How Long Do Trending AI Projects Stay Available and Maintained? Survival of 733 GitHub and Hugging Face Projects Picked by a Weekly Radar (June–October 2026)

Evgenii Arsentev, MD, PhD · ORCID 0000-0002-9120-7298 · Technical report, 10 November 2026 · DOI 10.5281/zenodo.24101933 · License CC BY 4.0

Summary

A public weekly AI radar picked 757 projects from 11 June to 31 October 2026; 733 GitHub and Hugging Face projects were observed in 50 daily snapshots (12 September to 31 October 2026). The study asked how many picked projects become unavailable, archived or abandoned, and whether this depends on source, pick month and stars at pick. The metric was pre-registered on 5 October 2026, before analysis. Unavailability (HTTP 404/451 on GitHub, 401/404 on Hugging Face, two snapshots in a row) gave Kaplan–Meier survival of 96.7 % at week 4, 94.6 % at week 8 and 93.0 % at week 12. In total, 23 projects became unavailable, and 5 of them later came back. The composite also counts abandonment (no push for 90 days) and archiving: 90 events, 67 of them abandonment. It falls from 91.9 % at week 12 to 61.5 % at week 16. This drop is produced by the 90-day rule and is not evidence of projects disappearing. GitHub projects in the top star quartile at pick (4,988 stars or more) had 1 event in 144. The difference holds when stratified by pick month (log-rank $p = 0.004$), but stars and the radar's changing selection cannot be fully separated. No difference between Hugging Face and GitHub was shown ($p = 0.21$ for unavailability).

1. Question and pre-registration

The question was fixed in a written protocol on 5 October 2026, before any analysis: what share of AI projects picked by a public weekly radar (GitHub repositories and Hugging Face models, picks from 11 June to 31 October 2026) becomes unavailable, archived or abandoned in the weeks after the pick, and does this depend on source, pick month and stars at pick. The question was chosen from three candidates (availability, star concentration, licence drift) because availability already showed a signal and licence drift did not.

2. Data

- Picks: 757 items picked from 11 June to 31 October 2026 (584 GitHub, 163 Hugging Face, 10 Hacker News). For each item the pick date and stars at pick are recorded.
- Snapshots: 50 daily snapshots from 12 September to 31 October 2026. Each record holds the HTTP status, archived flag, last push date (GitHub) and public counters.
- Analysed: 733 projects observed in at least one snapshot (575 GitHub, 158 Hugging Face). Excluded: 24 items, namely 10 Hacker News items that the snapshotter does not

track, and 14 projects (9 GitHub, 5 Hugging Face) picked on 31 October 2026, which had no follow-up.
- Only public repositories and models; no personal data. The snapshots belong to the public radar dataset series (DOI 10.5281/zenodo.22730450); the exact copy used here is in the supplementary file with SHA-256 sums.

3. Metric (as pre-registered)

- Unavailable: HTTP 404 or 451 on GitHub, or 401 or 404 on the Hugging Face metadata endpoint without a token, in two consecutive snapshots; the event date is the first of the pair. On that endpoint a gated model returns 200, so 401 means the model was deleted or made private.
- Archived: archived flag true.
- Abandoned: GitHub last push older than 90 days on the snapshot date.
- Survival by weeks since pick: Kaplan–Meier with delayed entry (projects picked before 12 September 2026 enter the risk set on that date), censoring at 31 October 2026. Group differences by log-rank test. No difference is called significant without a test.

4. Results

4.1 Unavailability

Under the unavailability rule alone (abandonment and archiving do not stop follow-up), 23 projects became unavailable: 20 on GitHub and 3 on Hugging Face.

Week since pick	**4**	**8**	**12**	**16**
Survival, unavailability	96.7 %	94.6 %	93.0 %	91.3 %
Projects at risk	315	276	193	122
Survival, permanent unavailability only	97.0 %	94.9 %	94.1 %	93.6 %

5 of the 23 unavailable projects later returned to HTTP 200 (picked on 15 June, 1 July, 5 July, 21 July and 10 September 2026); 18 stayed unavailable to the end of the window. 11 of the 23 were already unavailable on 12 September 2026, the first snapshot, so their event date is the entry date and the true date is earlier. Hugging Face versus GitHub: 3 events observed versus 5.41 expected for Hugging Face, log-rank chi2 = 1.57, df = 1, p = 0.21. No difference is shown.

4.2 Pre-registered composite (unavailable, archived or abandoned)

90 events: 67 abandoned (GitHub), 19 unavailable on GitHub, 3 unavailable on Hugging Face, 1 archived. One GitHub project counted as abandoned in the composite became unavailable later, which is why the unavailability rule above finds 20 GitHub events and the composite 19.

Week since pick	4	8	12	16	20
Survival, composite	95.9 %	93.5 %	91.9 %	61.5 %	54.5 %
Projects at risk	315	275	192	84	5

The drop after week 12 is produced by the 90-day rule and is not evidence of projects disappearing: a project can cross 90 days without a push only about 13 weeks after its last push, and in this window only projects picked in June and July had that much follow-up. The week-20 value rests on 5 projects and is not interpretable; the last event occurred at week 17.6. 17 composite events (11 unavailable, 5 abandoned, 1 archived) fall on the entry date. Excluding them leaves 716 projects and 73 events, with survival 97.1 %, 96.5 %, 95.6 % and 66.1 % at weeks 4, 8, 12 and 16. Two projects picked in October were already more than 90 days past their last push one day after the pick.

