Generative AI and the Future of Work in Technology Workers' Discourse: A Topic Modelling and Clustering Analysis of Hacker News Discussions From 2023 To 2026

Fazidah Wahit

Faculty of Artificial Intelligence and Frontier Technologies, UNITAR International University, Malaysia

DOI: https://doi.org/10.47772/IJRISS.2026.101100137

Received: 21 October 2026; Accepted: 26 October 2026; Published: 21 November 2026

ABSTRACT

Generative artificial intelligence (AI) is reshaping the organisation of work, yet relatively little is known about how the workers closest to the technology interpret its consequences for their own occupations. This study examines how technology practitioners discuss the effects of generative AI on jobs, careers and hiring, and how this discourse changed between January 2023 and September 2026. A new corpus was assembled from Hacker News, a large online forum for software and technology professionals. Keyword retrieval was combined with a zero-shot relevance classifier, which yielded 34,270 comments written by 13,576 contributors with an estimated precision of 85.3% on a held-out random sample. Four topic models, namely LDA, NMF and two BERTopic variants, were compared on coherence, diversity and coverage, and non-negative matrix factorisation with 15 topics provided the best balance. Sentence embeddings were then clustered with k-means to identify broader discourse frames, and the monthly topic composition was segmented with contiguity-constrained hierarchical clustering to detect discourse phases. Six frames emerged, covering corporate layoffs and *AI-washing*, AI as a tool or a threat to the programming craft, software careers and the junior pipeline, creative and knowledge-work displacement, AI in recruitment, and macro-level automation, economy and society. Labour-related discussion intensified from about 31 to 47 comments per 1,000 AI-related comments. The only robust change point, in May 2025, marks a shift towards layoffs, careers and the entry-level pipeline and away from debates about AI in recruitment, while macro-level debate about automation remained stable. The findings indicate that practitioner attention has turned towards how employers restructure work and entry routes, with implications for human resource management, career guidance and workforce policy.

Keywords generative artificial intelligence, future of work, topic modelling, clustering, online discourse

INTRODUCTION

Since the public release of ChatGPT in late 2022, generative artificial intelligence (AI) has moved rapidly from research laboratories into everyday work. Experimental and field studies report substantial productivity gains in professional writing (Noy & Zhang, 2023), customer support (Brynjolfsson, Li, & Raymond, 2025) and consulting tasks (Dell'Acqua et al., 2023), and occupational exposure analyses suggest that most occupations contain tasks that AI (Felten et al., 2021), and large language models (LLMs) in particular (Eloundou et al., 2024), could perform or assist. The evidence on employment, however, remains mixed. Studies of online freelancing platforms report declining earnings for exposed freelancers (Hui et al., 2024) and falling demand for automation-prone work (Demirci et al., 2024). Administrative data from Denmark show negligible effects on earnings and hours (Humlum & Vestergaard, 2025). In contrast, recent payroll data from the United States indicate a relative decline in employment among early-career workers in the most exposed occupations, including software development (Brynjolfsson, Chandar, & Chen, 2025). These contrasting results show that the labour-market consequences of generative AI are still unfolding and that they are unevenly distributed across occupations and career stages.

Software developers and other technology workers occupy a distinctive position in this transition, because they build and adopt generative AI tools, use them in their daily work, and belong to the occupations whose tasks are

most exposed to them (Eloundou et al., 2024). The way these workers interpret what is happening to their jobs, careers and hiring practices is therefore of direct interest to the social sciences. Sociological accounts of technology at work emphasise that the consequences of new technologies are negotiated through the interpretations of workers and managers rather than determined by technical capability alone (Kellogg et al., 2020). Economic accounts similarly stress that the balance between displacement and the creation of new tasks depends on organisational choices (Acemoglu & Restrepo, 2019), a pattern also observed in earlier waves of automation (Autor, 2015). Practitioner discourse can also act as an early indicator, because concerns raised by the people who perform the work may anticipate developments that official statistics register only later.

Existing studies of public discourse on generative AI have relied mainly on Twitter and Reddit and have concentrated on the first months after the release of ChatGPT. They have also treated employment as one theme among many, as in the analyses of Koonchanok et al. (2024), Leiter et al. (2024) and Miyazaki et al. (2024). More recent work on technology-career communities on Reddit shows that labour-market disruption became a salient concern after 2022 (Azevedo et al., 2026), and interview research describes how generative AI absorbs the developmental pathway through which junior engineers become senior engineers (Yu & Moon, 2026). To the best of the author's knowledge, which is based on literature searches conducted on 23 October 2026, no study has systematically mapped how technology workers discuss generative AI and work on Hacker News, one of the largest and most influential forums for software practitioners. Nor has any study followed this discourse into the 2025 to 2026 period in which agentic coding tools became widely available.

This study addresses this gap by combining topic modelling with clustering on a newly constructed corpus of Hacker News comments posted between January 2023 and September 2026. It is guided by three research questions. The first research question (RQ1) asks which topics structure technology workers' discussion of generative AI and work. The second research question (RQ2) asks into which broader discourse frames these discussions cluster and how the frames relate to the topics. The third research question (RQ3) asks how the discourse has changed over time and whether distinct discourse phases can be identified. In answering these questions, the study makes three contributions. Substantively, it documents a shift in practitioner attention towards layoffs, hiring and the erosion of entry-level career pathways. Methodologically, it combines validated corpus construction, a comparison of four topic models and two layers of clustering, including a statistically tested periodisation. Practically, it derives implications for human resource management, career guidance and workforce policy.

The remainder of the paper is organised as follows. The next section reviews research on generative AI and the labour market, on public and practitioner discourse about generative AI, and on topic modelling and clustering in the social sciences. The methodology section then describes the data collection, the construction and validation of the corpus, and the modelling procedures. The results are presented for each research question in turn, and the discussion relates them to prior studies before setting out implications and limitations. The final section concludes the paper.

LITERATURE REVIEW

Generative AI and the Labour Market

Debates about automation and employment have long contrasted displacement with complementarity. Task-based frameworks argue that technology displaces workers from some tasks while reinstating labour in new ones, so that net employment effects depend on the balance between the two processes (Acemoglu & Restrepo, 2019). Historical evidence also shows that automation has often complemented rather than replaced labour (Autor, 2015). Early estimates of the risk of computerisation (Frey & Osborne, 2017) were followed by exposure measures tailored to AI (Felten et al., 2021) and, more recently, to LLMs. Eloundou et al. (2024) estimate that around 80% of workers in the United States have at least 10% of their tasks exposed to LLMs. Exposure is highest in high-wage cognitive occupations, which reverses the pattern observed in earlier waves of automation. Whether such exposure leads to automation or to augmentation is, in this view, a matter of choices made by firms and technology developers rather than a technical given (Brynjolfsson, 2022).

Evidence on realised effects is accumulating but remains heterogeneous. Randomised and field studies show sizeable productivity gains that are often largest for less experienced workers, both in writing tasks (Noy & Zhang, 2023) and in customer support (Brynjolfsson, Li, & Raymond, 2025). These gains, however, disappear or even reverse for tasks that lie outside the frontier of AI capability (Dell'Acqua et al., 2023). On online labour platforms, freelancers in exposed occupations lost work and earnings after the release of ChatGPT (Hui et al., 2024), and demand for automation-prone freelance work declined (Demirci et al., 2024). In contrast, economy-wide estimates of near-term productivity effects are modest (Acemoglu, 2024), and Danish administrative data reveal little change in earnings or hours (Humlum & Vestergaard, 2025). The most recent payroll evidence nevertheless points to concentrated effects on young workers. The employment of workers aged 22 to 25 in the most exposed occupations declined by about 13% relative to comparable groups, mainly through reduced hiring rather than layoffs (Brynjolfsson, Chandar, & Chen, 2025). Qualitative evidence from software engineering is consistent with this pattern, as it suggests that generative AI absorbs entry-level tasks into workflows performed by senior engineers with AI assistance and thereby threatens the formation of expertise (Yu & Moon, 2026).

Public and Practitioner Discourse on Generative AI

A growing body of research analyses social-media discourse about generative AI. Studies of Twitter in the months after the release of ChatGPT identified enthusiasm about productivity alongside concerns about education, misinformation and employment (Koonchanok et al., 2024), a pattern also documented in the meta-analysis by Leiter et al. (2024). These studies also showed that attitudes varied by occupation (Miyazaki et al., 2024). Studies of Reddit extend this work to particular communities. In education communities, career anxiety emerged as one of several themes, and the discourse moved over time from detection and enforcement towards integration (Yüce et al., 2026). In technology-career communities, career-related discussion grew after 2022 and labour-market disruption became a distinct theme (Azevedo et al., 2026). Hacker News has been used to study developers' security concerns about coding assistants (Díaz Ferreyra et al., 2026), but not to examine how practitioners discuss the consequences of AI for work. Most existing studies also end before 2025, when agentic coding tools became widely available.

