Microsoft | AI Economy Institute

Global AI Diffusion

Q2 2026 Trends and Insights



September, 2026

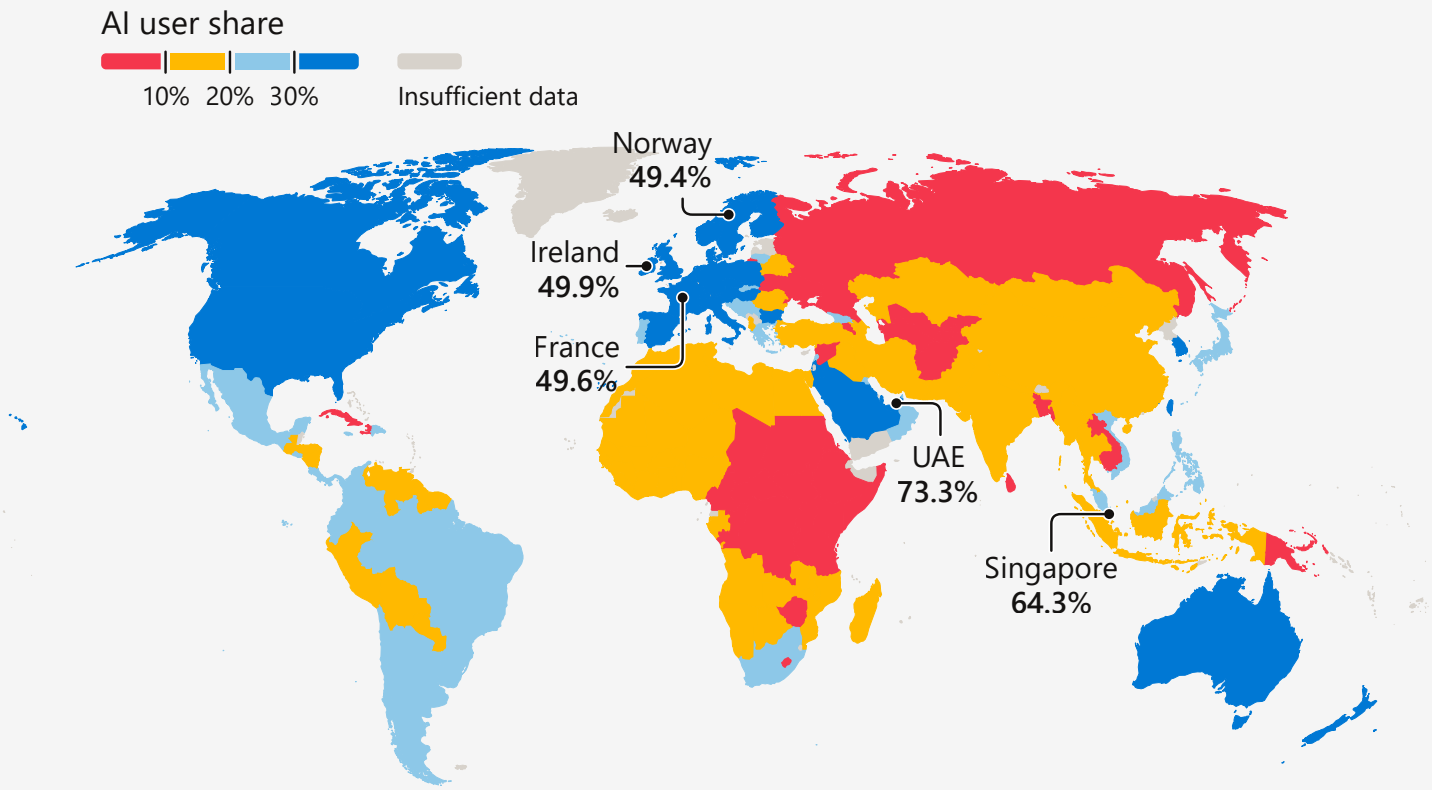
Executive Summary

Global AI diffusion continued to climb in the second quarter of 2026. In June, AI usage was 18.8% of the working-age population worldwide, an increase of about 1% from the first quarter of this year. Growth was broad-based: almost every economy saw increased usage.

Notably, the UAE (73.3%) and Singapore (64.3%) continue to top the leaderboard. South Korea again saw the largest absolute change, with an increase of 3.5 percentage points. Saudi Arabia climbed five places from 30th to 25th, the largest increase in rank. Japan recorded the largest relative change, growing 2.2 percentage points, or roughly 10% relative to its Q1 usage of 22.5%. Worldwide, the quarter brought continued widening of the AI gap between the Global North and South, with usage now at 28.8% in the North and 16.2% in the South.

To track all these trends, we continue to measure AI diffusion as the share of people worldwide between ages of 15 and 64 who have used a generative AI product during the reported period. This measure is derived from aggregated and anonymized Microsoft telemetry data and is adjusted to reflect differences in operating system and device-market share, internet penetration, and population. Additional details on the methodology are available in our AI Diffusion technical paper. **[1]**

AI diffusion by economy, June 2026



No single metric is perfect, and this one is no exception. Through the Microsoft AI Economy Institute, we continue to refine how we measure AI diffusion globally, including how adoption varies across countries in ways that best advance priorities such as scientific discovery and productivity gains. For this report, we rely on the strongest cross-country measure available today. Over the past year, there has been a significant rise in new AI tools and models. For our next report we will expand our measure to include usage of these new tools and expect this update to increase AI User Share across almost all economies, with larger projected increases in China.

