

(Chinese) Media (in Africa)


Alberto Alesina Workshop in Political Economy and Culture

Emily Silcock, Nan Xiang

17th September 2026

Motivation

- In 2026, for the first time, Pew found that China is now seen more positively than the U.S. in most of 36 countries surveyed.
- Both an increase in favourability towards China and a decrease towards U.S.



“It is necessary to make friends, unite and win over the majority, and constantly expand the circle of friends [when it comes to] international public opinion” – Xi Jinping, 2021.

- China's global influence has traditionally relied on hard infrastructure investments.
- However, they have recently started expanding their soft influence.
- In 2021, Chinese President Xi Jinping called on state-controlled media to “accelerate the construction of a Chinese discourse and Chinese narrative system”

Research Questions

- In this paper, we focus on one specific part of China's soft power strategy: news media.
 - We focus on the idea that in order to exert influence, China needs people to be willing to expose themselves to Chinese content.
1. How does China balance exposure and slant?
 2. What happens to this equilibrium when competition from other providers of news decreases?
 3. How has this affected global public opinion over time?

- ***Media slant***: why media chooses particular slants and how audiences respond. (Mullainathan and Shleifer 2005; Gentzkow and Shapiro 2006; Gentzkow and Shapiro 2010; Durante and Knight 2012; Martin and Yurukoglu 2017; Alonso and Padró i Miquel 2025)
- ***Media competition and news supply***: how news competition affects news content. (Anderson and Coate 2005; Anderson and McLaren 2012; Drago, Nannicini, and Sobbrío 2014; Gentzkow, Shapiro, and Sinkinson 2014; Cagé 2020; Perego and Yuksel 2022; Angelucci, Cagé, and Sinkinson 2024)
- ***Foreign media exposure and persuasion***: effects of international media on opinions. (DellaVigna, Enikolopov, et al. 2014; Peisakhin and Rozenas 2018; Mattingly et al. 2025)
- ***State-controlled media and propaganda***: how governments design media to influence domestic audiences. (Besley and Prat 2006; Edmond 2013; Gehlbach and Sonin 2014; Yanagizawa-Drott 2014; Qin, Strömberg, and Wu 2018; Y. Chen and Yang 2019)
- ***Strategic manipulation of international coverage***: how political actors influence their representation in foreign media. (Qian and Yanagizawa-Drott 2017; Durante and Zhuravskaya 2018)

- Most media outlets are reliant on news wires for international news.
- Most multi-home - they face a menu of content from multiple newswires, and choose which articles to republish.
- Local outlets often use Xinhua news for free or close to free.
- China calls this strategy “*borrowing boats to reach the sea*”: by placing Chinese newswire copy inside trusted local outlets, the message reaches a much wider - and less self-selected - audience.
- Content can be laundered through several hands: sold or given to a national news agency, which is then reprinted by local outlets.

Project Outline

1. Descriptive evidence of China's growing presence in African news markets.
2. Local news decisions (i.e. Xinhua's constraints):
 - Xinhua newswires with explicitly pro-China sentiment are less likely to be republished by local media outlets
 - Xinhua articles are more likely to be republished if they are the only newswire to cover that event
3. Xinhua's strategy:
 - When Xinhua faces lower competition from other newswires, it will produce stories with more positive sentiment towards China
 - Xinhua is more likely to cover an event/topic/location that no other newswire will cover, all else equal
4. Future work:
 - Theoretical additions
 - Downstream effects

Data

Data: Local news websites

- BBC has a media guide for every country, with a list of news outlets.
- We take all written media outlets (e.g. newspapers, news sites, news agencies).
- Scrape websites of these outlets for complete archives.
 - Africa: 178/217 sites, 5.5 million/15.5 million articles banked.
 - MENA: 212 sites, 142.7 million articles.
 - Latin America: 135 sites, 55.1 million articles.
 - Asia outlets: 109 sites.
- Many sites have per-story engagement (e.g. views, likes, comments).

Eswatini media guide

© 12 April 2023



This page is no longer being updated. It was last updated on 12 April 2023



The government exercises total control over the broadcast media, including the only privately owned TV channel, which belongs to the royal family. Almost all media outlets are controlled, directly or indirectly by the monarch.

Journalists are liable to be prosecuted for criticising the government.

Media from South Africa are available.

Journalists cannot work freely and are liable to be prosecuted for any criticism of the government, says Reporters Without Borders. Self-censorship is routine.

There were 665,000 internet users by December 2021, comprising 56% of the population (Internetworldstats.com). The government does not restrict access.

• [Media guides by BBC Monitoring](#)

Press

- [The Times of Swaziland](#) - privately-owned daily
- [The Swazi Observer](#) - pro-government daily

Television

- [Eswatini TV](#) - state-run

Radio

- [Eswatini Broadcasting and Information Services](#) - state-run, operates networks in Swazi and English
- [Trans World Radio-Voice of the Church](#) - Christian network

Data: Newswire articles

- All Xinhua articles 2017–2026, scraped directly from Xinhua's website: 2.9 million articles.
 - 9 languages: Arabic, English, French, German, Japanese, Korean, Portuguese, Russian, Spanish.
- In the process of getting Reuters, AP and AFP.



