**Frugal AI Hub**

Research Report · July 2026

# MINDS, MODELS & MARKETS

*The Next Chapter of the UK–India AI Partnership: from DeepMind to districts, connecting India's deployment strength with UK frontier AI*

TWO CURRENCIES

TRUST

REACH

UNITED KINGDOM

certifiable, auditable AI

INDIA

offline, in-language AI at population scale

FRUGAL AI HUB at UNIVERSITY OF CAMBRIDGE Judge Business School

Centre for India & Global Business

UNIVERSITY OF CAMBRIDGE Judge Business School

A Frugal AI Hub Research Report · July 2026

**Prachetas Bhatnagar**
Head of Strategy, Global Affairs, Tony Blair Institute for Global Change

**Serish Venkata Gandikota**
Co-founder & Co-Lead, Frugal AI Hub · Visiting Fellow, Cambridge Judge Business School

**Elizabeth Osta**
Co-founder & Co-Lead, Frugal AI Hub · Visiting Fellow, Cambridge Judge Business School

**Arjuna Sathiaseelan**
Technology Lead, Frugal AI Hub · Visiting Fellow, Cambridge Judge Business School

**Jaideep Prabhu**
Professor of Marketing and Nehru Professor of Indian Business, Cambridge Judge Business School · Co-Director, Centre for India and Global Business

## ABOUT THE FRUGAL AI HUB

The Frugal AI Hub at Cambridge Judge Business School is a research and practice hub advancing the principles, methods and adoption of Frugal AI: intelligent systems that are high performing yet resource efficient, and that can be understood and governed by the institutions that deploy them. The Hub works with industry, public sector and civil society partners across sectors including fintech, healthcare, education and sustainability. Further information is available at https://frugalai.org.

## ABOUT THE CENTRE FOR INDIA AND GLOBAL BUSINESS

The Centre for India and Global Business (CIGB) at Cambridge Judge Business School, University of Cambridge, is a platform for research and engagement with partners across industry, academia and policy in India, the United Kingdom and the wider world. Its focus is to understand and promote innovators and innovation in India and globally, through research, publication and the events and programmes it convenes at Cambridge. The Frugal AI Hub sits within the Centre, which is co-directed by Jaideep Prabhu. Further information is available at https://www.jbs.cam.ac.uk/centres/india-global-business.

## DISCLAIMER

The views expressed in this report are those of the authors and do not necessarily represent the institutional positions of their affiliated organisations. Prachetas Bhatnagar contributed in a personal capacity.

## SUGGESTED CITATION

Bhatnagar, P., Gandikota, S.V., Osta, E., Sathiaseelan, A. and Prabhu, J., 2026. Minds, Models and Markets: The Next Chapter of the UK–India AI Partnership. Frugal AI Hub Research Report. Frugal AI Hub at Cambridge Judge Business School, July 2026.

# Contents

# Executive Summary

This report argues that the UK–India artificial intelligence (AI) relationship is misread as a race between a frontier research economy and a deployment economy. Examined function by function, each country supplies a different currency of AI value, trust and reach, and the binding constraint on the relationship is translation: the absence of standing mechanisms to move a researcher, a dataset, a certification or a product across the gap between them. India's reach capability is examined through Frugal AI, the design and deployment discipline developed through the Frugal AI Hub's research programme, in which engineering technique is one layer within a wider architecture of design, lifecycle cost, governance and institutional absorption. This is a policy research report based on documentary and public evidence from both countries; no primary data collection or interviews were undertaken.

## Key findings

- *Two independent currencies of AI value. The UK is among the world's strongest suppliers of trust: certifiable, auditable AI for consequential decisions. India is among the world's strongest suppliers of reach: AI that works cheaply, offline and in-language at population scale. Neither is a lesser version of the other.*
- *Selective sovereignty is the achievable goal. The serious strategic question for either country is which one or two of seven stack layers (energy, compute, data, models, assurance, deployment, distribution) to own outright, and where to responsibly depend on partners for the rest.*
- *Smaller models are frequently the sovereign choice. "Frontier-minus-N" models, deliberately built several generations behind the absolute frontier, are often more auditable, governable and cost-effective for the great majority of real public-service and productivity problems.*
- *The corridor already carries real traffic. Qure.ai's NHS deployment and Oxford Nanopore's Indian genomics partnerships show trust and reach combining productively today, without any dedicated bilateral mechanism, evidence the proposed architecture would accelerate rather than invent.*
- *Five sector intersections are investable now, in sequence. Life sciences ranks first on every readiness criterion; climate technology sits immediately behind it pending field validation; institutional fintech and deep mobility form a second tier pending regulatory bridges; semiconductors require the deepest, slowest ecosystem upgrade.*
- *Seven structural obstacles are individually solvable but currently unowned. Data governance ambiguity, compute access, unresolved IP allocation, mismatched procurement rules, RBI risk aversion, absent researcher-mobility routes, and, most fundamentally, no institution mandated to carry a stalled project across departmental silos.*

## Recommended next step

**Establish a first joint demonstrator in life sciences within twelve months, with a named institutional owner on each side, a defined budget and a success metric agreed before the project starts (see Section 8). Success criteria should include a jointly governed data pipeline, a shared evaluation protocol and at least one co-owned IP artefact or product. A single demonstrator judged honestly, including if judged a failure, would do more to establish the corridor's credibility than further joint statements**.

# Foreword

Artificial intelligence is rapidly becoming one of the defining strategic technologies of the 21st century. It is reshaping economies, public services, national security and industrial competitiveness at a pace few technologies have matched before. As AI becomes increasingly central to economic resilience and geopolitical influence, partnerships between trusted nations will matter as much as technological capability itself.

The United Kingdom and India are uniquely placed in this regard. They are vibrant democracies, longstanding strategic partners and increasingly aligned in their commitment to responsible innovation, AI safety and assurance, and the development of trusted digital ecosystems. Recent initiatives, including the UK–India Technology Security Initiative led by the National Security Advisers of both countries, alongside the landmark Comprehensive Economic and Trade Agreement, have established a stronger foundation of trust on which deeper technological collaboration can now be built.

The next chapter of this relationship should not simply be about exchanging technology. It should be about enabling the translation and diffusion of AI capabilities across two highly complementary innovation ecosystems. Each country brings distinctive strengths developed over decades. When combined through purposeful collaboration, those strengths become significantly greater than either could achieve independently. The opportunity lies not in building parallel capabilities, but in creating pathways through which innovation can move, adapt, scale and create value on both sides.

The life sciences sector offers a compelling illustration. The UK’s world-leading scientific research and discovery capabilities, combined with India’s scale in clinical deployment, manufacturing and constrained environment implementation, create opportunities that neither ecosystem could fully realise alone. Similar opportunities are emerging across advanced manufacturing, financial services, climate technologies and digital public infrastructure.

This paper makes an important contribution by examining how those complementary capabilities can be translated into practical economic partnerships. It moves beyond broad aspirations to explore the institutional, commercial and policy foundations needed to accelerate collaboration between researchers, startups, industry and government. In doing so, it presents a timely framework for strengthening one of the world’s most important technology partnerships.

The UK–India AI partnership has the potential to become a model for international collaboration in the age of AI, benefiting both our nations and contributing to a more secure, innovative and prosperous global technology ecosystem.

**Lord Nagaraju of Bloomsbury**

Member of the House of Lords, United Kingdom Parliament
& Founder, AI Policy Labs

# Abstract

Artificial intelligence is often described as a race to build the most powerful models, with the United States far ahead, China competing, and countries like the United Kingdom and India cast as followers. That view measures the wrong economy. Software and digital services, the layer this race is actually about, are a modest share of global output next to the far larger opportunity in manufacturing, energy, agriculture, health, logistics, defence and physical infrastructure. That domain rewards two different things: whether a system can be trusted with a consequential decision, and whether it can actually run, cheaply and reliably, under the real constraints of cost, connectivity and language most of the world lives under. This paper calls these the two currencies of the AI economy, trust and reach, and argues that the UK and India are, respectively, the world's most developed suppliers of each. Seen this way, the UK–India relationship is a partnership between two complementary paths.

A related misconception has already misdirected billions of pounds of public AI investment worldwide: that a country must own the entire AI stack, frontier compute, frontier models and full data residency, to participate seriously. Full-stack sovereignty is prohibitively expensive and, for nearly every country outside the US and China, unnecessary. The real strategic question is which one or two layers, such as assurance, data commons, deployment rails or distribution, a country can own well while safely depending on partners for the rest. Nor does a country without frontier-scale data lack meaningful things to build. Smaller "frontier-minus-N" models, deliberately several generations behind the absolute frontier and trained on what this paper calls a subcritical dataset, are frequently the more auditable, governable and useful choice for public services, health screening or agricultural forecasting. The UK's strength lies in certifying high-stakes systems, structural biology, physics-informed and Bayesian methods, safety and evaluation science and hardware–software co-design, built over centuries as the world's counterpart of trust. India's strength, which this paper terms Frugal AI, is a genuine design discipline, expressed through rigorous engineering, for making intelligence work under cost, connectivity and language constraints that a wealthy market never had to design for.

This paper argues that the binding constraint on the UK–India relationship is a two-way translation gap: UK systems achieve scientific excellence but stall before institutional or population-scale deployment, while Indian systems achieve massive domestic reach but stall before frontier-grade validation and premium international markets. Existing agreements give the relationship legal and diplomatic scaffolding, but they do not, by themselves, move a researcher, a dataset, a certification or a product across the gap. What follows maps both paths on their own terms, identifies where they combine into new economic value, and sets out the pathways, obstacles and sequencing that would make that combination real, before considering what widening this corridor to other trusted partners would, and would not, resolve.

Throughout, Frugal AI carries the fuller meaning developed through the Frugal AI Hub's research programme. What began as a response to resource asymmetry, the recognition that most of the world must reach comparable outcomes without comparable resources, has matured into a design philosophy for the deployment phase of AI. Engineering techniques such as compression, quantisation and edge architecture are its most visible layer. The discipline begins with the deployment environment and the outcome required, treats compute, model, data, orchestration and governance as a single design problem, and extends to the ecosystem of policy, procurement, skills and institutional absorption that determines whether a system endures. The corridor argument that follows rests on this definition.

## Scope and analytical approach

This is a policy research report rather than a peer-reviewed academic paper, and it is worth being explicit about what that means for the evidence beneath it.

The evidence base is documentary and public. It draws on government publications and official statistics from both countries, regulatory and parliamentary material, company disclosures and technical documentation, peer-reviewed research where available, and reporting from the specialist technology and policy press, current to July 2026. No primary data were collected and no interviews were conducted. That evidence base closes on 21 July 2026, the date on which the United Kingdom substantially rearranged its machinery of government for artificial intelligence; Sections 6 and 8 note the consequences where they bear on the argument.[32]

Several categories of figures carry known limits, and the report flags them where they appear. Deployment outcomes reported by companies, including the performance claims attached to India's sovereign models, are those companies' own and are not independently verified. Unit costs for several of the deployments examined in Section 7 are not publicly disclosed. The carbon sequestration multiplier underlying the climate technology intersection is a research proposition under field validation rather than an established result.

The comparative figures in Section 2 and Appendix A are qualitative. They position the United Kingdom and India relative to one another across a set of axes, on the basis of the documentary evidence assembled here. They are not a scored index, they carry no weighting methodology, and no numerical value should be read off them. Their purpose is to make the shape of the complementarity visible in a single view. Where a specific quantity is claimed anywhere in this report, it appears in the text with a source attached.

The analytical framework is drawn from the Frugal AI Hub's published research programme, in particular its work on Frugal AI as a design and deployment discipline[1, 2] and the total cost of ownership and return on investment framework developed with the United Nations International Computing Centre.[3] This report applies that framework to the UK–India relationship. It does not attempt to develop it, and readers wanting the underlying theory should turn to the Frugal AI Hub publications listed at the head of the sources.

# 1

# Introduction: The Two-Way Translation Gap

Two national conversations about artificial intelligence have been running in parallel for several years without quite meeting each other, and the shorthand each country reaches for about itself, and about the other, is wrong in specific and correctable ways. In the United Kingdom, the conversation is usually framed as frontier research: the country that produced AlphaFold, that hosts DeepMind, is racing to keep a compute-poor, capital-constrained economy relevant against laboratories with an order of magnitude more capital. That framing is not so much false as it is aimed at the wrong competitor. The United Kingdom is not a frontier-research economy in the sense the United States is, and will not out-build it on raw model capability. In India, the conversation is usually framed as deployment: a billion-plus population, twenty-two languages, patchy connectivity, and a genius for making digital services reach people other systems cannot. That framing, too, understates what is actually happening, because India is no longer only deploying other people's AI. It is training its own frontier models, on its own infrastructure, and putting real political and financial weight behind the claim that it can compete.

Both shorthand descriptions survive contact with reality about as well as most shorthand does, which is to say partially, and in ways that matter for exactly the reasons this introduction exists to correct. What each country actually has, examined at the level of specific function rather than national reputation, is a distinctive currency, and the two currencies differ from those the popular framing suggests. The first currency is trust: the capacity to make a system's behaviour provably safe, certifiable, and legally accountable enough that a regulator, an insurer, or a hospital will actually rely on it for a consequential decision. The second is reach: the capacity to make a system's capability actually arrive, cheaply, offline-capable, in the right language, for the billions of people a frontier laboratory's default design assumptions quietly exclude. Raw model capability that is neither trusted nor reachable remains a demonstration. The United Kingdom is the world's most developed supplier of the first currency. India is the world's most developed supplier of the second. Neither is a coincidence, and neither is well captured by the words "research" and "deployment."

