

Research Article

In GPT we trust: perceptions of the future of work with artificial intelligence on online forums

Szabina Fodor^{*}, Bettina Boncz

Department of Computer Science, Corvinus University of Budapest, Budapest, 1093, Hungary

ARTICLE INFO

Keywords:

Artificial intelligence
Topic modeling
Sentiment and emotional analysis
Social media
Reddit
Public perception

ABSTRACT

This study analyzes over 10 years of Reddit discussions on artificial intelligence (AI)'s impact on labor market using a state-of-the-art pre-trained robust optimized bidirectional encoder representations from transformers (BERT) model, or RoBERTa model, for sentiment analysis and emotional detection. Additionally, the BERTopic method is employed to identify the discussed topics. The results show generally mixed sentiments towards AI's impact on the labor market. Fueled by influencers and media sensationalism, negative emotions—fear, anger, and disgust—stem from concerns about job replacement, economic instability, and societal change. However, the release of chat generative pre-trained transformer's (ChatGPT) in 2022 has created a more positive outlook, with users expressing surprise and optimism about generative AI's capabilities. Discussions have shifted from fears of complete automation to the potential for augmentation. Emotional dynamics have been shaped by events such as the COVID-19 pandemic and the rise of generative AI, underscoring AI's dual role as both a threat and an opportunity.

1. Introduction

Artificial intelligence (AI), a machine or algorithm that can imitate human skills or abilities (Sheikh et al., 2023), has already left the pages of science fiction literature and has become closer to reality. Today, with the rise of large language models (LLMs), we can imitate some human intelligence with AI. However, the question remains whether this comes from excelling at pattern recognition or real human-like machine intelligence (Ilić and Gignac, 2024). This question is especially difficult to answer as human intelligence varies on an individual level and lacks a standard definition (Gignac and Szodorai, 2024). In addition, most AI applications specialize in performing specific tasks (e.g., generating text or images, analyzing visual inputs, or big datasets) instead of manifesting universal skills and intelligence (Kshetri et al., 2024; Kutyaurov et al., 2023; Revilla-León et al., 2023). Perhaps no single general AI application is available yet. Still, the space of current AI development is already remarkable (Bubeck et al., 2023; Triguero et al., 2024), including its impact on our daily lives. The emergence of chat generative pre-trained transformer (ChatGPT) in 2022 has made us rethink how an LLM, not even general AI, can reshape our lives and eventually become our most significant competitor in the labor market. Within a short

period, we can see substantial changes in the way content creators, translators, lawyers, scientists, and many other professionals work (Fui-Hoon Nah et al., 2023). The biggest threat of a machine, which can (imitate to) imitate thinking and behave like a human, is that it is also capable of working like a human, but faster, with fewer mistakes, and without the need for rest, rendering human beings irrelevant in the labor market. Without work, the entire human global society and ecosystem can change radically, as work has been a means to survive and fulfil human aspirations since the dawn of time.

Paradoxically, since we work, we always look for means to make work easier, faster, and less painstaking, and eventually look for ways not to work at all but to automate tasks. From inventing the wheel to modern robots, all technological innovations in the past have changed the way we work and, consequently, labor market dynamics in some way. Some jobs and tasks disappeared (e.g., switchboard operators), others came into existence (e.g., data scientists) (Wilson et al., 2017), tasks and required skills changed, and the labor market transformed together with humanity. These past changes triggered by technological advancements in the labor market have already been extensively examined (Acemoglu and Restrepo, 2018; Frey and Osborne, 2017; Garcia-Murillo et al., 2018; Mitchell and Brynjolfsson, 2017; Țițan et al.,

^{*} Corresponding author.

E-mail addresses: szabina.fodor@uni-corvinus.hu (S. Fodor), bettina.boncz@stud.uni-corvinus.hu (B. Boncz).

Peer review under responsibility of Xi'an Jiaotong University.

<https://doi.org/10.1016/j.dsm.2025.11.001>

Received 27 August 2024; Received in revised form 11 November 2025; Accepted 12 November 2025

Available online 20 November 2025

2666-7649/© 2025 Xi'an Jiaotong University. Publishing services by Elsevier B.V. on behalf of KeAi Communications Co. Ltd. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

2014). Extensive evidence suggests that technological progress usually improves human productivity and decreases routine, monotonous work which follow structured rules mainly because it can easily be taught to a machine (Atalay et al., 2018), whether it is a simple mechanical spinning jenny or complex machine learning algorithm.

