

Artificial Intelligence as a Utility in India by 2047

The promise and limits of the grid analogy

Devashish Pawar

4 October 2026 | Conceptual essay | Version 1.0 | CC BY 4.0

DOI: <https://doi.org/10.5281/zenodo.23130931>

Abstract

Artificial intelligence may become an ordinary input into economic and civic life, accessed through services as routinely as electricity. This essay examines what that possibility means for India as it approaches the centenary of independence in 2047. Building on the argument of *The Grid That Thinks*, it uses the history of electrification and selected research on productivity, AI costs and infrastructure to assess the utility analogy. The central argument is that affordable computational capability produces broad social value only when institutions reorganise around it and access includes language, reliability, accountability and effective redress. The analogy illuminates complementary investment and unequal diffusion, but weakens when applied to the measurement of intelligence, market structure and systems that interpret information or act on users' behalf. The essay proposes a plural infrastructure approach combining shared public capabilities, competing providers, local evaluation and human fallback. It treats 2047 as a planning horizon rather than a forecast and identifies observable outcomes by which the promise of an intelligence utility could be judged.

Keywords: artificial intelligence; India; digital public infrastructure; productivity; linguistic inclusion; AI governance; technological sovereignty; 2047

1 The public significance of ordinary technology

The most consequential technologies often become familiar enough to escape attention. A household switches on a fan; a farmer starts a pump; a shopkeeper turns on a refrigerator. Electricity is essential to each activity, yet the immediate concern is the service it enables. Artificial intelligence could follow a comparable path. People may increasingly encounter it within applications and institutions, without choosing to consult a separate system labelled AI. Whether this makes life better will depend on the arrangements behind that apparent simplicity.

The Grid That Thinks, the essay from which this discussion develops, places that question in India's future. Its productive move is to shift attention from spectacular demonstrations of machine capability to the terms on which ordinary people might use it. This essay extends that argument: AI should be assessed as a possible infrastructure for access to useful expertise, while recognising that automated output is not equivalent to expertise itself. A fluent response can be inexpensive to generate and expensive to trust.

The year 2047, India's centenary of independence, provides a horizon for institutional choices whose consequences unfold over decades. It does not establish a date by which AI will become universally available or produce a measurable economic transformation. The approach here is interpretive and prospective, drawing on selected historical, empirical and policy sources. It offers neither an original empirical study nor a systematic literature review. Examples of future services are illustrative scenarios, and the proposed governance arrangements are arguments for consideration rather than findings from an implemented programme.

2 What an intelligence utility would actually supply

An intelligence utility can be understood as widely available computational assistance delivered through dependable services. Its functions might include translation, document preparation, search, classification and support for decisions. This definition does not require human-level general intelligence or imply that a model possesses consciousness. It describes the organisation of access to capabilities, rather than a claim about the nature of thought.

There is evidence that some capabilities have become much cheaper to obtain. Stanford's 2025 AI Index reports that the inference price of systems reaching approximately GPT-3.5-level performance on the MMLU benchmark fell by more than 280 times between November 2022 and October 2024 [1]. That comparison supports a limited claim about the price of a specified benchmark performance over a specified period. It does not establish an enduring rate of decline, equivalent competence across tasks or the approaching disappearance of service costs.

The distinction matters because tokens are units of processing and billing, not units of verified understanding. Two models may generate similarly long answers with different accuracy, while the same model may perform differently across languages or settings. An electricity customer buys a relatively standardised physical quantity. An AI customer buys an output whose usefulness depends on the task, the evidence available and the consequences of error. A meaningful service promise must therefore specify what the system is expected to do and how its performance will be checked.

Consider a hypothetical farmer in Bundelkhand seeking advice about a damaged crop. A useful service would need more than a photograph and a fluent response in Bundeli. It would have to handle uncertain symptoms, local conditions and missing information, and refer difficult cases to qualified assistance. The cost of the complete service would include maintaining relevant information, checking advice and supporting escalation. As the price of generation falls, these surrounding responsibilities may account for a larger share of the value and cost. Universal access should consequently mean access to dependable help, rather than merely access to generated text.

