






Table of contents

Acronyms and abbreviations	6
Introduction	10
Important limitations and assumptions	10
Chapter 1. Executive summary	12
Solving for jobs	13
Why it matters	13
Key findings	13
Most significant evidence	14
The burning questions	14
Chapter 2. The state of the AI nation	16
The problem is jobs	16
Developing AI skills as a national priority	19
A solid AI foundation	20
Education is key	20
A human in the loop	21
Expert opinion and analysis	21
Chapter 3. Why two models? The context behind the comparison	22
The Goldilocks zone	22
The Western hyperscaler model	23
The Chinese open-weight tier	24
The composite picture	25
The Chinese tier is within single digits of the frontier, at a fraction of the cost	25
The critical African-language benchmarks	25
The language state of play in South Africa	27
The Te Hiku Media case study	29
South Africa already has these critical AI assets	30
The Non-Aligned Movement can inform South Africa's response	31
Value in non-alignment	31
A more sovereign path forward	32
Chapter 4. The published scientific evidence for an AI learning pathway in South Africa	34
The demand for AI tools is already here	35
AI supports teaching, it does not replace it	35
The entrepreneurship and employability evidence	37
Equity is a design choice, not a by-product	37
The dangers of AI in learning	37
Credibility of credentials: the assessment problem follows the learner	38
AI does not and should not replace teachers	38
Governance builds the pathway inside the system, not beside it	40
Chapter 5. Driving job creation with appropriate AI education	41
Training the teachers and trainers of tomorrow	45
Developing AI training that works	45
AI skills training to new AI jobs	46
Chapter 6. Conclusion	48
Chapter 7. Recommendations. How to make it work	51
AI use disclosure	54
Impressum	55
References	56
Additional studies screened in the PRISMA corpus, and further reading	57



Students at the We Think Code_ campus in Johannesburg, South Africa.

Acronyms and abbreviations

ACL Association for Computational Linguistics. The main international academic body and conference series for natural language processing research. Two of the African-language evaluation studies cited in this report were published in its proceedings.

AfriMGSM African Multilingual Grade School Math. The mathematics component of the IrokoBench benchmark, and the task on which AI models perform worst in African languages.

AfriMMLU African Massive Multitask Language Understanding. The general knowledge component of IrokoBench, adapted from the international MMLU test.

AfriXNLI African Cross-Lingual Natural Language Inference. The reasoning and inference component of IrokoBench, which tests whether a model can judge how two statements relate to one another.

AI Artificial Intelligence.

API Application Programming Interface. The paid technical connection through which developers plug an AI model into their own software, and one of the cost barriers that open-weight models remove.

B-BBEE Broad-Based Black Economic Empowerment. A South African government policy and legislative framework designed to redress the economic inequalities of the apartheid era.

BRICS The intergovernmental grouping originally comprising Brazil, Russia, India, China and South Africa, expanded from 2024 to include further members.

CSIR Council for Scientific and Industrial Research. South Africa's national research organisation, whose Meraka Institute Human Language Technologies Research Group was for two decades the primary academic home of South African speech technology work.

EdTech Education technology.

ERIC Education Resources Information Center. The United States Department of Education's repository of education research, the world's largest, and the database underpinning this report's literature review.

GLM General Language Model. The model line developed by the Chinese company Zhipu AI, notable because GLM-5.1 was trained without any NVIDIA hardware.

GPQA Graduate-Level Google-Proof Question and Answer. A science reasoning benchmark designed so the answers cannot be found by searching, and currently regarded as the most discriminating test of the leading models.

GPT Generative Pre-trained Transformer. The model architecture and product line developed by OpenAI.

HIV/AIDS Human Immunodeficiency Virus and Acquired Immunodeficiency Syndrome. South Africa's national response to the epidemic is used in this report as the model for a mass AI literacy campaign.

IDC Industrial Development Corporation of South Africa. The state-owned development finance institution funding the AI incubator described in chapter 2.

IMF International Monetary Fund.

IoT Internet of Things. Networked physical devices, one of the subjects of the Nigerian experiential learning workshops cited in chapter 4.

LLM Large Language Model. The class of AI system, trained on very large quantities of text, that underpins all the models discussed in this report.

MIND Institute Machine Intelligence and Neural Discovery Institute at the University of the Witwatersrand, directed by Professor Benjamin Rosman.

MIT licence Massachusetts Institute of Technology licence. A permissive open-source licence allowing free reuse, modification and commercial deployment, under which DeepSeek releases its models.

MLLU Massive Multitask Language Understanding. A standard general knowledge and reasoning benchmark, now largely saturated at the top of the market.

MRF Masakhane Research Foundation. The South African non-profit behind IrokoBench.

NAM Non-Aligned Movement. Founded at Belgrade in 1961, and the political tradition this report applies to South Africa's AI strategy.

NEET Not in Employment, Education or Training. The category covering 3.9 million South Africans aged 15 to 24.

NLP Natural Language Processing. The field of computer science concerned with machines handling human language.

POPIA Protection of Personal Information Act, Act 4 of 2013. South Africa's data protection law. Section 71 prohibits reliance on purely automated decision-making, which is the legal hook Monisha Prem identifies for AI accountability.

PRISMA Preferred Reporting Items for Systematic Reviews and Meta-Analyses. The international standard for conducting and reporting systematic literature reviews, and the method used for this report's evidence base.

SA South Africa.

SABC South African Broadcasting Corporation. The public broadcaster, whose zero-rated SABC+ platform is already carrying AI literacy videos.

SADiLaR South African Centre for Digital Language Resources. Founded in 2016, hosted at North-West University and funded by the Department of Science, Technology and Innovation. It digitises text, speech and multi-modal resources for all 12 official languages

and is the closest thing the country has to sovereign language-data infrastructure.

SARIR South African Research Infrastructure Roadmap. The national funding framework under which SADiLaR sits.

SARS South African Revenue Service.

SETA Sector Education and Training Authority. The sector bodies responsible for skills development levies and training in South Africa.

SWE-Bench Software Engineering Benchmark. A test of whether a model can resolve real software problems, used here in its Pro version.

TVET Technical and Vocational Education and Training. In South Africa, the college system most likely to anchor a post-school AI skills pathway.

WEO World Economic Outlook. The IMF's twice-yearly global economic report and database.

Technical terms

Agentic Describing an AI system that carries out multi-step tasks on its own rather than answering a single question, for example completing a task in a computer terminal.

Algorithmic sovereignty Control not only over a country's data but over the systems that turn that data into economic and social value. The term is Professor Benjamin Rosman's.

Benchmark A standardised test used to compare AI models. This report treats benchmark scores as establishing broad tiers rather than fine rankings.

Compute Computing power, usually specialised chips, needed to train and run AI models. Along with data and tools, one of the three access constraints Kuben Nair identifies.

Context window The amount of text a model can consider at once. A one-million-token window allows a model to process a very long document in a single pass.

Data sovereignty A country's control over data generated within its borders, including where that data is stored and processed and who may access it.

Fine-tuning Further training of an existing model on a specific dataset to improve its performance on a particular task or language.

Frontier model The most capable AI models currently available, generally the closed commercial systems from the largest American laboratories.

Hugging Face An online platform hosting freely downloadable AI models and datasets, and the source developers such as Lee Vandre Napier use to build local-language applications.

Hyperscaler One of the handful of technology firms operating data centres at global scale, principally Amazon, Microsoft, Alphabet and Meta.

Matric The South African National Senior Certificate, the school-leaving qualification.

Open-weight A model whose trained parameters are published, allowing anyone to download, run and adapt it on their own hardware. Used in this report in preference to “open-source”, which strictly refers to published code.

Parameter One of the internal values a model learns during training. Model size is usually described by parameter count, so Lelapa AI’s InkubaLM is a 400-million-parameter model.

Self-hosting Running an AI model on your own or your country’s computers rather than calling a provider’s servers, which breaks the data link to the provider’s jurisdiction.

Token The unit into which text is broken for processing, roughly a short word or part of a word. AI pricing is quoted per million tokens.

Zero-rated Data that is free to the user because the provider or a sponsor covers the cost. Zero-rating youth AI access is among the most frequently recommended interventions in this report.

Organisations and projects referenced

Cassava Technologies The pan-African technology group building fibre, data centre and AI computing infrastructure across Southern Africa, led in South Africa and Botswana by Ziaad Suleman.

Collective X South Africa’s national co-ordinating intermediary for digital skills development, which commissioned this report.

Harambee Harambee Youth Employment Accelerator, the South African non-profit that has connected millions of young work-seekers to jobs and training.

InkubaLM Lelapa AI’s deliberately small open-source multilingual model, the “dung beetle language model”, which runs on entry-level smartphones without cloud dependency.

IrokoBench The Masakhane benchmark measuring AI performance across 17 African languages, and the authoritative source on the language performance gap.

Karya The Indian language data company, originally a Microsoft Research project, that pays contributors well above local piecework rates and provides royalties on relicensing.


Lelapa AI The Johannesburg research and product laboratory founded in 2022, developer of the Vulavula service and the InkubaLM model.

Masakhane The grassroots, volunteer-driven research collective for African-language natural language processing, founded around 2019.

Te Hiku Media The Māori radio station in northern New Zealand whose community-governed language AI tools provide this report’s model for sovereign language data.

Vulavula Lelapa AI’s commercial text and speech service covering isiZulu, isiXhosa, Sesotho, Afrikaans and English.

We Think Code_ The South African software development training institution whose Johannesburg campus hosts the AI incubator described in this report.



Introduction



Three independent research modalities underpin the findings in this investigative research paper.

The first is a basic, systematic literature review conducted specifically for this study, drawing on the United States Department of Education’s Education Resources Information Center (ERIC), the world’s largest repository of open education research.

The second is an assessment of the current AI model landscape, covering capability, cost, licensing and African-language performance, benchmarked against independent evaluations including IrokoBench and AfroBench.

The third methodology is a journalistic-style set of interviews with more than a dozen primary experts, conducted face-to-face or by video call over two months in mid-2026. The subjects were chosen across disciplines for their AI thought leadership, supplemented by interviews with students, teachers and professors.

Important limitations and assumptions

A number of caveats apply to this research paper. First, the AI model landscape dates quickly. The benchmark figures in chapter 3 were accurate as of early July 2026 and will be substantially out of date within months of publication. They establish broad tiers of capability and cost, not fine rankings, and the benchmarks themselves are less reliable than their precision suggests, but they are offered here as an illustration of both the current power and cost of AI systems.

Second, the published literature review carries a probable positive bias, since studies showing an effect are more likely to be published than those showing none, and the African evidence base remains very thin relative to the pace of AI adoption on the continent. The fact is that there is a severe shortage of good Africa-based AI research in education and training from which to draw concrete conclusions. Several studies establish association rather than causation, and the report says so.

Third, the experts quoted for this report were selected purposely for their standing in the field, not sampled to be representative across all demographics. Three serve on the Collective X board and Collective X commissioned this research. Professor Benjamin Rosman is a co-founder of Lelapa AI, one of the organisations this report identifies as a potential national asset. Those interests are disclosed where the relevant views appear.

Fourth, unemployment figures vary by definition, and both the narrow and expanded measures of

unemployment in South Africa appear here where the source requires it. They are noted as such.

Lastly, this report operates under the assumptions that South Africa's binding constraints for AI educational solutions are internet access, educator capacity and language data rather than AI talent or appetite for adoption. The evidence supports them, but these points remain assumptions.



1

Executive summary



Like every country, South Africa is currently undergoing one of the biggest technological shifts in human history, but it is carrying one of the world's heaviest unemployment burdens. The country's official, or narrow, unemployment rate stands at 32.5%, with 8.1 million people out of work and youth unemployment at 45.8%. At the same time, artificial intelligence (AI) is already reshaping the very industries, like call centre work, that have traditionally produced most of the country's new jobs over the last decade. This report explores how AI can help more South Africans participate in the economy.

This report is built on a systematic review of the published Africa-focused scientific literature, a hard-nosed assessment of the AI models now reaching South Africa, and crucially, direct interviews with some of the country's leading AI thinkers and young developers. The report reaches a number of conclusions that can help inform and reframe the national conversation about how AI can be successfully adopted in the country so it can be a job creator and enabler, rather than an economic burden.

South Africa is at once one of the worst-prepared countries in the world for the AI economy, while also being one of the best positioned to benefit from it, according to a number of experts. The risks of AI removing entry-level jobs in the South African economy are very real, but the experts believe that there is also the potential to get more people working using AI skills training. This report explores these opinions and various options.

The gap between those two futures is not talent, appetite or technology. It is a commitment to change that needs concrete educational interventions, broader internet access, trained facilitators, a focus on AI literacy, language data and most of all, understanding. All these barriers are potentially surmountable, at a cost far below what it will cost to course-correct if the country doesn't seize the opportunity to engage its AI educational future as a national priority.

If decision-makers are to understand this potential opportunity and the country is to steady its AI adoption course, there are a number of facts to take on board

immediately. First, the demand problem is already solved because South Africans of all ages are adopting AI tools on their own, ahead of any policy, curriculum or set institutional framework.

Second, some of the evidence outlined in this report shows that low-cost human-driven instruction over the phones and platforms people already own and use daily can potentially be an effective AI instruction method, if done ethically and in line with the country's laws and institutions.

Third, the binding constraint is not learner appetite but educator capacity, data affordability and the near-total absence of South African languages from the world's AI systems.

Fourth, no foreign model was built for South Africa, and the country should not pick an AI "side", rather it should practise the digital non-alignment it already espouses politically, while building the sovereign language assets no other nation can build for it.

Fifth, the institutions needed to do this already exist. They need to be scaled and connected, not reinvented. Every month of delay narrows the window of opportunity to seize AI in educational development, steering and blending it into the new skills and jobs that the experts consulted for this report believe can come from an AI educational rollout.

Solving for jobs

South Africa needs new jobs, and lots of them. If AI can be leveraged to deliver them, excellent, but if AI puts the few jobs in the country at risk, that would be a disaster. Beyond the headline unemployment figures, roughly eight million unemployed high school graduates currently have little to no hope of finding work under the current employment regime.

Rising educational attainment has not translated into jobs, and the education system feeding this pipeline is itself in crisis. Like thousands of schools across the country, Hantam High in the Northern Cape town of Calvinia operates without reliable power or internet, and more than 90% of its students have no internet access at all. Layer onto this some of the most expensive data costs in the world, no national AI training pathway, and no AI training for the teachers in these schools and the scale of the current challenge comes into better focus.

Why it matters

"The stakes are generational," warns Professor Benjamin Rosman, chair of South Africa's national AI advisory panel. "No economy can easily absorb even a five-to-seven percent uptick in unemployment over a few months; that is civilisational-stress territory if you are not prepared." He notes that South Africa's current 270,000 call centre workers likely sit directly in the path of current AI capabilities.

Importantly, there is a clock running on the country's one unique asset. South African languages barely exist in the world's AI training data, and the big AI players are already paying to acquire this kind of language wealth across Africa, on terms that see the value of this language data leave the continent permanently. What South Africa holds in abundance, fluency in its own languages, concentrated precisely among the people the formal economy underserves, is a national asset that the experts say must be secured now, or lost forever.

I feel like AI is one of us, given that the correct parameters are given to it. In the next five years, if digital inclusion takes its course properly, we could use AI to better our economy.

Akhona Mohulo (21, Lenasia), We Think Code_ Incubator, Johannesburg

Key findings

The report's three research modalities, the literature, the model landscape and the expert interviews, converge on the same picture. Demand for AI learning among African youth needs no stimulation and adoption is running ahead of policy across all three evidence bases. Some African trials show AI tutoring delivering meaningful learning gains at low cost, over infrastructure as simple as WhatsApp.

One of the main binding constraints across most studies referenced in this report is educator capacity and AI skills training, not learner willingness. Equity in AI educational outcomes is a design choice, not a by-product. The experts interviewed for this report believe that now is the time to design new educational systems to teach people new AI skills. Conventional

educational assessment has also been destabilised by AI, which makes a new, AI-native credential system an opportunity for South Africa.

