



HOOVER
INSTITUTION

Stanford
University



Stanford University
Human-Centered
Artificial Intelligence

White Paper
June 2026

Update: DeepSeek AI and the Great Talent Competition

Amy Zegart and
Emerson Johnston



更多免费行业报告

并购家

www.ipopo.cn

Authors

Emerson Johnston is a full-time research assistant at the Hoover Institution's Technology Policy Accelerator. She is a graduate of Stanford's Ford Dorsey Master's in International Policy and received her B.S. and B.A. from Northeastern University.

Amy Zegart is the Morris Arnold and Nona Jean Cox Senior Fellow at the Hoover Institution, director of Hoover's Technology Policy Accelerator, and a senior fellow and associate director at Stanford University's Institute for Human-Centered AI.

How to cite this paper

Emerson Johnston and Amy Zegart, "Update: DeepSeek AI and the Great Talent Competition" Technology Policy Accelerator, Hoover Institution and the Stanford Institute for Human-Centered Artificial Intelligence, June 2026.
<https://doi.org/10.64576/0626>



When DeepSeek AI released its R1 paper in January 2025 that revealed a model rivaling OpenAI’s o1 for far less compute, media narratives were dominated by two misconceptions. First, that the Chinese company’s breakthroughs were the result of Chinese researchers who could imitate but not innovate; and second, that the research team was young, green, and generally not highly cited in the computer science field. [Our 2025 report](#), “A Deep Peek into DeepSeek AI’s Talent and Implications for US Innovation,” extensively examined the 223 researchers behind the company’s five foundational papers and found a very different picture.

We found that DeepSeek *did* innovate in significant ways, that its research team was far more experienced than most assumed, and that more than half of the team had never left China for school or work. Together, the findings pointed to a domestic AI talent pipeline in China that is larger and more capable than the prevailing narrative allowed. They also suggested that any U.S. strategy built on the assumption that China lacks the talent or know-how to surpass the U.S. on its own was already out of step with the evidence.

We argued then that the United States needs to compete much more aggressively to attract, welcome, and retain the world’s best and brightest while urgently growing domestic capabilities by improving K–12 and higher education in STEM fields at home.

This report asks how much has changed in the year since. We extended our analysis from five to seven papers—adding [DeepSeek V3.2](#) (December 2025) and [V4](#) (April 2026) to incorporate the company’s latest research. By expanding our original dataset, we find the author pool has grown from 223 to 356 and we were able to build comprehensive profiles tracking where they studied, lived, and worked for 282 of them.

“... a third of the most integral contributors to a Chinese model performing comparably to OpenAI’s o1 model never trained, published, or worked in an international research environment.”

China’s domestic AI talent pipeline is still going strong

- The share of DeepSeek AI researchers affiliated exclusively with Chinese institutions throughout their recorded careers stands at 53.5 percent.
- Ten of the 31 Key Team members—researchers listed on all seven DeepSeek AI papers published in the last two years—are also affiliated exclusively with Chinese institutions, meaning a third of the most integral contributors to a Chinese model performing comparably to OpenAI’s o1 model never trained, published, or worked in an international research environment.

New and better data reveals longer U.S. training periods, especially for top talent

A year ago, 63.3 percent of the U.S.-affiliated cohort had recorded only a single year of U.S. experience. That figure is now 35 percent, with nearly half in the 2- to 4- year range. This shift reflects both new DeepSeek talent and improved data about the original cohort we examined last year. Notably, the 141 net new

contributors to DeepSeek Papers 6 and 7 in 2026 are a more internationally mobile cohort than the original 223 that contributed to papers in 2025. In addition, we found that OpenAlex's affiliation coverage has improved across the board, surfacing prior U.S. positions for researchers we tracked last year that were not yet in the database when we ran our original queries.

The picture emerging from both the new researchers and our improved data is not of brief, exchange-style visits to the U.S., but of meaningful embedding in U.S. research environments followed by a return rate of 70.2 percent among internationally mobile DeepSeek contributors.

- Longer U.S. exposure did not produce higher retention. The 13 long-stay researchers in this dataset spent a combined 119+ recorded years at U.S. institutions, and most of them are now in China.
- Although nine of the 13 longest-stay researchers eventually returned to China, they did not

simply land in the U.S. stay for many years and then return to China. Instead, they made many international transitions over many years.

- The 80 DeepSeek AI researchers with any U.S. institutional affiliation are, on average, the most academically accomplished in the pool, and a meaningful share of them spent five or more years embedded in U.S. institutions before returning to China.

The implication is, at best, uncomfortable. Knowledge that has already moved with these researchers cannot be recovered by tightening visa policy or export controls going forward. Restricting future access addresses one talent problem but leaves the second untouched. The domestic Chinese pipeline is now producing frontier-model contributors without much U.S. or international exposure. At the same time, prior U.S. exposure continues to pay dividends to Chinese AI firms at the frontier. The 2026 data shows clearly that the United States is facing both problems simultaneously.

“... the domestic Chinese pipeline is now producing frontier-model contributors without much U.S. or international exposure. At the same time, prior U.S. exposure continues to pay dividends to Chinese AI firms at the frontier.”

Methodology

Full methodological details are in our original report. Briefly, we identified the research and engineering contributors across DeepSeek’s seven papers, drawn from arXiv (2024–2025) and Hugging Face (April 2026). Of the 356 named authors, 282 had usable profiles in OpenAlex, which we queried in April 2026 for publication records, citation metrics, and institutional affiliation histories dating back to 1989, the earliest year for which OpenAlex carries usable affiliation records for any researcher in this cohort. We then used Python to map each researcher’s complete institutional history, allowing us to track cross-border mobility.

Unlike our previous analysis, which counted all named contributors across each paper, this analysis restricts its scope to those identified on research and engineering

teams. We believe this methodological change focuses more centrally on research contributions and will enable better longitudinal analysis in the future.

Where papers distinguished contributor roles, ‘data annotation’, ‘business’, and ‘compliance’ contributors were excluded because they fall outside the population of researchers most relevant to understanding talent competition dynamics in AI. Where papers did not distinguish roles, all listed contributors were included. Full per-paper counts, including excluded contributor totals, can be found in Appendix B, Table 1.



A Bigger Team, Same Structure


A total of 356 people were identified as research and engineering contributors across DeepSeek’s seven papers, up from the 223 identified in our last analysis. The 57 percent expansion in a single year is large in absolute terms and large relative to the team’s prior trajectory: the first five papers covered roughly two years and added researchers steadily. The last two papers added roughly 141 net new contributors in 14 months.¹ Departures over that same window total 33 contributors (eight from

Paper 6, 10 from Paper 7, with the rest concentrated in Papers 4 and 5), meaning DeepSeek brought in over four researchers for every one that it lost.

We find, again, that 31 researchers (just under nine percent) contributed to all seven papers, which we refer to as the “Key Team”. Another 47 authors appeared on six papers, 17 on five, 41 on four, 5 on three, 79 on two, and 136 on just one (see Figure 1).

Figure 1: Authorship Distribution Across DeepSeek Papers

Authorship Distribution of 356 Contributors Across the Seven Most Recent DeepSeek AI papers



¹ The 223-author total from last year’s report included data annotation, business, and compliance contributors. Restricting to research and engineering, last year’s comparable count was 215. The growth figures here are calculated against 215, not 223.


The persistent shape of this distribution in 2025 and 2026 suggests something specific about how DeepSeek organizes research labor. Occasional contributors cluster at the edges: 215 researchers (roughly two-thirds of the pool) appear on only one or two papers. The middle drops out almost entirely. Researchers who move past two papers tend to keep going, while those who don't stop early. That pattern is consistent with a two-track workforce: track one being a large rotating cast of contributors pulled in for specific model releases, and track two being a committed core that ships everything. The 31-person Key Team and the 47 six-paper contributors are on one track, and the 215 one- and two-paper contributors are on the other. The three to five paper tier is thin because it is the outlier.

Figure 2 cuts the same population three ways: the Key Team (all included papers), the middle tier (contributed to more than one but less than all), and one-paper contributors. Between 2025 and 2026, the Key Team held flat at 31 in absolute terms but shrank as a share of the pool from 13.9 percent to 8.7 percent. The middle tier grew from 169 contributors to 189. The one-paper group grew from 23 to 136, representing a near sixfold increase, and a jump from 10.3 percent of the pool to 38.2 percent.

Of the 141 new researchers added between the 2025 and 2026 datasets, 71 entered as one-paper contributors, and the other 70 entered the middle-tier directly by appearing on both Papers 6 and 7. This suggests that

Figure 2: Change in Authorship Distribution of DeepSeek AI Researchers, 2025-2026

Share of authors contributing to one, some, or all papers in 2025 and in 2026, with one-hit wonders rising from 10% to 38%



DeepSeek's two-track structure has strengthened over the past year, with the company's growth coming entirely from the rotating-cast side of its workforce, while the committed core stays the same size.

One additional pattern worth noting is that 43 researchers appear on Papers 1, 2, 4, 5, 6, and 7, each

only absent on Paper 3 (VCoder 2). Combined with the 31 who appear on all seven, this suggests Paper 3 drew on a narrower specialist team that did not overlap fully with the main model development track, consistent with VCoder 2's more focused scope relative to the flagship V-series releases.

CHINA'S DOMESTIC AI TALENT PIPELINE — ONE YEAR LATER

Our original analysis covered 201 DeepSeek authors with known affiliation data. Our new dataset extends that analysis to cover 271 authors who had institutional affiliation records—out of a grand total of 356.²

The core finding from last year holds. China has a robust, and growing, domestic AI talent pipeline. Just over half—53.5 percent, or 145 authors—have been affiliated exclusively with Chinese institutions throughout their recorded careers, never holding a position or academic post outside China, down only slightly from 55.2 percent in our original analysis.³ We also found again that nearly all of the DeepSeek researchers (97 percent in the expanded dataset) have held at least one confirmed affiliation with a Chinese institution at some point in their careers.

The share of authors with any *U.S.* affiliation has also grown. 29.5 percent of authors in the current sample (n=80) have held a U.S. institutional tie at some point,

up from 24.4 percent in our original analysis.

The current snapshot tells much the same story as a year ago (Figure 3). Of the 271 authors with known affiliations, 229 (84.5 percent) are currently affiliated with Chinese institutions and 18 (6.64 percent) are currently at U.S.-based institutions. The rest are distributed across the United Kingdom, Taiwan, Japan, Singapore, and a handful of other countries, none exceeding 1.5 percent.

The composition of the 18 currently in the U.S. is less concentrated than last year's report suggested. Carnegie Mellon, UC Berkeley, the University of Pennsylvania, UT Austin, Purdue, and Stony Brook are represented, as is UT Southwestern Medical Center. The rest include Abraham Baldwin Agricultural College in Georgia, King University in Tennessee, Framingham State University in Massachusetts, and two small private companies, Intelligent Automation and Scientific Methods, indicating

² As we noted in last year's analysis, not every researcher with an OpenAlex profile has a fully cataloged career history. The gap between 282 and 271 reflects authors whose profiles exist but whose institutional affiliation records are missing.

³ This figure includes those who have been affiliated with Hong Kong SAR and Taiwan.

that the U.S.-affiliated tail of DeepSeek's author network includes several peripheral contributors rather than a concentrated presence at elite AI research institutions.


