From Artificial Intelligence as Automation to Generative AI as a Work Tool: A Text Network Analysis of Public Discourse Before and After ChatGPT

Katarina Kostelić
Juraj Dobrila University of Pula
Faculty of Informatics
Ulica Alda Negrija 6
52100 Pula, Croatia
katarina.kostelic@unipu.hr

Cristina Fleșeriu
Babeș-Bolyai University
Faculty of Business
Horea Street 7
400000 Cluj-Napoca, Romania
cristina.fleseriu@ubbcluj.ro

Vlad Bocăneț
Technical University of Cluj-Napoca
Faculty of Industrial Engineering, Robotics and Production Management
B-dul Muncii nr. 103-105
400000 Cluj-Napoca, Romania
vlad.bocanet@tcm.utcluj.ro

Abstract. *This paper examines how public-facing discourse on AI and work was reframed after the public release of ChatGPT in November 2022. Using text network analysis of articles from The Conversation, it shows that the debate shifted from automation-centred narratives toward a more differentiated view of generative AI as a practical work tool. The post-ChatGPT discourse connects work with education, writing, organisational adoption, productivity, copyright, language models and governance concerns. Sentiment patterns indicate thematic ambivalence. The paper contributes by demonstrating how text network analysis captures changing public meanings of AI and how work transformation is discussed across topic areas.*

Keywords. generative AI, ChatGPT, intelligent systems, text network analysis, public discourse

1 Introduction

Artificial intelligence has long been part of public debate, but the public launch of ChatGPT in November 2022 (OpenAI, 2022) created a new point of reference for discussions of generative AI. Although large generative language models had already been examined before ChatGPT, particularly in relation to the automatic production of textual and semantic artefacts (Floridi & Chiriatti, 2020), ChatGPT contributed to the wider public visibility of this class of systems by offering a conversational interface that could be used by non-specialist users. Since then, generative AI has become increasingly visible in discussions about work, education, productivity, creativity, regulation and professional knowledge (Dwivedi et al., 2023). However, it remains less clear how this visibility changed the semantic organisation of public-facing AI-and-work discourse.

The discussion is relevant because AI plays an important role in transforming work, learning, communication, and organisational processes. It also extends earlier AI-and-work debates, which often focused on automation and job displacement, toward questions of task change, augmentation and human–AI interaction (Huang & Rust, 2018; Warner & Wäger, 2019). These developments raise economic, organisational and educational questions, but also discursive ones. Public understanding of AI is shaped through media articles, expert commentary, educational explainers and research communication, which influence how intelligent technologies are interpreted in relation to work transformation across different domains.

This paper examines how public-facing discourse represents the role of artificial intelligence in transforming work processes across different domains before and after the public launch of ChatGPT. The study asks:

RQ1: How has the central vocabulary of public discourse on artificial intelligence changed in the context of transforming work processes across different domains before and after the public launch of ChatGPT?

RQ2: Which themes became structurally prominent in the post-ChatGPT discourse on AI and work transformation?

RQ3: How is sentiment distributed across thematic communities in the later stage of the discourse?

Rather than measuring the actual AI-related transformation of work processes across different domains, the study analyses how these changes are communicated, interpreted and semantically organised in expert-oriented public discourse. The study applies text network analysis, modelling words as nodes and

co-occurrences as ties to identify structurally important terms. In this sense, the study is exploratory: it maps changes in the vocabulary, thematic structure and sentiment of AI-and-work discourse.

2 Literature review

Research on AI and work has examined the implications of intelligent technologies for jobs, tasks, skills, decision-making and organisational processes. A recurring theme in AI-and-work research concerns automation, technological unemployment and the substitution of human labour, particularly in relation to robotics and machine intelligence. However, AI-and-work research also includes more differentiated perspectives. Huang and Rust (2018) argue that AI job replacement proceeds at the task level before it reaches whole occupations, while Raisch and Krakowski (2021) conceptualise automation and augmentation as interdependent dimensions of AI use in organisations. Zirar et al. (2023) further highlight worker–AI coexistence, including issues of distrust, skills, reskilling and augmentation.

