arXiv × NeurIPS

# Global Research Institution Trends

*arXiv × NeurIPS Research Report*

Coverage: 2024–2025 (Year prefix: 24/25)

Data: arXiv full corpus / High-impact papers (over_70) / NeurIPS acceptances

Lens: Volume × Quality × Research Portfolio

Blog: https://staka.jp/wordpress/?p=1164

**3,000+**
Institutions

**580K+**
Author
Contributions

**800K+**
Category Records

# Framework: Reading Research Through Three Complementary Metrics

## ① arXiv Full Corpus

### Volume Indicator

- Author-contribution scores aggregated by institution
- sum(score): total contribution to papers
- count(arxiv_id): author contribution instances
- → Shows research breadth and output volume

Leading indicator: 6–18 months ahead of conferences

## ② arXiv over_70

### Quality Indicator

- High-attention papers only (score ≥ 70)
- Collection of papers by influential authors
- Used for category-level portfolio analysis
- → Shows research depth and impact

Primary metric for this analysis

## ③ NeurIPS Acceptances

### Lagging Indicator

- Post-peer-review visibility — lagging signal
- Appears 6–12 months after arXiv over_70
- Joint affiliations introduce counting errors
- → Supplements track record and reputation

Supplementary: RIKEN overridden with official count

📌 *Categories are multi-label — one paper can contribute to multiple categories. Category totals therefore exceed institutional totals.*

# Executive Summary

**01** **LLM/Foundation Models Are the Universal Core**
Every top institution centers on LLM/Foundation Models. Competitive differentiation emerges in Efficiency, RL/Agents, Multimodal, and GenAI/Video.

**02** **Clear Role Division Across Sectors**
Companies: LLM + efficiency. Universities: evaluation + theory. Public labs: open model hubs. Startups: sharp specialization.

**03** **NVIDIA Makes the Biggest Leap: +4,172 pts**
Top over_70 score gainer: NVIDIA (+4,172). Alibaba (+1,979), ByteDance (+1,710), and NUS (+1,517) also surged.

**04** **Startups Diverge Into Two Poles**
StepFun (generalist), DeepSeek (reasoning/RL specialist), HiDream (diffusion/media specialist) — startups are bifurcating around clear strategic bets.

**05** **Always Verify NeurIPS Counts Against Official Sources**
RIKEN: aggregation table shows 22 (2024); official count is 41. Joint-affiliation institutions can show 2× divergence — never treat these tables as ground truth.

