

Pakistan's National AI Policy 2025: A Comparative Appraisal, Advantages, Risks, Execution Pathways, and Regional Benchmarks

Sohail Rao, MD, MA, DPhil¹

¹Senior Scientific Advisor to the Vice Chancellor, Dow University of Health Sciences, Baba-E-Urdu Road, Karachi 74200, Pakistan

¹Corresponding Author Email: srao@duhs.edu.pk

ABSTRACT:

Pakistan's National Artificial Intelligence (AI) Policy articulates a coherent scaffold for nationwide AI adoption built on four pillars: ring-fenced financing through a National AI Fund (NAIF); a geographically distributed network of Centers of Excellence in AI (CoE-AI); ambitious human-capital targets (e.g., nationwide awareness by 2026; ~1M trained learners and ~10k trainers by 2027); and a trust framework centered on an AI Regulatory Directorate with sectoral sandboxes. Using qualitative document analysis, we examine these instruments' funding architecture, institutional design, skills pipeline, and governance, and assess internal coherence, feasibility, and delivery risks against the policy's own baselines and timelines. On July 30, 2025, Pakistan's federal cabinet approved the National AI Policy 2025; approval announcements highlighted an AI Council, a master plan/action matrix, AI Innovation and Venture Funds, and a headline target to train one million AI professionals by 2030. To contextualize feasibility and ambition, we benchmark key instruments against regional peers India (IndiaAI Mission), China (Generative AI Measures), Sri Lanka (National AI Strategy and 2024 AI White Paper), Bangladesh (National AI Policy 2024), and Iran (National AI Plan 2025), and incorporate these comparators throughout the manuscript. (Dawn.com, 2025)

Our appraisal finds notable strengths: predictable capital that can crowd-in private and multilateral co-funding; a distributed, demand-driven CoE topology that couples R&D with incubation and workforce development; time-bound, measurable targets that create accountability; and an explicit "pro-innovation with guardrails" posture via sandboxes and rights-respecting oversight. Critical risks include: (1) NAIF portfolio governance (absence of stage-gate disbursement criteria, portfolio-mix guidance, and enhanced conflict-of-interest firewalls); (2) trainer capacity versus timelines (10k master trainers as the throughput bottleneck); (3) regulatory overlap between the AI directorate, sectoral regulators, and higher-education bodies; (4) under-specified data/compute reference architectures and access standards; and (5) ambiguous measurement of "awareness" and public-sector upskilling without protected training time.

We propose execution-ready remedies: stage-gated NAIF disbursements tied to outcomes (certifications, sandbox graduates, IP/startups); a funded national Train-the-Trainer corps anchored at CoEs; a published sandbox rulebook (eligibility, risk tiers, pre-deployment testing/red-teaming, exit-to-market); a national data reference architecture (metadata schemas, APIs, consent patterns, access tiers); and independent, pre/post awareness and skills surveys with disaggregated reporting. If implemented, these adjustments

materially raise implementation fidelity and the likelihood that Pakistan's AI ambitions translate into durable economic and social value.

KEYWORDS: *Pakistan National AI Policy; NAIF (National AI Fund); Centers of Excellence (CoE-AI); AI Governance; Regulatory Sandbox; Compute Infrastructure; Hyper-Converged Infrastructure (HCI); Human Capital; Train-The-tTainer (TTT); Public–Private Partnership (PPP); Data Reference Architecture; Inclusion and Equity; Regional Comparison; IndiaAI Mission; Generative AI Measures (China).*

INTRODUCTION:

Many Countries are racing to codify AI strategies that accelerate innovation while protecting rights and public trust. Pakistan's National AI Policy 2025 situates this challenge within the nation's demographic opportunity and its broader digital agenda, explicitly linking AI adoption to the Digital Pakistan Policy, the Cloud First Policy, and forthcoming personal data protection frameworks, while drawing normative cues from the UN SDGs and ITU's "AI for Good."