Large language models and generative AI add a further dimension to this debate because they are directly relevant to cognitive, communicative and knowledge-intensive work. Studies of occupational exposure suggest that LLMs may affect a substantial share of work tasks, especially when integrated into software systems, although exposure should not be equated with direct replacement (Eloundou et al., 2024; Gmyrek et al., 2023, 2025). Empirical studies also point to heterogeneous effects. Noy and Zhang (2023) found productivity and quality gains in professional writing tasks, while Dell'Acqua et al. (2026) show that generative AI can improve performance on some tasks but reduce it on others, depending on task characteristics and the limits of the technology.

Generative AI has also become relevant beyond workplace productivity narrowly defined. ChatGPT and related systems have been discussed in relation to research, education, organisations, policy, ethics, privacy, misinformation, bias and security (Budhwar et al., 2023; Dwivedi et al., 2023). In education, studies have identified both possible benefits, such as feedback, learning support and enhanced interaction, and concerns related to academic integrity, overreliance, privacy, truthfulness and assessment security (Tlili et al., 2023; Dempere et al., 2023). Broader discussions of generative AI also point to issues of bias, stereotypes and trust in AI-generated content (Ferrara, 2023; Vaccari & Chadwick, 2020).

This shows that AI and work can be examined through several overlapping lenses: automation, task exposure, augmentation, organisational adoption, education, knowledge work, ethics and governance. However, less is known about how these issues are connected in public-facing, research-informed discourse. Much existing research examines occupational exposure, productivity effects, organisational implications, educational uses or public attitudes. Fewer studies examine how AI-and-work discourse is semantically organised, or how its structure differs before and after the public release of ChatGPT. This is the gap addressed in the present paper. Rather than measuring public opinion or the actual labour-market effects of AI, the study examines how public-facing discourse organises the meaning of AI and work before and after ChatGPT.

3 Methods

3.1 Research Design and Corpus

This study applies a text-as-data approach to examine public-facing discourse on artificial intelligence and work before and after the public launch of ChatGPT. Such methods are useful for detecting patterns across a corpus, but their results require substantive interpretation rather than fully automated reading (Grimmer & Stewart, 2013; Welbers et al., 2017).

The corpus consists of articles retrieved from The Conversation using the search query *artificial intelligence work*. This source was selected because it publishes research-based commentary written by academics and researchers for wider audiences, making it an intermediate discourse space between academic literature and general public debate. Since The Conversation allows redistribution under a Creative Commons Attribution–NoDerivatives licence, article texts were used only as input for computational analysis and results are reported in aggregate form.

The scraping procedure collected article links, titles, authors, publication dates, and article bodies. Records were retained if they contained a URL, title and non-empty article body; texts of 500 characters or fewer were excluded as too short for meaningful analysis. The final corpus contained 200 articles. Documents were divided by the public launch date of ChatGPT, 30 November 2022, into a pre-ChatGPT period (78 articles) and a post-ChatGPT period (122 articles, ending in April 2026). Texts were cleaned and tokenised into individual words. Standard English stop words were removed, while negation terms were retained. Stemming and lemmatisation were not applied in order to preserve interpretability of terms in tables, figures and community labels. The analysis was conducted in R version 4.4.2.

3.2 Co-occurrence Network Construction

Word co-occurrence networks were constructed to analyse how terms are structurally connected in the corpus. In such networks, words are represented as nodes, while ties indicate co-occurrence within a defined textual context. This approach captures relationships between terms rather than only their

individual frequencies (Danowski, 1993; Diesner & Carley, 2005; Drieger, 2013).

Each token was assigned a sequential position within its document, and word pairs were created when two terms appeared within a ten-word window. Edge weights represent the number of times two words co-occurred within this window. Separate undirected weighted networks were constructed for the pre- and post-ChatGPT periods. Only word pairs with at least ten co-occurrences were retained, and the largest connected component was extracted from each network.