# Macro Metrics: Volume Surges, High-Impact Output Dips (2024 vs 2025)

## 📄 arXiv Full Corpus Score
Volume metric: total author-contribution score

2024
**1,666,537**

2025
**1,910,892**

**+14.7%**

## ⭐ arXiv over_70 Score
Quality metric: total high-attention paper score

2024
**310,192**

2025
**300,806**

**−3.0%**

## 🏆 NeurIPS Accepted Papers
Lagging indicator: accepted paper count (official)

2024
**4,497**

2025
**5,290**

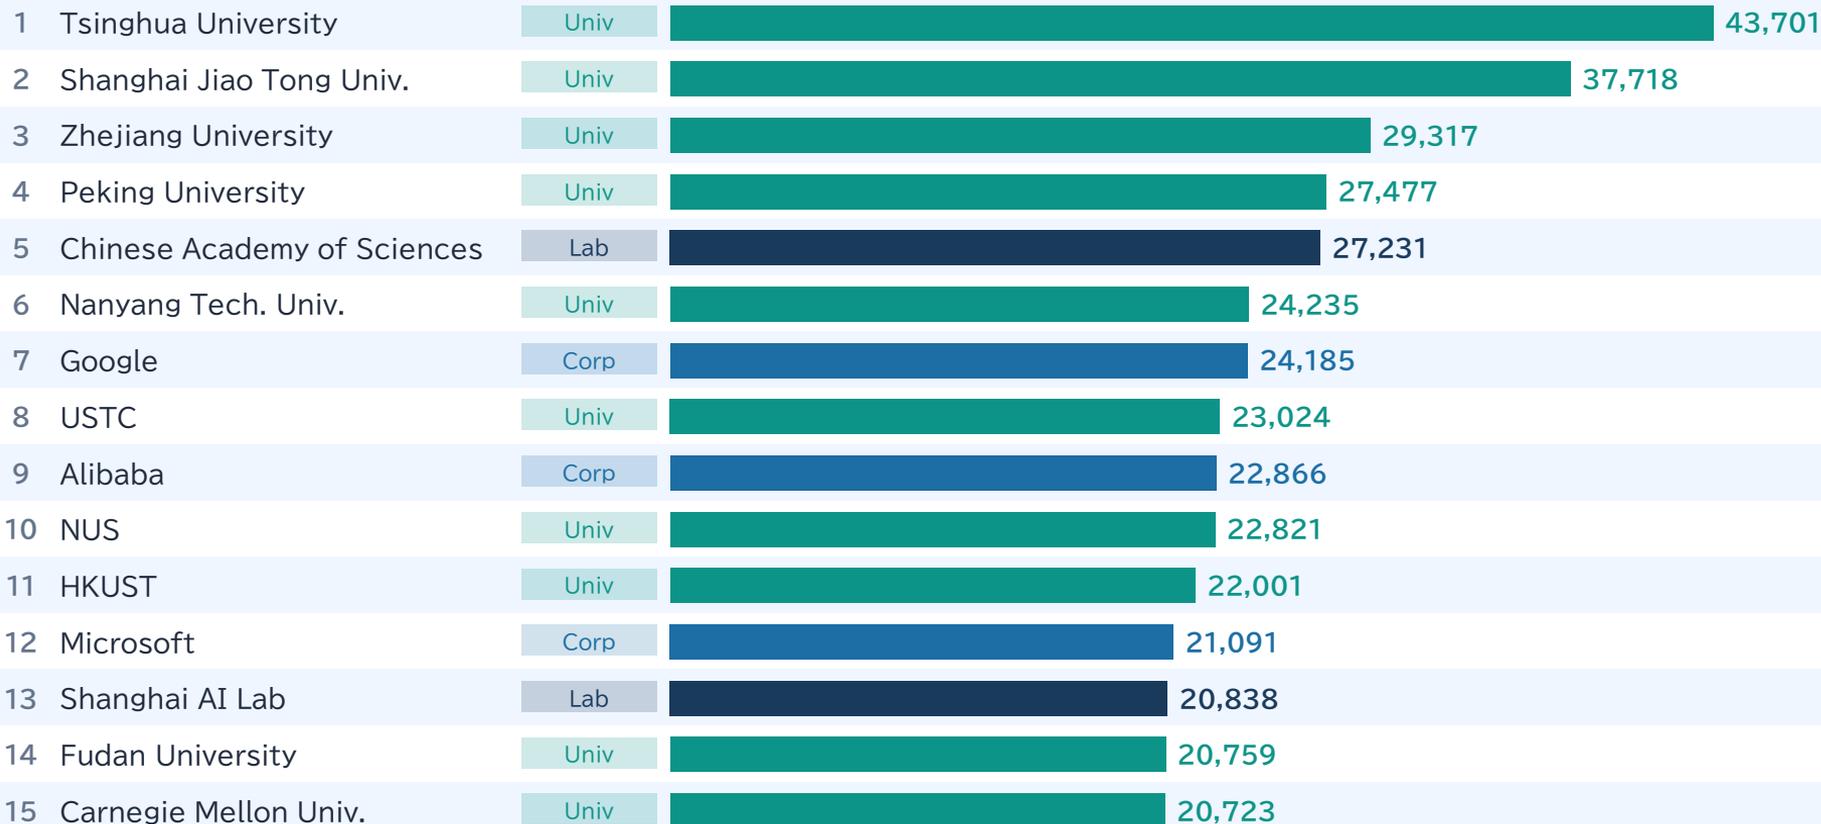**+17.6%**