Data processing: Locations and wire sources

- For each article we extract the wire source and bureau from the dateline.
- (Also extract byline, but not using this today.)

Tanzania's first fishing harbor to be completed in 2025: minister

The New Times (Rwanda) | newtimes.co.rw | 4 April 2024


DAR ES SALAAM, April 3 (**Xinhua**) – The construction of Tanzania's first-ever fishing harbor will be completed in 2025, boosting the country's fish catches in deep seas, a senior official said Wednesday...

 bureau DAR ES SALAAM → GeoNames 160263  source Xinhua → xinhua

- Disambiguate wire sources to the BBC media guide list of outlets.
- Disambiguate bureaus to GeoNames locations → list of operating bureaus for each newswire at each time.

Data processing: Locations and wire sources

- 4.4% local African articles credit a global newswire.
- Since 2021, Xinhua has been the second largest wire, after AFP.



Data processing: Matching local article to source

- Want to know which Xinhua article is being reprinted.
- Additionally, many reprinted articles are not credited, or credited to an intermediate news agency.
- So match on the text instead: fine-tune a multilingual bi-encoder to decide whether a local article was copied from a particular wire article.
- Train on credited reprints, with same-event articles credited to AFP, Reuters or AP as hard negatives.
- In production: retrieve within a $-7/+1$ day window.

Validation: reprint matching


Data processing: Xinhua content

- **Countries mentioned.** Validation: countries mentioned
 - NER extracts every country named.
 - Use a local LLM to determine whether it is a substantive subject of the article or only incidental.
- **Frames.** Xinhua copy is mostly neutral in tone, so we measure the *role* a country is given rather than its sentiment. Validation: frames
 - A local LLM turns each (article, country) pair into one anonymised sentence, e.g. “The country is a provider of aid”.
 - Those sentences are clustered into a frame taxonomy by agentic clustering (Löwe and Silcock 2026): specialised agents dispatched by an orchestrator LLM, rather than a fixed pipeline. Pipeline Benchmarks
 - Every pair is then classified against the frozen 69-frame taxonomy.

Data processing: Xinhua content

- Most Xinhua articles from African bureaus do not even mention China.
- Share mentioning China has doubled since 2017.

By dateline



Data processing: Xinhua content

- Reporting surface for economic, market and social indicators 9.2%
- Directs its own society through law, policy and administration 5.5%
- Directs its own economy through strategy and reform 4.1%
- Reporting surface for epidemic case and death counts 3.7%
- Meets a counterpart with no direction of travel marked 3.6%

Data processing: Xinhua content

- Reporting surface for economic, market and social indicators 9.2%
 - “Chicago agricultural commodities post losses over the week” — United States, 17 Dec 2017
 - “Swedish property market troubled by sharp drop in sales: mortgage lender” — Sweden, 2 Sep 2023
- Directs its own society through law, policy and administration 5.5%
- Directs its own economy through strategy and reform 4.1%
- Reporting surface for epidemic case and death counts 3.7%
- Meets a counterpart with no direction of travel marked 3.6%

Data processing: Xinhua content

- Reporting surface for economic, market and social indicators 9.2%
- Directs its own society through law, policy and administration 5.5%
 - “Sri Lanka warns fishing, naval communities of high-speed wind” — Sri Lanka, 3 Sep 2023
 - “Albanian gov’t pledges compensation to flood-hit citizens” — Albania, 3 Dec 2017
- Directs its own economy through strategy and reform 4.1%
- Reporting surface for epidemic case and death counts 3.7%
- Meets a counterpart with no direction of travel marked 3.6%

Data processing: Xinhua content

- Reporting surface for economic, market and social indicators 9.2%
- Directs its own society through law, policy and administration 5.5%
- **Directs its own economy through strategy and reform** 4.1%
 - “Economic Watch: Service sector vouchers boost consumption vitality, daily joy” — China, 25 Aug 2025
 - “Biden’s measures to relieve bottlenecks in supply chain to fail – experts” — United States, 19 Oct 2021
- Reporting surface for epidemic case and death counts 3.7%
- Meets a counterpart with no direction of travel marked 3.6%

Data processing: Xinhua content


- Reporting surface for economic, market and social indicators 9.2%
- Directs its own society through law, policy and administration 5.5%
- Directs its own economy through strategy and reform 4.1%
- Reporting surface for epidemic case and death counts 3.7%
 - “China’s Xinjiang reports 26 asymptomatic COVID-19 cases” — China, 15 Oct 2021
 - “UK records another 34,574 coronavirus cases” — United Kingdom, 11 Oct 2021
- Meets a counterpart with no direction of travel marked 3.6%

Data processing: Xinhua content

- Reporting surface for economic, market and social indicators 9.2%
- Directs its own society through law, policy and administration 5.5%
- Directs its own economy through strategy and reform 4.1%
- Reporting surface for epidemic case and death counts 3.7%
- Meets a counterpart with no direction of travel marked 3.6%
 - “Trump says he won’t “negotiate for Ukraine” ahead of Alaska summit with Putin” — United States, 15 Aug 2025
 - “China’s top legislator meets vice chairman of National Assembly of Vietnam” — China, 20 Aug 2025


Data processing: Xinhua content

- 38,571 articles with substantive China frames, among African datelines.
- Heavily focused on provision of aid.