The UK's trust currency has a longer history than most AI commentary acknowledges, and understanding that history is not a digression, it is the reason the UK's advantage is durable. London built its original global economic role as the world's counterparty of trust: in marine insurance, in reinsurance, in the legal architecture of cross-border contracts, in the slow accumulation of institutions, Lloyd's, the Bank of England, the common-law courts, whose entire function was to make a promise enforceable across a distance neither party could personally verify. The UK AI Security Institute, the Medicines and Healthcare products Regulatory Agency, the National Institute for Health and Care Excellence, and the Financial Conduct Authority's regulatory sandbox model are not new inventions bolted onto an AI strategy; they are the same institutional function, verifying a promise at a distance, applied to a new kind of promise. This is why the UK's advantage in AI safety and evaluation science is not a policy choice that could as easily have been made by another country with enough funding. It is downstream of two centuries of specialising in exactly this kind of institutional trust production.

India's reach currency has an equally specific and equally non-accidental history. It was not designed in the abstract; it was forced into existence by a domestic market that no off-the-shelf Western technology stack could serve, a market of 1.4 billion people, twenty-two scheduled languages and hundreds of dialects, a large population not comfortably literate in any single one of them, patchy and inconsistent broadband, and a willingness to pay for digital services measured in cents rather than dollars. India Stack, Aadhaar, the Unified Payments Interface, and the underlying frugal AI engineering discipline examined in the next section were built because nothing else would have worked at that population scale and that price point. This is, again, why the advantage is durable: it was earned by solving the hardest deployment problem in the world, and no development strategy assembled it.

**This is a translation gap; it runs in both directions, and closing it is an act of parallel engineering.**

The gap this paper is about follows directly from this. UK capability, however scientifically excellent, routinely stalls before it reaches commercial or institutional scale, precisely because trust without reach stops at the laboratory. Indian capability, however proven at scale, routinely stalls before it reaches frontier-grade international validation and premium market access, precisely because reach without trust remains a domestic success story. Existing cooperation, the Comprehensive Economic and Trade Agreement in force from 15 July 2026, the India-UK Joint Centre for AI, the wider Vision 2035 architecture,[12, 21] gives the relationship real legal and diplomatic scaffolding. It does not, by itself, move a researcher, a dataset, a certification, or a half-finished product across the gap. What follows maps both currencies on their own terms, with their real limits intact, before turning to where they combine, how capability should move, what blocks it today, where the value actually comes from, and what a working institutional architecture would look like.

# 2

# Two Independent Paths to AI Sovereignty

It is tempting, and wrong, to describe the UK-India AI relationship as a hierarchy, frontier science at the top, application engineering below it. The two countries have built genuinely different engineering paths, optimised for different objective functions, and each is sophisticated on its own terms. This section examines both in the technical detail that a hierarchy framing would skip past, because the specific mechanisms are what later sections build on. It begins, deliberately, with two corrections to how the wider AI-sovereignty debate is usually conducted, because both countries' own strategies, examined honestly, are better explained by these corrections than by the full-stack race the abstract opened by rejecting.

## Stack Sovereignty: Own One Layer, Secure the Rest

The instinct to build sovereign AI capability at every layer at once, compute, frontier models, and data residency alike, is widespread among governments and, on the evidence, largely misdirected. The Tony Blair Institute's own recent analysis of AI sovereignty states the correction bluntly: leaders should resist the urge to own every model and data centre, and instead prioritise securing access to frontier capabilities while building domestic strength only where it matters most, because full self-sufficiency is too expensive, too slow, and for most countries simply impossible.[8] This is not a fringe view. Analysts distinguishing "full-stack sovereignty" from narrower "compute sovereignty" or "application sovereignty" have made the same point using India's own strategy as the clearest illustration: India has deliberately emphasised application-led sovereignty, embedding AI within its digital public infrastructure to close local language and service gaps, rather than chasing frontier-model ownership for its own sake. Commentary on India's own compute build-out makes the point even more sharply: the layer of the AI stack that attracts the most political attention, the model, is also becoming the cheapest, most abundant and most disposable, while the layers that actually determine whether a country can be coerced, compute it does not yet manufacture, and data it already owns, sit beneath it and reward patient, unglamorous infrastructure investment rather than a single headline-grabbing model launch.[8]

The serious strategic question for any country, including both the UK and India, is therefore which one or two layers it can own well enough to matter, and where it can responsibly depend on partners for the rest, an approach some analysts have termed managed interdependence: pick the layers where genuine comparative advantage exists, build there, partner for the rest, and use that strength to negotiate favourable terms on everything else.[8] This is precisely the shape of the UK-India relationship examined throughout this paper: two countries each choosing to own a different, complementary layer well.

Making this operational requires breaking "the stack" into layers precise enough to assign ownership against. This paper uses seven: energy and sites, compute, data, models, assurance, deployment and procurement, and distribution. Table 1 below maps the UK and India's actual position on each, and states, for each layer, what a defensible ownership-versus-dependency choice looks like.

*Table 1 · The seven-layer AI stack: where the UK and India actually sit, an analytical comparison developed for the UK–India relationship rather than a universal AI architecture. Sources: 9–11, 18, 19.*

| Layer or capability | UK position | India position | Sovereignty logic: own, or depend? |
|---|---|---|---|
| **Energy & sites** | AI Growth Zones with co-located generation; mature grid, but slow interconnection queues | Rapidly scaling renewable capacity and lower site costs; regionally variable grid reliability | Each country should own its own siting and generation; the exportable asset is the AI-driven grid-forecasting method |

| Layer or capability | UK position | India position | Sovereignty logic: own, or depend? |
|---|---|---|---|
| **Compute** | Deep public investment (£28bn+ Growth Zones, Hardware Plan) but reliant on imported accelerators | IndiaAI Mission allocation, a fraction of frontier scale | Neither can realistically achieve near-term chip self-sufficiency; the defensible sovereignty choice is compute access and allocation policy |
| **Data** | High-quality, curated, longitudinal (UK Biobank, NHS) but smaller in volume and largely monolingual | Enormous volume and linguistic diversity (Aadhaar, UPI, Bhashini corpora) but less curated and fragmented across sectors | Each country should own and protect its own data commons; cross-border partnership is the dependency-management mechanism (Section 6) |
| **LLMs and frontier models** | Strong in scientific and frontier-adjacent models, with particular depth in life sciences, defence, robotics, media and real-world reasoning, though constrained in general-purpose frontier-scale compute | Sovereign models (Sarvam, BharatGen) competing on Indian-language tasks | For most civilian and public service workloads, this is the layer that matters least for sovereignty (see frontier-minus-N doctrine below); countries should own the model classes suited to the actual problem and depend on others for general-purpose capability. Defence and cyber security may legitimately prioritise frontier capability on security grounds |
| **Vertical AI models (domain-specialised systems)** | Strong in biotechnology and drug discovery, legal technology and compliance, regulatory technology, fraud detection and graph AI, and insurtech and computer vision (Tractable) | Strong in multilingual and voice-first AI, healthcare diagnostics (Qure.ai in radiology, SigTuple in pathology), agricultural technology and financial inclusion, across applications designed for population-scale or resource-constrained deployment | The two ecosystems specialise in different verticals because they serve different markets and institutional environments. UK firms integrate deeply into regulated European and US enterprise sectors, while Indian firms build cost-efficient systems domestically and scale them into global markets. |
| **Assurance** | World-first state-backed evaluation institute; MHRA/NICE-grade certification culture | Nascent; prioritises inclusion and deployment speed over frontier-risk evaluation infrastructure | The layer the UK should own on the corridor's behalf; India should depend on it for frontier-risk evaluation while building its own deployment-specific assurance capacity |
| **Deployment & procurement** | Sandbox model (FCA) is a genuine asset, but public-sector procurement is notoriously slow | Population-scale deployment machinery via DPI rails; institutional procurement less mature | Each country owns its own deployment rails; the dependency is procedural: mutual sandbox recognition |
| **Distribution** | Strong route into European and other developed-market distribution via regulatory recognition | Strong route into its own vast domestic market and other price-sensitive emerging economies | Each country owns its own distribution channel and grants the other access to it, the two-sided gateway argument of Section 3, rather than either building a parallel channel |

Read as a whole, Table 1 makes a specific claim: the UK's defensible ownership set is energy and sites, assurance, and distribution into developed markets; India's is data commons, deployment and procurement rails, and distribution into price-sensitive markets; and both countries should treat compute and models as layers to access and shape jointly rather than each trying to own alone. This is a materially different, and more actionable, strategic map than either country's current AI strategy documents typically offer.

## Frontier-minus-N: Why Smaller Models Are Often the Sovereign Choice

A related and equally consequential misconception is that a country without frontier-scale data or frontier-scale compute has nothing worth building. It does. Frontier capability earns its considerable cost in a narrow band of genuinely hard problems, principally advanced scientific research, long-horizon multi-step reasoning, and the cyber-offensive and cyber-defensive edge where marginal capability is itself the strategic asset. It is not required, and is frequently the wrong tool, for the great majority of problems that actually determine whether AI improves a country's public services and economic productivity. Independent analysis of enterprise and sovereign deployments increasingly bears this out directly: small, domain-specific models routinely outperform generalist frontier systems within the environment they are built for, at a fraction of the compute, energy and capital cost, and on the dimensions that most directly determine whether an organisation can safely deploy the system in front of customers or regulators, a smaller, purpose-built model has been shown to outperform a far larger generalist counterpart by a wide margin.[8] The same logic is now shaping sovereign strategy at the national level: serious technical assessments of what a country actually needs for AI sovereignty conclude that for most countries the achievable near-term objective is running smaller, open-weight or purpose-built models on infrastructure the country itself controls, which is sufficient for the overwhelming majority of real workloads.[8]

This paper terms the resulting design space frontier-minus-N: a model, or a family of models, deliberately built several generations behind the absolute frontier, trading raw general capability for lower cost, greater auditability, tighter alignment to a specific task, and genuine deployability under the connectivity and budget constraints most real institutions face. A country, or an institution within it, with what this paper calls a subcritical dataset, real and useful, but well short of the scale a frontier laboratory would train on, can still build something valuable; it is simply building at the correct point on the frontier-minus-N spectrum for the problem actually in front of it. India's own sovereign models are the clearest illustration available: Sarvam's 30-billion and 105-billion parameter systems, examined in full later in this section, do not attempt to match the frontier's absolute scale, and compete instead, credibly, on Indian-language tasks where a frontier generalist system's English-centric training gives it no natural advantage.[19] The five frugal AI deployments catalogued later in this paper, in health, in agriculture, in energy forecasting, run on a small fraction of frontier scale and are demonstrably more valuable for it, because the relevant test is whether the system could actually reach the person who needed it, at a cost someone was willing to pay.

This distinction is only useful, however, if it can be applied as a decision rule rather than invoked after the fact to justify whichever model a team already prefers. Applied honestly, the choice between frontier and frontier-minus-N models turns on a small number of questions about the task, its stakes, the available data, cost tolerance, auditability, connectivity and the pace at which the underlying capability changes. A national biosecurity assessment and a regional language government assistant sit at very different points on that spectrum. Much of the public value the UK and India could jointly create depends on recognising that distinction.

## What the UK actually does well, function by function

The UK's advantage lies in vertical AI models: concentration in a small number of mathematically demanding, long-R&D-window, highly certifiable functions. Five stand out on the evidence.[22] The UK leads in AI for science and structural biology, the discipline that produced AlphaFold and that continues, through Isomorphic Labs and the wider Oxford-Cambridge-London research corridor, to train models on centralised, longitudinal datasets, the UK Biobank's full-cohort whole-genome sequencing of 500,000 participants being the clearest example of an asset no other country's health system currently matches for scale and structure.[10] The UK leads in physics-informed and Bayesian methods, embedding physical laws directly into model architectures so that a system simulating a jet engine or a wind turbine needs far less brute-force data than a purely statistical model would, and producing both a prediction and a calibrated measure of how uncertain that prediction is, which is functionally close to mandatory for deploying AI in surgery, systemic financial clearing, or nuclear decommissioning. The UK leads in graph-based relational AI, the mathematics used to track international financial crime networks and to model molecular interactions, a genuinely different technical tradition from the sequence-processing transformer architectures that dominate the language-model conversation. The UK leads in safety, alignment and evaluation science, formalised through the world's first state-backed AI Security Institute, functioning as an internationally trusted referee for frontier-model risk. And the UK leads in hardware-software co-design, in silicon photonics and neuromorphic computing research aimed at breaking the thermal and electrical limits of conventional chips, and in the deep instruction-set and micro-architecture design lineage that produced ARM, a different and arguably more durable form of hardware leadership than owning fabrication capacity.

What unifies these five functions is risk tolerance and time horizon. Each requires the kind of multi-year, low-volume, high-margin research investment that a capital-constrained, wealthy, small country can sustain precisely because it does not need to compete on manufacturing scale to capture the value; the UK captures value by owning the intellectual property and the certification. This is also, read carefully, a description of the UK's translation problem in miniature: every one of these five functions is superb at generating a validated result and comparatively weak at turning that result into a population-scale deployed product, because turning it into a deployed product requires exactly the reach capability the next subsection describes.

## Frugal AI as a deliberate design discipline

Frugal AI is fundamentally a systems design philosophy that seeks to maximise social and economic value by matching technological capability to real-world constraints, institutional capacity and deployment needs. India's reach capability is frequently described, including by some of its own commentators, as jugaad, an improvised, resourceful workaround. That description undersells what has actually been built. India has developed a systematic engineering discipline for designing AI that delivers meaningful capability under severe constraints of cost, infrastructure, connectivity, language diversity and institutional capacity. This broader design philosophy, developed and articulated through the Frugal AI Hub as Frugal AI,[1, 2] begins with the deployment environment itself. It asks whether AI is the appropriate solution at all, what level of capability is actually required, how the system will be maintained, what infrastructure is realistically available, how human judgement will remain embedded, and whether the receiving institution can sustainably operate the technology over time. It extends across workflow design, model choice, infrastructure, lifecycle cost, governance, institutional capability and long-term operational resilience. Technical optimisation is therefore only one component of Frugal AI. Within this broader deployment philosophy sit four engineering strategies that have become especially characteristic of Indian AI development in low-resource environments.