Topic Modelling and Clustering in the Social Sciences

Topic models represent documents as mixtures of latent themes and have become a standard instrument in the social sciences (Blei, 2012), where they have been used to study culture (DiMaggio et al., 2013) and politics (Grimmer & Stewart, 2013) at scale. They are most informative when they are combined with interpretive reading in a computational grounded theory design (Nelson, 2020), which links text mining to the development of social theory (Evans & Aceves, 2016). Probabilistic models such as latent Dirichlet allocation (LDA) (Blei et al., 2003) and matrix factorisation approaches such as non-negative matrix factorisation (NMF) (Lee & Seung, 1999) operate on word counts. More recent approaches instead cluster contextual sentence embeddings (Reimers & Gurevych, 2019) after dimensionality reduction (McInnes et al., 2018) and describe the resulting clusters with class-based term weights, as in BERTopic (Grootendorst, 2022). Comparative studies find that no model family dominates across corpora (Abdelrazek et al., 2023) and that short, informal texts from social media are particularly sensitive to the choice of model (Egger & Yu, 2022). This motivates an evaluation on several criteria, including topic coherence (Röder et al., 2015), topic diversity (Dieng et al., 2020) and coverage of the corpus.

Clustering offers a complementary perspective on the same texts. Document clustering groups texts according to their overall semantic similarity rather than their shared vocabulary, and the clustering of time periods can reveal phases in the development of a discourse. Both applications require an explicit and transparent choice of the number of clusters. This choice is commonly supported by internal validity indices such as the silhouette coefficient (Rousseeuw, 1987), the Davies and Bouldin index (Davies & Bouldin, 1979) and the Calinski and Harabasz index (Caliński & Harabasz, 1974).

Research Gap

Taken together, the literature reveals three gaps. First, research on discourse about generative AI and work has focused on general social-media audiences and on the early period after the release of ChatGPT. The discourse of technology practitioners on their own professional forums after 2024 has not been examined systematically.

Second, social-science applications typically rely on a single topic model without comparing alternatives or validating the relevance of the analysed corpus. Third, temporal change is usually described by dividing time at external events chosen in advance rather than by detecting and testing change points in the data. The present study addresses these gaps by combining a validated corpus, a comparison of topic models and two layers of clustering on a practitioner forum over almost four years.

METHODOLOGY

Figure 1 summarises the research design, which proceeds from data collection and corpus construction to text representation, topic modelling, clustering and interpretation. Each step is described in the following subsections.

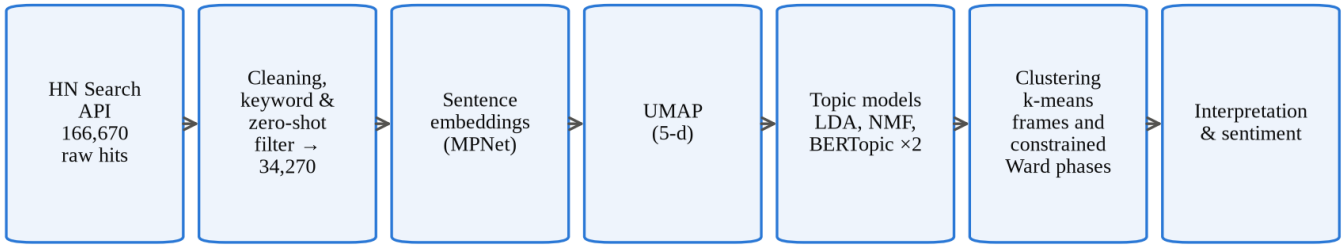


Figure 1. Research design, from data collection and corpus construction to topic modelling, clustering and interpretation.

Data Source and Collection

Hacker News is a high-traffic discussion forum operated by the start-up accelerator Y Combinator, and its users are predominantly software engineers, founders, researchers and other technology professionals. Its discussions are threaded and public, and its comments are typically longer and more argumentative than posts on microblogging platforms, which makes the forum well suited to the study of practitioner reasoning. All comments posted between 1 January 2023 and 30 September 2026 were searched through the public Hacker News Search API provided by Algolia. Each query paired an AI term, such as AI, LLM, ChatGPT, Copilot or agents, with a labour-market term, such as jobs, hiring, layoffs, career, junior, replace or salary, and 33 query strings were used in total. Because the API returns at most 1,000 results per request, each query was run over monthly windows that were split in half recursively whenever a window exceeded this limit. The collection took place on 23 October 2026 and returned 166,670 unique comments.

Corpus Construction and Validation

The raw comments were cleaned by removing HTML markup, quoted passages, code blocks and links. Because the search API also matches story titles and tolerates spelling variants, the corpus was defined on the text of each comment, and a comment was retained only if it contained both an AI term and a labour-market term. Ambiguous terms such as replace, interview, fired and recruit were counted only when they occurred within six tokens of a noun denoting workers, such as developers, workers or juniors. Idiomatic uses such as a good job or a cron job were masked before matching. Three groups of threads were excluded, namely the monthly job-advertisement threads, the monthly product-showcase threads, and the threads on the leadership crisis at OpenAI in November 2023, which concern corporate governance rather than work. Comments shorter than 20 words and exact duplicates were also removed, and Table 1 reports the number of comments retained at each step.

A manually assessed random sample showed that keyword matching alone favoured recall over precision. Only 63% of the sampled comments substantively discussed AI in relation to jobs, work, careers, hiring or the labour market, while many others mentioned a job only in passing. Each comment was therefore scored with a zero-shot natural language inference classifier based on DeBERTa-v3-large (Laurer et al., 2023, 2024), using the hypothesis that the comment discusses how AI affects jobs, work, workers, careers, hiring or employment. The decision threshold was fixed in advance by the rule of maximising the F1 score subject to a precision of at least

0.85 on a labelled tuning sample of 181 comments. This rule gave a threshold of 0.15 with a precision of 0.85 and a recall of 0.81. The precision of the final corpus was then estimated on an independent held-out random sample of 150 comments that was assessed against the same criterion, which gave 85.3% with a Wilson 95% confidence interval from 78.8% to 90.1%. The final corpus contains 34,270 comments written by 13,576 pseudonymised contributors in 11,890 discussion threads, with a median length of 88 words. Because the data are publicly available forum posts and the study involved no interaction with participants, usernames were replaced with salted one-way hashes before analysis, and the text reports only short paraphrased descriptions of comments, in line with guidance on research ethics for public online data (Zimmer, 2010) and on the expectations of platform users (Fiesler & Proferes, 2018).

Table 1. Construction of the corpus.

Step	Comments
Unique comments returned by the search API	166,670
After excluding job-advertisement and product-showcase threads	157,328
After removing comments shorter than 20 words	150,592
Comment text contains an AI term	138,852
Comment text also contains a labour-market term	61,632
After excluding OpenAI leadership-crisis threads	61,469
After removing exact duplicates	61,423
Zero-shot relevance score of at least 0.15 (final corpus)	34,270

Text Representation

For the models based on word counts, comments were lower-cased, tokenised and lemmatised, and standard English stop words were removed together with the terms that define the corpus, such as AI, LLM and job, because these terms occur in almost every document by construction. Frequent two-word phrases such as white collar and junior developer were detected automatically, and terms that occurred in fewer than ten comments or in more than half of the comments were discarded. For the analyses based on embeddings, each cleaned comment was encoded with the all-mpnet-base-v2 sentence transformer (Reimers & Gurevych, 2019) into a normalised 768-dimensional vector. The vectors were then projected onto five dimensions with UMAP using 15 neighbours and the cosine metric (McInnes et al., 2018).

Topic Modelling and Model Selection

Four topic-model specifications were compared, namely LDA (Blei et al., 2003), NMF applied to TF-IDF weights (Lee & Seung, 1999), BERTopic with HDBSCAN density clustering (Campello et al., 2013) in which the native topics were merged into K topics, and BERTopic with k-means clustering (Grootendorst, 2022). Each specification was estimated for K equal to 10, 15, 20, 25 and 30 with three random seeds, and all models shared the same vocabulary. The topics were evaluated with the C_V and normalised pointwise mutual information (C_NPMI) coherence of their ten most heavily weighted words (Röder et al., 2015). They were also evaluated with topic diversity, defined as the proportion of unique words among the 25 leading words of all topics (Dieng et al., 2020). Two further criteria were considered because they matter for substantive interpretation, namely coverage, which is the share of comments assigned to a topic, and the share of the largest topic, since a high value signals a residual topic that absorbs heterogeneous comments.