India's IndiaAI Mission (₹10,371.92 crore over five years) couples a public–private compute program of 10,000+ GPUs with datasets, model development, skills and safety streams; China's 2023 Generative AI Measures impose pre-deployment obligations on public-facing services; Sri Lanka's 2024 AI White Paper and National AI Strategy emphasise use-case pilots and central coordination; Bangladesh's 2024 draft National AI Policy aligns to the Smart Bangladesh 2041 vision; and Iran's parliament approved the outlines of a National AI Plan in May 2025 aiming for a top-10 global ranking by 2032. (Press Information Bureau [PIB], 2024; Cyberspace Administration of China [CAC], 2023; Ministry of Technology Sri Lanka, 2024a, 2024b; OneTrust DataGuidance, 2024; ICT Division, 2024; SpecialEurasia, 2025). (CAC, 2023). (SpecialEurasia, 2025).

The draft organizes its program through four developmental pillars that balance demand and supply side interventions: Market Enablement, Progressive and Trusted Environment, Awareness and Readiness, and Sectoral Transformation & Evolution. Concretely, it proposes (i) ring-fenced financing via a NAI with independent governance and perpetual inflows from the national ICT R&D fund; (ii) a two-tier network of CoE-AI complete centers in Islamabad, Karachi, and Lahore with auxiliary hubs in additional cities operating through public–private partnerships and evolving to self-sustaining models; (iii) quantified awareness and skills targets to build national capacity; and (iv) a trust and governance layer that includes regulator-supported sectoral sandboxes and data-standardization efforts anchored in lawful processing and citizen rights.

Against this backdrop, our manuscript examines the internal coherence and execution realism of the policy's instruments NAIF's financing architecture, the CoE-AI topology and mandates, the human-capital pipeline, and the governance/sandbox approach to surface where the policy is strongest and where delivery risks remain, and to offer actionable adjustments that can raise implementation fidelity and public value.


METHODS:

We conducted a qualitative policy analysis of the Government of Pakistan's National Artificial Intelligence Policy 2025 using the uploaded PDF as the primary source. Provenance and authorship were verified from the title/cover and ministry attributions.

Comparative benchmarking: We triangulated Pakistan's instruments with neighbouring policies and official notices PIB releases on the IndiaAI Mission (compute $\geq 10k$ GPUs, funding outlay), China's Interim Measures for Generative AI Services (CAC), Sri Lanka's National AI Strategy and 2024 AI White Paper (Ministry of Technology), Bangladesh's National AI Policy 2024 (draft; ICT Division notices and summaries), and reporting on Iran's National AI Plan approval (May 2025). These sources informed cross-section comparisons in Financing/Governance, Institutional Architecture, Human Capital, Governance/Sandboxes, and Compute/Data. (PIB, 2024; CAC, 2023; Ministry of Technology Sri Lanka, 2024a, 2024b; OneTrust DataGuidance, 2024; ICT Division, 2024; SpecialEurasia, 2025)

We structured analysis around five domains explicitly articulated in the draft: (i) financing (NAIF scope, governance), (ii) institutions (CoE-AI topology/mandates), (iii) human capital (awareness/skills targets), (iv) governance (AI directorate, sandbox, review cadence), and (v) compute/data infrastructure. These domains map to the policy's Policy Directives and Implementation sections and their associated targets/definitions (Figure 1).

Figure 1: Key Instruments for a Coherent Scaffold for Nationwide AI Adoption in Pakistan



Using two-cycle coding for data extraction (structural → axial), we tagged clauses for: instrument/mechanism, timelines, quantitative targets, stewardship entity, and verification method. Examples included the $\geq 30\%$ perpetual allocation to NAIF, CoE-AI PPP model, geography, and sectoral/academic compute-data access targets. Extracted items were placed into a matrix capturing: statement text, page/line range, implied theory of change, dependencies, risks, and candidate mitigations.

RESULTS:

Financing & Governance (NAIF):

The policy establishes the NAIF as a dedicated, perpetual vehicle by directing no less than 30% of Ignite's national ICT R&D fund to NAIF, complemented by initial PSDP working capital. NAIF is designed with independent oversight and an external Board of Directors (≤ 11 members) drawn from industry, academia, and government, and explicit powers to co-fund with multilateral and private partners. Operational integrity is reinforced through a CEO recruited on open merit and annual public audits by a top-rated firm. Collectively, these provisions create a ring-fenced and professionally governed mechanism to finance skills, R&D, and commercialization at scale.