3 Productivity depends on institutional change

The history of electrification explains why a powerful technology can spread before its full benefits become visible. Paul David's account of the dynamo and the computer emphasises the complementary changes through which electricity reshaped production [2]. Substituting an electric motor into a factory organised around an earlier power system was different from redesigning the factory to exploit distributed electric drive. The broader lesson concerns the interaction between technology, capital and organisation. It does not support a simple claim that all productivity remained unchanged for forty years.

Brynjolfsson, Rock and Syverson develop a related explanation in their work on the productivity J-curve [3]. General-purpose technologies require investments in new processes, skills and other intangible assets. Some of that investment is poorly captured by conventional measurement, which can obscure benefits during the transition. This is a reason to examine implementation and measurement carefully. It is not a guarantee that every organisation investing in AI will eventually become more productive.

For an Indian district office, adding a drafting assistant to an existing process may save time at one desk while leaving the citizen's experience unchanged. If files still wait for unavailable records or unclear approvals, faster drafting cannot remove the main delay. More substantial improvement might require simpler procedures, accessible records and a clear division of responsibility. This is a hypothetical

institutional example, but it shows why the relevant unit of assessment is often the completed service rather than the isolated task.

Employment outcomes also depend on those choices. Time saved could support more attentive service and skill development, or it could become a reason to reduce staffing and increase workloads. Neither outcome follows automatically from model capability. Institutions should evaluate changes in service quality, work intensity, wages and learning opportunities alongside throughput. A pilot that reduces processing time while making errors harder to contest would offer an incomplete account of productivity.

The same reasoning cautions against using electrification as a clock. Software distribution can be rapid, while procurement, professional practice and public trust may change slowly. Different sectors will move at different speeds. The defensible claim is that complementary investment matters; the timing and distribution of its returns remain empirical questions.

4 Lower prices do not settle who captures value

If routine outputs become cheaper, reliable interpretation and responsibility may become more valuable. A professional who checks an answer, understands its setting and accepts responsibility supplies something distinct from the output itself. Yet this does not mean that AI developers must become low-margin utilities or that most economic value must migrate to application providers. Utility-like use is a pattern of consumption, not a prediction about ownership or profitability.

Several competing mechanisms are plausible. Widely available models and interoperable services could make it easier for smaller firms to enter a market. Conversely, a supplier that combines models with cloud infrastructure, exclusive data or control over customer access could make switching difficult. A low introductory price could coexist with substantial dependence. These are possible market dynamics, not predictions that one structure will inevitably prevail. They qualify the original essay's suggestion that the history of electricity supplies a settled answer about where value goes.

For India, the relevant question is whether enterprises and public institutions can obtain useful capability on terms they can sustain and challenge. Can they move their records and workflows to another provider? Can they compare outcomes independently? Can a small supplier compete without rebuilding every underlying component? Public procurement could help by requiring usable data exports, documented interfaces and transition arrangements. Such measures would need testing and enforcement; calling a platform open would not establish that users can leave it.

Trust must also be earned through arrangements that survive commercial pressure. A provider's reputation is useful, but a service becomes more accountable when failures can be documented and corrected by someone with the authority to act. Responsibility should attach to identifiable organisations throughout deployment. A model's inability to accept moral responsibility does not remove the responsibility of the people and institutions that commission, operate or rely on it.

5 Inclusion begins after connection

Electrification offers a particularly important warning about how access is measured. In an April 2018 clarification, India's Ministry of Power explained that the village-electrification definition included a requirement to electrify at least ten per cent of households. It also stated that this was not a ceiling and that actual household coverage was considerably higher [4]. The distinction remains instructive: a classification applied to a village does not establish that every resident receives an adequate service.

The equivalent error for AI would be to count accounts, downloads or connected villages as proof of inclusion. A person can have a phone and still be unable to obtain useful assistance because of language, device limitations, connectivity, disability or lack of confidence. Access can also fail when the system answers confidently but incorrectly, or when there is no practical way to challenge a consequential outcome. These barriers call for an assessment of effective use rather than nominal availability.

Voice interfaces and translation could help, but performance should be tested where people actually use them. An evaluation should include regional speech, dialects, code-switching, background noise and the vocabulary of local occupations. Translation between widely used languages would not by itself represent every community. The original essay's Bundeli-speaking farmer is therefore a valuable test case: the interface must respond to the user's speech and context, rather than require the user to approximate a standard language.