Clustering

Clustering was applied at two levels. At the level of comments, k-means was applied to the UMAP projection of the embeddings in order to identify broad discourse frames, understood as groups of comments that are semantically similar regardless of their exact vocabulary. Values of k from 2 to 6 and from 8 to 40 in steps of two were evaluated with the silhouette coefficient (Rousseeuw, 1987), the Davies and Bouldin index (Davies & Bouldin, 1979) and the Calinski and Harabasz index (Caliński & Harabasz, 1974). Seed stability was measured as the mean pairwise adjusted Rand index across three seeds. The clusters were described by their most distinctive words and their most central comments, and they were related to the topics through cross-tabulation

and through the normalised mutual information (NMI) and adjusted Rand index (ARI) between the two partitions.

At the level of time, months were clustered into discourse phases. Each month was represented by its mean topic-proportion vector after a Hellinger transformation, and the months were segmented with Ward's agglomerative clustering (Ward, 1963) under a contiguity constraint that allowed only adjacent months to be merged. For each number of phases from two to seven, the proportion of variance explained (R^2) was compared with a null distribution obtained from 999 random permutations of the month order. The number of phases was then chosen to maximise the excess of R^2 over the permutation mean. The stability of each boundary was assessed with 200 bootstrap replicates in which comments were resampled within months, and the two periods separated by the most stable boundary were compared directly with bootstrap confidence intervals.

Temporal Intensity, Affect and Interpretation

To separate growth in labour-related discussion from the general growth of discussion about AI, the monthly volume of the corpus was divided by the number of Hacker News comments that matched the query AI in the same month, obtained from the same API as daily counts. The sentiment of each comment was scored with VADER (Hutto & Gilbert, 2014) and is reported descriptively. Topics and clusters were labelled after reading their leading terms and their most representative comments, following the logic of computational grounded theory (Nelson, 2020). Differences across phases and frames were examined with chi-square and Kruskal and Wallis tests, for which effect sizes are reported. All analyses were conducted in Python with scikit-learn, gensim, BERTopic and sentence-transformers, using fixed random seeds to ensure reproducibility.

RESULTS

Volume and Intensity of Labour-Related Discourse on AI

Figure 2 presents the monthly volume and intensity of the discourse. Discussion rose sharply in March 2023, the month in which GPT-4 was released, before returning to a lower level for the remainder of 2023 and most of 2024. From late 2024 onwards, both volume and intensity increased. Averaged over the phases identified below, the corpus grew from about 497 comments per month in the first phase to 1,204 in the last phase, and labour-related discussion rose from 30.6 to 46.8 comments per 1,000 AI-related comments on the platform. The growth therefore reflects a shift of attention within the discussion of AI towards questions of work, and not merely an increase in the overall volume of discussion about AI.

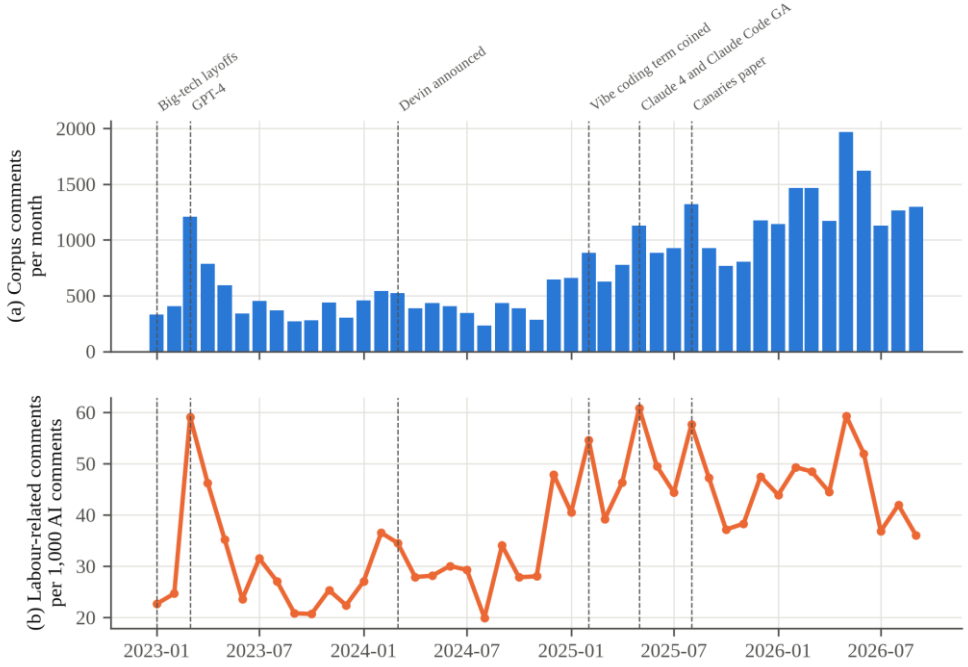


Figure 2. Monthly number of corpus comments (a) and corpus comments per 1,000 Hacker News comments matching the query AI (b). Dashed lines mark selected external events for orientation only.

Topics of the Discourse (RQ1)

Table 2 and Figure 3 summarise the comparison of topic models. LDA produced the lowest coherence and diversity at every value of K, for instance a C_V of 0.461 at K equal to 15. BERTopic with HDBSCAN reached the highest coherence, with a C_V of 0.548 and a C_NPMI of 0.074 at K equal to 30, but it assigned only 55% of the comments to topics and treated the remainder as outliers. It also produced a residual topic that contained about a quarter of the assigned comments. Among the specifications that cover the whole corpus, NMF with 15 topics reached the highest C_V (0.536), a C_NPMI (0.062) equal to the best value, and a diversity of 0.64, which is considerably higher than the 0.46 obtained by BERTopic with k-means. NMF with 15 topics was therefore selected for interpretation.

Table 2. Comparison of topic models at K equal to 15 (mean and standard deviation over three seeds, with the full grid shown in Figure 3).

Model	C_V	C_NPMI	Diversity	Coverage	Largest topic
LDA	0.461 (0.003)	0.013 (0.001)	0.37	100%	12.7%
NMF	0.536 (0.005)	0.062 (0.004)	0.64	100%	12.7%
BERTopic-HDBSCAN	0.520 (0.005)	0.059 (0.004)	0.57	55%	25.8%
BERTopic-KMeans	0.507 (0.006)	0.043 (0.002)	0.46	100%	12.1%

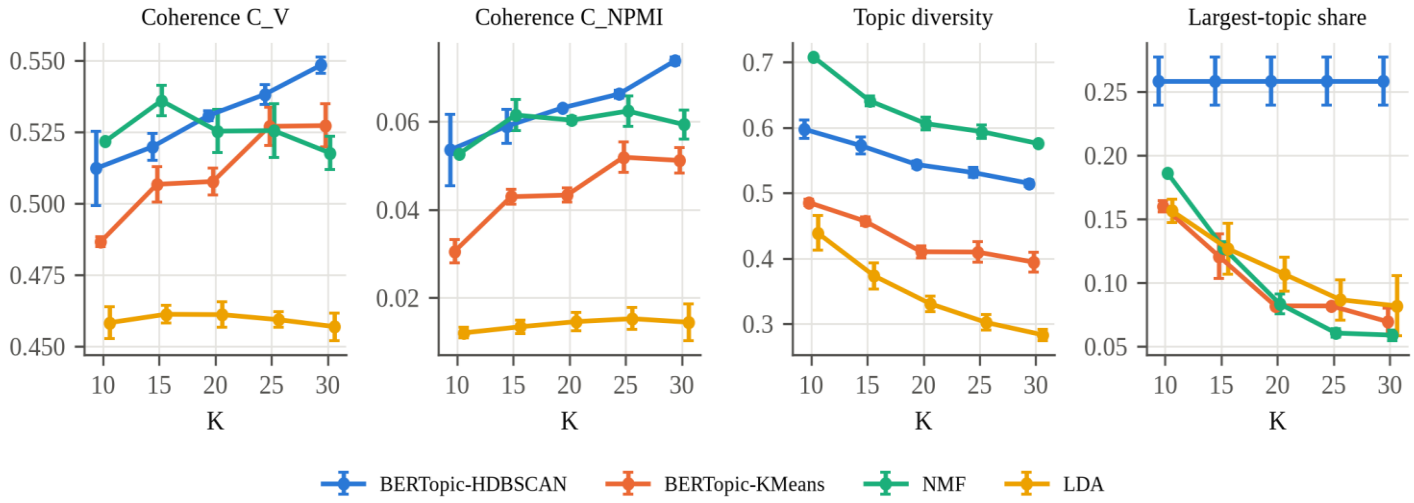


Figure 3. Coherence (C_V and C_NPMI), topic diversity and share of the largest topic by model and number of topics K (mean and standard deviation over three seeds).