Seven authors in the expanded dataset have no recorded affiliation with any Chinese institution at all. Three are exclusively U.S.-affiliated: Panpan Huang, whose only recorded affiliation is Purdue University West Lafayette in 2017; Jiewen Hu, who spent six years at Virginia Tech with a year at Apple in between, and is now at Carnegie Mellon; and Ruiling Xu, who trained at the University of Pennsylvania before moving to California University of Pennsylvania, where they remain. Two others split between the U.S. and one other country: Yifan Ding, who has spent most of their recorded career

at the University of Notre Dame with a brief stop at Amazon Germany; and Erhang Li, whose recorded affiliations span Intel's U.S. and U.K. offices.

The remaining two trained entirely outside both China and the U.S.: Yuduan Wang, who studied at the National University of Singapore and is currently affiliated with the National University Cancer Institute there; and Zizheng Pan, whose career has been at the University of Adelaide and Monash University in Australia. Y.Q. Wang, who appeared in this category in our original analysis, now shows a Chinese institutional affiliation at the 'Institute of High Energy Physics' in the updated data.

Figure 3: Geographic Distribution of Current Institution Affiliations

Distribution of 'Current Affiliations' of the 271 DeepSeek AI Researchers with Affiliation Data





Box 1: Where The New Researchers Have Been

Of the 141 brand-new researchers in this year's dataset, 69 had resolvable OpenAlex profiles, 68 had affiliation data, and 18 of those 68 (26 percent) had at least one U.S. affiliation past or present. That rate is essentially the same as the full author pool's 29 percent, suggesting DeepSeek's recent hiring has not skewed more or less toward U.S.-trained researchers than its earlier hires. The 18 researchers spread across 37 U.S. institutions (see Table 1).

Only Carnegie Mellon and UT Southwestern Medical Center host more than one new contributor. Every other institution in Table 1 represents a single researcher's career stop. The variation is in how long those single researchers stayed. One new contributor recorded eight years at Baylor College of Medicine, another six at Virginia Tech, another five at Beth Israel Deaconess Hospital. These high affiliation-year counts reflect long single-researcher embeddings rather than institutional pipelines hosting multiple DeepSeek-bound researchers.

Table 1 - U.S. Institutions where the new cohort of DeepSeek AI researchers have trained.

No. Authors	Institution	City	State
2	Carnegie Mellon University*	Pittsburgh	Pennsylvania
2	The University of Texas Southwestern Medical Center*	Dallas	Texas
1	Alcoa	Pittsburgh	Pennsylvania
1	Apple	Cupertino	California
1	Arconic	Pittsburgh	Pennsylvania
1	Auburn University	Auburn	Alabama
1	Baylor College of Medicine	Houston	Texas
1	Beth Israel Deaconess Hospital	Boston	Massachusetts
1	California University of Pennsylvania*	California	Pennsylvania
1	Child Mind Institute	New York	New York
1	Children's Cancer Center	Tampa	Florida
1	Dana-Farber Cancer Institute	Boston	Massachusetts
1	Entrepreneurial Ecosystems	Lincoln	Nebraska
1	Framingham State University*	Framingham	Massachusetts
1	Fred Hutch Cancer Center	Seattle	Washington
1	Howard Hughes Medical Institute	Chevy Chase	Maryland
1	International Craniofacial Institute	Dallas	Texas
1	Lehigh University	Bethlehem	Pennsylvania
1	Mars	McLean	Virginia
1	Massachusetts Institute of Technology	Cambridge	Massachusetts
1	Penn State Milton S. Hershey Medical Center	Hershey	Pennsylvania
1	Renal Research Institute	New York	New York
1	Rockefeller University	New York	New York
1	Seattle University	Seattle	Washington

Box 1: Where The New Researchers Have Been (cont'd)

Table 1 - U.S. Institutions where the new cohort of DeepSeek AI researchers have trained. (cont'd)


No. Authors	Institution	City	State
1	Stanford University	Stanford	California
1	Target	Minneapolis	Minnesota
1	The University of Texas Health Science Center at Houston	Houston	Texas
1	The University of Texas MD Anderson Cancer Center	Houston	Texas
1	University of California, Berkeley*	Berkeley	California
1	University of California, Davis	Davis	California
1	University of California, San Diego	La Jolla	California
1	University of Delaware	Newark	Delaware
1	University of Notre Dame*	Notre Dame	Indiana
1	University of Oklahoma	Norman	Oklahoma
1	University of Pennsylvania	Philadelphia	Pennsylvania
1	University of Southern California	Los Angeles	California
1	Virginia Tech	Blacksburg	Virginia
1	Yale University	New Haven	Connecticut

HOW THE US-CHINA PIPELINE ACTUALLY MOVES

Of the 80 DeepSeek researchers with any U.S. institutional affiliation, the largest revision from our original analysis appears in the duration data (see Figure 4). Last year's dataset recorded 63.3 percent of the 49 US-affiliated researchers as having only a single year of U.S. experience. In the expanded dataset, the equivalent share is 35 percent (n=28). Nearly half—48.8 percent (n=39)—now fall in the 2–4-year range, and 16.3 percent (n=13) spent five or more years at U.S. institutions.

Figure 4: U.S. Experience Duration

U.S. Experience Duration for the 80 DeepSeek Researchers with U.S. Affiliations



This new data revises last year's read of the U.S. as a mere stepping stone for elite Chinese AI researchers. A meaningful subset of DeepSeek's most cited contributors had longer, formative stays in the United States. For them, the U.S. is still a step, but a much bigger one than we reported last year.

The 13 long-stay researchers include some of the most academically accomplished in the entire dataset. Minghua Zhang, currently at Stony Brook University in the United States, has 28 recorded years of U.S. affiliations and nearly 30,000 career citations. Chen Dong, whose 51,823 citations and h-index of 112 account for a disproportionate share of the U.S.-affiliated cohort's citation totals, spent 12 years across three U.S. stays and remains U.S.-based at Framingham State University. On the other hand, Songyang Zhou's 18-year U.S. run through Beth Israel Deaconess, Rockefeller, MIT, Baylor College of Medicine, and UC San Diego ended in their return to China to Hainan Medical University.


Notably, what sets these 13 careers apart is how rarely they look like a single embedding followed by a clean return. Nine of the 13 made more than one international transition. Lei Xu's record runs through U.S. institutions, then the U.K., then Brunei, before settling back in China. Minghui Tang's includes Japan and the U.K. between U.S. and Chinese stays, with a return through Tsinghua. B. Zhang began in Sweden before alternating between U.S. and Chinese affiliations. For this group, U.S. institutions sit inside sustained transnational careers rather than serving as one-shot credentialing stops. The directional weight is important, but not absolute: eight of the 13 ended in China and only five in the U.S., and length of U.S. exposure does not predict the outcome. Minghua Zhang's 28 years and Tongzheng Ren's five years both ended in the U.S.; Songyang Zhou's 18 years and Ruoyu Zhang's six years both ended in China.

The broader trajectory data tells a similar story to last year (see Figure 5). Among the 80 US-affiliated

researchers, the most common mobility pattern is China → U.S. → China: 38.8 percent (n=31) started and ended in China, with U.S. experience in between. The second largest group, 23.8 percent (n=19), started in the United States and ended up in China. Only 10 researchers (12.5 percent) followed the China → U.S. → stayed in the U.S. path, down from 14.2 percent in our original analysis, and only five researchers (6.2 percent) started and ended in the United States with international travel in between.

Figure 5: U.S. Retention Rate

Full Retention Data for the 80 DeepSeek Researchers with U.S. Affiliations



The overall retention picture across all 282 authors reinforces this finding (see Figure 6). The single largest mobility category (excluding researchers who have only ever been affiliated with Chinese institutions) is contributors who started and ended in China with some international travel in between: 53 authors or 19.6 percent of the pool. Taken together, the trajectories that end up outside China account for roughly 15.13 percent of the full author pool. The rest, whether they stayed in China throughout or returned

after time abroad, ended up there, contributing to an unprecedented retention rate of 70.3 percent of international mobile contributors.

The geographic distribution of U.S. institutions in the expanded dataset is wider than last year's (see Figure 7). The 80 researchers connect to 113 U.S. institutions, with clusters in California, Texas, New York, and the Boston-to-DC corridor. Interestingly, there is a much more significant representation in the Midwest and South than last time. As in our original analysis, no single institution dominates, but Microsoft leads with seven total affiliated researchers across the network, consistent with the MSRA pipeline; Columbia, Carnegie Mellon, and Stanford each have four or five. For a full list of U.S. institutions and the number of authors affiliated with each, see Appendix B, Table 2.

Figure 6: Overall Retention Rate

Full Retention Data for the 271 DeepSeek Researchers with Affiliation Data







Figure 7: Geographic Distribution of US Institutions Affiliated with DeepSeek AI Researchers



These patterns complicate a policy debate that has tended to treat talent flows as a pipeline problem, one that better visa policy or stronger export controls might redirect. A year ago, we described American institutions as stepping stones, or places that equipped elite researchers with skills and credentials that were ultimately reinvested into China's AI ecosystem. The fuller 2026 data reframes that finding.

Some U.S. stays were short and instrumental, consistent with the stepping-stone description. But a significant share were longer than we previously documented, sometimes a decade or more, and the return rate did not decline with longer exposure. That is closer to brain gain for China than to brain circulation.

The 13 long-stay researchers in this dataset spent a combined 119+ recorded years at U.S. institutions, and most of them are now in China.

“The 13 long-stay researchers in this dataset spent a combined 119+ recorded years at U.S. institutions, and most of them are now in China.”

Further, of the 25 researchers across the full dataset who began their recorded careers in the United States, 19 (76 percent) ended up in China, and one ended up elsewhere abroad. The nine long-stay researchers who spent between five and eighteen years embedded in American institutions but ultimately returned, did so at roughly the same rate as the rest, meaning longer exposure did not produce higher retention. Songyang Zhou's 18 years in the United States and Guanting Chen's three years both ended in China.

This new data proves not only that China has developed its own researchers, but it has also been successfully recovering researchers who had trained in the United States, sometimes for over a decade, and who returned carrying methods, networks, and credentials built inside the American research system.

A DEEPER DIVE INTO INSTITUTIONAL AFFILIATIONS

The 282 authors in the expanded dataset connect to 812 unique institutions across their recorded careers, up from 499 last year. Chinese institutions account for 570 of those, or 70.2 percent of the total — down slightly from 74 percent in our original analysis. U.S. institutions account for 113 (13.9 percent). The remaining 129 are spread across the UK, Australia, Taiwan, Canada, Japan, and elsewhere.

The Chinese Academy of Sciences again has the broadest reach in the network. CAS directly hosts only four of the 282 authors, but its 170 affiliated institutions — research institutes, laboratories, and specialized centers that OpenAlex classifies as organizational children of CAS — connect to 104 researchers in total, about 37 percent of the full author pool. That figure has doubled from 53 a year ago, the largest absolute year-over-year gain of any institution in the dataset.

Outside the CAS network, the comparison with last year sharpens the institutional picture (see Table 2). Tsinghua University leads with 46 affiliated researchers, up from 16, representing a near 200 percent increase. Similarly, Peking University moved from 21 to 29, Zhejiang from 8 to 25, and the University of Chinese Academy of Sciences from 7 to 21.

Several institutions that were marginal in 2025 are now in the top tier: Southeast University grew from 1 affiliated researcher to 15, Beihang from 3 to 14, Lanzhou from 0 to 10, and Northeastern University (China) and Guangzhou Medical University each from 1 to 10. Two institutions moved the other way. Sun Yat-sen dropped from 10 to 4 and Tongji from seven to five, though both drops may be a result of how OpenAlex re-attributed how those affiliations rolled up to parent institutions elsewhere in the dataset.



