ITMAN INTELLIGENCE · INDEPENDENT RESEARCH

The Adaptation Differential:

Engineering Society for Rapid AI Capability Change

A systems perspective on technological acceleration, institutional response and practical adaptation capacity

Thabang Mokoena

Independent researcher · ITMan Intelligence

ORCID iD: 0000-0003-0689-7632

Status	Conceptual working paper	**Version**	0.1
Date	14 August 2026	**Review**	Not peer reviewed

Abstract

This conceptual working paper examines the possibility that artificial-intelligence capability can advance faster than societies can absorb it safely and productively. It treats exponential growth as a scenario rather than a universal empirical law, then defines an adaptation differential between a moving technical frontier and domain-specific readiness across institutions, work, education, governance, culture and infrastructure. From an ITMan Intelligence perspective, the paper proposes practical engineering responses: capability envelopes, staged deployment, traceable authority, continuous evaluation, incident learning, skills renewal and public-infrastructure planning. The model is a decision aid, not a forecast, causal estimate or claim of peer-reviewed evidence.

Keywords: artificial intelligence; societal adaptation; institutions; workforce; education; governance; culture; infrastructure; software engineering.

Epistemic status. This is a conceptual working paper and preprint-style draft. It does not claim that every AI capability follows an exponential law, forecast a date for general intelligence, estimate employment effects or report original empirical results. Its equations are scenario and measurement frames for future testing.

1 The adaptation problem

Technical capability and social readiness do not move on the same clock. A model can be updated in weeks; professional standards, curricula, procurement rules, public infrastructure and cultural expectations may require years. The resulting difficulty is not acceleration by itself. It is the possibility of a persistent rate mismatch between what a technology can do and what a society can responsibly absorb.

The phrase "exponential AI growth" should therefore be handled with care. Empirical scaling work has found power-law relationships between selected training inputs and loss within specific model families and regimes [5, 3]. That evidence does not establish that socially meaningful capability, adoption or impact grows exponentially without limit. In this paper, exponential growth is a stress-test scenario: a way to ask what happens when the capability frontier moves substantially faster than the institutions around it.

From an ITMan Intelligence perspective, adaptation is a systems-engineering problem. The unit of concern is not the model alone. It is the full chain joining compute, data, energy, software, organisations, skills, rules, culture and accountable action.

2 A conceptual adaptation differential

Let $C(t)$ represent a defined technical capability frontier for a particular task and context. Let $A_d(t)$ represent readiness in adaptation domain d, measured against explicit criteria. An aggregate adaptation capacity can be written as a weighted profile:

$$A(t) = \sum_{d=1}^{n} w_d A_d(t), \qquad \sum_{d=1}^{n} w_d = 1. \tag{1}$$

Thabang Mokoena · ORCID 0000-0003-0689-7632

The adaptation differential is then

$$\Delta(t) = \max\left[0, C(t) - A(t)\right]. \tag{2}$$

Equations (1) and (2) are not a forecast model. Their purpose is diagnostic. They force an analyst to define the capability, choose domains, expose weights and state readiness criteria. The scalar Δ should never erase the underlying profile: a society may be prepared in regulation yet constrained by electricity, or rich in connectivity yet weak in workforce transition.

An exponential stress case may be represented as

$$C_s(t) = C_0 e^{gt}. \tag{3}$$

Equation (3) is a scenario assumption, not an empirical conclusion. The policy question is whether adaptation mechanisms can increase their learning rate without abandoning due process, inclusion or safety.

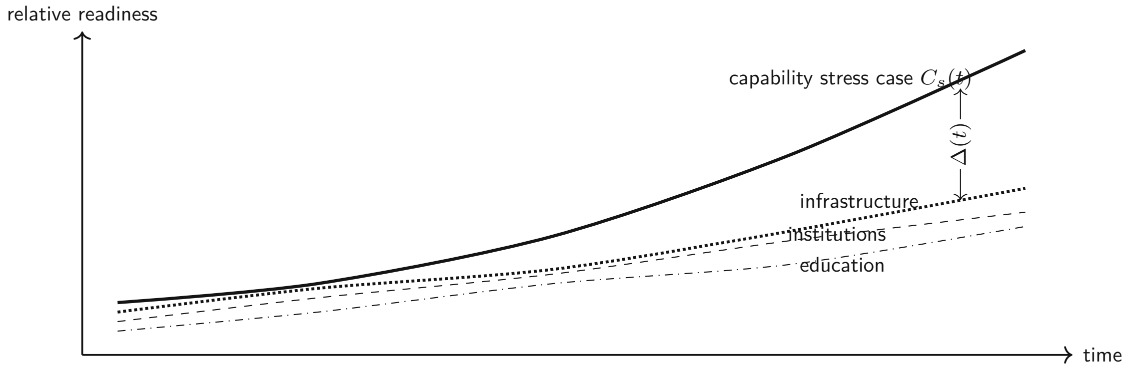

Figure 1: One frontier, multiple adaptation clocks. The capability curve is a stress scenario, not measured data; domain readiness is heterogeneous.