India's Mission explicitly links funding to large-scale public-private compute and indigenous model work; Bangladesh's draft policy sketches PPP funding but remains at consultation/early institutionalisation. Pakistan's ring-fenced NAIF compares favourably on predictability; adding public stage-gates and portfolio-mix guidance would align execution clarity with India's infra-and-model pillars. (PIB, 2024). (OneTrust DataGuidance, 2024; ICT Division, 2024).

Institutional Architecture (CoE-AI network):

To operationalize delivery, the policy proposes a two-tier network of CoE-AI: fully fledged CoEs in Islamabad, Karachi, and Lahore, supported by auxiliary hubs in additional cities. Centers are to be sited in public-private partnerships, evolve to self-sustaining operations within 5 (+2) years, and follow a "zero-equity"/minimal-fee model. Their mandate encompasses demand-driven R&D, startup incubation, and access to cloud/data resources. Additionally, it includes curriculum development across school and university levels, short-course skills programs, hands-on internships, and support for theses and academic research, thereby linking talent, research, and commercialization in a single topology.

Sri Lanka proposes a National AI Center and central coordination to link pilots with measurement; India's Mission includes an Innovation Centre for indigenous LMMs; Bangladesh's draft calls for a UGC-anchored AI hub. Pakistan's wider CoE footprint is an advantage, provided sustainability (sponsored research, fee-for-service, and limited equity via a public-interest vehicle) is codified. (Ministry of Technology Sri Lanka, 2024a, 2024b). (OneTrust DataGuidance, 2024; ICT Division, 2024).

Human-capital Targets:

The draft sets quantifiable national goals: 90% awareness among people with internet access by 2026, and 1 million learners trained with 10,000 new trainers by 2027 from a baseline in which $< 10\%$ of the current computing/IT workforce is AI-skilled. It introduces a public-sector program aimed at raising awareness

among all Grade 12–22 employees and training 70% of existing IT/AI-adjacent staff and 100% of new inductees by 2026. Higher-education measures include a 30% uplift in AI-focused scholarships and a financed pipeline featuring train-the-trainer bootcamps and a national internship scheme delivering ~20,000 placements annually (Table 1).

Table 1: Equity and Measurement Plan – Pakistan’s National AI Policy 2025

Metric	Instrument	Cadence	Disaggregation	Target / Success Threshold
Awareness	National pre/post survey; randomized sampling; validated instrument	Annual	Gender, region, disability	≥15-pt lift vs. baseline
Scholarships	Disbursal logs + audit; completion tracking	Quarterly	Gender, field, SES	≥40% women; ≥20% underserved regions
Public-sector Upskilling	HRIS + assessments; recertification records	Semi-annual	Cadre, ministry, region	≥70% incumbents; 100% new entrants (by 2026)
Employment Outcomes	Graduate tracer (6/12/24 months)	Annual	Gender, region, disability	≥60% in-field at 12 months
Sandbox Outcomes	Public registry; exit reports; incidents	Quarterly	Sector, risk tier	≥30% pilots → production

India’s ‘FutureSkills’ pillar pairs scaled training with access to compute, models, and startup financing; Sri Lanka’s strategy highlights talent retention and pilot-linked training; Bangladesh’s draft emphasises university-centric pipelines. Pakistan’s 1M-by-2030 headline target is competitive, but hinges on a funded Train-the-Trainer corps and independent, disaggregated measurement. (OneTrust DataGuidance, 2024; ICT Division, 2024)

Governance, Trust, and Sandboxes:

For responsible deployment, the policy creates an AI Regulatory Directorate (ARD) within the forthcoming National Commission for Personal Data Protection, charged with monitoring AI-automated functions, advising on sectoral regulation, and coordinating a sandbox approach for novel algorithms and applications. In parallel, the policy calls for an integrated regulatory sandbox to pilot research platforms and commercial systems, linking regulatory learning with real-world testing and international best-practice partnerships.