The analysis reports word frequencies, normalized degree, and normalized betweenness. Normalized degree indicates the share of possible direct connections a term has, while normalized betweenness indicates the extent to which a term bridges other parts of the network (Freeman, 1978; Wasserman & Faust, 1994). In addition, thematic communities were extracted with Louvain clustering applied to the 2025 network (Blondel et al., 2008). For community and sentiment analyses, a 2025 network was used as the most recent complete post-ChatGPT calendar year. Communities were interpreted as thematic clusters, and labels were assigned by inspecting the most central and substantively meaningful words in each community.

3.3 Sentiment Analysis

Lexicon-based sentiment analysis was used to examine whether thematic communities in the 2025 network differed in affective vocabulary. The Bing lexicon was used to classify sentiment-bearing words as positive or negative (Hu & Liu, 2004; Silge & Robinson, 2017). Sentiment is treated as an indicative measure of affective vocabulary, not as a definitive measure of article tone, because of the well-known limitations (not capturing irony, negation scope, domain-specific meanings, argument structure or author stance; Liu, 2012).

4 Results

4.1 Most Frequent Terms Before and After ChatGPT

The comparison of the most frequent terms indicates a shift in the vocabulary of AI-and-work discourse before and after ChatGPT (Table 1). The pre-ChatGPT period is marked by a vocabulary associated with general AI, human intelligence, jobs, data, technology companies and future-oriented technological change. In contrast, the post-ChatGPT period gives greater visibility to terms associated with practical use, including tools, students, research, workers and tasks, alongside the emergence of ChatGPT-specific and generative AI vocabulary.

Table 1. Most frequent terms before and after ChatGPT

Rank	Before ChatGPT: word	Before: n	After ChatGPT: word	After: n
1	ai	447	ai	1654
2	not	298	not	437
3	human	254	people	321
4	data	246	data	267
5	people	232	human	267
6	intelligence	230	technology	230
7	jobs	210	research	216
8	google	202	workers	206
9	artificial	172	time	203
10	world	168	tools	182

At the level of frequent terms, this suggests a movement from AI as a technology associated with transformation and possible replacement of work toward AI as a tool discussed in relation to workflows, classrooms, writing practices, research activities and organisational routines.

4.2 Structural Changes in the Co-occurrence Networks

The co-occurrence networks provide additional insight into how AI-and-work discourse is structurally organised. As shown in Table 2, the pre-ChatGPT network is organised around a relatively concentrated semantic core. The most central terms include *ai*, *jobs*, *data*, *intelligence*, *human*, *people*, *google* and *workers*, suggesting a discourse structured around AI as a technological and labour-market issue. The presence of *jobs*, *workers*, *machine*, *robots*, *skills* and *future* indicates an emphasis on automation, employment change and technological transformation.

Table 2. Most central terms in the pre- and post-ChatGPT co-occurrence networks

Metric	Pre-ChatGPT	Post-ChatGPT
Highest normalized degree terms	ai (0.48)	ai (0.9)
	jobs (0.23)	not (0.09)
	data (0.19)	people (0.06)
	intelligence (0.18)	human (0.05)
	not (0.18)	research (0.05)
	human (0.18)	data (0.05)
	people (0.15)	time (0.04)
	google (0.12)	chatgpt (0.04)
	workers (0.12)	students (0.04)
	artificial (0.11)	workers (0.04)
Highest norm	ai (0.5)	ai (0.92)
	data (0.18)	not (0.08)
	people (0.15)	simply (0.06)

	workers (0.14)	data (0.04)
	google (0.13)	algorithms (0.04)
	jobs (0.13)	chatgpt (0.04)
	intelligence (0.11)	academic (0.04)
	job (0.09)	business (0.04)
	time (0.09)	digital (0.04)
	computer (0.07)	trained (0.04)

The post-ChatGPT network is broader and more differentiated. The degree of *ai* increases strongly, indicating that the term becomes connected to a wider range of semantic contexts. Newer structurally visible terms, including *chatgpt*, *students*, *tools*, *language*, *productivity*, *generative*, *models*, *companies*, *employees* and *decision*, suggest a discourse more closely connected to practical use, education, organisational adoption, workplace processes and decision-making.