The first is extreme mathematical compression. A conventional neural network stores its weights as 16-bit or 32-bit floating-point numbers; Indian engineering practice, driven by the need to fit a model into the 2 to 50 megabyte memory cache of an entry-level smartphone processor costing only a few dollars, routinely compresses those weights to 4-bit or even 1-bit integer representations, a technique known as quantisation. Combined with knowledge distillation, in which a large, general "teacher" model trains a much smaller, task-specific "student" model that discards every capability outside its narrow purpose, this can shrink a multi-gigabyte model into a system of only a few tens of megabytes while retaining the great majority of its accuracy on the specific task it was built for. Layer-fusion techniques, standard in frameworks such as ONNX Runtime and TensorFlow Lite, then merge consecutive mathematical operations into single hardware instructions, reducing the memory read-write cycles that would otherwise overheat a smartphone with no active cooling system.

The second is network-agnostic architecture, designed on the explicit assumption that connectivity will fail rather than the Western default assumption that it will hold. Inference runs locally, on the device, using an embedded local database, without constant cloud round trips. When results must eventually reach a central system, the payload is compressed into the smallest possible packet and, if standard mobile data fails, the system can fall back to transmission over basic SMS protocols or opportunistic peer-to-peer synchronisation when another user's device reaches an area with connectivity. This architecture allows AI systems to remain functional even in environments where bandwidth is intermittent, expensive or unavailable.

The third is voice- and language-first interface design, a direct response to the reality that text-based interaction excludes a meaningful share of a population spanning hundreds of languages, dialects and uneven literacy levels. Without assuming English literacy or keyboard-based interaction, India's AI ecosystem has invested heavily in multilingual speech recognition, translation and language technologies through initiatives such as AI4Bharat and Bhashini.[18] These systems are purpose-built to parse heavily accented, code-mixed and noisy audio into usable text before any downstream model processes it, significantly expanding access to AI-enabled public services.

The fourth is algorithmic query routing. A lightweight local classifier resolves the majority of routine queries directly on the device and escalates only genuinely anomalous or complex cases to a larger, more computationally intensive system. This hierarchical architecture can reduce inference costs by an order of magnitude compared with blanket cloud-based deployment while maintaining comparable user experience for most interactions. It also enables scarce computational resources to be concentrated where they add the greatest value.

None of these four engineering strategies is a shortcut; each represents a substantial body of computer science with its own research literature and practical innovations. Together they illustrate one important technical dimension of Frugal AI, but they do not exhaust the concept. The cumulative effect is a cost structure and deployment model that a conventional cloud-first architecture cannot readily match at comparable population scale. This discipline constitutes India's genuinely transferable capability and its emerging comparative advantage in the global AI economy.

## The institutional base this discipline was built on

This discipline did not emerge from a single company; it emerged from roughly a decade of accumulated public and academic infrastructure. AI4Bharat, a research lab founded at IIT Madras in 2020, built the empirical foundation years before large language models existed as a mainstream concept, producing the Samanantar corpus, the largest publicly available parallel-sentence dataset spanning English and eleven Indic languages, and IndicTrans2, the first open-source model supporting high-quality translation across all twenty-two scheduled Indic languages.[18] When the government's Bhashini Mission, housed within the Ministry of Electronics and Information Technology, needed a data-management partner to make every digital public service available in every Indian language, it turned to AI4Bharat, which today supplies roughly eighty percent of Bhashini's underlying data.[18] Bhashini itself has crossed one hundred million monthly inferences and hosts more than three hundred AI language models, a scale of multilingual public infrastructure with no real equivalent anywhere else, including in the UK, whose own digital-identity and public-service infrastructure remains substantially monolingual by comparison.

Jugalbandi, a WhatsApp-based assistant built jointly by AI4Bharat, Microsoft Research India and the legal-technology non-profit OpenNyAI, illustrates the four techniques from the previous subsection working as a single deployed system rather than as separate techniques. A villager sends a voice message in one of ten Indian languages; AI4Bharat's speech-recognition model transcribes it, Bhashini's translation layer converts it to English, a language model retrieves the relevant answer from a database of government schemes, and the AI4Bharat text-to-speech model synthesises the reply back into the user's own language and dialect, the entire round trip running over ordinary WhatsApp messaging rather than a bespoke app. At launch the system covered around 171 of India's nearly 20,000 government welfare schemes;[18, 26] what is worth noting is the architecture, since it demonstrates that voice-first, multilingual, low-infrastructure design is now mature enough to be assembled from existing open components rather than requiring bespoke engineering for every new use case, which is precisely the transferable asset a UK public-service deployment would actually be licensing.

It would be a mistake, however, to treat any of this as a frictionless success story: Bhashini's own assessments acknowledge that several of India's lowest-resource languages remain poorly served because the underlying training data for them is simply thin, and no amount of engineering sophistication substitutes for data that does not yet exist in usable form.[18] Jugalbandi's own early documentation flagged the same honest limitation common to every retrieval-based assistant: government schemes change faster than any underlying database can be kept current, and a system this accessible can as easily deliver confident, fluent, out-of-date guidance as it can deliver correct guidance. Frugal engineering solves the reach problem; it does not, by itself, solve the trust problem, which is exactly why this paper treats the two currencies as complements rather than substitutes.

## The UK's own translation problem

If the paper is to take its own thesis seriously, the UK's translation failures need the same specificity as India's verification gap, rather than a passing mention. Isomorphic Labs, the AlphaFold-lineage company applying AI to structure-based drug discovery, raised a six-hundred-million-dollar first external funding round and expects its first clinical trials by the end of 2026.[23] Its constraint is not scientific, it is logistical: converting a computational lead into an approved therapy needs large, genuinely diverse trial populations and pharmaceutical manufacturing capacity at a scale and cost point a UK-only pathway to market cannot easily provide. This is precisely the trust-without-reach failure mode this paper has described throughout, a world-class predictive model with no population-scale route to validation, and it is precisely what India's pharmaceutical manufacturing base and trial-population scale could supply, provided the partnership is structured as genuine co-development rather than a one-way licensing sale that captures none of the downstream value for the Indian partner.

The UK's engineering-biology sector makes the same point at industry scale rather than single-company scale. The government's own sector analysis describes it as Europe's leading biotechnology hub, with more than a thousand active businesses and roughly £3.7 billion of equity financing in 2024 alone,[21] and the sector's own trade literature is candid about its bottleneck: promising science routinely gets stuck between laboratory success and commercial reality for want of infrastructure to scale beyond a pilot, made worse by development timelines that can run to eight years before first revenue, a horizon most venture capital will not underwrite.[21] Read against Table 3's climate-tech row, this is a precisely specified gap that India's agricultural deployment infrastructure, the same CropIn- and Wadhwani-class systems examined in Table 5, is functionally suited to close, since the missing ingredient in both cases is exactly the same: a proven route from validated science to field-scale, real-world deployment. The UK does not lack ambition or capital markets sophistication in the abstract; UK AI start-ups raised more than eight billion pounds in the first half of 2026 alone.[13] What it lacks, structurally and repeatedly, is the specific deployment infrastructure that converts a well-funded, well-evidenced laboratory result into a population-scale product, which is exactly the reach currency this paper has argued India supplies in abundance.

## India's sovereign AI ambitions, and their real limits

A fair account of India's AI ecosystem cannot stop at frugal deployment, because India's own policy establishment has explicitly rejected the idea that deployment is where its ambition ends. The IndiaAI Mission, backed by roughly £1.2 billion, treats sovereign frontier-model capability as a strategic objective in its own right.[11] Sarvam AI, a Bengaluru company founded by Vivek Raghavan and Pratyush Kumar, the latter a co-founder of AI4Bharat, was selected by government tender in April 2025 to build India's sovereign large language model. In February 2026, at the India AI Impact Summit in New Delhi, Sarvam unveiled two open-weight foundation models trained from scratch on Indian infrastructure using more than a thousand H100-class chips at a facility called the Shakti cluster: a thirty-billion-parameter mixture-of-experts model and a hundred-and-five-billion-parameter model with a hundred-and-twenty-eight-thousand-token context window, both released under an open licence and presented on stage alongside the Prime Minister.[19] BharatGen, an IIT Bombay-led consortium, launched Param2, a seventeen-billion-parameter model intended to sit inside government digital services; Tech Mahindra released an eight-billion-parameter Hindi-first educational model; and Krutrim, founded in 2023 by the Ola co-founder Bhavish Aggarwal and briefly India's first AI unicorn, had by the 2026 summit lost much of its early momentum, its attempt at a competitive general-purpose model not delivering what its early launch promised.[19] Together these form a genuine, layered national AI stack: a frontier story running in parallel with India's deployment story.

It is also, on the evidence available in mid-2026, a story with specific and honestly acknowledged limits. Independent technology journalism has pointed out that none of Sarvam's models appeared on any standard open, third-party leaderboard at launch; there was no accompanying academic paper with methodology or ablation studies; and the flagship benchmark used to justify the headline performance claims was designed by Sarvam, translated by Sarvam, and judged by a model Sarvam itself selected using prompts Sarvam itself had chosen.[19] This is a structural observation about evaluation infrastructure: India has begun building the evaluation infrastructure a serious sovereign-AI programme needs, but does not yet have an independent, nationally trusted scoreboard with the authority to arbitrate a claim this consequential, at a moment when the same models are already inside live public infrastructure, including Aadhaar's voice services and an insurer's systems serving roughly eighty million customers.[19] The second limit is more structural still: training a model genuinely competitive with the frontier requires sustained access to tens of thousands of advanced accelerators for months at a time, at a capital cost running into the hundreds of millions of pounds; India's Sarvam allocation ran to a few thousand chips for a period of months, a gap that no amount of quantisation or distillation efficiency fully closes, because those techniques reduce the compute needed for a given capability level without eliminating the advantage of simply having an order of magnitude more compute to begin with.

## Table 2 · A functional comparison

The table below restates the two paths at the level of specific function rather than national reputation, and is the basis for the comparative positioning in Figure 1.

> Table 2 · UK and India positioning across eight AI functions. Illustrative characterisation drawn from the evidence in this paper; no scoring methodology underlies it. Sources: 8–11, 13, 18, 19, 22.

| AI / Technology Function | UK Position | India Position |
|---|---|---|
| **Foundational model training & compute** | Strong in small/medium scientific and frontier-adjacent models; compute-constrained relative to the US frontier. | Emerging; sovereign models trained on allocated national compute, not yet frontier-scale or independently verified. |
| **AI safety, certification & evaluation** | World-first state-backed evaluation institute; internationally trusted referee function. | Nascent; prioritises inclusion and deployment speed over frontier-risk evaluation infrastructure. |
| **Deep-tech & physics-informed science AI, Vertical AI** | Global leadership in structural biology AI, physics-informed and Bayesian methods. | Applies open-source frontier models to manufacturing and clinical trial execution rather than building the underlying science. |
| **Quantum & hardware co-design** | Over 160 active quantum companies;[24] deep chip micro-architecture and photonics research base. | National Quantum Mission underway; assembly and fabrication scale, limited design-stage hardware R&D to date. |
| **Institutional / capital-markets AI** | Leads in algorithmic trading, AML graph networks, insurtech, and regulatory sandbox design. | World-class consumer and retail fintech; institutional and cross-border wholesale layer far less developed. |
| **Mass consumer deployment & digital public infrastructure** | Fragmented, still-modernising digital identity and payment rails. | World-leading sovereign DPI at population scale: Aadhaar, UPI, Account Aggregator. |
| **Frugal / edge / low-resource engineering** | Limited commercial incentive to build for the constrained case; little domestic capability. | One of the world's deepest capabilities in cost-, connectivity- and power-constrained AI. |
| **Multilingual / low-resource NLP** | Substantially monolingual public-service and identity infrastructure. | Deepest multilingual NLP infrastructure outside major US labs, purpose-built for 22 scheduled languages. |

## Figure 1 · Functional positioning: trust versus reach

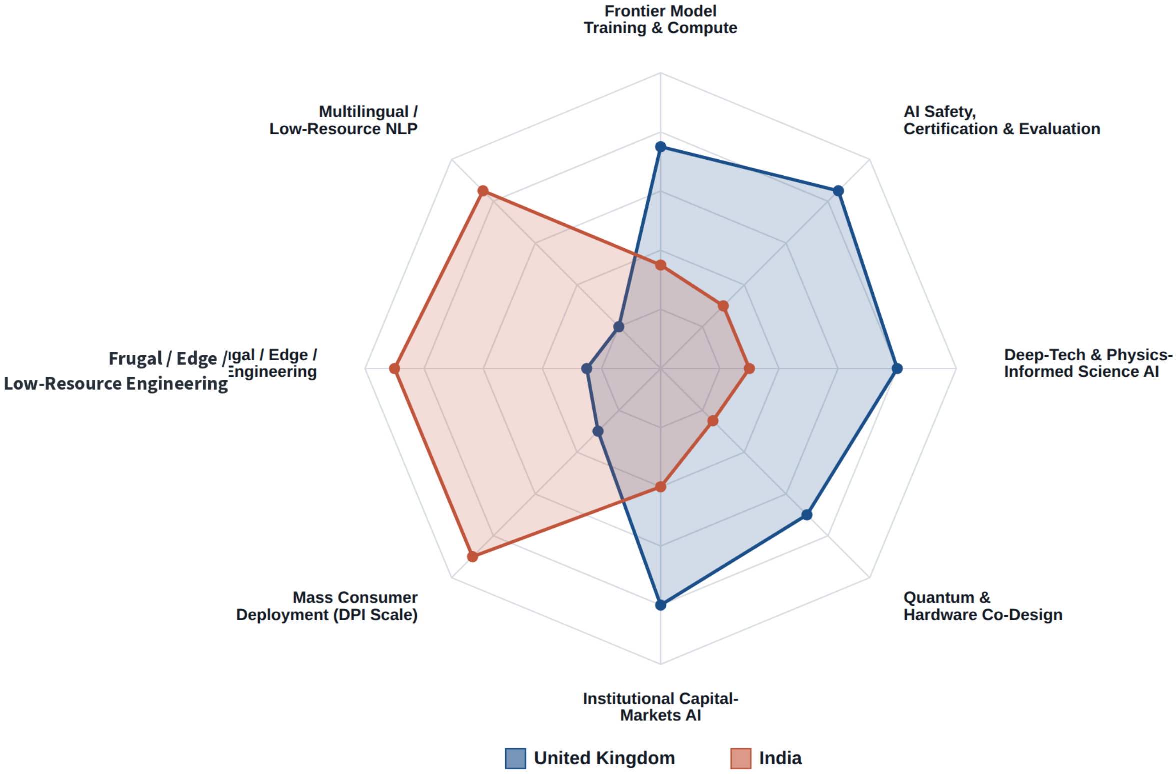


Illustrative positioning across the eight functions in Table 2, drawn from the qualitative evidence assembled in this paper rather than a precise scored index. The two shapes overlap very little: the UK's profile and India's profile read as near-complements. Vertices represent the authors' qualitative reading of the evidence cited in this report. No labelled numerical scale or weighted scoring methodology underlies this visual.

