Microsoft | AI Economy Institute

AI Diffusion Report: Where AI is most used, developed, and built

November, 2025

Executive Summary

The history of human progress is, in many ways, the history of technology's diffusion. Among these advances, a few—what economists call general-purpose technologies—have had the greatest impact on humanity. The printing press democratized knowledge. The steam engine powered industry. Electricity transformed daily life and enabled computing. And the internet connected the world, accelerating the exchange of ideas, trade, and collaboration across borders.

Each transformed society not through invention alone, but through diffusion—the way millions of people made these technologies part of how they live, work, and learn.

Artificial intelligence is the next great general-purpose technology—and the fastest-spreading technology in human history. In less than three years, more than 1.2 billion people have used AI tools, a rate of adoption faster than the internet, the personal computer, or even the smartphone. Yet, similar to other general-purpose technologies before, its benefits are not spreading evenly. AI use in the Global North is roughly double that in the Global South. Without focused effort, this gap will define who benefits from AI for decades to come.

The Three Forces of Diffusion

Every transformative technology advance through three forces:

- **Frontier builders** — the inventors and pioneers who push the boundaries of what is possible.
  In AI, these are the researchers and model creators expanding the limits of intelligence.

- **Infrastructure builders** — the engineers, entrepreneurs, and institutions who scale breakthroughs through networks, tools, and skills.
  In AI, infrastructure builders provide the compute and connectivity that make large-scale intelligence possible.

- **Users** — the individuals, companies, and governments who use, adapt and apply new technologies to solve real-world problems.

History teaches that progress accelerates when all three evolve together. Edison built the light bulb—but it took power grids and everyday users to make electricity universal. The same holds true for AI.

What the Data Shows

As a general-purpose technology, AI stands on the shoulders of three others—electricity, connectivity, and computing. Its adoption is fastest where these foundations exist, and slowest where they do not. Nearly four billion people—half the world—still lack the basics needed to use AI.

AI adoption in the Global North is roughly twice that of the Global South, with the gap widening sharply in countries where GDP per capita falls below $20,000. In some Global North countries, more than half of the working-age population uses AI, while in parts of Sub-Saharan Africa and Asia, some of the least developed nations have adoption rates below 10%.

A new barrier is also emerging: language. Nations where low-resource languages dominate—like Malawi or Laos—show lower adoption even after adjusting for GDP and internet access.

Some countries—Singapore, the UAE, Norway, and Ireland—stand out as leaders in AI Adoption, proving that strong access to technology, education, and policy coordination can drive rapid adoption even without frontier-level model development or data centers.

From a frontier builder perspective, the number of AI models continues to rise, while the performance gap between them keeps narrowing. The U.S., led by OpenAI's GPT-5, remains at the frontier, with China trailing by less than six months. Only seven countries—the U.S., China, France, South Korea, the U.K., Canada, and Israel—rank among the top 200 models, and the distance between the frontier (U.S.) and the last of these (Israel) is now just 11 months.

From an infrastructure builder perspective, the U.S. and China together host 86% of global data center capacity, underscoring how concentrated the foundation of AI remains.

Measuring AI's Global Progress

To understand this transformation, we propose three complementary indices:

- **AI Frontier Index** — measuring the world's leading frontier models by their performance and innovation.
- **AI Infrastructure Index** — capturing where the capacity to build, train, and scale AI exists.
- **AI Diffusion Index** — reflecting where AI is being adopted.

Together, they show not only who is building AI, but who can benefit from it.

Ultimately, the value of artificial intelligence will be judged not by the number of models produced, but by the extent to which they benefit society.

The Importance of Diffusion

Thirty years ago, the World Bank published The East Asian Miracle, [1] a landmark report that sought to explain how a handful of economies in East Asia achieved unprecedented growth. Its conclusion was clear: adoption and adaption of technologies developed elsewhere—rather than invention of these technologies—can drive national transformation. Few comparisons illustrate this more clearly than South Korea and the Philippines.

In 1960, the two countries were strikingly similar. Each had a per capita income of about $2,000 in 2025 dollars, with comparable demographics and education levels. Yet when modern manufacturing, export-oriented industrialization, and expanded access to higher education became available, South Korea embraced and scaled these drivers of growth.

In the late 1970's, the South Korean government identified semiconductors as a strategic priority and partnered with the private sector to make strategic investments. Through licensing agreements, engineering expertise, and partnerships—supported by government R&D, tax incentives, and infrastructure investment—Korean semiconductor manufacturers quickly established themselves as global leaders.

South Korea's economy surged, growing 6.2 percent annually and doubling living standards every eleven years. By contrast, the Philippine economy grew at 1.8 percent annually, close to the world average, as it remained more closely tied to primary commodities, particularly in agriculture, mining, and business process outsourcing services.

What began as two countries on similar footing became two very different economies. Today, South Korea is known not for inventing semiconductors but for producing them better, faster, and cheaper than anywhere else and continues to be a global powerhouse in the manufacturing of memory chips. Its success is largely attributable to mastering and scaling digital technology and building industries around it. It serves as an example of how technological adoption and collaboration between the public and private sectors has powered one of the world's most advanced nations.

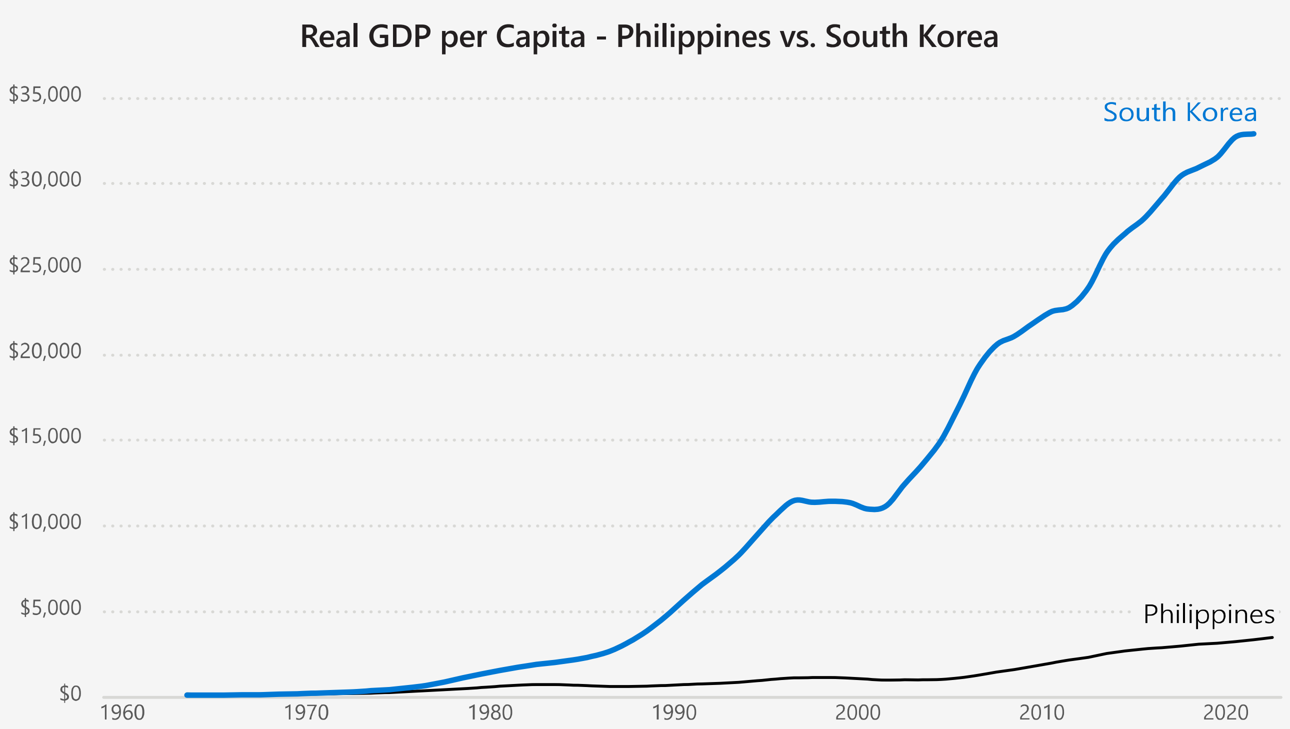

Data Source: GDP Per Capita, World Bank World Development Indicators. [2]

The Fastest Start to Diffusion in Human History

Artificial intelligence has emerged as the newest general-purpose technology of our era. Its adoption—surpassing a billion users in under three years—illustrates a speed of diffusion that few earlier technologies, such as the radio, Internet, or smartphone, have matched. [3,4]

And yet, beneath the headline numbers, the familiar patterns of uneven diffusion are already emerging. In some countries like UAE or Singapore over half of the working age population uses AI. In other regions, particularly in Sub-Saharan Africa and parts of Asia, adoption in many countries remains below 10%. [5] The divide is attributable not only to access to AI tools; it reflects wider disparities in infrastructure, education, and language.