For the models themselves, neither the American frontier systems nor the Chinese open-weight tier was built to benefit South Africans. The closed models are currently strongest but this may not continue long. They all process South African data offshore under a foreign national-security umbrella. The open-weight tier delivers near-frontier capability at a fraction of the cost and can be self-hosted, but degrades roughly twice as badly on African languages. That language gap is a data-scarcity problem, not a model-design one, which means South Africa can solve it, and then own these new language models for the benefit of the South African people.

Most significant evidence

A randomised controlled educational AI trial in Nigeria[1] found generative AI tutoring produced meaningful English-language gains when used as part of a regular curriculum, at low cost per student. A number of experts that we have interviewed for this report note that these learnings can be delivered on cellphone-based platforms already in the pockets of South African school students and leavers.

Interestingly, the Africa-centred, Global South literature review conducted by our team found scant evidence of AI being used to teach students and graduates AI skills or AI literacy itself successfully, but another more global literature review[2] called “AI Literacy in K-12 and Higher Education in the Wake of Generative AI: An Integrative Review”, conducted from 2020 to 2025, found that AI use itself, coupled strongly with in-person or over-the-internet educational initiatives, can help drive AI literacy in both K-12 students and high school graduates. The authors note that much more research is needed to confirm their results.

Probably the most important study outlining how AI can be used, alongside human instruction, to help teach students safe and responsible AI usage and prompting was published in March 2026 in the Proceedings of the AAAI Conference on Artificial Intelligence[3]. The fact that this is one of the only peer-reviewed and published studies into this usage of AI also makes it clear that much more research into using AI as a successful teaching tool needs to be conducted.

IrokoBench, the benchmark built by South Africa-based AI collective Masakhane, records a performance gap of up to 31 percentage points between English and African languages on the best AI models. The root cause is stark. IsiZulu, isiXhosa and Sesotho each have roughly 230 to 290 megabytes of monolingual text on the global internet against English’s nine terabytes, a more than ten-thousand-to-one difference, and a gap that South Africa can service and from which it can potentially derive employment.

Since the start of this research project, the leading Chinese open-weight models have come to score within single digits of, or better than, the best closed US-based frontier models on the main composite artificial intelligence index. They can do this at a drastically lower cost.

Te Hiku Media in New Zealand has proved there is a sovereign alternative to foreign-run language models. The Māori community produced 310 hours of training data in ten days and built its own speech tools under licences that keep that language model and its benefits within that community. The capability to duplicate this success is already here in Africa. When 490 competitors from 61 countries entered the 2025 challenge to compress Lelapa AI’s InkubaLM model, the winners, all African, cut its size, compute needs and cost by three quarters, without losing any performance.


From this research, the way forward for South Africa does not appear to be new institutions, rather, this report shows that it is scaling and connecting the ones that already exist, like SADiLaR’s state-funded data infrastructure, Masakhane’s benchmark ecosystem, Lelapa AI’s commercial proof that small local models can run on entry-level phones, and two decades of speech technology at the Council for Scientific and Industrial Research (CSIR). This proposed delivery must be a genuine public good, in partnership with business, but owned by the public, led by the government and delivered through the institutions people already trust. Some of the experts interviewed in this report outline in chapter 5 that this could be on the model of the country’s admirable HIV/AIDS response.

The burning questions

Some of the questions addressed in this report include: where can AI be used best? What opportunities are there for developing the skills that will drive job creation? Who teaches AI literacy in the country? Who can best

benefit from these skills in the near and long term?
What roles do government, industry and individuals have
to embrace to realise true broad-based development
from this powerful and rapidly evolving technology?
What are the risks of inaction?


Because AI is changing so rapidly, we must make every
effort to understand it, from what models to use to
what kinds of jobs and other employment outcomes
the country as a whole can benefit the most from. This
research report's direct expert interviews, the literature
review and the model evaluation all bring something to
this discussion. Together they create a rough baseline of
where South Africa's AI education space is today, and
where it can hopefully get.



**South Africans of
all ages are adopting
AI tools on their
own, ahead of any
policy, curriculum
or set institutional
framework.**

2

The state of the AI nation



The problem is jobs

South Africa's unemployment crisis is not a variation on a common developing-economy problem. The country ranks second worst in employment in the world out of the 103 economies the IMF tracks (2026 projections), behind only war-torn Sudan. Morocco, whose AI-in-education research features prominently in the literature reviewed here, carries an unemployment rate roughly two-fifths of South Africa's. Other economies South Africa is often measured against are in a different league altogether[4].


Brazil's unemployment rate is about a fifth of South Africa's, and both India and China sit at around a sixth, with China's rate of 5.1% close to the global median of 5.2%. Put another way, a South African worker is roughly six times as likely to be unemployed as a worker in China or India, which means every assumption imported from those countries' AI and skills strategies has to be stress-tested against a labour market with no real international parallel.

Rising educational attainment has not translated into jobs, since universities are producing more graduates each year even as job creation lags behind, with youth absorption into the labour market staying the lowest of any age group at just 10.1% for 15- to 24-year-olds in the first quarter of 2026. In total, 3.9 million young people aged 15 to 24, or 37.6% of that age group, are not in employment, education or training at all, showing a generation increasingly disconnected from both the classroom and the workplace[5].

The unemployment crisis and the successful upskilling of young learners and school leavers in AI literacy and training start with the national education system, according to the experts consulted for this report. Although the report itself is focused on how to upskill high school and post-schooling graduates, the country cannot chart a new course forward without talking about how these graduates got to where they are today.


Why the stakes are so high: the South African jobs crisis by the numbers

Unemployment rank, 103 economies tracked by the IMF, 2026




Unemployment rate, 2026

IMF projections. Each figure = 1 percentage point



Source for rank and rates: IMF World Economic Outlook, April 2026, unemployment rate, 2026 projections. imf.org/external/datamapper/LUR@WEO

Young people aged 15 to 24, Q1 2026



Youth figures: Statistics South Africa, Quarterly Labour Force Survey, Q1 2026 (released May 2026). statssa.gov.za

Hantam High School in South Africa's Northern Cape town of Calvinia is typical of many secondary schools in the country. Due to the twin problems of rolling blackouts and high power costs, the school does not have power all the time. There is also little to no internet on site at the school, so no way for students to ever experiment and learn with AI unless they have access at home.

According to the principal, more than 90% of the roughly 330 students do not have any internet access at all. If there is no power at the school, there is no internet, and in the hot summer months the students must be sent home because the classrooms are simply too hot, something that occurs at least monthly if not weekly.

Interview: Akhona Mohulo

With little power and no internet access it seems like a far-away dream for young people to be learning about AI, developing much-needed AI literacy and then taking those skills and getting a job in the workforce. Twenty-one-year-old Akhona Mohulo from Johannesburg notes that he was lucky to have the internet at home where he taught himself the AI skills he uses today. “We did not have AI whatsoever,” he said in a recent interview. “It was mainly just your standard subjects, and that was it. Everything AI? I had to learn that in my spare time.”

He is now getting further training through an AI incubator funded by the Industrial Development Corporation of South Africa (IDC) at the We Think Code_ campus in Rosebank, Johannesburg. The IDC is a state-owned development finance institution that funds and supports industrial and economic development in South Africa and elsewhere in Africa.

Even at high-end private schools in the big commercial centres like Cape Town and Johannesburg there is little to no AI training for students who will be leaving high school and there is little to no AI literacy training for teachers either, which is crucial given that they are the ones who are essentially the only possible AI training source for South Africa’s students.

Now that he knows how to use it and learn with it, Mohulo says he uses AI all the time. “I use AI every day. From sending an email, I have bad grammar, as much as I’d like to think I’m well spoken, so AI really helps me with that, as well as with keeping track of tasks. I am not even going to mention code debugging; it really helps with that. And especially studying, because when it comes to studying, AI can really break down problems for you and help you get a better sense of understanding.”



So, the problems are real and they are not going to go away without a real and concerted effort. But what this report uncovers is that experts and students, like Mohulo, all believe the issues with AI in South African education are surmountable, and South Africa has an almost unique opportunity to be a global leader, if the right actions are taken as soon as possible, like building AI skills and literacy through TVET or vocational training systems, and getting a nationally recognised accreditation or certificate for it, something which employers say they already prefer.

Developing AI skills as a national priority


According to the AI Index from the Stanford Institute for Human-Centered Artificial Intelligence (HAI), South Africa is also lagging badly behind its peers and major trading partners in AI skills development in the workforce at large. India ranks first globally in AI skill

penetration with a score of 2.8, ahead of the United States at 2.2 and Germany at 1.9. On AI hiring growth in 2024, India led at 33.4% year-on-year, followed by Brazil at 30.8%. By contrast, no African nation made the top 15 for AI skills penetration in 2022. South Africa, which had ranked 13th back in 2019, has dropped out of that group. The story here is that on actual AI skills, both India and Brazil are among the global front-runners while South Africa has slipped down the table over time.

Self-reported AI skill penetration, 2015–2023 (Stanford University HAI AI Index, 2024)

	India	2.8% of workers
	United States	2.2% of workers
	Germany	1.9% of workers
	South Africa	0.6% of workers

AI hiring growth, 2024, year-on-year



Source: Stanford University HAI AI Index, as cited in this report.

The South African employment crisis is concurrent with a problem which all countries are struggling with, the lack of forward planning in how to create new jobs and integrate AI into the current workforce through educational initiatives.

But South Africa also has another compounding issue. The country has an unequal internet penetration rate, coupled with expensive data costs when compared to the rest of the world. South Africa’s internet penetration stood at 78.9% of the population at the start of 2025, placing it a few points above the global average of 74% but well short of the 94% typical of high-income countries, and behind comparator economies such as Morocco (over 92%) and Brazil (mid-to-high 80s).

So South Africa is more connected than most of the world on paper, but affordability relative to income and a deeply unresolved rural-urban internet access divide means meaningful internet access remains highly uneven across the country. In effect, much of the rural population simply can’t afford to get online to use AI, or any kind of online learning tools at all[5,6].

Given these issues, it is fair to say that the country is almost uniquely placed to benefit from a national AI skills-development plan that leverages AI education, if it can increase internet access. Dr Rodney Manyike and his colleagues from the South African government’s Human Resource Development Council noted in a recent interview that this may enable it to leapfrog ahead of comparative economies and build a uniquely South African AI-driven skills revolution that can help revitalise the economy.

When asked whether AI will destroy the limited jobs South Africa has now or open up new opportunities, he noted that “both scenarios are plausible,” and that “AI is neither inherently job-destroying nor job-creating; its impact will depend largely on the policy choices South Africa makes today. The deciding factor is not the technology itself but the speed and seriousness of the country’s investment in education, reskilling, workforce transition programmes, and the policy choices and positioning that will accompany AI adoption.”

One thing is clear: no nation has the answer to these problems. The experts like Dr Manyike who agreed to be interviewed for this report all believe that AI can potentially unlock new jobs, help build new skills through AI literacy programmes for adult learners, reskill or upskill current workers, and help provide many other education opportunities, but only if it is properly harnessed, developed and delivered where it is needed the most. “Ultimately,” Dr Manyike noted, “South Africa’s ability to harness AI as a tool that augments human capabilities, rather than simply replacing labour, will determine whether AI exacerbates unemployment or becomes part of the solution.”

This report leverages these experts’ experience in order to outline a potentially better AI adoption path for the country, one that travels towards increased employment and job resilience using teacher and vocational training, as well as local language model development. All of the interviews drive this report from the experts’ perspective, leaning on the people from institutions, schools and businesses who seem to have the most informed opinions about this rapidly changing technological leap. They have shared their knowledge as they also seek to understand and adopt their own AI lessons in businesses, government, their classrooms and beyond.

A solid AI foundation

South Africa enjoys a unique place in the evolving world of AI skills development for job placement. It has a strong constitution, with systems and laws already in place that can address many of the systemic problems of AI outlined above, including a strong public media environment and free online learning opportunities. It has a young population who are extremely motivated to learn, integrate new ideas, and to develop into a new generation far more comfortable and prepared than most of their parents for living an AI-integrated life.

The country also has notable centres of excellence, where some of the top AI minds in the world have been putting their efforts into helping the country adapt to this ambiguous future in an ambitious way, and start solving the problems that already bedevil the AI rollout around the world. It also has a current non-aligned status within the global political power structure, allowing it to harness the best technology that is out there, regardless of who controls it.

This report highlights these efforts, helping everyone who is interested in understanding where the country is at, and where it needs to get, as it seeks to forge a successful path through what is unknown territory. The team behind this report has endeavoured to include both the big voices in AI, and lesser-heard opinions.

Education is key

Included here are the voices of the nation’s youth. Some of these young people are already building part of the AI education and language systems of tomorrow in startup businesses. The experts believe that this kind of entrepreneurship is crucial if the country is going to address the eight million high school graduates who currently have little to no hope of getting a job under today’s employment regime.

University of the Witwatersrand’s Professor Benjamin Rosman is currently the chair of South Africa’s national AI advisory panel. He is one of Time Magazine’s 100 most influential people in AI globally, and he is also a co-founder and the chief science officer of Lelapa AI, the Johannesburg language-technology laboratory profiled later in this report. He believes that despite the challenges, South Africa can rise to this vitally important occasion. “No economy can easily absorb even a five-to-seven percent uptick in unemployment over a few months; that is civilisational-stress territory if you are not prepared.”

But he also believes that “South Africa is different. We have more sensitivity to employment, mass layoffs are harder here, and there is genuine corporate care for the country.”

He also thinks that because there will be new jobs and skills in AI, South Africa could truly benefit from the technology because of the country’s structural makeup, especially if it focuses on AI education, helping it leapfrog ahead of other more economically entrenched nations. “Our complex economy, with its informal

sectors, has more nooks and crannies for people and AI to fill. I am not confident it will magically work out; it needs steering. There are infinite possible futures, most more negative than positive, but the positive ones exist.”

Dr Rodney Manyike from the South African government’s Human Resource Development Council and his colleague Samantha Senosha agree with Professor Rosman, but take his goals further, noting that “a realistic ambition over the next three to five years would be to position South Africa within the top 40–50 countries globally on AI readiness while strengthening its position as Africa’s leading AI innovation hub.”

Dr Manyike’s team believes that “achieving this will require finalising and implementing a co-ordinated national AI strategy, investing in digital infrastructure and advanced computing capacity, expanding AI education and skills development, supporting research and commercialisation, and fostering stronger collaboration between government, academia, industry, and civil society.” What kind of AI rollout can help support these ambitions?

A human in the loop

Monisha Prem is a lawyer, AI researcher and educator who works within AI governance and skills development through her legal firm in Johannesburg, South Africa. Like almost everyone interviewed for this report, she believes that AI will never fully replace humans in the workforce. “From the day the algorithm is being written until it is deployed and used by the end user and continuously being used, you need human intervention. The more checkpoints you have, the closer you are to controlling it. The closer you are to controlling it means that this fast-moving train is not going off the rails anytime soon.”

She believes that as long as leaders in the country, from education to government and the commercial sector, are working like this to integrate people within the AI development and operational chain and de-risk new employment solutions, AI can and will be a successful economic driver of skills development in South Africa. “If you have those checkpoints in place by humans and/or other machines that are being monitored by humans, irrespective of where downstream, the more of those you have, the more they can be regulated.”

Expert opinion and analysis


This report outlines and includes the opinions of fourteen primary experts who were interviewed over two months in mid-2026. These experts were chosen from a broad range of disciplines, but all were included because of their AI thought leadership, their direct work on the topic or their research and development of the technology itself. By conducting interviews either face-to-face or on video call, the report is able to bring this much-needed first-person expert perspective on this subject. Students, teachers and professors were also interviewed in order to gain their knowledge on AI and the paths to better AI literacy and skills development.

What is clear from these three research modalities (literature, model evaluation and first-person interviews) is that there is a growing consensus among the science, the officials and experts that there is a path to successfully integrate AI within the South African educational and job-training context (that ultimately impacts jobs), with much agreement from all three about what can constitute a successful AI skills training system that South Africa can scale urgently.