Data processing: Xinhua content

- 8,901 articles with substantive United States frames, among datelines.
- Focuses on accused wrongdoing and conflict.



Data: Wire journalists (LinkedIn)


- PDL data on LinkedIn career histories, scraped 2026.
- **69,004** job stints at AP, AFP, Reuters and Xinhua.
- **49.7%** editorial stints resolve to both a country and a year-month, giving a monthly head-count of each wire's journalists by country.

Reprinting of articles

- Xinhua can publish what they like, but no one will see it if it does not get reprinted.
- Local news editors choose which newswire articles to reprint.
- i.e. provide the constraints under which Xinhua must optimise.


Articles about China are less likely to be reprinted

- Not necessarily causal - China-mentioning articles may just be less interesting.




Most reprinted frames

- Frames on conflict, events, and natural disasters are most likely to be reprinted.



Most reprinted China frames

- For those frames on China, unlikely to be about conflict or natural disasters.
- Leaves frames on events as the most likely to be reprinted.



**What happens when Xinhua
faces lower competition?**


Detecting competitor drawdowns

- Want to detect breaks in the number of journalists at AP, AFP and Reuters bureaus.
- Use BIC (Bai–Perron) to detect breaks, ruling out surges that unwind (e.g. summer interns, extra journalists for the Olympics).
- Gives an ambiguous band around when the break happened.

Xinhua's strategy: competitor staffing

- AP, AFP and Reuters journalists in the country, in the months around a detected break.
- Baseline:** 10.84 journalists per country-month.

Table



Xinhua's strategy: event study specification

Stacked DiD, estimated with PPML. For country i , event stack e (a treated event plus its time matched controls) and calendar month t , with relative time $\tau = t - g_e$ estimate:


$$\mathbb{E}[Y_{iet}] = \exp\left[\alpha_{ie} + \lambda_{et} + \sum_{\tau} \beta_{\tau} T_{ie} \mathbf{1}\{t - g_e = \tau\}\right]$$

- α_{ie} stack $\times$ country and λ_{et} stack $\times$ month fixed effects, so each β_{τ} is a difference-in-differences against control countries living through the same calendar months.
- Donut: The event date is estimated, with treatment some time between $\tau = -2$ and $\tau = 0$.
- Africa, any-size drawdowns: 39 events, 34 stacks (as sometimes two wires cut at same time), ≈ 20 controls per stack.

Xinhua's strategy: total Xinhua output

- Total Xinhua output decreases in the months after a competitor drawdown.
- **Baseline:** 52 articles per country-month.


Table



Xinhua's strategy: articles mentioning China


- However, the number of articles that Xinhua produces mentioning China does not change.
- **Baseline:** 5 articles per country-month.

Table



Xinhua's strategy: which frames move after a drawdown


- The number of articles focusing on local economic conditions and policy frames decreases (mechanical effect of reduced non-China articles).
- The number of articles focusing on Belt-and-Road type articles increases (not mechanical).



Xinhua's strategy: Xinhua's own staffing

- These effects are not driven by Xinhua also reducing their editorial staff.


Table



Reprints

- Number of reprints of Xinhua articles appear unchanged (although noisy).
- **Baseline:** 2.17 credited reprints per country-month.


Table



Reprints

- Reprints of articles about China also unchanged (although very noisy).

Table



Next steps

Next steps

1. 8–9 months of scraping left to go.
2. Theory side: a lot is moving at once here, and would be good to put some structure on it.
3. Downstream outcomes: possibility for field work?

Conclusion

Thanks for your time!

References

- Alonso, Ricardo and Gerard Padró i Miquel (2025). **“Competitive Capture of Public Opinion”**. In: *Econometrica* 93.4, pp. 1265–1297. DOI: 10.3982/ECTA22072.
- Anderson, Simon P. and Stephen Coate (2005). **“Market Provision of Broadcasting: A Welfare Analysis”**. In: *The Review of Economic Studies* 72.4, pp. 947–972. DOI: 10.1111/0034-6527.00375.
- Anderson, Simon P. and John McLaren (2012). **“Media Mergers and Media Bias with Rational Consumers”**. In: *Journal of the European Economic Association* 10.4, pp. 831–859. DOI: 10.1111/j.1542-4774.2012.01069.x.

- Angelucci, Charles, Julia Cagé, and Michael Sinkinson (2024). **“Media Competition and News Diets”**. In: *American Economic Journal: Microeconomics* 16.2, pp. 62–102. DOI: 10.1257/mic.20220163.
- Besley, Timothy and Andrea Prat (2006). **“Handcuffs for the Grabbing Hand? Media Capture and Government Accountability”**. In: *American Economic Review* 96.3, pp. 720–736. DOI: 10.1257/aer.96.3.720.
- Cagé, Julia (2020). **“Media Competition, Information Provision and Political Participation: Evidence from French Local Newspapers and Elections, 1944–2014”**. In: *Journal of Public Economics* 185, p. 104077. DOI: 10.1016/j.jpubeco.2019.104077.
- Chen, Jianlyu et al. (2024). **“M3-Embedding: Multi-Linguality, Multi-Functionality, Multi-Granularity Text Embeddings Through Self-Knowledge Distillation”**. In: *Findings of the Association for Computational Linguistics: ACL 2024*, pp. 2318–2335.