Table 2 - Top Institutions by Researcher Count, 2025 vs 2026

Institution	2025	2026	Change
Chinese Academy of Sciences (network)	53	104	+51
Tsinghua University	16	46	+30
Zhejiang University	8	25	+17
University of Chinese Academy of Sciences	7	21	+14
Southeast University	1	15	+14
Beihang University	3	14	+11
Lanzhou University	0	10	+10
Northeastern University (China)	1	10	+9
Guangzhou Medical University	1	10	+9
Peking University	21	29	+8
Nanjing University	10	16	+6
Huazhong University of Science and Technology	5	11	+6
Xi'an Jiaotong University	4	10	+6
Shanghai Jiao Tong University	8	13	+5
Tongji University	7	5	-2
Sun Yat-sen University	10	4	-6

THE “JUNIOR TALENT” NARRATIVE, ONE YEAR LATER

A year ago, the prevailing narrative was that DeepSeek’s R1 model succeeded on the back of fresh PhD graduates with little applied experience. Liang Wenfeng, DeepSeek’s CEO, had said as much himself, telling interviewers that core technical positions were mostly filled by researchers with one to two years of work experience, and that DeepSeek recruited for “passion and curiosity” over credentials.⁴

Our original analysis pushed back on this account. We found that DeepSeek’s Key Team averaged 1554 citations per author with a median h-index of 10. That is not a picture of inexperience. A year later, the numbers have moved further in the same direction.

Across the full author pool, average citations rose from roughly 1,000 to 1,763, and the median more than doubled, from 249 to 681 — meaning the whole distribution shifted upward, not just the tail. The i10-index, which counts an author’s papers with at least 10 citations, moved more modestly: the average grew from just over 19 to 22.55, and the median from five to seven. The h-index moved in a similar register, with the average rising from 10.8 to 12.46 and the median from seven to eight.

A more striking change in metrics comes from the Key team. The 31 authors now average 2,470 citations, up from 1,554 a year ago and representing a 59 percent increase. Median citations rose from roughly 700 to 1,200, the average h-index from 13.5 to 17, the median h-index from 10 to 12, and the i10-index average reached 33.84. Much of this jump traces back to the DeepSeek papers themselves: R1, V3, and V2 have accumulated thousands of citations each since our last analysis, and the

Key Team are authors on all of them. The bibliometric shift is, in that sense, a measure of how quickly the field absorbed DeepSeek’s 2025 output.

The expanded dataset also allows us to isolate a group we could not cleanly analyze before: the 80 DeepSeek authors with any past or current U.S. institutional affiliation. They average 131 published works, 4,108 citations, and an h-index of 21.26—more than double the full pool’s citation average and well above even the Key Team’s. Median citations sit at 1,452, against 681 for the full pool. By every metric, this is the most academically accomplished cohort.

This strength becomes even more apparent when compared to a peer group from one of the world’s leading AI labs. The same U.S. lab analysis mentioned above shows OpenAI’s pool of 894 authors has an average citation count of 2,481.5, but a median of just 100.5, indicating a steep drop-off beyond a few highly cited individuals. Further, the group’s median h-index was only two, down from six in our previous reporting, and the median i10-index is the same, reflecting more limited consistency of impact for OpenAI. This pattern recurs for both Anthropic and Alphabet, whose median citation counts are only 15 percent and 20 percent of their averages, respectively. DeepSeek’s is 35 percent.

In our 2025 ‘Deep Peek’ paper, we described DeepSeek’s bibliometric profile as more comparatively balanced and less reliant on a small number of standout figures than its American counterparts. A year later, this characterization holds for the full author pool, the Key Team, and the group of authors with U.S. affiliations.

⁴ Tony Peng, “The Deep Roots of DeepSeek: How It All Began,” Recode China AI (blog), January 27, 2025, <https://www.recodechinaai.com/p/the-deep-roots-of-deepseek-how-it>; Jordan Schneider et al., “Deepseek: The Quiet Giant Leading China’s AI Race,” ChinaTalk (blog), November 27, 2024, <https://www.chinatalk.media/p/deepseek-ceo-interview-with-chinas>.

THE KEY TEAM REVISITED

In last year's report, we examined in depth DeepSeek's Key Team—the 31 researchers whose names appeared across all five foundational papers. This year's query of OpenAlex surfaced more complete data and revealed three errors in the original count. Runxin Xu and Yukun Li had appeared on all five papers but were not matched to their full-name profiles, because some papers recorded their names under initials (R. X. Xu and Y. K. Li, respectively). Yanhong Xu, by contrast, was listed last year but turns out not to have been on Paper 3, which means they should not have been on the all-papers count to begin with. Net of those three corrections, the 2025 Key Team should have totaled 32, not 31.

Extending the analysis to seven papers tightens the Key Team list further. Fuli Luo left DeepSeek after Paper 5 and is absent from Papers 6 and 7, dropping one from the all-papers count in 2025. Resultantly, the Key Team in 2026 holds at 31, but the names behind that number are not quite the same. For a full list of 2026 Key Team members across all seven papers, see Appendix C.

Of the 31 Key Team members, OpenAlex returns affiliation data for 28 of them. Eighteen have held at least one position outside China during their recorded careers, meaning that 10 Key Team researchers, or roughly 36 percent of them, never left China.

Eleven of the 18 internationally mobile Key Team members spent time at U.S. institutions, but most spent only one or two years there. Zhihong Shao's four years at the University of Michigan is the longest single U.S. stay; Qihao Zhu spent a year at Carnegie Mellon, Zhenda Xie two years at Columbia, and Runxin Xu a year at Ohio State before returning through

Imperial College London on their way back to China. Additionally, the timing for most of these researchers—mid-career rather than at the start—points to visiting researcher or postdoctoral appointments rather than degree-seeking, consistent with targeted knowledge acquisition rather than emigration.

The seven internationally mobile Key Team members without U.S. affiliations followed different routes. Daya Guo spent three years at Microsoft Research UK and a German university before returning to Sun Yat-sen. Damai Dai did a year at Microsoft Research India. Kai Dong has rotated between China, Canada, Italy, and Finland, and is currently based at the University of Oulu. Wenjun Gao has affiliations in Singapore, Denmark, and Hong Kong.

Taken together, this trajectory data suggests again that international experience for this cohort was brief, targeted, and almost universally followed by a return to China. In this context, the United States appears less as a destination than as one node among several, interchangeable, for some researchers, with Microsoft Research India or the Technical University of Denmark. The Key Team thus becomes a story of a domestic research base that selectively absorbed outside experience and then reconverged, rather than one of global recruitment.

TWO TALENT PROBLEMS, ONE POLICY DEBATE

A year ago, we framed the geopolitical implications of DeepSeek AI's talent pool around a single proposition: the United States can no longer assume a permanent talent advantage. The updated data suggests the U.S. faces two separate talent challenges, not one.

The first challenge is retaining China's top talent in the United States. The 80 researchers who passed through U.S. institutions are the most accomplished group in DeepSeek's author pool. They average 4,108 citations and a median h-index of 16.5. Most of them are now in China. Key Team members who went to the U.S. typically stayed one or two years. The highest-cited contributors outside the Key Team stayed much longer, sometimes more than a decade. Both groups returned to China at roughly the same rate. Length of exposure did not change the outcome.

The second challenge is the growing independence of China's AI talent supply chain. Ten of the 31 Key Team members have no recorded affiliation outside China. Across the full author pool, 145 researchers, or 53.5 percent, have never held a position outside China. These researchers contributed to a model that performed comparably to OpenAI o1 on reasoning benchmarks. They did not train in the United States, work at U.S. labs, or publish from U.S. institutions.

Visa policy changes can address the first challenge, but researchers who never came to the United States

are not reachable by rules about coming to the United States. A policy debate focused only on whether to enable Chinese talent to come to the U.S. and remain here does not address China's growing domestic talent supply chain.

Improving retention of Chinese talent in the United States is the more tractable of the two. Among DeepSeek's contributors with any U.S. affiliation, 70.3 percent of the internationally mobile ones ultimately ended up in China. U.S. policy currently treats that figure as a fixed feature of the landscape rather than the variable it is. Green-card backlogs, post-PhD visa uncertainty, and the absence of clear long-term pathways for STEM researchers to remain in the United States instead of returning to their home countries are all crucial parts of the talent retention problem, and all are addressable.

We should not assume Chinese researchers leave the U.S. because they want to. Many may return to China because they have to: the U.S. visa system makes it too difficult, costly, and uncertain to remain. A Chinese researcher who files for an employment-based green card today can expect to wait close to five years before a visa number even comes up, and that line is getting longer, not shorter.⁵ Meanwhile, a new \$100,000 fee on each H-1B petition and a reweighted lottery favoring high salaries push postdocs and recent PhDs to the back of the line.⁶

⁵ U.S. Department of State, "Visa Bulletin For June 2026," Travel.State.Gov, accessed May 28, 2026, <https://travel.state.gov/content/travel/en/legal/visa-lawO/visa-bulletin/2026/visa-bulletin-for-june-2026.html>

⁶ The White House, "Restriction on Entry of Certain Nonimmigrant Workers," *Presidential Actions*, September 19, 2025, <https://www.whitehouse.gov/presidential-actions/2025/09/restriction-on-entry-of-certain-nonimmigrant-workers/>; U.S. Citizenship and Immigration Services, "DHS Changes Process for Awarding H-1B Work Visas to Better Protect American Workers," USCIS, December 23, 2025, <https://www.uscis.gov/newsroom/news-releases/dhs-changes-process-for-awarding-h-1b-work-visas-to-better-protect-american-workers>

The competitive position of the U.S. research system is another important factor. A decade ago, elite Chinese researchers with international ambitions had few credible alternatives to U.S. institutions. The DeepSeek data shows that has changed. Microsoft Research India, the National University of Singapore, and University College London appear in trajectories where U.S. institutions once would have.

China's growing domestic AI talent supply chain is far harder to address. The 10 Key Team members who never left China contributed to a frontier model without any U.S. exposure. So, too, did 145 contributors in the full cohort. Understanding what Chinese institutions are doing to produce researchers of this caliber at scale, and how U.S. K–12 and higher education in computer science and other STEM fields would need to change to produce comparable outcomes, is the longer-horizon question.

A year ago, the DeepSeek findings were already enough to unsettle the assumption that the United States held a durable lead in AI talent. The 2026 data goes further. It shows a Chinese AI workforce that is growing fast and increasingly producing frontier-model contributors without routing them through American institutions during their careers as much as they once did.

None of these patterns are likely to reverse quickly or without concerted policy changes. The question is no longer whether the United States can maintain a permanent talent advantage. The question is whether it can compete on the terms the data now describes, which means treating the comparative advantage of U.S. research institutions as something to be earned rather than assumed, and treating the domestic Chinese pipeline as a development to be acknowledged.

Appendix A: DeepSeek Research Papers (2024–2026)

The following seven papers released by DeepSeek AI between January 2024 and April 2026 formed the basis for our institutional and author-trajectory analysis:

1. DeepSeek LLM: Scaling Open-Source Language Models with Longtermism

arXiv:2401.02954 — January 2024

Full Abstract: The rapid development of open-source large language models (LLMs) has been truly remarkable. However, the scaling law described in previous literature presents varying conclusions, which casts a dark cloud over the scaling of LLMs. We delve into the study of scaling laws and present our distinctive findings that facilitate scaling of large-scale models in two commonly used open-source configurations, 7B and 67B. Guided by the scaling laws, we introduce DeepSeek LLM, a project dedicated to advancing open-source language models with a long-term perspective. To support the pretraining phase, we have developed a dataset that currently consists of two trillion tokens and is continuously expanding. We further conduct supervised fine-tuning (SFT) and direct preference optimization (DPO) on DeepSeek LLM Base models, resulting in the creation of DeepSeek Chat models. Our evaluation results demonstrate that DeepSeek LLM 67B surpasses Llama-2 70B on various benchmarks, particularly in the domains of code, mathematics, and reasoning. Furthermore, open-ended evaluations reveal that DeepSeek LLM 67B Chat exhibits superior performance compared to GPT-3.5.