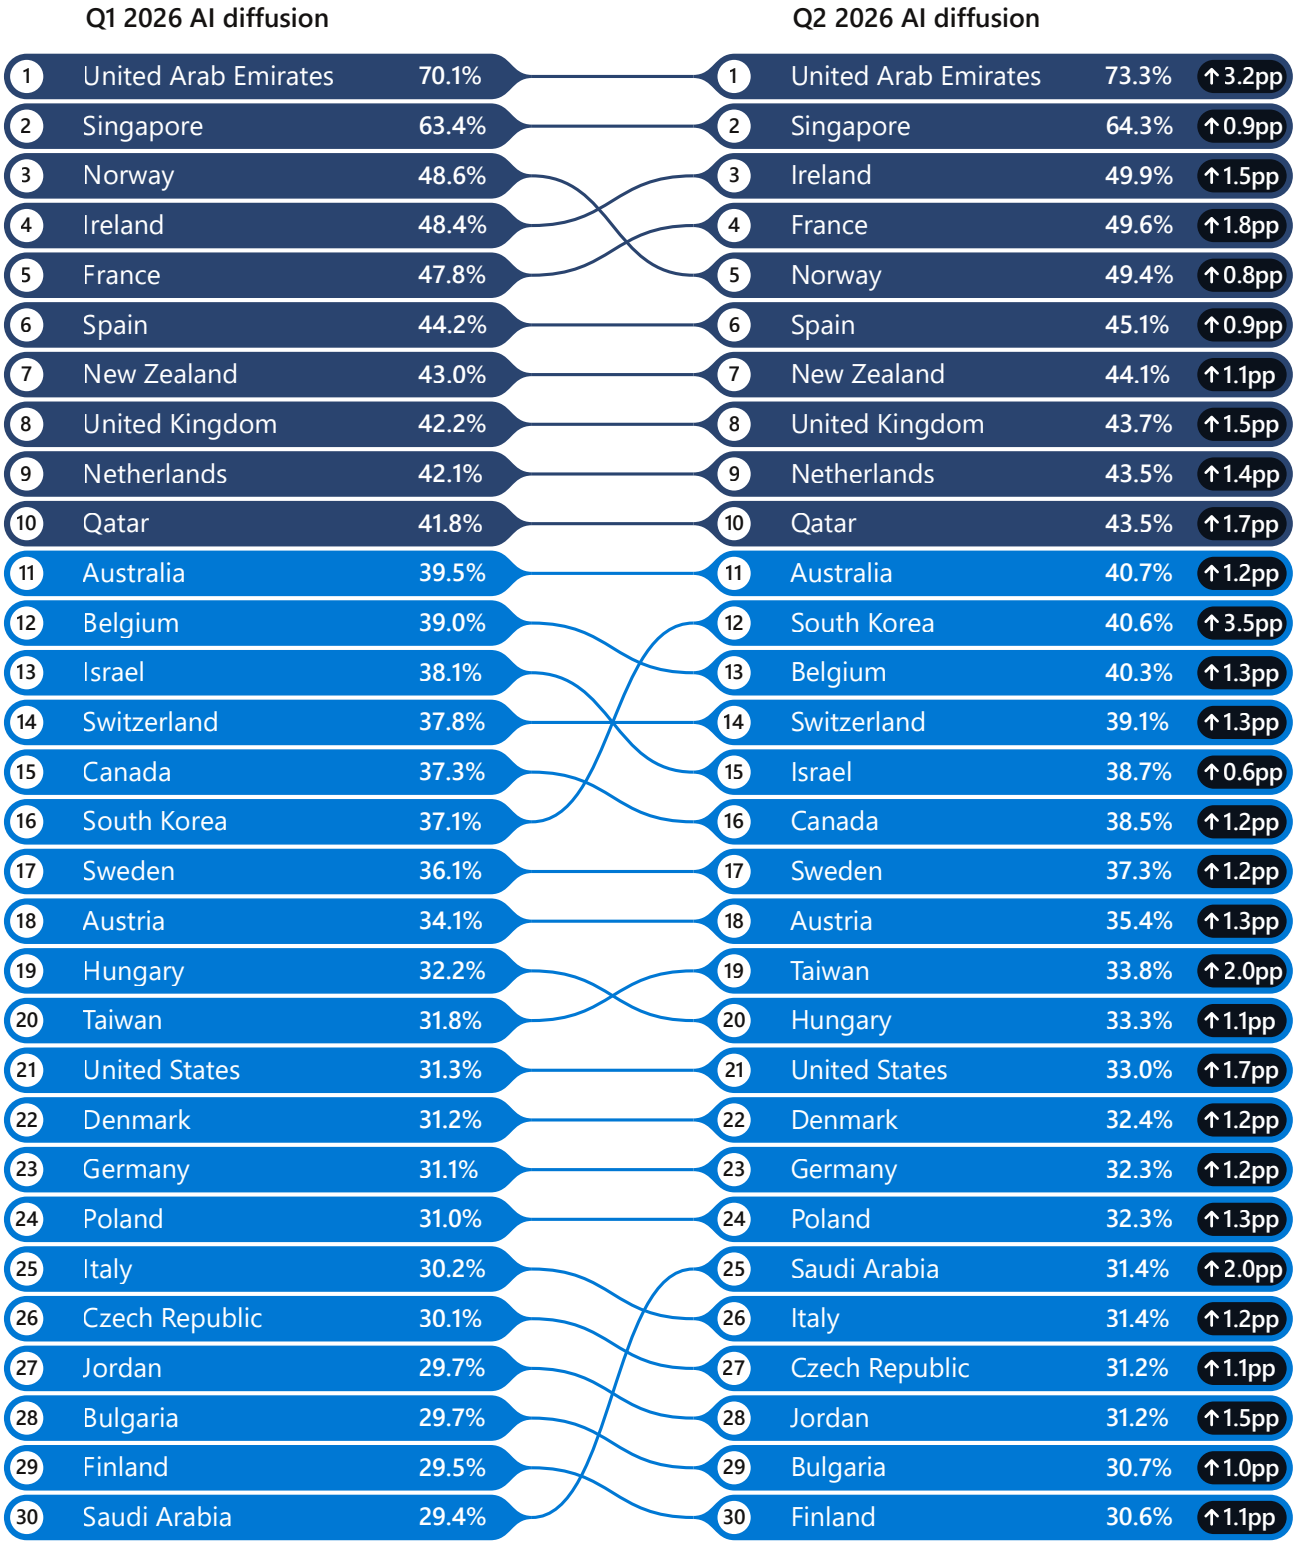
In addition to our comparison of AI diffusion between the Global North and the Global South, we also find important differences in how AI is being used. Using Consumer Copilot usage insights, we examined these differences. The most notable difference is a higher share of education and learning use in the Global South.