Table 3 lists the 15 topics with their most heavily weighted words and their prevalence, measured as the mean document-topic weight, and Figure 4 displays the prevalence of each topic. The most prevalent topics concern the broad societal and economic consequences of AI (13.9%), the automation of everyday work tasks (9.7%), skills and career paths (8.8%) and corporate layoffs (8.7%). A second group of topics concerns concrete labour-market mechanisms, namely recruitment and screening, pay and labour costs, hiring decisions and the pipeline from junior to senior positions. A third group concerns the practice of software work itself, including coding and code review, engineering roles and management, and the use of AI as a productivity tool. Taken together, the topics show that practitioners discuss the labour implications of generative AI across the entire employment relationship, from entry through recruitment and junior roles, through the labour process and its rewards, to exit through layoffs.

Table 3. Topics of the final NMF model with 15 topics, ordered by prevalence. The frame column indicates the discourse frame in Table 4 that holds the largest share of the topic's comments.

Topic	Leading words	Prevalence	Frame
AI, economy and societal futures	world, happen, create, come, change, lose, technology, society	13.9%	C6
Everyday work and task automation	work, day, project, hour, automate, hard, team, task	9.7%	C6
Skills, learning and career paths	learn, career, skill, experience, start, programming, hard, field	8.8%	C3
Corporate layoffs and tech-firm strategy	company, layoff, employee, tech, lay, product, big, ceo	8.7%	C1
Coding practice and code review	code, write, programmer, understand, coding, writing, review, read	7.7%	C3
AI in recruitment and screening	interview, candidate, resume, question, ask, application, answer, apply	7.6%	C5
Human versus machine labour and AGI	human, task, labor, machine, robot, agent, agi, cheap	6.5%	C6
Pay, salaries and labour costs	pay, money, salary, cost, spend, cheap, employer, low	6.4%	C6
Replacement claims	replace, programmer, software engineer, ceo, artist, ceos, writer, hype	6.1%	C6
Hiring decisions and talent	hire, person, instead, process, cheap, talent, team, startup	5.1%	C1
Engineering roles and management	engineer, engineering, team, software engineer, manager, prompt, role, product	4.4%	C3
AI as a productivity tool	tool, useful, help, productive, productivity, build, replacement, task	4.3%	C3
Workers, wages and class	worker, white collar, labor, wage, blue collar, tech, automation, robot	4.0%	C6
Demand for software developers	developer, software, dev, team, software developer, development, experience, productive	3.6%	C3
Junior–senior pipeline	junior, senior, level, experience, senior engineer, train, role, junior engineer	3.3%	C3

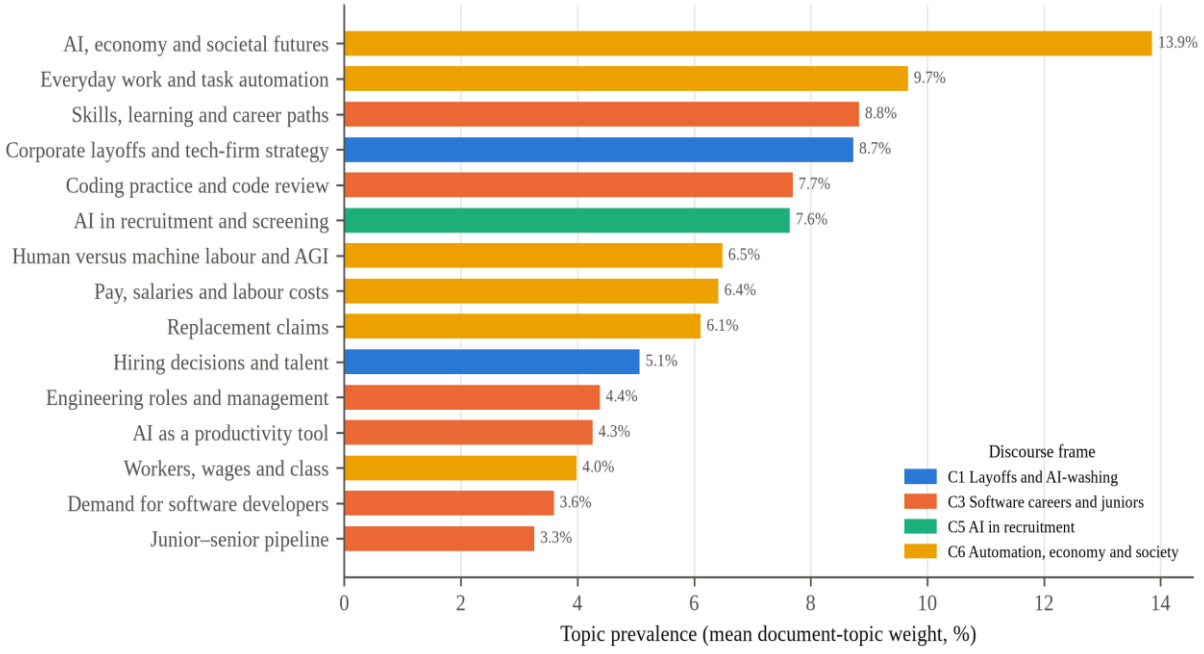


Figure 4. Topic prevalence measured as the mean document-topic weight, coloured by the discourse frame in which each topic is most concentrated.

Discourse frames identified by clustering (RQ2)

Figure 5 presents the cluster validity indices. Across the full range of k, the indices favour very coarse partitions, and k equal to 2 attains the highest silhouette (0.53), although such solutions separate only broad registers of discussion. Within the substantively useful range of k of at least 4, k equal to 6 obtained the best aggregate rank across the four criteria. This solution represents a local maximum of the silhouette (0.41) and is highly stable across seeds, with an ARI of 0.995, and it was therefore retained. The separation of the clusters in the original 768-dimensional embedding space is weak, with a cosine silhouette of 0.05, so the frames are best interpreted as regions of a continuous semantic space rather than as sharply bounded groups.

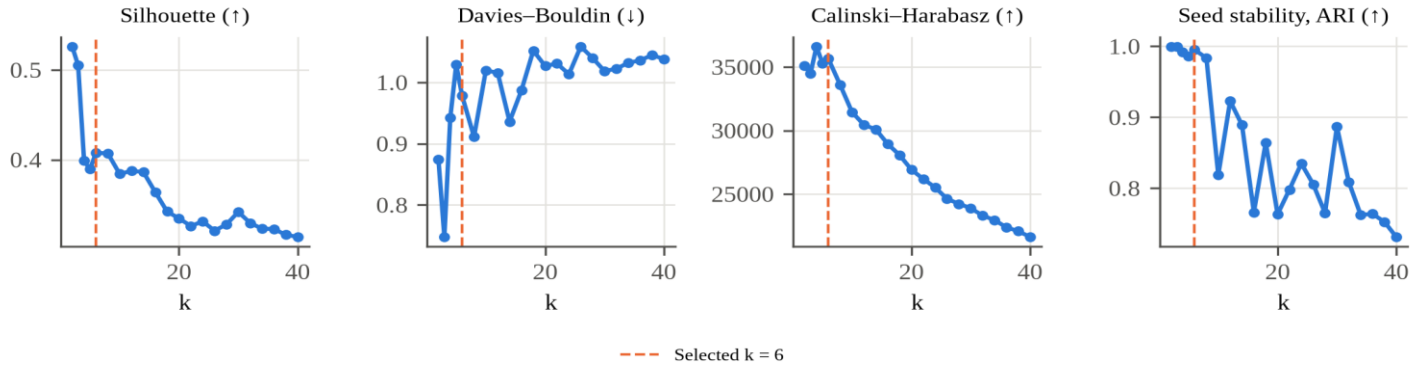


Figure 5. Internal validity of the k-means clustering of comment embeddings in the UMAP space by number of clusters k. The dashed line marks the selected value of six clusters.