3 Seven coupled domains

Table 1: Domains of adaptation and example readiness questions.

Domain	Adaptation concern	Example readiness question
Technological capability	Performance, reliability, access and integration	For which bounded tasks is capability repeatable under local conditions?
Institutions	Mandates, procurement, process and organisational memory	Can the institution change its operating model without losing accountability?
Workforce	Task redesign, bargaining power, mobility and professional identity	Who gains capability, who bears transition cost and who can contest a decision?

Domain	Adaptation concern	Example readiness question
Education	Curricula, educator development and assessment integrity	Are learners prepared to understand, use and challenge AI rather than merely consume it?
Governance	Rights, risk ownership, audit and remedy	Can harms be detected, attributed, corrected and compensated?
Culture	Trust, norms, language, status and meaning	Does adoption preserve human agency and legitimate local knowledge?
Infrastructure	Compute, energy, connectivity, data and public capacity	Are benefits and dependencies distributed sustainably and resiliently?

The domains are coupled. A school system cannot build AI competence without trained educators, accessible infrastructure and assessment rules. A firm cannot responsibly automate a process without data governance, role redesign and a route for affected people to challenge errors. A state cannot depend on computational services without understanding energy, connectivity, procurement and vendor concentration.

4 Institutions, governance and the lag of legitimate change

Institutional delay is not always failure. Public consultation, legal review, professional accreditation and democratic contestation can be protective. The aim should not be to make every institution move at software speed. It should be to reduce avoidable delay while retaining the functions that make change legitimate.

A useful distinction is between *learning latency* and *decision latency*. Institutions can shorten the time required to detect a new capability, gather evidence and run bounded trials even when final rule changes appropriately take longer. Sandboxes, sunset clauses, revision schedules and transparent incident registers can turn static governance into an evidence-producing process. The NIST AI Risk Management Framework provides a practical reference for linking governance, mapping, measurement and management [7].

Proposition 1 — Govern the deployment, not the demo. Readiness should be assessed against the full operating context, including data, users, incentives, failure recovery and appeal—not against a curated model demonstration.

Proposition 2 — Preserve contestability. A consequential AI-assisted process is institutionally incomplete if affected people cannot understand the decision path, question it and reach a human authority able to provide remedy.

5 Workforce and education: adaptation as capability distribution

Workforce analysis often collapses into a count of jobs gained or lost. That framing misses task redistribution, quality of work, wage bargaining, access to tools, the burden of monitoring automation and the possibility that expertise becomes harder to develop when entry-level tasks disappear. The International Labour Organization's task-based analysis of generative AI is valuable partly because it distinguishes exposure from inevitable displacement [2].

Practical adaptation needs occupational transition maps at the level of tasks and authority. For each role, an organisation should identify which tasks may be assisted, which require preserved human skill, which new verification duties arise and what evidence would justify changing the boundary. Training must be scheduled as part of deployment, not offered as a retrospective benefit.

Education faces a related challenge. UNESCO's competency frameworks emphasise human-centred judgement, ethics, foundational understanding and creative system design [8, 9]. Those capacities imply more than prompt technique. Learners need to test provenance, recognise uncertainty, understand when not to automate and retain the domain knowledge needed to detect plausible error.

6 Culture and infrastructure are part of the technical system

Culture determines whether people will report an AI failure, defer to a fluent answer, share local knowledge or interpret automation as support, surveillance or threat. Language coverage, professional status and histories of exclusion affect adoption even when the same software is available. These factors should be treated as design inputs rather than "change management" after the technical decision has been made.

Infrastructure creates a harder boundary. AI systems depend on electricity, data centres, networks, devices, cooling, security operations and skilled maintenance. The International Energy Agency's analysis makes energy dependence explicit and highlights substantial uncertainty in future demand [4]. Resilient adaptation therefore includes capacity planning, efficiency targets, geographic redundancy, procurement diversity, lifecycle cost and graceful degradation when cloud or network access fails.

For regions with unequal connectivity and energy reliability, the relevant frontier is not the capability of the largest remote model. It is the capability that can be operated affordably, lawfully and reliably in the local environment.

7 An engineering response: the adaptation control loop

The ITMan Intelligence lens can be summarised as: systems over silos; doctrine before product; human accountability in AI-assisted decisions; and records strong enough for knowledge to survive turnover. Applied to rapid AI change, this produces a six-stage control loop.

Concrete artefacts make the loop operable:

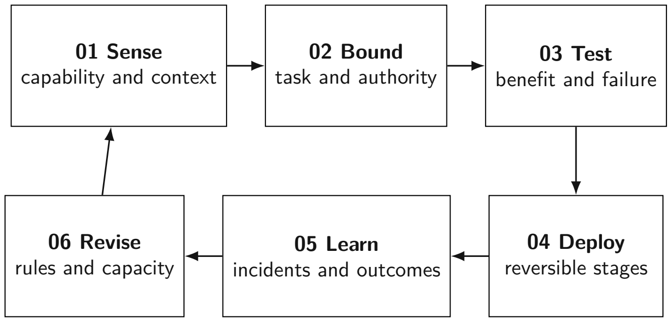

Figure 2: The adaptation control loop. Each deployment produces evidence that updates both the technical boundary and the institution's readiness model.