China’s 2023 Generative AI Measures create a pre-deployment license-and-guardrail regime for public-facing services, offering high compliance clarity. Pakistan’s ARD and sandbox model is pro-innovation; a public sandbox rulebook and regulator MoUs would close the predictability gap. (CAC, 2023).

Compute and Data Foundations:

Delivery is anchored in enabling infrastructure: upgraded hyper-converged infrastructure (HCI) in public data centers, national/provincial data repositories, and sandboxed access to curated data lakes for public and private use. The CoE-AI network is tasked to provide sectoral, hybrid/multi-cloud environments for pilots and model training, with targets to serve both industry (e.g., registration of 100 commercial models annually) and academia (e.g., access for 150 institutions and training 150 unique models per year).

India targets $\geq 10,000$ GPUs in public-private configurations and has begun empanelment/tendering. Pakistan should complement HCI/data-lake goals with a time-phased compute roadmap (on-prem + cloud credits) tied to CoE utilisation and quotas for academia/industry. (PIB, 2024)

DISCUSSION:

The policy's architecture exhibits several systemic strengths that, if executed well, could materially accelerate safe AI uptake. First, predictable capital with external leverage is baked in a ring-fenced NAIF with a perpetual $\geq 30\%$ allocation from Ignite, independent board oversight, co-funding authority, merit-recruited CEO, and annual public audits, which buffer programs from fiscal shocks and can crowd-in multilateral and private investment. Second, distributed, demand-driven capacity comes via a two-tier CoE-AI network, complete centers in Islamabad, Karachi, and Lahore with auxiliary hubs elsewhere delivered through PPPs that pair public goals with private execution while linking R&D, incubation, curricula, skills, and internships under one roof. Third, time-bound human-capital targets (90% awareness by 2026; 1M learners and 10k trainers by 2027, plus public-sector upskilling and scholarship uplift) create accountability and a visible pipeline, with explicit equity aims for women and persons with disabilities. Fourth, trust-by-design is signaled through an AI Regulatory Directorate within the data-protection commission and regulator-supported sandboxes, a stance that enables experimentation in high-risk domains while preserving rights. Finally, the policy is realistic about data/compute: it commits to HCI upgrades, national/provincial repositories, sandboxed data-lake access, and sectoral multi-cloud environments managed through the CoE network, an essential substrate for training and evaluation.

Against India's compute-forward mission and China's binding deployment rules, Pakistan's edge lies in perpetual financing (NAIF) and inclusion. Priority moves include compute phasing, sandbox codification, and trainer throughput, which will close the gap most efficiently. (PIB, 2024; CAC, 2023)

Balanced against these strengths are execution risks that could blunt impact without early course-correction. The NAIF portfolio lacks explicit stage-gates, outcome-tied disbursement criteria, and portfolio-mix guardrails (e.g., spend shares for skills vs compute vs R&D vs startups), and it would benefit from stronger conflict-of-interest controls beyond board composition; without these, capital can diffuse or concentrate sub-optimally (Table 2). The ambitious trainer throughput (10k in ~ 2 years) is the hard constraint behind the 1M-learner goal; absent a funded train-the-trainer (TTT) pipeline with stipends and service commitments, delivery risk is high. CoE financial sustainability may be challenged by a "zero-equity/minimal-fee" posture for deep-tech labs and testbeds; a phased revenue mix (sponsored research, services, limited equity via a public-interest holding vehicle) could be necessary post-PPP. The proposed AI directorate's remit may overlap with sector regulators and higher-education bodies; without codified MoUs, SLAs, and joint-review processes, uncertainty can slow deployment and deter industry.