# 3

# Where the Two Paths Actually Meet: Sector Intersections

A functional comparison is only useful if it identifies where the two profiles in Figure 1 combine into something neither could build alone. This section works through five specific intersections in enough technical detail to be commercially actionable, rather than resting at the level of general synergy. Each is structured the same way: the UK catalyst, the specific input only the UK ecosystem currently supplies at this quality; the India testbed, the specific scale or infrastructure only India currently supplies; the co-developed output, the new asset that results; and the India-side ecosystem upgrade required before the intersection can actually be built, because in each case one or both ecosystems still lack a specific regulatory, institutional, data or deployment capability.

Table 3 · The macro-sector value matrix: where UK catalyst meets India testbed. Sources: 10, 14, 20–22.

| Vertical | UK Catalyst (Input) | India Testbed (Target) | Co-Developed Output | India Ecosystem Upgrade Required |
|---|---|---|---|---|
| **Life sciences & TechBio** | Structural biology AI, UK Biobank-scale genomic data, physics-informed molecular design | Mass clinical-trial execution, high-throughput biologics manufacturing, multi-morbidity patient pools | Precision-medicine and next-generation biologics pipelines with compressed discovery timelines | Unified FHIR-standard health-data interoperability across ABDM; stronger bio-IP enforcement |
| **Institutional capital markets & fintech** | FCA sandbox model, AML/graph-network fraud detection, cross-border compliance AI | UPI and Account Aggregator real-time transaction data at population scale | Automated cross-border B2B settlement and fraud-detection engines trained on live transaction graphs | RBI sandbox liberalisation for foreign fintech experimentation; GIFT City legal alignment |
| **Climate technology & carbon removal** | Synthetic biology and plant genomics for enhanced-sequestration planting stock; earth-system modelling | Tissue-culture propagation at commercial scale; afforestation and agroforestry deployment channels | Field-validated high-sequestration planting stock with an associated measurement and verification protocol | A credible carbon measurement, reporting and verification registry; settled land title and benefit-sharing rules |
| **Deep mobility & autonomous edge systems** | End-to-end reinforcement learning for navigation; physics-informed safety verification | Dense, chaotic real-world telemetry; extreme quantisation and edge-compute engineering talent | Affordable autonomous fleets for logistics, agri-drones and warehousing in unstructured, low-infrastructure settings | Statutory liability and insurance frameworks for autonomous machinery |
| **Semiconductors & AI hardware** | Chip micro-architecture design IP, photonics and neuromorphic research lineage | Semiconductor Mission fabrication capacity and assembly-scale manufacturing | Sovereign next-generation chips combining UK design with Indian fabrication | Curriculum shift from VLSI physical layout toward micro-architecture co-design; expanded design-linked incentives |

## Life sciences: moving beyond generics

The CETA framework is explicit that its intent is to shift the trade dynamic from low-cost supply toward joint intellectual property development,[12, 14] and life sciences is the sector where that shift is most concretely available. The UK's structural biology and genomic AI capability, anchored in the UK Biobank's full-cohort whole-genome data and the AlphaFold research lineage,[10] produces predictive models of extraordinary quality but needs large, genuinely diverse trial populations and high-throughput manufacturing to convert a computational lead into an approved therapy at global scale. India's pharmaceutical sector, the world's largest supplier of generic medicines and vaccines by volume,[22] sits at exactly the opposite end of the same value chain: unmatched execution capacity, historically reliant on chemical blueprints designed elsewhere. The bottleneck preventing the two from combining is not scientific, it is data architecture: India's health information remains siloed across thousands of public and private providers, and until it is genuinely unified under interoperable standards, a UK-trained diagnostic or discovery model cannot safely ingest it at the scale either side would need to make this partnership more than a licensing transaction.

## Deep mobility: the edge-compute opportunity

UK autonomous-systems companies, Wayve being the clearest example, have pioneered end-to-end reinforcement learning for navigation, teaching a system to drive from raw sensor experience rather than from expensive, location-specific high-definition maps.[25] That architecture is scientifically elegant and commercially heavy: it still assumes access to substantial onboard or cloud compute that makes it unviable for mass-market vehicles, agricultural drones, or low-cost warehouse robotics. This is precisely where India's quantisation and edge-compilation discipline, described in Section 2, becomes a direct commercial input: applying INT4/INT8 quantisation and knowledge distillation to a UK reinforcement-learning model can compress it to run on hardware costing a small fraction of the original architecture's requirement, opening an entirely new addressable market of affordable autonomous fleets for logistics, agriculture and warehousing across markets that cannot afford Western-grade compute. The obstacle is legal: neither country has yet settled a clear statutory liability and insurance framework for autonomous machinery operating in unstructured public environments, and that gap is what currently prevents this intersection from moving past pilot stage.

## Climate technology: the least proven of the five

The climate intersection pairs UK synthetic biology and plant genomics with India's tissue-culture propagation infrastructure and its afforestation and agroforestry deployment channels. The proposition is that engineered planting stock with materially higher sequestration rates could be propagated and planted at a scale and cost point no other pairing could reach. It belongs in this section because the complementary fit is real and the deployment infrastructure on the Indian side already exists at commercial scale. It is set out with more caution than the other four because its central scientific claim, the sequestration multiplier itself, has not been field-validated, and because the verification infrastructure that would allow any resulting carbon to be credibly counted does not yet exist in either country. The ecosystem upgrade required here is therefore twofold and unusually demanding: a measurement, reporting and verification registry with enough authority to be trusted by a carbon market, and settled rules on land title and benefit sharing for the communities on whose land the planting would happen. Both are prerequisites rather than refinements, which is why Section 7 places this intersection behind life sciences in the sequencing despite the strength of the underlying fit.

## Semiconductors and institutional fintech: the shorter cases

Two further intersections are worth stating briefly rather than at length, because the pattern is now familiar. In semiconductors, the UK's chip micro-architecture design lineage, in the tradition that produced ARM, meets India's rapidly scaling Semiconductor Mission fabrication capacity;[27] the constraint is that Indian engineering curricula remain oriented toward physical layout assembly rather than the micro-architecture co-design work needed to make Indian fabs anything more than a contract manufacturer for foreign-designed chips. In institutional fintech, the UK's regulatory sandbox model and its graph-network fraud and anti-money-laundering capability meet India's uniquely deep real-time consumer transaction data through UPI and the Account Aggregator framework; the constraint here is regulatory rather than technical, with the Reserve Bank of India's institutional risk-aversion, discussed further in Section 6, currently limiting how far foreign compliance and risk-assessment technology can actually plug into Indian capital pools. It is worth being precise about the limits of India's fintech strength. UPI's scale is primarily a retail and consumer achievement, involving tens of billions of low-value transactions processed at minimal cost. It does not by itself demonstrate institutional readiness for the far smaller number of far larger, far more complex cross-border wholesale and trade-finance transactions the UK side of this intersection is built for. Treating retail transaction volume as evidence of institutional-fintech maturity would be a category error this paper is careful to avoid.

## Table 4 · Maturity staging across the five intersections

The five intersections in Table 3 are not equally mature, and a credible commercial partner needs to know which stage each has actually reached, lab-validated science with no deployment partner yet identified, an active pilot, or a genuinely scaled, revenue-generating joint operation, before committing resources. Table 4 below stages each intersection against that test.

> Table 4 · Where each intersection actually sits today, versus where Table 3 says it could go.

| Intersection | Stage today | What is validated | What is not yet validated |
|---|---|---|---|
| **Life sciences & TechBio** | Institutional bridge exists (Joint Centre health workstream); no joint clinical programme yet operating at scale | UK discovery science; India manufacturing and trial-execution capacity, separately | Whether Indian health data can be safely ingested by UK models at the standard required |
| **Institutional fintech** | Pre-pilot; regulatory sandbox linkage not yet built | UK compliance AI in UK markets; India's retail transaction rails at consumer scale | Whether UK institutional tools function on India's wholesale/trade-finance volumes at all |
| **Climate tech / engineered trees** | Early research stage; sequestration multiplier not yet field-validated | India's tissue-culture propagation infrastructure, independently, at scale | The core scientific claim itself, and any credible verification registry |
| **Deep mobility / autonomous edge** | Component pilots exist separately; no joint UK-India deployment yet | UK reinforcement-learning navigation; India's quantisation and edge-compilation methods | Liability and insurance frameworks for autonomous machinery in either jurisdiction |
| **Semiconductors** | Earliest stage; ecosystem upgrade (curriculum, incentives) not yet begun | UK design lineage; India fabrication-scale investment, separately | Whether Indian design talent can be built at the pace fabrication capacity is scaling |

# 4

# Startups as the Transmission Mechanism

Everything in Sections 2 and 3 describes a structural complementarity. It says nothing about how that complementarity actually moves between the two countries in practice, today, before any of the institutional architecture in Section 8 exists. The honest answer is that it already moves, through individual companies, and examining two of the clearest cases in detail matters for a reason beyond illustration: both were built by the market itself, without a purpose-designed corridor, which is exactly why the corridor this paper proposes is an attempt to widen a channel that has already proven it carries real traffic.

## India to UK: Qure.ai and the NHS

Qure.ai, the Mumbai-founded health-AI company examined here and included among the five documented deployments in Table 5, is the clearest and most extensively verified case of Indian AI capability moving into UK institutional use. Its qXR chest X-ray triage tool holds CE Class IIb certification under the EU Medical Device Regulation and the corresponding MHRA approval required for commercial deployment in the UK,[6] and it is not a pilot confined to one enthusiastic trust: it is live, on a commercial basis, procured through the NHS Shared Business Services framework, and running inside a genuinely wide spread of NHS organisations, including Bolton NHS Foundation Trust, which became the first UK trust to deploy the tool in April 2020, during the acute phase of the Covid-19 pandemic, University Hospitals of Leicester, East Kent Hospitals, Frimley Health, University College London Hospitals, and a multi-trust deployment across the North East London Cancer Alliance spanning Barts Health, Barking Havering and Redbridge, and Homerton Healthcare.[6]

The reported clinical results are specific rather than promotional. At Frimley Health, a pilot found qXR correctly identified normal chest X-rays with 99.7% accuracy, a result projected to redirect up to 58% of a consultant radiologist's caseload to more complex work, saving as much as two hours a day per radiologist.[6] In North East London, the tool is targeted at cutting the wait for a significant chest X-ray finding from three weeks to three days, against a backdrop of chest X-ray volumes that rose 94% between 2017 and 2022 and a national radiologist shortfall the Royal College of Radiologists projects will reach 40% by 2027.[6] The company's own account of its route to this position is instructive for the paper's wider argument: it entered the UK through the unglamorous, cumulative mechanism of NHS research and innovation funding, an SBRI Healthcare and NHS Cancer Programme-funded multi-site study called LungIMPACT, real-world evaluation against NICE's Evidence Standards Framework, and only then commercial procurement at scale.[6] Reach, in other words, still had to clear a trust gate, exactly as this paper's central thesis predicts, and the UK's evaluation and procurement architecture is what converted a frugal Indian product into a credible NHS asset.

## UK to India: Oxford Nanopore and the genomics corridor

Oxford Nanopore Technologies, the Oxford-founded sequencing company built on a fundamentally different, portable, real-time approach to reading DNA and RNA, is the clearest current case running in the opposite direction. The company has operated in India since 2017, but the relationship changed register in April 2025, when it signed Letters of Intent with two of India's Biotechnology Research and Innovation Council institutes, BRIC-CDFD in Hyderabad and BRIC-NIBMG in Kolkata, to establish Genomic Centres of Excellence, alongside the opening of a new Oxford Nanopore office in Bengaluru.[7] The agreements are specific rather than aspirational: the BRIC-CDFD partnership targets rare-disease research and clinical characterisation using Oxford Nanopore sequencing validated for the Indian clinical context, while the BRIC-NIBMG partnership targets multi-omics applications including newborn screening for maternal and child health and genomic research into oral cancer, alongside antimicrobial-resistance surveillance under the wider One Health agenda. The UK's own Deputy High Commissioner for Western India framed the agreement explicitly as delivery under the UK-India Technology Security Initiative,[7] the same institutional architecture referenced throughout this paper's discussion of the wider bilateral relationship.

