinated responses.
- Technology developers and infrastructure providers can use it to understand dependencies between data, models, and applications, and identify where shared components or standards are needed.
- Philanthropies, multilaterals, and enablers can use it to direct funding and technical assistance toward system-level gaps like evaluation protocols, compute access, or talent pipelines.
- Researchers and civil society organisations can use it to generate evidence on what works and why, grounding AI policy debates in practical insight rather than abstraction.

# Design Principles of Use Case Adoption Framework

The framework starts with real use cases. From these, it identifies what enables adoption and what holds it back. This approach highlights patterns that recur across sectors and geographies, showing where progress depends on shared infrastructure, data, or governance. The intent behind the framework is not to score or rank AI use cases, but to identify and prioritize bottlenecks that hinder adoption and to guide efforts toward their alleviation.

**Design Principles of Adoption:**

1. Simplicity and transparency: Our first design principle is grounded in building a bottom-up, practitioner-oriented framework rather than a theoretical one meant only for academic study. The goal is to create something simple, transparent, and usable by diverse personas and cohorts working directly with AI, from developers and data scientists to policymakers and implementers. We believe that less is more: the simpler and more intuitive the framework, the greater its utility, adoption, and adaptability across contexts. A highly technical or academic model would exclude those who are actually developing or deploying AI. Hence, our focus is on simplicity and accessibility, ensuring that even small or core functional teams can use, adapt, and contribute to this framework meaningfully.
2. Continuous Feedback Loops: this approach is not limited to pilot or early adoption stages but it applies equally from the 0 to 1 stage of experimentation to the 1 to 1 billion stage of large scale development. Its core purpose is to establish a systematic mechanism for identifying and addressing bottlenecks as use cases evolve. Rather than limiting the analysis of challenges to the initial ideation stages, the framework embeds ongoing evaluation at every step from product market fit and pilot implementation to replication and scaling.
3. DPI Thinking: This framework applies the same design principles that form the foundation of India's Digital Public Infrastructure: inclusivity, interoperability, safety, and openness. These principles ensure that systems can work together, reach diverse users, and remain accountable as they scale. When applied to AI, they encourage models, data, and applications to be built on shared foundations rather than isolated systems.

This approach aligns with the FAIR principles already discussed, creating infrastructure that is findable, accessible, interoperable, and reusable. Together, these ideas make AI adoption safer, more transparent, and easier to extend across sectors and geographies.

4. Iterative and Evolving Model: This framework is intentionally conceived as a working and evolving model which is a first attempt that will continue to improve through iteration and collective learning. As a design principle, it acknowledges that developing an adoption framework for AI is an ongoing process rather than a final product. The framework is expected to mature and refine over time through engagement with diverse stakeholders across the ecosystem, including governments, multilateral organizations, practitioners, and researchers. By positioning it as an iterative and adaptive construct, the intent is to ensure that the framework remains responsive to emerging insights, evolving technologies, and contextual realities of AI adoption at scale.

## Measurable Impact of Use Cases

This framework is diagnostic and improvement-oriented, not a scoring tool. Its purpose is to make impact observable and comparable across contexts, so teams can learn, adjust, and scale responsibly.

To keep measurement practical, we organize indicators into core dimensions. Each dimension has a focus and example indicators you can adapt to the sector and context.

| Dimension | Focus | Example indicators |
|---|---|---|
| Scale of adoption | Reach & inclusion | Number of end users reached; share of underserved users; usage frequency and retention. |
| Tangible outcomes | Sectoral results | Measurable improvements in target outcomes (e.g., yields, learning scores, turnaround time, error rates). |
| Reliability & safety | Trust & stability | Advisory/answer accuracy; model robustness; bias reduction; uptime; grievance resolution time. |
| Local capacity built | Ecosystem enablement | Number of practitioners trained; availability of local support; documentation and playbooks adopted. |
| Governance & guardrails | Fairness & accountability | Consent/privacy controls in use; audit trail coverage; conformance to evaluation standards. |
| Geographical reach | Diversity & inclusion | Spread across regions and contexts; performance consistency across diverse user groups. |

Table 1: Core Dimensions of Impact

The goal is not to "pass a test," but to identify gaps, improve design, and build confidence that a use case is delivering public value - safely, reliably, and at scale.