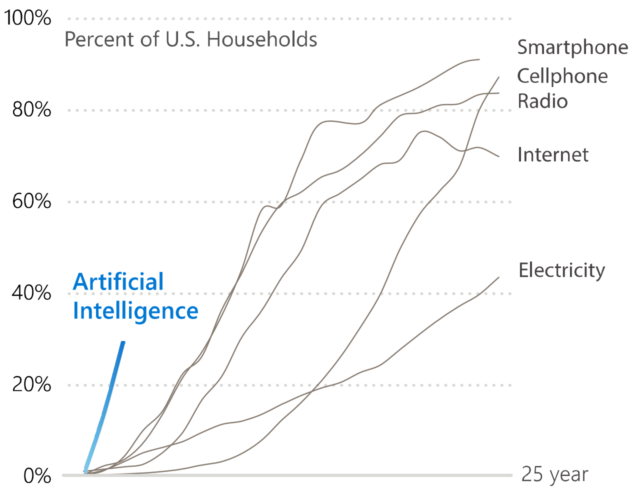

Data Source: The New York Times; [3]
Our World in Data (for Internet). [4]

The **UAE (59.4%)** and **Singapore (58.6%)** lead in AI use among working-age adults, reflecting their long-term investment in digital connectivity and skills

AI Diffusion by Economy

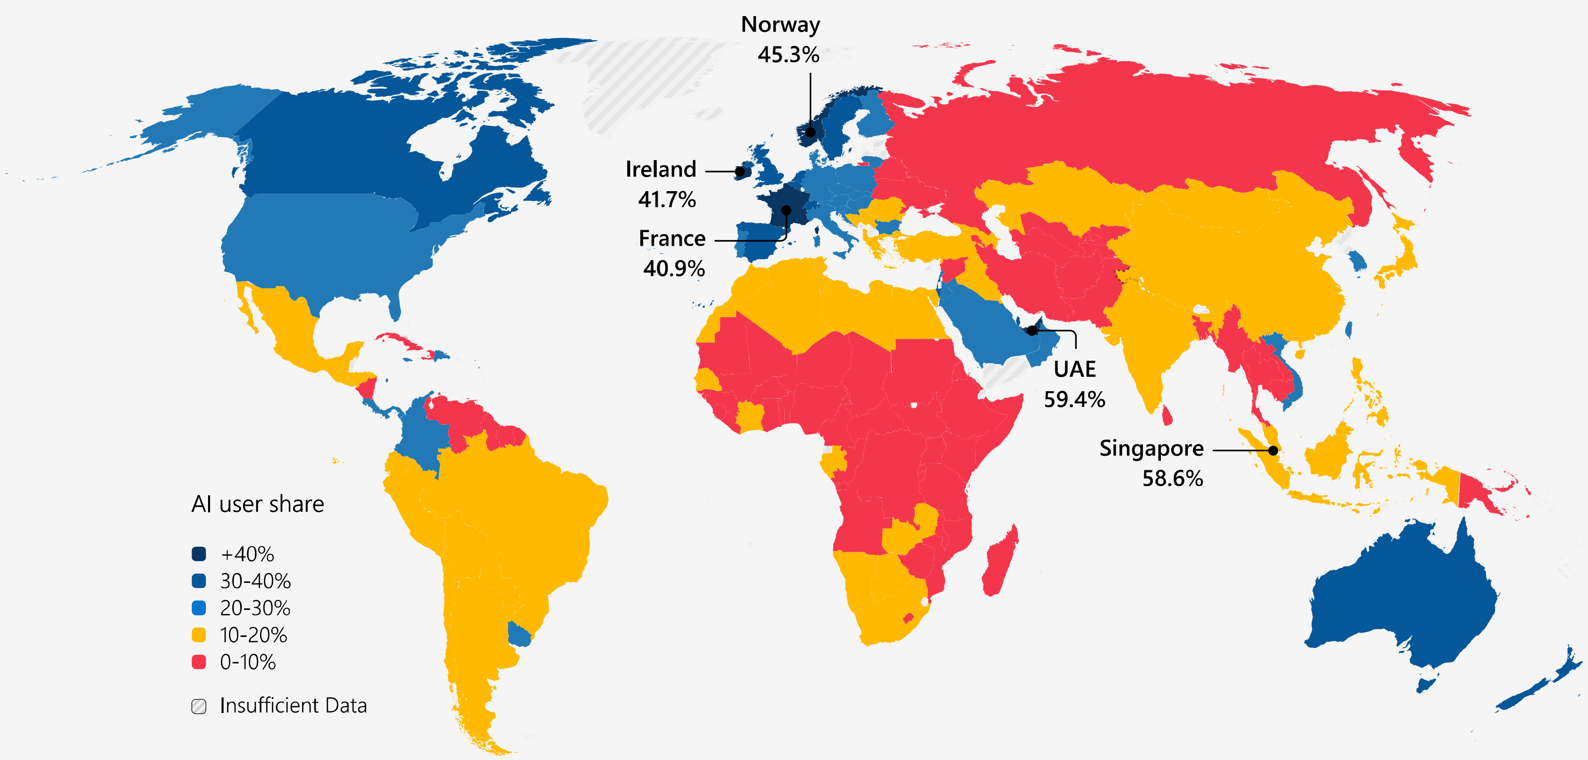

Data Source: AI Diffusion Technical Report, AI for Good Lab. [5]

The Building Blocks of AI

When we look at the diffusion of AI, we see a clear pattern: adoption strongly correlates with GDP. Countries in the Global North show significantly higher uptake, while many in the Global South are falling behind.

To understand this more clearly, it helps to look at the "building blocks" of AI

- **Electricity** powers the devices and data centers that make AI possible.
- **Data centers** form the backbone of global networks – including the Internet – and are essential to train and run inference on AI models.
- **The internet** connects users to AI and enables the flow of data.
- **Digital and AI skills** begin with the ability to use a computer and navigate the digital world.
- **Language** determines who gets to use AI and how it is shaped.

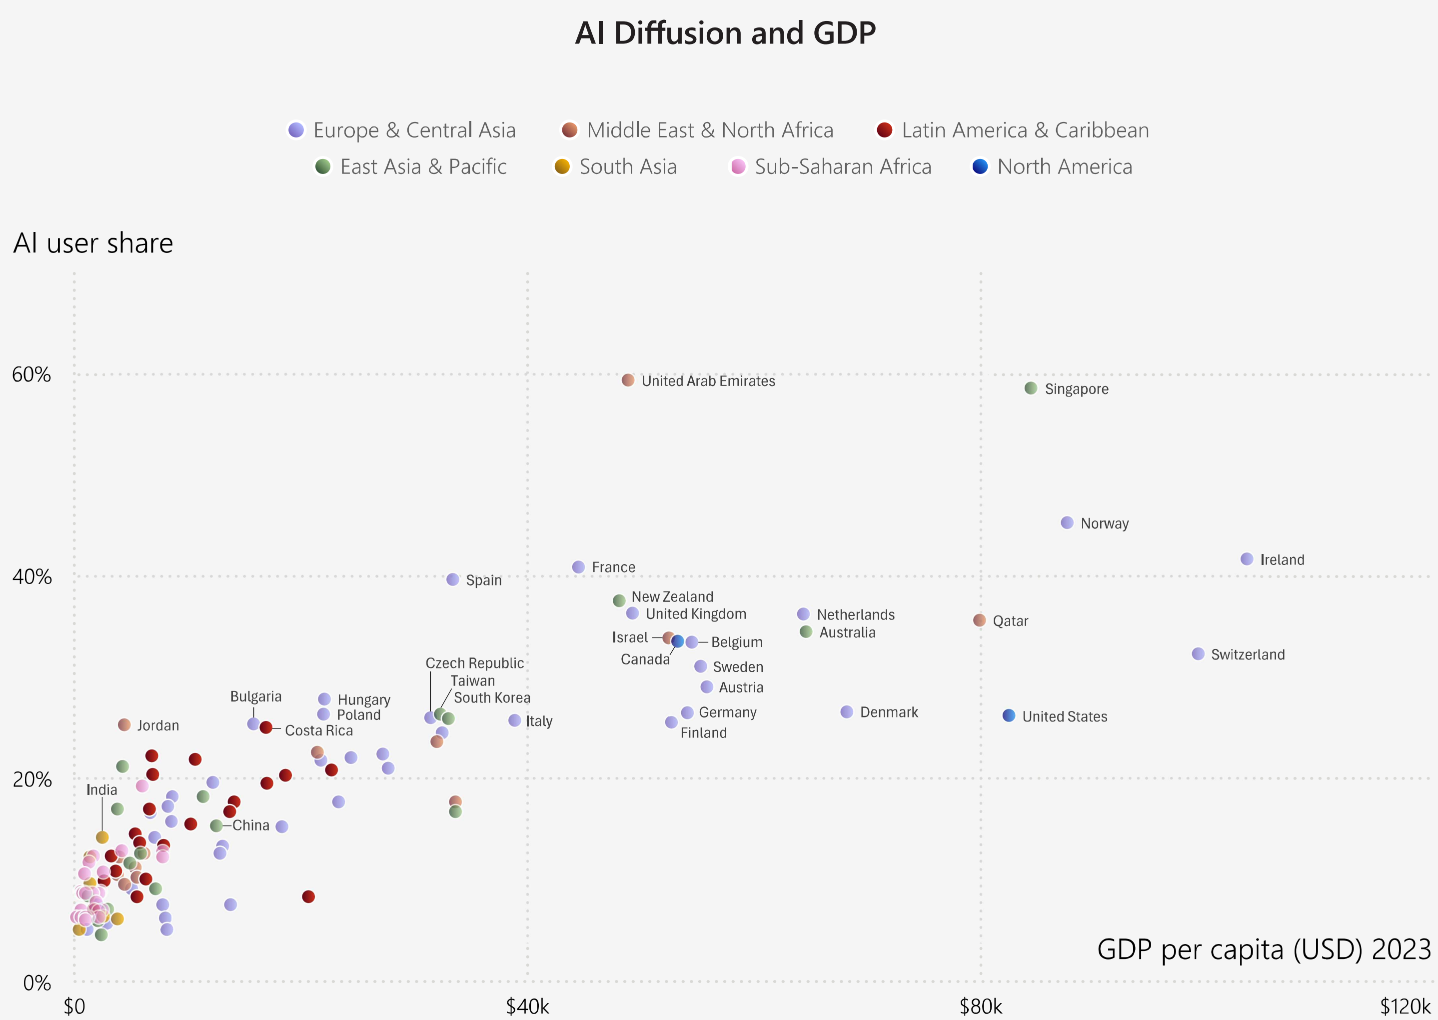

Data Sources: World Bank GDP per capita 2023; **[6]** Taiwan GDP per capita: USD | Economic Indicators | CEIC. **[7]**