This supports the interpretation that public-facing discourse shifted from AI as a future-oriented automation technology toward generative AI as a practical work and learning tool. The post-ChatGPT network does not replace the automation narrative entirely; instead, it embeds it within a broader set of debates about institutional adoption, human-AI collaboration, education, productivity and governance.

4.3 Thematic Communities in the 2025 Network

The 2025 co-occurrence network reveals a differentiated post-ChatGPT discourse. Louvain community detection identified 19 communities (Table 3, Fig. 1). These clusters of words are more densely connected to one another than to the rest of the network, so their interpretation combines the algorithmic grouping of terms with qualitative reading of the representative words.

Table 3. Louvain communities in the 2025 co-occurrence network

Community	Representative words	Interpretive label
1	ai, systems, employees, literacy, training, generative, job, creative	Workplace adoption and AI literacy
2	future, time, intelligence, artificial, participants, world	Future-oriented societal framing
3	students, tasks, writing, education, university, public, thinking	Education and academic writing
4	people, not, workers, college, technology, knowledge	Workers, knowledge and technological change
5	research, review, peer, consciousness, social, academic	Research, peer review and academic discourse
6	learning, bio, machine, signal, learn, skills, algorithm	Machine learning and technical systems
7	tools, found, design, chatgpt, digital, process, business, software	Tools, business processes and digital transformation
8	data, human, quantum, collaboration, video, scale, chips, judgement	Data, human judgement and technical infrastructure
9	jobs, roles, entry, level, low, positions, replace, robots, income	Labour-market automation and job restructuring
10	language, models, text, produce, word, native	Language models and text generation
11	information, google, company, million, prompt, users, active	Platform use, prompting and user adoption
12	companies, copyright, publishers, authors, industry, tech	Copyright, publishing and industry conflict
13	calls, complex, chickadees, mountain, chick, call, simple	Non-human communication theme
14	health, monitoring, care	Health monitoring and care
15	prompting, rhetorical, method	Prompting methods and interaction design
16	ability, belief, person	Human ability and belief
17	capacity, fitness	Capacity and performance
18	loss, mourned	Loss and mourning
19	reform, week, roundtable	Policy discussion

Note: Representative words are shown as returned by the community-level summary. Interpretive labels were assigned by the authors based on the most central and substantively meaningful words in each Louvain community. Smaller communities should be

interpreted cautiously because they may reflect specific article-level topics rather than broad discourse themes.

The largest and most substantively coherent communities show that AI-and-work discourse had become multidimensional by 2025. Several communities concern workplace adoption, AI literacy, business processes and digital transformation. These clusters frame generative AI as a practical tool entering organisational systems, employee training, software processes and workplace routines. This supports the interpretation that post-ChatGPT discourse increasingly treats AI as work infrastructure rather than only as a future automation technology.

A second group of communities links AI to education, academic writing and knowledge production. The presence of communities centred on students, writing, universities, research and peer review shows that the discourse extends beyond current workplaces. It also concerns the preparation of future knowledge workers and the transformation of academic practices. This is consistent with the broader visibility of ChatGPT in debates about learning, authorship, academic integrity and scholarly evaluation.

The network also retains technological and labour-market themes. Communities related to machine learning, algorithms, data, language models and text generation show that public-facing discourse still includes explanations of how AI systems work. At the same time, the labour-market automation community shows that job replacement, entry-level roles, income and robots remain part of the discussion. However, this theme appears as one community among many, rather than as the dominant organising frame.