## Real-world examples

To illustrate how vertical sectors and horizontal enablers interact within the Use Case Adoption Framework, we present two representative case studies. A framework gains real value only when it is tested, refined, and validated through actual use cases. These examples therefore serve not just as illustrations, but as practical demonstrations of how the framework can help identify enablers, reveal bottlenecks, and guide adoption strategies in real-world contexts.

## CASE STUDY 1: AI for Agriculture

Consider the agriculture sector which forms the backbone of India's economy, employing nearly 45% of the country's workforce and contributing 18% to GDP.(12) Farmers in India face persistent challenges such as fragmented market access, technological issues, weather volatility, and limited access to timely, localized information - further compounded by barriers of language diversity, literacy gaps, including digital literacy. Across India's rural regions, millions of smallholder farmers work on small plots of land, often in low-connectivity environments, and speak in more than 20 regional languages.(13) Digital literacy remains limited, with many farmers still relying on peers, local agricultural officers or input retailers who travel from village to village to share advice on application.

The OpenAgriNet (OAN) initiative bridges last-mile advisory gaps by using AI to deliver voice-first, multilingual advisories with localized weather, market, and crop guidance in regional languages. OAN's Maharashtra pilot, MahaVISTAAR, is an open-network platform that advances the state's digital inclusion goals for farmers. Powered by Generative AI, it provides tailored, reliable, location-specific support that farmers can access by speaking in Marathi, Hindi, or Telugu—no literacy barrier.

**How the framework works on the ground:**

A farmer in Nanded asks in Marathi, *"Tur dal (pigeon pea) la ropan kadi karu? Fusarium wilt kasa talu?"* The system (1) recognizes the language and crop, (2) pulls district-scale weather nowcasts and 7–10-day forecasts, (3) retrieves crop packages of practices from ICAR institutes/KVKs and state agri universities, and (4) checks mandi trends before advising a sowing window, wilt-resistant varieties, seed treatment, and a safe, label-compliant spray plan—then reads it back in Marathi and offers a short SMS recap.(14)

(12) Economic Survey 2022-23

(13) Digital agriculture in action – Selected case studies from India

(14) Journey of MahaVISTAAR

This use case operationalizes four horizontal enablers:

- **Multilingual & Voice Accessibility**
  - Speech-in/speech-out for Marathi, Hindi, and Telugu; domain vocab for crops like tur dal, soybean, cotton; handles dialectal variants (e.g., "tur," "arhar") and code-mixing.
  - Offline/low-bandwidth fallbacks: compressed audio prompts, IVR callbacks.
- **Democratised Data**
  - Structured ingestion from ICAR/KVK packages of practices, vetted state advisories, IMD weather feeds, and market data (e.g., APMC/Agmarknet) via open APIs.
  - Provenance is retained; every answer cites its data source and timestamp.
- **Interpretability**
  - Farmer-facing: simple "Why this advice?" card showing key drivers (rain probability, soil type, pest risk index) with confidence bands.
  - Expert-facing: rule traces and source snippets from ICAR documents so agronomists can audit or update the logic.
- **AI Safety & Alignment**
  - A moderation module screens outputs for unsafe, non–label-compliant, or off-label chemical use, dosage errors, and misinformation; triggers fallback to a human expert/KVK when confidence is low.
  - Region- and crop-specific guardrails (e.g., PHI/REI checks, restricted-chemical lists) and continuous post-harvest feedback loops reduce harm.

**Impact signals**: early deployments report fewer pesticide errors, better-timed sowing/harvest decisions, improved price discovery, and higher confidence in digital advisories—especially among low-literacy users who rely on voice.

## CASE STUDY 2: AI for Safe and Fair Migration

Across Asia, low-wage migration is marked by structural imbalances. In places like Hong Kong - where thousands of migrants sustain households and local economies - workers (especially women in domestic and caregiving roles) face exploitative intermediaries, excessive recruitment fees, and deceptive job promises. Many arrive indebted and under-informed, with limited access to reliable information, legal support, or fair dispute resolution.

Migrasia is a social enterprise reshaping this journey through transparency and technology. Combining research, advocacy, and digital tools, Migrasia advances fair recruitment and worker protection. Its AI-powered PoBot acts as a trusted digital companion for migrant workers - answering questions on employment rights, assisting with contract checks, and routing users to credible support channels. By reducing misinformation and agency dependence, PoBot helps lower vulnerability to abuse and strengthens worker autonomy.