Countries that have invested in digital infrastructure and whose citizens have the skills to use AI tools in a language they speak are adopting AI at notably higher rates. For example, Singapore's rapid embrace of AI—with 59% of the working-age population using AI **[5]**—builds upon decades of deliberate investment in digital infrastructure and education. Beginning in the 1980s, the government launched successive efforts to wire the nation with high-speed connectivity and expand access to computers in schools. **[8]** By the 1990s and 2000s, national initiatives positioned digital technology as central to economic growth and public services. At the same time, Singapore invested heavily in human capital, strengthening STEM education and fostering partnerships between universities, industry, and government. **[9]** This consistent, long-term strategy created a highly connected, digitally literate population—an essential foundation that has enabled Singapore to quickly adopt AI and compete in the new digital economy.

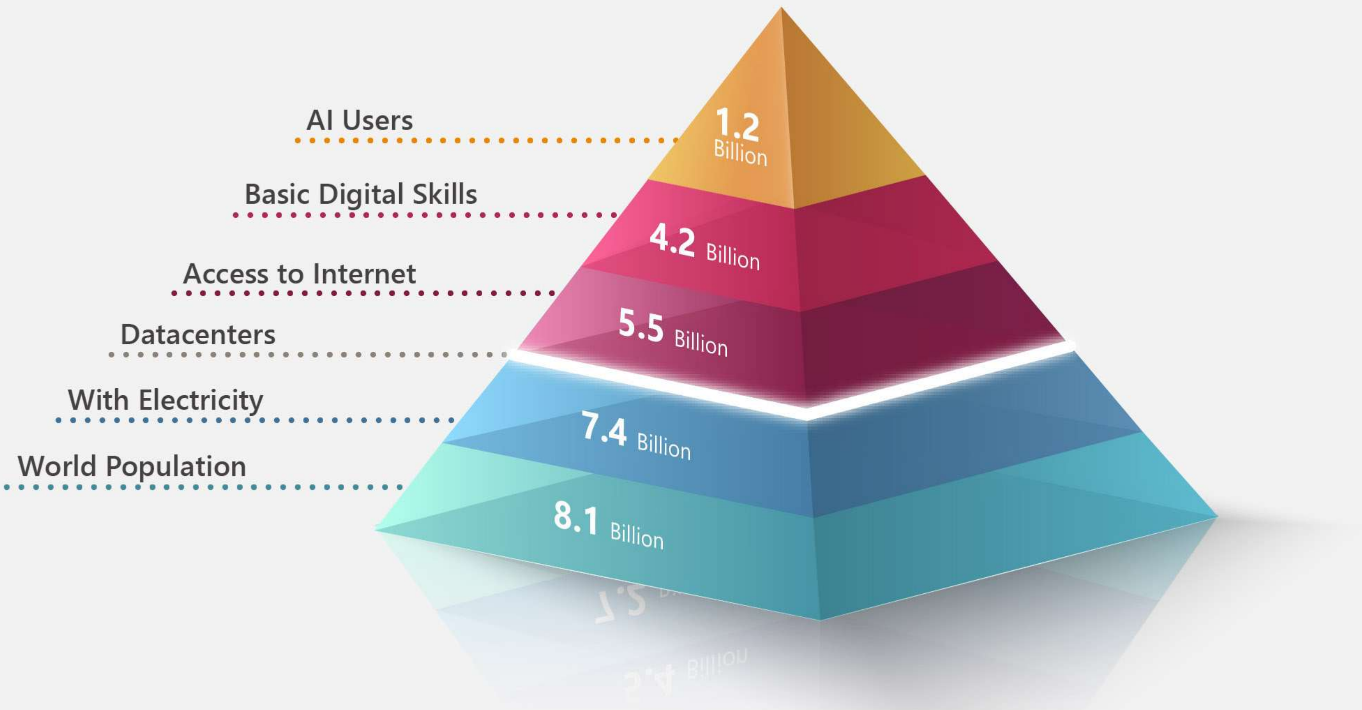

Data Sources: World Bank Total Population 2024; **[10]** Share of the population with access to electricity, Our World in Data; **[11]** Individuals using the Internet 2024, ITU; **[12]** Internet users with basic information and data literacy skills, ITU. **[13]**

Electricity

Access to reliable electricity is strongly correlated with GDP per capita and is foundational for AI adoption. Reliable energy and electricity access are critical foundations for life and livelihoods, powering homes, devices, data centers, and businesses. Without consistent electricity, communities are unable to participate fully in the digital economy, connect to the internet, access modern healthcare, or support educational advancement.

While many of the world's countries have achieved universal electricity coverage, 18 of the 20 countries with the largest electricity access deficits are in Sub-Saharan Africa. **[14]** As a result, Sub-Saharan Africa now accounts for 85 percent of the global population without electricity—up from 50 percent in 2010. **[15]** Globally, over 750 million people lack access to electricity, limiting their ability to participate in the digital economy. **[16]**

Access to Electricity by Economy

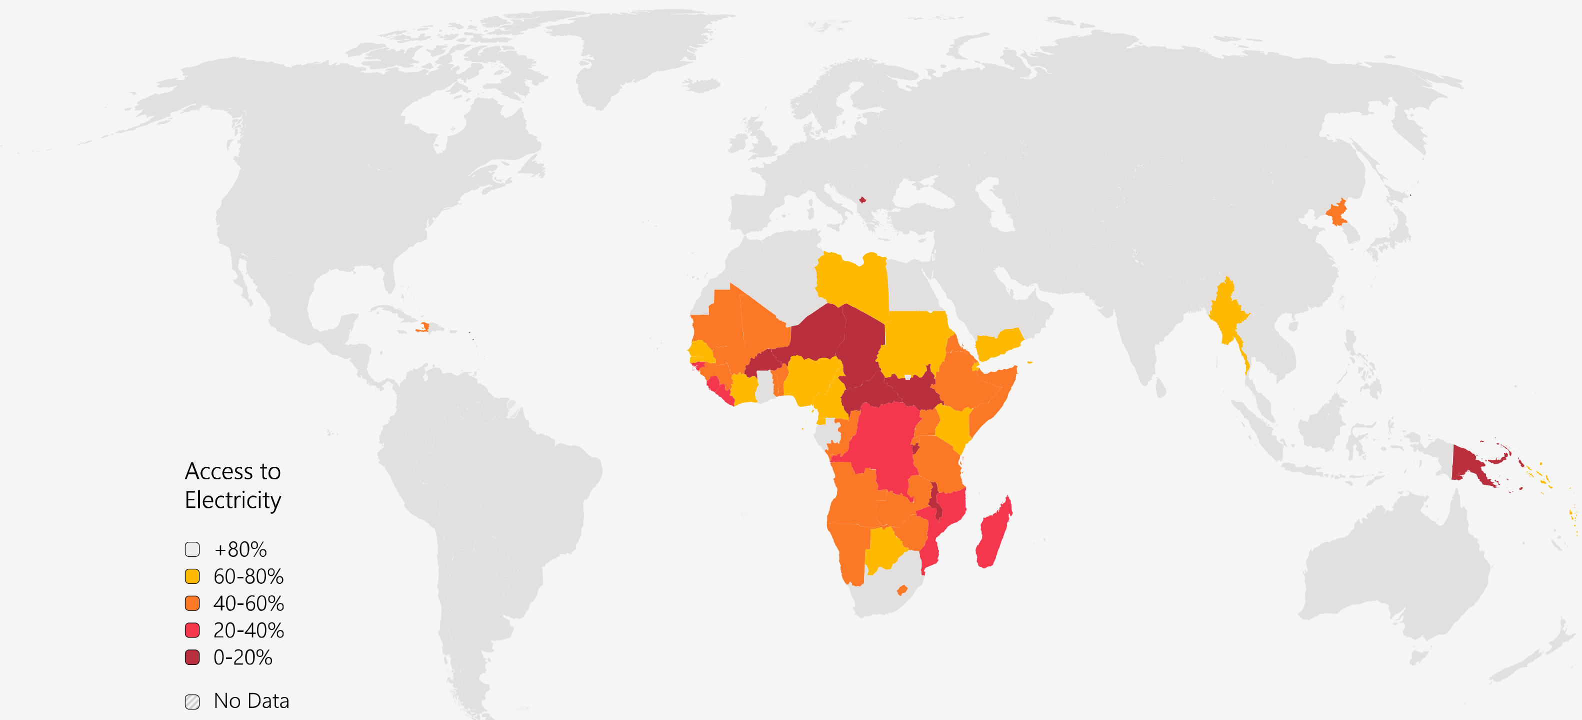

Data Source: Our World in Data. **[17]**

Data Centers

The new generation of large language models require not only computing power to train, but also significant infrastructure to run. Inference—the moment when people actually use these models—typically takes place in the cloud. That doesn't mean every country needs its own data center. Technically, people everywhere can still access AI services hosted abroad. But proximity matters.