This report also includes a substantive discussion of open-weight AI models. Open-weight AI model use is increasingly visible through application programming interface (API) usage of AI, where open models have captured a growing share of token volume. This shift could especially benefit the Global South by lowering access costs and enabling adaptation to local languages and needs, although infrastructure, connectivity, skills, and other foundational barriers remain.

Leading Economies

The United Arab Emirates (73.3%) and Singapore (64.3%) again lead the world by a wide margin, followed by Ireland (49.9%) and France (49.6%). Within the top 25, South Korea was again the standout riser, gaining 3.5 percentage points, the largest quarterly increase in the group, and climbing four places from 16th to 12th. South Korea has now posted the largest gain of any major economy for three consecutive reporting periods from H2 2025 to Q1 2026 to Q2 2026, adding nearly 15 points to its AI user share over the past 12 months. Saudi Arabia saw the largest rise in rank, climbing five places from 30th to 25th and entering the top 25.

AI diffusion over time by economy



The Global North-South divide

The underlying gap between the Global North and Global South continued to widen in Q2 2026. In the second quarter of 2026, 28.8% of the population in the Global North used generative AI, up from 27.5% in Q1 2026, a gain of 1.3 percentage points. In the Global South, usage reached 16.2%, up from 15.4%, a gain of 0.8 percentage points.

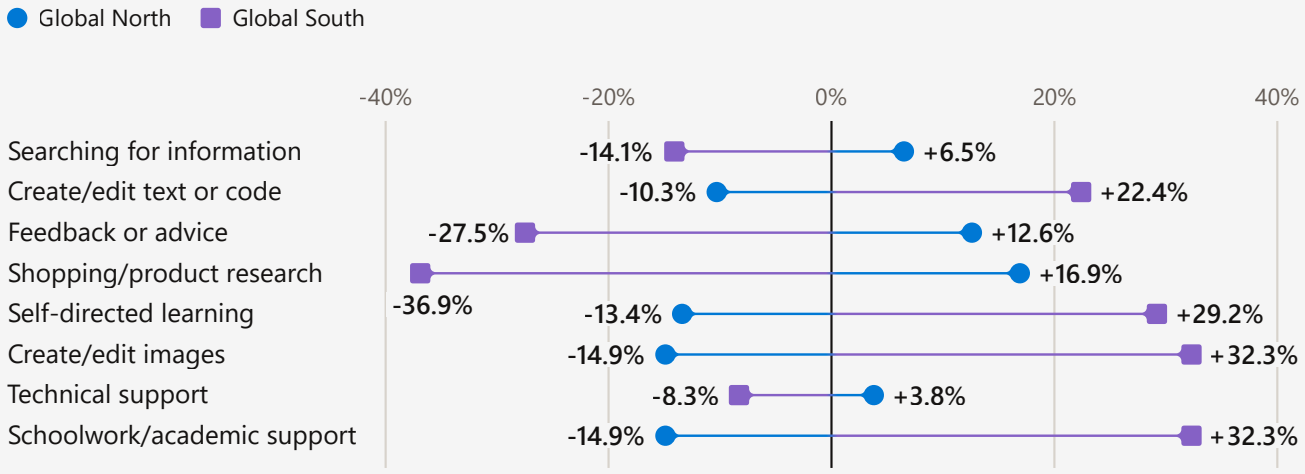
Underlying this gap in adoption are important differences in how AI is being used. Learning stands out in particular in the Global South. Our analysis of a subset of consumer Copilot conversations uses a privacy-preserving classifier to identify the broad "intent" of a conversation – the primary task in which Copilot users engage. **[2]** We grouped AI usage by economy into the Global North and South, as above, and compared the share of conversations for each intent with its share globally.

In the Global South, self-directed learning and skill development represent a 29.2 percent larger share of conversations than they do globally, and schoolwork and academic support accounts for a 32.3 percent

Region	H2 2025 AI Diffusion	Q1 2026 AI Diffusion	Q2 2026 AI Diffusion	Change
Global North	24.7%	27.5%	28.8%	**+1.3**
Global South	14.1%	15.4%	16.2%	**+0.8**
World	16.3%	17.8%	18.8%	**+1.0**
Gap between Global North and Global South	**10.6pp**	**12.1pp**	**12.6pp**	

Global North and South intent prevalence relative to the global mix

Intents ordered by global prevalence; North is 68.5% of traffic and South is 31.5%



larger share. Importantly, the Global South has a much younger population than the North. Therefore, it is likely that demographic differences are at least partially driving these results.

The remaining intents tilt in more familiar directions. The Global North leans toward consumer uses such as shopping, search, and feedback. And in the Global South, image, text, and code creation and editing take an outsized share. Taken together, these results suggest that learning support may be a particularly important entry point to AI in regions that are earlier in their adoption curve. This is consistent with recent studies of Bing Copilot, Gemini, and Claude, which also find schooling or coursework to be a larger part of observed AI use in lower-income countries. **[3][4][5]**

However, we suggest cautious interpretation of these findings. These results apply only to the early adopters using AI in these regions, and may not extrapolate as AI usage expands to the broader population. Additionally, the data do not explain why learning features more prominently in the Global South. Even so, the pattern adds a dimension to the North-South gap discussed above: in regions where diffusion still lags, a larger share of existing use is devoted to learning.