Table 4 describes the six discourse frames. The largest frame is C6 Automation, economy and society (22.8%), in which automation is discussed in relation to society, the economy, robotics, artificial general intelligence and universal basic income. The second largest frame, C3 Software careers and juniors (22.3%), concerns the careers and learning of developers and the relationship between junior and senior engineers. The frame C1 Layoffs and AI-washing (17.5%) focuses on firms, chief executives and layoffs. Its most typical comments question whether announced job cuts attributed to AI reflect genuine productivity gains or serve as a narrative that justifies cost reduction, a practice that commentators describe as *AI-washing*. The frame C2 Tool or threat to the craft (14.1%) weighs the usefulness of AI as a tool against its threat to the craft of programming. The frame C4 Creative and knowledge-work displacement (13.0%) extends the discussion to artists, writers, translators and other creative and knowledge workers. Finally, C5 AI in recruitment (10.3%) concerns the use of AI on both sides of recruitment, by candidates who prepare applications and interviews with AI and by employers who screen applicants with it.

Table 4. Discourse frames obtained from the k-means clustering of comment embeddings with six clusters. The last column gives the share of comments with negative sentiment.

Frame	Share	Distinctive terms	Negative
C1 Layoffs and AI-washing	17.5%	company, layoff, employee, hire, tech, ceo, worker, salary, ceos	33%
C2 Tool or threat to the craft	14.1%	code, junior, task, output, replace, understand, prompt, human, developer	28%
C3 Software careers and juniors	22.3%	code, developer, software, engineer, learn, junior, career, coding, programmer	27%
C4 Creative and knowledge-work displacement	13.0%	artist, art, human, writer, creative, agent, voice, content, tool	31%
C5 AI in recruitment	10.3%	interview, candidate, resume, application, hire, recruiter, apply, applicant, question	28%
C6 Automation, economy and society	22.8%	human, labor, society, economy, worker, robot, technology, automation, agi	36%

Figure 6 cross-tabulates topics and frames. Several topics are concentrated in a single frame. The clearest example is AI in recruitment and screening, of which 62% of comments fall in C5 AI in recruitment. Similarly, 55% of the comments on AI, economy and societal futures fall in C6 Automation, economy and society, and 53% of those on the junior-senior pipeline fall in C3 Software careers and juniors. Other topics, such as replacement claims, pay and labour costs, and everyday task automation, are distributed across several frames. The agreement between the 15 topics and a k-means partition of the embeddings into 15 clusters is low, with an NMI of 0.13 and an ARI of 0.07, which indicates that the two methods capture different properties of the texts. The topics describe what is discussed, in the sense of the shared vocabulary of an issue, whereas the embedding clusters describe how it is framed. The vocabulary of replacement, for example, appears both in macro-level debates about society and in debates about the layoffs of particular firms.

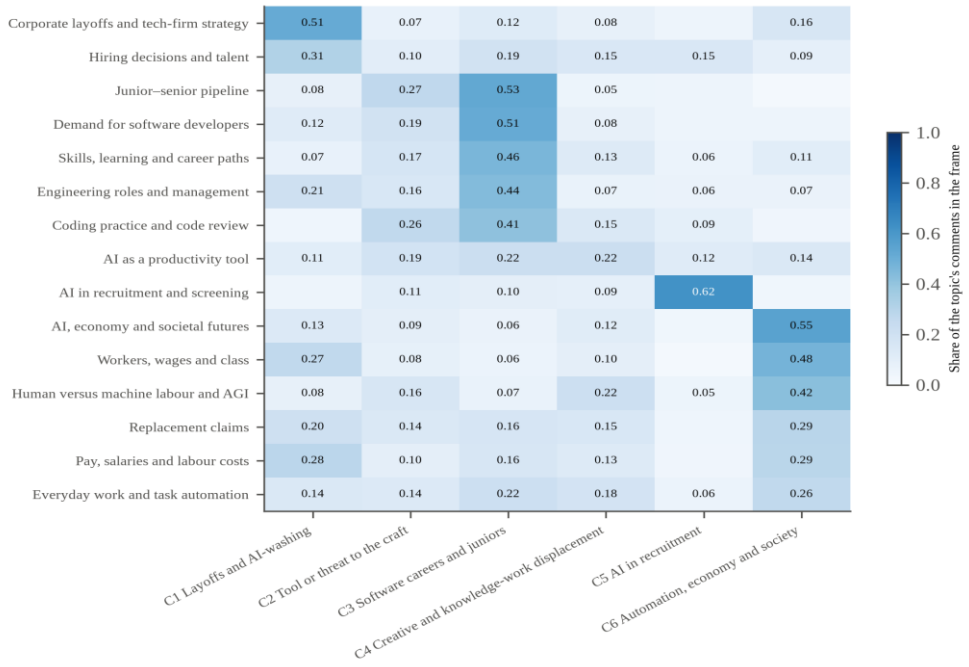


Figure 6. Distribution of the comments of each topic, assigned by dominant topic, across the six discourse frames. Cells below 0.05 are left blank.

Temporal Change and Discourse Phases (RQ3)

Table 5 reports the periodisation of the discourse. For every number of phases, the contiguity-constrained segmentations explained significantly more variance in the monthly topic composition than segmentations of randomly ordered months, with all permutation p values at or below 0.001. The excess over the permutation null was highest for four phases, with an R^2 of 0.36 against a null mean of 0.13. Solutions with five or six phases performed almost as well, whereas unconstrained clustering of the same monthly profiles did not yield contiguous periods. The boundaries differ markedly in robustness. The break between April and May 2025 appears in every solution from two to seven phases and recurs in 68% of the bootstrap replicates within one month. By contrast, the boundaries in December 2023 and August 2024 recur in only 35% and 36% of the replicates. The analysis therefore supports one robust change point in May 2025, which separates two broad periods, and the earlier sub-periods are reported as tentative.

Table 5. Discourse phases obtained from the contiguity-constrained clustering of months. The frame columns give the mean monthly share of comments in selected frames.

Phase	Months	n	Comments per month	Per 1,000 AI	C6 Society	C5 Recruitment	C1 Layoffs	C3 Careers
Phase 1	2023-01 to 2023-11	11	497	30.6	25%	18%	12%	17%
Phase 2	2023-12 to 2024-07	8	424	29.5	24%	14%	15%	20%
Phase 3	2024-08 to 2025-04	9	548	37.6	18%	14%	15%	22%
Phase 4	2025-05 to 2026-09	17	1,204	46.8	23%	7%	20%	23%

Because only the May 2025 change point is robust, the main temporal comparison contrasts the two periods that it separates, January 2023 to April 2025 and May 2025 to September 2026, using monthly averages with 95% confidence intervals obtained by resampling comments within months (Table 6). Across this break, the layoffs frame grew from 14.1% to 20.0% of comments (+5.9 percentage points, 95% CI +5.0 to +6.6) and the software-careers frame grew from 19.7% to 23.5% (+3.8 percentage points, 95% CI +2.8 to +4.7). The recruitment frame roughly halved, from 15.5% to 7.1% (-8.4 percentage points, 95% CI -9.1 to -7.6), and the frame on creative and knowledge-work displacement also declined (-3.1 percentage points, 95% CI -3.9 to -2.3). In contrast, the macro-level frame on automation, economy and society remained essentially unchanged, at 22.5% before and 22.7% after the break (+0.2 percentage points, 95% CI -0.7 to +1.1). It thus forms a persistent background to the discourse rather than an early concern that receded.

The same pattern appears at the level of topics. The junior-senior pipeline rose from 2.6% to 3.7% of topic weight, an increase of 43%, and corporate layoffs rose from 7.6% to 9.4%, whereas AI in recruitment and screening fell from 9.1% to 6.9%. The more speculative topics declined modestly, namely human versus machine labour and AGI (-1.2 percentage points, 95% CI -1.5 to -0.8) and replacement claims (-1.0 percentage points, 95% CI -1.3 to -0.8). The robust boundary coincides with the general release of agentic coding tools, such as Claude 4 and Claude Code on 22 May 2025, but the design does not permit causal attribution to any single event.

Table 6. Comparison of the periods before and after the robust change point in May 2025, the mean monthly share of comments in each frame and the mean weight of selected topics. Differences are in percentage points with 95% bootstrap confidence intervals.