Reducing the physical distance between datacenters and customers means data has less distance to travel, which minimizes the response time. This is critical for user experience and application performance; time and again, studies have shown that even modest delays reduce how often people use online services. [18] Proximity can also reduce bandwidth costs and drive network efficiencies. And increasingly countries and regions have laws requiring that certain data, such as government, healthcare, and financial data, be stored and processed within their borders.

That's why governments and companies alike have long invested in placing data centers closer to their citizens and customers. Yet today, the map of AI infrastructure reveals a sobering truth: most of the world's AI data centers are concentrated in the Global North. Only a handful exist in the Global South.

Where AI Data Center Regions are Located

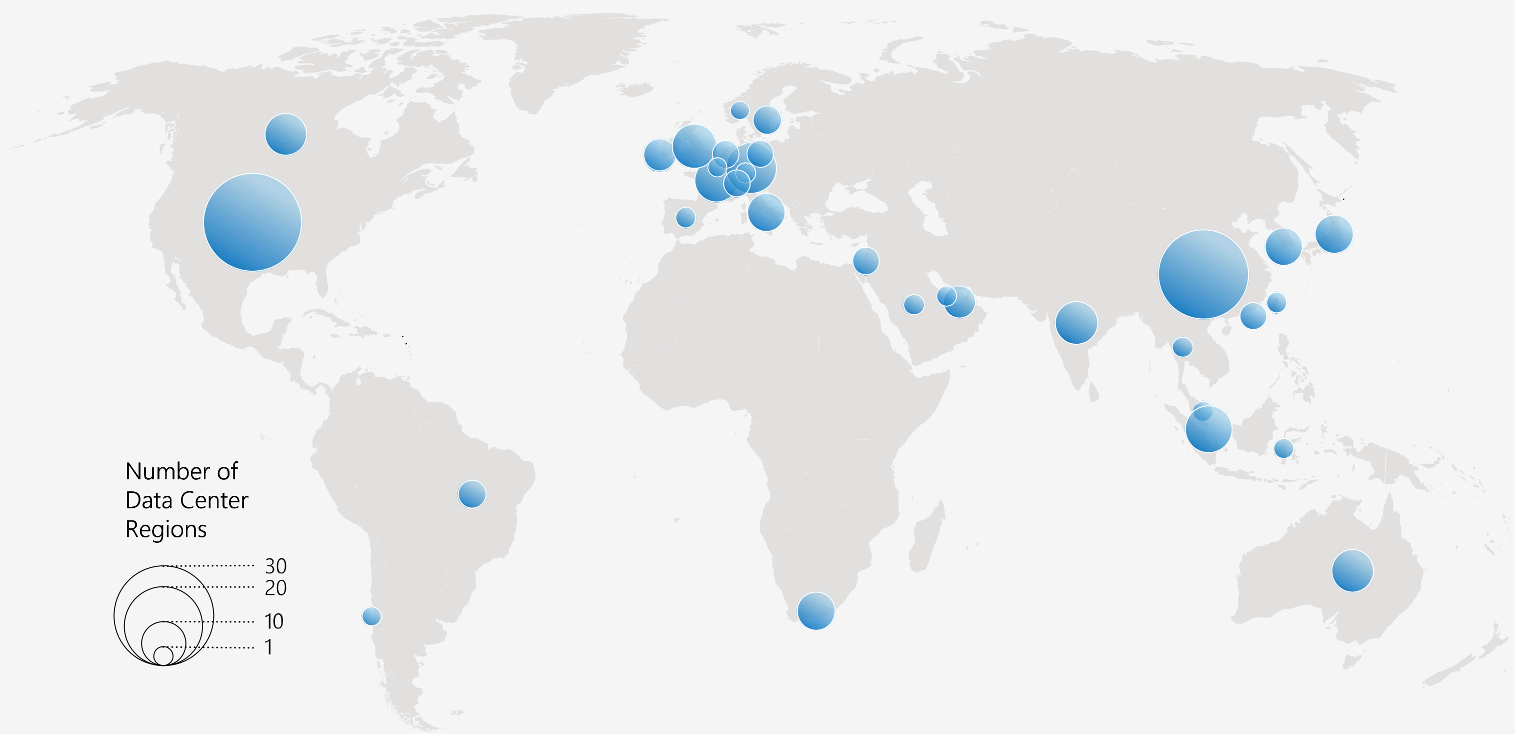

Data Source: University of Oxford. [19]

The Internet

Access to the internet is another one of the key enablers of AI adoption. Across the world, we see that when connectivity expands, adoption follows. In Zambia, for example, the national AI adoption rate is only 12 percent. But among those with internet access, it rises to 34 percent—almost three times as high. This pattern holds across countries with similar levels of internet connectivity. From Pakistan to Côte d'Ivoire, from Zimbabwe to The Gambia, from Guatemala to Kenya, Nepal, and Honduras, the story is the same. Connectivity is a gateway—an essential step to participation in the AI economy.

Internet Connectivity by Economy

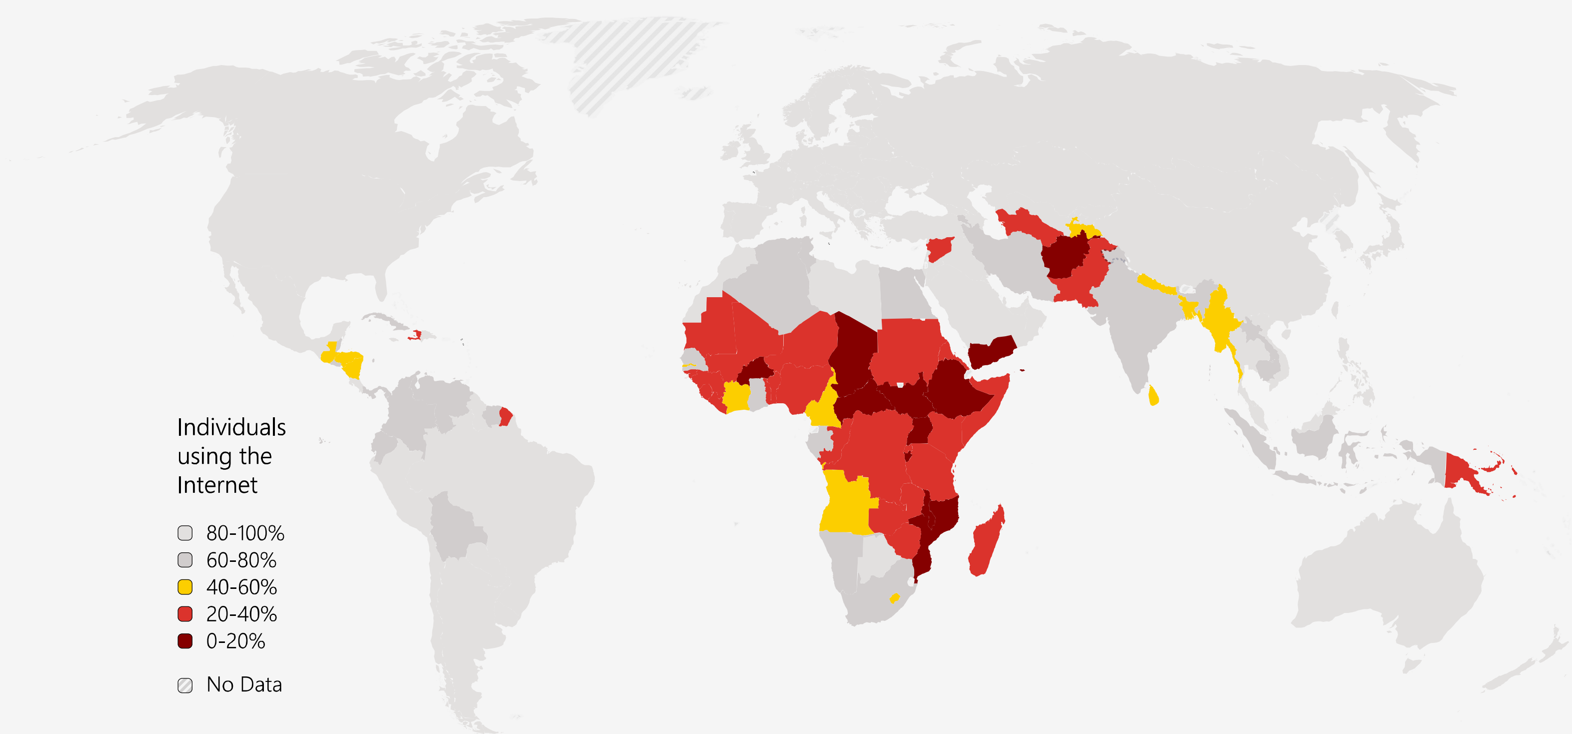

Data Source: Individuals using the Internet 2024, ITU. **[20]**

Digital and AI Skills

To fully participate in an AI-driven economy, people need the digital skills and technical proficiencies to use AI tools productively and responsibly. Without this foundation, AI risks becoming a technology that only a segment of society can leverage—deepening inequality rather than broadening opportunity.