The experts interviewed for this report say that with good planning, public institutions and outcomes-based public-private partnerships, South Africa may be perfectly placed to benefit from AI. To do so it must incorporate the lessons of the country’s past, harness today’s working AI institutions, develop and nurture new thought leaders, and integrate these people and systems into a new AI educational architecture.

“Failing to adopt a co-ordinated and inclusive AI strategy is likely to exacerbate South Africa’s already significant digital divide,” explained Dr Manyike. He and his colleague Samantha Senosha noted that new AI educational interventions are crucial for the country, but people must have basic internet access too. “Without deliberate public investment and targeted interventions to expand digital infrastructure, improve affordable connectivity, strengthen digital and AI literacy, and build inclusive education and training pathways, the benefits of AI will remain concentrated among a relatively small segment of society.”

3 Why two models? The context behind the comparison



There are plenty of different AI systems in the world, some well-known and some not. South Africa is part of the BRICS system, which of course includes China, and in the few short months since the reporting for this project began, Chinese open-source models, which will be referred to here by the term “open-weight”, are already gaining parity or leading the more well-known “frontier” commercial models from the USA like OpenAI’s ChatGPT and Anthropic’s Claude. Open-weight models are leaner, can be hosted locally and some of them are already finding traction within the South African educational environment. Others are already being harnessed locally to provide a homegrown South African language model approach to AI.

So what kind of AI should South Africa be looking at adopting in order to seize this remarkable technology? South Africa’s choice between American and Chinese artificial intelligence systems is really about two theories of how the technology gets built, financed and controlled. These two distinct economic and infrastructural visions

are competing for adoption, but neither was designed with South Africa in mind. Evaluating the manufacturers’ claims requires understanding where each model comes from, what it can and cannot do, and if it can be used to achieve better education, more jobs and a more AI-proof workforce without compromising national and personal security.

The Goldilocks zone

Among the experts interviewed for this report, many believe that South Africa needs to embrace an “all of the above” approach to AI integration, including Microsoft South Africa CEO Vukani Mngxati, who is also on the board of Collective X, which commissioned this report. “You are going to see all sorts of technology here, and that is fine. What matters is how we use the tools, how we skill our people, and how we reimagine our country with AI in mind. Get that right and South Africa will be okay.”

Lawyer and specialist in AI governance Monisha Prem agreed and noted that South Africa is in a good place to

regulate AI integration and is already governed by the country's current legal framework. She says that where there are already potential conflicts with current South African law, there are already legal remedies available. For instance, she outlined in a recent interview how many AI companies are already in breach of South African law because they "may not transfer personal data without the express consent of the data subject if it's being shared, particularly cross-border." She says it is unlikely anyone could get into trouble for this because right now "controlling that is almost impossible."

She also explained how South African laws already protect citizens from these kinds of potential breaches, placing the country at the forefront of legal guardrails when it comes to AI. "There is recourse for managing most of these risks, and POPIA (South Africa's Protection of Personal Information Act) is one of them. POPIA says very clearly, in section 71, that you may not rely on a decision which has been automated. That speaks directly to artificial intelligence and machine learning. If I am dismissing an employee, I need to justify that I made that decision, and not the machine. If I can't, I'm in breach of section 71, and I'm in breach of the Labour Relations Act, because I'm unable to show due process as to how I got to that decision. What I'm trying to illustrate is that we have a sophisticated network of laws in South Africa that we can still rely on to protect us. It may be slower, it doesn't move at the pace of AI, and it may be traditional, in that you still have to look at interdicts and approaching the courts, but the laws are in place, and we must not forget that."

There are also other serious questions about who controls these models, which include concerns around data, algorithmic and processing sovereignty.

Keeping these real concerns in mind amid the urgency of this increasingly hurried AI rollout, what are the current options for AI integration into South Africa's much-needed jobs and skills generation?

The Western hyperscaler model

The Western frontier of AI development in the United States, which includes OpenAI's GPT line, Anthropic's Claude and Google's Gemini, is the product of an investment cycle without precedent. Cumulative AI infrastructure spending by the major hyperscalers, the handful of technology firms operating data centres at global scale, has passed a trillion US dollars in two years alone[8]. The four largest, Amazon, Alphabet,

Microsoft and Meta, are estimated to burn through roughly \$725 billion in capital expenditure in 2026, up 77% on 2025's record \$410 billion, with Goldman Sachs projecting \$5.3 trillion from the four between 2025 and 2030. That capital has flowed despite the absence of any demonstrable market for these large language models (LLMs), which is not subsidised by state or commercial institutions.

These current American frontier AI systems are priced well below what it costs to run them[9]. They are loss leaders, meaning that they are sustained by investor capital rather than customer revenue. The industry's own narrative, of imminent machine sentience, existential risk and mass job displacement, cannot indefinitely sustain the expectation of future financial returns on which the current investment case for more capital into these models depends. This means many will likely fail in the medium to long term, unless they are either supported by state-sponsored money or they scale back their ambitions, or a combination of both.

Any prospective client of an LLM, whether it is Chinese or American, should be asking how long current low pricing can last, and what replaces these models should they fail. What is at stake is the future of the vitally important data that they manage once the current round of financial subsidy ends. But data loss is only one danger of adopting these models within the South African context. Chinese models are indeed cheaper, but there is no guarantee they will remain open-weight, or cheap.

There is also the question of where the data goes to be stored and parsed, and whether that is even legal under South African law. Under the American hyperscaler model, a client's information is processed on servers controlled offshore in places like California and Indiana. The companies' terms of service have stated that OpenAI and Anthropic staff with appropriate access can, in principle, view that data. For law firms, government health departments or universities holding student records, this is a serious material and legal risk because it raises direct legal compliance questions under the Protection of Personal Information Act (POPIA), South Africa's data protection law[10].

What goes largely unstated is the reason the funding for these models continues, despite the arithmetic. That reason is national security. The United States has imposed export controls on advanced AI chips[11], drafted with the active involvement of the major Silicon Valley firms. Also, the leading American laboratories

are not standalone commercial actors. Anthropic's chief executive has said publicly that the company's technology will be used against US adversaries[12], and the firm, like all of the major hyperscalers, holds active Pentagon contracts and, notably, has already been used to target civilian infrastructure in Iran.

The point is not to single out Anthropic, which has shown some limited insistence on independence. It is to establish the pattern that Western frontier AI laboratories are institutionally embedded in the US national-security apparatus in ways that shape their strategic behaviour and their willingness to absorb losses.

Other foreign AI companies are more deeply entrenched in the surveillance and security state still[13], with associated free-speech issues that may also run afoul of South African law. The fact is that the current AI leaders are all a product of concentrated capital aligned with state power, something that South African decision-makers should be very aware of.

The Chinese open-weight tier

China and the United States are the world's two main centres of AI research. While other countries have produced large language models, like France's Mistral and Germany's Aleph Alpha, it is China that remains the only meaningful competitor to Silicon Valley. The current Chinese models that now equal or rival American frontier capability include DeepSeek, Kimi, GLM and MiMo. These systems emerged under a different set of social and political pressures.

The American chip export controls were designed by American policymakers to constrain Chinese AI development and secure American dominance, but in practice they spurred innovation[11]. After being cut off from the most advanced NVIDIA chip hardware that is vital to the US effort to develop AI, Chinese firms developed techniques that reduced the cost of training models and the chips required to run them. The result has been a succession of open-source models with near-frontier or better-than-frontier capability at a fraction of the energy and financial cost of their American peers.

DeepSeek, the most visible example, is financed by a Chinese quantitative-trading firm with its strategy prioritising market disruption over immediate revenue[14]. Its models are released under permissive

open-weight licences like MIT, a standard open-source licence permitting free reuse. Unlike US frontier AI models, DeepSeek removed the financial and technical barriers to uptake like having to own a credit card, paying for upgrades, or leasing the application programming interface (API) key.

Zhipu's GLM-5.1 went further: it was trained entirely on roughly 100,000 Huawei Ascend chips, with no NVIDIA hardware at all[15]. Zhipu has been on Washington's trade blacklist since January 2025. Frontier-level training is now demonstrably possible on a non-Western compute stack, which is already disrupting the far more expensive Western AI systems.

The practical result is a tier of models that is near-free at the point of use, because the strategic incentives of the organisations that built them differ from those of the Western hyperscalers.

None of this makes the Chinese tier independent of the Chinese state, or makes for a more secure operating environment. DeepSeek is privately owned but it operates within China's AI regulatory framework. Zhipu, a spinout from Tsinghua University, has deeper state links[15], which is why it sits on the US Entity List. The Huawei chips on which GLM-5.1 was trained are likewise a product of state industrial policy. Chinese models, like Western ones, carry the imprint of the system that produced them, and that system likely has complete access to whatever material a user might access through these AI products.

The difference is in how they operate. Open-weight AI systems can run by themselves on private computer architecture, and once downloaded, do not, as far as is known, phone home. Self-hosting breaks the data link to the provider's jurisdiction, whether that jurisdiction is Washington or Beijing. The openness is not guaranteed to last, however.

Reuters reported on 7 July 2026 that Chinese authorities had met with technology companies about restricting the availability of open-weight models overseas[16]. Given these circumstances, relying on China to keep releasing increasingly capable open-weight models is not a viable strategy for any country.

The comparison that follows is focused on measured capability, on African-language performance, and on cost and access.

With measured model efficacy, the benchmarks are less trustworthy than they look. The leaderboard shifts with every AI product release cycle, which is usually measured in weeks. The figures below are accurate as of early July 2026. What matters is the broad picture of capability and what the trends suggest, which is that Chinese open-weight models already have parity with their American peers, as well as being far cheaper to run.

The composite picture

On Artificial Analysis's Intelligence Index (version 4.1[7], early July 2026), the top closed models are Anthropic's Claude Fable 5, OpenAI's GPT-5.6 Sol and Claude Opus 4.8, with the top ten models all scoring within roughly nine points of one another. The frontier is genuinely close. No single model wins every category, and the top performer changes with almost every release.

The Chinese tier is within single digits of the frontier, at a fraction of the cost

This is the most consequential measured finding for the South African question. The leading open-weight models like Zhipu's GLM-5.2, Kimi K2.6 and DeepSeek V4 Pro, among other newer ones, score around 51 on the same Intelligence Index where the top closed frontier model scores roughly 60. A single-digit gap means that the closed models do not have an insurmountable advantage, and as of publication they may well have gained parity or passed the US frontier models.

DeepSeek V4 is an open model launched in April 2026[17] with an update following in July. It is under an MIT licence with a one-million-token context window, meaning it can process very long documents in a single pass. According to DeepSeek's published usage rates, the flagship V4 Pro costs about \$1.60 per million input tokens and \$3.20 per million output tokens. The comparable rates for OpenAI's GPT-5.5 and Claude Opus 4.8 are roughly \$5 for input and \$25 to \$30 for output. The specific figures will change but the scale of the gap is robust. Put simply, it is way cheaper to use Chinese models, and likely to remain that way for some time.

Zhipu's GLM line shows the same pattern at the high end, and adds a compute-sovereignty dimension. GLM-5.1 (March 2026) scored 58.4% on the SWE-Bench Pro coding benchmark, narrowly ahead of OpenAI's GPT-5.4. Its successor GLM-5.2 (June 2026) is now the leading open-weight model, scoring 62.1% on

SWE-Bench Pro and 91.2% on GPQA Diamond, within three points of Claude Opus 4.8 on the most discriminating reasoning benchmark in use.

Western closed models still hold the very top on the hardest reasoning and agentic (agent-driven) benchmarks. But the open-weight tier now sits within single digits, and is gaining steadily on the best closed model on the main composite index. They deliver this competitiveness at a small fraction of the cost, increasingly without Western chips, and under licences that permit self-hosting.

This means that they can run natively on a user's machine, or in the case of South Africa, on a locally hosted and controlled server system. The development trajectory across these new generations of AI open-weight models indicates fast-increasing parity. One caveat is that Alibaba's frontier Qwen line went closed in May 2026, which is a reminder that openness is a strategic choice its makers can reverse, unless of course one is harnessed and developed locally.

The critical African-language benchmarks

All current top-tier AI models are LLMs, which stands for large language models. Although, as experts have noted, some models need bigger language datasets than others, all of them need language datasets, which in many ways is what drives AI's impressive abilities. What happens with LLMs when they encounter users who are not English speakers is that because the languages they speak are under-represented on the internet, the models perform a lot less well for these users. In short, models do not perform well in languages that their datasets are not based on.

Given that it is largely American and Chinese companies who created the current AI models for their own purposes, how do the current set of AI models perform in a South African context? It turns out that, according to experts like twenty-five-year-old AI education entrepreneur Lee Vandred Napier, not great.

Interview: Lee Vandre Napier

AI cannot perform language tasks in every one of South Africa's eleven official spoken languages because it is not trained on large datasets of these languages. His team is working to solve this problem so they can bring AI-assisted learning into South African classrooms.

His new startup to fix this problem is being incubated at the We Think Code_ campus in northern Johannesburg, South Africa. He explains that this new business is currently focused on teaching English literacy using AI, "because South Africa has a literacy crisis. We looked at the statistics behind it, and our diagnosis of the problem was that it starts around Grade 3, with mother-tongue instruction."

He believes that a child should read for meaning in their home language first, and this will improve vital literacy deficits in English and other languages later. But he also notes that South Africa doesn't have big enough datasets in young people's own languages for AI to be useful, so to solve that problem they are now training their own language models of South Africa's languages. "We hope that by doing that we'll improve outcomes for children in general."

The biggest problem his team of seven developers is facing is this lack of good datasets of African languages to train their own open-source model for their educational startup. There are simply not enough digitised versions of African languages for entrepreneurs like Napier to develop solutions for learning on, so they are waiting for South African AI groups like Lelapa AI, which is outlined later in this chapter, to help develop this resource.

He explains why this constraint has made it impossible to incorporate all eleven official spoken languages of the country. "So we have four languages currently available," he explained in a recent interview. "That would be Zulu, and we



currently support Xhosa as well, plus Afrikaans and English. For the other languages, we're partly building those databases ourselves and partly waiting for organisations like Lelapa AI, who are working on South African languages, because a team of seven can't do it all."

Since it is clear that home-language instruction for education is so crucial to overcome South Africa's AI education divide, the South African IrokoBench language dataset offers a rare glimpse into what is now available for home-language learning in the country.

The language state of play in South Africa

IrokoBench, developed by the South African non-profit Masakhane Research Foundation (MRF)[18], evaluates models on natural-language inference, knowledge and mathematics across 17 African languages. It found a substantive performance gap between African languages and English of up to 31 percentage points on the best AI models they tested.

For their latest research, that model was OpenAI’s GPT-4o, averaging 59% across tasks, and the languages that scored worst were those with the least monolingual text on the global internet, which makes sense given that the greater the amount of language training data, the better these models seem to perform.

IrokoBench covers four Southern African languages relevant here, including three of South Africa’s official languages, isiZulu, isiXhosa and Sesotho, plus chiShona, which is Zimbabwean but regionally relevant. The verified per-language scores for GPT-4o, the best-performing model prompted in these languages are:

Language	AfriMMLU (knowledge)	AfriMGSM (mathematics)	AfriXNLI (inference)
English (reference)	87.4	84.0	89.2
isiZulu	68.2	57.2	67.5
isiXhosa	70.2	52.4	72.0
Sesotho	67.4	60.8	71.8
chiShona	67.6	58.4	71.3

A further complication affects any plan to adapt an open model, as Napier and his startup team already pointed out. For the open models, machine-translating the question into English before answering beats local language prompting across the board. This is likely because again, there is a serious lack of datasets for these models to utilise these languages.

The reason is that the open models remain heavily English-centric because the models are not working natively in African languages. This is only a concrete limitation for South Africa now, but as explored later in the paper, this opens up the possibility of the country building its own language models, like the ones Napier and his team are trying to build in their startup.