- Chen, Yuyu and David Y. Yang (2019). **“The Impact of Media Censorship: 1984 or Brave New World?”** In: *American Economic Review* 109.6, pp. 2294–2332. DOI: 10.1257/aer.20171765.
- Conneau, Alexis et al. (2020). **“Unsupervised Cross-lingual Representation Learning at Scale”**. In: *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 8440–8451.
- DellaVigna, Stefano, Ruben Enikolopov, et al. (2014). **“Cross-Border Media and Nationalism: Evidence from Serbian Radio in Croatia”**. In: *American Economic Journal: Applied Economics* 6.3, pp. 103–132. DOI: 10.1257/app.6.3.103.
- DellaVigna, Stefano, John A. List, and Ulrike Malmendier (Feb. 2012). **“Testing for Altruism and Social Pressure in Charitable Giving”**. In: *The Quarterly Journal of Economics* 127.1, pp. 1–56. DOI: 10.1093/qje/qjr050. URL: <https://doi.org/10.1093/qje/qjr050>.

- Drago, Francesco, Tommaso Nannicini, and Francesco Sobbrío (2014). **“Meet the Press: How Voters and Politicians Respond to Newspaper Entry and Exit”**. In: *American Economic Journal: Applied Economics* 6.3, pp. 159–188. DOI: 10.1257/app.6.3.159.
- Durante, Ruben and Brian Knight (2012). **“Partisan Control, Media Bias, and Viewer Responses: Evidence from Berlusconi’s Italy”**. In: *Journal of the European Economic Association* 10.3, pp. 451–481. DOI: 10.1111/j.1542-4774.2011.01060.x.
- Durante, Ruben and Ekaterina Zhuravskaya (2018). **“Attack When the World Is Not Watching? U.S. News and the Israeli-Palestinian Conflict”**. In: *Journal of Political Economy* 126.3, pp. 1085–1133. DOI: 10.1086/697202.
- Edmond, Chris (2013). **“Information Manipulation, Coordination, and Regime Change”**. In: *The Review of Economic Studies* 80.4, pp. 1422–1458. DOI: 10.1093/restud/rdt020.

- Gehlbach, Scott and Konstantin Sonin (2014). **“Government Control of the Media”**. In: *Journal of Public Economics* 118, pp. 163–171. DOI: 10.1016/j.jpubeco.2014.06.004.
- Gentzkow, Matthew and Jesse M. Shapiro (2006). **“Media Bias and Reputation”**. In: *Journal of Political Economy* 114.2, pp. 280–316. DOI: 10.1086/499414.
- (2010). **“What Drives Media Slant? Evidence from U.S. Daily Newspapers”**. In: *Econometrica* 78.1, pp. 35–71. DOI: 10.3982/ECTA7195.
- Gentzkow, Matthew, Jesse M. Shapiro, and Michael Sinkinson (2014). **“Competition and Ideological Diversity: Historical Evidence from US Newspapers”**. In: *American Economic Review* 104.10, pp. 3073–3114. DOI: 10.1257/aer.104.10.3073.

- Löwe, Simon and Emily Silcock (2026). **“Agentic Clustering: Controllable Text Taxonomies via Multi-Agent Refinement”**. In: *arXiv preprint arXiv:2606.01255*.
- Martin, Gregory J. and Ali Yurukoglu (2017). **“Bias in Cable News: Persuasion and Polarization”**. In: *American Economic Review* 107.9, pp. 2565–2599. DOI: 10.1257/aer.20160812.
- Mattingly, Daniel C. et al. (2025). **“Chinese State Media Persuades a Global Audience That the “China Model” Is Superior: Evidence from a 19-Country Experiment”**. In: *American Journal of Political Science* 69.3, pp. 1029–1046. DOI: 10.1111/ajps.12887.
- Mullainathan, Sendhil and Andrei Shleifer (2005). **“The Market for News”**. In: *American Economic Review* 95.4, pp. 1031–1053. DOI: 10.1257/0002828054825619.

- Peisakhin, Leonid and Arturas Rozenas (2018). **“Electoral Effects of Biased Media: Russian Television in Ukraine”**. In: *American Journal of Political Science* 62.3, pp. 535–550. DOI: 10.1111/ajps.12355.
- Perego, Jacopo and Sevgi Yuksel (2022). **“Media Competition and Social Disagreement”**. In: *Econometrica* 90.1, pp. 223–265. DOI: 10.3982/ECTA16417.
- Qian, Nancy and David Yanagizawa-Drott (2017). **“Government Distortion in Independently Owned Media: Evidence from U.S. Cold War News Coverage of Human Rights”**. In: *Journal of the European Economic Association* 15.2, pp. 463–499. DOI: 10.1093/jeea/jvx007.
- Qin, Bei, David Strömberg, and Yanhui Wu (2018). **“Media Bias in China”**. In: *American Economic Review* 108.9, pp. 2442–2476. DOI: 10.1257/aer.20170947.