Basic digital literacy is the first step: understanding how to navigate digital platforms, evaluate information, and engage safely online. But effective use of AI requires acquiring additional skills. People must develop fluency in how AI works, how it can be used in daily life and at work, and how to apply it creatively and critically.

For many, this means reskilling or upskilling to adapt to new job requirements, while for others it opens pathways to entirely new professions in fields like data science, software development, or AI engineering. Just as industrial machinery once required mechanical know-how, today's AI tools demand new forms of digital competency.

Global Digital Adoption by Economy

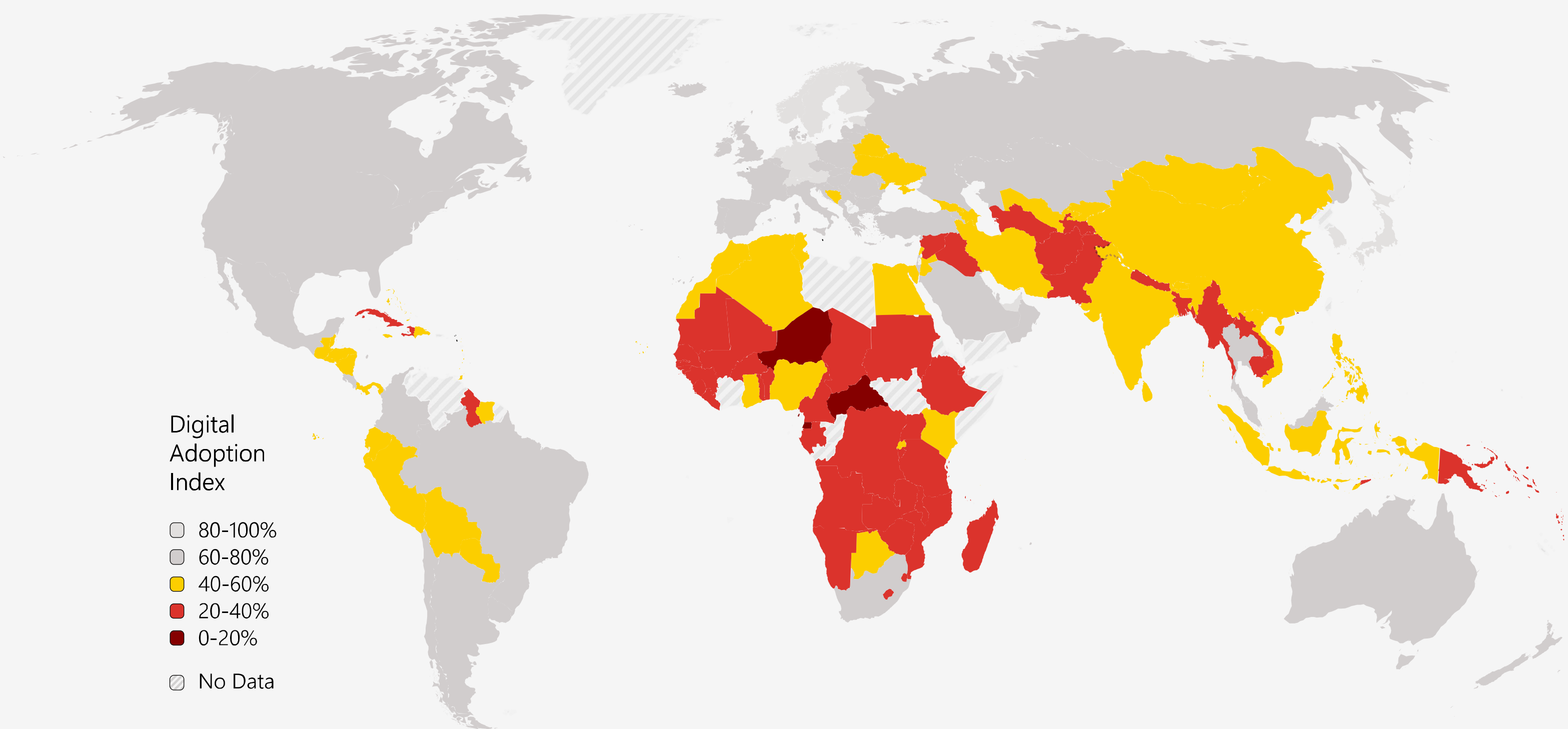

Data Source: World Bank Digital Adoption Index. [21]

Language

Unlike earlier technologies, AI is built on data—particularly human language. The open web, where most of humanity's knowledge is stored, serves as the single most important dataset for AI. Yet half of this content is in English, even though English is spoken natively only by just 5% of the world's population. **[22, 23]** For low-resource languages, AI models are often less proficient—or even nonexistent. **[24]**

This mismatch excludes millions: in Malawi, fewer than 4% speak English fluently, with the majority speaking Chichewa, Chitumbuka, or Chiyao, languages almost absent from the web and LLMs. **[25]** The result is a systemic barrier: even with access to electricity, the digital infrastructure and digital skills, AI remains inaccessible when people cannot engage with it in their own language.

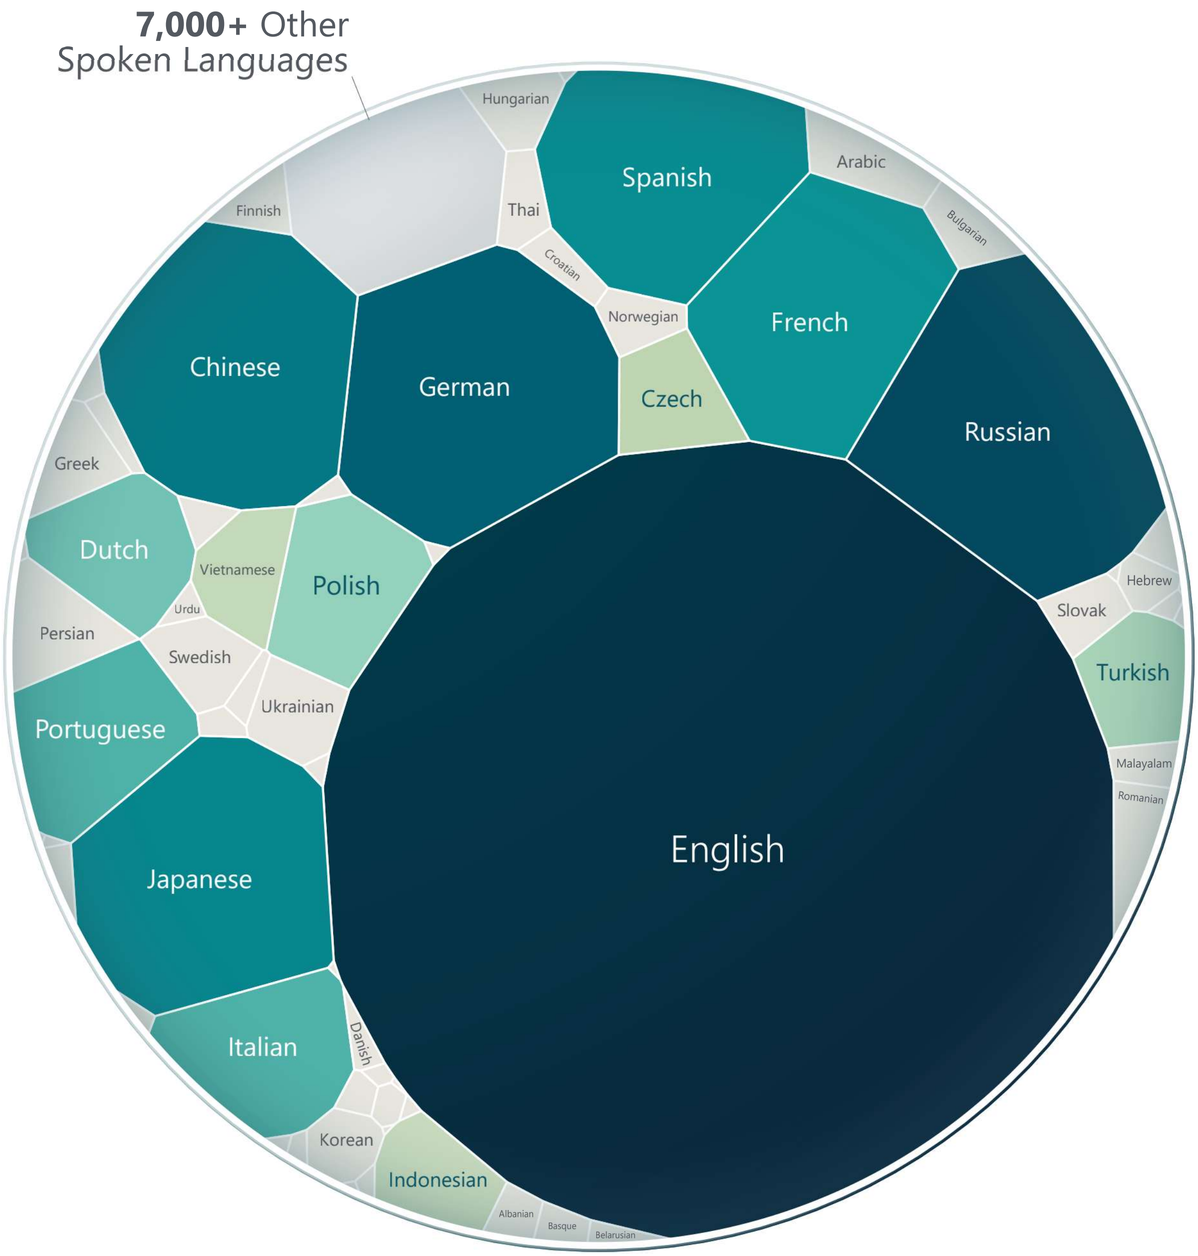

Data Source: Common Crawl Foundation **[23]**; UNESCO. **[26]**

The volume of open web content differs greatly between languages. Swahili, spoken by over 200 million people, has over 500 times less digital content than German, despite comparable numbers of speakers.[23] Our analysis shows that low-resource language countries adopt AI at rates 20% lower than high-resource language countries—even under similar GDP and connectivity conditions. [27] Model performance across languages, as shown in the study by Xuan et al, [28] indicates that even state-of-the-art LLMs achieve around 80% accuracy in English but drop below 55% for some low-resource languages such as Yoruba (YO), one of Nigeria's three major languages, spoken by more than 50 million people across Africa.