Open-weight models are broadening AI Diffusion

Many AI models are "open" in the sense that they release their trained weights, though the training process and code may or may not be open. These open-weight models can be downloaded, modified, and run by other organizations, giving developers much greater freedom to deploy and adapt them than closed models accessed only through a proprietary service. Their growing importance was reflected in a July 2026 open letter signed by Microsoft, NVIDIA, Meta, IBM, Hugging Face, the Linux Foundation, and other companies and organizations, which argued that open weights can make advanced AI more accessible and adaptable while allowing organizations to match "the right model to the right job at the right cost." **[6]**

Open-weight models also highlight an important dimension of AI diffusion that is difficult to capture through consumer-facing measures alone. Currently, our primary AI diffusion measure in this report estimates the share of the working-age population that has used a generative AI product. But open-weight models are especially likely to be encountered indirectly: through an application, coding agent, or other service that accesses a model programmatically through an API, rather than through a branded chatbot (though some chatbots do build on open models). Some organizations also download and run open weights on their own infrastructure, where their usage may not appear in any centralized data source. This makes the API market a useful complementary window into diffusion. One of the best public sources of data on open-weight model usage is OpenRouter, which provides a common interface through which developers can access hundreds of open and closed models and route requests among them. **[7]** Its position between users and many different inference providers allows an unusual view of which models developers are choosing in practice, although its developer-oriented user base should not be interpreted as representative of all AI users.

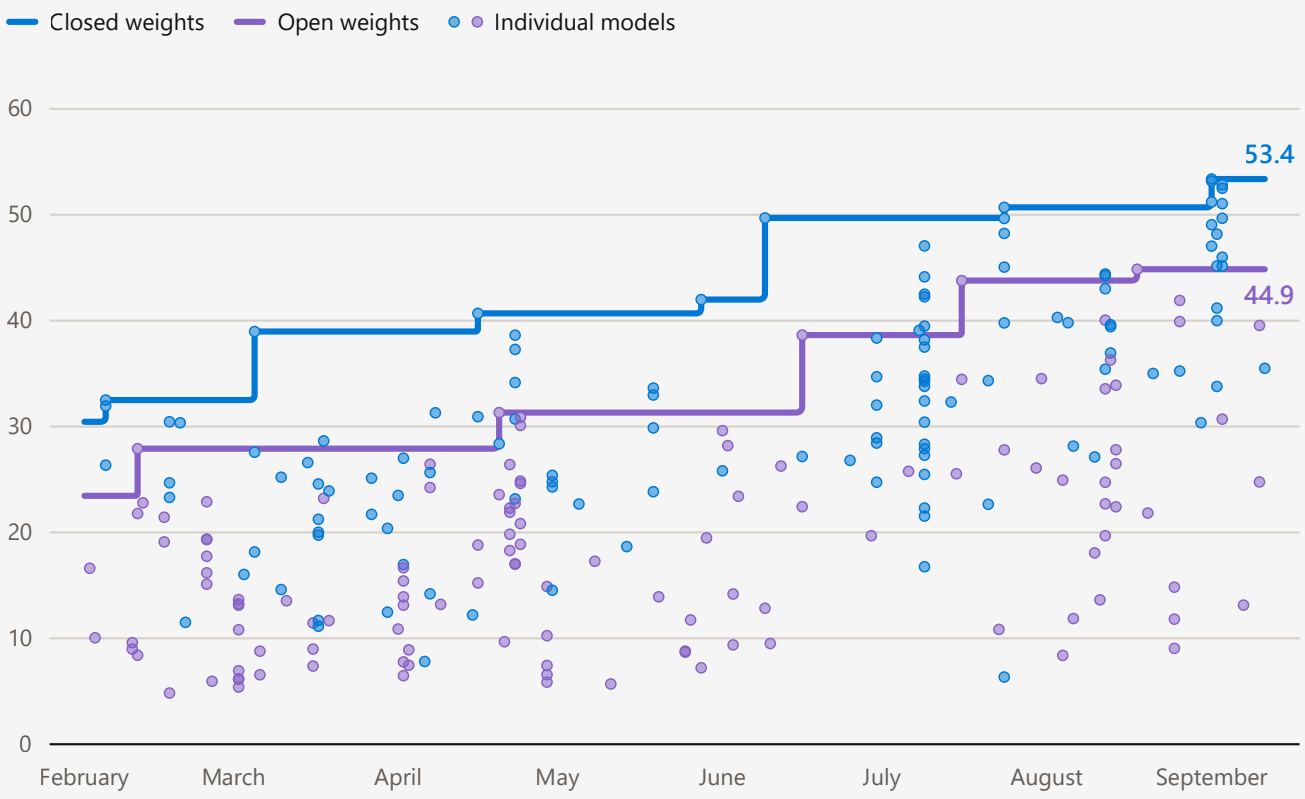
This market is becoming increasingly important because the capability gap between open-weight and closed models has narrowed rapidly. Research comparing successive generations of models across independent

benchmarks shows a recurring pattern: new closed models push the capability frontier forward, but leading open-weight models often reach similar performance within a matter of months. **[8]** Recent releases suggest that even this distinction is becoming increasingly task-specific. Moonshot AI's Kimi K3, whose full weights were released in July, still trails the strongest proprietary models on some broad evaluations, including in Moonshot's own technical report. **[9]** But on other measures it is already competitive with—or ahead of—leading closed models. As of September 14, for example, Kimi K3 Max ranked in the top 5 on Arena's WebDev coding leaderboard, ahead of a number of proprietary alternatives at a fraction of the cost. **[10]** These rankings should not be interpreted as showing that Kimi K3 is universally superior, but they illustrate how quickly open-weight models have moved from lower-cost alternatives to competitors at parts of the capability frontier.