**over_70 Share** | 2024: 18.6% | 2025: 15.7% (▼ -2.9pt)

### Key Insight: Volume Growth Is Outrunning Quality — The Dilution Effect
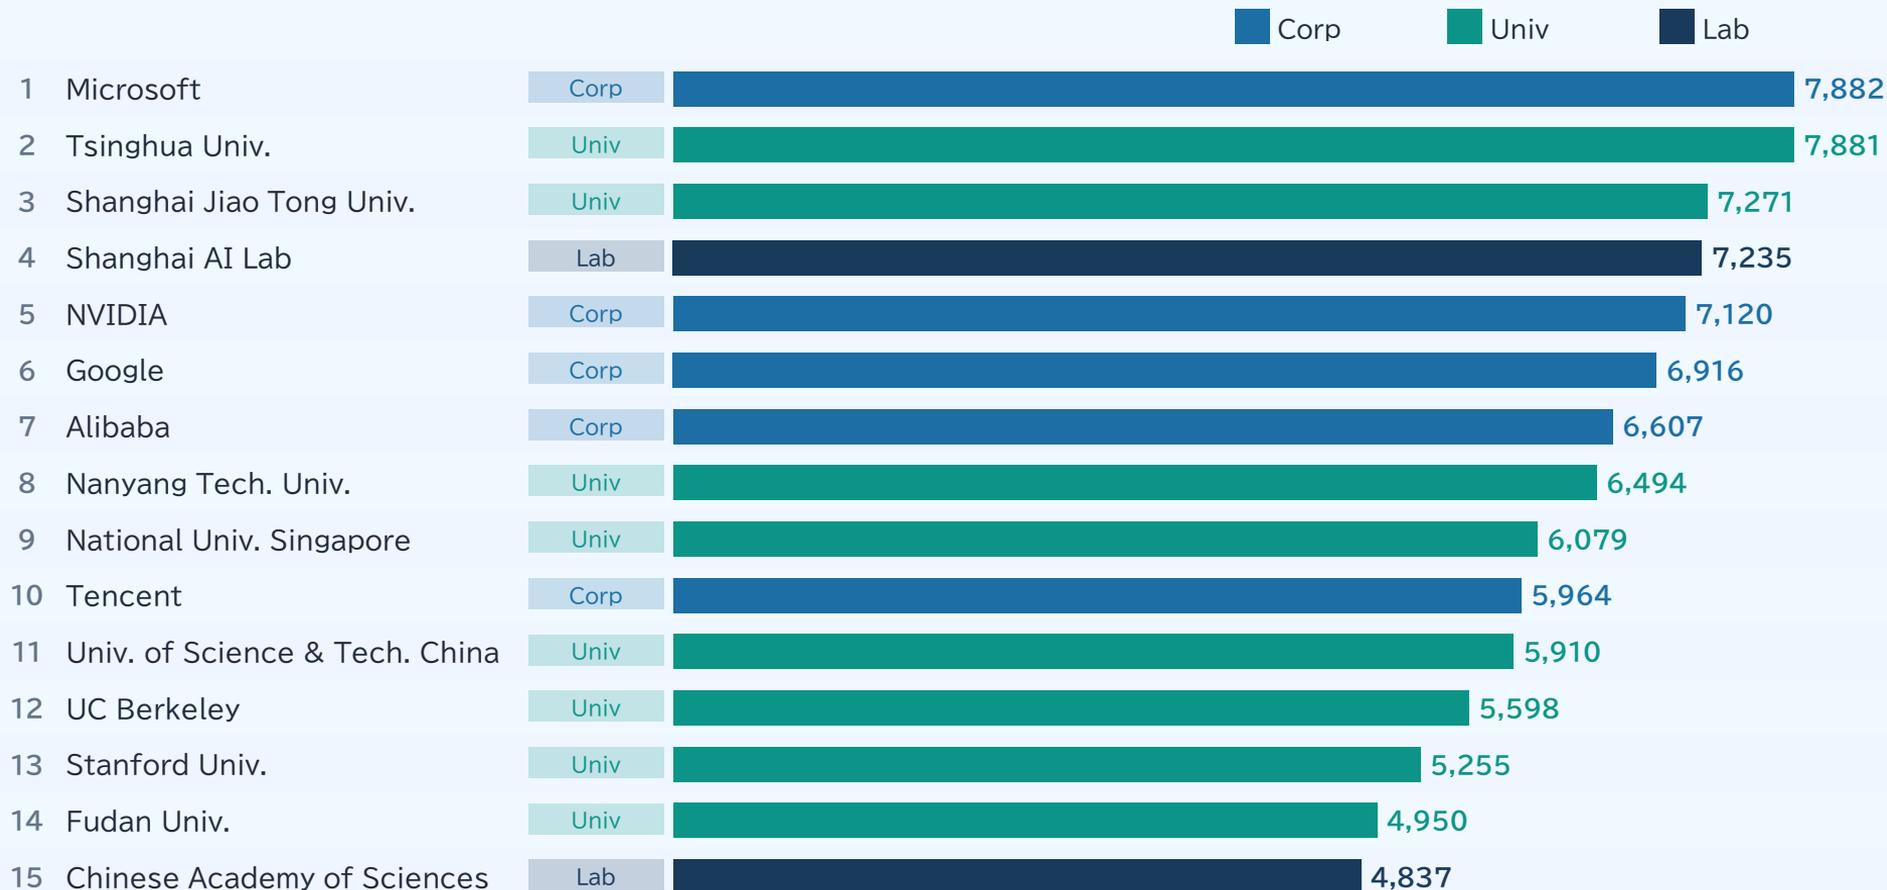Full-corpus output grew +14.7%, but high-attention paper score fell −3.0%. The over_70 share declined from 18.6% to 15.7% — a clear signal that volume growth is diluting the density of globally noticed work. NeurIPS grew +17.6%, reflecting the conference's expanding scope.