These gaps are not just technical; they determine who can access AI's benefits in education, health, agriculture, finance, and public services. For some communities, like Māori speakers in New Zealand, bilingualism with English makes the issue one of cultural preservation. But for others—such as Hausa in Nigeria and Niger, or Guarani in Paraguay—unlocking AI in local languages is the only path to meaningful access.

While the language divide remains a challenge, it is also one of AI's greatest opportunities. Large language models (LLMs) now make it possible to build high-quality translation systems, and even train new models, with far less data than before. Because LLMs learn shared semantic representations, knowledge gained in one language can benefit another, a phenomenon known as cross-lingual transfer. In this sense, bridging the language gap is not just a technical task, but a chance to ensure AI truly serves all of humanity, in every language.

Language Coverage in AI Resources

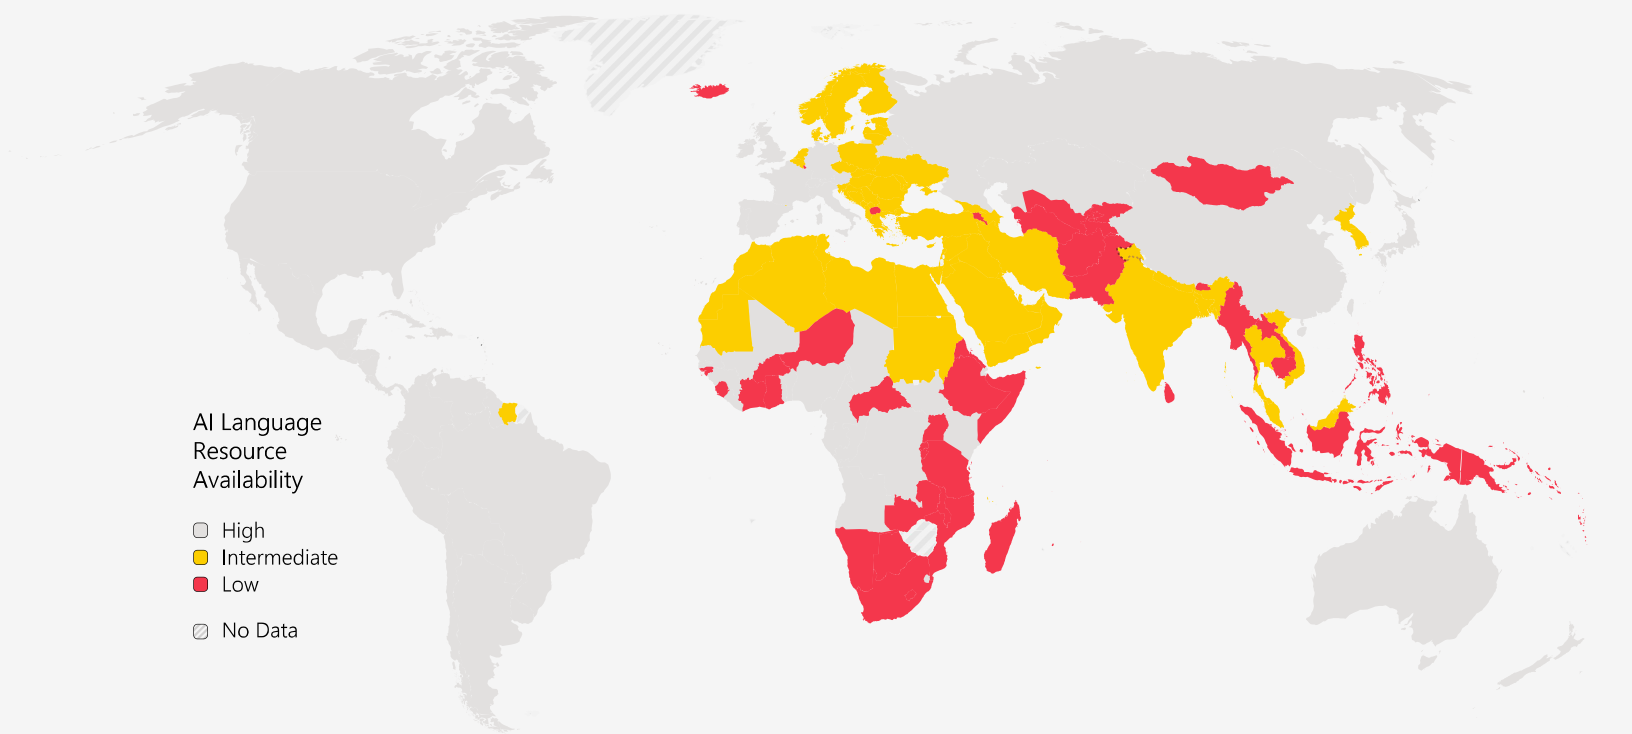

Data Source: AI Diffusion in Low-Resource Language Countries, AI for Good Lab. [27]

General Purpose Technologies Need Builders and Users to Evolve

Builders and Users: Two Sides of Every GPT

Every general-purpose technology depends on two groups: builders and users. Builders expand the technology's capabilities—creating, refining, and scaling it—while users apply it to solve real problems and push it into new contexts. This interplay creates ubiquity: as users experiment, they guide builders' focus, and as builders advance the frontier, they open new opportunities for users.

General-purpose technologies thrive when they're useful and flexible—meeting common needs while adapting to many applications. Computers, once confined to rooms, now exist all around us.

For AI, builders include frontier researchers and infrastructure providers, while users range from individuals to enterprises. Both are essential: without builders, progress stalls; without users, innovation loses direction.

Builders

On the builder side, there are two types of organizations: frontier builders and infrastructure builders.

Frontier Builders

Frontier Builders are organizations at the cutting edge of AI model development. They include the major AI labs designing foundation models (e.g. OpenAI, Microsoft, Anthropic,Google, Mistral, DeepSeek, Alibaba, etc.) who push the state-of-the-art in model architectures, algorithms, and training techniques.

In addition, an open-source AI community has emerged as a vital builder of the ecosystem. Open builders include researchers, hobbyists, and companies who release AI tools, models, and libraries under open-source licenses. They advance model accessibility and customization, often fine-tuning large models into smaller, more efficient versions for anyone to use.

For instance, Meta's release of the LlaMA open-weights models catalyzed a wave of community innovation – over 7,000 variant models were developed by the community on platforms like Hugging Face within months of LlaMA's debut. [29] Other labs followed: OpenAI now offers open-weights versions of GPT, and Google, through its Deep Mind subsidiary, offers the Gemma family of models.

To measure the frontier, we look at the provenance of major AI models and datasets. One of our first observations is that the number of models and their quality increased quickly over time, while the price per token decreased substantially from earlier models.

To calculate the frontier index, we first calculated a blended score of benchmarks:

- **Coding:** average of SciCode and LiveCodeBench benchmarks
- **Knowledge:** average of the MMLU-Pro and HLE benchmarks