In addition, technology changes the required education level, skills, and capabilities (Chigbu and Makapela, 2025; Mirbabaie et al., 2022) needed to succeed in the labor market. A change in the required human skill set and education level in the labor market is bringing inevitable social changes along with it (Khogali and Mekid, 2023). It is sparking, among others, a fierce debate on technology in education (Bond et al., 2024) and education about technology (Casal-Otero et al., 2023). Therefore, drastic labor market changes bring about drastic societal changes. Guilds were replaced by manufacturers, opening up doors to less skilled workers in the labor market; early industrial workers were replaced by engine-powered machines, opening the door to the creation of white-collar jobs. Moreover, the improvement in agricultural practices has meant that, almost 90% of the Earth's population worked in agriculture in the early 19th century, this number is close to 2% today (Ahlqvist, 2005; Segal, 2018) without creating a global unemployment crisis. However, the timescale is the difference between examples of the past and the predictable future with AI. Although changes have occurred over decades, we can now observe rapid and extreme changes in a matter of years (e.g., the emergence of the LLM). Radical changes in labor needs, and the required skills and abilities can occur within the same human generation, possibly even several times. As humans cannot be reprogrammed easily, many can lose their jobs as they are unable to keep pace with technological advancements, causing technological unemployment (Lima et al., 2021). In 2014, research from the OECD countries (Arntz et al., 2016; OECD, 2021) revealed that on average, 14% of jobs were susceptible to automation. Five years later, in 2021, the growth of these jobs decreased. They lost their stability, although the potential worker numbers also decreased. According to Frey and Osborne (2017), the real automation danger can reach up to 47% in the USA in the future, although in a recent revamp of their research, they represent a more conservative standpoint. However, we should not only examine technology-induced labor market changes through dark lenses. Studies have pointed out that technological changes can positively impact the labor market (Mutascu, 2021) by increasing productivity (Al-Emran et al., 2025) and wages, or even decreasing product prices (Noy and Zhang, 2023).

When discussing AI's impact only, extant studies focus on task automation instead of job automation. Thus, whether a specific AI application augments or replaces a particular human task is what ultimately matters (Furendal and Jebari, 2023; Raisch and Krakowski, 2020). This implies that jobs have not ceased to exist. Rather, they are transforming, allowing humans to increase their productivity and concentrate on higher value-added tasks that require more human skills. In recent years, the biggest game changer has been the appearance and democratization of generative AI. Generative AI refers to algorithms trained on large datasets that are capable of generating texts or images. While past technological breakthroughs (e.g., the production line or steam engine) were mostly capable of automating manual jobs, generative AI is now also capable of automating some intellectual jobs (AbuMusab, 2024). Thus, all sectors and levels of the labor market are, to some extent, susceptible to this job transformation. Hence, AI's impact on the labor market is a critical research topic, as it affects all levels of society and the labor market.

Moreover, AI is a type of technology that is not only heavily influencing one's work but also private life. On human-AI relations at work, we cannot separate one's workplace "self" from their personal "self", as AI applications are getting increasingly widespread in all aspects of our lives. Workplace behaviors related to AI are heavily influenced by underlying everyday perceptions about AI, technology phobia at work or in

private life, technological understanding and knowledge of AI, and the locus of control over AI (Cheng et al., 2023; Kim and Im, 2023; Sinha et al., 2020). Conversely, work-related AI perceptions influence general AI perceptions, as human well-being is closely related to job safety (Nazareno and Schiff, 2021). People feel less threatened by AI when they are convinced that AI only replaces dangerous jobs and not theirs. If people work with AI and understand it well, they know that it does not threaten their job safety. Indeed, they can build better trust with technology in general (Albarrán Lozano et al., 2021).

This study advances our understanding of the current societal perceptions towards AI, labor market, and associated societal changes. We have collected a large textual dataset from one of the largest online forums, Reddit. Online forums emerged as platforms for AI discussion more than a decade ago when technology became a potential reality for the public. Concurrently, these began to trigger fears and hopes for the future of work and life (Neri and Cozman, 2020). Forums and, in general, social media have already been used in different studies to measure perception towards COVID-19 (Melton et al., 2021), health-related issues (Jeon et al., 2023), product usage (Singh and Glińska-Neweś, 2022), and other social (Botzer et al., 2023) or political (He et al., 2020) phenomenon and financial market dynamics (Long et al., 2021). From 37 Reddit forums, we have gathered data encompassing more than 100,000 comments about labor market changes with AI to analyze people's perceptions of AI through sentiment analysis and topic extraction using state-of-the-art technologies (bidirectional encoder representations from transformers, or BERT) developed by Google (Devlin et al., 2019; Venugopalan and Gupta, 2022).

2. Literature review

As this study discusses AI perceptions over 10 years, we need to recreate, to some extent, the evolution of the scientific standpoint about the future of work with AI over this period, providing sufficient context for later findings. AI-induced changes in the labor market have been studied for decades. Early research in the 2010s mainly focused on predicting the possible modifications that an intelligent machine or algorithm could bring to the labor market, as no tangible technologies were available yet. The general concern about AI's drastic impact on the labor market was already present in these studies, suggesting various measures to keep humans competitive in the long term in the labor market against machines, including restructuring education (Coates, 1998; Crawford and Calo, 2016; Kim et al., 2017), our current economic policy (David, 2017; Dirican, 2015; Harari, 2017), and the way we do business (Allen, 2017; Garcia-Murillo et al., 2018).