Table 2: Train-the-Trainer Throughput Plan to 2023 – Pakistan’s AI Policy 2025

Year	Master Trainers	Learners per Trainer (Annual)	Learners Trained (Annual)	Target Pass Rate	Milestones	Cumulative Graduates
2025	2000	100	200000	≥80%	TTT cohorts launched; proctoring vendor onboarded	200000
2026	5000	100	500000	≥80%	Provincial quotas met; internship pipeline live	700000
2027	8000	100	800000	≥82%	TTT retention ≥85%; employer partnerships 300+	1500000
2028	10000	90	900000	≥82%	Rural/remote share ≥35%; women ≥35%	2400000
2029	10000	90	900000	≥84%	Recertification cycle v1 completed	3300000
2030	10000	80	0	≥85%	1M cumulative graduates achieved	3300000

Additionally, data access and standards are underspecified; reference architectures, interoperability baselines, consent patterns, and open-data tiers are prerequisites for lawful, proper access across domains. Finally, targets for awareness and public-sector upskilling require operational definitions and protected training time with recertification; otherwise, metrics risk becoming vanity indicators rather than capability gains.

To mitigate these risks, several execution moves should be front-loaded. Stage-gate NAIF by publishing outcome-linked disbursement gates (e.g., certifications earned, sandbox pilots graduating to production, IP events, startup survival) on a public dashboard audited annually, leveraging the audit provision already envisioned. Stand up a National TTT Corps anchored at CoEs/auxiliaries with funded fellowships, micro-credentials, and two-year service obligations to guarantee classroom and workplace delivery at scale. Sustain CoEs through a clear revenue strategy: broaden fee-for-service menus for labs/testbeds; create a capped, ethics-guarded equity vehicle for spinouts; and formalize sponsored-research pipelines with industry. Codify the regulatory operating model via MoUs between the AI directorate and sector regulators, and publish a sandbox rulebook that covers eligibility, risk tiers, pre-deployment testing and red-teaming, incident reporting, and exit-to-market criteria. Adopt a national data reference architecture with standard metadata schemas, secure APIs, access tiers, and consent patterns, resourcing data stewards at CoEs, and maintaining a national catalog to curb duplication. For outcomes measurement, institutionalize awareness and skills KPIs through independent, nationally representative pre/post surveys (disaggregated by gender, region, and disability), and embed public-sector training within HR policy with protected time and periodic recertification to meet 2026 targets (Table 3).

Taken together, the design’s strengths, ring-fenced capital, distributed capacity, measurable human-capital goals, trust mechanisms, and enabling infrastructure provide a credible scaffold. The discussion above argues that the decisive factor will be delivery discipline: tightening NAIF governance, solving the trainer bottleneck, clarifying regulatory roles, operationalizing data/compute access, and measuring what matters.

If the recommended mitigations are adopted in the first implementation cycle, the probability that Pakistan's AI policy generates durable economic and social value rises substantially.

Table 3: Phased Execution Roadmap – Pakistan's National AI Policy 2025

Timeline	Priority Actions	Primary Owners	KPIs/Evidence
Phase I: (0–6 months)	Stand up NAIF governance & public dashboard; publish Sandbox Rulebook; activate cloud credits & shared CoE access; appoint data stewards; finalize Train-the-Trainer (TTT) design; launch baseline awareness survey.	MoITT/NAIF, ARD, CoE-AI, HEC, Sector Regulators	NAIF bylaws & dashboard live; Sandbox Rulebook published; # orgs using cloud credits; # data stewards appointed; national baseline survey completed
Phase II: (6–18 months)	PPP procurements to commit $\geq 10k$ GPU-equivalent capacity; run first sandbox cohorts; launch TTT cohorts; publish national data catalog & APIs; sign MoUs/SLAs with regulators; start sponsored-research programs at CoEs.	NAIF, MoITT/PC, CoE-AI, ARD, Regulators, Industry	Committed GPU-equivalent $\geq 10k$; # sandbox pilots $\rightarrow$ production; # master trainers deployed; data catalog live; # MoUs/SLAs in force; sponsored-research \$/projects
Phase III: (18–36 months)	Regionalize compute clusters; scale sector accelerators (health, agri, finance); expand internships/apprenticeships; publish yearly audit & outcomes; refresh rulebook based on incident/impact data.	NAIF, CoE-AI, ARD, Provinces, Ministries	Utilization $\geq 70\%$; sector pilots at scale; # apprenticeships; audited outcomes published; updated rulebook (v2) released