The direction of value here mirrors Qure.ai's almost exactly, reversed. Oxford Nanopore supplies the trust currency, a sequencing platform, and the scientific and regulatory credibility behind it, that neither BRIC-CDFD nor BRIC-NIBMG could build from nothing on a comparable timeline. India supplies the reach currency: population-scale genetic diversity that UK-based genomic research cannot access at home, a route into national public-health priorities, newborn screening, oral cancer, antimicrobial resistance, that only operate at the scale India's health system provides, and a Bengaluru office as the operational base for exactly the kind of scaled clinical validation this paper has argued UK life-sciences companies structurally lack. Read together with Section 3's life-sciences row in Table 3, this is that intersection, already under construction, roughly a year ahead of this paper's publication.

### What the two cases say about the corridor this paper proposes

Both companies reached the other market the hard way, through years of regulatory validation, institutional relationship-building and iterative deployment rather than through any dedicated UK-India mechanism, because no such mechanism yet exists. That is precisely the gap the institutional architecture proposed in Section 8 is designed to address: making the path that Qure.ai and Oxford Nanopore each had to build over many years available to the next company in months. A joint demonstrator modelled on the LungIMPACT study, regulatory sandbox linkages that reduce duplicative approvals, and a dedicated carrier institution to coordinate projects across both ecosystems are not untested innovations. They are the deliberate, repeatable institutionalisation of pathways that the market has already demonstrated can succeed. Such a carrier institution would also form one element of the broader institutional absorption capacity required to move projects from bilateral agreement to sustained deployment at scale.

# 5

# How Capability Should Move: Four Pathways

The intersections in Section 3 identify where value could be created. They do not, by themselves, explain how a specific piece of UK or Indian capability physically and institutionally moves from where it exists today to where the intersection needs it. This paper organises that movement around four pathways, chosen because each corresponds to a distinct stage of the journey from an idea to a commercially or institutionally deployed system, and because each currently exists, at best, in fragmented, ad hoc form.

The first pathway runs from research to ventures, and it is where the relationship is weakest today. Both countries generate more promising research than either commercialises; the UK's own innovation-policy literature has long lamented a valley of death between publicly funded university research and venture-backed commercialisation, and India faces a structurally similar gap between its research institutes and its venture ecosystem. Closing this pathway needs reciprocal fellowships and secondments, letting a UK postdoctoral researcher spend eighteen months embedded in an Indian deployment-focused lab and an Indian researcher spend an equivalent period inside a UK frontier-safety or synthetic-biology group, paired with a light-touch joint seed fund willing to back the venture that results.

The second pathway runs from models to applications, and is fundamentally a compute-access and co-location problem. UK frontier models and UK compute capacity, and India's applied-development talent and uniquely large deployment data, are complementary halves of a single production process that currently runs as two disconnected assembly lines; a functioning pathway looks less like a memorandum of understanding and more like a specific co-location arrangement granting Indian application teams priority, cost-effective access to UK compute in exchange for the resulting applications being validated first against Indian deployment conditions, the harshest robustness test either ecosystem can offer.

The third pathway runs from pilots to institutional deployment, and recognises that moving from pilot to sustained deployment requires far more than technical validation. Institutions must be able to redesign workflows, govern data, assign accountability, finance ongoing operations and continuously learn from implementation. This paper refers to these enabling conditions collectively as institutional absorption capacity, a concept developed further in Gandikota (forthcoming) and through the Frugal AI Hub's ongoing research. India's experience of deploying AI through digital public infrastructure at genuinely national scale speaks directly to the problem of moving NHS and local government AI pilots past proof of concept. The mechanism is procurement and regulatory design rather than new capital: mutual recognition between the UK's emerging cross-economy AI sandbox and India's sector-specific sandboxes, and a small number of jointly funded, explicitly time-bound deployment missions designed to run from pilot to sustained institutional use rather than stopping at demonstration.[2, 5]

The fourth pathway carries a solution proven domestically in either country out into third markets, and is examined at length in Section 8, because it is the mechanism through which a bilateral corridor becomes something larger. A diagnostic tool validated in Indian deployment and then hardened against UK regulatory and evaluation standards becomes credible for other resource-constrained health systems across the Commonwealth and beyond in a way that a purely domestic product from either country is not.

# 6

# The Structural Obstacles

None of the four pathways will build itself, and the obstacles are specific enough to deserve more precision than the general language of friction that policy papers often reach for.

The most consequential and least understood obstacle concerns data governance. India's Digital Personal Data Protection Act, passed in 2023 and brought into phased effect through rules notified in November 2025, is considerably more permissive on cross-border transfer than its reputation suggests: rather than the strict localisation regime earlier drafts proposed, the final Act adopts a blacklist model in which personal data may generally be transferred outside India by default unless the government specifically restricts a destination country or data category.[20] Most UK-India AI collaboration involving personal data is not, in principle, blocked by Indian law. What complicates this cleaner picture is that sector-specific regulators impose materially stricter rules on top of it: the Reserve Bank of India requires certain categories of payment data to be stored exclusively within India regardless of what the general statute permits, and equivalent sectoral rules apply to insurance and securities data, meaning a UK fintech partnership faces a genuinely more restrictive environment than a UK health-AI partnership even under the same overarching law. A further complication sits inside the implementing rules rather than the headline Act: companies designated as Significant Data Fiduciaries, a category the government has not yet fully defined, may be required to keep specific categories of sensitive data, health and biometric information being the most frequently discussed candidates, inside India altogether.[20] Indian legal commentary has been sharply critical of this provision as currently drafted, arguing that an open-ended government power to designate data types as non-exportable creates exactly the regulatory uncertainty that deters the investment India's own digital economy strategy is trying to attract.[20]

Compute dependency is the second major obstacle, and cuts in a direction opposite to the usual assumption: it is not that India lacks compute in absolute terms, it is that India's compute is concentrated in government-allocated national missions rather than being commercially available at the depth and flexibility a fast-moving joint venture needs. Intellectual property allocation is the third obstacle, and remains genuinely unresolved: CETA's source-code protection is a meaningful floor, but it says nothing about who owns the output of a jointly trained model, or how royalties should be split when UK algorithmic design and Indian deployment-data validation are each essential but difficult to value against each other by any existing convention, a problem this paper's own Table 3 makes concrete in the life-sciences and climate-tech rows. Procurement rules on both sides, designed for conventional software purchasing rather than for systems that improve iteratively after deployment, slow adoption even where technology and appetite are both ready. The Reserve Bank of India's traditionally conservative, risk-averse regulatory posture, in contrast to the FCA's pioneering sandbox model, is a specific and named obstacle to the institutional-fintech intersection identified in Section 3, and closing it requires a deliberate policy choice inside India rather than a bilateral instrument alone.

A fifth obstacle sits specifically inside the first pathway from Section 5, research to ventures, and is worth naming precisely rather than folding into a generic mobility comment: moving a researcher between the two ecosystems for a meaningful embedded placement, eighteen months being the horizon this paper has suggested as useful, currently has no purpose-built route on either side. The CETA's Double Contribution Convention solves a real but narrower problem, double social-security taxation for corporate secondees, and was never designed for an early-career researcher moving labs rather than moving jobs.[12] The absence of a dedicated instrument here is the single most direct explanation for why the research-to-ventures pathway remains, on the evidence of this paper, the weakest of the four.

Beneath all of these sits the most fundamental obstacle: responsibility for the relationship is currently dispersed across trade officials, science and innovation agencies, sector regulators and individual company relationships on both sides, with no single institution whose explicit mandate is to identify when a promising project is stalled between these silos and carry it through to deployment. This is not a hypothetical risk. It is the pattern visible across almost every intersection identified in Table 3: the life-sciences intersection stalls on health-data interoperability that sits with neither trade officials nor either country's AI ministry; the climate-tech intersection stalls on carbon-verification infrastructure that sits with neither country's environment ministry acting alone; and the deep-mobility intersection stalls on liability frameworks that fall outside the remit of AI policymakers altogether. Each obstacle is individually solvable. None is clearly owned. The United Kingdom supplied a sharp illustration of the problem on 21 July 2026, when the Department for Science, Innovation and Technology was dissolved and its functions divided three ways: science, innovation and research funding to the newly created Department for Business, Innovation, Science and Trade; artificial intelligence strategy, public sector adoption and the AI Security Institute to the Cabinet Office; and online harms and digital identity to the Department for Digital, Culture, Media and Sport. Artificial intelligence gained its first Cabinet-attending minister in the same announcement, a post held jointly across two of those departments. A corridor project stalled between three departments now has further to travel before anyone is accountable for carrying it.[32] It is this governance gap, more than any other, that will determine whether the pathway architecture set out in Section 5 survives contact with the institutional realities of two large economies. The institutional design proposed in Section 8 responds directly to this challenge. In particular, a dedicated carrier institution would provide one element of the broader institutional absorption capacity required to move projects from bilateral agreements and successful pilots to sustained deployment at scale.

# 7

# The Value Case and Sector Sequencing

A full econometric estimate of the value at stake is worth commissioning and this paper does not attempt to substitute for one; producing a precise-looking number without primary UK and Indian unit-cost data would offer the appearance of rigour alone. What can be done responsibly is to identify with precision where the value comes from and how sectors should be sequenced, so that a future modelling effort knows exactly what it is trying to measure.

Six sources of value follow directly from the analysis so far. Reduced deployment cost, since pairing UK-origin technology with India's frugal-engineering method lowers the absolute cost of reaching scale in either country, as the individual case studies below already demonstrate at the unit level. Faster commercialisation, since access to India's deployment environments and manufacturing base shortens the distance between a UK laboratory result and a paying customer, precisely the bottleneck the UK's own engineering-biology and life-sciences sectors describe as the reason promising science stalls before revenue. Market access in both directions, India gaining a credible route into regulated developed markets, the UK gaining a credible route into India's domestic market and other price-sensitive economies. Shared infrastructure, since co-locating compute and deployment activity raises utilisation on otherwise under-used, expensive assets on both sides. Talent mobility, treated as commercial and research infrastructure rather than an immigration question, raising the rate of new venture formation on an already large UK AI venture base that raised over eight billion pounds in the first half of 2026 alone.[13] And joint development explicitly aimed at third markets, a revenue line that exists only if the corridor is built to look outward from the start, examined fully in Section 8.

## Five documented deployments, and what they are honestly worth

Five independently documented Indian deployments illustrate the reduced-deployment-cost and faster-commercialisation arguments concretely rather than abstractly, and each maps onto a specific, named UK application. Two caveats apply throughout: several unit prices are not publicly disclosed and would need direct commercial negotiation to establish, and the UK-side value figures must be sourced from NHS, DEFRA and Ofgem data rather than modelled from Indian proxies, since labour costs, regulatory requirements and existing infrastructure all differ.

Table 5 · Five Indian AI deployments and potential UK adaptation contexts. Outcomes as publicly reported; see sources.

| India solution | Reported outcome | Cost signature | UK target application | Value logic |
|---|---|---|---|---|
| **Qure.ai (qXR)** | Flagged ~90% of missed or mislabelled chest X-rays[6, 15] | Software overlay on existing imaging hardware; no new capital | NHS diagnostic-imaging backlog | Earlier catch of missed pathology; radiologist time freed |
| **Niramai** | Portable breast-screening alternative to fixed mammography[15] | Handheld device; lower capex than fixed units | NHS screening in rural / underserved areas | Avoided late-stage treatment cost from earlier detection |
| **Wadhwani CottonAce** | Cut pesticide costs ~25%; raised farmer income ~20%[16] | Offline smartphone app; $3.3m grant to scale | UK arable / horticulture disease detection | Avoided crop loss and reduced input spend per hectare |
| **CropIn (SmartFarm)** | ~30% yield uplift for groundnut farmers in Karnataka[17] | SaaS; near-zero marginal cost per added farmer | UK smallholder / tenant precision agriculture | Yield uplift per hectare against low deployment cost |
| **Solar forecasting (Gujarat)** | ~30% smaller gap between forecast and actual solar output[17] | Software; pilot-stage, cost not disclosed | UK grid balancing; community energy schemes | Avoided backup-generation and curtailment cost per MWh |

## A worked example of the reduced-deployment-cost logic

The reduced-deployment-cost case can draw on the Frugal AI Hub and UNICC's published Total Cost of Ownership and Return on Investment framework.[3, 4] That framework measures value through time savings, cost avoidance, error reduction and additional capacity, while treating deployment cost across the full lifecycle rather than as a single headline figure. It then carries that measurement through to social impact, mapping deployment outcomes onto development indicators, which is the level at which a policy audience will want a corridor's value expressed. For the population-scale interventions examined here, one part of the value case can be expressed as $N \times p \times \delta$: the number of units reached, the rate at which the intervention changes the outcome, and the monetised value of that change. A credible calculation would also require a decomposed UK-side TCO covering infrastructure, data, models and software, personnel, integration and governance.

