A semantic atlas of highly liked alphaXiv papers

Libor “Burny” Burian · 26 November 2026

Abstract

We collected 3,703 unique research papers with more than 100 public likes from alphaXiv. We embedded their titles and abstracts, projected them for display, and grouped them in a separate ten-dimensional space. Across 10 seeds, two-dimensional UMAP preserved 26.34% of each paper’s ten original neighbors on average. PCA preserved 2.86%. The paired difference was 23.49 percentage points, with a 95% seed interval of 23.36 to 23.61. At the display seed, 1,242 papers remained unassigned.

Motivation and contribution

A ranked feed answers which papers attracted attention. A researcher looking for related work also needs to know which papers discuss similar ideas and how much a two-dimensional view distorts those relationships. This release combines a verified selection boundary, source-space recommendations and inspectable topic labels in a downloadable atlas. The contribution is the audited collection and evaluation of this map.

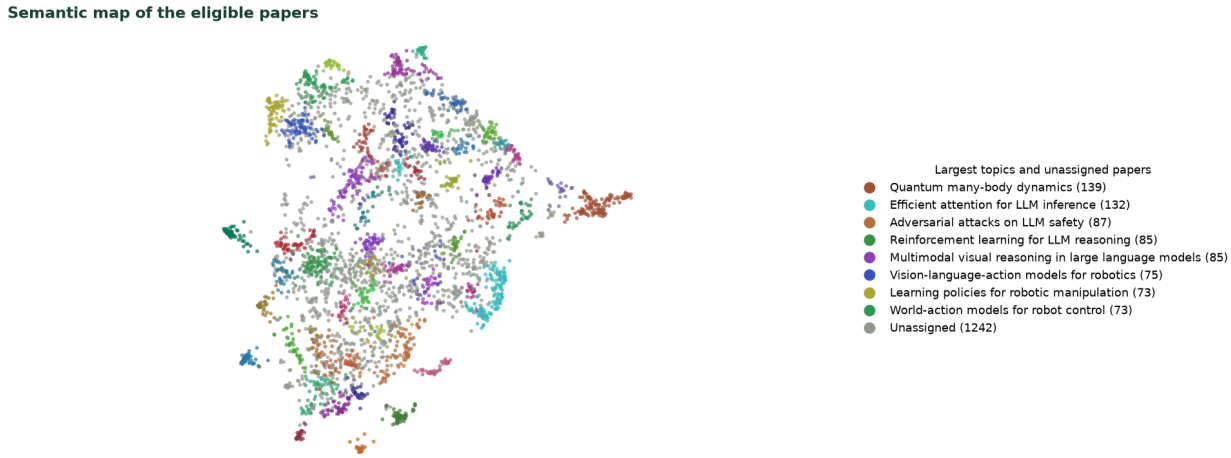

Figure 1. One point represents one paper. The legend lists the largest assigned groups and the unassigned set. Distances between distant regions have no calibrated interpretation.

Collection and methods

Census

The two accepted passes ran between 2026-11-26T11:13:01+00:00 and 2026-11-26T11:14:03+00:00. We sorted by all-time public votes, disabled both blog flags, and removed records explicitly marked as external posts. The first pass encountered 3,723 entries above the threshold, excluded 20 marked posts and deduplicated 0 group occurrences. The independent pass reproduced the eligible IDs. 0 papers used the title-only fallback. alphaXiv API specification

Representation and algorithms

The hosted text-embedding-3-large encoder received the title, a blank line and the abstract. We normalized its vectors and calculated exact cosine neighbors, excluding each paper itself. UMAP used 30 neighbors and cosine distance. Display coordinates used two dimensions and minimum distance 0.1. A separate ten-dimensional fit used minimum distance zero, followed by Euclidean excess-of-mass HDBSCAN with minimum cluster size 20 and minimum samples 5, including the point. Seed 42 determined the public map. OpenAI embeddings · UMAP clustering guide · HDBSCAN implementation