CONCLUSION:

Regionally, India has signaled sheer scale through the IndiaAI Mission: a five-year, centrally funded push that pairs a public–private compute fabric exceeding 10,000 GPUs with a national dataset platform, an Innovation Centre for indigenous multimodal models, a FutureSkills pipeline, startup finance, and a Safe & Trusted AI stream. China complements long-running industrial policy with binding Generative-AI Measures (2023) that require security, content, and accountability checks before public release, effectively licensing high-risk deployments and creating strong ex-ante predictability for providers and regulators alike. Sri Lanka has opted for a coordinated, use-case–driven path: its National AI Strategy and 2024 White Paper emphasize measured pilots overseen by a centrally guided National AI Center, with early attention to institutional coordination and talent sustainability. Bangladesh's 2024 National AI Policy aligns to the Smart Bangladesh 2041 vision, tasks the University Grants Commission with standing up a national AI hub, and foregrounds ethics and PPPs; much of the machinery is still moving from consultation to institutionalization. Iran, meanwhile, approved the outlines of a national plan in May 2025 that aims for a top 10 global ranking by 2032, with sovereign compute/data ambitions and evolving governance designed to knit sector pilots into national priorities. (PIB, 2024). (PIB, 2024; CAC, 2023; Ministry of Technology Sri Lanka, 2024a, 2024b; OneTrust DataGuidance, 2024; ICT Division, 2024; SpecialEurasia, 2025)

Against this backdrop, Pakistan's National AI Policy is competitively comprehensive on institutional design, perpetual, ring-fenced capital via NAIF; a geographically distributed CoE-AI network; quantified human-capital targets, including a headline aim of one million AI professionals by 2030; and a trust architecture anchored by an AI Regulatory Directorate and regulator-supported sandboxes with lawful data access. Where regional peers set a higher bar is in two areas: first, compute at scale, where India has published explicit, time-bound supply signals; second, codified deployment rules, where China's ex-ante licensing provides clarity for high-risk releases. Closing this gap without losing Pakistan's inclusion-first posture means translating policy architecture into near-term proof points. A pragmatic path is a three-phased approach: Phase I (0–6 months) activates cloud credits and shared access to a few thousand data-center-class accelerators through CoEs, with quotas for universities and SMEs and a visible booking system; Phase II (6–18 months) executes PPP procurements that bring national capacity into the $\geq 10\text{k}$ -GPU range with fair-use scheduling, inference capacity for civic-scale projects, and sector carve-outs; Phase III (18–36 months) regionalizes capacity and scales specialized accelerators for health, agriculture, and finance, with cost-recovery through metered services and sponsored research (Table 4).

Financing must signal the same discipline. NAIF should publish stage-gated disbursements linked to verifiable output certifications issued, sandbox pilots graduating to production, IP events and licensing, startup survival at 12/24 months with quarterly dashboards and annual independent audits. Portfolio guidance should be explicit (for example, minimum shares for skills, research, compute access, and commercialization), and conflict-of-interest firewalls should extend beyond board composition to cover screening of affiliated grantees and co-investment partners. The CoE-AI network will remain the delivery engine. Still, durability requires a three-stream revenue mix beyond the current “minimal-fee/zero-equity” posture: transparent rate cards for testbeds and labs, sponsored-research pipelines with industry and government, and limited, guard-railed equity participation in spinoffs via a public-interest holding vehicle—while preserving academic openness and avoiding capture.

On the human-capital side, the binding constraint is trainers, not learners. To make “1M by 2030” a reality, Pakistan should establish a National Train-the-Trainer Corps of 10,000 master trainers, anchored at CoEs and auxiliary hubs, with each trainer committing to two years of service, receiving stipends, and earning standardized micro-credentials assessed by independent proctoring. Set a 1:100 annual throughput ratio per master trainer, publish province-wise targets, and embed a placement service that routes completers into apprenticeships, internships, and funded fellowships. For the public sector, make progress verifiable by mandating protected training time, baseline and recertification assessments, and an inventory of AI-relevant roles with competency tiers that map to promotion pathways.