Today, with the widespread use of generative AI technologies, researchers can conduct more profound quantitative research in the field and, crucially, experience changes firsthand, which has changed the scientific discourse. Before the recent AI breakthroughs, the general perception of the impact of AI on the labor market was twofold: either AI is going to replace all jobs or is not going to change anything, but rather transform the labor market. Most research fully or partially supported one or the other standpoint at the time (although the context changed, partially even today). Complete human substitution theories were rooted in the perception that machines do not sleep, take sick leave, or demand paid leave. They performed tasks faster and made fewer mistakes. Therefore, AI can significantly increase productivity (Su et al., 2022) and significantly threaten human workplaces. The underlying logic was simple. Substitution generates immediate quality improvement and cost savings in many cases, which are attractive to companies (Fast-Berglund et al., 2020; Ray and Mookherjee, 2022). The substitution of technologies also requires less time and effort to implement than those that collaborate well with human co-workers, as the latter require significantly longer development and design time, adoption, and more costly maintenance. Additionally, workers' adaptation to new market

circumstances has human limits in our rapidly changing world (and work). Reeducating employees or hiring new employees with new skillsets requires time and money. Moreover, workers can also show resistance (or an inability to adapt) to changes. Therefore, companies may favor using robots or software that can easily be reprogrammed to execute new tasks more efficiently and in a cost-saving manner, making humans replaceable (Eldakruri and Senyurek, 2025; Korinek and Juelfs, 2024).

Theories on the opposite side of human substitution argue that AI will never be able to completely substitute humans because new machines and products create new jobs for humans based on the assumption that new machines need people (and professionals) to develop them, take care of them, or undertake maintenance and other tasks (Choi and Leigh, 2024; Pianta and Vivarelli, 2000). Additionally, the more we dive into substitution, the more new jobs will emerge on the other side. For instance, AI explainers, who help people understand, say algorithm-based decisions made about their lives (e.g., credit checks), is a job that did not exist before (Saeed and Omlin, 2023; Wilson et al., 2017). Human wage fluctuations can also positively contribute to the preservation of human workforce. Hughes (2014) believed that even if wages decrease due to technology, they will never reach zero; therefore, as long as humans are willing to be employed for that specific low wage, they will be used instead of machines (if the profit and loss analysis works). Conversely, if wages increase because machines make it possible to work more efficiently with lower costs, and human intelligence is used for more value-added (and salary-increasing) tasks, surplus capital will be created in the economy. This can be invested (Allen, 2017) or spent on extra consumption, which will create a place for new demand and workplaces (Jazdauskaite et al., 2021; Upadhyay, 2021), thereby creating a self-enforcing circle.

Technology also enables faster innovation by linking people's minds and sharing knowledge. Furthermore, AI can produce valuable insights. An innovative idea, even if AI-enabled, can lead humans to enhance their entrepreneurship by creating new products or industries. If people can eventually start and succeed with their business ideas, then they would not be forced to entirely rely on the labor market. They can simply become their own employer and the employer of others (Gama and Magistretti, 2025; Jorzik et al., 2024) while also creating social value (Battisti et al., 2022). AI-based product innovation (Cooper, 2024) can eventually create new jobs even in existing companies or industries. Today, with the growing adoption of generative AI technologies, the past 10 years of unknown "scientific guesswork" have moved in a more tangible, empirical direction. Moreover, AI's impact on the labor market has now, in some way, a different outlook. Recent studies have suggested that the reality lies between the two aforementioned theories. AI will subtly change the labor market, softening extreme utopian or dystopian theories. Earlier research has already suggested that digital technologies can cause job polarization: The ratio of highly educated, high-skilled workers to mainly low-skilled and educated workers performing non-routine tasks gradually grows. Meanwhile, middle-skilled workers primarily working in manufacturing, doing clerical, craft, or routine-intensive jobs significantly lose jobs (Goos et al., 2014; Kolade and Owoseni, 2022). We also call this phenomenon skill-biased technical change because the skills of higher- and lower-educated non-routine jobs are increasing in demand. Concurrently, the middle sector of the labor market hollows out.

Job polarization negatively affects the middle sector of jobs because machines can easily substitute for their workflows. These jobs involve predictable, easy-to-describe, and understandable processes that can be translated more easily into machine languages (Frey and Osborne, 2017). As technology advances, the proportion of workplaces at risk of being automated grows as intelligent machines can perform and learn a wider variety of tasks. If it reaches or even exceeds human intelligence,

AI will be able to automate most, if not all, of the aforementioned middle-sector jobs. Only jobs that require non-automatable human skills, such as creativity, human intelligence, and social skills, might survive. Empirical research shows that job polarization and/or skill-biased technological change are already occurring in several countries (Bárány and Siegel, 2020; Cavaglia and Etheridge, 2020) and will probably persist in the age of AI. The most significant novelty that current generative AI algorithms bring to the labor market is the shift in predicted job polarization tendencies: generative AI is shattering the safe haven of high-skilled workers, who were previously considered unsusceptible to automation. Recent research suggests that new technologies increasingly threaten high-skilled jobs (AbuMusab, 2024; Felten et al., 2023).