# 2025 arXiv Full-Corpus Score — Global Top 15

| | | | Corp | Univ | Lab |
|---|---|---|---|---|---|

| Rank | Institution | Type | Score |
|---|---|---|---|
| 1 | Tsinghua University | Univ | 43,701 |
| 2 | Shanghai Jiao Tong Univ. | Univ | 37,718 |
| 3 | Zhejiang University | Univ | 29,317 |
| 4 | Peking University | Univ | 27,477 |
| 5 | Chinese Academy of Sciences | Lab | 27,231 |
| 6 | Nanyang Tech. Univ. | Univ | 24,235 |
| 7 | Google | Corp | 24,185 |
| 8 | USTC | Univ | 23,024 |
| 9 | Alibaba | Corp | 22,866 |
| 10 | NUS | Univ | 22,821 |
| 11 | HKUST | Univ | 22,001 |
| 12 | Microsoft | Corp | 21,091 |
| 13 | Shanghai AI Lab | Lab | 20,838 |
| 14 | Fudan University | Univ | 20,759 |
| 15 | Carnegie Mellon Univ. | Univ | 20,723 |

# 2025 High-Impact Score (over_70) — Global Top 15

| Legend | | |
|---|---|---|
| Corp | Univ | Lab |

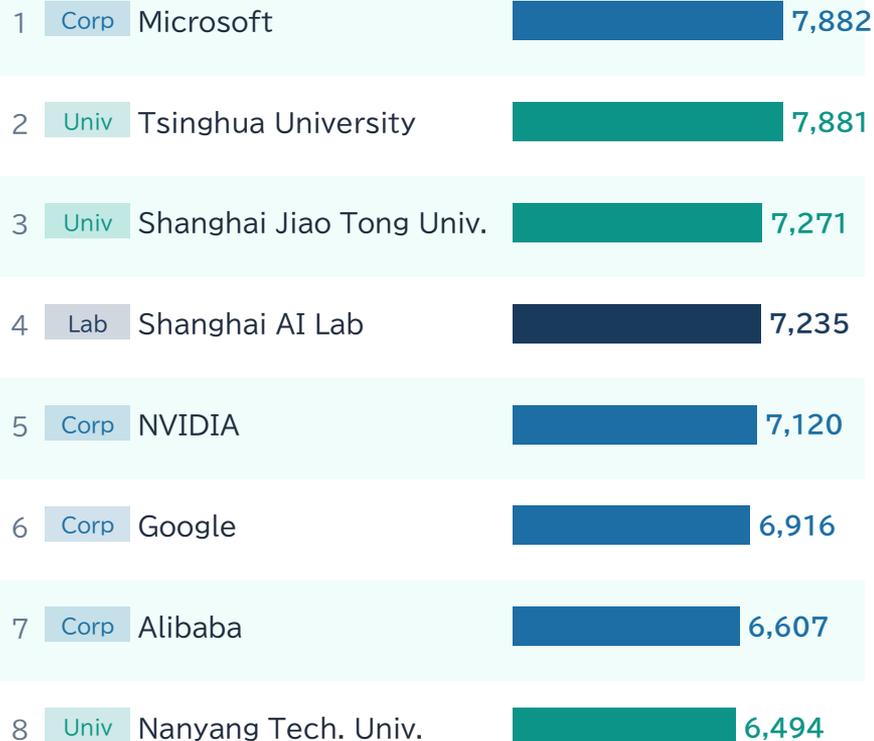| Rank | Institution | Type | Score |
|---|---|---|---|
| 1 | Microsoft | Corp | 7,882 |
| 2 | Tsinghua Univ. | Univ | 7,881 |
| 3 | Shanghai Jiao Tong Univ. | Univ | 7,271 |
| 4 | Shanghai AI Lab | Lab | 7,235 |
| 5 | NVIDIA | Corp | 7,120 |
| 6 | Google | Corp | 6,916 |
| 7 | Alibaba | Corp | 6,607 |
| 8 | Nanyang Tech. Univ. | Univ | 6,494 |
| 9 | National Univ. Singapore | Univ | 6,079 |
| 10 | Tencent | Corp | 5,964 |
| 11 | Univ. of Science & Tech. China | Univ | 5,910 |
| 12 | UC Berkeley | Univ | 5,598 |
| 13 | Stanford Univ. | Univ | 5,255 |
| 14 | Fudan Univ. | Univ | 4,950 |
| 15 | Chinese Academy of Sciences | Lab | 4,837 |