AfroBench is the first comprehensive LLM benchmark for African languages, studying natural language processing (NLP)[19]. The group is also evaluating proprietary and open models across 15 NLP tasks, 22 datasets, and 64 languages. The team there published their results in 2025 in the Findings of the Association for Computational Linguistics (ACL). It covered the same four Southern African languages as IrokoBench. Across all tasks,

GPT-4o (59.6 average) and Gemini-1.5-Pro (58.5) beat the best open model, Gemma 2 27B (47.7), by more than 10 percentage points, again probably due to better English-language translation, as noted above.

Performance correlates closely with how much text exists in a language on the global internet or other training data. The monolingual web text available for these four Southern African languages runs to a remarkably small 230 to 290 megabytes each.

Swahili, the best-performing African language across AI evaluation studies, has 2.4 gigabytes. English has roughly nine terabytes. This is a more than ten-thousand-to-one difference, so there needs to be more African language data for better LLM capabilities. The cheap, self-hostable open-weight tier that wins on cost and sovereignty loses most on indigenous languages, according to AfroBench and IrokoBench’s findings.

[SADiLaR] runs
a digitisation
programme creating
digital text, speech
and multi-modal
resources for all
12 official South
African languages,
including the
recently added
South African
Sign Language.



The 2025 paper “Where Are We? Evaluating LLM Performance on African Languages” by Ife Adebara *et al.*[20], published at the 63rd Annual Meeting of the Association for Computational Linguistics (ACL), provides a third independent evaluation and confirms this pattern. This peer-reviewed evidence shows Western closed models leading and open models doing measurably worse on African languages, because the differentiator is training-data representation, not AI model architecture or price. The claim that the open-weight tier is adequate holds for English-medium reasoning tasks. The research clearly shows that it does not automatically transfer to a successful isiZulu or Sepedi AI language model deployment.

Because the gap in a model’s ability to parse local language inputs is a data-scarcity problem, it can be solved by creating bigger language datasets through open-source systems like those being created by developers like Napier and his team, as well as other larger groups busy with this task in South Africa.

Small groups like Napier’s take open-source models that can run natively on local servers, sourced from online AI model warehouses and then integrate them into a working application that can be run on smartphones.

“These [free and open-weight] models already exist online from places like Hugging Face, which is an open-source platform where you can get an AI model and train it on whatever dataset you can acquire to increase its accuracy,” Napier explains. “We get models from there and package them into the app. Essentially it’s transcription. Once speech has been transcribed, we can go back and forth, from Zulu to English or straight Zulu. We can send that to a model, get information back, and then use text-to-speech to explain it better to the user.”

To understand better how this kind of open-weight model can be developed and implemented for African and other indigenous languages in order to solve this critical online language-data shortage problem, there is an excellent, illustrative open-source local-language LLM case study that is already up and running.

The Te Hiku Media case study

Te Hiku Media is a non-profit Māori radio station based in Kaitiaki, in northern New Zealand. Over the past decade it has built a working suite of Māori-language AI tools: speech recognition, automated transcription, and pronunciation feedback for learners.

Te Hiku began with an open-source speech recognition model from the Mozilla Foundation. The community then produced the training data itself. In ten days, roughly 2,500 people submitted 200,000 recordings, yielding 310 hours of paired speech and text, enough to fine-tune the new language model to 86% accuracy. The data was contributed under a relationship of trust, with Te Hiku acting on the community’s behalf, rather than being scraped from the internet like most other English-language LLM datasets.

Te Hiku developed a licence that permits access to the data only on condition that Māori values are respected, the bounds of consent are maintained, and any benefits derived flow back to the Māori community. The tools are accessible through Te Hiku’s own platform but are not downloadable for unrestricted reuse. This allows the community to benefit from this language “sovereignty” which they developed and control.

Chief executive Peter-Lucas Jones explained that creating their own language models for Māori was not only about making AI more useful for the community but also about claiming back their own heritage and securing it from further exploitation by global AI. That community was concerned that the powerful foreign AI companies seemed to copy the inequality that drove the colonialism that nearly wiped out their nation. “Our data would be used by the very people that beat that language out of our mouths to sell it back to us as a service,” Jones says.

Keoni Mahelona, Te Hiku’s chief technology officer, says, “It’s just like taking our land and selling it back to us.”[21]

Language resources built under sovereign governance can serve as a base for regional adaptation, rather than waiting for an English-centric model to trickle down, or a big AI company to sell services that could be part of a new publicly funded AI governance and equality-based solution.

Te Hiku is a third category of model which is community-governed adaptation built on an open-weight foundation. It demonstrates that the sharp performance drop on under-resourced languages, worse for open models, is not a permanent feature of the landscape. It is a product of data scarcity and data governance, both of which are policy shortcomings, not technical inevitabilities. They are also a potential opportunity to secure sovereignty and create jobs.

South Africa already has these critical AI assets

South African groups are already nurturing and growing these new local LLMs. Lelapa AI is a Johannesburg-based research and product laboratory founded in 2022 by six people drawn from South African academia and industry, including chief executive Pelonomi Moiloa and Jade Abbott, who also co-founded Masakhane. It has two main products. Vulavula is a commercial text and speech service, used by South African firms, that transcribes, translates and analyses isiZulu, isiXhosa, Sesotho, Afrikaans and English.

Lelapa AI's InkubaLM, the “dung beetle language model”, is a deliberately small 400-million-parameter multilingual model trained on 2.4 billion tokens, like word fragments, covering isiZulu, isiXhosa, Yoruba, Hausa, Swahili, English and French. It is open-source, self-hostable, and runs on entry-level smartphones without cloud dependency.

Masakhane is a grassroots, distributed, community-driven research collective developing NLP on African languages. Masakhane helped build the benchmark infrastructure by which African-language AI is now judged globally, including IrokoBench, cited above as one of the authoritative sources on the African language gaps. The collective currently has no permanent institutional home, no dedicated funding line, and depends on volunteer labour and small grants.

SADiLaR, the South African Centre for Digital Language Resources, was founded in 2016 and

is hosted at North-West University[22]. It is funded by South Africa's Department of Science, Technology and Innovation as part of the South African Research Infrastructure Roadmap (SARIR). It runs a digitisation programme creating digital text, speech and multi-modal resources for all 12 official South African languages, including the recently added South African Sign Language.

The output is hosted in the Language Resource Catalogue and Language Resource Index, freely available under research licences, with reusable resources also licensed for commercial applications in local languages. SADiLaR is under-resourced relative to its potential, but it is the closest thing South Africa has to publicly developed sovereign language-data infrastructure. Any serious policy response to the African-language AI gap, crucial as noted above to any real AI educational effort, likely runs through scaling SADiLaR and supporting groups like Masakhane, not through creating a new institution.

Another critical South African language resource is also already owned by the public. For two decades the Council for Scientific and Industrial Research (CSIR) and its Meraka Institute Human Language Technologies Research Group have been the primary academic home of South African speech-technology work[23]. A domestic AI strategy that begins with publicly owned speech technology, such as public-service voice interfaces, court transcription and medical dictation in indigenous languages, has more to build on than one that is only now beginning to develop the data for text-based large language models.

The Non-Aligned Movement can inform South Africa's response

The Non-Aligned Movement (NAM) emerged from the 1955 Bandung Conference. It was formally founded at Belgrade in 1961 by Nasser, Tito, Nehru, Nkrumah and Sukarno. Its core principles were mutual respect for sovereignty and territorial integrity, non-aggression, non-interference in internal affairs, equality and mutual benefit, and peaceful coexistence. These were a specific response to the Cold War's demand that every state align with a superpower.

At the time, exiled elements of the current South African government were part of this movement, and in the current global geopolitical framework, South Africa is still non-aligned, particularly with regard to technology adoption.

Fidel Castro summarised the Non-Aligned Movement's purpose in the 1979 Havana Declaration, adopted at the summit Castro hosted. It was to ensure "the national independence, sovereignty, territorial integrity and security of non-aligned countries."

The dominant AI players around the globe want South Africa locked into a contract with them, and the pressure to choose is considerable. The Te Hiku case shows how the principal South African AI efforts outlined above could map onto the current AI educational development landscape. Given the current unstable global AI development environment, South Africa's likely best path is to focus on developing its own language models, and other AI educational systems, in order to maintain its data, language and algorithmic sovereignty, as much as possible.

Value in non-alignment

The published literature noted in this chapter comparing American frontier systems and Chinese open-weight models shows that neither was built to benefit South Africans. The closed frontier models currently deliver the strongest capability, but they process data offshore and answer to a national-security establishment that is not South Africa's.

The Chinese open-weight tier delivers near-frontier capability at a fraction of the cost and can be self-hosted in South Africa, but there is no guarantee that future releases will remain open at all.

If the country can develop its own language datasets, South Africa's advantages are immediate. As already noted, the components of a sovereign language-AI capability already exist. They include the two decades of speech and multi-language databases and pronunciation dictionaries at the CSIR and North-West University, the state-funded data infrastructure of SADIaR, and the benchmark and dataset ecosystem.

Lastly, in Lelapa AI, there is commercial proof that small, open, locally governed models can serve South African clients on entry-level phones, which everyone already uses and the experts agree will have to be the de facto AI delivery system for most South Africans. These four assets are under-connected and under-funded, but they provide a uniquely South African pathway to sovereign AI adoption.

The way forward is to scale and connect institutions that already exist.

The skill these models need most is fluency in South African languages. This skill is held in abundance by precisely the people who the formal economy currently underserves, like rural speakers, township speakers, speakers without matric, and above all the elders, whose vocabulary is the deepest. This knowledge of South Africa's languages may well hold the key that can help unlock a sovereign response to educationally relevant AI.

Because of these factors, a national programme to develop local language datasets is a race against time. Every year of delay developing these language sets means that South Africa is closer to losing its language wealth.

The big AI players know language wealth is worth real money and are already paying people to acquire it.

India has demonstrated that this kind of AI language work can be organised and is already getting that country's own indigenous languages integrated into large language models. Karyā, which began as a Microsoft Research project and now supplies language data to Google DeepMind and Microsoft, has paid more than fifty thousand workers for over fifty million tasks, at wages many multiples of the standard piecework rate in that country.

Unlike other projects of this sort, in some cases India has created a way for people to also get some ongoing royalties that pay local language contributors again when their data is relicensed, which helps mitigate the danger that this kind of work becomes gig work that is simply sold on to the highest bidder from a foreign country. This distributive intervention is one of the most important components of a viable approach and should not be sidelined or diminished.

An alternative language data-acquisition model also operates on this continent, and it should be a warning. Annotation platforms developing African language datasets for large foreign AI companies pay one to three dollars an hour for African language work whose value leaves permanently, on terms the workers never see. Rarely mentioned in the promotional material of these efforts are Africans like the Kenyan AI data trainers, who are often paid extremely poorly. The language data leaves Africa forever, becoming an asset on a foreign company's balance sheet.

A more sovereign path forward

"The issue is broader than data sovereignty," notes Collective X's board chairman and Cassava CEO Ziaad Suleman. "It is infrastructure, regulatory and economic sovereignty, and all of it leads to one thing: control. The real question is how much risk is palatable for the environment you operate in? Companies sitting in Amazon or Microsoft clouds have a problem, because US law obligates any US company, wherever it operates, to hand data back to the US. Even a data centre in Cape Town or Johannesburg does not give you regulatory sovereignty."

He takes sovereign AI development even further, explaining that South Africa should not only control the language datasets used to train models but as much of the architecture of AI as possible. "Controlling your own network, data centres, compute and AI factory as a South African company is a big differentiator; done that way, data and economic sovereignty need not leave the country's shores."

Surprisingly, he also mentions something most current AI users and regulators aren't aware of. "Using something like Meta for internal messaging can already be illegal depending on your business, because the moment a message leaves to be processed elsewhere, you have lost sovereignty, and most people do not know this."

Supporting his expert opinion is the research in this chapter that clearly illustrates that in order to manage its AI educational outcomes, South Africa could license the operational AI machinery where it exists, but keep the AI governance structures separate and as sovereign as possible.

How this could look in practice is data contributed into community-governed trusts on the model Te Hiku Media has proven for the te reo Māori language, held under licences that keep benefit flowing back to the communities the language is derived from, with SADiLaR, or something like it, as custodian on behalf of the nation.

Structured this way, the South African AI language data development programme would become simultaneously an employment programme, creating thousands of dignified part-time jobs concentrated where unemployment is deepest, and the construction of a permanent national asset. This programme could be delivered through the public sector using existing national initiatives such as Operation Vulindlela or MyMzansi. There are current implementation structures in the public sector from where this sort of initiative could be driven and that will make adopting this kind of programme much simpler.

In a September 2025 Wits University article, Professor Benjamin Rosman agreed with Suleman's feelings about the need for South Africa to develop its own sovereign AI resources. In the article, Professor Rosman, now chair of South Africa's national AI advisory panel and himself a co-founder and the chief science officer of Lelapa AI, argued that data sovereignty is insufficient without "algorithmic sovereignty", meaning Africa must also control the systems that turn its data into economic and social value. His warning was stark. "Without these capabilities, we become little more than exporters of raw digital resources, dependent on foreign companies and countries to refine our data for us."

Trust in this potential new source of South African educational strength will be earned by delivery through institutions people already trust. The first deployments, according to the experts, should be in people's own languages. From a government grant query answered correctly in spoken Sesotho, to a clinical or evaluative tool that works in isiXhosa, these solutions should be delivered through the banks, radio stations, clinics and community structures that hold the public's trust.

These small successes will help build the public case that South Africa should be in charge of its own AI education systems. The transparency the foreign providers cannot offer, like open training data, published failure modes, and community-held licences driven by youth education and development, is a domestic AI project's structural advantage, explains Kuben Nair. Nair works as the chief technology officer for Harambee, a South African educational non-profit that has already placed millions of young South Africans into the country's workforce.

"Our biggest untapped potential in South Africa is our young talent, our youth. We see this all the time. The problem is that we see young people as beneficiaries, people whose problems you have to solve and to whom you give things. But when you free them up, when you give them real access to information, tools and resources, they find the solutions," he explained in a recent interview.

Nair knows young AI talent exists. When 490 competitors from 61 countries entered the 2025 challenge to compress Lelapa's InkubaLM, the winners, all African, cut the model's size, its required compute power and its cost by three quarters, without losing any performance. So the capability is already here. It is simply not concentrated in African institutions or paid for by public money or no-strings-attached philanthropy. These are choices this country can reverse.

For South Africa to build the technology locally, as a national project, and support, educate and train the young AI engineers of tomorrow there needs to be a disruption to the current educational paradigm, says Nair. "My bent with education has always been to be more disruptive, and I think the opportunity now is to be disruptive with AI." He is passionate that "for things to really get done, it can't be government alone and it can't be industry alone; it's a partnership, and when that comes together, great things happen. You get speed, innovation, scale and depth."

None of this requires South Africa to choose an AI development "side", and the country's own history helps inform the reasoning. The tradition running from the Bandung Conference through to BRICS holds that the state may engage foreign powers on transparent commercial terms, while refusing to become a client of theirs.

Apply this politically non-aligned standpoint to AI and that means South Africa can run American infrastructure where it serves, self-host open-weight AI models where they serve, license data outward on terms set by South African regulators, and treat covert interference from any quarter as grounds for disqualification.

India's AI mission and Brazil's national programme are already framed in these terms, and the African Union's continental strategy gives the position a regional architecture. South Africa does not need to invent a new digital non-alignment in AI. It can practise the non-alignment it already espouses on the political stage for the AI digital age.

Nair believes this is achievable through a few simple steps. "The first item is access to AI, AI infrastructure, let me call it that, and data, tools and compute, because that access really is what's lacking across the board." To address this he goes on to explain, "whether that's some standard national infrastructure that multiple companies, Microsoft, Amazon, the cloud infrastructure providers and different tools, can plug into, rather than companies going one by one, with one company sponsoring a hundred schools and another sponsoring another hundred: a real national, country-level approach. I think that is achievable; there are a number of companies, actors and funders that would want to do this if we can bring them together."