Paradoxically, AI skill demand among highly skilled workers is growing, which the labor market cannot fulfil. Therefore, a significant wage premium is associated with AI-related jobs (Alekseeva et al., 2021; Duch-Brown et al., 2022; Hui et al., 2024; Mäkelä and Stephany, 2024), deepening the job polarization effect. Essentially, those who are AI-skilled secure a stronger labor market position, while AI can replace those who are less experienced in AI. A positive trend is that current AI technologies do not aim to replace all jobs. They aim to transform employment; furthermore, AI's job destruction effect is currently compensated for by its job creation effect (Damioli et al., 2024). First, because even the most intelligent generative AI remains a machine at its core, it cannot obtain purely human skills, such as creativity, social skills (Benvenuti et al., 2023), and empathy. This makes it impossible to completely substitute human workers. Second, AI is more efficient in reinforcing human capabilities (human augmentation) than in actually achieving substitution (Focacci, 2021). Thus, it will not replace humans but rather cooperate with them. A good machine collaborator can perform all mundane or dangerous tasks that people are unwilling to or should not do, and can share new and interesting insights to improve work efficiency. By 2025, 24% of the global gross domestic product (GDP) will come from AI (Dennehy et al., 2023). Modern generative AI technologies are capable of automating 27 million jobs globally (Gmyrek et al., 2023) and will augment approximately 46 million of them. We can already see that generative AI is used in diagnostics, content creation, academics, education, computer science, and agriculture (Fui-Hoon Nah et al., 2023). Moreover, human-AI collaboration indeed increases productivity in many economic sectors (Sowa et al., 2021).

Although generative AI is capable of improving, for example, the recruitment processes and finding the best candidates for jobs (De Obesso Arias et al., 2023), a human must still make the hiring decision. Although an increasing number of business-to-business (B2B) companies rely on AI-powered sales bots to arrange business transactions without human intervention (Frey and Osborne, 2023; Yin and Yuan, 2022), human workers still make complex business decisions. AI still requires human supervision (Downey, 2021), which can increase user trust (Aoki, 2021) in AI applications. If AI needs humans just as much as humans need AI, then the next most significant challenge in the labor market is ensuring that human-AI collaboration occurs such that it benefits both parties. AI must adapt to humans through good design, interface, and built-in security factors (Sanfilippo et al., 2025). Most importantly, humans must adapt to AI, or the discussion will lose meaning.

Human-AI interactions must be carefully orchestrated from a human end user's perspective. Human-AI augmentation can negatively influence human self-esteem, increase feelings of job insecurity (Bhargava et al., 2021), and decrease team performance. This is especially true when humans are faced with the fact that in specific tasks, they can never be as good as machines, discouraging them from striving for excellence (Felten et al., 2023; Flathmann et al., 2023). AI must also be adapted to humans. Generative AI applications are prone to




Fig. 1. An overview of our research methodology.

hallucinations (Orchard and Tasiemski, 2023), wherein they can create false information. Worryingly, they can create harmful information, fake news, or cause data privacy issues, potentially misleading or harmfully influencing human counterparts. Creating and retraining AI, especially on representative data, requires enormous capital employment, which only a few countries can afford. This can restrict access to technology, increasing inequality, and negative societal impacts, such as a lack of access to upskilling opportunities to work with AI (Bhargava et al., 2021). AI also has high energy consumption, which increases its negative environmental impact (Frey and Osborne, 2023; Fui-Hoon Nah et al., 2023; Khogali and Mekid, 2023), an important point to consider for all organizations and humans engaging with the technology. Every AI application must deal with the inherent human factor: people trust AI (or machines in general) less than their human counterparts (De Obesso Arias et al., 2023), especially when the stakes are high or the task to be executed is complex (Chang, 2022). When people feel that their jobs are at risk of automation due to new AI technology, they can also manifest resistance to change (Cheng et al., 2023). This resistance can also be triggered when the understanding or trust in the new technology is low (Chiu et al., 2021; Leichtmann et al., 2023), or communication channels with or about AI are restricted or not user-friendly (Flathmann et al., 2023). Human-likeness also influences adoption. Anthropomorphism can trigger the rejection of collaboration in humans if situated in the “Uncanny valley” but can also increase attachment (Kim and Lee, 2023). Furthermore, underlying technology phobia is vital for workplace adoption (Sinha et al., 2020).

This research aims to address the question of what people feel and think about AI at work, revealing how the perception of AI and work has changed over the past 10 years. This period encompasses the moment when AI was merely a science fiction topic to today when almost everyone works with AI daily. Undoubtedly, we can assume that AI is present and will remain present in our lives in the long term. Artificial Intelligence likely lacks any inherent emotions or subjective thoughts regarding collaboration with humans. However, humans do crucially, our perceptions can make or break the success of the next AI-augmented chapter of our human history.