Inclusion also requires ways to ask for help without surrendering unnecessary information. A service supporting access to public benefits should have a clearly explained purpose for collecting data and an assisted route for people who cannot use the automated channel. Public agencies should retain responsibility for resolving errors. Otherwise, a system intended to reduce the cost of expertise may transfer the work of correcting mistakes to people least able to bear it.

The appropriate ambition is not simply a free interface. It is affordable access to a service whose quality can be assessed, whose limitations are communicated and whose failures can be repaired. Subsidies may support that ambition, but they need an explicit account of what is being subsidised and who remains responsible when the service fails.

6 Public infrastructure can support a plural ecosystem

The source essay draws on India's Unified Payments Interface as a model for shared infrastructure. NPCI describes UPI as an instant payment system that it developed as an RBI-regulated entity [5]. Its relevance lies in the possibility of a common system supporting different user-facing services. It should not be described as an ownerless network, and its success does not prove that AI can be governed through an identical institutional design.

A payment system can specify transaction states with comparatively clear rules. An AI service producing advice must also address uncertain meaning and contested judgments. The useful inference from UPI is consequently narrower: shared interfaces and common operational expectations may allow providers to compete without forcing every user into an isolated system. AI still requires task-specific evaluation and explicit responsibility for the content and actions delivered through those interfaces.

India's 2024 approval of the IndiaAI Mission illustrates that shared capability is already a policy concern. The announced programme included computing access, data quality, indigenous capabilities, skills and safe and trusted AI, supported by an approved outlay of Rs 10,371.92 crore [6]. This establishes the programme's announced scope and allocation, not proof of its eventual effectiveness. Evaluation must distinguish resources committed from capabilities delivered and benefits realised.

A workable approach need not choose between one national model and unrestricted dependence on a few foreign providers. Public investment could support computing access, lawful and well-documented datasets, language resources and independent evaluation, while universities, public bodies and private firms develop different applications. Where disclosure is lawful and appropriate, publishing evaluation methods and results would help users compare services. Some applications may be better served by smaller specialised models or simpler software; scale alone is not a public-interest objective.

Domestic ownership also does not guarantee accountability. Concentration within one country can restrict choice and expose users to opaque decisions just as concentration abroad can. A stronger conception of technological sovereignty is the practical ability to choose, inspect, replace and maintain critical services under accountable institutions. International collaboration is compatible with that aim when dependence is understood and alternatives remain usable.

7 The intelligence utility depends on physical systems

Digital services remain dependent on electricity, hardware and communications. In its 2025 Energy and AI report, the International Energy Agency estimated that data centres consumed approximately 415 terawatt-hours of electricity in 2024 and projected consumption of around 945 terawatt-hours in 2030 [7]. These figures cover data centres generally, with AI identified as a major driver of growth. The projection is a dated scenario, not an observed outcome or a measure of AI's electricity use alone.

The policy implication is that computing plans should be assessed alongside power availability and the demands of other users. A centre's effects on its locality may matter even when its share of national demand looks small. Decisions should consider water requirements, reliable supply and the allocation of infrastructure costs, rather than treating installed computing capacity as an unqualified benefit. The electricity and intelligence systems would remain distinct, but their planning would become more closely connected.

Resource constraints also strengthen the case for matching a system to a task. A larger model need not be used where a smaller one or a conventional database performs adequately. Service-level evaluation can make that comparison possible by asking how much energy and expenditure are required for a successful, checked outcome. The objective is useful capability with a defensible resource burden.

8 A utility that interprets and acts needs different safeguards

The analogy with electricity becomes weakest when it treats generated answers as interchangeable supplies. AI systems select, summarise and frame information. Their outputs reflect training material and design choices, and they may present errors persuasively. NIST's Generative AI Profile identifies risks including confabulation and harmful bias, providing a framework for considering such failures across the system lifecycle [8]. Widespread use can make these problems institutional rather than merely individual.

This need not imply that a model has intentions or a unified point of view. The original title's suggestion of a grid that thinks works as a metaphor, but it should not become a technical assertion. The important issue is that outputs influence what users believe and do. An incorrect interpretation repeatedly embedded in public services could become difficult to detect precisely because it appears routine.