That improvement in capability is increasingly translating into usage. OpenRouter's public analysis of more than 100 trillion tokens found that open-weight models had grown to roughly one-third of token volume by late 2025, with major releases such as DeepSeek, Kimi, Qwen, and GPT-OSS producing increases that persisted beyond their initial launch periods—evidence that adoption reflected production use rather than only experimentation. **[11]** The shift accelerated further in 2026. More recently our own analysis of public OpenRouter data shows that open-weight models now represent 75% of OpenRouter token usage. These figures are not global market shares: OpenRouter users are unusually likely to experiment across providers, while direct usage of closed frontier models as well as locally deployed open models are missing from the data entirely. Nevertheless, the magnitude and direction of the change provide a

Open-weight models catch up to the closed frontier within a few months

Artificial Analysis Intelligence Index by model release date



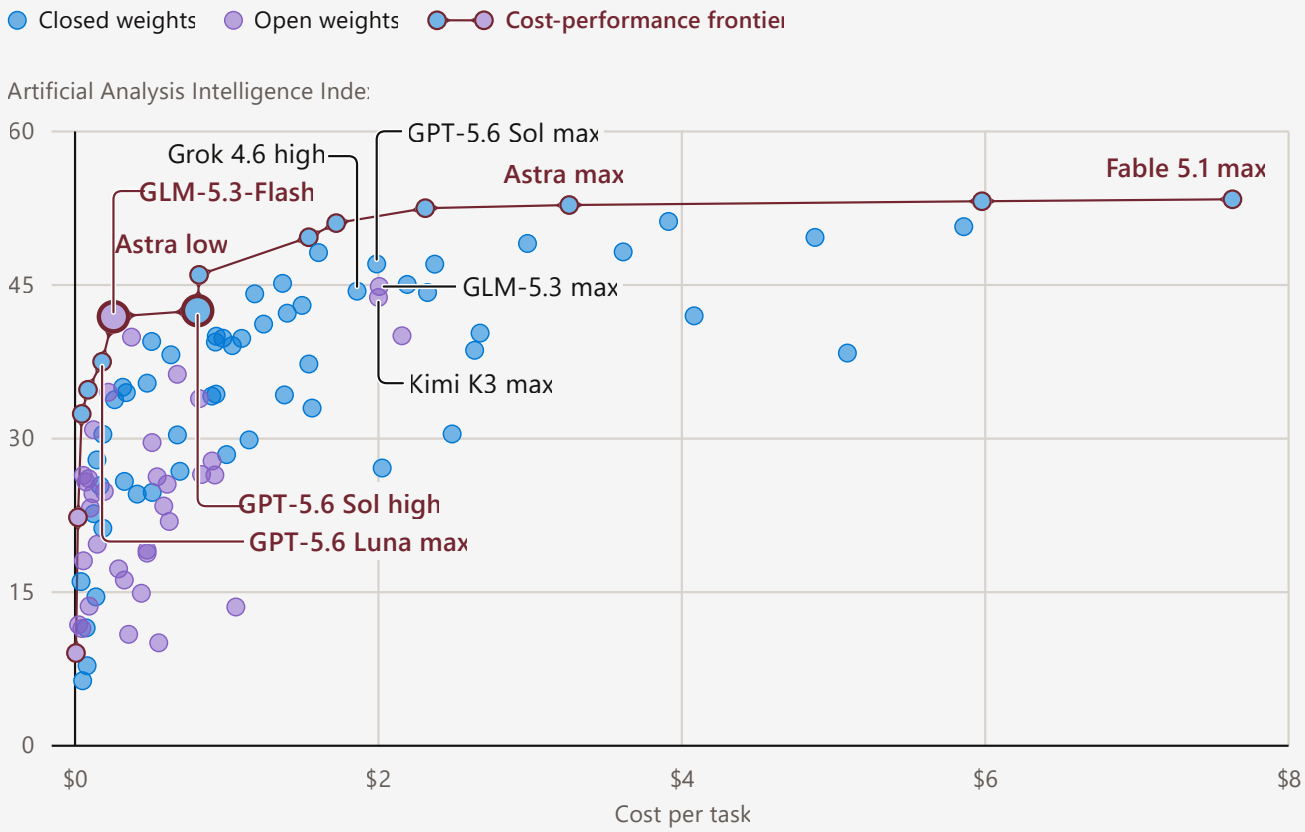
Sources: Microsoft AI Economy Institute; Artificial Analysis.

clear signal that open-weight models are becoming an increasingly important part of the API-based AI economy. This pattern is also consistent with emerging research using more granular OpenRouter data, which finds a broad movement toward greater use of open models as their capabilities improve. **[12]**

The economics of that shift could become increasingly important as AI moves from experimentation to large-scale deployment. Analyses show that the capability gap between open and closed models is not only decreasing, but is decreasing at a more rapid pace. **[12]** Further, the most recent public data from Artificial Analysis shows that some open-weight models deliver similar capabilities at substantially lower cost per task. For example, the recently released GLM 5.3 Flash has similar performance to GPT 5.6 Sol at 'high' reasoning level at roughly one third the cost. The Sol models from OpenAI were, as of just a few weeks ago, near the frontier for intelligence. At the same time, businesses are becoming more attentive to the cost of rapidly growing AI workloads. The once-popular practice sometimes described as "tokenmaxxing"—treating rising token consumption as a measure of successful AI adoption—can become expensive as agents and other applications make repeated model calls. Companies are therefore increasingly routing different tasks to different models rather than automatically using the most capable frontier system for every request; Reuters reported open-weight models gaining ground as businesses scrutinized AI spending and searched for less costly alternatives. **[13]** The emerging principle is not simply to minimize token use or always choose the cheapest model, but to match the model to the task: reserving expensive frontier intelligence for problems where its additional capability creates value while using lower-cost models for the much larger set of workloads where they perform sufficiently well. This is the same economic logic emphasized in the recent open-weight letter.