Capability envelope.
Tasks, contexts, data conditions and performance boundaries for which use is currently permitted.

Human authority map.
Named decision owners, escalation paths, review thresholds and remedy channels.

Evaluation register.
Versioned tests, known failure modes, subgroup checks and local acceptance criteria.

Deployment ladder.
Offline evaluation, supervised pilot, bounded production and only then wider use.

Incident and learning record.
A route from observed failure to corrective action, training and design change.

Adaptation balance sheet.
Benefits, transition costs, infrastructure dependencies and groups carrying risk.

These artefacts bring established software-engineering disciplines into the social adaptation problem. Research on machine-learning software practice shows that data, models and system components create lifecycle and integration difficulties beyond conventional application code [1]. The response is not paperwork for its own sake; it is an executable institutional memory.

8 Measurement without false precision

The differential becomes useful only when each dimension is operationalised. A readiness score should be accompanied by its definition, evidence date, uncertainty and owner. Ordinal states may be more honest than a fine-grained number: absent, designed, piloted, operational and independently reviewed. Weighting should be public within the decision process and sensitivity-tested.

Three measurements are especially important: time-to-detection of a capability change; time-to-revise a deployed boundary; and time-to-remedy a harmful or contested outcome. Together they capture whether an institution can perceive, learn and correct. They still do not measure justice or legitimacy by themselves; qualitative review and participation remain necessary.

Organisations must also balance exploration and exploitation. March's classic formulation shows why learning systems can over-invest in familiar returns or in perpetual search [6].

Applied here, adaptation needs bounded experiments that generate evidence while preserving reliable service and established rights.

9 Limitations and research agenda

This paper does not validate the existence, magnitude or direction of a global adaptation gap. Capability is task-dependent, readiness is contextual and aggregation can hide inequality. The equations omit political economy, market concentration, international dependency and strategic behaviour except where these are represented in domain criteria. No causal relationship is estimated.

Empirical work could test the framework through comparative case studies of bounded deployments. Researchers could measure whether capability envelopes, authority maps and incident learning reduce reconstruction time, severity of repeated errors or unplanned expansion of use. Distributional research should examine who participates in boundary setting and who bears adaptation cost. Infrastructure studies should connect service reliability, energy demand and local value creation.

The central claim is therefore modest: when technical capability moves rapidly, adaptation should be engineered as a continuous, evidence-producing system rather than treated as an announcement, a training event or a policy document.

10 Conclusion

Society cannot be "updated" as if it were software, and it should not be. Legitimate institutions contain histories, rights and relationships that require deliberation. Yet societies can improve how quickly they sense change, learn from bounded use, revise unsafe boundaries, distribute capability and repair harm. The adaptation differential is offered as a disciplined way to locate those gaps. Its value will depend on transparent criteria, local evidence and the willingness to revise the model itself.

References

[1] S. Amershi et al., "Software Engineering for Machine Learning: A Case Study," *ICSE-SEIP 2019*, pp. 291–300, 2019. `https://doi.org/10.1109/ICSE-SEIP.2019.00042`.

[2] P. Gmyrek, J. Berg, and D. Bescond, *Generative AI and jobs: A global analysis of potential effects on job quantity and quality*, ILO Working Paper 96, 2023. `https://doi.org/10.54394/FHEM8239`.

[3] J. Hoffmann et al., "Training Compute-Optimal Large Language Models," *arXiv*, 2022. `https://doi.org/10.48550/arXiv.2203.15556`.

[4] International Energy Agency, *Energy and AI*, IEA, Paris, 2025. `https://www.iea.org/reports/energy-and-ai`.

[5] J. Kaplan et al., "Scaling Laws for Neural Language Models," *arXiv*, 2020. https://doi.org/10.48550/arXiv.2001.08361.

[6] J. G. March, "Exploration and exploitation in organizational learning," *Organization Science*, vol. 2, no. 1, pp. 71–87, 1991. https://doi.org/10.1287/orsc.2.1.71.

[7] E. Tabassi, *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*, NIST AI 100-1, 2023. https://doi.org/10.6028/NIST.AI.100-1.

[8] F. Miao, K. Shiohira, and N. Lao, *AI competency framework for students*, UNESCO, 2024. https://unesdoc.unesco.org/ark:/48223/pf0000391105.

[9] F. Miao and C. Cukurova, *AI competency framework for teachers*, UNESCO, 2024. https://unesdoc.unesco.org/ark:/48223/pf0000391104.

Working paper v0.1 · 14 August 2026 · Thabang Mokoena · ORCID 0000-0003-0689-7632