*over_70 score = sum of institutional contribution scores in high-attention papers (2025)

# Full Corpus vs over_70 Rankings: Universities Lead on Volume, Companies Reverse on Quality (2025)

## 📄 arXiv Full-Corpus Score Ranking (2025)

| # | Type | Institution | Score |
|---|------|-------------|-------|
| 1 | Univ | Tsinghua University | 43,701 |
| 2 | Univ | Shanghai Jiao Tong Univ. | 37,718 |
| 3 | Univ | Zhejiang University | 29,317 |
| 4 | Univ | Peking University | 27,477 |
| 5 | Lab | Chinese Academy of Sci. | 27,231 |
| 6 | Univ | Nanyang Tech. Univ. | 24,237 |
| 7 | Corp | Google | 24,185 |
| 8 | Univ | USTC | 23,024 |
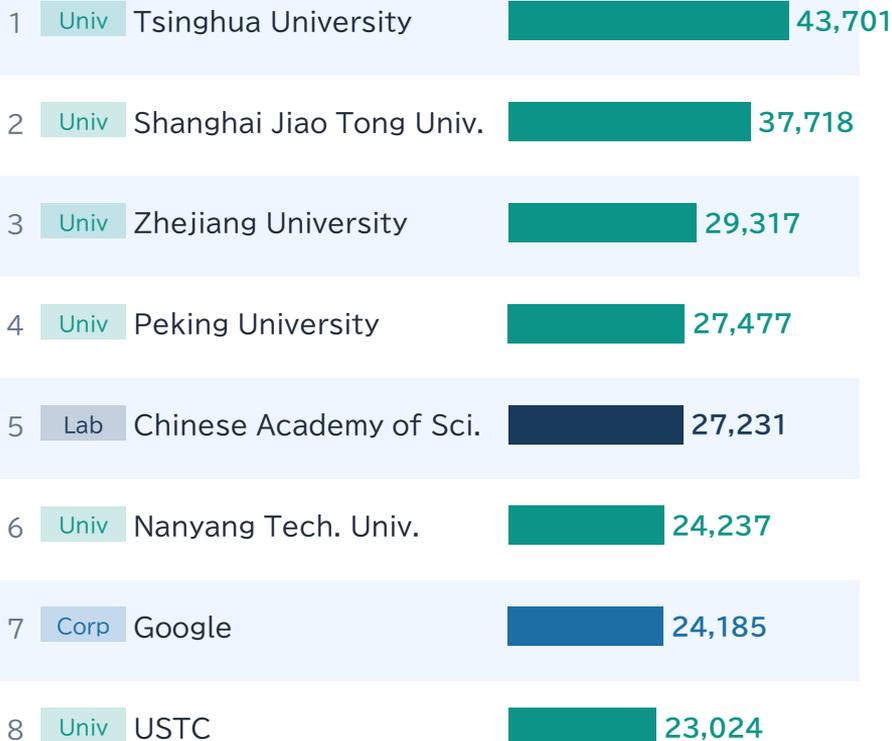
## ⭐ over_70 Score Ranking (2025)

| # | Type | Institution | Score |
|---|------|-------------|-------|
| 1 | Corp | Microsoft | 7,882 |
| 2 | Univ | Tsinghua University | 7,881 |
| 3 | Univ | Shanghai Jiao Tong Univ. | 7,271 |
| 4 | Lab | Shanghai AI Lab | 7,235 |
| 5 | Corp | NVIDIA | 7,120 |
| 6 | Corp | Google | 6,916 |
| 7 | Corp | Alibaba | 6,607 |
| 8 | Univ | Nanyang Tech. Univ. | 6,494 |

💡 Tsinghua (full-corpus #1) holds over_70 #2. But Zhejiang, Peking, and CAS (full-corpus #3–5) fall outside the over_70 top 10. Companies (Microsoft, NVIDIA, Google) move significantly higher when quality is the lens.