AI Models by Performance, Cost, and Release Date

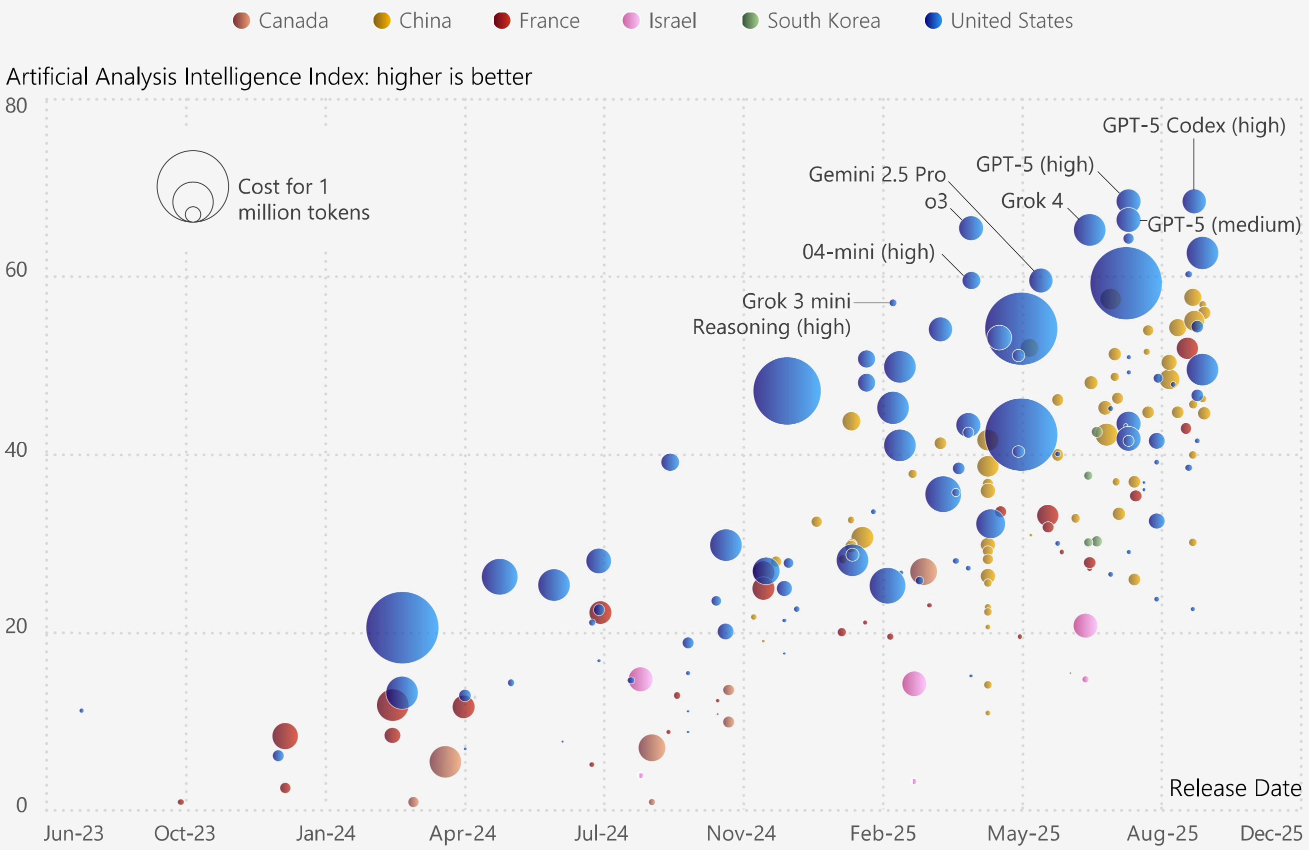

Data Source: artificialanalysis.ai [30]

- **Reasoning:** Long-context reasoning (LCR benchmark)
- **Instruction following:** Used the ifbench benchmark
- **Information retrieval:** Used the GQPA benchmark

Our metric is the simple average of the five scores above. Benchmarks we used include the following:

- SciCode – Tests scientific programming by having models write Python that passes unit tests on research-style tasks (scored by pass@1).
- LiveCodeBench – Evaluates general coding ability on live programming tasks with hidden tests; complements SciCode.
- MMLU-Pro – A tougher version of MMLU with ~12k questions across disciplines to assess broad knowledge and reasoning.
- HLE (Humanity's Last Exam) – Adversarial academic questions testing high-level reasoning; rewards semantically correct answers.
- AA-LCR – Measures long-context reasoning, requiring retrieval and synthesis across lengthy documents.
- IFBench – Assesses precise instruction following on open-ended tasks using rule-based evaluation.
- GPQA Diamond – Graduate-level science Q&A benchmark in biology, physics, and chemistry; designed to test deep reasoning without web retrieval.

Using the result of our metric, we calculated which model was the best performing at each given date in the past. Then, for each country, we selected the best-performing model they have ever created. Using that model's release date, we checked how long ago a model with the same performance was on the frontier. This determined our "months to frontier" metric. For example, when DeepSeek V3.1 was released in September 2025, its performance was closest to OpenAI's GPT-o3, which had been released in April, 5 months before.

Country	Best Model	Frontier Index	Months to Frontier
United States	GPT-5 (high)	1.000	0.0
China	DeepSeek V3.1 Terminus (Reasoning)	0.841	5.3
South Korea	EXAONE 4.0 32B (Reasoning)	0.824	5.9
France	Magistral Medium 1.2	0.789	7.0
United Kingdom	Gemma 3 27B Instruct	0.768	7.7
Canada	Command A	0.767	7.8
Israel	Jamba 1.7 Large	0.651	11.6

Infrastructure builders

AI infrastructure includes the physical and energy systems that enable model training and inference, mainly data centers, connectivity, and power. Because a data center requires both reliable electricity and network access, its presence serves as a strong proxy for infrastructure.

Not all data centers are equal; they differ in size, purpose, and compute capacity. The best indicator of scale is electricity use. While AI facilities increasingly rely on GPUs, the slowest part of infrastructure growth remains building and operating data centers.

To measure regional capacity, we use International Energy Agency (IEA) estimates of data center capacity in gigawatts, providing the most comprehensive mid-2025 view of global compute resources.

IEA Region	Installed Capacity (GW)
United States	53.7
China	31.9
European Union	11.9
Japan and South Korea	6.9
India	3.5
Other Asia Pacific	3.1
United Kingdom	2.6
Australia and New Zealand	1.6
Other North America	1.5
Africa	1.5
Central and South America	1.4
Eurasia	1.2
Middle East	1.1
Brazil	0.6

AI Users (Diffusion)

Measuring AI diffusion—the extent to which AI is adopted and used across sectors and populations—is inherently complex. Traditional survey-based approaches often lack precision or coverage.

Microsoft holds a unique vantage point to assess this diffusion at scale. By analyzing aggregated and anonymized telemetry from over one billion Windows devices, we can estimate the prevalence of AI-related activity across regions. Although this dataset excludes non-Windows devices, we corrected for this limitation by incorporating third-party data on Windows' market share and the relative distribution of desktop versus mobile usage.

While not perfect, this methodology provides a robust and consistent proxy for measuring real-world AI diffusion globally over time. This data includes data from all major AI models

What Can We Learn From the Data Cross-Analysis

- **Fastest Adoption in History:**
  With more than 1.2 billion users in under 36 months, AI has become the fastest-adopted technology in human history. The data show that AI adoption is not difficult when people have access to a connected computer or smartphone. However, once electricity, connectivity, and digital skills are factored in, nearly half of humanity, or about four billion people—still lack the basic capabilities needed to use AI. While the next billion users may come relatively easily, progress will plateau as we reach populations constrained by these structural barriers.

- **North–South Divide:**
  The AI diffusion map reveals a stark contrast between the Global North and the Global South. AI adoption in the Global North is approximately 23%, compared with only 13% in the Global South. Adoption correlates strongly with GDP per capita, with the largest disparities emerging among countries below $20,000 GDP per capita—a range that includes all least-developed countries (LDCs).

- **Language Barrier:**
  Countries where low-resource languages are predominant exhibit significantly lower AI adoption, even after controlling for GDP and internet access. This suggests that language inclusion is an independent and powerful driver of AI diffusion.

- **AI Diffusion Leaders:**
  The highest rates of AI adoption are observed in the UAE, Singapore, Norway, and Ireland—countries that aren't all traditionally known as AI powerhouses. While the UAE has begun developing its own models, these are not yet at the technological frontier, indicating that strong infrastructure, policy coordination, and digital readiness can drive rapid adoption even without frontier development.

- **Infrastructure Concentration:**
  Datacenter capacity remains heavily concentrated, with the United States and China accounting for roughly 86% of global compute. These same countries also lead in the number and performance of frontier AI models, ranking #1 and #2 respectively.

- **Frontier Narrowing:**
  Only seven countries currently host frontier-level AI models—the U.S., China, South Korea, France, the U.K., Canada, and Israel—yet the performance gap is shrinking. Based on our comparative metrics, China trails the U.S. by less than six months, and Israel, ranked seventh, is less than a year behind the frontier. This acceleration suggests that diffusion at the frontier is faster than in prior technological revolutions.

AI Diffusion Data Source

Economy	AI Diffusion
United Arab Emirates	59.4%
Singapore	58.6%
Norway	45.3%
Ireland	41.7%
France	40.9%
Spain	39.7%
New Zealand	37.6%
United Kingdom	36.4%
Netherlands	36.3%
Qatar	35.7%
Australia	34.5%
Israel	33.9%
Belgium	33.5%
Canada	33.5%
Switzerland	32.4%
Sweden	31.2%
Austria	29.1%
Hungary	27.9%
Denmark	26.6%
Germany	26.5%
Poland	26.4%
Taiwan	26.4%
United States	26.3%
Czechia	26.0%
South Korea	25.9%
Italy	25.8%
Finland	25.6%
Bulgaria	25.4%
Jordan	25.4%
Costa Rica	25.1%
Lebanon	24.8%
Slovenia	24.6%
Saudi Arabia	23.7%
Oman	22.6%
Portugal	22.4%
Jamaica	22.2%
Slovakia	22.1%

Economy	AI Diffusion
Dominican Republic	22.0%
Croatia	21.8%
Vietnam	21.2%
Lithuania	21.0%
Uruguay	20.9%
Colombia	20.4%
Panama	20.3%
Serbia	19.7%
Chile	19.6%
South Africa	19.3%
Malaysia	18.3%
Bosnia and Herzegovina	18.2%
Argentina	17.8%
Greece	17.7%
Kuwait	17.7%
Georgia	17.3%
Philippines	17.1%
Ecuador	17.0%
Japan	16.7%
Mexico	16.7%
Moldova	16.6%
Albania	15.8%
Brazil	15.6%
China	15.4%
Romania	15.3%
El Salvador	14.6%
Azerbaijan	14.2%
India	14.2%
Guatemala	13.7%
Peru	13.4%
Türkiye	13.4%
Namibia	13.0%
Botswana	12.8%
Kazakhstan	12.7%
Libya	12.7%
Mongolia	12.6%
Egypt	12.5%