If the state abandons its duty of care and leaves citizens and their data to the private sector, he feels the social issues in the country will likely deepen. The path forward for South Africa looks like having the state grow and foster the AI ecosystem by tailoring it to meet the needs of its citizens. If AI is simply introduced as a tool for profit extraction, as the rest of the economy is currently arranged, then it will likely produce the same results as the current economy, driving both inequality and exclusion.

This is an opportunity to do things differently and create jobs through a human-focused AI training system and model that benefits South Africa's citizens first. The most meaningful intervention is a dignified, well-paid salary for the people who have the knowledge, and critically the country's languages.

The languages, speakers and institutions are South African. The models can be South African too.

4 The published scientific evidence for an AI learning pathway in South Africa



The purpose of this report is to develop a rough baseline of where South Africa's AI educational efforts to create more jobs are today, and where they should be. That baseline also includes an AI in education literature review, done using best-practice techniques, called PRISMA. This review identified more than 160 different recent studies focused on AI in education in the Global South published since 2025. Every study has been individually cross-checked by a human and either included in, or excluded from, the list of resources, based on the human's judgement.

Within this section we break down how those studies can inform the South African response broadly. The full review dataset is available online[27], and each study it draws on is listed in the references at the end of this report for anyone who would like to dig into that literature and learn more. The authors intended it to be just a starting point towards a larger, more comprehensive review, hopefully taken up by South African academics.

As noted above, South Africa has a real chance to develop its own national AI education and job training strategies. To do so it is critical to understand what works and what does not in the relevant countries facing similar problems. To create an understanding of the published scientific literature, this report includes a basic, human-developed literature review of all the publicly available published science relating to educational outcomes in AI, focused tightly on the Global South and published in the last two years.

This report outlines the relevant public social science studies published in the United States Department of Education's Education Resources Information Center (ERIC), the world's largest repository of education-related literature. The outlined studies are retained in a full report after a full-text review by humans in what is called a PRISMA-style literature search. It was conducted for Collective X especially for this national baseline report on AI, skills development, and job creation.

Of 483 studies identified, 118 concerned Africa and were carried through detailed human screening, with the large majority retained across roughly fifteen African countries. The organising questions were practical. What does the published African evidence tell us about building education and training in AI and for AI skills, particularly for people who have already left school? What helps equip them with AI skills so they can find jobs?

The included studies span early childhood to university, but they are read here through a post-school lens, because that is what this report is focused on. They are school leavers, matriculants without a university place, TVET students, working adults and unemployed young people.

Moroccan researcher Farah Chkarka, whose work on AI literacy in education is covered in the corpus of research, says employers are already looking for AI skills in the workplace. If they are looking for “good prompting skills and good ethical understanding of how AI should be used then yes, it will make people a lot more hireable because nowadays employers will look for someone who knows how to use and engage with AI.”

The demand for AI tools is already here

The most consistent finding in the literature is that demand for AI tools among African learners needs no outside stimulation. Study after study using technology-acceptance models across Nigerian, Ghanaian and Egyptian post-secondary populations finds that people adopt AI tools like ChatGPT because they likely see them as immediately useful and easy to use, not because an institution told them to. Adoption, therefore, is running ahead of policy, curriculum and formal training, with gender and age shaping how people engage with the technology. Young people are more likely to experiment with AI tools, with women reporting use slightly more often than men.

For a post-school skills training system this is an unusual starting position. Most skills interventions spend heavily on recruitment and motivation; here the motivation already exists. What is missing is structure, quality and a route from informal use to certified, employable competence. A person who has taught themselves to prompt a chatbot on a phone has taken the first step on an AI skills pathway. The policy task is building the rest of the ladder so these skills can be translated into jobs, which is not simple.



Xolisile Zwane, a 26-year-old software engineering student at the We Think Code_ campus in Johannesburg, South Africa, says that she thinks developing AI skills can help her find a job. She noted that “when it comes to coding, you can use AI to upskill yourself. Learn new skills, have projects to showcase your skills, and you can get in the industry with AI.”

However, Kuben Nair, chief technology officer at Harambee, is more cautious. He doesn’t believe that developing AI skills alone is enough to find a job. “For me, in some cases we’re missing a trick in thinking about AI skills as just how to be a prompt engineer and how to use this or that AI tool. It’s a lot more than that.”

AI supports teaching, it does not replace it

There is strong evidence in the literature review and in other recent studies like the one published in March 2026 from a team that worked in both China and the United States, that noted that AI can help teachers and other trainers deliver effective, low-cost instruction, even in AI literacy[3].

A World Bank randomised controlled trial in Nigeria[1] found generative AI tutoring produced meaningful English-learning gains when used alongside traditional teaching at low cost per student, with the evidence holding across gender groups. The caveat is that this was simple instruction, not AI skills learning.

The Chinese/American authors of one of the latest AI teaching AI studies, published in the Proceedings of the AAAI Conference on Artificial Intelligence, noted that “with the learning-by-doing prompting activities with LLMs, students gained more experience in including contexts in their prompts and increased confidence in prompt writing for learning.” The authors also clearly note that there has to be a teacher in the loop: “the AI’s tendency to misinterpret the prompt intention or overly focus on language mechanics highlights where human oversight is necessary.” They also say that more studies should be done to replicate their research.


In some of the studies, chatbot teaching assistants improved achievement in Ghanaian higher education in controlled comparisons, although they were not teaching AI skills. And a Nigerian study[24] found a WhatsApp chatbot matched some conventional teaching for economics students, notable because WhatsApp is the platform already in the pockets of South African school leavers. One randomised trial found AI-immersion training improved researchers’ English skills comparably to human-only workshops, but that study has not been duplicated yet.

The findings of this critical research into AI education matter for a country trying to reach millions of people

outside formal institutions. Traditional education is constrained by lecturer supply, campus capacity and cost per student. According to the literature, AI-delivered instruction can loosen all three constraints at once. But for the foundational skills school leavers most often lack in AI, that increase in educational outcomes is not supported by the literature, mostly due to the shortage of AI training studies in the review.

In total though, the studies appear to show that a post-school AI learning pathway does not need to begin with new buildings. It can begin with structured, quality-assured, AI-supported, human-led tutoring delivered over internet infrastructure people already use, like WhatsApp, with the human facilitation concentrated where it adds the most value.

At the We Think Code_ campus in Johannesburg, South Africa, the twenty-five-year-old startup developer Lee Vandré Napier was upskilled through a programme that trains young high school leavers to work with AI in coding, merging online skills development with human trainers. “After high school I did some carpentry and then I came to learn and work in code, and I learned how to program, and now I’m working at a startup.”



Students at the We Think Code_ campus in Johannesburg, South Africa.

He said he decided to do something to further his own AI training because “schools are not performing at the rate we need. Like we need better than eighty percent literacy, right?” Now he is part of a seven-person team who have created a startup business developing the educational AI program to address the reading literacy problem in South Africa among eight-to-ten-year-old school learners, highlighted in the previous chapter. It was this training programme that made it possible for him to be building his educational AI tool for local languages like isiZulu and isiXhosa.

The entrepreneurship and employability evidence

A small but directly relevant cluster of studies connects AI learning to labour-market outcomes. A Nigerian experiential-learning intervention[25] running hands-on workshops in AI, IoT and cybersecurity found measurable gains in entrepreneurial skill and opportunity identification, providing some basic evidence that practical AI exposure can help drive entrepreneurship.

An Egyptian study found that embedding AI into work-based learning materially increased adults’ intention to enrol, suggesting AI content is itself a recruitment draw, at least for higher education.

A machine-learning analysis of returns to education across African regions adds a targeting insight: the payoff to additional education varies sharply by region and gender, so a post-school AI programme generates the most value where returns are currently weakest, typically among women and in under-served regions, rather than where enrolment is easiest.

Together, these findings support a design principle: teach AI through doing, in contexts tied to real economic activity, with entrepreneurship as a first-class outcome alongside employment. The Nigerian workshop model, short, practical and project-based, is a good illustration of how to make a shorter qualification course for AI literacy that South Africa might emulate, if it can also plan for ethics and privacy concerns.

Equity is a design choice, not a by-product

Monisha Prem is a Johannesburg attorney specialising in AI governance and one of the few South African lawyers internationally certified in the field.

The one skill you’ve got to teach right now, and you can teach a lot of kids how to do this, is how to control AI. Governance is about self-regulation; it’s about implementing controls within yourself, your team, your environment, your groups, in the little economies of society.

Monisha Prem, Johannesburg attorney specialising in AI governance

Because a post-school education system targets people the formal system has generally under-served, the equity findings in the studies carry particular weight. The Seeing-AI application outlined in the research produced meaningful mathematics retention gains among visually impaired Nigerian learners, showing AI-assisted learning can help reach groups that conventional educational provision fails, if the educational system is matched to a particular need. A Ghanaian intelligent tutoring intervention built around gender equity in under-resourced schools narrowed, rather than widened, computing participation gaps between male and female learners.

The dangers of AI in learning

But the claim that AI is straightforwardly good for teaching, and that South Africa should adopt it as a blanket plan, does not reflect the evidence in the literature review. Controlled research in 2025 found that learners who leaned on large language models to write showed measurably lower neural engagement, an effect its authors called “the accumulation of cognitive debt.” Add that to a large study that linked frequent AI use to poorer critical thinking. Even more concerning is that another systematic review concluded that over-reliance on AI dialogue systems also eroded analytical reasoning and creativity among participants. These are critically important findings and cannot be easily dismissed.

And these are not the only negative educational outcomes that illustrate the danger of employing AI in education. A Nigerian early-childhood study[26] found AI-linked gains were realised unevenly across rural and better-resourced settings, and South African work on inclusive AI course design also warned that algorithmic bias, especially uneven digital literacy which has already been noted in this report, can undermine equity in education, unless corrected at the course design stage.

The overall pattern in the data, which the authors noted is already biased towards the positive, is not consistent. So an educational pathway that simply posts content online can reproduce and amplify existing divides in connectivity, devices and confidence, and AI can do nothing to prevent this failure. But an educational pathway that tries to budget deliberately for data costs, low-bandwidth delivery, and targets disability access and recruitment of women and rural participants can potentially help plan for the opposite outcome, if it is worked into the course.

Given that most of the reported benefits within the literature are self-reported and short-term, and that rigorous South African or other directly relevant trials remain scarce, the responsible course is not wholesale AI adoption in education but conditional, safeguarded integration into educational systems that already work, utilising real teachers in working oversight roles. This integration must also first fix infrastructure, then capacity, then equity/access and after that, assessment.

Credibility of credentials: the assessment problem follows the learner

If the AI skills development pathway is to lead to jobs, employers must trust its certificates. Here the research sounds another warning. South African lecturers report that generative AI has destabilised conventional assessment validity, which in simple English means certificates and degrees are not necessarily what they once were even two years ago, because AI has offered students a chance to “cheat” and pass exams which they are simply not able to pass without its help.

This is not just a South African problem. Moroccan researcher Farah Chkarka, whose research is included in the literature, noted in a recent interview that “degrees no longer hold the value they once did, because students lean on AI in ways that diminish their critical thinking and even very basic skills. That is very alarming, which is why AI literacy and critical thinking should start from the very beginning.”

South African researchers cited in the study agree and argue assessment needs a fundamental redesign rather than patching. Nigerian and Moroccan integrity studies find learners’ willingness to misuse AI is governed more by peer norms and perceived detectability than by written rules which are largely disregarded.

According to Chkarka, who authored the Moroccan study, her university students struggle to understand

how to use AI ethically in her classroom. “It’s an exam. We’re handing out papers to students. The first thing before even reading the questions is to pull out a phone, take a picture, and cheat to get the answers.”

For a new post-school system the fact that AI is upsetting traditional education qualifications is both concerning and a potential opportunity to do things differently. Unlike legacy institutions retrofitting old examination models, a new pathway can be built AI-native from the start, assessing students through supervised practical demonstration, portfolios of real work, workplace verification and oral defence, formats in which AI assistance is either transparent and permitted or irrelevant in the case of supervised assessment.

AI does not and should not replace teachers

Farah Chkarka is a researcher at the Applied Linguistics and Artificial Intelligence Lab at Université Moulay Ismail in Meknès, Morocco, whose peer-reviewed work on academic integrity in the age of AI, including one of the studies captured in this report’s own literature review, makes her one of Africa’s leading voices on AI in education.

Invest in teachers and trainers before students. Teachers are the ones who transmit these skills, so if they do not have the readiness or the training, there will be a problem. I would suggest investing heavily in teacher training on AI first.

Farah Chkarka, researcher at the Applied Linguistics and Artificial Intelligence Lab at Université Moulay Ismail in Meknès, Morocco

The second most consistent theme in the research is that educator capacity, not learner appetite, is the binding constraint. South African pre-service teachers show strong intention to use AI but are held back by resources and lack of training rather than unwillingness. Ghanaian teacher studies identify AI literacy gaps as a barrier to teaching. Moroccan academics also report institutional barriers to learning AI outweighing personal willingness. South African distance-learning research finds that staff lack structured development pathways for AI tools they already want to use. Studies from Sub-Saharan Africa, Zimbabwe, Kenya and Rwanda document the same asymmetry.



South Africa has a real chance to develop its own national AI education and job training strategies. To do so it is critical to understand what works and what does not in the relevant countries facing similar problems.

A post-school AI educational pathway faces the same constraint, but if designed for these drawbacks, it can first upskill and then utilise human teachers and facilitators, mentors and TVET instructors as well as traditional lecturers. According to experts like Chkarka, training this facilitation workforce is not a supporting activity but the core investment, and common sense dictates that it should precede any learner-facing rollout.

Governance builds the pathway inside the system, not beside it

Several studies frame this AI teacher-training readiness gap as a governance failure. A South African bibliometric review found that the domestic research base is fragmented relative to the pace of adoption, and a World Bank analysis of EdTech across Sub-Saharan Africa concludes generative AI is outrunning ministries' capacity to produce guidance, leaving institutions to improvise. The lesson is to anchor the pathway inside existing structures such as the Department of Higher Education and Training and the TVET colleges from the outset, rather than launching a parallel scheme whose certificates employers cannot place. The governance vacuum is filled fastest by whoever moves first with a credible framework and to fix this, a co-ordinating intermediary, led by government, would be well placed to be that first mover.

Pieter de Villiers is the co-founder and CEO of Clickatell. He is also a board member of Collective X, which commissioned this report. De Villiers built a global chat commerce business from Cape Town that now moves billions of messages a month, and has long argued that South Africa's skills gap is a bigger problem than its jobs gap. He says that South Africa should "democratise [AI] access. Our young people are resilient, resourceful and curious, so start by giving them access, phone-first. Use the mechanisms that already exist: telcos already contribute a percentage of income toward access, and we can zero-rate [make free] a youth AI programme, partly state- and partly telco-sponsored."

De Villiers says: First, build the pathway around AI-delivered instruction on South African-developed AI models that are owned by the government. Develop this instruction with human facilitation, on platforms school leavers already use like SAYouth.

Second, make the offer practical and enterprise-oriented. Create short project-based modules in the mould of the Nigerian workshops, with self-employment and entrepreneurship treated as primary outcomes.

Third, train the trainers. Invest first in the facilitator workforce, since educator capacity is the binding constraint everywhere the evidence looks.

Fourth, design assessment AI-native from day one, certifying demonstrated ability to work with AI through practical and portfolio formats rather than policing AI out of examinations.

Fifth, engineer for equity deliberately, budgeting for data, low-bandwidth delivery, disability access and targeted inclusion of women and rural participants.

Sixth, anchor credentials in the national qualifications architecture so certificates convert into hiring decisions.

The African evidence shows a population already reaching for AI tools on its own initiative. What no African country has yet built is the structure converting this momentum into recognised, job-linked competence for people outside formal institutions. That is the gap a South African post-school AI skills pathway could help fill, and the evidence suggests it is both feasible and affordable. These findings are reflected in the expert interviews conducted for this report, which are explored in more depth in the next section.

This report's full PRISMA dataset is available online[27].