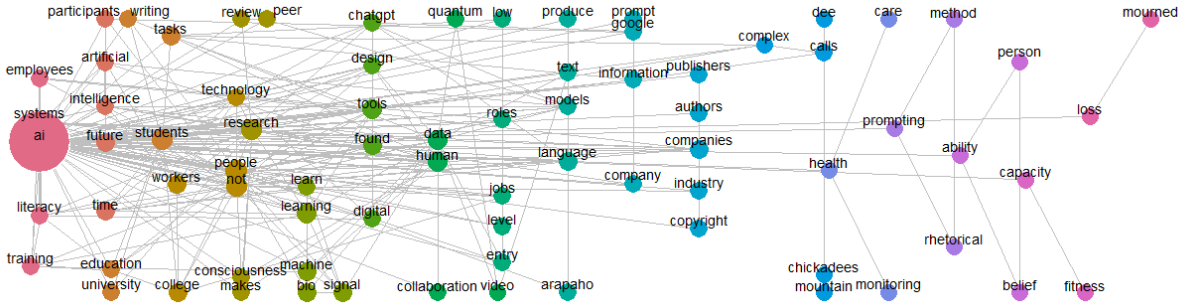

Figure 1. Thematic communities in the giant component of 2025 network

Notes: The graph shows the largest connected component of the 2025 co-occurrence network, using up to five top words per community ranked by normalized degree. For readability, the force-directed Fruchterman-Reingold layout was horizontally constrained by Louvain community; node proximity should therefore be interpreted visually rather than metrically.

Finally, several communities capture newer conflicts and practices associated with generative AI, including prompting, platform use, copyright, publishing and industry conflict. Smaller communities, such as those related to health monitoring, non-human communication or mourning, appear to reflect specific article-level topics and are therefore not considered as major discourse-level themes.

The 2025 communities show that the post-ChatGPT discourse is differentiated and multidimensional. AI is no longer discussed primarily as a general technological force or job-replacement threat. Instead, it is embedded in several interconnected debates about organisational adoption, education, knowledge work, language models, labour-market restructuring, copyright, prompting and digital transformation. This supports the broader interpretation of a shift from AI as future-oriented automation toward generative AI as an everyday work and learning tool.

4.4 Sentiment as Thematic Ambivalence

The community-level sentiment analysis for 2025 shows that AI-and-work discourse is affectively uneven across thematic areas (Fig. 2; community numbers correspond to the communities reported in Table 3). Communities related to future-oriented societal change, tools and business processes, education and academic writing, workplace adoption, AI literacy, and technical infrastructure are ambivalent, but predominantly positive.

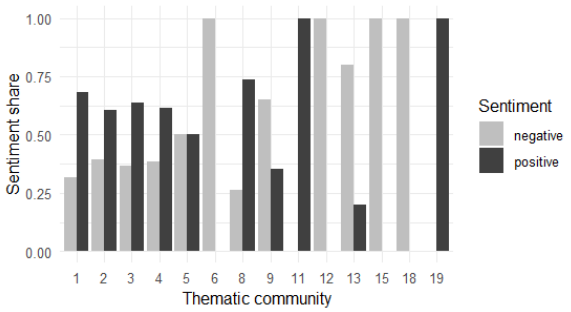

Figure 2. Sentiment distribution by 2025 thematic community

In contrast, communities related to research and peer review, workers and technological change, and smaller risk-oriented themes show more negative or more ambivalent sentiment patterns. These differences

suggest that post-ChatGPT discourse does not form a single positive or negative narrative. AI is framed more positively when discussed as a tool for innovation, productivity, education or digital transformation, and more negatively when connected to labour uncertainty, institutional tensions, legal conflict or broader social risk.

5 Discussion

The findings suggest that public-facing discourse on AI and work changed substantially after the release of ChatGPT. The change is both lexical and structural: before ChatGPT, the discourse was organised mainly around general AI, jobs, data, machines, robots and large technology companies, while the post-ChatGPT discourse became broader and more differentiated, connecting AI to tools, students, writing, productivity, language models, organisational adoption, copyright, prompting and business processes.