Economy	AI Diffusion
Honduras	12.4%
Senegal	12.4%
Gabon	12.3%
Nepal	12.3%
Tunisia	12.3%
Indonesia	11.7%
Zambia	11.7%
Algeria	11.3%
Bolivia	10.9%
Côte d'Ivoire	10.8%
Gambia	10.6%
Morocco	10.5%
Iraq	10.3%
Paraguay	10.1%
Nicaragua	10.0%
Pakistan	9.7%
Iran	9.6%
Thailand	9.1%
Ukraine	9.1%
Angola	8.9%
Madagascar	8.9%
Mozambique	8.9%
Malawi	8.9%
Lesotho	8.8%
Benin	8.7%
Burkina Faso	8.7%
Ghana	8.7%
Guinea	8.7%
Guinea-Bissau	8.7%
Liberia	8.7%
Mali	8.7%
Mauritania	8.7%
Niger	8.7%
Nigeria	8.7%
Sierra Leone	8.7%
Togo	8.7%
Myanmar	8.4%
French Guiana	8.3%
Guyana	8.3%
Suriname	8.3%

Economy	AI Diffusion
Venezuela	8.3%
Kenya	7.8%
Russia	7.6%
Belarus	7.6%
Kyrgyzstan	7.6%
Papua New Guinea	7.2%
Haiti	7.1%
Central African Republic	7.0%
Cameroon	7.0%
Congo (DRC)	7.0%
Congo	7.0%
Chad	7.0%
Zimbabwe	6.9%
Syria	6.7%
Bangladesh	6.5%
Burundi	6.4%
Eritrea	6.4%
Ethiopia	6.4%
Sudan	6.4%
Somalia	6.4%
South Sudan	6.4%
Tanzania	6.4%
Uganda	6.4%
Armenia	6.2%
Sri Lanka	6.2%
Laos	6.0%
Rwanda	6.0%
Cuba	5.7%
Uzbekistan	5.7%
Afghanistan	5.1%
Tajikistan	5.1%
Turkmenistan	5.1%
Cambodia	4.6%

References

[1] B. Eugene Nancy M. ,Campos, Jose Edgardo L. ,Kim, Chang-Shik,Corden, W. Max,MacDonald, Lawrence [editor],Pack, Howard,Page, John,Sabor, Richard,Stiglitz, Joseph, "The East Asian miracle : economic growth and public policy (Vol. 1 of 2) : Main report," World Bank. Accessed: Sept. 10, 2025. [Online]. Available: https://documents.worldbank.org/en/publication/documents-reports/documentdetail/975081468244550798

[2] World Bank, "GDP per capita, PPP (current international $)," World Development Indicators, World Bank Open Data, May 30, 2024. [Online]. Available: https://data.worldbank.org/indicator/NY.GDP.PCAP.PP.CD

[3] "Opinion | You Are What You Spend - The New York Times." Accessed: Sept. 11, 2025. [Online]. Available: https://www.nytimes.com/2008/02/10/opinion/10cox.html

[4] H. Ritchie, E. Mathieu, M. Roser, and E. Ortiz-Ospina, "Internet," Our World in Data, Apr. 2023, Accessed: Sept. 11, 2025. [Online]. Available: https://ourworldindata.org/internet

[5] A. Misra and J. Wang et al., "Measuring AI Diffusion: A Population-Normalized Metric for Tracking Global AI Usage.", Microsoft AI for Good Lab, 2025. https://aka.ms/AI_Diffusion_Technical_Report

[6] "World Development Indicators | DataBank." Accessed: Sept. 11, 2025. [Online]. Available: https://databank.worldbank.org/reports.aspx?source=2&series=NY.GDP.PCAP.CD&country=

[7] "Taiwan GDP per Capita: USD | Economic Indicators | CEIC." Accessed: Sept. 11, 2025. [Online]. Available: https://www.ceicdata.com/en/taiwan/sna-08-reference-year2021-gdp-by-expenditure-current-price-annual/gdp-per-capita-usd

[8] "Our milestones," Smart Nation Singapore. Accessed: Sept. 11, 2025. [Online]. Available: https://www.smartnation.gov.sg/about/milestones/

[9] "(PDF) STEM Education in Singapore," in ResearchGate. doi: 10.1007/978-981-16-1357-9_3.

[10] World Bank, "Population, total," World Development Indicators, World Bank Open Data, May 30, 2024. [Online]. Available: https://data.worldbank.org/indicator/SP.POP.TOTL

[11] "Share of the population with access to electricity," Our World in Data. Accessed: Sept. 12, 2025. [Online]. Available: https://ourworldindata.org/grapher/share-of-the-population-with-access-to-electricity

[12] "Facts and Figures 2024 - Internet use." Accessed: Sept. 12, 2025. [Online]. Available: https://www.itu.int/itu-d/reports/statistics/2024/11/10/ff24-internet-use

[13] ITU. (n.d.). SDG-related ICT indicators. International Telecommunication Union. Retrieved October 29, 2025, from https://www.itu.int/en/ITU-D/Statistics/Pages/SDGs-ITU-ICT-indicators.aspx

[14] "Tracking SDG 7 | Progress Towards Sustainable Energy." Accessed: Sept. 15, 2025. [Online]. Available: https://trackingsdg7.esmap.org/

[15] "ieg.worldbankgroup.org/sites/default/files/Data/reports/Ap_WBG_Elec_Africa_2015-24.pdf?utm_source=chatgpt.com." Accessed: Sept. 14, 2025. [Online]. Available: https://ieg.worldbankgroup.org/sites/default/files/Data/reports/Ap_WBG_Elec_Africa_2015-24.pdf

[16] "Access to electricity – SDG7: Data and Projections – Analysis," IEA. Accessed: Sept. 10, 2025. [Online]. Available: https://www.iea.org/reports/sdg7-data-and-projections/access-to-electricity

[17] "Share of the population with access to electricity," Our World in Data. Accessed: Sept. 10, 2025. [Online]. Available: https://ourworldindata.org/grapher/share-of-the-population-with-access-to-electricity?tab=line&country=~OWID_WRL

[18] O. Nilsson and N. Yngwe, "API Latency and User Experience: What Aspects Impact Latency and What are the Implications for Company Performance?".

[19] Z. Hawkins, V. Lehdonvirta, and B. Wu, "AI Compute Sovereignty: Infrastructure Control Across Territories, Cloud Providers, and Accelerators," June 20, 2025, Social Science Research Network, Rochester, NY: 5312977. doi: 10.2139/ssrn.5312977.

[20] "Individuals using the Internet - ITU DataHub." Accessed: Sept. 11, 2025. [Online]. Available: https://datahub.itu.int/data/?i=11624

[21] "Digital Adoption Index," World Bank. Accessed: Sept. 10, 2025. [Online]. Available: https://www.worldbank.org/en/publication/wdr2016/Digital-Adoption-Index

[22] "Languages by number of native speakers | List, Top, & Most Spoken | Britannica." Accessed: Sept. 10, 2025. [Online]. Available: https://www.britannica.com/topic/languages-by-number-of-native-speakers-2228882

[23] "Statistics of Common Crawl Monthly Archives by commoncrawl." Accessed: Sept. 11, 2025. [Online]. Available: https://commoncrawl.github.io/cc-crawl-statistics/plots/languages

[24] "Mind the (Language) Gap: Mapping the Challenges of LLM Development in Low-Resource Language Contexts | Stanford HAI." Accessed: Sept. 10, 2025. [Online]. Available: https://hai.stanford.edu/policy/mind-the-language-gap-mapping-the-challenges-of-llm-development-in-low-resource-language-contexts

[25] E. Kayambazinthu, "The Language Planning Situation in Malawi," Journal of Multilingual and Multicultural Development, vol. 19, no. 5, pp. 369–439, Sept. 1998, doi: 10.1080/01434639808666363.

[26] "Multilingual education, the bet to preserve indigenous languages and." Accessed: Sept. 12, 2025. [Online]. Available: https://www.unesco.org/en/articles/multilingual-education-bet-preserve-indigenous-languages-and-justice

[27] A. Misra, S. W. Zamir, W. Hamidouche and J. Lavista Ferres, "AI Diffusion in Low Resource Language Countries", Microsoft AI for Good Lab, 2025. https://aka.ms/AI_Diffusion_Low_Resource_Language_Countries

[28] W. Xuan et al, "MMLU-ProX: A Multilingual Benchmark for Advanced Large Language Model Evaluation," arXiv preprint arXiv:2503.10497, 2025. [Online]. Available: https://arxiv.org/abs/2503.10497. Accessed: Oct. 29, 2025.

[29] S. Goldman, "Forget ChatGPT, why Llama and open source AI win 2023," VentureBeat. [Online]. Available: https://venturebeat.com/ai/forget-chatgpt-why-llama-and-open-source-ai-win-2023. Accessed: Oct. 29, 2025.

[30] Artificial Analysis, "AI Model & API Providers Analysis," ArtificialAnalysis.ai, 2025. [Online]. Available: https://artificialanalysis.ai/.

Microsoft