Frame or topic	**Jan 2023 to Apr 2025**	**May 2025 to Sep 2026**	**Difference (pp)**	**95% CI**
C1 Layoffs and AI-washing	14.1%	20.0%	+5.9	+5.0 to +6.6
C3 Software careers and juniors	19.7%	23.5%	+3.8	+2.8 to +4.7
C2 Tool or threat to the craft	13.0%	14.6%	+1.6	+0.8 to +2.4
C6 Automation, economy and society	22.5%	22.7%	+0.2	-0.7 to +1.1
C4 Creative and knowledge-work displacement	15.2%	12.1%	-3.1	-3.9 to -2.3
C5 AI in recruitment	15.5%	7.1%	-8.4	-9.1 to -7.6
Corporate layoffs and tech-firm strategy (topic)	7.6%	9.4%	+1.8	+1.5 to +2.2
Junior–senior pipeline (topic)	2.6%	3.7%	+1.1	+0.9 to +1.3
Pay, salaries and labour costs (topic)	6.0%	6.7%	+0.7	+0.4 to +1.0
Human versus machine labour and AGI (topic)	7.2%	6.0%	-1.2	-1.5 to -0.8
Replacement claims (topic)	6.7%	5.7%	-1.0	-1.3 to -0.8
AI in recruitment and screening (topic)	9.1%	6.9%	-2.2	-2.6 to -1.9

The finer phases add descriptive detail (Table 5), and the composition of frames differs across them ($\chi^2(15) = 1234$, $p < .001$, Cramér's $V = 0.11$), although their earlier boundaries are not robust. Figure 7 traces the frames month by month. Recruitment was prominent in 2023, when debates centred on candidates who used ChatGPT in applications and interviews, and the recruitment topic peaked in the third phase at 10.4% of topic weight before falling sharply. The macro-level frame on society was highest in 2023, declined in late 2024 and early 2025, and rose again in 2026, and Figure 8 shows the corresponding changes in topic prevalence relative to the whole period.

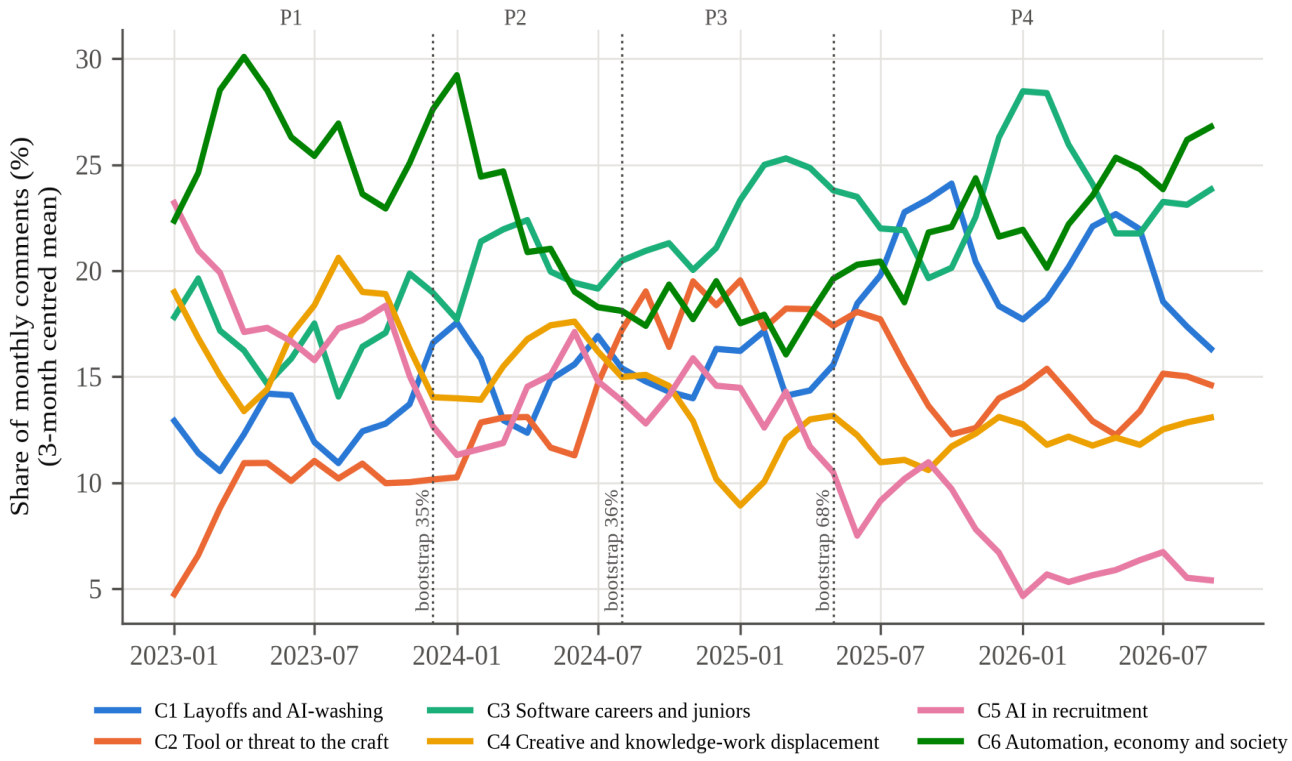


Figure 7. Monthly share of comments in each discourse frame as a three-month centred mean. Dotted lines mark phase boundaries and are labelled with the share of bootstrap replicates that reproduce each boundary within one month.

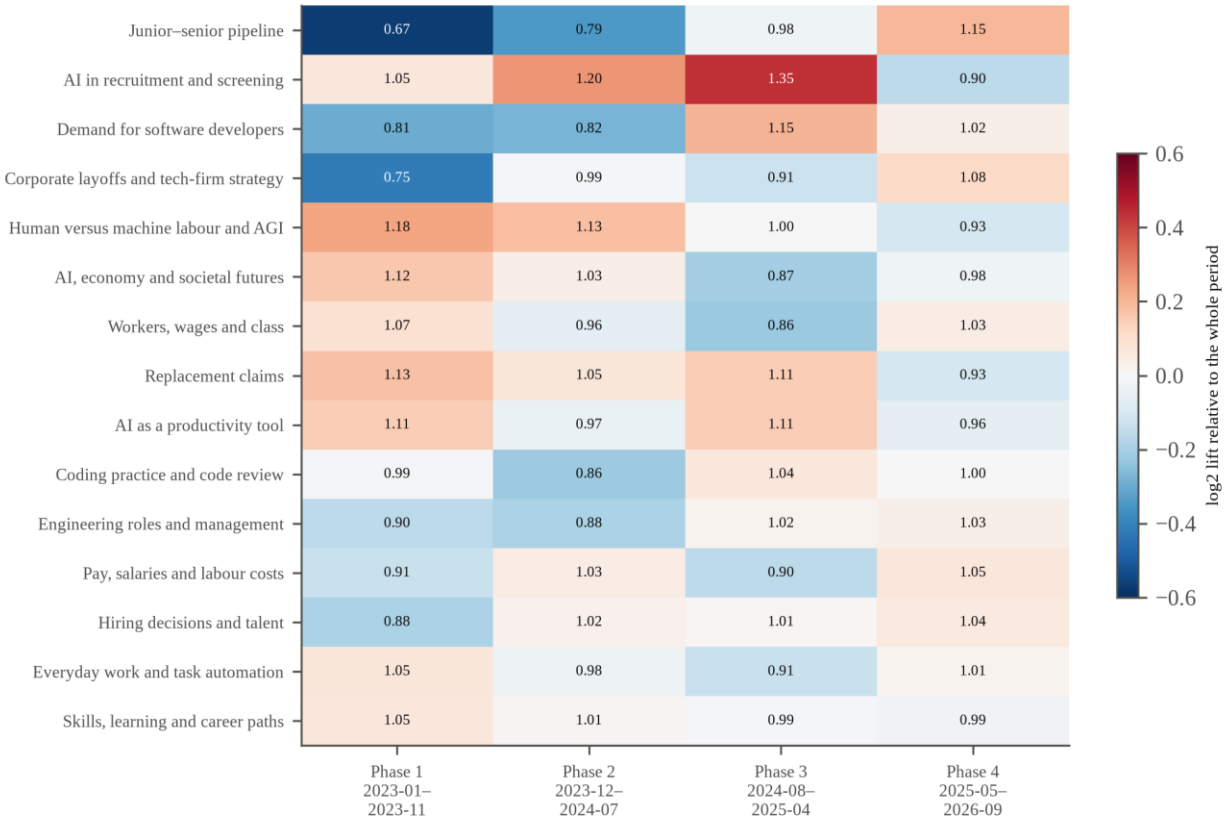


Figure 8. Topic prevalence in each phase relative to the whole period, where a value of 1 denotes the average. Rows are ordered by variability across phases.