2. DeepSeek-V2: A Strong, Economical, and Efficient Mixture-of-Experts Language Model

arXiv:2405.04434 – May 2024

Full Abstract: We present DeepSeek-V2, a strong Mixture-of-Experts language model characterized by economical training and efficient inference. It comprises 236B total parameters, of which 21B are activated for each token, and supports a context length of 128K tokens. DeepSeek-V2 adopts innovative architectures, including Multi-head Latent Attention and DeepSeekMoE. MLA guarantees efficient inference through significantly compressing the Key-Value (KV) cache into a latent vector, while DeepSeekMoE enables training strong models at an economical cost through sparse computation. Compared with DeepSeek 67B, DeepSeek-V2 achieves significantly stronger performance, and meanwhile saves 42.5 percent of training costs, reduces the KV cache by 93.3 percent, and boosts the maximum generation throughput to 5.76 times. We pretrain DeepSeek-V2 on a high-quality and multisource corpus consisting of 8.1T tokens, and further perform supervised fine-tuning and reinforcement learning (RL) to fully unlock its potential. Evaluation results show that, even with only 21B activated parameters, DeepSeek-V2 and its chat versions still achieve top-tier performance among open-source models.

3. DeepSeek-Coder-V2: Breaking the Barrier of Closed-Source Models in Code Intelligence

[*arXiv:2406.11931*](#) – June 2024

Full Abstract: We present DeepSeek-Coder-V2, an open-source Mixture-of-Experts code language model that achieves performance comparable to GPT4-Turbo in code-specific tasks. Specifically, DeepSeek-Coder-V2 is further pretrained from an intermediate checkpoint of DeepSeek-V2 with an additional six trillion tokens. Through this continued pretraining, DeepSeek-Coder-V2 substantially enhances the coding and mathematical reasoning capabilities of DeepSeek-V2, while maintaining comparable performance in general language tasks. Compared to DeepSeek-Coder-33B, DeepSeek-Coder-V2 demonstrates significant advancements in various aspects of code-related tasks as well as reasoning and general capabilities. Additionally, DeepSeek-Coder-V2 expands its support for programming languages from 86 to 338, while extending the context length from 16K to 128K. In standard benchmark evaluations, DeepSeek-Coder-V2 achieves superior performance compared to closed-source models such as GPT4-Turbo, Claude 3 Opus, and Gemini 1.5 Pro in coding and math benchmarks.

4. DeepSeek-V3 Technical Report

[*arXiv:2412.19437*](#) – December 2024

Full Abstract: We present DeepSeek-V3, a strong Mixture-of-Experts language model with 671B total parameters with 37B activated for each token. To achieve efficient inference and cost-effective training, DeepSeek-V3 adopts Multi-head Latent Attention and DeepSeekMoE architectures, which were thoroughly validated in DeepSeek-V2. Furthermore, DeepSeek-V3 pioneers an auxiliary-loss-free strategy for load balancing and sets a multitoken prediction training objective for stronger performance. We pretrain DeepSeek-V3 on 14.8 trillion diverse and high-quality tokens, followed by supervised fine-tuning and reinforcement learning stages to fully harness its capabilities. Comprehensive evaluations reveal that DeepSeek-V3 outperforms other open-source models and achieves performance comparable to leading closed-source models. Despite its excellent performance, DeepSeek-V3 requires only 2.788M H800 GPU hours for its full training. In addition, its training process is remarkably stable. Throughout the entire training process, we did not experience any irrecoverable loss spikes or perform any rollbacks. The model checkpoints are available at this URL.

5. DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning

[*arXiv:2501.12948*](#) – January 2025

Full Abstract: We introduce our first-generation reasoning models, DeepSeek-R1-Zero and DeepSeek-R1. DeepSeek-R1-Zero, a model trained via large-scale reinforcement learning without supervised fine-tuning as a preliminary step, demonstrates remarkable reasoning capabilities. Through RL, DeepSeek-R1-Zero naturally emerges with numerous powerful and intriguing reasoning behaviors. However, it encounters challenges such as poor readability and language mixing. To address these issues and further enhance reasoning performance, we introduce DeepSeek-R1, which incorporates multistage training and cold-start data before RL. DeepSeek-R1



achieves performance comparable to OpenAI o1-1217 on reasoning tasks. To support the research community, we open-source DeepSeek-R1-Zero, DeepSeek-R1, and six dense models (1.5B, 7B, 8B, 14B, 32B, 70B) distilled from DeepSeek-R1 based on Qwen and Llama.

6. DeepSeek-V3.2: Pushing the Frontier of Open Large Language Models

[arXiv:2512.02556](#) – December 2025

Full Abstract: We introduce DeepSeek-V3.2, a model that harmonizes high computational efficiency with superior reasoning and agent performance. The key technical breakthroughs of DeepSeek-V3.2 are as follows: (1) DeepSeek Sparse Attention (DSA): We introduce DSA, an efficient attention mechanism that substantially reduces computational complexity while preserving model performance in long-context scenarios. (2) Scalable Reinforcement Learning Framework: By implementing a robust reinforcement learning protocol and scaling post-training compute, DeepSeek-V3.2 performs comparably to GPT-5. Notably, our high-compute variant, DeepSeek-V3.2-Special, surpasses GPT-5 and exhibits reasoning proficiency on par with Gemini-3.0-Pro, achieving gold-medal performance in both the 2025 International Mathematical Olympiad (IMO) and the International Olympiad in Informatics (IOI). (3) Large-Scale Agentic Task Synthesis Pipeline: To integrate reasoning into tool-use scenarios, we developed a novel synthesis pipeline that systematically generates training data at scale. This methodology facilitates scalable agentic post-training, yielding substantial improvements in generalization and instruction-following robustness within complex, interactive environments.

7. DeepSeek-V4: Towards Highly Efficient Million-Token Context Intelligence

[Link](#) – April 2026

We present a preview version of the DeepSeek-V4 series, including two strong Mixture-of-Experts (MoE) language models — DeepSeek-V4-Pro with 1.6T parameters (49B activated) and DeepSeek-V4-Flash with 284B parameters (13B activated) — both supporting a context length of one million tokens. DeepSeek-V4 series incorporates several key upgrades in architecture and optimization: (1) a hybrid attention architecture that combines Compressed Sparse Attention (CSA) and Heavily Compressed Attention (HCA) to improve long-context efficiency; (2) ManifoldConstrained Hyper-Connections (mHC) that enhance conventional residual connections; (3) and the Muon optimizer for faster convergence and greater training stability. We pre-train both models on more than 32T diverse and high-quality tokens, followed by a comprehensive post-training pipeline that unlocks and further enhances their capabilities. DeepSeek-V4-ProMax, the maximum reasoning effort mode of DeepSeek-V4-Pro, redefines the state-of-the-art for open models, outperforming its predecessors in core tasks. Meanwhile, DeepSeek-V4 series is highly efficient in long-context scenarios. In the one-million-token context setting, DeepSeekV4-Pro requires only 27 percent of single-token inference FLOPs and 10 percent of KV cache compared with DeepSeek-V3.2. This enables us to routinely support one-million-token contexts, thereby making long-horizon tasks and further test-time scaling more feasible. The model checkpoints are available at <https://huggingface.co/collections/deepseek-ai/deepseek-v4>.

Appendix B: Additional Tables

Table 3 - Author Count and Contributor Annotations

In the previous analysis, contributor counts included all listed authors across each paper. Here, counts are restricted to those identified on research and engineering teams. Data annotation, business, and compliance contributors were excluded on the grounds that they are unlikely to hold publication profiles in the databases we scanned and fall outside the population of researchers this analysis is tracking. The table below shows, for each included paper, the total number of authors meeting that criterion, whether the paper distinguished a core contributor tier, and if so, how many of those were included.

Paper	Authors	No. 'Core'	No. Other	No. Departed
LLM	86	-	53: Data Annotation (36), Compliance (7), Business (8), Design Team (2)	-
V2	107	-	51: Data Annotation (32), Business & Compliance (19)	-
VCoder 2	39	4	-	-
V3	150	-	49: Data Annotation (31), Business & Compliance (18)	10
R1	196	18	-	5
V3.2	212	-	51: Data Annotation (27), Business & Compliance (24)	8
V4	270	-	48: Business & Compliance (48)	10

Across the seven papers, 356 unique authors were identified. Of these, 282 had OpenAlex profiles available for analysis. The 75 unresolved authors reflect in part the recency of Paper 7, released in mid-April 2026 and not yet indexed in OpenAlex: 66 appeared solely on Paper 7, and six appeared only on Papers 6 and 7. The remaining two, Fucong Dai and Xiangyue Jin, were unresolved in our first analysis as well.

Table 4 - U.S. Institutions Affiliated with DeepSeek Researchers

The following list includes the 113 US-based academic, research, medical, and industry institutions where DeepSeek authors have held prior or current affiliations. This includes both educational and professional roles. Asterisk (*) indicates a current affiliation based on the most recent OpenAlex data.

No. Authors	Institution	City	State
7	Microsoft (United States)	Redmond	Washington
5	Columbia University	New York	New York
4	Carnegie Mellon University*	Pittsburgh	Pennsylvania
4	Stanford University	Stanford	California
4	State University of New York	Albany	New York
3	Auburn University	Auburn	Alabama
3	The University of Texas at Austin*	Austin	Texas
3	University of California, Davis	Davis	California
2	Dana-Farber Cancer Institute	Boston	Massachusetts
2	Intel (United States)	Santa Clara	California
2	Iowa State University	Ames	Iowa
2	Johns Hopkins Medicine	Baltimore	Maryland
2	Johns Hopkins University	Baltimore	Maryland
2	Massachusetts Institute of Technology	Cambridge	Massachusetts
2	New York University	New York	New York
2	Northwestern Polytechnic University*	Fremont	California
2	The Ohio State University	Columbus	Ohio
2	The University of Texas Health Science Center at Houston	Houston	Texas
2	The University of Texas MD Anderson Cancer Center	Houston	Texas
2	The University of Texas Southwestern Medical Center*	Dallas	Texas
2	University at Buffalo, State University of New York	Buffalo	New York
2	University of California, San Diego	La Jolla	California
2	University of California, Berkeley*	Berkeley	California
2	University of Maryland, Baltimore	Baltimore	Maryland
2	University of North Carolina at Charlotte*	Charlotte	North Carolina
2	University of Notre Dame	Notre Dame	Indiana
2	University of Pennsylvania*	Philadelphia	Pennsylvania
2	University of Southern California	Los Angeles	California
2	Veterans Health Administration	Various	Various
2	Virginia Tech	Blacksburg	Virginia
1	Abraham Baldwin Agricultural College*	Tifton	Georgia
1	Alcoa (United States)	Pittsburgh	Pennsylvania
1	Alibaba Group (United States)	San Mateo	California
1	Ames National Laboratory	Ames	Iowa



Table 4 - US Institutions Affiliated with DeepSeek Researcher (cont'd)