# 2024→2025 Movers: over_70 Score Change
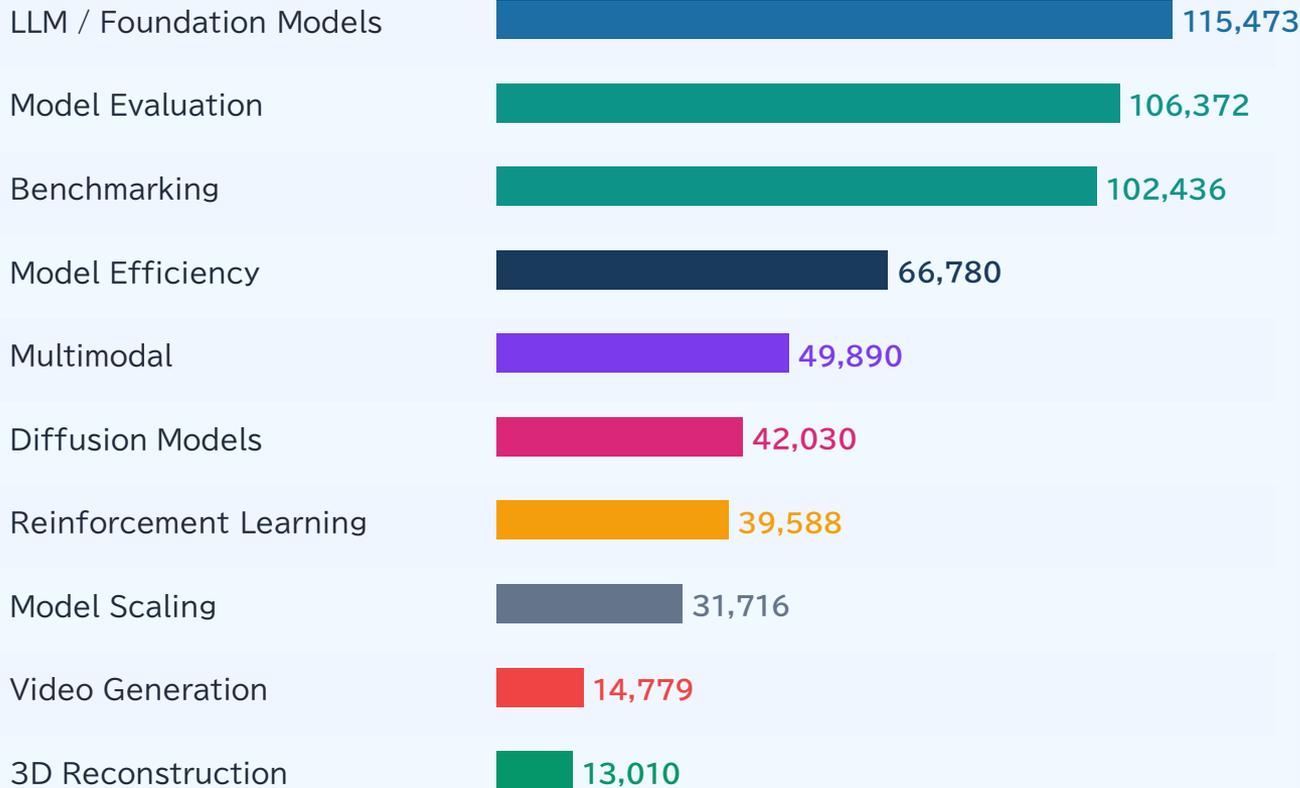
## Top 7 over_70 Score Gainers (2024→2025)

NVIDIA           2,949 →

Alibaba          4,628 →  +1,979  → 6,607

ByteDance        1,844 →  +1,710  → 3,555

NUS              4,562 →  +1,517  → 6,079

Stanford         4,458 →  +797  → 5,255

Harbin Inst. Tech.   2,990 →  +797  → 3,786

Zhejiang Univ.   3,421 →  +631  → 4,051

🚀 **Biggest Leap: NVIDIA**

## +4,172

2024: 2,949  →  2025: 7,120

Hybrid architecture, long-context
& inference efficiency drive the leap

# Research Theme Distribution (2025 over_70, by Category Score)

## Global Top 10 Categories (total score)

| Category | Score |
|---|---|
| LLM / Foundation Models | 115,473 |
| Model Evaluation | 106,372 |
| Benchmarking | 102,436 |
| Model Efficiency | 66,780 |
| Multimodal | 49,890 |
| Diffusion Models | 42,030 |
| Reinforcement Learning | 39,588 |
| Model Scaling | 31,716 |
| Video Generation | 14,779 |
| 3D Reconstruction | 13,010 |

*Multi-label: category totals exceed institutional totals

### Key Insights

- LLM / Eval / Benchmark dominate — the universal core of AI research
- Efficiency ranks 4th — growing demand for practical deployment
- Diffusion / Video Gen growing fast, led by startups
- RL ranks 7th but rising fast on reasoning / agentic wave
- 3D / Video: smaller volume but high score efficiency