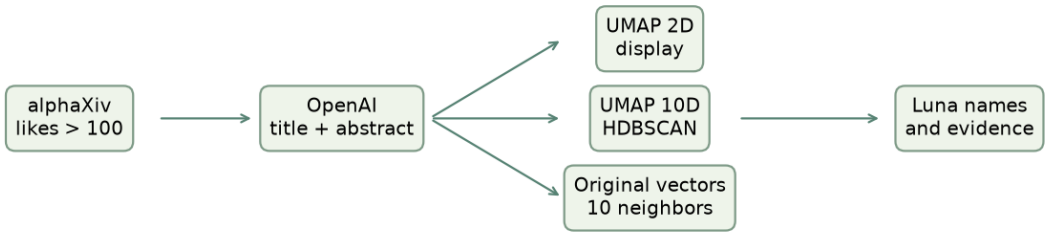

Figure 2. Display positions, group assignments and nearest-paper recommendations use three different transformations of the same hosted vectors.

Primary analysis

For each paper we divided the intersection of its original and displayed ten nearest neighbors by ten. We averaged across papers and calculated the UMAP minus PCA difference for seeds 42–51. The two-sided Student t interval uses seed-level differences. The shuffled-position control randomly reassigned the UMAP coordinates to paper IDs once per seed. The ablation replaced ten-dimensional UMAP with ten principal components while retaining the clustering settings. Preregistered analysis · Huang et al. on neighborhood preservation

Labeling

Luna received distinctive TF–IDF phrases and excerpts from members nearest their cluster centroid in the original vector space. Each displayed name links to supporting member IDs validated against the prompt. These are generated descriptions. Murray et al. on labeling scientific documents

Results and controls

Projection / assignment	Recall at 10	Clusters	Unassigned	Shared subject*
UMAP / UMAP–HDBSCAN	26.29%	56	33.54%	70.36%
PCA / primary assignments	2.86%	56	33.54%	70.36%
Shuffled / primary assignments	0.24%	56	33.54%	70.36%
PCA / PCA–HDBSCAN	2.86%	10	75.21%	68.66%

Table 1. Display seed 42. *Fraction of tagged within-cluster pairs sharing an arXiv category. The shuffled-tag control for primary assignments was 48.22%. Tags were available for 3,690 of 3,703 papers. The middle two rows reuse the primary labels to isolate changes in projection.