Governance should convert intent into operating clarity. Codify MoUs and SLAs between the AI directorate and sectoral regulators (health, finance, education, transport), define who authorizes what (and in how many days), and publish a sandbox rulebook that sets eligibility thresholds, risk tiers, pre-deployment testing and red-teaming requirements, incident reporting, and exit-to-market criteria. In parallel, adopt a national data reference architecture with standard metadata schemas, consent patterns, privacy-preserving access tiers, and secure APIs; appoint data stewards at each CoE; and publish a national catalog so repositories and HCI upgrades yield lawful, proper access rather than siloed infrastructure. Because FX constraints and vendor lock-in can derail affordability, prioritize hybrid/multi-cloud designs, open standards, and sovereign controls over sensitive data to keep total cost of ownership predictable and sovereignty intact.

Table 4: It highlights where Pakistan leads (perpetual financing; distributed CoEs; inclusion targets) and where to catch up (explicit compute targets; codified sandbox rules).

Country	Policy Status / (year)	Capital & Compute	Institutions	Skills Targets	Governance / Rules	Notes for Pakistan
Pakistan	National AI Policy 2025 (approved)	NAIF ring-fenced ($\geq 30\%$ of Ignite); HCI upgrades; data lakes; hybrid/multi-cloud	CoE-AI network (full CoEs in Islamabad, Karachi, Lahore + auxiliaries) via PPPs	1M AI professionals by 2030; public-sector upskilling; scholarships uplift	AI Regulatory Directorate; sectoral sandboxes; data reference architecture	Strong on institutional completeness and inclusion; needs time-phased compute roadmap , sandbox rulebook, and NAIF stage-gates.
India	IndiaAI Mission (2024)	~₹10,371.92 cr (~US\$1.25B); 10,000+ GPUs public-private compute; IndiaAI Datasets Platform	Innovation Centre for indigenous models; national tools	FutureSkills skilling pillar; startup financing	Safe & Trusted AI track; standards uplift	Clear, compute-forward pipeline from infra → models → deployment; useful benchmark for Pakistan's compute phasing.
China	New Generation AI Plan (2017→); Generative-AI Measures (2023)	Multi-city compute; state/provincial funding	National/municipal AI hubs	STEM pipelines; industrial upskilling	Pre-deployment license-and-guardrail regime for public-facing GenAI	Sets bar on regulatory clarity; Pakistan can mirror via a public sandbox rulebook + MoUs.
Sri Lanka	National AI Strategy + AI White Paper (2024)	Programmatic budgets via line ministries; pilot funding	National AI Center; central coordination	Use-case-linked training; talent retention	Principle-based governance; measurement focus	Emphasis on pilots + measurement complements Pakistan's CoE model.
Bangladesh	National AI Policy 2024 (draft)	UGC-anchored research hub; PPP resourcing (emerging)	University-centric hub	Early pipeline development	Ethics guidance; initial governance framing	Still institutionalizing; Pakistan is ahead on approval & financing certainty.
Iran	National AI Plan outlines approved (2025)	Sovereign compute/data ambitions (details evolving)	New national AI bodies forming	School-to-work programs (reported)	Sector pilots under development	Sets long-term ambition (top 10 by 2032); specifics still maturing.

Finally, treat equity and resilience as non-negotiable metrics, not slogans. Use independently administered, nationally representative pre/post surveys (disaggregated by gender, region, and disability) to measure awareness, scholarship reach, and employment outcomes; localize tools and curricula for Urdu and regional

languages to widen access; and set civil-service KPIs for completion rates and recertification. If, within the next 12–18 months, Pakistan makes visible progress on a staged compute build-out, publishes and uses a sandbox rulebook, launches a funded trainer corps with province-wise quotas, and brings NAIF’s stage-gates and dashboards online, it will credibly match India’s scale signals and China’s clarity signals while remaining true to its own strengths perpetual financing, distributed capacity, and inclusion turning a strong policy scaffold into durable scientific, economic, and social value.