# Research Portfolio Comparison — Companies & Public Labs (2025)

## Microsoft — 7,882pt

- LLM/Foundation 36%
- Eval/Bench 21%
- Efficiency 15%
- RL/Agents 9%

Phi-4-Mini / MoE / 1-bit LLM: practical model family triad

## Google — 6,916pt

- LLM/Foundation 30%
- RL/Agents 12%
- Multimodal 11%
- Eval/Bench 13%

Gemini 2.5 Pro / long-context / agentic multimodal

## NVIDIA — 7,120pt

- LLM/Foundation 34%
- Eval/Bench 15%
- Efficiency 14%
- Multimodal 8%

Mamba-Transformer / 1M context / inference efficiency

## Shanghai AI Lab — 7,235pt

- Eval/Bench 30%
- LLM/Foundation 24%
- RL/Agents 14%
- Multimodal 12%

Open model & benchmark hub: Eval/Bench 30% is striking

## Alibaba — 6,607pt

- LLM/Foundation 35%
- Eval/Bench 24%
- RL/Agents 18%
- Efficiency 10%

Qwen series: long-context reasoning / tool-use / RFT

## Tencent — 5,964pt

- LLM/Foundation 30%
- Eval/Bench 22%
- Efficiency 13%
- Multimodal 9%

Adaptive CoT / Hybrid Mamba-Transformer

# University Research Portfolio Comparison (2025 over_70 Categories)

## Tsinghua Univ. — 7,881pt

- LLM/Foundation 20%
- Eval/Bench 31%
- Efficiency 15%
- Multimodal 11%

Full-stack: LLM through safety and generative AI. Eval/Bench 31% is the largest theme.

## SJTU — 7,271pt

- LLM/Foundation 14%
- Eval/Bench 38%
- Efficiency 14%
- Multimodal 14%

Scientific AI and evaluation infra stand out. Eval 38%, GenAI/Video, Multimodal roughly equal.

## NUS — 6,079pt

- LLM/Foundation 20%
- Eval/Bench 36%
- Efficiency 13%
- Multimodal 12%

Multimodal reasoning and recommendation/agent applications. Eval-heavy like SJTU.

## UC Berkeley — 5,598pt

- LLM/Foundation 20%
- Eval/Bench 34%
- Efficiency 12%
- Multimodal 6%

RL/Agents 11% and Scaling 8% are high. Embodied AI, sim-to-real, parallel reasoning.

## Stanford — 5,255pt

- LLM/Foundation 19%
- Eval/Bench 38%
- Efficiency 11%
- Multimodal 5%

Evaluation design and Preference Optimization. Eval/Bench 38% rivals SJTU.

## Fudan Univ. — 4,950pt

- LLM/Foundation 17%
- Eval/Bench 44%
- Efficiency 7%
- Multimodal 14%

Eval/Bench 44% is the highest among universities. Vision/robotics/AI-generated image detection.

💡 Universal pattern: all universities make Eval/Bench their largest theme (31–44%). Differentiation lies in GenAI/Video, RL/Agents, and Scaling weights. Berkeley and Stanford lean into RL/Scaling; Fudan maximizes Eval specialization.

# Startup Research Portfolio Comparison (2025 over_70 Categories)

## StepFun — NEW 2025

### Generalist Frontier

- LLM/Foundation 15%
- Eval/Bench 35%
- Efficiency 21%
- GenAI/Video 19%

Simultaneous push across LLM, efficiency, and video/3D. Broadest startup portfolio.

## DeepSeek-AI — +143

### Reasoning / RL Specialist

- LLM/Foundation 26%
- Eval/Bench 37%
- RL/Agents 37%

Eval + RL = 74%. Hyper-focused on reasoning, attention, MoE. Only 3 themes.

## Kimi/Moonshot — NEW 2025

### Long-Context / Agentic

- LLM/Foundation 32%
- Eval/Bench 23%
- Efficiency 25%
- Multimodal 7%

MoE, long-context, agentic benchmarks. Efficiency 25% signals deployment focus.

## HiDream.ai — +135

### Diffusion / Media Specialist

- LLM/Foundation 7%
- Eval/Bench 17%
- Efficiency 7%
- GenAI/Video 63%

GenAI/Video 63% — most extreme specialization. Fully committed to image/video generation.

## Inclusion AI — NEW 2025

### Large-scale MoE / Systems

- LLM/Foundation 30%
- Eval/Bench 10%
- Efficiency 15%
- RL/Agents 20%

Scaling 25%, RL 20%. 1T-scale MoE and system co-design. Pre-NeurIPS.

## Zhipu AI — +34

### GLM / Agentic Eval

- LLM/Foundation 41%
- Eval/Bench 47%
- GenAI/Video 10%
- Multimodal 2%

LLM + Eval = 88% — a two-theme concentration. GLM-family agentic reasoning at the core.