Yanagizawa-Drott, David (2014). **“Propaganda and Conflict: Evidence from the Rwandan Genocide”**. In: *The Quarterly Journal of Economics* 129.4, pp. 1947–1994. DOI: 10.1093/qje/qju020.

Validation: dateline NER

◀ Back

A fine-tuned `xlm-roberta-large` (550M) tags `LOC`, `SOURCE` and `AUTHOR` spans. We evaluate on the extracted *strings* rather than the token tags, since the strings are what the rest of the pipeline consumes.

- 1,347 train / 76 validation / 76 test articles, stratified by corpus; validation and test drawn only from the hand-labelled sample.
- Base model, hyperparameters and input format chosen by grid search on the validation set; the 76 test articles were never seen.
- Training is not exactly reproducible even at a fixed seed, so performance is the mean over ten runs of the same configuration.

Held-out test set (76 articles)	F1
Micro-averaged	0.956 (s.d. 0.018)
LOC	0.966
SOURCE	0.956
AUTHOR	0.949

Validation: bureau geocoding

◀ Back

Each LOC span is resolved to an exact GeoNames identifier, which supplies the coordinates, country and administrative division behind the bureau panels.

- Rule-based retrieval builds the candidate list: exact match on names and alternate names, restricted by dateline country where one is given; otherwise an approximate match, up to 50 candidates. The list contains the correct place for **95.8%** of test spans that have one.
- A fine-tuned cross-encoder (xlm-roberta-base) reranks, trained on 11,600 span–candidate pairs from 1,094 spans labelled by Claude Opus; settings chosen on a separate 430-span validation set.
- Test set: 101 dateline spans hand-labelled to an exact GeoNames id — 34 Reuters, 38 African local, 29 Xinhua — stratified by span, so each appears once.

Test set (101 spans)	Precision	Recall	F1
Full pipeline	90.8%	90.8%	90.8%

Remaining errors are mostly granularity rather than misidentification — e.g. returning a capital city instead of the island it sits on.

Validation: reprint matching

◀ Back

- **Model.** bge-m3 (568M) (J. Chen et al. 2024), fine-tuned as a bi-encoder: embed the local body and the Xinhua body, both credit-stripped and truncated to 512 tokens, and match above 0.88 cosine similarity.
- **Labels.** 2,631 local articles crediting Xinhua, each paired with up to two Xinhua articles from the $-7/ + 1$ window by five-gram containment — usually the true parent and a hard near-miss. All 4,480 pairs labelled by Claude Sonnet 5 against a written rubric, shown only the two bodies and the two dates; Claude Opus relabelled 60 blind and agreed on 58.
- **Splits** are disjoint by outlet as well as by wire story — 5,047 train / 647 validation / 786 test pairs, with five outlets held back entirely. Base model and hyperparameters chosen over 182 configurations on the validation split.
- **Scoring** reproduces the production task: match each article against every Xinhua article in the window, take the best candidate if it clears 0.88, no match otherwise.

Held-out test set (786 pairs)	Correct source	Precision	Recall	F1
Full pipeline	99.0%	0.997	0.990	0.994

Every labelled article credited Xinhua, so this is conditional on the article deriving from a wire at all — not an estimate of production precision.

Validation: reprint matching, same-event negatives

◀ Back

The test set above cannot see the failure that matters: it holds almost no articles from the general population.

- Run over the full African corpus, a first version matched *events* rather than derivation. Its uncredited matches had median five-gram containment 0.039 — 56% shared essentially no wording with the article they matched — against 0.978 for its credited matches.
- The cause was a train/inference mismatch: every training candidate had been mined by containment, so the model was never shown a zero-overlap same-event pair and never learned to reject one.
- We added 2,000 hard negatives — articles carrying an AFP, Reuters or AP credit that v1 nonetheless matched to Xinhua, and so same-event coverage by construction.


10,730 held-out same-event negatives	v1	v2
Rejected at threshold 0.88	0.0%	99.72%
Threshold needed for $\geq 99.5\%$ rejection	0.98	0.87
Validation recall at that threshold	0.890	0.931

v2 is the shipped model: the same control on the failure class, at higher recall and a far less brittle operating point. Validation recall at 0.88 falls from 0.965 to 0.928.

Xinhua content: China mentions by dateline

◀ Back

- Share of Xinhua articles mentioning China by dateline.



Validation: countries mentioned

◀ Back

Each article–country pair is sorted into a substantive or an incidental class by `gemma-4-26B-A4B-it`. Only the substantive class is used downstream: incidental pairs are recorded and then discarded, and no frame is extracted for them, so performance is reported on the substantive class alone.

- 400 articles drawn at random from the corpus, stratified across it, and hand-labelled.
- The dateline is not discarded: the filing location is supplied back to the model as context.

Held-out test (252 articles, 352 labels)	Precision	Recall	F1
Substantive class	94.8%	94.0%	94.4%

Validation: country frames

◀ Back


Frames are induced from the anonymised descriptions rather than imposed, using the agentic clustering method of Löwe and Silcock (2026): an orchestrator inspects the state of the discovery process at each step and dispatches a proposer, synthesizer, auditor, investigator or critic accordingly.