No. Authors	Institution	City	State
1	Apple (United States)	Cupertino	California
1	Applied Materials (United States)	Santa Clara	California
1	Arconic (United States)	Pittsburgh	Pennsylvania
1	Baylor College of Medicine	Houston	Texas
1	Beth Israel Deaconess Hospital	Boston	Massachusetts
1	Block Engineering (United States)	Southborough	Massachusetts
1	Boston Consulting Group (United States)	Boston	Massachusetts
1	California Institute of Technology	Pasadena	California
1	California University of Pennsylvania*	California	Pennsylvania
1	CAO Group (United States)	West Jordan	Utah
1	Capital University	Bexley	Ohio
1	Child Mind Institute	New York	New York
1	Children's Cancer Center	Tampa	Florida
1	City College	New York	New York
1	Colorado State University	Fort Collins	Colorado
1	Cornell University	Ithaca	New York
1	Engineering Systems (United States)	Omaha	Nebraska
1	Entrepreneurial Ecosystems	Lincoln	Nebraska
1	Fairchild Semiconductor (United States)	San Jose	California
1	Framingham State University*	Framingham	Massachusetts
1	Fred Hutch Cancer Center	Seattle	Washington
1	Goodyear (United States)	Akron	Ohio
1	Google (United States)	Mountain View	California
1	Graduate School USA	Washington	D.C.
1	Helix (United States)	San Mateo	California
1	Howard Hughes Medical Institute	Chevy Chase	Maryland
1	Hunter College	New York	New York
1	Illumina (United States)	San Diego	California
1	Indiana University Melvin and Bren Simon Comprehensive Cancer Center	Indianapolis	Indiana
1	Intelligent Automation (United States)*	Rockville	Maryland
1	International Craniofacial Institute	Dallas	Texas
1	John D. Dingell VA Medical Center	Detroit	Michigan
1	King University*	Bristol	Tennessee
1	Lehigh University	Bethlehem	Pennsylvania
1	Loyola University Medical Center	Maywood	Illinois
1	Mars (United States)	McLean	Virginia
1	MEI Research (United States)	Edina	Minnesota
1	Michigan State University	East Lansing	Michigan
1	Michigan Technological University	Houghton	Michigan



Table 4 - US Institutions Affiliated with DeepSeek Researcher (cont'd)

No. Authors	Institution	City	State
1	Mississippi State University	Starkville	Mississippi
1	National Institute of Diabetes and Digestive and Kidney Diseases	Bethesda	Maryland
1	New York Institute of Technology	Old Westbury	New York
1	North Carolina State University	Raleigh	North Carolina
1	Oklahoma State University, Oklahoma City	Oklahoma City	Oklahoma
1	Otterbein University	Westerville	Ohio
1	Penn State Milton S. Hershey Medical Center	Hershey	Pennsylvania
1	Public Health Dayton & Montgomery County	Dayton	Ohio
1	Purdue University, West Lafayette*	West Lafayette	Indiana
1	Renal Research Institute	New York	New York
1	Rockefeller University	New York	New York
1	S.P.E.C.I.E.S.	Ventura	California
1	Scientific Methods (United States)*	Granger	Indiana
1	Seattle University	Seattle	Washington
1	Sidney Kimmel Comprehensive Cancer Center	Baltimore	Maryland
1	Southwest Florida Research	Naples	Florida
1	Stony Brook University*	Stony Brook	New York
1	Structural Integrity Associates (United States)	San Jose	California
1	Target (United States)	Minneapolis	Minnesota
1	The University of Texas at Arlington	Arlington	Texas
1	The University of Texas at San Antonio Health Science Center	San Antonio	Texas
1	Union University	Jackson	Tennessee
1	University of Baltimore*	Baltimore	Maryland
1	University of Chicago	Chicago	Illinois
1	University of Delaware	Newark	Delaware
1	University of Houston	Houston	Texas
1	University of Illinois Chicago	Chicago	Illinois
1	University of Illinois Urbana-Champaign	Champaign	Illinois
1	University of Kentucky	Lexington	Kentucky
1	University of Memphis*	Memphis	Tennessee
1	University of Michigan	Ann Arbor	Michigan
1	University of Minnesota	Minneapolis	Minnesota
1	University of Missouri	Columbia	Missouri
1	University of North Dakota	Grand Forks	North Dakota
1	University of Oklahoma	Norman	Oklahoma
1	University of Wisconsin-Madison	Madison	Wisconsin
1	William & Mary	Williamsburg	Virginia
1	Yale University	New Haven	Connecticut
1	Yankton Rural Area Health Education Center	Yankton	South Dakota

Table 5 - Scholarly Output and Citation Metrics of DeepSeek Teams

		All Authors (no. 282)	Key Team (no. 31)	US Affiliated (no. 80)
Works Count	Average	59.99	76.68	131.40
	Median	23.50	56.00	78.00
Citation Count	Average	1763.85	2470.65	4108.50
	Median	681.00	1200.00	1452.00
h-Index	Average	12.46	17.00	21.26
	Median	8.00	12.00	10.49
i10-Index	Average	22.55	33.84	23.19
	Median	7.00	16.00	16.50
Unique Country Affiliations	Average	1.77	2.00	2.76
	Median	1.00	2.00	2.00

Appendix C - List of Key Team Researchers, 2026

The core contributor list holds at 31 names, though the composition has shifted slightly from our last publication. Fuli Luo and Yanhong Xu, who appeared in the previous list, have been removed: Yanhong Xu was absent from paper 3, thus appearing on only six of the seven papers analyzed, while Fuli Luo appeared on only five, missing Papers 6 and 7. This is consistent with Fuli Luo being listed as departed in Paper 5.

Runxin Xu and Yukun Li replace them. Both appeared across all seven papers, but because their names were recorded under initials in some papers — R. X. Xu and Y. K. Li — we did not match those entries to their full-name profiles, and they were therefore excluded from the original count of ‘Key Team’ researchers.

- | | | |
|-------------------|-------------------|------------------|
| 1. Bingxuan Wang | 12. Junxiao Song | 23. Xiao Bi |
| 2. Chenggang Zhao | 13. Kai Dong | 24. Xin Xie |
| 3. Chengqi Deng | 14. Kang Guan | 25. Yaohui Wang |
| 4. Chong Ruan | 15. Liyue Zhang | 26. Yishi Piao |
| 5. Damai Dai | 16. Peiyi Wang | 27. Yukun Li |
| 6. Daya Guo | 17. Qihao Zhu | 28. Yuxiang You |
| 7. Dejian Yang | 18. Qiushi Du | 29. Zhenda Xie |
| 8. Deli Chen | 19. Runxin Xu | 30. Zhewen Hao |
| 9. Hanwei Xu | 20. Shirong Ma | 31. Zhihong Shao |
| 10. Huazuo Gao | 21. Wenfeng Liang | |
| 11. Jiashi Li | 22. Wenjun Gao | |



更多免费行业报告

并购家

www.ipoiipo.cn

白皮书

更新：DeepSeek AI 与人才大竞争

2026年6月

Amy Zegart 与 Emerson Johnston

作者简介

Emerson Johnston，胡佛研究所技术政策加速器全职研究助理，斯坦福大学福特·多西国际政策硕士，东北大学理学学士及文学学士。

Amy Zegart，胡佛研究所Morris Arnold与Nona Jean Cox高级研究员，胡佛技术政策加速器主任，斯坦福大学以人为本AI研究院高级研究员及副主任。

引用格式：Emerson Johnston and Amy Zegart, "Update: DeepSeek AI and the Great Talent Competition," Technology Policy Accelerator, Hoover Institution and the Stanford Institute for Human-Centered Artificial Intelligence, June 2026. <https://doi.org/10.64576/0626>

更新：DeepSeek AI与人才大竞争

2025年1月，DeepSeek AI发布R1论文，揭示了一个以远低于OpenAI o1计算量即可实现同等水平的模型，媒体叙事随即被两种误解主导。第一种误解是：这家中国公司的突破源于只会模仿、不会创新的中国研究人员；第二种误解是：研究团队年轻稚嫩，在计算机科学领域普遍缺乏广泛引用。我们在2025年的报告《深度透视DeepSeek AI的人才及其对美国创新的影响》中，对支撑该公司五篇基础论文的223名研究人员进行了深入考察，发现的图景截然不同。

我们发现，DeepSeek在多个重要方面确实实现了创新；其研究团队的经验远比大多数人预想的丰富；超过一半的团队从未出国求学或工作。这些发现共同揭示：中国国内的AI人才培养管道规模更大、实力更强，远超主流叙事所承认的程度。这些发现还表明，任何建立在“中国缺乏自主超越美国所需人才或技术积累”这一假设之上的美国战略，都已与现实严重脱节。

我们当时的政策建议是：美国必须更积极地吸引、欢迎并留住全球顶尖人才，同时通过改善国内K-12及高等教育中的STEM教育，紧迫培育本土能力。

本报告探讨：过去一年来，情况发生了多少变化。我们将分析范围从五篇扩展至七篇论文——纳入DeepSeek V3.2（2025年12月）和V4（2026年4月），以涵盖该公司最新研究成果。通过扩展原始数据集，作者池从223人增至356人，并对其中282人建立了追踪其求学、生活和工作轨迹的完整档案。

中国国内AI人才管道依然强劲

- 在有记录的职业生涯中，始终仅与中国机构有关联的DeepSeek AI研究人员比例为53.5%。
- 31名核心团队成员（即在过去两年发表的全部七篇DeepSeek AI论文上均有署名的研究人员）中，有10人同样仅与中国机构有关联。这意味着，在这个能够产出与OpenAI o1模型性能相当的中国模型的团队中，约三分之一的核心贡献者从未在国际研究环境中接受训练、发表论文或从事工作。

".....在这个能够产出与OpenAI o1模型性能相当的中国模型的团队中，约三分之一的最核心贡献者从未在国际研究环境中接受训练、发表论文或从事工作。"

新数据与更完善的数据揭示了更长的美国培训周期，顶尖人才尤为突出

去年，63.3%的美国关联研究人员仅有一年美国经历。这一数字现已降至35%，近半数（48.8%）集中在2至4年区间。这一转变既反映了新增DeepSeek人才的构成，也体现了我们对原有队列数据的改进。值得注意的是，第6和第7篇论文（2026年）的141名新增贡献者，是比最初参与2025年各论文的223人更具国际流动性的群体。此外，OpenAlex的机构覆盖率也全面提升，为我们去年追踪的部分研究人员补充了此前数据库中尚未收录的美国任职记录。

新研究人员与改进数据共同描绘出的图景，并非短暂的交流访问，而是在美国研究环境中的深度嵌入，随后70.2%的国际流动DeepSeek贡献者选择回流。

- 更长的美国停留时间并未带来更高的留美率。本数据集中的13名长期留美研究人员，在美国机构合计有记录的时间超过119年，但其中大多数人现在都在中国。
- 尽管13名最长期留美研究人员中有9人最终回到中国，但他们并非简单地来到美国、停留多年后便返回。实际上，他们在多年间经历了多次国际转场。
- 80名与美国机构有任何关联的DeepSeek AI研究人员，平均而言是整个研究人员群体中学术成就最高的，其中相当一部分人在回到中国之前，曾在美国机构深度嵌入五年或更长时间。

".....中国国内管道正在培育出无需经过太多美国或国际经历的前沿模型贡献者。与此同时，先前的美国经历仍在持续为中国AI前沿企业带来回报。"

这一含义，往好了说也令人不安。已随这些研究人员迁移的知识，无法通过此后收紧签证政策或出口管制来追回。限制未来准入能解决一种人才问题，但第二种问题依然存在。中国国内管道正在培育出无需经过太多美国或国际经历的前沿模型贡献者。与此同时，先前的美国经历仍在持续为处于前沿的中国AI企业带来红利。2026年的数据清楚地表明，美国正同时面临这两种问题。