3. Research design: questions and methods

We do not seek to understand the labor market changes caused by AI, regardless of whether it is generative or general. We want to understand the perceptions of the studied population regarding AI and the future work. As such, the following research questions (RQs) seek to achieve the objectives of (1) understanding the general perception and emotions of tech-enthusiast online communities towards AI and its impact on the future of work, and (2) studying the most important topics of discussion around AI and the future of work.

RQ1. What are the most interesting topics concerning AI and work in forums?

RQ2. What are the general sentiments and emotions of forumers towards AI and work?

RQ3. What have been the turning points in the past ten years from the perspective of forumers concerning AI and work?

Data collection, data cleaning, and three distinct text analysis techniques—sentiment analysis, emotion detection, and BERT-based topic modelling—were used in the research process, as illustrated in Fig. 1. RQ1 was addressed in the subsection titled “The topics discussed and their associated emotions”. RQ2 and RQ3 were the responses to a discussion of chronological alterations in sentiments and emotions.

3.1. Data collection

Data from social media platforms are often used for topic modeling (TM) and sentiment analysis due to the vast amount of real-time user-generated content on these platforms. Several billion users share their opinions daily on social media (Yadav and Vishwakarma, 2020). 63% of the population with an Internet connection used Facebook weekly, while 23% used it daily.¹ Social media platforms such as X (formerly Twitter), Facebook, and Reddit offer rich sources of unstructured data that reflect public opinion, emotions, and trends, making them invaluable for understanding consumer behavior, market dynamics, and societal trends.

Reddit was selected as the source for the analysis. Reddit is an online forum where topics are categorized under “reddits”. In 2022, the forum had 430 million users; of which 47% lived in the USA. Regarding the population, 63% of the users are male, and approximately half of all users, otherwise called redditors, are less than 35 years old. According to a survey conducted in 2021, the majority of Reddit users have a higher education degree. Reddit has, on average, 1.7 billion visitors per month (Sun et al., 2021). According to official Reddit Statistics, 38% of users are intensely interested in technological novelties. These statistics indicate that Reddit is a good data source for our analysis (Aldous et al., 2023). Essentially, we could access the opinion of a population more involved in technology, and therefore, the AI topic than the average person.

This research was exempt from institutional ethical review by Corvinus University of Budapest, as it exclusively analyzed publicly available online content. No personal or identifying information was collected, and the data were processed in aggregate form to represent general discourse patterns rather than individual statements. All materials originated from the public domain, and no interaction with users occurred during data collection or analysis.

We collected Reddit social media data from January 2013 to January 2024 from the subreddits communities associated with “artificial intelligence” or “AI”. We identified a list of search terms, such as “artificial

¹ <https://www.statista.com/topics/1164/social-networks/#dossierSummary>.

intelligence + work/job” or “artificial intelligence + automation”, which we used on the Reddit platform to retrieve forums associated with the future of work with AI topic. We used the search terms with and without standard abbreviations, such as “AI”. Then, we manually reviewed and filtered the search results to ensure that only forums with relevant discussions were included. If, for example, a forum was asking for help in using an AI application to write a resume for a job application, we excluded it. This is because it concentrated on a special use case and not a general discussion about the implications of AI on the future of work. For forums created after 2022, we paid special attention to exclude any forums that only discussed ChatGPT’s special use cases as a standalone topic, such as how to write a good prompt. We included forums on ChatGPT’s overall impact on the future of work and trending topics where job transformation in some sectors was discussed. This allowed us to identify industries where generative AI’s impact is visible and collect testimonials. Our general guideline was to consider AI and what forum users consider AI. Finally, we only considered reddit posts or forums with more than 40 comments among the search results to ensure that we have sufficient data per forum. Finally, we selected 37 different forums that were scraped from top-level comments using the Python package PRAW through an API,² and organized them into one database of 117,834 comments of various lengths. For the data cleaning, first, we deleted the comments labelled “removed” or “deleted”, meaning that the user or platform’s moderators have unpublished them after posting. Second, comments that were too short to contribute to our analysis were removed. We decided not to define a character minimum per comment. We instead manually reviewed comments with the lowest character count and filtered out the ones that were only answering a closed question with a “yes”, “no”, “yeah”, or “indeed”. This labor-intensive process allowed us to more accurately clean the data. We also deleted comments that were not in English (the specific steps are shown in Appendix Fig. A1). Comments were stripped of special characters, URLs, and emojis. Following data cleaning, we conducted TM, sentiment analysis, and emotion recognition on the 114,377 tidy comments.