5 Driving job creation with appropriate AI education



One way the AI hyperscalers in the United States built market attention for their language models was by positing that AI was coming to take everyone's job. That claim helped fuel the largest single economic investment in human history, but it has done little to endear the technology to working people anxious about their job security, and little to nothing to train and educate workers in how to use AI to secure a brighter financial future.

I think over time we'll get models that are viable throughout the South African context, and obviously expanding to all of Africa.

Lee Vandré Napier, 25, founder of an AI home-language literacy startup

AI systems produce outputs that look like learning while bypassing the work that is learning. Automated grading is fluent but hollow and unreliable, often mimicking the process of marking without the experience, empathy and

solidarity of a human teacher. The opportunity is real, but conditional. The defensible and enabling use of AI in education is narrow, must be supervised by a human and must be labour-augmenting, not labour-cancelling.

It is worth noting here that many of those building the technology, as well as a number of globally focused economists, signed an open letter, released on 14 July 2026, warning the public about the real danger AI poses to jobs across the global economic system. The letter, which has only four sentences, says leaders must “build the incentives, guardrails, and institutions needed to steer AI in a direction that complements humans and benefits society.”

AI hyperscaler proponents often overlook the exploitation already present in the AI industry, such as the wages paid to AI language data creators in Kenya or the appropriation by AI of all online global intellectual property, from songs to videos to entire literature collections, without consent.

The Northern Cape is sunny, and solar energy costs have come down immensely in the last ten years. The cost of putting the town of Calvinia's Hantam High School onto a 100% solar and battery-powered system is less than half a million rand, roughly thirty-two thousand US dollars. That is nearly the same amount the school pays on its electricity bill in about seven months, so there is money to pay for a permanent power solution, but it just needs to be funded up-front.

That saved capital cost could go towards developing a technology lab with internet and AI training for Hantam High School. This double intervention would see both power restored and urgently needed training take place with internet access for young people whose future employment likely depends upon AI training and skills development.

This is not a difficult win, explains Ziaad Suleman, who leads Cassava Technologies' AI strategy in South Africa. He explains that Cassava is expanding the infrastructure, data centres, and high-performance computing needed to accelerate sovereign AI development across Southern Africa. Cassava is one of the only companies doing this kind of work in the whole region and Suleman believes that "the infrastructure is not the problem. My business runs about thirty thousand kilometres of terrestrial fibre, plus subsea cables and microwave, and together with competitors there is real connectivity."

He notes that the schools are there, the kids are excited and eager to learn, and the technology exists to fix these problems, but the current cost of data, and the training to use it with AI, is too expensive to implement for students to learn on, or for most people who are looking for AI to help upskill them for new jobs. "The issue is affordability, the opportunity cost. A young person in a rural area has to choose whether the thirty rand sent home buys food, data, or transport to school. So the real question is how do we make it easy for people to connect?"

He believes that the opportunity for AI to advance learning requires that we motivate all sectors of society so that it is easy for young people to learn and for learners and other people who may lose their jobs to AI to get trained. "People cannot be left behind, especially in education, where there is so much to benefit from, whether it is a video on long division or a lesson on photosynthesis. Creating affordable access and hubs is the bridge we are trying to cross."

So he believes that the way to fix these problems is through partnerships with businesses and the government working together, as well as finding alternative funding structures to help drive these kinds of AI educational outcomes. He is not alone.

Professor Rosman, from Wits University, explained what an AI literacy model would look like to him, one that would not only bridge the gap in current education, but one that could also upskill current job-seeking young people who don't have any work at the moment. This training, he notes, "has to start with in-person, human-facilitated learning, but built by experts across disciplines." He notes that getting delivery right is the hard part. "I would bring experts together to design it honestly, train facilitators, and build sustainability by having programme graduates become facilitators for the next group."

"Crucially," Professor Rosman says, "it starts with the local population, which varies place by place. I would leverage the experience of existing community programmes, like our HIV and AIDS interventions, for example, which already work with vulnerable communities, and I would translate AI courses into that kind of model. That bridging between worlds only happens if it is done very deliberately."

AI governance and legal expert Monisha Prem also thinks that South Africa's successful HIV/AIDS clinics in rural communities, built to address that historic crisis, are an excellent example of how to motivate this kind of education from the community level on up. She says the country's much-lauded and decades-long response to the HIV epidemic illustrates how South Africa already knows how to create an integrated campaign that any AI educational initiative will need. It would include radio ads and programming that promote AI literacy, printed posters about the dangers and benefits of AI, as well as content delivered free through government-run media outlets like the South African Broadcasting Corporation (SABC).

"I think it has a multi-pronged approach and it's going to take a couple of campaigns; it's going to take a strategy. It's not something that we can bury our heads in the sand about and it's not something we can do in isolation with a single policy." She says this is not only doable, it is in fact already happening within the commercial sector where she works.

“We have to get to people every way we can,” she says, “on their telephones, on social media, in spaces and places where they’re not using devices.” Her legal firm is already addressing questions like “how do we talk about subject matter such as your labour law rights under AI?”

She says that any public-private effort for AI education can follow her successful AI training approaches that she is facilitating for South African enterprises. “We are including it [AI] in employment handbooks. We’re including it in their weekly communication. We do it not only to protect the enterprise, but to train the families of the employees associated.”

She says good and successful AI training percolates throughout the community through this kind of family involvement too. “Take this information home to let your kids know how and when to use AI and how it actually can be harmful, and how it can actually enhance and work for you.”

It turns out that South Africa has already set out on a national AI training programme, much like the one Prem is suggesting. Vukani Mngxati is the chief executive officer of Microsoft South Africa. He noted that through a new partnership with the South African government, AI skills are being offered through the SABC in exactly this manner. “On our AI tour in January this year, together with the SABC, we announced a programme of short videos made available on the zero-rated (free) SABC+ platform.”

He says this vital training covers the most basic fundamentals, like “what AI is, how it works, what a large language model is.” He notes that already, “any South African with a phone can sign up and teach themselves on demand.”

Mngxati is adamant that AI skills training is vital for every sector of society, but “there are levels of skilling. At the top are government and company leaders, the top echelons of running South Africa Inc. They must be skilled first, because AI is not a bottom-up change; it can only be driven top-down, since it is such a significant change in how the world works and lives.”

He notes that one level down from this top echelon of business in corporates and government institutions is where there is already friction against AI use. “Fear is the problem,” he says. “People just don’t understand it and that drives fear.” He says education can help. “We need the mass of today’s workers to stop fearing AI and instead imagine working with it, so it helps them do their jobs better rather than displacing them.”

One of the experts consulted for this report is South African AI thought leader Johan Steyn. He writes a popular newsletter about AI called Johan’s Substack and runs an AI advisory business working with universities on AI in education. When asked if he thinks AI itself can help bridge this large AI education divide he said without missing a beat, “one hundred percent, if we do it right. Self-education through AI can be amazing: a personalised tutor at low or no cost, like the (US-based online learning system) Khan Academy.”

He agreed that schools simply aren’t currently teaching young people AI skills but notes that working, learning and teaching yourself or your children AI systems is already working within South African families. He explains that for instance he built his son “a platform in [Google] Gemini that ingests his schoolwork, picks up that our home language is Afrikaans and his age, and takes him through the material by voice using a Socratic method, never giving the answer but guiding him.”

Because he is an AI expert, this is possible, but not for everyone. It is one solution to the big issue of AI in education, as noted in chapter 2 and only relevant when used with the right amount of education and literacy. “I can do that, but most parents have no idea where to start. AI can also give faculty better course design, assessments and training. The trouble is the digital divide. It will only reach the few who can afford devices and connectivity unless the government and others do far more than they are now.”

Ziaad Suleman, the CEO of Cassava Technologies, is chairman of the board of Collective X, which commissioned this study. He agrees that AI systems can help bridge the education divide around AI and help upskill people for the AI jobs of the future that the country so greatly needs. “People cannot be left behind, especially in education, where there is so much to benefit from, whether it is a video on long division or a lesson on photosynthesis.”

He notes that “many children do not come home to food, a stable family or a desk,” and these problems of the country cannot be ignored, and that includes the systemic issues around poverty and lack of internet access. “Creating affordable access and hubs is the bridge we are trying to cross, and that needs a proper government strategy with both public and private participation, because the public sector does not have the infrastructure and the private sector cannot run at a loss.”



Students at the We Think Code_ campus in Johannesburg, South Africa.

Training the teachers and trainers of tomorrow

Many of the experts consulted for this report noted that one way to drive educational change is what is called “training the trainers”. This includes AI literacy training for teachers throughout the educational system, AI skills training for people who already are in the workforce, as well as AI training for job-seeking and educated youth. This is crucial since these youth, according to the experts, can become the AI trainers of tomorrow, one job that appears to have long-term sustainability.

Microsoft South Africa CEO Mngxati notes that in his educational hierarchy, after governments, corporate heads and current workers need urgent AI training, “the third level is the masses of South Africans, and here we must be smart and train the trainers, and ensure the fundamentals are taught at school level.” This means teachers and professors and other more informal AI trainers need AI literacy training, teaching materials and above all, support.

“Today,” he says, “those who have AI are learning, and those who don’t have it aren’t.” He says we also need to use the publicly available systems already in place in the country and supercharge them with resources and good AI training materials and then push them to the population through free “zero-rated” internet and TV-based AI training modules. “We have to democratise learning through zero-rated platforms so that anyone with a phone can access it.”

He goes on to say, “If I were in the seat of building the nation, I would train millions of South Africans on AI, transform institutions using AI, start with the government, and grow the AI economy.”

Developing AI training that works

Teachers need to understand both what AI can help with, such as research, course design, assessments and difficult student questions, and the immense amount it can never do. Standing in front of a class, the empathy, discipline, example and trust a teacher brings is not automatable.

Johan Steyn, founder of AIforBusiness.net

All of the experts interviewed for this report agreed that countries like South Africa need to train and upskill job seekers, current workers, students as well as teachers in AI.

Farah Chkarka is a researcher and instructor at Université Moulay Ismail in Meknès, Morocco. Her work focuses tightly on the intersection of AI and education. She has found that simply teaching an AI literacy class isn’t enough; it needs to be ongoing AI education driven from a literacy standpoint with the understanding that learning is constant.

Chkarka’s approach to AI education is not just teaching people how to use tools. She says students and teachers need AI literacy that includes “technical, critical evaluation, creative application, communicative competence, and ethical competence.” In her own class she says people learn by creating with AI, criticising it, using it ethically, and understanding its drawbacks. This is important because many students already use AI but “don’t know what is ethical, what counts as plagiarism, what does not.”

Chkarka says AI literacy should not be treated as a one-off online course but as a constant element of learning itself. “A single online course is doable and still good,” she says, “but it should be standardised and woven into every module, not taught in isolation.”

She argues that AI should be integrated into normal teaching. “When we teach writing, we should teach what is ethical with AI and what is not; when we teach communication, we should teach how to communicate with AI.” In her view, this must also be paired with deeper education in “cognitive skills, critical thinking and creativity,” because teachers and students need those human skills in order to properly evaluate and challenge AI.

South African tech entrepreneur and former First National Bank CEO Michael Jordaan has been publicly vocal on AI’s impact on banking, innovation, and the future of work. He thinks that South Africa does not need a new kind of “course” for AI literacy training, but “a national, outcome-funded credential layer on existing distribution.”

He imagines there would have to be at least two tiers: “universal AI literacy for schools and tertiary education, applied AI for working adults like bookkeepers, teachers, nurses.” The teaching systems would also be “anchored

to real employers who commit to hire or upskill, so the credential means a job, not a certificate.”

According to Jordaan, the training would have to include “the 45.6% [expanded] of young South Africans currently not in employment, education or training.” He notes that if the training system does not include these young people, “it is a vanity project.”

But Jordaan himself is also worried about the viability of his proposal. “What keeps me up at night is inaction by institutions and policymakers who care only about the next three years, and the assumption that this will solve itself. It will not.”

AI skills training to new AI jobs

Jordaan expounds his belief that South Africa’s call centre systems actually are a kind of English-based “empathy delivery system” that can catapult the country to a greater AI skills opportunity, turning what looks like a potential AI job loss liability into a huge AI jobs opportunity that he believes can scale quickly.

He says that isn’t the only way the country can embrace AI jobs for current workers. He says to imagine “AI-enabled diaspora work: our time-zone and English let a Cape Town graduate bill a London rate from a laptop. The constraints (bandwidth and skills) are fixable. Move now and we become the back-office brain of the global service economy, not its switchboard.”

Microsoft South Africa CEO Mngxati also sees a wealth of new jobs in South Africa. He groups the emerging roles into three tiers. “First, the builders, which include AI engineers, machine-learning experts, data scientists, solutions architects and product managers, including the brand-new head-of-AI role. Obviously that includes the data engineers, analysts and governance specialists they depend on, because there is no AI without data.”

Second, he believes that “AI gets paired with every profession, the AI-enabled banker, doctor, lawyer or teacher, and the agent economy of people who build and train these agents. Third, the trust-and-teaching layer of cybersecurity and compliance specialists and the AI trainers and curriculum designers who build and rebuild how people learn.”

In government, he believes that the paper-pushers have to go, replaced by digital-transformation and process-automation specialists who redesign processes

from first principles. “As Edward Kieswetter, the recently retired SARS commissioner, put it to me, imagine if we discarded the current processes and kept only the law. What do tax processes look like in the age of AI? Identity and access management alone has already changed, because your phone can authenticate who you are without you having to bring an ID. If we don’t train people for these jobs now, they will never get there.”

In reality, the problem is that no country or business has yet made this considerable leap in creating new and emerging AI jobs. The scientific data that would support this assertion, or shoot it down, is not yet published and until it is it remains highly speculative.

Critics of this kind of thinking about AI rollouts include Canadian author Cory Doctorow, who brought the world the term “enshittification”, which means “how digital platforms and services gradually decay and become worse over time to maximize corporate profits”.

He says that the point of big American AI companies is to replace expensive workers with their own AI products that companies have to pay for, not upskill people in call centres. In a June 2026 interview with the journalists at Democracy Now![28], he explained that AI companies’ subscribers only netted the companies \$50 billion a year. But he points out that as of 2026, “they’re spending \$1.4 trillion a year.”

He and many other analysts believe this loss is obviously completely unsustainable and what the big American hyperscaler AI companies “ultimately want to do is fire high-waged workers and replace them with substandard algorithms, and then make us accept substandard products. And that’s where they’re going to realise the return on this gigantic investment.”

This kind of job loss would of course be a disaster for the South African economy, notes Kuben Nair, chief technology officer at Harambee. “If you look at the roles that are more likely to be automated, those tend to be entry-level work that young people are predominantly in, and particularly young women.”

If these jobs were to disappear, he says, he is deeply concerned about the country’s future. “You see a lot of protests and marches already. I think that might escalate and then eventually you know, we’d have to do something because there’s a crisis at our doorstep and the country is burning.”


But Nair also believes that this can be avoided, if the right path forward is taken and the right investments in educational AI can be developed. “I’m an AI optimist and I see the opportunity for inclusive growth if we invest and drive our AI efforts intentionally.”

If South Africa is to actually develop these new jobs, which do not yet exist in any meaningful way, a new AI-centred educational order is needed, according to Dr Rodney Manyike from the South African government’s

Human Resource Development Council. “The challenge for South Africans is that the current basic education, TVET, and higher education systems are not yet adequately equipped to develop these skills at the scale and speed required. This presents an opportunity to adopt a leapfrog approach by embedding AI and digital skills into school curricula, modernising TVET programmes, expanding work-integrated learning opportunities and investing in large-scale reskilling and lifelong learning initiatives.”

Many of the experts consulted for this report noted that one way to drive educational change is what is called “training the trainers”. This includes AI literacy training for teachers throughout the educational system, AI skills training for people who already are in the workforce, as well as AI training for job-seeking and educated youth.

6 Conclusion



So what does all of this mean for South Africa looking ahead?

Kuben Nair has spent years building the systems that link millions of unemployed young South Africans to opportunity, making him one of the few people who understand both the technology and the ground-level reality of the country's jobs crisis.