研究方法

完整方法论细节请参见原报告。简言之，我们识别了DeepSeek七篇论文的研究与工程贡献者，数据来源为arXiv（2024–2025年）和Hugging Face（2026年4月）。

356名署名作者中，282人在OpenAlex中有可用档案——我们于2026年4月查询了其发表记录、引用指标及机构关联历史，追溯至1989年（该年为OpenAlex对本队列任何研究人员均有可用关联记录的最早年份）。随后，我们用Python绘制每位研究人员的完整机构历史，以追踪跨境流动轨迹。

与先前分析统计各论文所有署名贡献者不同，本次分析将范围限定于研究和工程团队中的成员。我们认为，此方法论调整更聚焦于研究贡献本身，并将有助于未来开展更好的纵向分析。

在区分贡献者角色的论文中，“数据标注”、“商务”和“合规”贡献者被排除在外，原因在于他们不属于最相关于理解AI人才竞争动态的研究人员群体。在不区分角色的论文中，所有署名贡献者均被纳入。各论文的完整计数（包括被排除贡献者的总计）见附录B表1。

团队规模更大，结构不变

七篇论文共识别出356名研究与工程贡献者，比上次分析识别的223人有所增加。在一年内增长57%，就绝对数量而言规模可观，相对于团队此前的增长轨迹而言同样显著：前五篇论文涵盖约两年时间，研究人员稳步增加；后两篇论文则在14个月内净增约141名新贡献者。同期离职人数为33人（第6篇论文8人，第7篇10人，其余集中于第4和第5篇），这意味着DeepSeek每流失一名研究人员，就会引进逾四名新人。

我们再次发现，31名研究人员（略低于9%）参与了全部七篇论文，我们称之为“核心团队”。另有47名作者出现在六篇论文中，17人出现在五篇，41人出现在四篇，5人出现在三篇，79人出现在两篇，136人仅出现在一篇（见图1）。

图1：DeepSeek论文的作者分布

356名贡献者在七篇最新DeepSeek AI论文中的作者分布情况

论文数	作者数	占比
1篇	136	38.20%
2篇	79	22.19%
3篇	5	1.40%
4篇	41	11.52%
5篇	17	4.78%
6篇	47	13.20%
核心团队（7篇）	31	8.71%

更新：DeepSeek AI与人才大竞争

2025年和2026年这一分布形态的持续存在，揭示了DeepSeek组织研究劳动力的具体方式。偶发贡献者聚集于两端：215名研究人员（约占总数的三分之二）仅出现在一至两篇论文中。中间层几近消失。跨越两篇论文门槛的研究人员往往持续参与，而未能突破者则止步于此。这一模式与双轨制劳动力结构相吻合：第一轨是为特定模型发布而调入的大批轮换贡献者，第二轨是参与所有发布的稳定核心。31人组成的核心团队和47名六论文贡献者属于第一轨，215名一至两论文贡献者属于第二轨。三至五论文层极为稀薄，因为这本就是异常情况。

图2从三个维度切割同一总体：核心团队（参与全部论文）、中间层（参与多于一篇但少于全部）以及仅参与一篇的贡献者。2025年至2026年间，核心团队在绝对数量上维持31人不变，但占总体比例从13.9%降至8.7%。中间层从169人增至189人。一论文组从23人增至136人，近六倍增长，占总体比例从10.3%跃升至38.2%。

图2：DeepSeek AI研究人员作者分布变化，2025-2026年

类别	2025年	2026年
核心团队（全部论文）	31人（13.90%）	31人（8.71%）
中间层（多于1篇，少于全部）	169人（75.78%）	189人（53.09%）
仅一篇贡献者	23人（10.31%）	136人（38.20%）

在141名新增研究人员中，71人作为一论文贡献者入场，另外70人直接进入中间层（同时出现在第6和第7篇论文中）。这表明DeepSeek在过去一年中双轨结构得到强化：公司的增长完全来自轮换贡献者一侧，而稳定核心保持规模不变。

还有一个值得关注的规律：43名研究人员出现在第1、2、4、5、6和7篇论文中，仅缺席第3篇（VCoder 2）。结合参与全部七篇的31人，这表明第3篇依托于一支与主模型开发轨道不完全重叠的专业团队，与VCoder 2相对于旗舰V系列发布的更聚焦定位相吻合。

中国国内AI人才管道——一年之后

我们原始分析涵盖201名具有已知关联数据的DeepSeek作者。新数据集将分析范围扩展至271名拥有机构关联记录的作者（共356名作者中）。

去年的核心发现依然成立：中国拥有一条强劲且持续壮大的国内AI人才管道。略超半数——53.5%，即145名作者——在有记录的职业生涯中始终仅与中国机构有关联，从未在中国以外担任职务或学术职位，仅略低于原始分析中的55.2%。我们再次发现，几乎所有DeepSeek研究人员（在扩展数据集中占97%）至少曾与一家中国机构有确认的关联。

与美国机构有任何关联的作者比例也有所上升：当前样本中有29.5%（n=80）曾与美国机构有关联，高于原始分析中的24.4%。

目前的快照呈现的故事与一年前大致相同（见图3）。在271名有已知关联的作者中，229人（84.5%）目前与中国机构有关联，18人（6.64%）目前在美国机构。其余分布在英国、中国台湾、日本、新加坡及少数其他国家，均未超过1.5%。

图3：当前机构关联的地域分布

271名具有关联数据的DeepSeek AI研究人员当前关联的分布情况

国家/地区	作者数	占比
中国	229	84.50%
美国	18	6.64%
英国	4	1.48%
中国台湾	3	1.11%
日本	3	1.11%
新加坡	2	0.74%
波兰	2	0.74%
加拿大	2	0.74%
澳大利亚	2	0.74%
韩国	1	0.37%
沙特阿拉伯	1	0.37%
爱尔兰	1	0.37%
印度	1	0.37%
中国香港	1	0.37%
芬兰	1	0.37%

专栏1：新研究人员的学术背景

在今年数据集中的141名全新研究人员里，69人有可解析的OpenAlex档案，68人有关联数据，其中18人（26%）曾与美国机构有过或现有关联。这一比例与全部作者池的29%基本相同，表明DeepSeek近期招募并未比早期招募更多或更少地倾向于有美国培训背景的研究人员。18名研究人员分布于37所美国机构（见表1）。

仅卡内基梅隆大学和德克萨斯大学西南医学中心各有逾一名新贡献者。表1中所有其他机构均代表单一研究人员的职业节点。差异在于这些研究人员在各机构停留的时长：一名新贡献者在贝勒医学院有记录的时间长达八年，另一名在弗吉尼亚理工大学停留六年，还有一名在贝斯以色列女执事医院停留五年。这些高关联年数反映的是单一研究人员的长期深度嵌入，而非向多名DeepSeek目标研究人员输送人才的机构管道。

表1：新一批DeepSeek AI研究人员曾就读/工作的美国机构（部分）

人数	机构	城市	州
2	卡内基梅隆大学*	匹兹堡	宾夕法尼亚州
2	德克萨斯大学西南医学中心*	达拉斯	德克萨斯州
1	奥本大学	奥本市	阿拉巴马州
1	贝勒医学院	达拉斯	德克萨斯州
1	贝斯以色列女执事医院	波士顿	马萨诸塞州
1	斯坦福大学	斯坦福	加利福尼亚州
1	麻省理工学院	剑桥	马萨诸塞州
1	加州大学伯克利分校*	伯克利	加利福尼亚州
1	加州大学圣地亚哥分校	圣地亚哥	加利福尼亚州
1	宾夕法尼亚大学	费城	宾夕法尼亚州
1	耶鲁大学	纽黑文	康涅狄格州
1	弗吉尼亚理工大学	布莱克斯堡	弗吉尼亚州
1	圣母大学*	布卢明顿	印第安纳州
1	南加州大学	欧文	加利福尼亚州

注：*表示当前关联机构。

美中人才管道的实际流向

在80名与美国机构有任何关联的DeepSeek研究人员中，与原始分析相比，最大修正出现在时长数据上（见图4）。去年的数据集记录显示，49名美国关联研究人员中有63.3%仅有一年美国经历。在扩展数据集中，相应比例降至35%（n=28）。近半数——48.8%（n=39）——现在集中在2至4年区间，16.3%（n=13）在美国机构停留五年以上。

这一新数据修正了去年对美国不过是顶尖中国AI研究人员“中转站”的判断。在DeepSeek被引用最多的贡献者中，有相当一部分人在美国度过了更长、更具塑造性的时光。对他们而言，美国仍是一个节点，但比我们去年报告的重要得多。

13名长期留美研究人员中，有些是整个数据集中学术成就最高的。张鸣华目前在美国石溪大学，有28年美国机构关联记录，职业引用近30,000次。陈东的51,823次引用和112的H指数在美国关联队列引用总量中占据不成比例的份额，曾三次在美国停留共12年，目前仍在弗雷明汉州立大学。另一方面，周松阳在贝斯以色列女执事医院、洛克菲勒大学、麻省理工学院、贝勒医学院和加州大学圣地亚哥分校度过了18年后，回到中国海南医科大学。

图4：美国经历时长

80名与美国机构有关联的DeepSeek研究人员的美国经历时长分布

时长	人数	占比
1年	28	35.0%
2-4年	39	48.8%
5年以上	13	16.3%

值得注意的是，使这13段职业生涯与众不同之处，在于它们极少呈现出单次嵌入后直接返回的形态。13人中有9人经历了不止一次国际转场。徐磊的记录历经美国机构，再到英国，再到文莱，最终回到中国。唐明晖的记录包含在美国和中国停留之间穿插日本和英国，并经由清华大学回国。张B则从瑞典起步，随后在美中机构之间交替。对这一群体而言，美国机构嵌入于持续的跨国职业生涯之中，而非一次性的资历加盖之旅。方向性权重重要但非绝对：13人中有8人最终落脚中国，仅5人留在美国，且美国经历时长并不能预测结果——张鸣华的28年和任同政的五年都终止于美国，周松阳的18年和张若宇的六年都终止于中国。

图5：美国留存率

80名与美国机构有关联的DeepSeek研究人员的完整留存数据

流动轨迹	人数（占比）
始于中国，止于中国（经历美国）	31（38.8%）
始于美国，止于中国	19（23.8%）

仅与美国有关联	7 (8.8%)
始于中国，止于美国	10 (12.5%)
始于美国，止于美国	5 (6.3%)
始于中国，止于海外（非美）	4 (5.0%)
始于海外，止于中国	3 (3.8%)
始于美国，止于海外（非中）	1 (1.3%)

图6：整体留存率

271名具有关联数据的DeepSeek研究人员的完整留存数据

类别	人数（占比）
仅与中国有关联	145（53.5%）
始于中国，止于中国（经历国际流动）	53（19.6%）
始于美国，止于中国	19（7.0%）
始于美国，止于美国	13（4.8%）
始于中国，止于美国	10（3.7%）
始于海外，止于中国	20（7.4%）
仅与美国有关联	5（1.8%）
始于中国，止于海外	3（1.1%）
仅与海外（非中非美）有关联	2（0.7%）
始于美国，止于海外（非中）	1（0.4%）

上述模式使一场惯于将人才流动视为“管道问题”的政策辩论变得更加复杂——有人认为更好的签证政策或更强的出口管制便可改变方向。一年前，我们将美国机构描述为“跳板”，即让顶尖研究人员获取技能和资历、最终反哺中国AI生态系统的场所。2026年更完整的数据重构了这一判断。