## Figure 2 · Where each side's surplus sits

**UNITED KINGDOM**

*Capital & compute surplus*

**£28bn**
AI Growth Zone investment

**£8.2bn**
AI venture capital, H1 2026

**£1.1bn**
AI Hardware Plan

**INDIA**

*Deployment & reach surplus*

**2,000+**
Applied AI use cases

**1,200+**
AIKosh datasets

**217+**
Pre-trained models

*Two different units of account, deliberately not converted to one. Figures as reported publicly; see Sources.*

> The two ecosystems' distinctive surpluses are set side by side without forcing them onto a shared scale, because no shared scale exists. The United Kingdom's surplus is capital and compute, denominated in pounds: confirmed private investment behind its AI Growth Zones, AI-specific venture capital raised in the first half of 2026, and the government's AI Hardware Plan.[9, 13] India's surplus is deployment and reach, denominated in count of use: live applied AI use cases, open datasets hosted on AIKosh, and the pre-trained models available through it.[11] Reading the two columns as competing magnitudes would miss the point of placing them together; the value of each figure lies in what kind of asset it represents. That is this paper's central claim rendered as a figure: the UK and India hold two different currencies of AI value, and the corridor this paper proposes exists precisely to let each currency purchase what the other cannot.

## Sequencing: which sectors first, and why

Not every intersection identified in Table 3 is equally ready to be built on today. Four criteria determine sequencing: the strength of the underlying complementary fit; the readiness of deployment infrastructure on both sides; whether an institutional bridge already exists to deepen rather than construct from nothing; and the commercial and social urgency of the underlying problem. Life sciences ranks highest on every criterion: the India-UK Joint Centre for AI already gives it a partial institutional home,[21] the complementary fit demonstrated in Table 3 is unusually precise, and both sides bring validated capability. Climate tech has an equally precise fit and the same partial institutional home, but its central scientific claim is not yet field-validated and the verification infrastructure a carbon market would require does not exist in either country, so it belongs immediately behind life sciences rather than alongside it. Institutional fintech and deep mobility have equally strong complementary fit but lack any institutional bridge today, and belong in a second tier until the sandbox linkages and liability frameworks described in Sections 3 and 6 are actually built. Semiconductors, while strategically important, requires the deepest and slowest ecosystem upgrade, a curriculum and incentive shift rather than a regulatory adjustment, and belongs furthest out in the sequence. The discipline here is a judgement about where a bridge already exists to be deepened against where one must be constructed before Section 5's pathways can be used at all.

# 8

# Building the Corridor, and Beyond Bilateral

The obstacles catalogued in Section 6, above all the absence of any institution whose explicit job is to carry a project across the gap, point toward a specific institutional design rather than a general call for more cooperation. Seven instruments, each with a named owner and a realistic first step, would turn the pathways of Section 5 into standing infrastructure. On the United Kingdom side, the counterpart for each instrument now sits with the Minister of State for Artificial Intelligence, a post created on 20 July 2026 and held jointly across the Cabinet Office, which holds artificial intelligence strategy, public sector adoption and the AI Security Institute, and the Department for Business, Innovation, Science and Trade, which holds science, innovation and trade.[32] Naming that counterpart matters for the same reason the rest of this section does: an instrument without an owner is a proposal rather than standing infrastructure.

Joint demonstrators are the unit of delivery: a twelve-month project, jointly funded and jointly owned by named institutions on each side, health being the natural first choice given the existing Joint Centre workstream,[21] with a defined budget and a success metric agreed before the project starts rather than negotiated after it finishes, behaves entirely differently from an open-ended workstream that can drift without ever being judged a success or a failure. A first demonstrator judged honestly, including if it is judged to have failed, would do more to establish the credibility of everything that follows than several further years of joint statements. Shared testing environments follow directly: common testbeds and the specific regulatory-sandbox mutual recognition identified in Sections 3 and 6, most immediately between the UK's emerging cross-economy AI sandbox and the Reserve Bank of India's fintech sandbox, so that a product cleared as safe in one jurisdiction does not restart its safety case from nothing in the other.

Researcher and founder mobility needs a purpose-built instrument rather than reliance on the general skilled-worker migration systems either country already operates, precisely because, as Section 6 set out, the CETA's existing Double Contribution Convention was designed for corporate secondment and does not solve the specific, narrower problem of an early-career researcher wanting an eighteen-month embedded placement in a partner institution's lab.[12] Procurement pathways address a problem that sounds administrative but is currently one of the largest sources of delay: agreed routes by which public buyers on each side, NHS trusts, Indian state health departments, can adopt a corridor-developed solution without each individually rebuilding the entire commercial and safety case from first principles. Co-investment mechanisms should target the missing-middle funding gap directly, the stage between a grant-funded result and a venture-backable company that neither country's existing funding architecture serves well, particularly for the eight-year pre-revenue horizons the UK's own engineering-biology sector describes, and should be capitalised jointly by both governments on terms that explicitly accept that long horizon rather than default to conventional venture return timelines.

Language-data partnerships should draw explicitly on the AI4Bharat and Bhashini infrastructure described in Section 2 rather than attempting to duplicate it, since the comparative advantage in multilingual data architecture sits overwhelmingly on the Indian side of this particular instrument; the UK's contribution here is less about building parallel capability and more about co-funding the specific low-resource languages, Section 2's example of Bhashini's own thin coverage of languages such as Bodo and Santali being the clearest gap, [18] where investment has not yet followed need. And a small number of named carrier institutions, whose sole mandate is to notice when a project is stuck between two bureaucratic silos of exactly the kind catalogued in Section 6, and to carry that project through to resolution rather than merely convene another meeting about it, is the instrument this paper regards as most important precisely because it is the least glamorous. Every other instrument in this list can exist on paper without producing outcomes if no one is accountable for making the connections between them actually happen; this seventh instrument is the answer to that risk, and its absence, more than any other single factor, is why the relationship's considerable diplomatic and legal scaffolding has not yet converted into a working corridor. The corridor itself should be understood as more than a channel for moving technology, researchers and investment between two countries. It must also become a mechanism for transferring institutional capability. Successful AI adoption depends on access to models and compute and, equally, on the capacity to evaluate systems, redesign procurement, govern data responsibly, train users, maintain deployed systems and continuously learn from implementation. In this sense, the corridor should carry the institutional practices that enable AI to be absorbed as effectively as it carries the technologies themselves. Building institutional absorption capacity is therefore not a parallel objective to technology transfer; it is the condition that allows technology transfer to create lasting economic and public value.

The fourth pathway from Section 5, carrying a domestically proven solution into third markets, is where the argument ultimately points. The strongest version of this claim is not that the UK and India should build a closed bilateral arrangement and stop there; it is that a working method for pairing trust with reach, once genuinely tested inside this relationship, becomes a template other countries with a comparable pairing of strengths could plausibly adopt. Australia and Canada would bring research depth and agricultural deployment contexts structurally similar to India's; Singapore an already proven cross-border sandbox model; Japan advanced manufacturing and robotics; the European Union market scale and standards weight; Gulf partners sovereign capital and deployment appetite; and across Africa, ASEAN and the wider Commonwealth sit the emerging markets where India's deployment method has the most to offer and, in areas where Indian infrastructure itself remains imperfect, low-resource language coverage being the clearest example, the most still to learn.

It is important to be precise about what this closing argument does not claim. It is deliberately not a containment strategy directed against any other country's AI ecosystem, and it deliberately avoids labelling the aspiration a fixed banner such as a democratic AI corridor, because countries that are democracies in the conventional political sense continue to diverge, sometimes sharply, on data governance philosophy, on the acceptable balance between sovereignty and openness, and on whether open release of frontier-capable models is a public good or a security risk, a live and unresolved debate inside both the UK and Indian policy establishments themselves. A label asserting an agreement that does not yet exist would be a diplomatic overstatement dressed as strategy. The more modest and, this paper argues, more durable framing is an affirmative, exportable method of AI diffusion built on documented complementary strength, examined here with as much attention to its real limits, Sarvam's unverified benchmarks, the UK's own translation bottleneck, India's data-localisation ambiguities, as to its genuine promise. The UK-India relationship is the proof of concept for a repeatable form. Whether it widens further is a question for its partners to answer, on terms they set for themselves.

## Final considerations

## Next-generation sustainable data centres

Several notable startups and scale-ups in both the UK and India are actively tackling the data centre sustainability crisis. While UK companies tend to leverage waste-heat harvesting and AI-driven grid orchestration, Indian startups focus heavily on tropicalised liquid cooling and low-water thermal design

## United Kingdom: Heat Reuse & Grid-Aware Intelligence

The UK startup ecosystem focuses on hardware-software solutions that re-integrate data centres into energy grids and local urban communities.

### 1. Deep Green

Deep Green builds immersion-cooled micro-data centres that double as community heating units.

High-density AI GPUs are submerged in non-conductive mineral oil. The captured waste heat is pumped directly into public infrastructure—such as municipal swimming pools and district heating schemes—reducing the cooling energy footprint to near-zero while saving local councils tens of thousands of pounds in heating bills. [28]

#### 2. Emerald AI

Developed in collaboration with entities like NVIDIA, National Grid, and Nebius, Emerald AI provides a software platform designed to make AI workloads "grid-flexible".

Training large language models requires massive energy spikes that strain local power grids. Emerald AI uses machine learning algorithms to orchestrate, pause, or dynamically shift non-time-sensitive batch jobs depending on live grid stress and energy prices—transforming energy-guzzling data centres into dispatchable grid assets. [29]

#### 3. Phaidra

Phaidra builds autonomous control systems (AI agents) specifically tailored for complex industrial operations like data centres.

Their AI control agents continuously monitor temperature, flow rates, and GPU thermal limits across liquid-cooling loops and chillers. By dynamically fine-tuning cooling controls in real-time, they lower Power Usage Effectiveness (PUE)[29] and maximise "tokens per watt".

### India: Extreme Heat Management

In India, where ambient temperatures often exceed 40 degrees Celsius and municipal water stress is high, startups are innovating around tropicalised liquid cooling

#### 1. Refroid Technologies

- *Refroid manufactures direct-to-chip liquid cooling and immersion cooling hardware tailored for high-density AI infrastructure in tropical regions.*

> *Traditional air-based cooling fails under high heat and heavy GPU loads. Refroid builds modular liquid-cooling pods (like their Sentraflo Coolant Distribution Units) that handle GPUs pulling more than 500 W per chip, maintaining thermal stability without draining local water supplies through evaporative cooling.*[30]

#### 2. Emerging Deep-Tech & Climate-Tech Innovators (e.g., CoolPact Ecosystem)

- *Supported by climate-focused funds like CoolPact, a cohort of deep-tech startups in India is building specialised software-sensor networks and closed-loop thermal systems.*[31]
- *These startups integrate IoT sensors and machine learning algorithms into data centre MEP (Mechanical, Electrical, and Plumbing) architectures to eliminate water loss and predict thermal throttles before they occur in hot climates.*

The room for collaboration between the UK and India in the data centre space spans hardware co-design, green building standards, and software efficiency. Governed by formal bilateral frameworks—such as the UK-India Technology Security Initiative (TSI) and the India-UK Vision 2035—both nations are positioning themselves to lead frugal, sustainable compute infrastructure.[21]

# Sources

## Frugal AI Hub publications


1. *Frugal AI Hub, 2025. Frugal AI: Technical and Socio-Economic Foundations. Cambridge Judge Business School.*
2. *Frugal AI Hub, 2025. Accelerating a Frugal AI Ecosystem. Cambridge Judge Business School.*
3. *Frugal AI Hub and United Nations International Computing Centre (UNICC), 2025. From Total Cost of Ownership to Social Impact: A Frugal AI Framework to Measure AI Portfolios as Strategic Assets. Cambridge Judge Business School and UNICC.*
4. *Frugal AI Hub, 2026. Applications of Frugal AI in the Fintech Industry: Measuring Total Cost of Ownership and Return on Investment in Real-Time Fraud Detection. Cambridge Judge Business School.*
5. *Frugal AI Hub and Global Education Lab, 2026. Frugal AI and the Executive Agenda. Cambridge Judge Business School.*

FORTHCOMING Gandikota, S.V., forthcoming. Frugal AI. Book in preparation.