"The most critical part of AI in education is access to AI infrastructure, let me call it that, and data, tools and compute, because that access really is what's lacking across the board ... we need a real national, country-level approach. I think that is achievable; there are a number of companies, actors and funders that would want to do this if we can bring them together. The next thing is talent and the curriculum. I love the story about teaching the teachers, because I think that's going to be the critical thing, rather than just thinking about the curriculum and what we do for the kids. The real intervention is going to be with the teachers: helping them understand AI."

A national AI strategy is meaningless in a school with no electricity. Hantam High School in Calvinia could be put on permanent solar and battery power for around half a million rand. Multiply that by the hundreds or thousands of schools in similar condition and the shape of the first intervention becomes clear.

The choice of what AI systems to adopt, which appears to be pressed on the country, is a false one because neither the American frontier systems nor the Chinese open-weight tier was built for South Africans. The evidence in this report shows why picking a side would be a mistake because the closed models deliver the strongest raw capability but process South African data offshore under a foreign national-security umbrella. The open-weight models can be self-hosted on local infrastructure at a fraction of the cost, but are less effective for many South African purposes.

Applied to AI, South Africa's non-aligned tradition opens up a way forward. Running American infrastructure where it safely serves, self-hosting open-weight models where they serve, licensing data outward on terms set by

South African regulators, and refusing to become a client of any single country's system.

But non-alignment is only a good posture to adopt if it rests on something owned. As Professor Rosman argues, data sovereignty is insufficient without algorithmic sovereignty, and as Ziaad Suleman adds, neither means much without infrastructure and economic sovereignty. Right now the sovereign core of a new South African AI strategy is the wealth of South African languages.

SADiLaR's state-funded language digitisation infrastructure, Masakhane's globally used benchmark ecosystem, two decades of speech technology at the CSIR, and Lelapa AI's commercial proof that small local models can run on entry-level phones together constitute the foundations of a sovereign capability. They are under-funded and under-connected, not absent.

The way forward is to scale and link them, not to staff new institutions, and that requires the public-private partnership every expert consulted for this report independently called for, because as Suleman puts it, the public sector lacks the infrastructure and the private sector cannot run at a loss.

How that reaches ordinary South Africans is the question on which everything else turns, and the country has already run the campaign that it can copy for this rollout. Professor Rosman and Monisha Prem reached, separately and unprompted, for South Africa's decades-long response to HIV and AIDS. That response worked because it was multi-pronged rather than a single policy and it was implemented through institutions people already trusted: the clinics, the radio stations, the SABC, the churches and community structures. Most importantly it reached people in their own languages, in the places where they live, where they got healthcare, where they worshipped and where they watched TV, whether or not they owned an internet-connected device.

It also worked because information travelled through families, which is precisely the percolation Prem describes in her own AI training work, where employees are asked to take what they learn about AI home to their children. The HIV/AIDS template is directly transferable, and the beginnings of it are already visible in the free AI literacy videos Microsoft South Africa and the SABC have put on television. These experts believe that this is the case study South Africa should be copying for the AI training and literacy rollout, and it is its own.

Which brings the point back to jobs. The evidence assembled here points to a sequence rather than a scatter of initiatives. Train the trainers first, because educator capacity, not learner appetite, is the main constraint and because as Kuben Nair notes, the real intervention is with the teachers, so train them. Deliver instruction over the platforms people already own, like cellphones and television.


Anchor the resulting credentials to employers who commit to hire, so that a certificate means a job, and place these credentials inside the existing qualifications architecture so they are recognised by industry. Then let the sovereign government-run language programme do double duty, creating thousands of dignified jobs transcribing, translating and evaluating South African languages exactly where unemployment is deepest, while building a permanent national asset. The skills are already here.

"Ultimately," the government's Dr Manyike says, "these new opportunities will not emerge automatically or be accessible to everyone. Replacing the traditional rungs on the ladder requires deliberate policy interventions, co-ordinated investment and strong partnerships between government, education institutions and industry. If the education and training system is intentionally redesigned to prepare people for AI-enabled work, AI can create new pathways into employment rather than simply eliminating existing ones. The objective should therefore be to ensure that the AI transition happens with the workforce rather than to it."

AI cannot be asked to make decisions about medical conditions or university admissions. A synthetic text generator built on languages will not, by itself, solve poverty and vital human-driven social services cannot be automated away. AI lacks the judgement to decide who passes and who fails a course or who is sick and who is well. It lacks the wisdom and empathy needed to carry out discipline or any kind of administrative justice. What it can do however is assist in learning. It can free up a teacher or a trainer's time to help them focus on teaching. It can provide a step-up for people to interact with technology on their own terms, in their own language.

Applied to AI, South Africa's non-aligned tradition opens up a way forward. Running American infrastructure where it safely serves, self-hosting open-weight models where they serve, licensing data outward on terms set by South African regulators, and refusing to become a client of any single country's system.

7 Recommendations. How to make it work



Professor Benjamin Rosman works at Wits University, is the chair of South Africa's national AI advisory panel and a TIME100 AI honouree. He also co-founded the MIND Institute, the Deep Learning Indaba and Lelapa AI, the institutions through which Africa is working to become a builder of AI rather than just a consumer of it.

For the government, go all in. Frame the future you want, then work backwards to the pathway. Prioritise a sector such as healthcare or employment and really commit. On employment, run pilot studies: get top experts to design a business- and a life-upskilling programme using AI.

Professor Benjamin Rosman, chair of South Africa's national AI advisory panel and a TIME100 AI honouree

These recommendations should most likely be placed under one of the national economic transformation initiatives within the Presidency such as Operation Vulindlela or the Presidential Youth Employment Intervention (PYEI). They can be aligned to other subset plans including in the Department of Communications and Digital Technologies (DCDT) Master Plan and the Master Skills Strategy within the Human Resource Development Council. Together they could be co-ordinated as a business and government partnership with Collective X playing an orchestration role across aspects of the national plan, specifically those related to information and communication technology (ICT) roles.

1. Training

Develop a set of new strategic frameworks for AI skilling. Build it from day one to grow from hundreds of thousands to millions of AI-educated job seekers. Government's job, potentially through the Departments of Basic Education and Higher Education and Training, is basic credentialing with a mandate no single university or SETA can provide. This will embed foundational AI capability across the education and skills system and enable recognition of agile micro-credentialing across universities, TVETs and SETAs.

Build the capacity for senior government and business leaders to help drive this skills development opportunity. Train the teachers, educators, and other instructors.

The scientific evidence shows that what is blocking uptake isn't learner appetite, it's educator capacity.

Deliver this training on the phones people already carry, with access designed in, including for women and rural learners, rather than left to chance. This means delivering data, some devices, and the necessary connectivity with free internet access, for example through a platform like SAYouth, South Africa's largest free online network designed to connect young, unemployed job-seekers with learning, training, and earning opportunities. Put the human facilitators trained within the context of where they add the most value into this educational chain.

2. Relevant AI qualifications

Make these new "micro-credentials" as recognised as possible. The DCDT noted that certificates take too long and the qualification authorities don't currently recognise these "credits". These new training initiatives can be delivered through the proposed and current new national frameworks under review, including the Digital Skills Framework One (DSF1), which describes the digital competencies people need from basic literacy through to advanced user and sector-specific capabilities, as well as through the Skills for the Information Age (SFIA) framework. This will yield trusted, portable and recognised evidence of competence, with formal recognition reserved for cases where it genuinely adds value.

Build the assessment of these new credentials as AI-supported from the start. Certify what people can actually do with AI through practical work, portfolios, workplace sign-off and oral defence. That's the thing employers will probably pay to trust, according to this research. Anchor these new credentials and training in places like the TVET colleges, universities, community training centres and tech hubs.

Fortunately the Digital Skills Impact Fund already has a new working mechanism to implement this recommendation and should be scaled up for this purpose.

3. Financing

Focus primarily on scaling this effort as a demand-led approach to skilling and workplace-based learning, using agile and outcomes-based financing mechanisms to respond to the changing skills landscape. Fund it properly through instruments like SETA discretionary grants, the Skills Development Levy, the National Skills

Fund (NSF) and B-BBEE skills-development spend, not new money.

Fund the facilitators, TVET instructors and mentors before the learner rollout, not after. AI can help reduce the costs of this training but it must be teacher-led, and those teachers and trainers must be supported.

Get supporting finance from industry to help cover the real-world costs of what has traditionally shut people out of online access, which is crucial to developing these skills. That means affordable access to data, lower-cost devices, sponsored connectivity in rural and township areas, as well as disability access, all delivered through zero-rated (or free) structures.

4. Languages as a national asset

Develop the South African language commons. Treat indigenous-language AI as public infrastructure, which South African businesses can build on, but which remains in the hands of the public. Public datasets and models for our languages are both a public good and something AI-skilled South Africans can build and sell.

Seed it through data work, and work through procurement that rewards local language capability and training, creating downstream employment, entrepreneurship and other income-generating opportunities. It keeps the intrinsic value of South African language datasets onshore and prevents locking the majority of South Africans out of English-only tools. Turn South African language and cultural data into productive national infrastructure.

5. Drive job creation from the outset

Aim at entrepreneurship, independent contracting, and freelance work, not just formal jobs. The economy can't absorb everyone into employment, and AI often does most for the smallest operators. A one-person business can now get sales strategy, legal review and an advisory board it could never afford. Build the national programme so self-employment and small-business productivity count as much as a payslip from some large employer, or more.


6. Keep AI development as sovereign as possible

Choose AI deliberately, with sovereignty in mind. The cheapest or most familiar model isn't automatically the right one to run nationally. DCDT, which is currently developing the National AI Policy, will have to weigh AI model capability against data sovereignty, and most importantly take into account laws that are already in place like the Protection of Personal Information Act (POPIA).

Crucial considerations like overall model cost and long-term dependency on providers who do not have the same values and responsibilities as the national government, which are outside the country's jurisdiction, must be taken into account. Establish a publicly owned entity that can hold the private data of the country's citizens and decide how to develop it openly, instead of drifting into whatever's easy, or paid for by private business.

To close:

- **This is a jobs plan, not a technology plan.**
- **Train the trainers first, and every rand spent reaches many more learners.**
- **Deliver the instruction on the phones people already carry, so nobody waits for a campus to be built before they can start learning.**
- **Scale and accelerate existing demand-led micro-credentialing initiatives.**
- **Direct and scale existing financing sources towards outcomes-based, demand-led models.**
- **Use AI to start businesses, not just improve them, and self-employment counts as much as, or more than, a payslip.**
- **Build AI in our own languages, so that valuable work takes place exactly where unemployment is worst.**
- **Choose AI deliberately with sovereignty in mind, so that South Africa controls its own destiny, because our livelihoods depend upon it.**
- **Scale what is already in place in the Digital Skills Impact Fund (DSIF).**



AI use disclosure

Because this is a report on AI, the team behind it, from Materra Research, are sharing how we used AI to assist in the compilation of this report, and crucially, how we did not.

Claude from Anthropic was used to spell-check the document and create the list of acronyms, which was all then rechecked by multiple humans. It was also used as a web search tool to research the internet for some of the facts and noted research in some parts of the report. It performed much better than Google or another search engine like Bing in this regard and saved large amounts of human search time. All links and citations provided were checked with the principal source in the bibliography.

Riverside AI, a professional transcription and interview tool, was used to take the audio of the expert interviews and create timed and searchable transcripts of them. Each one was also read line by line and double-checked by the researchers as Microsoft Word documents. Then, the primary interviews conducted by journalist Jeffrey Barbee for this research, and other interviews for this report, were stored and accessed as searchable documents through Google's NotebookLM AI system.

The original comparative analysis between AI models was supplied by a researcher who works with AI systems from Anthropic (Claude), OpenAI, Google, DeepSeek, Moonshot (Kimi), Zhipu (GLM) and Xiaomi (MiMo) and other Chinese models. Further analysis was sometimes provided by using these systems to evaluate each other, but the relevant comparative analysis was done by Lelapa and the noted evaluation systems.

The PRISMA study was conducted by a human researcher searching the ERIC database within the stated parameters of the study, using prompts that were both Claude-derived as well as researcher-derived. The study itself was directly compiled with every single inclusion or exclusion for the study checked by a human, and each study was sourced in its entirety either through the author, the ERIC database or the university/institution. Web addresses for all included studies are listed in the references, and were correct at the time of publication. The researchers involved in the PRISMA study are not academics and further research is required to confirm our analysis but the study followed all PRISMA principles.

Impressum

This report was created by Materra Research for Collective X and published in October 2026.

For more information contact jeff@materraresearch.com
www.materraresearch.com

References (works cited in the text, numbered in order of appearance)

1. De Simone, M., Tiberti, F., Barron Rodriguez, M., Manolio, F., Mosuro, W., & Dikoru, E.J. (2025). From Chalkboards to Chatbots: Evaluating the Impact of Generative AI on Learning Outcomes in Nigeria. World Bank Policy Research Working Paper 11125. <https://eric.ed.gov/?id=ED676624>
2. ACM ICER (2025). AI Literacy in K-12 and Higher Education in the Wake of Generative AI: An Integrative Review. Proceedings of ICER 2025. <https://dl.acm.org/doi/10.1145/3702652.3744217>
3. Learning to Use AI for Learning: Teaching Responsible Use of AI Chatbots to K-12 Students Through an AI Literacy Module (2026). Proceedings of the AAAI Conference on Artificial Intelligence. arXiv:2508.13962. <https://arxiv.org/abs/2508.13962>
4. International Monetary Fund (2026). World Economic Outlook, April 2026: unemployment rate (2026 projections). <https://www.imf.org/external/datamapper/LUR@WEO>
5. International Telecommunication Union (2025). Facts and Figures 2025 – Internet Use. <https://www.itu.int/itu-d/reports/statistics/2025/10/15/ff25-internet-use/>
6. DataReportal (2025). Digital 2025: South Africa. <https://datareportal.com/reports/digital-2025-south-africa>
7. Artificial Analysis. Artificial Analysis Intelligence Index v4.1. <https://artificialanalysis.ai/articles/artificial-analysis-intelligence-index-v4-1>. Independent model benchmarking, the source for the report’s capability-and-cost tiers (accurate as of early July 2026).
8. Goldman Sachs. Tracking Trillions: The Assumptions Shaping the Scale of the AI Build-Out. <https://www.goldmansachs.com/insights/articles/tracking-trillions-the-assumptions-shaping-scale-of-the-ai-build-out>. Industry-analyst estimate of hyperscaler capital expenditure, cited to establish the scale of the Western AI build-out.
9. Adams, S.P. Your AI Code Subscription Is Heavily Subsidised. Substack. <https://stephenpauladams.substack.com/p/your-ai-code-subscription-is-heavily>. Commentary arguing consumer AI pricing sits below operating cost, supporting the point that frontier LLM pricing is subsidised.
10. Protection of Personal Information Act (POPIA). <https://popia.co.za/>. South Africa’s data-protection statute, cited on the regulatory-sovereignty risks of foreign-hosted AI.
11. The Diplomat (2026). Competing Interests, Common Narrative: How Industry Shapes the AI Race. <https://thediplomat.com/2026/08/competing-interests-common-narrative-how-industry-shapes-the-ai-race/>. Analysis of how commercial interests shape the framing of the global AI race, supporting the non-alignment argument.
12. Anthropic. Statement on the Department of War. <https://www.anthropic.com/news/statement-department-of-war>. Primary-source company statement, cited on the entanglement of a frontier AI firm with US defence interests.
13. Palantir Technologies (company website). <https://www.palantir.com/>. Company reference, cited as an example of AI firms embedded in the surveillance and security state.
14. New York Times (2025). Report on the Ownership of DeepSeek. <https://www.nytimes.com/2025/01/28/business/deepseek-owner-china-ai.html>. Reporting on DeepSeek’s ownership by a Chinese quantitative-trading firm, cited on political-control questions around Chinese models.
15. Convly (2026). Zhipu GLM Explained. <https://convly.ai/zhipu-glm-explained-2026/>. Explainer on Zhipu’s GLM models and its Tsinghua and state links, used to describe the Chinese open-weight tier.
16. Reuters (2026). Beijing Is Looking at Curbing Overseas Access to China’s Top AI Models. <https://www.reuters.com/world/beijing-is-looking-curbing-overseas-access-chinas-top-ai-models-sources-say-2026-07-07/>. News report on possible Chinese limits to overseas model access, cited on the risk of depending on Chinese AI.
17. Fortune (2026). DeepSeek V4: Price and Performance. <https://fortune.com/2026/04/24/deepseek-v4-ai-model-price-performance-china-open-source/>. Coverage of DeepSeek V4’s price and performance, cited on the cost-competitiveness of Chinese open models.
18. IrokoBench benchmark paper (Masakhane Research Foundation). arXiv:2406.03368. <https://arxiv.org/abs/2406.03368>. Masakhane’s African-language benchmark, central to the report’s evidence on model performance in South African languages.
19. AfroBench benchmark paper. arXiv:2311.07978. <https://arxiv.org/abs/2311.07978>. A comprehensive African-language LLM benchmark, cited alongside IrokoBench.
20. Adebara, I., et al. (2025). Where Are We? Evaluating LLM Performance on African Languages. Findings of the ACL 2025. <https://aclanthology.org/2025.acl-long.1572/>. Peer-reviewed evaluation of LLM performance across African languages, supporting the language-gap argument.
21. Hao, K. (2022). Artificial Intelligence for the People. MIT Technology Review. <https://www.technologyreview.com/2022/04/22/1050394/artificial-intelligence-for-the-people/>. Feature on community-controlled AI (the Māori and Te Hiku case), cited as a model for data sovereignty.
22. SADiLaR (South African Centre for Digital Language Resources). New Infrastructure. <https://sadilar.org/en/new-infrastructure/#content>. South Africa’s digital-language-resources centre, cited as an existing national AI-language asset.