4.3 Source, pick month and stars

- Source (composite): GitHub 87 events versus 83.5 expected, Hugging Face 3 versus 6.5; log-rank chi2 = 2.75, df = 1, p = 0.098. No difference is shown.
- Pick month (composite): events 27 of 70 (June), 46 of 134 (July), 8 of 200 (August), 7 of 178 (September), 2 of 151 (October); log-rank chi2 = 18.77, df = 4, p = 0.0009. This is mostly mechanical: abandonment can only fire for the cohorts with long follow-up.
- Stars at pick, GitHub quartiles (cut points 374, 630 and 4,988 stars): events 29 of 143, 33 of 144, 24 of 144 and 1 of 144; log-rank chi2 = 13.33, df = 3, p = 0.004. Stratified by pick month, chi2 = 13.36, df = 3, p = 0.004. Most of the difference comes from abandonment (25, 22, 19 and 1 abandonment events by quartile). Unavailability events by quartile were 4, 10, 6 and 0, too few to test. Only 5 top-quartile projects had 16 weeks of follow-up.
- The radar’s selection changed over time. Median stars at pick on GitHub were 502 (June), 362 (July), 591 (August), 3,236 (September) and 4,217 (October), and the top-quartile share of picks went from 7 % in June and 3 % in July to 23 % in August, 46 % in September and 47 % in October. Stratification removes the month effect on the test, but stars and the selection rule cannot be fully separated in these data.
- Four log-rank tests were run on the composite and on unavailability. At a Bonferroni threshold of 0.0125, the pick-month and star-quartile tests pass, and both are shaped by the 90-day rule and by the changing selection described above.

5. What this means

For someone citing or building on a project found through a trending list, roughly 1 in 14 projects was gone 12 weeks after the pick (93.0 % available), and 5 of the 23 disappearances were temporary. Abandonment is the more common outcome than deletion: 67 projects crossed 90 days without a push, versus 23 that became unavailable. Low-star picks carried almost all events. The top quartile almost never left within the window, but only 5 top-quartile projects had 16 weeks of follow-up. The data describe one radar and one selection rule over 20 weeks, not GitHub or Hugging Face in general. Statuses before 12 September 2026 are

unknown. In practice, when citing fast-moving AI projects, a link to a DOI or an archived copy is safer than a link to a repository alone.

6. Limitations

1. Left truncation: statuses before 12 September 2026 are unknown, and 482 of the 733 projects were picked before the first snapshot.
2. The 90-day abandonment rule fires only after a long follow-up. The composite curve after week 12 and the pick-month difference mainly reflect this rule.
3. Renamed repositories are followed by the API and are not counted as deletions. Deletion, takedown (451) and a switch from public to private all produce the same unavailable status and are not separated.
4. One radar and one selection rule over 20 weeks. The results describe its picks, not GitHub or Hugging Face in general.

7. Deviations from the protocol

The protocol named Kaplan–Meier curves by source, pick month and star quartile. The first analysis (9 November 2026) published curves by source and event counts by month and quartile. After an internal adversarial review on 10 November 2026 the following were added without changing the metric: curves by pick month and by star quartile, numbers at risk at each reported week, the unavailability rule followed on its own, the count of temporary unavailability, a log-rank test of star quartiles stratified by pick month, and the composition of the excluded items. The pre-registered results are unchanged.

8. Files and reproduction

Supplementary archive: the 50 daily snapshots and the final pick list with SHA-256 sums; analyze.py (pre-registered analysis) and analyze_extra.py (review additions), both Python standard library only; results.json, extra.json and Kaplan–Meier tables (overall, by source, by pick month, by star quartile, unavailability rule). Running the two scripts on the snapshots reproduces every number in this report.

About the author

Evgenii Arsentev, MD, PhD (Евгений Арсентьев) — creator of Arsminal; Founder & CEO of Dad4Kids and ArsentevAI; founder of ReelsZavod. Physician and medical scientist who builds AI-agent systems in production and publishes the measurements.

ORCID 0000-0002-9120-7298 · Wikidata Q141398931 · Scopus Author ID 57202031127 · Google Scholar _nnPkaAAAAAJ · OpenAlex A5087068667 · Semantic Scholar 65769164

Projects and publications

- Arsminal: free terminal for AI agents, built as a messenger (DOI 10.5281/zenodo.24079643)

- AI Content Engine Economics: 1,497 short videos in 101 days by AI agents vs. people (DOI 10.5281/zenodo.22802228)
- Context economy of agent sessions, the U-curve of context clearing (DOI 10.5281/zenodo.22759217)
- AI managers in a live company, case series (DOI 10.5281/zenodo.23926850, 10.5281/zenodo.23927437)
- contextburn: open-source meter of agent-run cost
- hireburn: open method for AI-assisted hiring cost (payrollburn series)
- IETF Internet-Drafts: draft-arsentev-agent-run-metrics, draft-arsentev-llm-context-discovery
- Participant, W3C Agent Conformance and Benchmarking Community Group
- ReelsZavod: AI content factory for Russian-speaking businesses