3.2. TM

TM algorithms are a class of unsupervised machine learning techniques (Kherwa and Bansal, 2018) that can discover and understand topics in a collection of documents, where a topic is a set of statistically related words. Determining the number of topics for a TM algorithm presents challenges, as the “correct” number is often not predetermined or explicitly defined. This challenge stems from the unsupervised nature of the methodology. Several algorithms can be used for identifying latent semantic patterns, such as latent semantic analysis (LSA) (Hofmann, 2001) and latent dirichlet allocation (LDA) (Blei et al., 2003). Recently, newly developed algorithms such as non-negative matrix factorization (NMF) (Lee and Seung, 2000), Top2Vec, and BERTopic (Grootendorst, 2022) have received increasing attention (Obadimu et al., 2019; Sánchez-Franco and Rey-Moreno, 2022).

These clustering algorithms can be evaluated using several metrics. TM solutions have two types within them. First, the intrinsic properties of the clustering result can be compared with the internal cluster validity values, which only consider structural aspects such as the degree of separation of clusters and do not rely on input data. Second, clustering results can be compared by using an external knowledge source, such as a known classification of the document space (performed through traditional content analysis, for example). These are referred to as the external cluster validity values. This study used internal cluster validity values. Scholars (Egger, 2022; Liu et al., 2016) typically evaluate their performance using internal clustering values when comparing topic-modeling algorithms. The most commonly used measure is perplexity (Zhao et al., 2021), which compares the theoretical word distributions of topics with their actual distribution. The value is not to be

interpreted in isolation; however, the model with the lowest perplexity value is considered better than the other models. The other relevant metrics are word- or document-based topic coherence measures (such as PMI,³ CUMASS⁴), and KL⁵ topic divergence (Rosner et al., 2014) that aim to distinguish between semantically meaningful topics and the artefacts of statistical inference. Numerous studies have indicated that the PMI exhibits the strongest correlation with human judgment (Lau et al., 2014; Röder et al., 2015). Recent studies (Egger, 2022) have shown that BERTopic outperforms LDA, particularly for non-long corpora. Additionally, several researchers have strongly criticized the effectiveness of the LDA algorithm (Egger, 2022; Sánchez-Franco and Rey-Moreno, 2022). A document usually has multiple topics and LDA tends to neglect co-occurrence relationships (Jaradat and Matskin, 2019). Another criticism is that it is particularly relevant to noisy and sparse datasets. In such cases, the LDA algorithm is unsuitable owing to the lack of features required for statistical learning (Chen et al., 2019). The NMF algorithm relies on the co-occurrence patterns of words that may miss contextual meanings. Recent research compared the performance of the BERTopic and NMF algorithms on different datasets, and found that BERTopic generally achieved a higher topic coherence score than the NMF algorithm on all datasets, especially noisy ones (Grootendorst, 2022). Considering these criticism, we used the BERTopic solution proposed by extant (Thompson and Mimno, 2020) using a quantized LLM technique (Grootendorst, 2022) to automatically label topics. Our TM approach comprised four steps.

- First, the text was transformed using a pre-trained sentence transformer language model called “all-MiniLM-L12-v2” to obtain a 384-dimensional vector space.
- Second, we reduced these 384 dimensions to optimize the clustering process. Principal Component Analysis is perhaps the most popular method for dimension reduction. We used the uniform manifold approximation and projection (UMAP) method, which better preserves the local and global structure of datasets (Grootendorst, 2022; Thompson and Mimno, 2020). Based on the results of related literature, we used the five nearest neighbors in UMAP to prefer local structures.
- Third, the reduced dimensional embeddings were clustered using the hierarchical density-based spatial clustering of applications with noise (HDBSCAN) algorithm that defines clusters of different densities. This algorithm allows noise to be modelled as an outlier, preventing unrelated document parts from being assigned to any cluster and improving topic identification. We set the minimum cluster size to 45 to identify topics with a sufficiently large number of comments, thereby increasing the reliability of our method.
- Fourth, to characterize the resulting twenty-three clusters, we combined the comments belonging to the clusters and created topic vectors characterizing the clusters using three different methods: the class-based frequency-inverse document frequency approach (c-TF-IDF⁶), KeyBERT (Grootendorst, 2020; Sammet and Krestel, 2021) extraction technique, and quantized LLM technology. The quantized LLM method was initially employed to refine and improve topic

³ Pointwise mutual information (PMI), quantifying the probability of two words occurring together. If the two words are a and b , then $PMI(a, b) = \log \frac{P(a, b)}{P(a) \cdot P(b)}$.

⁴ UMass-coherence (CUMASS) calculates how often two words, w_i and w_j appear together in the corpus, and it’s defined as $CUMASS(w_i, w_j) = \log \frac{D(w_i, w_j) + 1}{D(w_i)}$, where $D(w_i, w_j)$ indicates how many times words w_i and w_j appear together in documents, and $D(w_i)$ is how many times word w_i appeared alone.