**How the framework works on the ground:**
A worker preparing to depart asks, "Is this contract legal if my agency keeps my passport?" The system (1) recognizes the query and language, (2) checks contract clauses against verified policy/rights resources, (3) flags red-line risks (e.g., passport retention, fee deductions), and (4) provides next steps - how to document the issue and contact approved support services - along with a short, plain-language recap.

This use case operationalizes four horizontal enablers:

- **Multilingual & Accessible**
  - Chat in workers' preferred languages with plain-language responses and low-bandwidth options. Clear, jargon-free "what this means for you" summaries.
- **Democratized Data**
  - Curated, versioned knowledge base from verified legal/policy sources and partner guidance.
  - Source provenance and timestamps attached to answers.

- **Interpretability**
  - Worker-facing "Why this advice?" card showing which clause or policy triggered the guidance.
- **AI Safety & Alignment**
  - Moderation to avoid unsafe or unauthorized legal guidance; escalation to human support when confidence is low.
  - Jurisdiction-aware checks to prevent misleading or out-of-scope advice.

**Impact signals:** Early use indicates fewer misinformation incidents, reduced reliance on intermediaries, faster navigation of contract and rights queries, and greater confidence in seeking help especially among first-time migrants.

# Appendix

## Appendix 1

AI adoption follows a steady path. It begins with 0→1, where pilots and proof-of-concepts test if an idea works and creates real value. Once validated, projects move toward scale, where systems must stay reliable, safe, and inclusive as they reach more users.

**Figure 1 shows how value grows along this path:**

- The POC phase is where organisations experiment, test assumptions, and decide whether to invest further.
- The Adopt phase requires integrating AI into existing workflows, building information architecture, and reskilling teams, often the hardest but most critical stage.
- The Transform phase represents maturity, where AI begins to generate exponential value, opening new services and efficiencies at scale.

## Appendix 2

**CASE STUDY 3: AI for Education**

India has one of the most diverse education systems with hundreds of millions of students across various levels, institutions, and mediums of instruction. The system serves students from diverse socio-economic, linguistic, and cultural backgrounds across rural and urban areas. In classrooms, students often face barriers of language diversity, rote learning and limited access to personalized feedback, especially in low-resource environments where equitable access to facilities, quality education, and effective teaching pedagogy continue to remain major challenges.

EkStep Foundation in collaboration with AI4Bharat developed the Assisted Language and Math Learning (AxL) project aimed to address these pain points through an AI-enabled, open platform that supports continuous, formative assessment in a "safe space". The initiative helps educators and learners track progress in real time and receive personalized learning insights in multiple Indian languages. It is now being scaled from 540 to 5,000 primary schools across 33 districts, extending its reach to thousands of learners and teachers, offering foundational learning in Telugu, English, and Mathematics for students in Classes 3 to 5.(15) The AxL toolkit is built on four interconnected assistants - a student assistant for home learning, a school practice assistant for lab-based exercises, a teacher assistant for personalised instruction, and a community learning assistant for diagnostics and reading support.

**How the framework works on the ground:**

- **Multilingual voice system**
  - AXL uses speech recognition and natural language processing (NLP) to enable interactive, voice-based learning in multiple Indian languages.
  - Students read aloud or respond verbally, and the AI provides real-time pronunciation and comprehension feedback in the same language.
  - This allows children to learn in their mother tongue while gradually building English proficiency, ensuring accessibility across linguistically diverse classrooms.

(15) Govt. signs MoUs with six NGOs to transform school education - The Hindu

- **Interpretability**
  - The platform includes teacher-facing dashboards that explain why certain exercises are assigned or where a student is struggling.
  - These dashboards display progress metrics such as reading accuracy, time spent, and repeated errors, offering transparent insights into the AI's reasoning.
- **AI safety and alignment**
  - AXL embeds AI safety principles through moderation layers that filter inappropriate or biased content before reaching learners.
  - The system is aligned with local curricula and pedagogical standards, ensuring culturally relevant and age-appropriate learning materials.
  - Student data is anonymized and securely stored, maintaining privacy and ethical use of AI while continuously refining the system based on feedback.

Once scale is achieved through these horizontal unlocks, the impact of different use cases can be measured through multiple dimensions of AI adoption which create a holistic picture.

*With inputs from: Keyzom Ngodup, UNDP AI Hub, Jagadish Babu, Gates Foundation.*