Near-frontier intelligence, far lower cost per task

Artificial Analysis Intelligence Index vs average cost per benchmark task



Sources: Microsoft AI Economy Institute; Artificial Analysis.

These economics could be particularly consequential for the Global South. As discussed earlier in this report, AI adoption reached 28.8 percent of the working-age population in the Global North in the second quarter of 2026 compared with 16.2 percent in the Global South, where gaps in electricity, connectivity, digital skills, and local-language support continue to constrain diffusion. Open-weight models cannot eliminate these foundational barriers. But they can reduce two additional ones: the cost of accessing capable AI and the difficulty of adapting models to local needs. Because weights can be modified, non-profits, businesses, governments, and universities can fine-tune or otherwise adapt existing models rather than train frontier systems from scratch. For example, companies have fine-tuned open-weight models to increase their expertise in specific domains like law. **[14]** Further, research from Microsoft's AI for Good Lab illustrates the potential of this approach: its Bring Your Own Language framework improved performance for Chichewa and Māori by approximately 12 percent on average across a set of benchmarks by refining data and adapting existing language models, while also developing a pathway for extremely low-resource languages such as Inuktitut. **[15]**

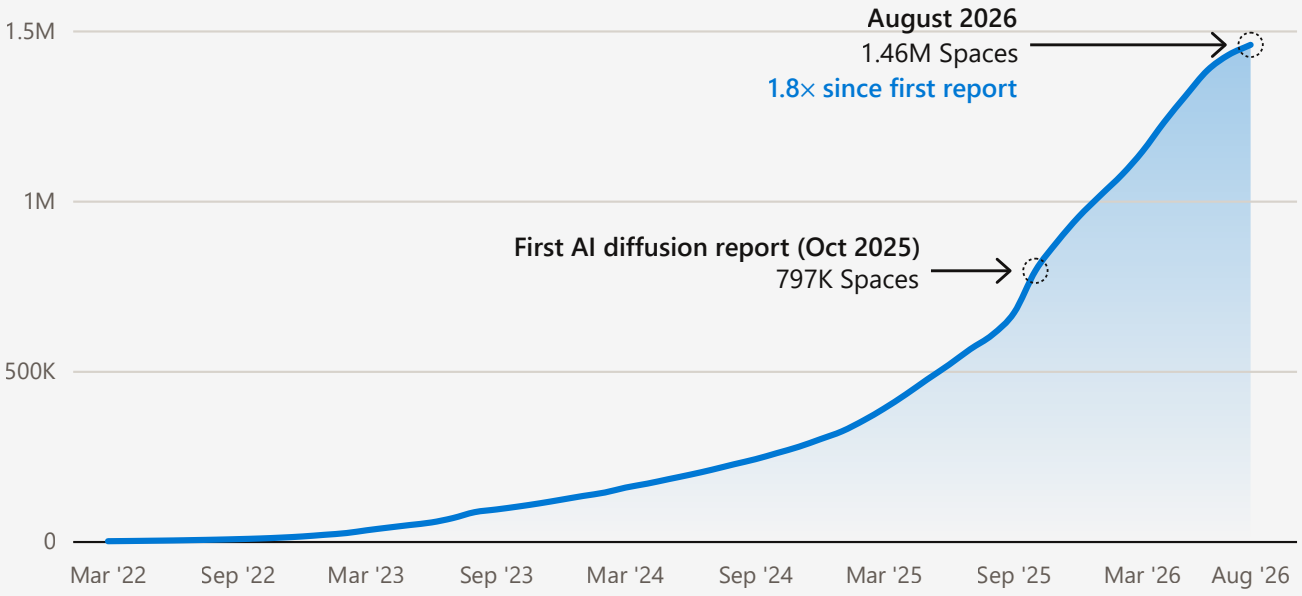
Open weights make this type of localization easier and allow models to be deployed through competing API providers—or, where infrastructure permits, locally—while giving users greater control over how they are customized. If capable open-weight models continue to trail the frontier by months rather than years while remaining available at a fraction of the inference price, they could become an important pathway for extending increasingly capable AI to populations and economies that have so far benefited least from AI's diffusion.

Growth in AI Tools

Alongside the growth of open-weight models, developer activity around AI applications has expanded rapidly. Hugging Face Spaces, which hosts runnable AI applications and demos ranging from personal experiments and research projects to more production-oriented applications, grew from approximately 124,000 Spaces at the end of 2023 to 1.46 million by September 2026. Since our first AI Diffusion report in November 2025, the total has increased by approximately 83%. Although Spaces should not be interpreted as distinct

Developer activity around AI applications has expanded rapidly

Cumulative surviving projects hosted on Hugging Face Spaces



Sources: Microsoft AI Economy Institute analysis of Hugging Face.

commercial AI products, their growth provides evidence of broadening experimentation and development at the application layer.

This growth in AI applications also has implications for how AI diffusion is measured. Many tools introduced since our first report have already attracted significant usage but are not included in our current measure. To address this gap, we plan to expand the list of AI tools covered in our next report, providing a more comprehensive view of AI usage worldwide. Initial analysis suggests that broader coverage will raise measured diffusion across nearly all countries, with the largest effect in China, where several widely used domestic tools are currently not included. These increases will primarily reflect improved measurement of existing usage rather than a sudden change in underlying adoption.