⁵ Kullback-Leibler (KL) divergence measure of how much an approximating probability distribution Q is different from a true probability distribution P , and it is defined as $D_{KL}(P||Q) = \sum_{x \in \mathcal{X}} P(x) \log \left(\frac{P(x)}{Q(x)} \right)$.

⁶ The c-TF-IDF value of a term x in a given class c , $w_{x,c} = \|tf_{x,c}\| \cdot \log \left(1 + \frac{A}{f_x} \right)$, where $tf_{x,c}$ is the frequency of word x in class c ; f_x frequency of word x in all classes and A is the average number of words in the classes.

² <https://github.com/praw-dev/praw>.

representation (Grootendorst, 2022). Combining BERTopic with quantized LLMs offers a practical and efficient TM method. Quantization is crucial for using LLMs. This involves reducing the precision of the model's weights by assigning smaller approximations, such as 4- or 8-bit values, instead of the original 32-bit floating points. Although accuracy may slightly decline, this approach effectively reduces the model's memory requirements. This study employed the pre-trained language model "OpenHermes-2.5-Mistral-7B-GGUF" and utilized the LlamaCPP (Betlen, 2023) representation model. A prompt is created for the LLM to be used when generating topic labels. These labels were obtained by clustering comments related to each topic.

3.3. Sentiment analysis and emotion detection

Sentiment analysis is widely used in academic research and business contexts to classify sentiments expressed in text based on the expressed viewpoints (Adoma et al., 2020; Brett et al., 2019). Researchers can easily access large amounts of textual data for sentiment analysis on the internet, which has become a preferred source for marketing, social, and political studies (Zhang and Zhang, 2022). Existing sentiment analysis techniques can be categorized into three main approaches: lexicon-based, machine-learning-based, and hybrid. Lexicon-based sentiment analysis mainly depends on predetermined dictionaries or lexicons that associate words with sentiment scores (positive, negative, or neutral) without considering the context of the words. This approach has difficulty in capturing negations, sarcasm, or complex linguistic structures which may alter sentiment meaning.

Machine learning-based methods, mainly supervised algorithms, leverage labelled datasets to train models that can classify sentiments. Recent natural language processing advances, such as deep learning models (e.g., BERT), have further enhanced the ability to capture contextual sentiments and offer a more precise understanding of the emotional intensity conveyed in text (Rodríguez-Ibáñez et al., 2023). These methods are highly accurate when ample labelled data are available. However, they are computationally expensive and require significant preprocessing. Hybrid approaches combine the strengths of lexicon-based and machine learning methods using lexicons to guide or enhance machine learning models. These methods attempt to address each technique's shortcomings but may still inherit their computational complexity. To perform sentiment analysis and emotion detection, we employed a pre-trained deep learning model called RoBERTa (Liu et al., 2019). Recent studies revealed that RoBERTa consistently outperforms many other techniques in fine-grained emotion classification (Zhang et al., 2025). RoBERTa enhances BERT by adjusting essential hyperparameters, including eliminating the next-sentence pretraining target, utilizing larger mini-batches, and increasing learning rates during training. This improvement seeks to improve the model's resilience and efficacy. RoBERTa can be used to compute the anticipated probability of three sentiment categories: negative, positive, and neutral (Barbieri et al., 2021). Additionally, it can determine the prevalence of the eight most frequent emotion categories (joy, confidence, fear, surprise, sadness, anticipation, anger, and disgust) in a given text. To ascertain a given comment's overall sentiment and emotion categories, we identified the sentiment or emotion with the most significant percentage value. We used the "cardiffnlp/twitter-roberta-base-sentiment"⁷ (Barbieri et al., 2020) trained model for sentiment analysis, and "j-hartmann/emotion-english-roberta-large"⁸ for emotion analysis.

⁷ This is a RoBERTa-base model trained on approximately 58 M tweets and finetuned for sentiment analysis with the TweetEval benchmark.

⁸ This is a RoBERTa-based model for classifying emotions in English text data. The model was trained on six different datasets, and can predict six basic Ekmanian emotions (anger, disgust, fear, joy, sadness, surprise) and one neutral class.

3.4. The topics discussed and their associated emotions

Next, we present the TM results using BERTopic. This analysis was conducted to answer RQ1: "What are the most interesting topics concerning AI and work among the forums?"

We extracted 23 topics, which are outlined in Table A1 in the Appendix. To represent each topic based on the content of the comments clustered for each topic, an LLM tag and a set of 10-10 keywords defined by KeyBERT or c-TF-IDF were defined, as described in the methodology section. After determining the emotions of the sentences associated with each topic, we constructed an emotion map for each topic, as summarized in Fig. A2 in the Appendix.