- Target range of 30–70 clusters. Induction runs on a subsample capped at 200 pairs per country, so the taxonomy is not driven by the few countries the wire covers most.
- It settled at **69 frames**, plus a reserved label for text fitting none of them.
- On a held-out audit sample, **99.3%** of descriptions were assignable.
- Assignment is by `gemma-4-26B-A4B-it` at the same revision and decoding, reading only the anonymised description, with the answer constrained to the set of frame identifiers.

[to complete] — per-frame assignment accuracy against human labels is not yet measured.

Agentic clustering: pipeline

◀ Back



Agentic clustering: benchmarks

◀ Back

- State of the art on 7 clustering benchmarks (Löwe and Silcock 2026)

Method	Banking77				CLINC150				MASSIVE-Intent				MASSIVE-Domain			
	$\bar{k}$	ARI	NMI	ACC	$\bar{k}$	ARI	NMI	ACC	$\bar{k}$	ARI	NMI	ACC	$\bar{k}$	ARI	NMI	ACC
<i>Given k</i>																
LDA	77	0.20	0.55	0.30	150	0.06	0.56	0.30	60	0.17	0.44	0.29	18	0.17	0.26	0.31
SBERT+ k -means	77	0.55	0.82	0.65	150	0.21	0.83	0.66	60	0.39	0.71	0.52	18	0.42	0.63	0.56
BERTopic	77	0.56	0.83	0.64	150	0.26	0.82	0.65	60	0.45	0.73	0.58	18	0.48	0.64	0.61
LLM-embedding+ k -means	77	0.62	0.85	0.72	150	<u>0.28</u>	0.87	0.70	60	0.40	0.73	0.51	18	0.49	0.70	0.62
ClusterLLM	77	0.60	0.84	0.70	150	0.21	0.85	0.68	60	0.47	0.77	0.59	18	0.44	0.68	0.59
Agentic clustering (ours)	77	<u>0.70</u>	<u>0.88</u>	<u>0.79</u>	150	0.28	<u>0.88</u>	<u>0.72</u>	60	<u>0.62</u>	<u>0.81</u>	<u>0.69</u>	18	<u>0.71</u>	<u>0.78</u>	<u>0.78</u>
<i>Discover k</i>																
BERTopic	90	0.60	0.84	0.69	166	0.26	0.82	0.70	67	0.43	0.73	0.57	67	0.27	0.64	0.36
TopicGPT	5	0.01	0.16	0.04	14	0.05	0.46	0.15	15	0.10	0.47	0.31	12	0.18	0.45	0.43
Huang & He	131	0.57	0.84	0.64	766	0.13	0.83	0.59	172	0.37	0.77	0.48	476	0.12	0.61	0.17
Agentic clustering, $k \pm 20\%$ (ours)	85	<u>0.69</u>	<u>0.88</u>	<u>0.80</u>	169	<u>0.39</u>	<u>0.91</u>	0.78	72	0.58	0.82	0.67	20	0.69	<u>0.78</u>	0.78
Agentic clustering, no k (ours)	81	0.57	0.85	0.66	231	0.29	0.88	<u>0.79</u>	76	<u>0.61</u>	<u>0.83</u>	<u>0.74</u>	14	<u>0.71</u>	0.77	<u>0.80</u>

Method	GoEmotions				20 Newsgroups				StackExchange				Cost	Token counts (M)		
	$\hat{k}$	ARI	NMI	ACC	$\hat{k}$	ARI	NMI	ACC	$\hat{k}$	ARI	NMI	ACC	total \$	Frontier	Cheap - in	Cheap - out
Given k																
LDA	27	0.01	0.04	0.10	20	0.19	0.36	0.34	121	0.01	0.31	0.08	\$0	–	–	–
SBERT+ k -means	27	0.02	0.07	0.11	20	0.42	0.58	0.61	121	0.16	0.52	0.29	\$0	–	–	–
BERTopic	27	0.01	0.03	0.24	20	0.31	0.53	0.39	70	0.17	0.51	0.30	\$0	–	–	–
LLM-embedding+ k -means	27	0.04	0.13	0.17	20	0.49	0.63	0.64	121	0.34	0.67	0.47	\$0.39	–	–	–
ClusterLLM	27	0.06	0.21	0.25	20	0.48	0.64	0.67	121	0.31	0.66	0.45	\$1.53	–	–	–
Agentic clustering (ours)	27	<u>0.08</u>	<u>0.24</u>	<u>0.35</u>	20	<u>0.55</u>	<u>0.67</u>	<u>0.73</u>	121	<u>0.44</u>	<u>0.73</u>	<u>0.56</u>	\$100 + \$22	4.31	210	17
Discover k																
BERTopic	429	0.01	0.09	0.22	272	0.17	0.51	0.20	70	0.17	0.51	0.30	\$0	–	–	–
TopicGPT	27	0.00	0.04	0.11	18	0.28	0.53	0.39	30	0.13	0.54	0.28	\$100 + \$66	4.13	100	43
Huang & He	207	0.03	0.17	0.16	377	0.22	0.53	0.29	2764	0.02	0.69	0.10	\$100 + \$44	0.68	384	33
Agentic clustering, $k \pm 20\%$ (ours)	27	0.08	<u>0.24</u>	0.32	18	0.53	<u>0.68</u>	0.72	130	<u>0.42</u>	<u>0.73</u>	<u>0.56</u>	\$100 + \$23	5.27	193	17
Agentic clustering, no k (ours)	15	<u>0.09</u>	0.23	<u>0.33</u>	19	<u>0.57</u>	0.68	<u>0.75</u>	21	0.14	0.56	0.21	\$100 + \$21	4.22	157	15