部分美国停留时间短暂且工具性强，与“跳板”描述相符。但相当一部分停留时间比我们此前记录的更长——有时长达十年甚至更久，且回流率并未随停留时间延长而下降。这更接近于中国的“人才回流获益”，而非单纯的“人才循环”。本数据集中的13名长期留美研究人员，在美国机构合计有记录的时间超过119年，其中大多数人现在都在中国。

"本数据集中的13名长期留美研究人员，在美国机构合计有记录的时间超过119年，其中大多数人现在都在中国。"

更新：DeepSeek AI与人才大竞争

此外，在全数据集25名职业生涯起点在美国的研究人员中，19人（76%）最终落脚中国，1人落脚其他海外地区。九名在美国机构深度嵌入五至十八年后最终回国的长期留美研究人员，与其余研究人员的回流率大体相当，这意味着更长的停留时间并未带来更高的留美率。周松阳在美国的18年与关廷陈的三年，都以回到中国告终。

这一新数据不仅证明中国培育了自己的研究人员，还表明中国成功地"招募回"了那些曾在美国受训（有时长达十年以上）的研究人员——他们回国时携带着在美国研究体系内积累的方法、人脉与资历。

机构关联深度分析

扩展数据集中282名作者的职业生涯共涉及812家独立机构，多于去年的499家。中国机构占570家，即70.2%，略低于原始分析中的74%。美国机构占113家（13.9%）。其余129家分布于英国、澳大利亚、中国台湾、加拿大、日本等地。

中国科学院再次在网络中覆盖最广。中科院直接招收了282名作者中的4人，但其170家附属机构——研究院、实验室和专业中心，OpenAlex将其归类为中科院的下属组织——共与104名研究人员有关联，约占全部作者池的37%。这一数字是一年前53人的近两倍，是数据集中任何机构年度增量最大的。

中科院体系之外，与去年的对比使机构图景更加清晰（见表2）。清华大学以46名关联研究人员居首，从16人跃升，增幅近200%。同样，北京大学从21人增至29人，浙江大学从8人增至25人，中国科学院大学从7人增至21人。

2025年处于边缘地位的几所机构如今跻身顶层：东南大学从1名关联研究人员增至15名，北京航空航天大学从3人增至14人，兰州大学从0人增至10人，东北大学（中国）和广州医科大学各从1人增至10人。两所机构出现反向变动：中山大学从10人降至4人，同济大学从7人降至5人，尽管这两处下降可能源于OpenAlex对机构关联在数据集其他位置的重新归属方式。

表2：按研究人员数量排列的顶级机构，2025年与2026年对比

机构	2025年	2026年	变化
中国科学院（体系）	53	104	+51
清华大学	16	46	+30
浙江大学	8	25	+17
中国科学院大学	7	21	+14
东南大学	1	15	+14
北京航空航天大学	3	14	+11
兰州大学	0	10	+10

东北大学（中国）	1	10	+9
广州医科大学	1	10	+9
北京大学	21	29	+8
南京大学	10	16	+6
华中科技大学	5	11	+6
西安交通大学	4	10	+6
上海交通大学	8	13	+5
同济大学	7	5	-2
中山大学	10	4	-6

"初级人才"叙事——一年之后

一年前，主流叙事认为DeepSeek R1模型的成功依赖于应用经验匮乏的应届博士。DeepSeek CEO 梁文峰曾亲口这样说，在接受记者采访时表示，核心技术职位大多由工作经验一至两年的研究人员担任，DeepSeek在招募中更看重"热情与好奇心"而非资历。

我们原始分析对此提出了质疑。我们发现DeepSeek核心团队平均每位作者被引用1,554次，h指数中位数为10——这绝非资历浅薄的图景。一年后，数字进一步印证了这一方向。

在全部作者池中，平均引用次数从约1,000次升至1,763次，中位数翻逾一倍，从249次升至681次——这意味着整个分布向上移动，而非仅尾部如此。i10指数（统计作者被引用至少10次的论文数量）变动较为温和：平均值从略超19升至22.55，中位数从5升至7。h指数的变动幅度相近，平均值从10.8升至12.46，中位数从7升至8。

核心团队的指标变化更为显著。31名作者现在平均引用2,470次，高于一年前的1,554次，增幅59%。引用中位数从约700次升至1,200次，平均h指数从13.5升至17，h指数中位数从10升至12，i10指数平均值达到33.84。这一跃升很大程度上源于DeepSeek论文本身：R1、V3和V2自上次分析以来各自积累了数千次引用，而核心团队是全部论文的作者。从这个意义上说，文献计量学的变化是衡量该领域吸收DeepSeek 2025年产出速度的一个指标。

扩展数据集还使我们能够对此前无法清晰分析的一个群体进行单独审视：80名曾与美国机构有过或现有关联的DeepSeek作者。他们平均发表131篇作品，被引用4,108次，h指数为21.26——超过全体作者池引用平均值的两倍，远超核心团队的水平。引用中位数为1,452次，而全体作者池为681次。从每项指标来看，这都是学术成就最高的群体。

与全球顶尖AI实验室的同行群体相比，这种优势更为突出。对同一美国实验室的分析显示，OpenAI的894名作者平均引用次数为2,481.5次，但中位数仅为100.5次，表明高引用个体之外存在断崖式下跌。此外，该群体h指数中位数仅为2（较我们上次报告的6有所下降），i10指数中位数相同，反映出OpenAI影响力一致性较为有限。这一模式在Anthropic和Alphabet同样存在，两者引用中位数分别仅为各自平均值的15%和20%。DeepSeek为35%。

在我们2025年的《深度透视》报告中，我们将DeepSeek的文献计量学档案描述为：与美国同行相比更均衡，对少数突出人物的依赖更少。一年后，这一描述对全部作者池、核心团队以及与美国机构有关的作者群体均依然成立。

核心团队再审视

在去年的报告中，我们深入考察了DeepSeek的核心团队——那31名名字出现在全部五篇基础论文中的研究人员。今年对OpenAlex的查询呈现了更完整的数据，并揭示了原始计数中的三处错误：徐润欣和李玉坤出现在全部五篇论文中，但由于部分论文以缩写记录其姓名（分别为R. X. Xu和Y. K. Li），未能与全名档案匹配。相反，许燕红去年被列入名单，但实际上未出现在第3篇论文中，意味着她本不应计入全论文贡献者名单。经这三处更正，2025年核心团队应为32人，而非31人。

将分析延伸至七篇论文使核心团队名单进一步收紧。吕福利在第5篇论文后离开DeepSeek，未出现在第6和第7篇中，使2025年全论文计数减少一人。由此，2026年核心团队维持在31人，但这个数字背后的具体名单已有所变化。2026年核心团队全部七篇论文的完整名单见附录C。

31名核心团队成員中，OpenAlex为其中28人返回了关联数据。18人在有记录的职业生涯中曾在中国以外担任过至少一个职位，这意味着10名核心团队研究人员（约占36%）从未离开中国。

18名具有国际流动经历的核心团队成员中，11人曾在美国机构停留，但大多数仅停留一两年。邵志宏在密歇根大学停留四年，是单次美国停留时间最长的；朱启浩在卡内基梅隆大学停留一年，谢振达在哥伦比亚大学停留两年，徐润欣在俄亥俄州立大学停留一年，随后经伦敦帝国理工学院回国。此外，大多数研究人员的时间节点——职业中期而非职业起步阶段——指向访问学者或博士后任命，而非学位攻读，这与有针对性的知识获取而非移民目的相吻合。

七名无美国关联的国际流动核心团队成员走了不同的路径：郭大雅在微软研究院英国分部和一所德国大学工作了三年，后回到中山大学；戴大迈在微软研究院印度分部停留一年；董凯在中国、加拿大、意大利和芬兰之间轮换，目前在奥卢大学任职；高文俊有新加坡、丹麦和中国香港的关联。

综合来看，这些轨迹数据再次表明，这一队列的国际经历短暂、有针对性，且几乎无一例外地以回到中国告终。在这一背景下，美国更像是若干节点之一，而非终点——对某些研究人员而言，它可与微软研究院印度分部或丹麦技术大学互换。核心团队由此呈现的是：一个选择性汲取外部经验后再度聚合的国内研究基地，而非全球招募的产物。

两个人才问题，一场政策辩论

一年前，我们将DeepSeek AI人才池的地缘政治影响围绕一个核心命题展开：美国不能再假设拥有永久性的人才优势。更新后的数据表明，美国面临的是两个独立的人才挑战，而非一个。

挑战一：留住顶尖人才

第一个挑战是在美国留住中国顶尖人才。经过美国机构的80名研究人员，是DeepSeek作者池中成就最高的群体，平均引用4,108次，h指数中位数为16.5，但其中大多数人现在都在中国。核心团队中曾赴美的成员通常只停留一两年；核心团队以外的高引用贡献者停留时间则长得多，有时超过十年。两组人回流中国的比率大体相当，停留时长并未改变结果。

挑战二：中国AI人才供应链的日趋独立

第二个挑战是中国AI人才供应链的日趋独立。31名核心团队成员中有10人没有任何在中国以外的记录关联；在全部作者池中，145名研究人员（53.5%）从未在中国以外任职。这些研究人员为一个在推理基准测试上与OpenAI o1性能相当的模型作出了贡献。他们没有在美国接受培训，没有在美国实验室工作，也没有从美国机构发表论文。

签证政策的改变可以应对第一个挑战，但从未来过美国的研究人员，不在任何关于来美规则的约束范围之内。一场仅聚焦于是否允许中国人才来美并留下的政策辩论，无法解决中国日益壮大的国内人才供应链问题。

提高中国人才在美留存率，是两者中更易着手解决的。在DeepSeek贡献者中有任何美国关联的人员里，70.3%的国际流动者最终落脚中国。美国当前政策将这一数字视为固定的风景线，而非可变的变量。绿卡积压、博士后签证的不确定性，以及缺乏明确的长期途径让STEM研究人员留在美国而非返回母国——这些都是人才留存问题的关键组成部分，且均可加以解决。

我们不应假设中国研究人员离开美国是因为他们想离开。许多人可能是因为不得不离开：美国签证体系使留下来变得过于困难、成本高昂且充满不确定性。一名中国研究人员今天申请就业类绿卡，可能需要等待近五年时间，排队号才能排到，而这条队伍还在变长。与此同时，每份H-1B申请新增10万美元费用，以及重新加权后倾向高薪申请者的抽签机制，将博士后和应届博士推到了队伍最末端。

美国研究体系的竞争地位是另一个重要因素。十年前，有国际抱负的顶尖中国研究人员几乎没有美国机构之外的可信选择。DeepSeek数据表明这一局面已经改变：微软研究院印度分部、新加坡国立大学和伦敦大学学院，已出现在原本属于美国机构的职业轨迹中。

中国日益壮大的国内AI人才供应链则更难应对。从未离开中国的10名核心团队成员，在没有任何美国经历的情况下，为前沿模型作出了贡献；全队列中同样如此的还有145名贡献者。理解中国机构如何批量培育出这一水准的研究人员，以及美国K-12和高等教育在计算机科学及其他STEM领域需要如

何变革才能实现相当的产出，是一个更长远的问题。

一年前，DeepSeek的研究发现就已足以动摇美国在AI人才领域拥有持久领先优势的假设。2026年的数据走得更远：它呈现出一支快速壮大的中国AI劳动力，正越来越多地在职业生涯中无需经由美国机构便能培育出前沿模型贡献者。