Citation and data availability

This report is based on the following technical paper:

Misra, A., Wang, J., McCullers, S., White, K., & Lavista Ferres, J. (2025). Measuring AI Diffusion: A Population-Normalized Metric for Tracking Global AI Usage. arXiv. **https://doi.org/10.48550/arXiv.2511.02781**
Also available at: **https://aka.ms/AI_Diffusion_Technical_Report**

The underlying AI Diffusion data used in this report is publicly available:

AI Diffusion dataset (Q2 2026 update):
https://github.com/microsoft/ai-diffusion-report/blob/main/data/AI_Diffusion_Q22026_Update.csv

When using this data, please cite the technical paper above and include a link to the dataset to ensure reproducibility and version transparency.

References

[1] A. Misra, J. Wang, S. McCullers, K. White, and J. L. Ferres, "Measuring AI Diffusion: A Population-Normalized Metric for Tracking Global AI Usage," Nov. 4, 2025, arXiv:2511.02781. doi: 10.48550/arXiv.2511.02781. **https://arxiv.org/abs/2511.02781**

[2] B. Costa-Gomes, S. Chen, C. Hsueh, D. Morgan, P. Schoenegger, Y. Shah, S. Way, Y. Zhu, T. Adeline, M. Bhaskar, M. Suleyman, S. Spielman, It's About Time: The Copilot Usage Report 2025: The Temporal and Modal Dynamics of Copilot Usage, Dec. 2025. Microsoft AI. **https://microsoft.ai/wp-content/uploads/2025/12/What_people_do_with_Copilot-8.pdf**

[3] M. I. G. Daepp and I. Slaughter, "How Early Adopters Used Generative AI Worldwide: Variation by Country Income and Language," May 2026, arXiv:2605.30685. doi: 10.48550/arXiv.2605.30685. **https://arxiv.org/pdf/2605.30685**

[4] Z. Iscenko, S. Strand, Y. Chen, G. Aimard, M. Codreanu, V. Sampathkumar, et al., "Google's AI & Economy ATLAS v1.0: Mapping Gemini Usage in the Economy," Jul. 2026. **https://ai.google/static/documents/GoogleATLASv1.pdf**

[5] R. Appel, M. Massenkoff, P. McCrory, M. McCain, R. Heller, T. Neylon, and A. Tamkin, "Anthropic Economic Index report: Economic primitives," Anthropic, Jan. 15, 2026. **https://www.anthropic.com/research/anthropic-economic-index-january-2026-report**

[6] Microsoft, "Open Weights and American AI Leadership," 2026. **https://www.microsoft.com/en-us/corporate-responsibility/topics/open-weight/**

[7] OpenRouter, "OpenRouter," [Online]. Available: **https://openrouter.ai/**

[8] Artificial Analysis, "Artificial Analysis AI Trends," 2026. [Online]. Available: **https://artificialanalysis.ai/trends**

[9] Kimi Team et al., "Kimi K3: Open Frontier Intelligence," Jul. 2026, arXiv:2607.24653. **https://arxiv.org/abs/2607.24653**

[10] Arena.AI, "Code Arena | WebDev Leaderboard," 2026. [Online]. Available: **https://arena.ai/leaderboard/code/webdev**

[11] M. Aubakirova, A. Atallah, C. Clark, J. Summerville, and A. Midha, "State of AI: An Empirical 100 Trillion Token Study with OpenRouter," 2026, arXiv:2601.10088. **https://arxiv.org/abs/2601.10088**

[12] F. Nagle and D. Yue, "The Latent Role of Open Models in the AI Economy," 2025. **https://papers.ssrn.com/sol3/Delivery.cfm?abstractid=5767103**

[13] A. Soni, "American AI Model Makers Smell an Opportunity," Reuters, Aug. 12, 2026. **https://www.reuters.com/technology/artificial-intelligence/american-ai-model-makers-smell-an-opportunity-2026-08-12/**

[14] C. Qi, V. Rengarajan, J. Pereyra, N. Grupen, and G. Pereyra, "Harvey Tenet Research Preview," Harvey, Aug. 20, 2026. **https://www.harvey.ai/blog/post-training-update-harvey-tenet**

[15] S. W. Zamir, W. Hamidouche, B. B. Amor, L. Marotti, I. Becker-Reshef, and J. L. Ferres, "BYOL: Bring Your Own Language Into LLMs," 2026, arXiv:2601.10804. **https://arxiv.org/abs/2601.10804**

AI Diffusion Data Source

Economy	Q1 2026	Q2 2026	Q2 Change
United Arab Emirates	70.1%	73.3%	3.2
Singapore	63.4%	64.3%	0.9
Ireland	48.4%	49.9%	1.5
France	47.8%	49.6%	1.8
Norway	48.6%	49.4%	0.8
Spain	44.2%	45.1%	0.9
New Zealand	43.0%	44.1%	1.1
United Kingdom	42.2%	43.7%	1.5
Netherlands	42.1%	43.5%	1.4
Qatar	41.8%	43.5%	1.7
Australia	39.5%	40.7%	1.2
South Korea	37.1%	40.6%	3.5
Belgium	39.0%	40.3%	1.3
Switzerland	37.8%	39.1%	1.3
Israel	38.1%	38.7%	0.6
Canada	37.3%	38.5%	1.2
Sweden	36.1%	37.3%	1.2
Austria	34.1%	35.4%	1.3
Taiwan	31.8%	33.8%	2.0
Hungary	32.2%	33.3%	1.1
United States	31.3%	33.0%	1.7
Denmark	31.2%	32.4%	1.2
Germany	31.1%	32.3%	1.2
Poland	31.0%	32.3%	1.3
Saudi Arabia	29.4%	31.4%	2.0
Italy	30.2%	31.4%	1.2
Czech Republic	30.1%	31.2%	1.1
Jordan	29.7%	31.2%	1.5
Bulgaria	29.7%	30.7%	1.0
Finland	29.5%	30.6%	1.1
Slovenia	29.0%	30.1%	1.1
Costa Rica	28.5%	29.7%	1.2
Vietnam	26.5%	27.9%	1.4
Lebanon	27.3%	27.7%	0.4
Oman	26.5%	27.6%	1.1
Portugal	26.4%	27.4%	1.0
Croatia	26.1%	27.2%	1.1
Slovak Republic	26.1%	27.1%	1.0
Colombia	24.5%	25.9%	1.4
Dominican Republic	24.8%	25.8%	1.0

