

COLM 2026

Conference on Language Modeling | San Francisco, October 6–9, 2026



Actions Speak Louder than Words

Measuring Cross-Lingual Policy Retention in Tool-Using Agents

Sourabrata Mukherjee

Kalika Bali

Sunayana Sitaram



Microsoft Research, India

8

models

6

benchmarks

41

languages

5,776

tasks

2.38M

rollouts



arXiv:2608.11110

THE PROBLEM • WHY IT MATTERS

For an agent, the route is the product



Cost

Extra steps change the price in each language.



Failure


Steps off the English path fail where tests never look.



Governance

Permissions and audits are written against actions, not answers.

THE BLIND SPOT



Our question: do agents run the same action policy in every language?

The obvious measurement fails for five reasons



No baseline

0.63–0.80

agreement between
two runs in
one language

✓ Run every
cell twice



Trace length

$R^2 = 0.67$

of raw agreement
explained by length

✓ Length-match
both ways



Empty traces

0.79 → 0.02

one model,
once blank pairs
are dropped

✓ Drop
empty pairs



Ceiling

$r = +0.97$

gap against
self-consistency,
 $T=0$

✓ Divide by
the ceiling



Chance floor

0.56


already scored
by traces of
different tasks

✓ Measure it
by permutation

On our own data, four of the five reverse a sign or a ranking.

Run every cell twice, then divide

TWO RUNS PER CELL



POLICY RETENTION

$$\tilde{I} = \frac{I_{\text{cross}}}{I_{\text{within}}}$$

The share of a model's own consistency that survives a change of language.

- ✓ Length-matched
- ✓ Empty pairs dropped
- ✓ Task-level bootstrap
- ✓ Chance floor measured


The gap is real, and sampling was hiding it

Every correction made it larger

length-matched gap, pooled over frontier models



Remove sampling: one line moves




At $T=0$ the gap is positive in all 24 cells, every interval above zero.

Frontier agents keep 71–73% of their action policy

71–73%

of a model's own
consistency survives a
change of language

four frontier models, greedy
decoding ($T=0$)




Four very different models land within 2.6 points of each other.

Which model it is barely matters


What explains the variation in $\tilde{I}$?

share of variance across 24 cells



Ranked by the raw gap


rank 1 = largest cross-lingual gap




A ranking on the raw gap measures determinism, not multilinguality.

Below roughly 10B, the constant breaks

Retention by model size



The chance floor flips it



Corrected, Qwen3-8B falls below the 4B model.

Below about 10B, only the chance-corrected scale can order models.

Agents plan in English, and will not stop



We observe it

#1

Translate is the most-used tool

≈99%

of reasoning is ASCII, even on हिन्दी, தமிழ், ଓଡ଼ିଆ tasks




We intervene

Remove  Translate

Length-matched, agreement **drops**

Mandate  Translate first


Gains track head-room, as predicted **in advance** 




We instruct

"Reason in the task's language."

Gemma-3-27B

 0.79%

Sarvam-M

 0.08%


<1% comply

The English pivot is load-bearing, and prompts cannot switch it off.

A regex, not the model, faked a multilingual failure


Parser expects Action: Translate(text=...)

GPT-OSS writes **"We will use Translate."** ✗ no action



English still wins on accuracy

English advantage, four gold benchmarks



The Indic-specialised Sarvam-M has the largest English edge.

Report parse-failure rates; above about 20%, leave the model unranked.

CONCLUSION

Answer agreement is not behavioural agreement

WHAT WE SHOWED

- 1 A ceiling-corrected estimand that prices five confounds
- 2 Frontier agents keep 71–73% of their action policy under greedy decoding
- 3 Below about 10B, the regularity breaks
- 4 An English pivot drives it, and prompts cannot remove it

WHEN YOU EVALUATE AN AGENT

- ✓ Log the **action trace**
- ✓ Run every condition **twice**
- ✓ Rank on **retention $\tilde{I}$**
- ✓ **Measure** the chance floor
- ✓ Report **parse failures** and T

“ An agent equally accurate in two languages is not thereby equally trustworthy in both. ”

Thank you

 Microsoft Research, India

Questions and discussion welcome

 connect.souro@gmail.com

 aka.ms/multilingual_agents_evaluation

 Code, prompts, both benchmark suites and all 2.38M traces



 Paper



 Website



 LinkedIn