23. Louw, J.A., Davel, M., & Barnard, E. (2005). A General-Purpose IsiZulu Speech Synthesizer. *South African Journal of African Languages*, 25(1). <https://www.tandfonline.com/doi/abs/10.1080/02572117.2005.10587252>. Peer-reviewed article on South African speech-technology research.
 24. Anyanwu, C.C., Ononiwu, P.N., & Isiozor, G.N. (2024). Comparative Impact of WhatsApp Chatbot Technology and Glaser's Teaching Approaches on the Academic Performance of Education Economics Students in Tertiary Institutions in Nigeria. *Education and Information Technologies*, 29(17). <https://eric.ed.gov/?id=EJ1453991>
 25. Nweke, L.O., Okebanama, U.F., & Mba, G.U. (2025). Enhancing Entrepreneurial Skills through Experiential Learning in IoT, AI, and Cybersecurity. *Discover Education*, 4. <https://eric.ed.gov/?id=EJ1473291>
 26. Akinrotimi, A.A., Popoola, T.J., & Udo, I.E. (2025). The Impact of Artificial Intelligence on Early Childhood Learning Outcomes in Nigeria. *International Journal of Education and Development using Information and Communication Technology*, 21(2). <https://eric.ed.gov/?id=EJ1488989>
 27. Materra Research (2026). PRISMA Study Dataset: ERIC Search Corpus for This Report. https://docs.google.com/spreadsheets/d/170QZddJeCwhWGuHlwIql_0-Kf_X_W1VhjkFZBTgxwF4/edit
 28. Democracy Now! (2026). Interview with Cory Doctorow ('Reverse Centaur'), 25 June 2026. https://www.democracynow.org/2026/6/25/reverse_centaur_doctorow
- Additional studies screened in the PRISMA corpus, and further reading (not all individually cited)**
29. Arkorful, V., *et al.* (2025). Exploring Artificial Intelligence Literacy among Basic School Teachers in Ghana. *Discover Education*, 4. <https://eric.ed.gov/?id=EJ1478500>
 30. Arthur, F. (2026). Modelling the Nexus between Artificial Intelligence Usage and Graduate Capital Dimensions: A PLS-SEM Approach. *European Journal of Education*, 61(1). <https://eric.ed.gov/?id=EJ1497865>
 31. Asanre, A.A., Oguntola, J.I., Ramatea, M.A., & Lawani, A.O. (2026). Enhancing Mathematics Retention through Seeing-AI: A Quasi-Experimental Study among Visually Impaired Junior Secondary Students in Ogun State, Nigeria. *European Journal of STEM Education*, 11(1). <https://eric.ed.gov/?id=EJ1505169>
 32. Ayanwale, M.A., & Molefi, R.R. (2024). Exploring Intention of Undergraduate Students to Embrace Chatbots: From the Vantage Point of Lesotho. *International Journal of Educational Technology in Higher Education*, 21. <https://eric.ed.gov/?id=EJ1417563>
 33. Chkarka, F., & Fatmi, H. (2026). Comparative Examination of Master's Students' and Faculty Members' Maintenance of Academic Integrity in the Age of AI. *Journal of Academic Ethics*, 24(1). <https://eric.ed.gov/?id=EJ1491803>
 34. Dewa, A. (2024). Artificial Intelligence for Educational Sustainability in the South African School System: A Bibliometric Analysis and Literature Review. *International Journal of Education and Development using Information and Communication Technology*, 20(3). <https://eric.ed.gov/?id=EJ1461864>
 35. Du Plessis, E., & Bayeck, R.Y. (2025). South African Lecturers' Views of ChatGPT: An AI Technology Used for Designing Online Assessments. *Journal of Teaching and Learning*, 19(4). <https://eric.ed.gov/?id=EJ1487153>
 36. Ekwam, E.O., Heloo, E.S., Okebukola, P.A., Kuagbedzi, F.N., & Oladejo, A.I. (2026). Enhancing Computing Education and Gender Equity through the Generalized Intelligent Framework for Tutoring (GIFT) in a Developing Country. *Discover Education*, 5. <https://eric.ed.gov/?id=EJ1504784>
 37. Enakrire, R.T., Fombad, M.C., & Morodi, L. (2025). Skills Required of Academics to Use Digital Technologies in Open Distance Learning Institutions. *Innovative Higher Education*, 50(3). <https://eric.ed.gov/?id=EJ1475113>
 38. Essel, H.B., Vlachopoulos, D., Nunoo-Mensah, H., & Amankwa, J.O. (2025). Exploring the Impact of VoiceBots on Multimedia Programming Education among Ghanaian University Students. *British Journal of Educational Technology*, 56(1). <https://eric.ed.gov/?id=EJ1457021>
 39. Essel, H.B., Vlachopoulos, D., Tachie-Menson, A., Johnson, E.E., & Baah, P.K. (2022). The Impact of a Virtual Teaching Assistant (Chatbot) on Students' Learning in Ghanaian Higher Education. *International Journal of Educational Technology in Higher Education*, 19. <https://eric.ed.gov/?id=EJ1355304>
 40. Gerlich, M. (2025). AI Tools in Society: Impacts on Cognitive Offloading and the Future of Critical Thinking. *Societies*, 15(1), 6. <https://doi.org/10.3390/soc15010006>. Study of 666 participants linking frequent AI use to weaker critical thinking via cognitive offloading (counterpoint source).
 41. Iddrisu, H.M., & Iddrisu, S.A. (2025). Teacher Readiness to Use AI Tools in the Classroom. *International Journal of Research in Education and Science*, 11(3). <https://eric.ed.gov/?id=EJ1480991>
 42. Kamdjou, H.D.T. (2023). Estimating the Returns to Education Using a Machine Learning Approach: Evidence for Different Regions. *Open Education Studies*, 5(1). <https://eric.ed.gov/?id=EJ1394840>
 43. Kosmyna, N., *et al.* (2025). Your Brain on ChatGPT: Accumulation of Cognitive Debt When Using an AI Assistant for an Essay-Writing Task. *arXiv:2506.08872*. <https://arxiv.org/abs/2506.08872>. Experimental study finding weaker neural engagement and "cognitive debt" among users who rely on LLMs to write (counterpoint source).
 44. Kwakye, I.N., Kwakye, K.A.P., Adom-Fynn, D., Mensah, F.E., & Agyei, S.D. (2025). Exploring the Use of Generative Artificial Intelligence Tools among University Students in Ghana. *Discover Education*, 4. <https://eric.ed.gov/?id=EJ1492124>
 45. Maimela, C., & Mbonde, P. (2025). Artificial Intelligence in South African Universities: Curriculum Transformation and Decolonisation, Aid or Obstacle? *Frontiers in Sociology*, 10, 1543471. <https://doi.org/10.3389/fsoc.2025.1543471>. Examination of AI, curriculum transformation and decolonisation in South African universities (counterpoint source).
 46. Mjahad, R., Boukranaa, A., El Karfa, A., & Sandy, K. (2025). The Implementation of Artificial Intelligence in Moroccan Higher Education: Benefits and Challenges from the Perception of Academics. *SAGE Open*, 15(3). <https://eric.ed.gov/?id=EJ1487282>
 47. Mnguni, L. (2025). A Qualitative Analysis of South African Pre-Service Life Sciences Teachers' Behavioral Intentions for Integrating AI in Teaching. *Journal for STEM Education Research*, 8(2). <https://eric.ed.gov/?id=EJ1467897>
 48. Negm, E. (2025). Developing Enrollment Intentions for Work-Based Learning That Uses Artificial Intelligence (AI): Aspects for Marketing Interventions for Higher Education Institutions. *Higher Education, Skills and Work-based Learning*, 15(5). <https://eric.ed.gov/?id=EJ1486935>

49. Nwokocho, G.C., & De Villiers Morkel, J. (2025). Navigating Online Learning and Artificial Intelligence: Identifying and Managing Assessment Risks in Private Higher Education in South Africa. *International Journal of E-Learning & Distance Education*, 40(2). <https://eric.ed.gov/?id=EJ1494273>
50. Ofem, U.J., Owan, V.J., Iyam, M.A., Udeh, M.I., Anake, P.M., & Ovat, S.V. (2025). Students' Perceptions, Attitudes and Utilisation of ChatGPT for Academic Dishonesty: Multigroup Analyses via PLS-SEM. *Education and Information Technologies*, 30(1). <https://eric.ed.gov/?id=EJ1457968>
51. Omeh, C.B., Olelewe, C.J., & Hu, X. (2025). Application of Artificial Intelligence (AI) Technology in TVET Education: Ethical Issues and Policy Implementation. *Education and Information Technologies*, 30(5). <https://eric.ed.gov/?id=EJ1465744>
52. Owan, V.J., Mohammed, I.A., Bello, A., & Shittu, T.A. (2025). Higher Education Students' ChatGPT Use Behavior: Structural Equation Modelling of Contributing Factors through a Modified UTAUT Model. *Contemporary Educational Technology*, 17(4). <https://eric.ed.gov/?id=EJ1488802>
53. Oyetade, K., & Zuva, T. (2025). Advancing Equitable Education with Inclusive AI to Mitigate Bias and Enhance Teacher Literacy. *Educational Process: International Journal*, 14. <https://eric.ed.gov/?id=EJ1466584>
54. Plantinga, P., Sinha, S., Eke, D.O., & Marivate, V. (2024). Responsible Artificial Intelligence in Africa: Towards Policy Learning. *Data & Policy*, 6, e22. https://www.researchgate.net/publication/386556474_Responsible_artificial_intelligence_in_Africa_towards_policy_learning. Analysis of responsible-AI policy learning in Africa, cited on governance and capacity constraints (counterpoint source).
55. Pramjeeth, S., & Ramgovind, P. (2025). Reconceptualizing Assessment Design in the Age of AI: Is It Genuine or Faux? *New Directions for Teaching and Learning*, 182. <https://eric.ed.gov/?id=EJ1475093>
56. Reuters (2025). China's Zhipu AI Launches Free AI Agent, Intensifying Domestic Tech Race. <https://www.reuters.com/technology/artificial-intelligence/chinas-zhipu-ai-launches-free-ai-agent-intensifying-domestic-tech-race-2025-03-31/>. News report on Zhipu's free AI agent launch, illustrating the pace of Chinese AI competition.
57. Salifu, I., Arthur, F., Acquah, B.Y.S., Opoku, E., Nortey, S.A., & Boateng, E. (2025). Exploring Graduate Students' Use of Generative Artificial Intelligence in Ghana. *Discover Education*, 4. <https://eric.ed.gov/?id=EJ1481554>
58. Salimi, F. (2025). Aligning Policy and Practice: The World Bank's Approach to EdTech in Sub-Saharan Africa. *Policy Futures in Education*, 23(6). <https://eric.ed.gov/?id=EJ1479176>
59. Saal, Chetty, Ntshayintshayi, Moosa & Masuku 'A scoping review of the integration of artificial intelligence in primary and secondary schools from 2020 to 2024: policy implications for South Africa'. *SciELO South Africa* (2025). https://www.scielo.org.za/scielo.php?script=sci_arttext&pid=S2520-98682025000100005. South African scoping review of AI in schools, cited on local policy gaps (counterpoint source).
60. Sebihi, A., Schoelen, L., & Uwamwezi, B. (2025). Unveiling AI's Impact in Rwandan Higher Education: Navigating Trends, Challenges, and Imperatives at the University of Rwanda. *Discover Education*, 4. <https://eric.ed.gov/?id=EJ1473995>
61. Soko, J., Pete, J., & Mureithi, G. (2025). Assessing Awareness and Usage of AI Tools among Faculty and Learners in Higher Education Institutions in Sub-Saharan Africa. *Asian Journal of Distance Education*, 20(2). <https://eric.ed.gov/?id=EJ1492783>
62. Strzelecki, A., & ElArabawy, S. (2024). Investigation of the Moderation Effect of Gender and Study Level on the Acceptance and Use of Generative AI by Higher Education Students: Comparative Evidence from Poland and Egypt. *British Journal of Educational Technology*, 55(3). <https://eric.ed.gov/?id=EJ1419158>
63. Tarisayi, K.S., & Manhibi, R. (2025). Revolutionising Education in Zimbabwe: Stakeholder Perspectives on Strategic AI Integration. *Journal of Learning and Teaching in Digital Age*, 10(1). <https://eric.ed.gov/?id=EJ1459941>
64. Ugwu, N.F., *et al.* (2025). Comparing the Efficacy of Artificial Intelligence Immersion and Human-Led Workshops for Enhancing Researchers' English Language Skills: A Randomised Control Trial. *Higher Learning Research Communications*, 15(1). <https://eric.ed.gov/?id=EJ1468504>
65. Wainaina, P.K., & Sun, Y. (2025). Educators' Perceptions and Willingness to Integrate Generative Artificial Intelligence in Teaching and Research: Evidence from Kenyan Higher Education. *Discover Education*, 4. <https://eric.ed.gov/?id=EJ1489433>
66. Xia, Q., Zhang, P., Huang, W., & Chiu, T.K.F. (2025). The Impact of Generative AI on University Students' Learning Outcomes via Bloom's Taxonomy: A Meta-Analysis. *Asia Pacific Journal of Education*. https://www.researchgate.net/publication/393867767_The_impact_of_generative_AI_on_university_students_learning_outcomes_via_Bloom's_taxonomy_a_meta-analysis_and_pattern_mining_approach. Meta-analysis finding AI's learning gains concentrate at lower cognitive levels and weaken for higher-order skills (counterpoint source).
67. Yakubu, M.N., David, N., & Abubakar, N.H. (2025). Students' Behavioural Intention to Use Content Generative AI for Learning and Research: A UTAUT Theoretical Perspective. *Education and Information Technologies*, 30(13). <https://eric.ed.gov/?id=EJ1480880>
68. Zhai, C., Wibowo, S., & Li, L.D. (2024). The Effects of Over-Reliance on AI Dialogue Systems on Students' Cognitive Abilities: A Systematic Review. *Smart Learning Environments*, 11(1), 28. <https://doi.org/10.1186/s40561-024-00316-7>. Systematic review concluding that over-reliance on AI dialogue systems can erode analytical reasoning and creativity (counterpoint source).