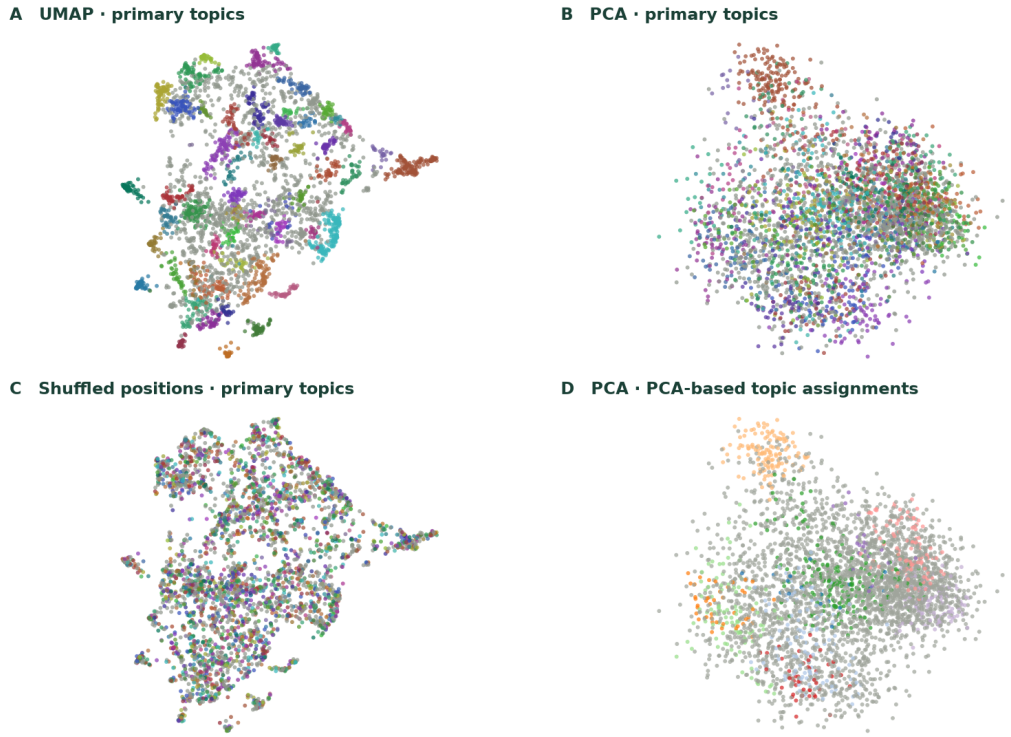

Figure 3. Panels A–C retain topic identity while changing screen coordinates. Panel D shows the assignments recovered by clustering PCA coordinates. Colors in D identify its own groups.

What the map loses

Even with the primary projection, a mean of 73.66% of the original ten neighbors fell outside the ten closest displayed points. The neighbor list exposes those relationships directly. The directional primary hypothesis was supported under the preregistered criterion.

Stability, limitations and reproducibility

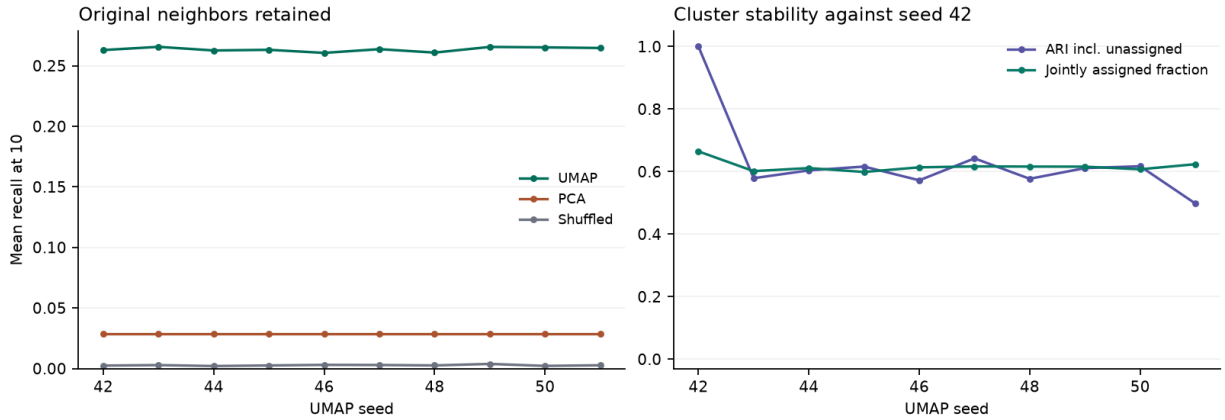

Figure 4. Each marker is one full-corpus run. The two stability series measure agreement with seed 42 and the fraction jointly assigned. Seed 42 compares with itself.

Across the other seeds, adjusted Rand agreement including unassigned points ranged from 0.498 to 0.642. The original and permuted subject agreement estimates pool within-cluster pairs, giving larger groups more weight. Cross-listed subjects are a coarse, incomplete external reference.

Example topics

Quantum many-body dynamics (139 papers). 2403.00091 · 2505.09707

Efficient attention for LLM inference (132 papers). 2505.21487 · 2504.06319

Adversarial attacks on LLM safety (87 papers). 2503.21598 · 2504.00218

Scope and limitations

The census represents public feed membership during the recorded window. Likes condition selection on platform attention. Older papers, topic popularity and exposure can affect inclusion. The analysis is descriptive and transductive. The interval quantifies variation from projection seeds conditional on these inputs. It does not estimate uncertainty across future papers. Hosted model revisions may alter a fresh run, so the exact vectors and returned labels accompany the release.

Reproduction and review

Restore the dated snapshot and run the offline build to reproduce tables, figures and the website. The archive includes checksums, paper metadata, embeddings, coordinates, supporting label IDs and per-seed scores. Dataset and numerical artifacts · Code and instructions · Run matrix

AI statement

A human provided the exact specification. Codex wrote the codebase and this paper to execute it. Human cleanup and rewriting are in progress.