This supports the interpretation of a shift from AI as imagined automation toward generative AI as a practical work tool. Earlier debates on AI and work often centred on automation, job replacement and technological unemployment, while more recent research emphasises task-level effects, augmentation, skills, organisational processes and human–AI collaboration (Huang & Rust, 2018; Raisch & Krakowski, 2021; Noy & Zhang, 2023). The present findings reflect this transition: employment concerns remain visible, but they are embedded in a wider discussion about how generative AI is used, adopted, governed and integrated into everyday work.

The post-ChatGPT network also confirms the multidisciplinary and institutionally distributed character of generative AI, which previous studies have linked to research, education, organisations, policy and society more broadly (Dwivedi et al., 2023; Budhwar et al., 2023). Therefore, the public meaning of generative AI is not organised around a single question, such as whether AI will replace jobs, but across several areas in which work is performed, taught, evaluated, regulated and reorganised.

The visibility of education and academic writing is particularly relevant because generative AI became directly accessible for writing, learning, coding, summarising and explanation tasks. Previous studies similarly describe ChatGPT in education as ambivalent: it can support feedback, personalised learning and productivity, but also raises concerns about academic integrity, overreliance, privacy and critical thinking (Tlili et al., 2023; Su & Yang, 2023). The present findings show that this educational discourse is part of AI-and-work discourse, because education shapes future knowledge work and professional expectations about AI use.

The sentiment results further show that post-ChatGPT discourse is neither simply optimistic nor pessimistic. Tool-oriented, education-related, workplace adoption and digital transformation communities are predominantly positive, while communities related to labour uncertainty, research and peer review, legal conflict and social risk are more negative or ambivalent. This is consistent with literature that describes generative AI as combining productivity, business, education and research opportunities with risks related to misinformation, bias, privacy, security, worker uncertainty and ethics (Dwivedi et al., 2023; Budhwar et al., 2023; Ferrara, 2023). The contribution of this study is to show that these framings are theme-dependent rather than evenly distributed across discourse.

Methodologically, the study shows the usefulness of combining frequency analysis, co-occurrence networks, community detection and lexicon-based sentiment analysis. In this case, AI became more frequent and more embedded in a wider semantic space after ChatGPT. Community detection revealed thematic differentiation within this space, while sentiment analysis added an affective layer to the interpretation.

The findings also have practical implications. For organisations, the discourse suggests that generative AI adoption is now discussed less as an abstract automation threat and more through AI literacy, training, workflow redesign, output evaluation and responsible use. This aligns with the European policy context, where the AI Act addresses trustworthy and risk-based AI governance (Regulation (EU) 2024/1689). For educators, the results support treating AI literacy as preparation for knowledge work, while the visibility of copyright, platform use and governance themes points to the need for sector-specific guidance in concrete work settings.

The study contributes by connecting AI-and-work research with public discourse analysis. It does not measure occupational exposure, productivity effects or adoption outcomes directly, but examines how these issues become publicly connected in research-informed media discourse. This matters because public-facing discourse shapes expectations, concerns and institutional responses to emerging technologies.

The findings suggest that the public launch of ChatGPT did not simply intensify existing discourse on AI and work, but marked a turning point in the broader public visibility of generative AI. AI remains connected to automation and labour-market risk, but it is increasingly framed as a practical work tool embedded in writing, learning, research, productivity, organisational transformation and digital work.

6 Conclusion

This paper examined changes in public discourse on artificial intelligence and work before and after the public launch of ChatGPT. Using articles from The Conversation, the study applied text preprocessing, word frequency analysis, co-occurrence networks,

centrality measures, community detection and sentiment analysis.

The results suggest that post-ChatGPT discourse is broader and more differentiated than earlier AI-and-work discourse. The 2025 thematic network further shows that AI discourse had differentiated into multiple fields, including workplace adoption, education, labour-market automation, copyright, prompting and business transformation.