## Sources cited in this report

*All figures drawn from publicly available sources as of July 2026. Government-official holders, current figures and policy status may change; readers should verify against the primary source before relying on any specific number. Several performance claims relating to India's sovereign AI models are the companies' own reported figures and are noted in the text as not yet independently verified. The engineered-tree carbon-sequestration multiplier cited in Section 3 is a research proposition still under active field validation.*


6. *Qure.ai and NHS partner reporting: SBRI Healthcare case study on the LungIMPACT study; North East London Cancer Alliance, "Using Artificial Intelligence for faster Chest X-ray results," 2024–2026; Qure.ai, "The impact of AI at The Royal Bolton NHS Hospital, UK" and "AI chest X-ray triage boosts radiology efficiency" (Frimley Health pilot, 99.7% normal-CXR accuracy, up to 58% caseload reduction), 2023–2026; Digital Health News on East Kent Hospitals NHS SBS framework deployment, 2023.*
7. *Oxford Nanopore Technologies, "Indian Research Institutions and Oxford Nanopore commit to collaborate on Genomic Centres of Excellence" (BRIC-CDFD, BRIC-NIBMG Letters of Intent; Bengaluru office; UK-India Technology Security Initiative), April 2025.*
8. *Tony Blair Institute for Global Change, "Sovereignty in the Age of AI: Strategic Choices, Structural Dependencies and the Long Game Ahead," January 2026; Tech Policy Press, "Rethinking Sovereign AI as Strategy," March 2026; CipherTalk, "Sovereign Compute," June 2026; SPMRF, "Owning the Stack, Not Just the Model," June 2026; Forbes, "Sovereign AI: Why Owning The Full Stack Is The New Strategic Imperative," April 2026; Leonardo Borges, "Sovereignty Is an Inference Problem," July 2026; Stanford HAI, "AI Sovereignty's Definitional Dilemma," 2026.*
9. *UK Government, "A decisive shift to power British AI: new £1.1 billion plan" and "UK AI Hardware Plan," GOV.UK, 8 June 2026; five AI Growth Zones and £28.2bn private investment figures per Tech Nation / DSIT reporting, 2026. DSIT was dissolved on 21 July 2026 and its functions transferred to DBIST, the Cabinet Office and DCMS; see source 32.*
10. *UK Biobank, "Whole genome sequencing data" release notes (500,000 participants; Wellcome Sanger Institute), 2023–2025; Science, "UK Biobank releases half a million whole-genome sequences," 2023.*
11. *IndiaAI Mission and MeitY materials on AIKosh and IndiaAI Mission 2.0, 2025–2026; India AI Impact Summit 2026 materials.*
12. *ClearIAS and UK Parliament (Business & Trade Committee) analyses of the India-UK CETA and the Double Contribution Convention, 2025.*
13. *AIFundingTracker, GrowthList, TechRound, Waveup and EU Tech Future reporting on London/UK AI venture capital rankings and scale, 2025–2026 (London ranks third among global cities for AI VC behind the San Francisco Bay Area and New York; UK ranks fourth globally among countries/regions behind the US, EU27 and China).*
14. *UK Government, "UK-India CETA Chapter 12: Digital Trade," GOV.UK, July 2025; Trade Notes (Substack) analysis of CETA Article 12.20 forward-review mechanism, August 2025.*
15. *BW Healthcare World on Niramai; WEF/industry analysis of Qure.ai qXR detection rates; ORF, "From Scale to Efficiency: India's Frugal AI Opportunity," March 2026.*
16. *howays.in / Google.org reporting on Wadhwani AI CottonAce ($3.3m grant; ~25% pesticide-cost reduction; ~20% income uplift), October 2025.*
17. *Reporting on CropIn SmartFarm groundnut yield uplift (~30%), Karnataka, 2025; SoInce Energy reporting on the Gujarat AI solar-forecasting pilot, 2025.*

18. *Forbes India, "Mitesh Khapra's AI4Bharat is closing the AI gap between English and Indian languages," June 2025; IMPRI, "Digital India Bhashini," 2025; AI4Bharat technical documentation (Samanantar, IndicTrans2), 2025–2026.*

19. *Rest of World, "India's frugal AI startups Sarvam and Krutrim build sovereign models," April 2026; explainx.ai, "India Sovereign AI Status 2026," 2026; Forbes, "India Can Train A Sovereign Model But Still Cannot Prove It Works," March 2026; Forbes India, "Five new sovereign AI models signal India's bold leap at AI Summit," February 2026; Inc42, "India's Sovereign AI Reality Check," May 2026.*

20. *Lexology, IAPP, Vidhi Legal Policy Centre, and American Bar Association analyses of India's Digital Personal Data Protection Act 2023 and DPDP Rules 2025, cross-border transfer and sectoral localisation provisions, 2025–2026.*

21. *digit.fyi and Technology Networks on UK engineering biology sector scale and pilot-to-commercial bottleneck, 2025–2026; India-UK Joint Statement, October 2025 (India-UK Joint Centre for AI; health, climate, fintech, engineering biology).*

22. *Institute of Development Studies, "Comparison of the UK and India's Research and Innovation Strengths," December 2025; UK-India FTA and Life Sciences sector analyses, BioPharma APAC, July 2025.*

23. *Isomorphic Labs, "Isomorphic Labs announces $600m external investment round," March 2025; Reuters and TechCrunch coverage of the round; revised clinical-trial timeline to end-2026 announced at the World Economic Forum, January 2026.*

24. *UK Government, National Quantum Strategy (DSIT) and "Quantum technologies" sector pages, business.gov.uk, 2023–2026. At least 160 companies were active in the UK quantum sector at the time of the National Quantum Strategy; the current government figure is over 180.*

25. *Wayve, "AV2.0" technology documentation and "A new approach to self-driving: AV2.0," 2023–2026; Sequoia Capital interview with Alex Kendall on end-to-end learning, November 2025.*

26. *Microsoft, "With help from next-generation AI, Indian villagers gain easier access to government services," Source Asia, 2023; TechCrunch reporting on Jugalbandi (171 of approximately 20,000 government programmes; ten Indian languages at launch).*

27. *India Semiconductor Mission and MeitY factsheets; Press Information Bureau and Prime Minister's Office releases on ISM project approvals, 2025–2026 (twelve approved projects, cumulative investment around Rs 1.64 lakh crore; ISM 2.0 announced in Union Budget 2026-27).*

28. *Deep Green and Octopus Energy Generation announcements on immersion-cooled edge data centres and heat reuse; Data Centre Dynamics and CIBSE Journal reporting on the Exmouth Leisure Centre deployment, 2023–2026.*

29. *NVIDIA, "Omniverse DSX Blueprint" and "How AI Factories Can Help Relieve Grid Stress," 2025–2026; Phaidra, "AI agents for gigawatt-scale AI factories," November 2025; Emerald AI and Nebius London grid-flexibility trial, December 2025.*

30. *Refroid Technologies, "SentraFlo Series" launch release, May 2025, and Refroid–Technavious partnership release, March 2026; Data Center Dynamics and Network World coverage.*

31. *CoolPact Capital and the Global Innovation Lab for Climate Finance, "CoolPact Capital India Fund," 2024–2026; ImpactAlpha reporting on India's cooling-technology startup base.*

32. *UK Government, written ministerial statement on machinery of government changes, Baroness Smith of Basildon, House of Lords, 21 July 2026; Bloomberg, UKAuthority and Computer Weekly reporting on the dissolution of the Department for Science, Innovation and Technology, the creation of the Department for Business, Innovation, Science and Trade and the appointment of the first Cabinet-attending Minister of State for Artificial Intelligence, 20–21 July 2026.*


## Sources for Appendix A


33. *KPMG International / University of Queensland, "Trust in Artificial Intelligence: A Global Study," 2025–2026 (India 71% general trust in AI; most advanced Western economies below 50%).*

34. *Pew Research Center, "Views of AI Around the World," October 2025 (89% of Indian adults trust their own country to regulate AI).*

35. *Stanford HAI, "2026 AI Index Report," Public Opinion chapter, 2026.*

36. *CSET Georgetown and Statista/UNESCO Institute for Statistics on global STEM graduate distribution, 2023–2025; StudyIQ, "India's STEM Future," December 2025 (2.5–2.55 million STEM graduates/year; ~260 researchers per million population; R&D spend ~0.67% of GDP); Airswift, "Top Countries for STEM Talent," 2025–2026 (~1.5 million engineering graduates/year).*

37. *UK Parliament POST, "UK STEM Skills Pipeline," 2025 (9.4 million people in UK STEM-related employment, 2023).*

38. *STFC / UKRI, "Particle Physics" and "CERN and the UK," 2025–2026 (20+ UK universities across all four LHC experiments: ATLAS, CMS, LHCb, ALICE).*

A

# Appendix A · Three Views of Complementary Strength

## The UK–India AI Capability Radar

The functional radar in Section 2 (Table 2, Figure 1) makes one comparative argument: that the UK and India occupy near-opposite positions across eight AI functions. This appendix extends that approach across two further dimensions sectoral application strength and innovation/capital architecture and concludes by synthesising what the three perspectives imply for where the UK–India AI corridor should be built. The radar charts are non-numerical conceptual visualisations, intended to show the relative areas of comparative strength rather than precise measurements. They are illustrative, directional representations informed by the evidence assembled throughout this paper and should not be interpreted as weighted indices, rankings or quantitative scores.

### RADAR A1 · AI FUNCTIONAL AND INFRASTRUCTURE PILLARS

#### Where the Underlying Capacity Sits

This radar deliberately sets aside sector applications, health, agriculture, finance, and looks only at the raw infrastructure and institutional pillars any AI system needs to be built and trusted: compute capacity, data storage and centre infrastructure, data residency and sovereignty control, application-layer maturity, the presence of a national tech stack, dataset availability, public trust, and the scale of the STEM and engineering talent pool behind it all.

The UK leads on the pillars that take years of capital and institutional patience to build: compute (its AI Growth Zones and Research Resource),[9] data residency and sovereignty control (a mature regulatory regime built on decades of financial and data-protection law), and dataset quality, if not volume, through assets like the UK Biobank.[10] India leads just as clearly on the pillars that take population scale and necessity to build: application-layer maturity and national tech-stack presence, where India Stack has no real UK equivalent, dataset availability at population scale, and, strikingly, public trust. Independent global studies converge on this: India records among the highest public trust in AI of any country surveyed (KPMG and University of Queensland found 71%,[33] and a separate Pew study found 89% of Indians trust their own country to regulate AI),[34] while the UK sits with most other wealthy Western economies in a considerably more sceptical band, well under half in several comparable surveys.[35]

The STEM talent axis deserves a more careful reading than a single number allows. India produces roughly 2.5 million STEM graduates a year, and around 1.5 million engineering graduates specifically,[36] a volume no other country except China approaches. But that scale sits alongside a real and separate constraint: only around 260 researchers per million population, R&D spending still near 0.67% of GDP,[36] and a widely reported pattern of top AI PhDs leaving the country. The UK's STEM base is smaller in absolute terms but converts a larger share of it into sustained research careers.[37] This is the same complementary pattern the whole paper describes, expressed now at the level of raw national capacity rather than applied sectors: India's scale advantage is real, the UK's conversion advantage is real, and neither substitutes for the other.

**Figure A1 · UK and India across eight infrastructure pillars**

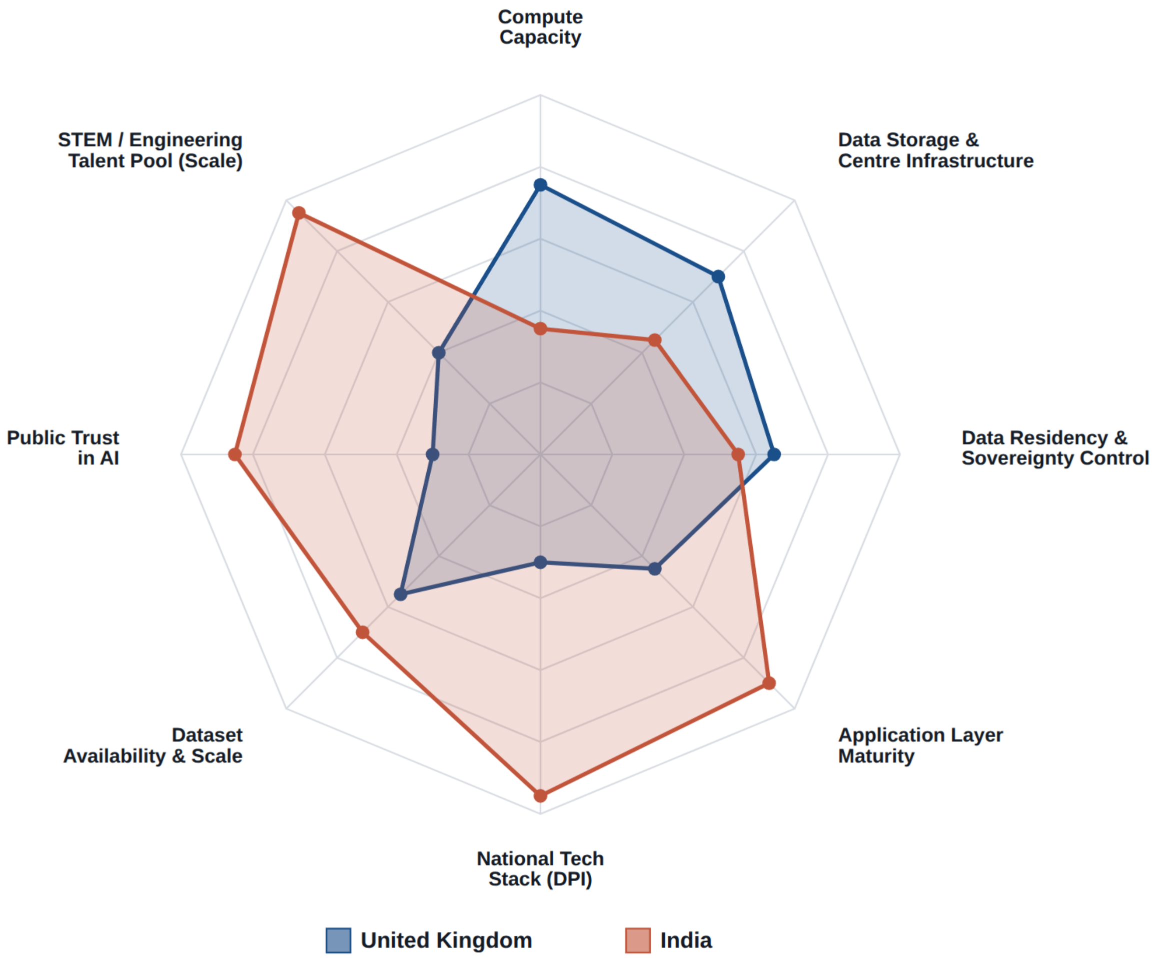


Illustrative positioning across eight infrastructure pillars, drawn from the sources listed in the Sources section. Not a precise scored index.

## RADAR A2 · SECTORAL APPLICATION STRENGTH

### Frontier Science Versus Subcritical, Small-Model Service Delivery

This radar tests the frontier-minus-N argument from Section 2 directly, by sector rather than by infrastructure layer. Four of the eight axes are sectors where genuine frontier capability, deep compute, decades of accumulated fundamental science, large research teams, is still the gating factor: fundamental and particle physics, life-sciences and genomic discovery, climate and earth-system supercomputing, and advanced industrial and manufacturing precision AI. The other four are sectors that this paper's central argument says do not need frontier capability at all, only a well-built, subcritical, often small model, deployed close to the user: agri-tech real-time advisory, language translation and multilingual services, public-service delivery efficiency, and financial fraud or economic-crime detection at population scale.