Affective Tone

The lexicon-based sentiment scores were predominantly positive in every frame, which reflects the argumentative and often ironic register of technical discussion more than genuine enthusiasm, and the scores are therefore interpreted only comparatively. The share of negative comments was highest in C6 Automation, economy and society (36%) and C1 Layoffs and AI-washing (33%) and lowest in C3 Software careers and juniors (27%). It rose from 28% in the first phase to 32% in the last. These differences are statistically significant but very small, with Kruskal and Wallis effect sizes (ε^2) of 0.011 across frames and 0.003 across phases.

DISCUSSION

The analysis reveals a practitioner discourse on generative AI and work that is broad, internally structured and changing over time. The first finding concerns the direction of change. Macro-level debate about automation, the economy and society was present throughout the period, but around this stable core the discourse moved towards organisational practice. From 2025 onwards, layoffs and their justification, the career consequences for developers and the junior-senior pipeline gained ground. Meanwhile, the debates about AI in recruitment that characterised the period after the release of ChatGPT receded and the more speculative topics about replacement declined. This trajectory parallels the development of the economic evidence, which has moved from estimates of exposure (Eloundou et al., 2024) to measured effects on hiring (Brynjolfsson, Chandar, & Chen, 2025). It also suggests that practitioners increasingly perceive generative AI as a managerial choice about labour rather than as a technological inevitability, in line with arguments that the direction of AI deployment is a choice (Brynjolfsson, 2022) and that algorithmic technologies reshape managerial control (Kellogg et al., 2020). The prominence of comments that question whether layoffs attributed to AI reflect real productivity gains further indicates that workers contest the explanations that firms give for their decisions. This pattern is consistent with the view of technology at work as contested terrain (Kellogg et al., 2020).

The second finding concerns the growing salience of the junior-senior pipeline, which provides a discursive counterpart to recent labour-market and qualitative evidence. Payroll data show declines in employment that are concentrated among early-career workers in exposed occupations (Brynjolfsson, Chandar, & Chen, 2025), and interview evidence describes AI absorbing the tasks through which junior engineers used to build expertise (Yu & Moon, 2026). The present data show practitioners articulating the same concern, namely that the automation of entry-level work leaves it unclear how the next generation of senior engineers will be trained. Because productivity studies report the largest gains for less experienced workers in customer support (Brynjolfsson, Li, & Raymond, 2025) and in professional writing (Noy & Zhang, 2023), these findings point to a tension in which generative AI may support individual novices while reducing the incentives of firms to hire them.

The third finding is methodological. The topics and the embedding-based frames agreed only weakly, and this low agreement does not indicate a lack of validity, because the two methods describe different properties of the texts. The same vocabulary of replacement, pay and automation appears in macro-societal, corporate and creative-work frames, so that researchers who rely on a single representation may conflate an issue with the way in which it is framed. In addition, the contiguity-constrained periodisation, tested against a permutation null and checked with bootstrap resampling of its boundaries, offers a transparent alternative to dividing time at external events chosen in advance. It also shows how strongly conclusions about discourse phases depend on the robustness of the detected boundaries.

Implications for Practice and Policy

The findings have several practical implications. For human resource management, the visibility of suspicions about AI-washing suggests that transparent communication about workforce decisions, and the alignment of claims about AI-driven productivity with evidence, are important for maintaining organisational trust. For employers and policymakers, the concern about the junior pipeline points to the value of deliberately preserving entry roles similar to apprenticeships, together with mentoring and structured learning on the job, even where AI could perform many entry-level tasks. For career guidance and computing education, the shift in the discourse suggests that students and early-career workers need preparation in judgement, code review, problem formulation and domain knowledge in addition to fluency with AI tools. For recruitment practice, the extended

debate about the use of AI by both candidates and employers calls for assessment designs that remain valid when both parties rely on AI.

Limitations and Future Research

Several limitations qualify these conclusions. First, the users of Hacker News are not representative of the workforce or of all technology workers. They are likely to be predominantly English-speaking, North American and experienced software professionals, and the moderation and ranking of the platform shape what is said. The findings therefore describe practitioner discourse rather than public opinion. Second, the corpus was constructed by keyword retrieval, so topic shares partly reflect the design of the queries. Although the zero-shot filter raised precision to about 85%, some irrelevant comments remain and recall is incomplete, possibly at the expense of long comments that the classifier truncates. Third, the observation window begins in January 2023, one month after the release of ChatGPT, so the very first reactions are not covered. Fourth, the denominator of the intensity measure relies on search counts for the query AI, which also match story titles. A drift towards other terms does not, however, explain the observed rise, because the share of corpus comments that use the term AI itself increased from 83% to 87%.

Fifth, the labels of topics and frames involve interpretive judgement, and the k-means frames are only weakly separated in the original embedding space. Sixth, only the May 2025 phase boundary is robust to resampling, and no causal link to particular events can be established. Seventh, lexicon-based sentiment is unreliable for technical and ironic text and is reported only descriptively. Future research could compare platforms and languages, link the discourse to labour-market indicators over time, and apply supervised classification of stance to distinguish optimistic, critical and resigned positions within each frame.

CONCLUSION

This study examined 34,270 Hacker News comments on generative AI and work posted between January 2023 and September 2026, using a validated procedure for corpus construction, a comparison of four topic models and two layers of clustering. The results show that technology practitioners discuss the labour consequences of AI across the entire employment relationship, from recruitment and entry roles to pay and layoffs, and that their discussion is organised into six discourse frames. Over time, attention within the discussion of AI shifted towards work. A robust change point in May 2025 marks a move towards layoffs, careers and the erosion of entry-level pathways and away from debates about recruitment processes, against a persistent background of macro-societal debate about automation.

The attention of practitioners has thus turned increasingly to how organisations are restructuring work and the routes into it. Employers, educators and policymakers are therefore advised to treat the preservation of entry-level learning opportunities, transparency about AI-related workforce decisions and the redesign of recruitment assessment as priorities. Continued monitoring of practitioner discourse, combined with labour-market data, can help to identify such emerging concerns before they become visible in official statistics.

ACKNOWLEDGEMENT

The author thanks UNITAR International University for supporting the publication of this paper. Artificial intelligence (AI) tools were used for editing and proofreading of this paper; the authors reviewed all content and take full responsibility for it.

REFERENCES

1. Abdelrazek, A., Eid, Y., Gawish, E., Medhat, W., & Hassan, A. (2023). Topic modeling algorithms and applications: A survey. Information Systems, 112, 102131. https://doi.org/10.1016/j.is.2022.102131
2. Acemoglu, D., & Restrepo, P. (2019). Automation and New Tasks: How Technology Displaces and Reinstates Labor. Journal of Economic Perspectives, 33(2), 3–30. https://doi.org/10.1257/jep.33.2.3
3. Acemoglu, D. (2024). The Simple Macroeconomics of AI. NBER Working Paper. https://doi.org/10.3386/w32487