The main contribution of the paper is methodological and interpretive. It shows that text network analysis can reveal how public discourse is structurally organised. The results also show that sentiment is not evenly distributed across AI discourse: themes related to productivity, education and adoption are more positively framed, while themes related to labour-market displacement, institutional tensions and legal conflict are more negative or ambivalent.

The findings should be interpreted with limitations. The corpus is limited to articles from The Conversation, which is suitable for analysing research-informed public commentary but does not represent public opinion, social media, mainstream journalism or academic literature as a whole. The use of a single search query may also have shaped the corpus, and the post-ChatGPT period contains more articles than the pre-ChatGPT period. Methodologically, co-occurrence does not prove conceptual agreement, causality or author intention; community interpretation remains researcher-assigned; and lexicon-based sentiment analysis cannot fully capture irony, negation, domain-specific meanings or whether a concern is endorsed or merely reported. Some smaller 2025 communities may also reflect article-specific themes rather than broad discourse-level patterns.

Future research could expand the corpus to other media outlets, compare research-informed media with mainstream journalism or social media, and use multiple search strategies. Cross-linguistic comparisons would also be valuable, since AI adoption, labour-market concerns and educational debates may differ across national and institutional contexts. Methodologically, future studies could combine co-occurrence networks with topic modelling, transformer-based embeddings or qualitative close reading.

Acknowledgments

This work was funded by the EU NextGeneration under the Juraj Dobrila University of Pula institutional research project number IIP_UNIPU_010133. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the Juraj Dobrila University of Pula.

References

Blondel, V. D., Guillaume, J.-L., Lambiotte, R., & Lefebvre, E. (2008). Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2008(10), P10008. https://doi.org/10.1088/1742-5468/2008/10/P10008

Budhwar, P., Chowdhury, S., Wood, G., Aguinis, H., Bamber, G. J., Beltran, J. R., Boselie, P., Cooke, F. L., Decker, S., DeNisi, A., Dey, P. K., Guest, D., Knoblich, A. J., Malik, A., Paauwe, J., Papagiannidis, S., Patel, C., Pereira, V., Ren, S., Rogelberg, S., Saunders, M.N.K., Tung, R.L., Varma, A. (2023). Human resource management in the age of generative artificial intelligence: Perspectives and research directions on ChatGPT. *Human Resource Management Journal*, 33(3), 606–659. https://doi.org/10.1111/1748-8583.12524

Danowski, J. A. (1993). Network analysis of message content. *Progress in Communication Sciences*, 12, 197–221.

Dell'Acqua, F., McFowland, E., Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Krayer, L., Candelon, F., & Lakhani, K. R. (2026). Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality. *Organization Science*. https://doi.org/10.1287/orsc.2025.21838

Dempere, J., Modugu, K., Hesham, A., & Ramasamy, L. K. (2023). The impact of ChatGPT on higher education. *Frontiers in Education*, 8, 1206936. https://doi.org/10.3389/feduc.2023.1206936

Diesner, J., & Carley, K. M. (2005). Revealing social structure from texts: Meta-matrix text analysis as a novel method for network text analysis. In V. K. Narayanan & D. J. Armstrong (Eds.), *Causal mapping for research in information technology* (pp. 81–108). IGI Global.

Drieger, P. (2013). Semantic network analysis as a method for visual text analytics. *Procedia - Social and Behavioral Sciences*, 79, 4–17. https://doi.org/10.1016/j.sbspro.2013.05.053

Dwivedi, Y. K., Kshetri, N., Hughes, L., Slade, E. L., Jeyaraj, A., Kar, A. K., Baabdullah, A. M., Koohang, A., Raghavan, V., Ahuja, M., et al. (2023). So what if ChatGPT wrote it? Multidisciplinary perspectives on opportunities, challenges and implications of generative conversational AI for research, practice and policy. *International Journal of Information Management*, 71, 102642. https://doi.org/10.1016/j.ijinfomgt.2023.102642