Event study: competitor journalists

[◀ Back](#)

τ	Est. (SE)	τ	Est. (SE)
-12	0.0442** (0.0179)	+1	-0.1248*** (0.0307)
-11	0.0535*** (0.0173)	+2	-0.1374*** (0.0316)
-10	0.0305** (0.0119)	+3	-0.1500*** (0.0336)
-9	0.0237** (0.0118)	+4	-0.1493*** (0.0349)
-8	0.0240** (0.0119)	+5	-0.1564*** (0.0353)
-7	0.0231* (0.0127)	+6	-0.1529*** (0.0323)
-6	0.0082 (0.0065)	+7	-0.1580*** (0.0330)
-5	-0.0024 (0.0054)	+8	-0.1536*** (0.0337)
-4	0.0019 (0.0041)	+9	-0.1296*** (0.0346)
-3	-0.0077 (0.0071)	+10	-0.1243*** (0.0345)
-2	0.0009 (0.0091)	+11	-0.1332*** (0.0331)
-1	-0.0203** (0.0094)	+12	-0.1405*** (0.0372)
0	-0.1279*** (0.0317)		
Stack $\times$ country FE: Yes		Stack $\times$ month FE: Yes	
Observations: 17,883		Pseudo R^2 : 0.693	

Event study: total Xinhua output

[◀ Back](#)

τ	Est. (SE)	τ	Est. (SE)
-12	0.0066 (0.0688)	+1	-0.1003 (0.0749)
-11	0.0194 (0.0631)	+2	-0.0725 (0.0684)
-10	-0.0289 (0.0516)	+3	0.0098 (0.0549)
-9	-0.0141 (0.0547)	+4	-0.1124* (0.0627)
-8	0.0206 (0.0999)	+5	-0.2225*** (0.0644)
-7	-0.0325 (0.0506)	+6	-0.1791** (0.0870)
-6	0.0130 (0.0426)	+7	-0.1852*** (0.0697)
-5	-0.0292 (0.0509)	+8	-0.1430 (0.1127)
-4	-0.0044 (0.0211)	+9	-0.2501*** (0.0894)
-3	0.0206 (0.0740)	+10	-0.3177*** (0.0864)
-2	-0.1172** (0.0536)	+11	-0.2846** (0.1113)
-1	-0.0459 (0.0445)	+12	-0.3365*** (0.0908)
0	-0.0927 (0.0629)		

Stack $\times$ country FE: Yes Stack $\times$ month FE: Yes
Observations: 17,325 Pseudo R^2 : 0.748

Event study: articles mentioning China

[◀ Back](#)

τ	Est. (SE)	τ	Est. (SE)
-12	-0.0165 (0.1113)	+1	0.0882 (0.1089)
-11	-0.0126 (0.1101)	+2	-0.1512 (0.1491)
-10	0.1181 (0.1692)	+3	0.2485 (0.2807)
-9	0.1842 (0.1823)	+4	0.4894*** (0.1860)
-8	0.4561 (0.3643)	+5	-0.0800 (0.1411)
-7	-0.0060 (0.1764)	+6	-0.0370 (0.1875)
-6	0.0137 (0.1055)	+7	0.0124 (0.1421)
-5	-0.0405 (0.1095)	+8	-0.0357 (0.1434)
-4	0.0725 (0.1046)	+9	-0.1607 (0.1527)
-3	-0.0457 (0.1347)	+10	0.0731 (0.1916)
-2	-0.1868 (0.1365)	+11	0.0041 (0.2018)
-1	-0.0553 (0.1449)	+12	0.0090 (0.1812)
0	0.1750 (0.1510)		
Stack $\times$ country FE: Yes		Stack $\times$ month FE: Yes	
Observations: 17,150		Pseudo R^2 : 0.491	

Event study: Xinhua's own staffing

[◀ Back](#)

τ	Est. (SE)	τ	Est. (SE)
-12	-0.0646* (0.0371)	+1	0.0080 (0.0325)
-11	-0.0563 (0.0369)	+2	0.0244 (0.0302)
-10	-0.0522 (0.0363)	+3	0.0164 (0.0364)
-9	-0.0347 (0.0271)	+4	0.0091 (0.0450)
-8	-0.0300 (0.0249)	+5	0.0129 (0.0460)
-7	-0.0283 (0.0229)	+6	0.0142 (0.0463)
-6	-0.0288 (0.0200)	+7	0.0142 (0.0467)
-5	-0.0034 (0.0128)	+8	0.0143 (0.0489)
-4	0.0195* (0.0117)	+9	0.0102 (0.0518)
-3	0.0128 (0.0154)	+10	0.0139 (0.0562)
-2	0.0187 (0.0178)	+11	0.0369 (0.0532)
-1	0.0250 (0.0219)	+12	0.0311 (0.0524)
0	0.0080 (0.0247)		
Stack $\times$ country FE: Yes		Stack $\times$ month FE: Yes	
Observations: 10,663		Pseudo R^2 : 0.243	