4. Autor, D. H. (2015). Why Are There Still So Many Jobs? The History and Future of Workplace Automation. Journal of Economic Perspectives, 29(3), 3–30. https://doi.org/10.1257/jep.29.3.3
5. Azevedo, M. S., Oliveira, G. S., Ferreira, C. H. G., & Lima, H. C. S. C. (2026). Voices from the community: A multi-dimensional analysis of tech career discourse on Reddit. Proceedings of the 32nd Brazilian Symposium on Multimedia and the Web (WebMedia), in press, 516–525. https://sol.sbc.org.br/index.php/webmedia/article/view/46126
6. Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent Dirichlet allocation. Journal of Machine Learning Research, 3, 993–1022. https://www.jmlr.org/papers/v3/blei03a.html
7. Blei, D. M. (2012). Probabilistic topic models. Communications of the ACM, 55(4), 77–84. https://doi.org/10.1145/2133806.2133826
8. Brynjolfsson, E. (2022). The Turing Trap: The Promise & Peril of Human-Like Artificial Intelligence. Daedalus, 151(2), 272–287. https://doi.org/10.1162/daed_a_01915
9. Brynjolfsson, E., Chandar, B., & Chen, R. (2025). Canaries in the coal mine? Six facts about the recent employment effects of artificial intelligence. Stanford Digital Economy Lab Working Paper. https://digitaleconomy.stanford.edu/publications/canaries-in-the-coal-mine
10. Brynjolfsson, E., Li, D., & Raymond, L. (2025). Generative AI at Work. The Quarterly Journal of Economics, 140(2), 889–942. https://doi.org/10.1093/qje/qjae044
11. Caliński, T., & Harabasz, J. (1974). A dendrite method for cluster analysis. Communications in Statistics, 3(1), 1–27. https://doi.org/10.1080/03610927408827101
12. Campello, R. J. G. B., Moulavi, D., & Sander, J. (2013). Density-Based Clustering Based on Hierarchical Density Estimates. Lecture Notes in Computer Science, 160–172. https://doi.org/10.1007/978-3-642-37456-2_14
13. Davies, D. L., & Bouldin, D. W. (1979). A Cluster Separation Measure. IEEE Transactions on Pattern Analysis and Machine Intelligence, PAMI-1(2), 224–227. https://doi.org/10.1109/tpami.1979.4766909
14. Dell'Acqua, F., McFowland III, E., Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Krayer, L., Candelon, F., & Lakhani, K. R. (2023). Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of Artificial Intelligence on Knowledge Worker Productivity and Quality. SSRN Electronic Journal. https://doi.org/10.2139/ssrn.4573321
15. Demirci, O., Hannane, J., & Zhu, X. (2024). Who is AI Replacing? The Impact of Generative AI on Online Freelancing Platforms. SSRN Electronic Journal. https://doi.org/10.2139/ssrn.4991774
16. DiMaggio, P., Nag, M., & Blei, D. (2013). Exploiting affinities between topic modeling and the sociological perspective on culture: Application to newspaper coverage of U.S. government arts funding. Poetics, 41(6), 570–606. https://doi.org/10.1016/j.poetic.2013.08.004
17. Dieng, A. B., Ruiz, F. J. R., & Blei, D. M. (2020). Topic Modeling in Embedding Spaces. Transactions of the Association for Computational Linguistics, 8, 439–453. https://doi.org/10.1162/tacl_a_00325
18. Díaz Ferreyra, N. E., Gurupathi, M. S., Codabux, Z., Arachchilage, N., & Scandariato, R. (2026). Security concerns in generative AI coding assistants: Insights from online discussions on GitHub Copilot. arXiv preprint arXiv:2604.08352. https://doi.org/10.48550/arXiv.2604.08352
19. Egger, R., & Yu, J. (2022). A Topic Modeling Comparison Between LDA, NMF, Top2Vec, and BERTopic to Demystify Twitter Posts. Frontiers in Sociology, 7. https://doi.org/10.3389/fsoc.2022.886498
20. Eloundou, T., Manning, S., Mishkin, P., & Rock, D. (2024). GPTs are GPTs: Labor market impact potential of LLMs. Science, 384(6702), 1306–1308. https://doi.org/10.1126/science.adj0998
21. Evans, J. A., & Aceves, P. (2016). Machine Translation: Mining Text for Social Theory. Annual Review of Sociology, 42(1), 21–50. https://doi.org/10.1146/annurev-soc-081715-074206
22. Felten, E., Raj, M., & Seamans, R. (2021). Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. Strategic Management Journal, 42(12), 2195–2217. https://doi.org/10.1002/smj.3286
23. Fiesler, C., & Proferes, N. (2018). "Participant" Perceptions of Twitter Research Ethics. Social Media + Society, 4(1). https://doi.org/10.1177/2056305118763366
24. Frey, C. B., & Osborne, M. A. (2017). The future of employment: How susceptible are jobs to computerisation?. Technological Forecasting and Social Change, 114, 254–280. https://doi.org/10.1016/j.techfore.2016.08.019

25. Grimmer, J., & Stewart, B. M. (2013). Text as Data: The Promise and Pitfalls of Automatic Content Analysis Methods for Political Texts. Political Analysis, 21(3), 267–297. https://doi.org/10.1093/pan/mps028
26. Grootendorst, M. (2022). BERTopic: Neural topic modeling with a class-based TF-IDF procedure. arXiv preprint arXiv:2203.05794. https://doi.org/10.48550/arXiv.2203.05794
27. Hui, X., Reshef, O., & Zhou, L. (2024). The Short-Term Effects of Generative Artificial Intelligence on Employment: Evidence from an Online Labor Market. Organization Science, 35(6), 1977–1989. https://doi.org/10.1287/orsc.2023.18441
28. Humlum, A., & Vestergaard, E. (2025). Large Language Models, Small Labor Market Effects. SSRN Electronic Journal. https://doi.org/10.2139/ssrn.5219933
29. Hutto, C., & Gilbert, E. (2014). VADER: A Parsimonious Rule-Based Model for Sentiment Analysis of Social Media Text. Proceedings of the International AAAI Conference on Web and Social Media, 216–225. https://doi.org/10.1609/icwsm.v8i1.14550
30. Kellogg, K. C., Valentine, M. A., & Christin, A. (2020). Algorithms at Work: The New Contested Terrain of Control. Academy of Management Annals, 14(1), 366–410. https://doi.org/10.5465/annals.2018.0174
31. Koonchanok, R., Pan, Y., & Jang, H. (2024). Public attitudes toward ChatGPT on Twitter: Sentiments, topics, and occupations. Social Network Analysis and Mining, 14(1). https://doi.org/10.1007/s13278-024-01260-7
32. Laurer, M., van Atteveldt, W., Casas, A., & Welbers, K. (2023). Building efficient universal classifiers with natural language inference. arXiv preprint arXiv:2312.17543. https://doi.org/10.48550/arXiv.2312.17543
33. Laurer, M., van Atteveldt, W., Casas, A., & Welbers, K. (2024). Less Annotating, More Classifying: Addressing the Data Scarcity Issue of Supervised Machine Learning with Deep Transfer Learning and BERT-NLI. Political Analysis, 32(1), 84–100. https://doi.org/10.1017/pan.2023.20
34. Lee, D. D., & Seung, H. S. (1999). Learning the parts of objects by non-negative matrix factorization. Nature, 401(6755), 788–791. https://doi.org/10.1038/44565
35. Leiter, C., Zhang, R., Chen, Y., Belouadi, J., Larionov, D., Fresen, V., & Eger, S. (2024). ChatGPT: A meta-analysis after 2.5 months. Machine Learning with Applications, 16, 100541. https://doi.org/10.1016/j.mlwa.2024.100541
36. McInnes, L., Healy, J., & Melville, J. (2018). UMAP: Uniform manifold approximation and projection for dimension reduction. arXiv preprint arXiv:1802.03426. https://doi.org/10.48550/arXiv.1802.03426
37. Miyazaki, K., Murayama, T., Uchiba, T., An, J., & Kwak, H. (2024). Public perception of generative AI on Twitter: an empirical study based on occupation and usage. EPJ Data Science, 13(1). https://doi.org/10.1140/epjds/s13688-023-00445-y
38. Nelson, L. K. (2020). Computational Grounded Theory: A Methodological Framework. Sociological Methods & Research, 49(1), 3–42. https://doi.org/10.1177/0049124117729703
39. Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. Science, 381(6654), 187–192. https://doi.org/10.1126/science.adh2586
40. Reimers, N., & Gurevych, I. (2019). Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), 3980–3990. https://doi.org/10.18653/v1/d19-1410
41. Rousseeuw, P. J. (1987). Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. Journal of Computational and Applied Mathematics, 20, 53–65. https://doi.org/10.1016/0377-0427(87)90125-7
42. Röder, M., Both, A., & Hinneburg, A. (2015). Exploring the Space of Topic Coherence Measures. Proceedings of the Eighth ACM International Conference on Web Search and Data Mining, 399–408. https://doi.org/10.1145/2684822.2685324
43. Ward, J. H. (1963). Hierarchical Grouping to Optimize an Objective Function. Journal of the American Statistical Association, 58(301), 236–244. https://doi.org/10.1080/01621459.1963.10500845
44. Yu, S., & Moon, T. (2026). Who will become the next senior? How generative AI erodes the development pathway in software engineering. arXiv preprint arXiv:2607.17067 (accepted to AIES 2026). https://doi.org/10.48550/arXiv.2607.17067

45. Yüce, P., Dai, X., Owens, R., & Elmas, T. (2026). ChatGPT vs teachers vs students: Large-scale analysis of generative AI discourse in education communities on Reddit. arXiv preprint arXiv:2605.17712. https://doi.org/10.48550/arXiv.2605.17712
46. Zimmer, M. (2010). "But the data is already public": on the ethics of research in Facebook. Ethics and Information Technology, 12(4), 313–325. https://doi.org/10.1007/s10676-010-9227-5