Economy	Q1 2026	Q2 2026	Q2 Change
Uruguay	24.6%	25.7%	1.1
Serbia	24.1%	25.3%	1.2
Lithuania	24.3%	25.1%	0.8
Jamaica	24.0%	25.0%	1.0
Japan	22.5%	24.7%	2.2
Panama	23.3%	24.2%	0.9
Chile	22.7%	24.0%	1.3
South Africa	23.1%	23.9%	0.8
Argentina	21.9%	23.4%	1.5
Bosnia And Herzegovina	22.1%	23.2%	1.1
Malaysia	21.8%	22.6%	0.8
Kuwait	21.1%	22.6%	1.5
Georgia	20.5%	21.6%	1.1
Greece	20.8%	21.5%	0.7
Mexico	20.1%	21.4%	1.3
Philippines	20.1%	20.7%	0.6
Ecuador	19.5%	20.3%	0.8
Brazil	19.1%	20.2%	1.1
El Salvador	18.3%	20.1%	1.8
Albania	18.5%	19.2%	0.7
Moldova	18.5%	19.1%	0.6
Turkey	17.4%	18.9%	1.5
Azerbaijan	17.7%	18.7%	1.0
India	17.6%	18.5%	0.9
Romania	17.5%	18.0%	0.5
Mongolia	16.7%	17.9%	1.2
China	16.4%	17.5%	1.1
Peru	16.4%	17.3%	0.9
Guatemala	16.4%	17.3%	0.9
Kazakhstan	15.9%	16.9%	1.0
Gabon	15.0%	15.7%	0.7
Libya	15.0%	15.5%	0.5
Egypt	14.8%	15.5%	0.7
Namibia	15.1%	15.4%	0.3
Botswana	14.8%	15.0%	0.2
Indonesia	14.1%	14.8%	0.7
Nepal	14.2%	14.6%	0.4
Honduras	14.0%	14.5%	0.5
Senegal	13.9%	14.4%	0.5
Tunisia	13.5%	13.9%	0.4

Economy	Q1 2026	Q2 2026	Q2 Change
Algeria	13.3%	13.8%	0.5
Cote D'Ivoire	13.1%	13.7%	0.6
Thailand	12.4%	13.4%	1.0
Bolivia	12.7%	13.3%	0.6
Zambia	13.1%	13.3%	0.2
Iraq	12.5%	13.0%	0.5
Iran	12.6%	12.8%	0.2
Paraguay	12.2%	12.6%	0.4
Nicaragua	11.8%	12.5%	0.7
Morocco	11.7%	12.0%	0.3
Pakistan	11.4%	11.9%	0.5
Gambia	11.5%	11.5%	0.0
Mozambique	10.9%	11.4%	0.5
Angola	10.9%	11.4%	0.5
Madagascar	10.9%	11.4%	0.5
Malawi	10.9%	11.4%	0.5
Guyana	10.3%	11.3%	1.0
Suriname	10.3%	11.3%	1.0
Venezuela	10.3%	11.3%	1.0
French Guiana	10.3%	11.3%	1.0
Myanmar	10.0%	10.4%	0.4
Togo	10.1%	10.4%	0.3
Guinea	10.1%	10.4%	0.3
Niger	10.1%	10.4%	0.3
Liberia	10.1%	10.4%	0.3
Nigeria	10.1%	10.4%	0.3
Benin	10.1%	10.4%	0.3
Mali	10.1%	10.4%	0.3
Ghana	10.1%	10.4%	0.3
Guinea-Bissau	10.1%	10.4%	0.3
Sierra Leone	10.1%	10.4%	0.3
Mauritania	10.1%	10.4%	0.3
Burkina Faso	10.1%	10.4%	0.3
Belarus	9.6%	10.3%	0.7
Kyrgyzstan	9.5%	10.0%	0.5
Russia	9.5%	9.9%	0.4
Lesotho	9.8%	9.9%	0.1
Ukraine	9.5%	9.5%	0.0
Chad	8.7%	9.1%	0.4
Central African Republic	8.7%	9.1%	0.4

Economy	Q1 2026	Q2 2026	Q2 Change
Cameroon	8.7%	9.1%	0.4
Democratic Republic Of The Congo	8.7%	9.1%	0.4
Congo	8.7%	9.1%	0.4
Haiti	8.5%	8.9%	0.4
Kenya	8.7%	8.9%	0.2
Zimbabwe	8.5%	8.9%	0.4
Laos	7.8%	8.4%	0.6
Bangladesh	7.8%	8.0%	0.2
Uganda	7.6%	8.0%	0.4
Ethiopia	7.6%	8.0%	0.4
Somalia	7.6%	8.0%	0.4
Eritrea	7.6%	8.0%	0.4
South Sudan	7.6%	8.0%	0.4
Sudan	7.6%	8.0%	0.4
Burundi	7.6%	8.0%	0.4
Tanzania	7.6%	8.0%	0.4
Armenia	7.4%	7.7%	0.3
Papua New Guinea	7.7%	7.6%	-0.1
Sri Lanka	7.3%	7.6%	0.3
Rwanda	7.2%	7.6%	0.4
Uzbekistan	7.2%	7.6%	0.4
Syria	7.5%	7.4%	-0.1
Cuba	6.7%	6.8%	0.1
Tajikistan	6.1%	6.3%	0.2
Afghanistan	6.1%	6.3%	0.2
Turkmenistan	6.1%	6.3%	0.2
Cambodia	5.8%	6.1%	0.3