Event study: credited reprints

[◀ Back](#)

τ	Est. (SE)	τ	Est. (SE)
-12	0.0025 (0.1376)	+1	-0.0844 (0.2096)
-11	0.1477 (0.1448)	+2	-0.1076 (0.1783)
-10	0.0251 (0.1362)	+3	0.0013 (0.2054)
-9	0.1437 (0.1554)	+4	0.0973 (0.2312)
-8	0.0425 (0.1582)	+5	-0.1847 (0.2298)
-7	0.0838 (0.1531)	+6	-0.1505 (0.1889)
-6	-0.2160 (0.1501)	+7	-0.2957* (0.1629)
-5	-0.0543 (0.1494)	+8	0.0360 (0.2393)
-4	0.0361 (0.1153)	+9	-0.0524 (0.1826)
-3	0.2342* (0.1351)	+10	-0.3457* (0.2091)
-2	0.0874 (0.1145)	+11	-0.0803 (0.2231)
-1	0.3756* (0.1966)	+12	-0.0068 (0.1451)
0	0.0764 (0.1593)		
Stack $\times$ country FE: Yes		Stack $\times$ month FE: Yes	
Observations: 16,431		Pseudo R^2 : 0.506	

Event study: reprints of China articles

[◀ Back](#)

τ	Est. (SE)	τ	Est. (SE)
-12	0.0108 (0.4586)	+1	-0.3377 (0.8470)
-11	0.4887 (0.6943)	+2	-0.7560* (0.4167)
-10	-0.5066 (0.4115)	+3	0.3933 (0.4451)
-9	0.1185 (0.4997)	+4	0.5734* (0.3131)
-8	0.5244 (0.7899)	+5	0.0115 (0.5166)
-7	0.6933 (0.4671)	+6	0.3031 (0.4103)
-6	-0.0142 (0.2224)	+7	-0.6633 (0.5577)
-5	-0.1330 (0.3219)	+8	0.4792 (0.3418)
-4	0.3719 (0.3814)	+9	0.3018 (0.6669)
-3	-0.2247 (0.3078)	+10	-0.1196 (0.7244)
-2	-0.0749 (0.6107)	+11	-1.0006 (1.0975)
-1	0.3411 (0.4106)	+12	0.9434*** (0.2429)
0	-0.2482 (0.4279)		
Stack $\times$ country FE: Yes		Stack $\times$ month FE: Yes	
Observations: 9,213		Pseudo R^2 : 0.296	

Event study: mean frame popularity

[◀ Back](#)

τ	Est. (SE)	τ	Est. (SE)
-12	-0.0017 (0.0148)	+1	0.0095 (0.0127)
-11	0.0150 (0.0151)	+2	0.0035 (0.0119)
-10	-0.0003 (0.0173)	+3	0.0166 (0.0202)
-9	0.0026 (0.0178)	+4	-0.0140 (0.0152)
-8	0.0066 (0.0213)	+5	-0.0021 (0.0192)
-7	-0.0274** (0.0108)	+6	-0.0047 (0.0132)
-6	-0.0020 (0.0115)	+7	-0.0098 (0.0142)
-5	-0.0045 (0.0109)	+8	0.0145 (0.0103)
-4	-0.0161 (0.0105)	+9	0.0213 (0.0134)
-3	0.0226** (0.0110)	+10	0.0055 (0.0189)
-2	-0.0072 (0.0129)	+11	0.0044 (0.0175)
-1	0.0042 (0.0156)	+12	-0.0151 (0.0142)
0	-0.0114 (0.0147)		
Stack $\times$ country FE: Yes		Stack $\times$ month FE: Yes	
Observations: 16,087		Pseudo R^2 : 0.512	

Event study: frame pairs (the offset denominator)

[◀ Back](#)

τ	Est. (SE)	τ	Est. (SE)
-12	-0.0301 (0.0994)	+1	-0.0550 (0.0753)
-11	0.0157 (0.0773)	+2	-0.1202 (0.0852)
-10	-0.0538 (0.0599)	+3	0.1011 (0.0839)
-9	-0.0167 (0.0623)	+4	-0.0884 (0.0567)
-8	0.1112 (0.1490)	+5	-0.1942*** (0.0603)
-7	-0.0565 (0.0476)	+6	-0.1559 (0.0950)
-6	-0.0032 (0.0465)	+7	-0.1440* (0.0867)
-5	-0.0526 (0.0550)	+8	0.0873 (0.1877)
-4	0.0078 (0.0230)	+9	-0.1715 (0.1099)
-3	0.0481 (0.0762)	+10	-0.3037*** (0.0850)
-2	-0.1209** (0.0540)	+11	-0.2856** (0.1234)
-1	-0.0213 (0.0506)	+12	-0.2906*** (0.1007)
0	0.0054 (0.0991)		

Stack $\times$ country FE: Yes Stack $\times$ month FE: Yes
Observations: 17,325 Pseudo R^2 : 0.770