The UK's profile is concentrated almost entirely in the first group. Its participation in fundamental physics runs through the Science and Technology Facilities Council, which funds UK researchers across all four major Large Hadron Collider experiments at CERN, alongside more than twenty UK universities and research institutes;[38] its life-sciences strength runs through the UK Biobank and the AlphaFold lineage discussed throughout this paper;[10] and its climate and industrial modelling draws on the same physics-informed AI capability. India's profile is the near-mirror image, concentrated in the second group: its agri-tech advisory and multilingual services strength was documented at length in Section 3, and its public-service delivery efficiency runs directly through Aadhaar- and UPI-scale digital public infrastructure.

The one axis that does not split cleanly, and is worth flagging rather than smoothing over, is financial fraud and economic-crime detection. The UK's strength here is institutional and wholesale, graph-neural-network-based anti-money-laundering tools built for complex, high-value, cross-border transaction webs. India's strength is retail and population-scale, fraud detection built to run across UPI's enormous real-time consumer transaction volume. Both are genuinely strong; they are strong at different points on the transaction-value spectrum, which is precisely the nuance Section 3 flagged when it warned against treating UPI's retail scale as evidence of institutional-fintech readiness.

**Figure A2 · UK and India across eight sectoral applications**

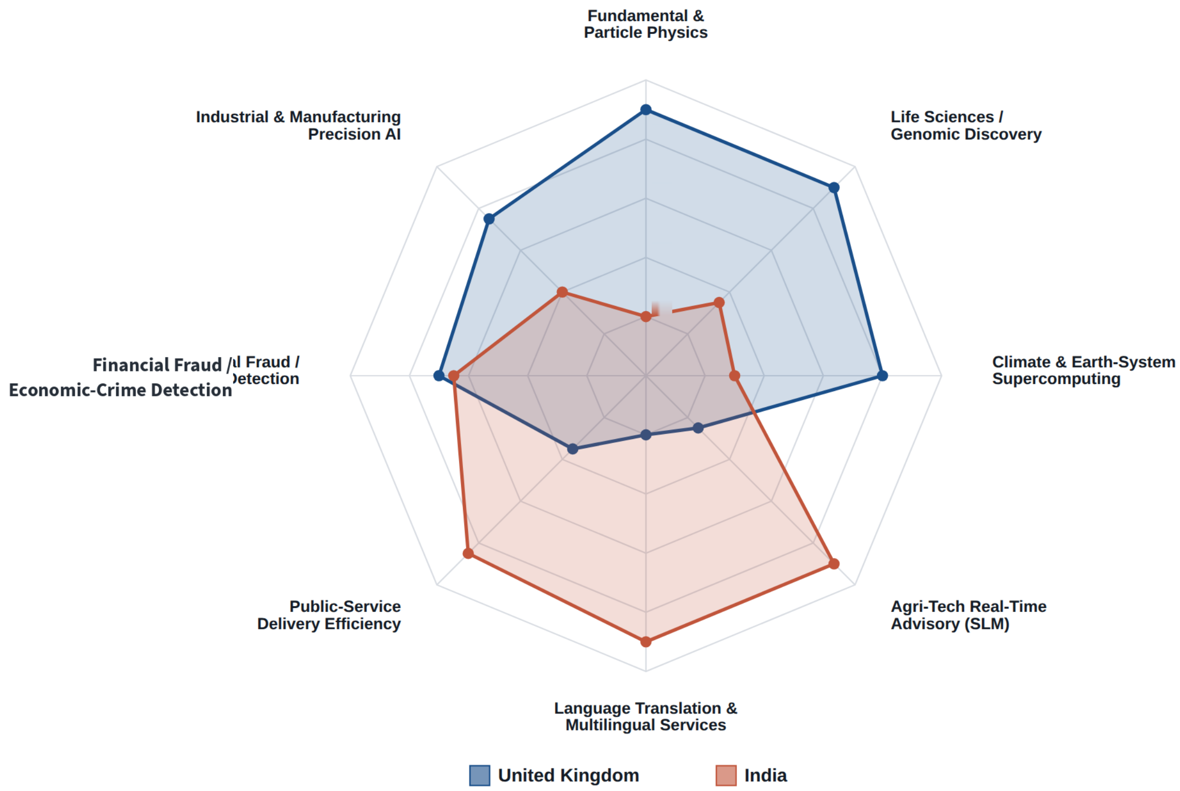


> Illustrative sectoral positioning: frontier-dependent sectors (physics, genomics, climate, industrial AI) versus subcritical/small-model sectors (agri-advisory, translation, public service, retail fraud detection).

## RADAR A3 · INNOVATION AND CAPITAL COMMERCIALISATION ARCHITECTURE

### A Third Pillar: Who Turns a Result Into a Company

The first two radars describe capacity and application. This third addresses a question the main paper raises repeatedly, in Section 2's account of the UK's own translation problem and in Section 4's startup case studies, but had not yet placed on its own chart: once a result exists, which country's architecture actually turns it into a company, a product, or a deployed system, faster and more cheaply. Eight axes: venture capital depth specific to AI, patent and IP output, the rate of unicorn and scale-up formation, public R&D investment intensity as a share of GDP, time-to-commercialisation, founder and talent retention, regulatory sandbox maturity, and the capital efficiency of deployment, how cheaply a working system can actually be built and run.

The UK leads on the institutional and capital-markets side of this pillar: London's venture capital depth, patent output,[13] and, notably, the Financial Conduct Authority's sandbox model, which pioneered the regulatory-testing approach this paper has repeatedly argued India's own regulators could usefully mirror. But the UK is positioned markedly lower on time-to-commercialisation, which reflects the translation problem Section 2 documented in detail: UK engineering-biology and deep-tech ventures routinely have the capital and the science but still take up to eight years to reach revenue,[21] precisely because deployment infrastructure is the bottleneck.

India's profile inverts this. Its public R&D intensity remains low (still near 0.67% of GDP, against considerably higher levels in China, South Korea and Israel),[36] and founder and talent retention is a genuine, openly acknowledged concern, a widely reported majority of India-trained AI PhDs move abroad. But India is positioned most strongly on the one axis that arguably matters most for translating the UK's own stalled science into a product: capital efficiency of deployment. The frugal engineering discipline detailed in Section 2 is, read through this radar, a capital-architecture advantage as much as a technical one, and it is the axis where India's profile most directly completes the gap in the UK's own, exactly as Section 4's Qure.ai and Oxford Nanopore cases already demonstrate in practice.

**Figure A3 · UK and India across eight innovation and capital dimensions**

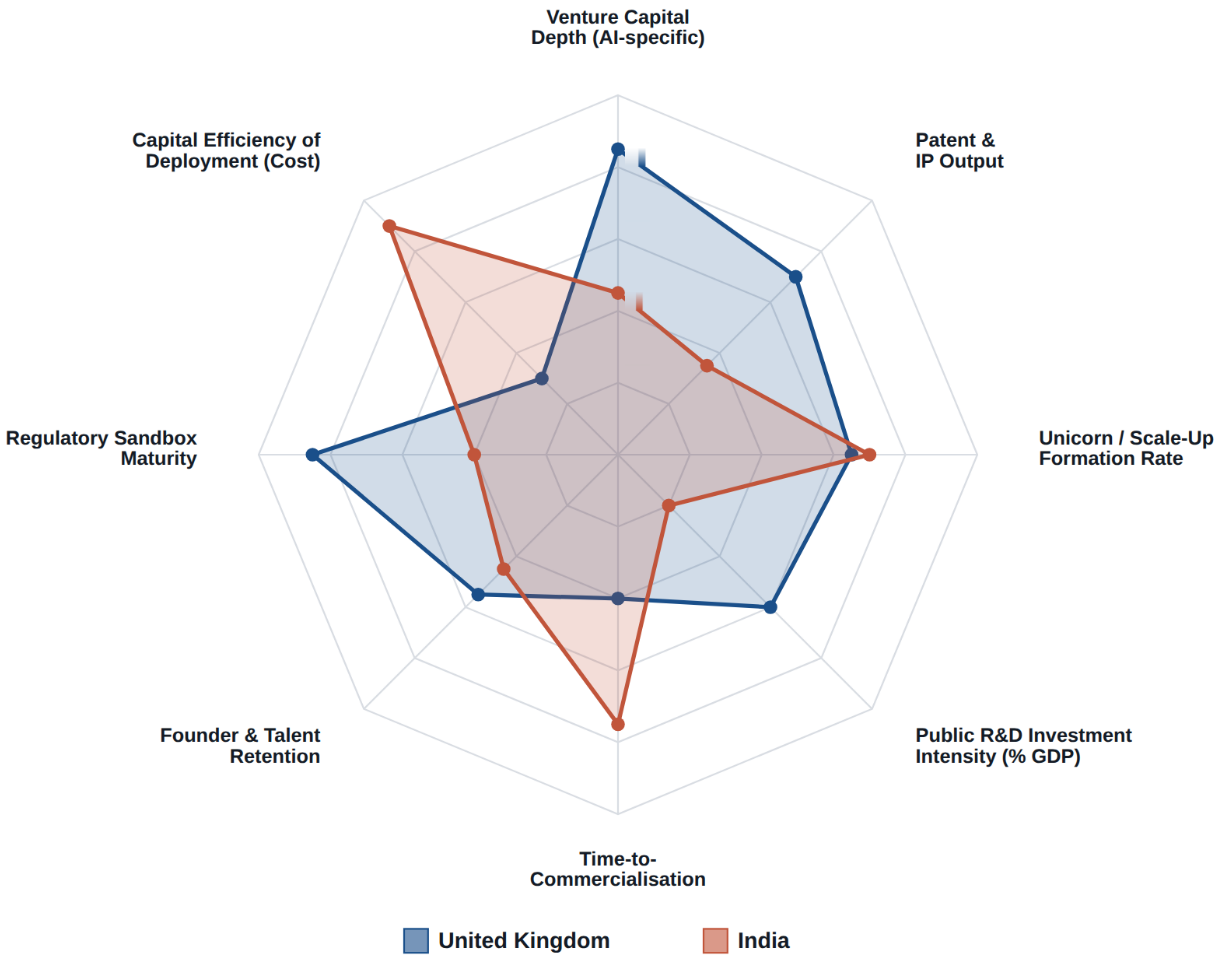

Illustrative positioning across eight innovation and capital-commercialisation dimensions.

## SYNTHESIS · READING THE THREE RADARS TOGETHER

### Where the Three Shapes Actually Overlap

Each radar on its own makes a point about capacity, application, or capital. Read together, they converge on the same handful of realisable intersections Section 3's macro-sector matrix already identifies, and this page exists to make that convergence explicit rather than leaving three separate charts to speak for themselves.

The pattern across all three radars is consistent: wherever the UK's shape peaks, in compute, data residency, fundamental science, venture capital and patents, India's shape is at its narrowest, and wherever India's shape peaks, in application maturity, national tech stack, agri-advisory and translation, capital efficiency, the UK's shape is at its narrowest. This is the specific, actionable finding. An intersection is realisable exactly where one country's peak sits opposite the other country's trough, because that is where a joint venture supplies what neither could build alone.

| Radar | UK peak / India trough | India peak / UK trough |
|---|---|---|
| **A1. Infrastructure** | Compute capacity; data residency and sovereignty control | National tech stack (DPI); application-layer maturity; public trust |
| **A2. Sectoral** | Fundamental physics; genomic discovery; climate supercomputing | Agri-tech advisory; multilingual services; public-service efficiency |
| **A3. Capital & innovation** | Venture capital depth; patent output; sandbox maturity | Capital efficiency of deployment; commercialisation speed once translated |

Reading the three radars together, the realisable intersections are exactly the five verticals Section 3 develops in full: life sciences pairs UK genomic discovery with India's trial and manufacturing scale, and is already visible in practice in Section 4's Oxford Nanopore case; climate tech pairs UK synthetic-biology science with India's tissue-culture and agricultural deployment infrastructure; deep mobility pairs UK reinforcement-learning navigation with India's quantisation and edge-compilation discipline; institutional fintech pairs the UK's sandbox model and compliance AI with India's transaction-scale data, once the Reserve Bank of India's sandbox is liberalised; and semiconductors pair UK chip design lineage with India's fabrication scale, once its engineering curricula shift toward micro-architecture co-design. These three radars do not add a sixth intersection to that list; they confirm, from three independent angles, capacity, sector, and capital, that the same five constitute this paper's principal opportunity set, although they differ substantially in readiness.

One honest caveat carries across all three charts and this synthesis: every position here is an illustrative, directional characterisation drawn from the qualitative and quantitative evidence assembled in this report and its companion research; no weighted index underlies it. Where a specific figure is cited in the text above, public trust percentages, STEM graduate counts, R&D intensity, it is sourced and can be checked; the radar positions themselves are a reading of that evidence, carry no scale and no scoring methodology, and should be presented to any external audience with that distinction intact.

**Frugal AI Hub**

at Cambridge Judge Business School

Trumpington Street
Cambridge CB2 1AG
United Kingdom

https://frugalai.org

**Contact**

Elizabeth Osta e.osta@jbs.cam.ac.uk
Serish Venkata Gandikota s.gandikota@jbs.cam.ac.uk
Arjuna Sathiaseelan a.sathiaseelan@jbs.cam.ac.uk
Jaideep Prabhu j.prabhu@jbs.cam.ac.uk

FRUGAL AI HUB at UNIVERSITY OF CAMBRIDGE Judge Business School

Centre for **India & Global Business**
UNIVERSITY OF CAMBRIDGE Judge Business School