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

Ferrara, E. (2023). Should ChatGPT be biased? Challenges and risks of bias in large language models. *First Monday*, 28(11). https://doi.org/10.5210/fm.v28i11.13346

Floridi, L., & Chiriatti, M. (2020). GPT-3: Its nature, scope, limits, and consequences. *Minds and Machines*, 30(4), 681–694. https://doi.org/10.1007/s11023-020-09548-1

Freeman, L. C. (1978). Centrality in social networks: Conceptual clarification. *Social Networks*, 1(3), 215–239. https://doi.org/10.1016/0378-8733(78)90021-7

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

Gmyrek, P., Berg, J., & Bescond, D. (2023). Generative AI and jobs: A global analysis of potential effects on job quantity and quality (ILO Working Paper No. 96). *International Labour Organization*. https://doi.org/10.54394/FHEM8239

Gmyrek, P., Winkler, H., & Garganta, S. (2025). Generative AI and jobs: A refined global index of occupational exposure (ILO Working Paper No. 140). *International Labour Organization*. https://www.ilo.org/publications/generative-ai-and-jobs-refined-global-index-occupational-exposure

Hu, M., & Liu, B. (2004). Mining and summarizing customer reviews. In Kohavi, R., Gehrke, J., DuMouchel, W. and Ghosh, J (Eds), *Proceedings of the Tenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 168–177). Association for Computing Machinery. https://doi.org/10.1145/1014052.1014073

Huang, M.-H., & Rust, R. T. (2018). Artificial intelligence in service. *Journal of Service Research*, 21(2), 155–172. https://doi.org/10.1177/1094670517752459

Liu, B. (2012). *Sentiment analysis and opinion mining*. Morgan & Claypool.

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

OpenAI. (2022, November 30). Introducing ChatGPT. https://openai.com/index/chatgpt/

Raisch, S., & Krakowski, S. (2021). Artificial intelligence and management: The automation–augmentation paradox. *Academy of Management Review*, 46(1), 192–210. https://doi.org/10.5465/amr.2018.0072

Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending Regulations and Directives (Artificial Intelligence Act). (2024). Official Journal of the European Union, L 2024/1689. https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng

Silge, J., & Robinson, D. (2017). *Text mining with R: A tidy approach*. O'Reilly Media.

Su, J., & Yang, W. (2023). Unlocking the power of ChatGPT: A framework for applying generative AI in education. *ECNU Review of Education*, 6(3), 355–366. https://doi.org/10.1177/20965311231168423

The Conversation. (2026). The Conversation. https://theconversation.com

Tlili, A., Shehata, B., Adarkwah, M. A., Bozkurt, A., Hickey, D. T., Huang, R., & Agyemang, B. (2023). What if the devil is my guardian angel: ChatGPT as a case study of using chatbots in education. *Smart Learning Environments*, 10, 15. https://doi.org/10.1186/s40561-023-00237-x

Vaccari, C., & Chadwick, A. (2020). Deepfakes and disinformation: Exploring the impact of synthetic political video on deception, uncertainty, and trust in news. *Social Media + Society*, 6(1). https://doi.org/10.1177/2056305120903408

Warner, K. S. R., & Wäger, M. (2019). Building dynamic capabilities for digital transformation: An ongoing process of strategic renewal. *Long Range Planning*, 52(3), 326–349. https://doi.org/10.1016/j.lrp.2018.12.001

Wasserman, S., & Faust, K. (1994). *Social network analysis: Methods and applications*. Cambridge University Press.

Welbers, K., Van Atteveldt, W., & Benoit, K. (2017). Text analysis in R. *Communication Methods and Measures*, 11(4), 245–265. https://doi.org/10.1080/19312458.2017.1387238

Zirar, A., Ali, S. I., & Islam, N. (2023). Worker and workplace Artificial Intelligence (AI) coexistence: Emerging themes and research agenda. *Technovation*, 124, 102747. https://doi.org/10.1016/j.technovation.2023.102747