REFERENCES

1. Cyberspace Administration of China. (2023, August 15). Interim measures for the management of generative artificial intelligence services. China Law Translate. <https://www.chinalawtranslate.com/en/generative-ai-interim/>
2. Dawn.com. (2025, July 30). Federal cabinet approves National AI Policy 2025. <https://www.dawn.com/news/1927634>
3. Government of Pakistan, Ministry of IT & Telecom. (2025). National Artificial Intelligence Policy 2025 (approved policy). Islamabad, Pakistan
4. ICT Division, Government of Bangladesh. (2024). Draft – National strategy for artificial intelligence of Bangladesh (Version 3.0). https://ictd.portal.gov.bd/sites/default/files/files/ictd.portal.gov.bd/page/6c9773a2_7556_4395_bbec_f132b9d819f0/Draft%20-%20Mastering%20National%20Strategy%20for%20Intelligence%20-%20Bangladesh.pdf
5. Ministry of Technology, Sri Lanka. (2024a). Artificial intelligence in Sri Lanka: White paper. <https://mot.gov.lk/assets/files/AI%20White%20Paper%20March%202024-c09aa49f7990358ad1442103b804511d.pdf>
6. Ministry of Technology Sri Lanka. (2024b). National AI strategy for Sri Lanka. https://mot.gov.lk/assets/files/National_AI_strategy_for_Sri_Lanka-08a88c78d541a0746aeb8c71ed312231.pdf
7. One Trust Data Guidance. (2024, April 4). Bangladesh ICT Department publishes draft AI policy. <https://www.dataguidance.com/news/bangladesh-ict-department-publishes-draft-ai-policy>
8. Press Information Bureau (PIB), Government of India. (2024, March 7). Cabinet approves ambitious IndiaAI Mission to strengthen the AI innovation ecosystem. <https://www.pib.gov.in/PressReleasePage.aspx?PRID=2012355>
9. SpecialEurasia. (2025, May 20). Iran’s national artificial intelligence plan was approved by parliament. <https://www.specialeurasia.com/2025/05/20/iran-ai-technology/>
10. Press Information Bureau (PIB), Government of India. “Cabinet Approves Ambitious IndiaAI Mission to Strengthen the AI Innovation Ecosystem.” 7 Mar 2024. <https://www.pib.gov.in/PressReleasePage.aspx?PRID=2012355>
11. Cyberspace Administration of China (via China Law Translate). “Interim Measures for the Management of Generative Artificial Intelligence Services.” Effective 15 Aug 2023. <https://www.chinalawtranslate.com/en/generative-ai-interim/>
12. Ministry of Technology, Government of Sri Lanka. “Artificial Intelligence in Sri Lanka – White Paper.” Mar 2024. <https://mot.gov.lk/assets/files/AI%20White%20Paper%20March%202024-c09aa49f7990358ad1442103b804511d.pdf>
13. Ministry of Technology, Government of Sri Lanka. “National AI Strategy for Sri Lanka.” 2024. https://mot.gov.lk/assets/files/National_AI_strategy_for_Sri_Lanka-08a88c78d541a0746aeb8c71ed312231.pdf

14. One Trust Data Guidance. “Bangladesh ICT Department publishes draft AI policy.” 4 Apr 2024.
<https://www.dataguidance.com/news/bangladesh-ict-department-publishes-draft-ai-policy>
15. ICT Division, Government of Bangladesh. “Draft – National Strategy for Artificial Intelligence of Bangladesh (Version 3.0).”
https://ictd.portal.gov.bd/sites/default/files/files/ictd.portal.gov.bd/page/6c9773a2_7556_4395_bbec_f132b9d819f0/Draft%20-%20Mastering%20National%20Strategy%20for%20Artificial%20Intelligence%20-%20Bangladesh.pdf
16. Special Eurasia. “Iran’s National Artificial Intelligence Plan Approved by Parliament.” 20 May 2025.
<https://www.specialeurasia.com/2025/05/20/iran-ai-technology/>
17. Government of Pakistan, Ministry of IT & Telecom, National Artificial Intelligence Policy 2025 (Approved Policy) (Islamabad, 2022/consultation text).