💡 Bifurcation: generalist (StepFun) vs. single-theme specialists (HiDream=GenAI 63%, DeepSeek=RL 37%). Kimi and Inclusion AI tilt toward efficiency/scaling, signaling deployment focus. arXiv reveals these strategic bets months before NeurIPS.

# University Portfolios & Startup Bifurcation

## Top Universities: Research Profile (2025 over_70)

| Institution | Score | Primary Strength |
|---|---|---|
| Tsinghua | 7,881 | Full-stack: LLM through safety and generative AI |
| SJTU | 7,271 | Scientific AI, evaluation infra, video/multimodal |
| NUS | 6,079 | Multimodal reasoning, recommendation, agent applications |
| UC Berkeley | 5,598 | RL, embodied AI, sim-to-real transfer |
| Stanford | 5,255 | Evaluation design, Preference Optimization, agents |
| Fudan | 4,950 | Vision / robotics / AI-generated image detection |

## Startup Bifurcation (2025 over_70)

**StepFun**
Generalist frontier
Score: 1183 (NEW)

LLM 15% / GenAI/Video 19% / Efficiency 21%
→ LLM + efficiency + video/3D simultaneously

**DeepSeek-AI**
Reasoning / RL specialist
Score: 367 (+143)

LLM 26% / Eval 37% / RL 37%
→ Extreme focus: reasoning, attention, RL

**Kimi/Moonshot**
Long-context / agentic
Score: 449 (NEW)

LLM 32% / Efficiency 25% / Eval 23%
→ MoE, long-context, agentic benchmarks

**HiDream.ai**
Diffusion / media specialist
Score: 446 (+135)

GenAI 63% / Eval 17% / LLM 7%
→ Image editing & video generation specialist

*arXiv over_70 reveals this bifurcation months before any NeurIPS signal*

# Key Takeaways & Strategic Implications

### ① Move from Volume to Portfolio Analysis

Submission counts alone do not reveal strategic differentiation. The over_70 category distribution shows where institutions are actually making their bets. The shared core is LLM/Foundation; the competitive edge lives in Efficiency, RL/Agents, Multimodal, or GenAI.

### ② Each Sector Plays a Different Role

Companies (frontier/efficiency), universities (eval/theory/exploration), public labs (open model hubs), startups (sharp specialization) — ranking them on the same axis without role context produces misleading conclusions.

### ③ NVIDIA's Rise and Chinese Startup Emergence

NVIDIA leads 2024→2025 over_70 growth (+4,172). Alibaba, ByteDance, NUS also surged. Chinese startups split into reasoning specialists (DeepSeek) and diffusion specialists (HiDream), each commanding a distinct frontier.

### ④ Japan: Role Differentiation Is Emerging

Three profiles are emerging: RIKEN/UTokyo (theory/eval/RL), Sony (open generative publisher), Science Tokyo (conference-first). Developing the second tier of universities and startups is the key to deepening national capacity.

### ⑤ Always Cross-Check NeurIPS Counts

Joint-affiliation-heavy institutions (e.g., RIKEN) can show 2× divergence between aggregated tables and official counts. Always validate against official news, annual reports, or conference websites before using these as openness proxies.

# High-Attention Rate Ranking: over_70 Share of Total arXiv Score (2025)

High-attention rate = over_70 score ÷ full-corpus score × 100%. Shows what fraction of an institution's output is globally noticed. Filtered to count ≥ 100.

| | Institution | Hit Rate | | | Institution | Hit Rate |
|---|---|---|---|---|---|---|
| 1 | Corp Salesforce | 72.9% | | 10 | Univ UNC Chapel Hill | 29.9% |
| 2 | Inst Allen Inst. for AI | 56.1% | | 11 | Univ Univ. of Washington | 29.8% |
| 3 | Corp NVIDIA | 47.6% | | 12 | Corp Alibaba | 28.9% |
| 4 | Corp Snap | 44.9% | | 13 | Corp Google | 28.6% |
| 5 | Corp Microsoft | 37.4% | | 14 | Corp Adobe | 28.3% |
| 6 | Univ UC Berkeley | 37.1% | | 15 | Univ UCLA | 28.0% |
| 7 | Lab Shanghai AI Lab | 34.7% | | 16 | Univ Stanford | 27.8% |
| 8 | Univ Renmin Univ. | 34.3% | | 17 | Univ NTU | 26.8% |
| 9 | Corp Tencent | 33.8% | | | | |

💡 Salesforce leads at 72.9% — over 70% of their output is high-attention. NVIDIA (47.6%) and AI2 (56.1%) also show strong quality concentration. Large-volume institutions (Google, Tsinghua) cluster in the 20–29% range.