The stakes change again when a system can submit forms, transfer information or purchase goods. Competence at answering a question does not establish authority to act. A service should distinguish advice from execution, set boundaries on delegated authority and require appropriate confirmation before consequential actions. It should also record what was done in a form that affected people can understand. These are proposed design requirements, not claims that current systems consistently satisfy them.

Dependence creates a further obligation. If an institution reorganises around a remote service, an outage, contractual dispute or loss of access could disrupt its work. Continuity requires usable records, tested alternatives and staff who can perform essential tasks when automation is unavailable. Human fallback has to be practised; a procedure that exists only on paper offers little resilience. Education and professional training

should therefore preserve the ability to assess evidence and recognise errors while teaching people to use the new tools effectively.

9 A credible test for India in 2047

An intelligence utility should ultimately be judged by what people can accomplish through it. For a public office, the test might include the time and total cost required to resolve a citizen's case, together with the rate of consequential errors and successful appeals. For a language service, it might include task completion across dialects and the availability of competent human assistance. For an institution, it might include its ability to move providers and continue essential work during an interruption. These measures would need to be designed for each setting and reported separately for groups whose experiences differ.

Such evaluation would make the utility proposition open to challenge. If apparent savings disappear after verification and repair, the service is less economical than its generation price suggests. If a system works well only for already advantaged users, it has not delivered universal access. If organisations cannot replace it without losing core capabilities, widespread adoption has created dependence as well as value. Evidence of those outcomes should change deployment decisions.

The promise of the grid analogy is that it directs attention to a shared foundation for ordinary life. Its danger is that familiarity can conceal the decisions through which that foundation is built. India's opportunity is to make useful computational assistance widely available while preserving the ability to question it, correct it and work without it when necessary. By 2047, success would mean that dependable assistance has become ordinary and that the institutions responsible for it remain visible.

Acknowledgement and provenance

This essay develops the ideas in *The Grid That Thinks: Intelligence as a utility, and the world it may build by 2047*, credited to @Devashish and dated 4 October 2026 in the source manuscript. AI assistance was used for research support, drafting and editing. The present essay adds references, qualifies historical and prospective claims, and extends the discussion of public infrastructure and accountability. No original dataset was generated or analysed.

License

© 2026 Devashish Pawar. This work is licensed under the Creative Commons Attribution 4.0 International License (CC BY 4.0). <https://creativecommons.org/licenses/by/4.0/>

References

- [1] Stanford Institute for Human-Centered Artificial Intelligence. (2025). The 2025 AI Index Report. <https://hai.stanford.edu/ai-index/2025-ai-index-report>
- [2] David, P. A. (1990). The dynamo and the computer: An historical perspective on the modern productivity paradox. *American Economic Review*, 80(2), 355-361. <https://www.jstor.org/stable/2006600>
- [3] Brynjolfsson, E., Rock, D., & Syverson, C. (2018; revised 2020). The productivity J-curve: How intangibles complement general purpose technologies. NBER Working Paper 25148. <https://doi.org/10.3386/w25148>
- [4] Ministry of Power, Government of India. (2018, April 30). Clarification on designation of villages as 'electrified'. Press Information Bureau. <https://www.pib.gov.in/newsite/PrintRelease.aspx?relid=179022>
- [5] National Payments Corporation of India. (n.d.). UPI frequently asked questions. Accessed 4 October 2026. <https://www.npci.org.in/what-we-do/upi/faqs>
- [6] Ministry of Electronics & IT, Government of India. (2024, March 7). Cabinet approves ambitious IndiaAI Mission to strengthen the AI innovation ecosystem. Press Information Bureau. <https://www.pib.gov.in/Pressreleaseshare.aspx?PRID=2012357>
- [7] International Energy Agency. (2025). Energy and AI. <https://www.iea.org/reports/energy-and-ai>
- [8] Autio, C., Schwartz, R., Dunietz, J., Jain, S., Stanley, M., Tabassi, E., Hall, P., & Roberts, K. (2024). Artificial intelligence risk management framework: Generative artificial intelligence profile. NIST AI 600-1. <https://doi.org/10.6028/NIST.AI.600-1>