这些模式都不太可能在没有集中政策调整的情况下迅速逆转。问题不再是美国能否保持永久性的人才优势，而是它能否在数据所描述的格局下展开竞争——这意味着将美国研究机构的比较优势视为需要赢得的东西，而非理所当然的存在；并将中国国内管道的发展视为需要正视而非回避的现实。

附录A：DeepSeek研究论文（2024–2026年）

以下七篇由DeepSeek

AI于2024年1月至2026年4月发布的论文，构成了本次机构关联及作者轨迹分析的基础：

1. DeepSeek LLM：以长期主义扩展开源语言模型

arXiv:2401.02954 — 2024年1月

摘要：开源大型语言模型（LLM）的快速发展令人瞩目，但此前文献中描述的扩展定律呈现出不同的结论，为LLM的扩展蒙上阴影。本文深入研究扩展定律，并呈现我们在两种常用的开源配置（7B和67B）下促进大规模模型扩展的独特发现。以扩展定律为指引，我们推出DeepSeek LLM，致力于以长远眼光推进开源语言模型发展。为支持预训练阶段，我们开发了目前包含两万亿个词元且持续扩展的数据集，并对DeepSeek LLM基础模型进行了监督微调（SFT）和直接偏好优化（DPO），形成DeepSeek Chat模型。评测结果显示，DeepSeek LLM 67B在各项基准测试中超越Llama-2 70B，尤其在代码、数学和推理领域表现突出。此外，开放式评测表明，DeepSeek LLM 67B Chat在性能上优于GPT-3.5。

2. DeepSeek-V2：一个强大、经济且高效的混合专家语言模型

arXiv:2405.04434 — 2024年5月

摘要：本文介绍DeepSeek-V2，一个以经济训练和高效推理为特点的强大混合专家语言模型。它共有2360亿参数，每个词元激活210亿参数，支持128K词元的上下文长度。DeepSeek-V2采用多头潜在注意力（MLA）和DeepSeekMoE等创新架构。MLA通过将键值（KV）缓存压缩为潜在向量，显著保障了高效推理；DeepSeekMoE则通过稀疏计算，实现了以经济成本训练强大模型。与DeepSeek 67B相比，DeepSeek-V2性能显著提升，同时节省了42.5%的训练成本，KV缓存减少93.3%，最大生成吞吐量提升至5.76倍。评测结果表明，即便仅激活210亿参数，DeepSeek-V2及其对话版本仍在开源模型中达到顶尖性能。

3. DeepSeek-Coder-V2：突破代码智能领域闭源模型的壁垒

arXiv:2406.11931 — 2024年6月

摘要：本文介绍DeepSeek-Coder-V2，一个在代码特定任务上达到GPT4-Turbo水平的开源混合专家代码语言模型。具体而言，DeepSeek-Coder-V2从DeepSeek-V2的中间检查点出发，额外预训练了6万亿个词元，大幅增强了DeepSeek-V2的编程和数学推理能力，同时在通用语言任务上保持了相当的性能。在标准基准评测中，DeepSeek-Coder-V2在编程和数学基准上的性能优于GPT4-Turbo、Claude 3 Opus和Gemini 1.5 Pro等闭源模型。

4. DeepSeek-V3技术报告

arXiv:2412.19437 — 2024年12月

摘要：本文介绍DeepSeek-V3，一个拥有6710亿参数的强大混合专家语言模型，每个词元激活370亿参数。DeepSeek-V3率先提出无辅助损失的负载均衡策略，并设置多词元预测训练目标以提升性能。我们在由14.8万亿高质量多样化词元构成的数据集上对DeepSeek-V3进行预训练。尽管性能卓越，DeepSeek-V3的全量训练仅需277.8万H800 GPU小时，训练过程高度稳定。

5. DeepSeek-R1：通过强化学习激励LLM的推理能力

arXiv:2501.12948 — 2025年1月

摘要：本文介绍我们的第一代推理模型DeepSeek-R1-Zero和DeepSeek-R1。DeepSeek-R1-Zero是一个无需监督微调前置步骤、通过大规模强化学习训练的模型，展现出卓越的推理能力。DeepSeek-R1在推理任务上达到与OpenAI o1相当的性能。为支持研究社区，我们开源了DeepSeek-R1-Zero、DeepSeek-R1，以及从DeepSeek-R1蒸馏的六个密集模型（1.5B、7B、8B、14B、32B、70B）。

6. DeepSeek-V3.2：推动开放大型语言模型的前沿

arXiv:2512.02556 — 2025年12月

摘要：本文介绍DeepSeek-V3.2，一个将高计算效率与卓越推理和智能体性能相融合的模式。关键技术突破包括：DeepSeek稀疏注意力（DSA）机制、可扩展强化学习框架（性能可与GPT-5相媲美，高算力变体DeepSeek-V3.2-Speciale在2025年国际数学奥林匹克和国际信息学奥林匹克上均取得金牌级表现）、以及大规模智能体任务合成管道。

7. DeepSeek-V4：迈向高效百万词元上下文智能

链接 — 2026年4月

摘要：本文介绍DeepSeek-V4系列预览版，包含DeepSeek-V4-Pro（1.6万亿参数，激活490亿）和DeepSeek-V4-Flash（2840亿参数，激活130亿），均支持100万词元的上下文长度。在100万词元上下文设置中，DeepSeek-V4-Pro仅需DeepSeek-V3.2单词元推理FLOPs的27%和10%的KV缓存，使长视野任务更具可行性。

附录B：附加表格

表3：作者数量与贡献者标注

在先前的分析中，贡献者计数包含各篇论文的所有署名作者。本次分析将范围限定于研究和工程团队中的成员。数据标注、商务和合规贡献者被排除，原因在于他们不太可能在我们扫描的数据库中拥有发表档案，也不在本次分析追踪的研究人员群体范围内。下表展示了各收录论文中符合该标准的作者总数、论文是否区分核心贡献者层级，以及（如区分）其中被纳入分析的人数。

七篇论文共识别出356名独特作者，其中282人有OpenAlex档案可供分析。75名未解析作者一定程度上反映了第7篇论文的新近性——该论文于2026年4月中旬发布，尚未被OpenAlex收录：其中66人仅出现在第7篇，6人仅出现在第6和第7篇。另外2人（戴福聪和金香月）在我们的首次分析中同样未能解析。

论文	作者数	核心人数	其他贡献者	离职人数
LLM	86	-	53：数据标注(36)、合规(7)、商务(8)、设计团队(2)	-
V2	107	-	51：数据标注(32)、商务&合规(19)	-
VCoder 2	39	4	-	-
V3	150	-	49：数据标注(31)、商务&合规(18)	10
R1	196	18	-	5
V3.2	212	-	51：数据标注(27)、商务&合规(24)	8
V4	270	-	48：商务&合规(48)	10

表4：与DeepSeek研究人员有关联的美国机构（部分）

以下为113所DeepSeek作者曾或现在与之有关联的美国机构（含高校、研究机构、医疗机构和企业机构）。标注*表示基于OpenAlex最新数据的当前关联机构。

人数	机构	城市	州
7	微软（美国）	■■■	■■■
5	哥伦比亚大学	■■	■■
4	卡内基梅隆大学*	■■■	■■■■■
4	斯坦福大学	■■■	■■■■■
4	纽约州立大学	■■■■	■■
3	奥本大学	■■	■■■■
3	德克萨斯大学奥斯汀分校*	■■■	■■■■

3	加州大学戴维斯分校	■■■■	■■■■■■■
2	麻省理工学院	■■■	■■■■■
2	纽约大学	■■■	■■■
2	俄亥俄州立大学	■■■■	■■■■
2	约翰斯·霍普金斯大学	■■■■■	■■■■
2	加州大学伯克利分校*	■■■■	■■■■■■■
2	加州大学圣地亚哥分校	■■■■	■■■■■■■
2	宾夕法尼亚大学*	■■■	■■■■■■■
2	石溪大学*	■■■	■■■
2	南加州大学	■■■■	■■■■■■■
2	弗吉尼亚理工大学	■■■■■■■	■■■■■
2	英特尔（美国）	■■■■■	■■■■■■■
1	谷歌（美国）	■■■■	■■■■■■■
1	康奈尔大学	■■■■	■■■
1	耶鲁大学	■■■■	■■■■■
1	伊利诺伊大学香槟分校	■■■	■■■■■
1	密歇根大学	■■■■	■■■■
1	威斯康星大学麦迪逊分校	■■■■	■■■■■

表5：DeepSeek各团队的学术产出与引用指标

指标	所有作者 (n=282)	核心团队 (n=31)	美国关联 (n=80)
论文数量 (平均值)	59.99	76.68	131.40
论文数量 (中位数)	23.50	56.00	78.00
引用次数 (平均值)	1,763.85	2,470.65	4,108.50
引用次数 (中位数)	681.00	1,200.00	1,452.00
H指数 (平均值)	12.46	17.00	21.26
H指数 (中位数)	8.00	12.00	10.49
i10指数 (平均值)	22.55	33.84	23.19
i10指数 (中位数)	7.00	16.00	16.50
独特国家关联数 (平均值)	1.77	2.00	2.76
独特国家关联数 (中位数)	1.00	2.00	2.00

附录C：2026年核心团队研究人员名单

核心贡献者名单维持在31人，但与上次发布相比，名单构成略有调整。吕福利和许燕红从上一份名单中被移除：许燕红未出现在第3篇论文中，因此仅参与了七篇论文中的六篇；吕福利仅参与了五篇论文，未参与第6和第7篇，与其被列为第5篇离职者相吻合。

徐润欣和李玉坤替代了上述两人。两人均参与了全部七篇论文，但由于某些论文中以缩写记录其姓名（分别为R. X. Xu和Y. K. Li），我们未能将这些条目与其全名档案匹配，因此他们被排除在原来的"核心团队"计数之外。

1. 王冰轩 (Bingxuan Wang)	17. 朱启浩 (Qihao Zhu)
2. 赵成刚 (Chenggang Zhao)	18. 杜秋实 (Qiushi Du)
3. 邓成琦 (Chengqi Deng)	19. 徐润欣 (Runxin Xu)
4. 阮冲 (Chong Ruan)	20. 马诗荣 (Shirong Ma)
5. 戴大迈 (Damai Dai)	21. 梁文峰 (Wenfeng Liang)
6. 郭大雅 (Daya Guo)	22. 高文俊 (Wenjun Gao)
7. 杨德坚 (Dejian Yang)	23. 毕潇 (Xiao Bi)
8. 陈德理 (Deli Chen)	24. 谢鑫 (Xin Xie)
9. 徐汉威 (Hanwei Xu)	25. 王耀辉 (Yaohui Wang)
10. 高花佐 (Huazuo Gao)	26. 朴易施 (Yishi Piao)
11. 李嘉仕 (Jiashi Li)	27. 李玉坤 (Yukun Li)
12. 宋俊晓 (Junxiao Song)	28. 游宇翔 (Yuxiang You)
13. 董凯 (Kai Dong)	29. 谢振达 (Zhenda Xie)
14. 管康 (Kang Guan)	30. 郝哲文 (Zhewen Hao)
15. 张立悦 (Liyue Zhang)	31. 邵志宏 (Zhihong Shao)
16. 王沛怡 (Peiyi Wang)	

本文件为《Update: DeepSeek AI and the Great Talent Competition》白皮书的中文翻译版本，原文由Amy Zegart与Emerson Johnston撰写，发布于胡佛研究所技术政策加速器及斯坦福大学以人本AI研究院（2026年6月）。原文DOI: <https://doi.org/10.64576/0626>