Fig. 2(a) shows the cumulative percentage of comments by topic. Notably, 87.2% of the comments collected were identified as BERTopic. The most commented topic was Topic 0 (87,347), while Topic 22 (275 comments) had the fewest comments. We examined which comments were not classified (17,494); that is, those identified as noise by BERTopic. These were generally comments containing short emotional expressions, such as "You just won't let logic get in the way, huh?", or specific information the algorithm could not interpret, such as "I'll take option 3, the Star Wars universe!" Fig. 2(b) chronologically shows the five most common topics. Topic 1, for self-driving cars, peaked when the first automated vehicle autonomously drove from the west to the east coast of the USA in 2015. Doctors being replaced by AI (Topic 3) peaked in 2019 and 2023, suggesting the importance of technological breakthroughs on perceptions suggesting the importance of technological breakthroughs on perceptions. Universal basic income (UBI; Topic 2) usually strengthens when some countries or states conduct experiments (e.g., the US or Finland). ChatGPT (Topic 4) has gained extreme popularity since its launch in 2022, making it a dominant topic. This demonstrates the influence of "breaking news" on the topics of interest.

After analyzing the content and emotions of the 23 topics extracted from the database, we created four meta-topics. As shown in Fig. 3, the meta-topics are general perceptions about the future of work together with AI; discussions about the economic implications of AI and the future of work; job transformation, which covers all the topics that were related to specific job changes or general labor market changes; and finally, politics and influencers show us topics about the people and cultural items that influence the trending conversations.

Next, we analyzed the emotions linked to the comments categorized under these meta-topics. The results are presented in Fig. 4. Next, we elucidate each meta-topic in greater detail, informed by the keywords, LLM labels, and emotions of the 23 topics (see Table A1, Fig. A2 in Appendix).

3.4.1. Meta-topic: general AI and the future of work

This meta-topic broadly discusses the future of AI and human workforce. Therefore, it can provide a general understanding of the feelings and focus points on the main topic, the future of work (see Fig. 4). First, two significant topics (Topics 4 and 15) discuss certain technologies, such as LLM and ChatGPT, which strongly influence the future of work. These topics are generally followed by emotions of surprise (see Topic 4 in Fig. A2 in the Appendix). While this topic typically praises ChatGPT, strong concerns (design, understanding, and user's level of knowledge) have also been raised. It states that even with a robust tool such as ChatGPT, humans cannot be replaced because it requires human supervision at the very least. Second, some topics (Topics 0, 11, and 17) are generally about AI and the future of work (see Table A1 in the Appendix). They were less optimistic about the outcome than the previous topics, which were primarily concerned with ChatGPT. They were filled with fear, anger or disgust. The general perception is that AI can replace all jobs, and chaos and societal collapse can follow. Humans are not competitive with machines. Further, companies have an economic interest in automation and process improvements, which can lead to complete human replacement. This difference is also reflected in the literature review. Before AI became a reality and




Fig. 2. Distribution and temporal trends of topics. (a) The cumulative percentage distribution of posts by topics; (b) chronological change in posting frequency of the top five most commented topics.




Fig. 3. The 23 topics are described using three words chosen by KeyBERT. The coloring illustrates the four meta-topics we manually defined, with cyan representing “general AI and future of work”, purple representing “job transformation”, gold representing “economy”, and dark pink representing “politics and influencers” meta-topics.

more democratized, general perceptions were quite extreme regarding labor market outcomes: either nothing would change, or machines and automation would replace all human jobs. This shift is also confirmed by general perceptions. Newer ChatGPT-related topics are more favorable than older comments, which are somewhat negative and predict a dystopian future.

3.4.2. Meta-topic: economy

This meta-topic contains topics related to wealth, income distribution, taxes, and UBI. UBI is a periodic (weekly, monthly, or yearly) given set amount of money whose use is not linked to any conditions. UBI appears twice: once alone (Topic 2) and a second time associated with automation (Topic 20). In most cases, anger and disgust are the most potent emotions linked to the meta-topic (see Fig. 4 and Fig. A2 in the Appendix). Looking closely at the comments, we can see that redditors consider UBI a potential solution to AI and automation, as it is presented as a guarantee for humanity’s survival in a jobless future. UBI topics are associated with negative emotions, such as disgust and anger. Forumers generally do not believe in the feasibility of the project, as it would, for example, only enforce work avoidance.

Another topic is taxation (Topic 13) and whether automated companies should be taxed for not using human labor (Gasteiger and Pre-ttner, 2022). Associated keywords include economic and shareholder interests, competition, and profit-seeking. The central discussion concerns a future in which human labor is no longer needed. Therefore, nations should obtain financial resources from consumption, company income tax, or even taxing robots to finance public services and potentially universal basic income.

3.4.3. Meta-topic: politics and influencers

The dominant emotions in this meta-topic are mixed (Fig. 4). The most anger and disgust-driven topics (Topic 14-Bill Gates, Topic 16-Elon Musk), and most positive (surprise and joy) topics about Andrew Yang (Topic 19) and Stephen Hawking (Topic 21) can also be found here (see Fig. A2 in the Appendix). These topics are unsurprisingly emotion-driven, discussing the role and impact of Elon Musk and Bill Gates on AI development, and their wealth as influencing powers. US politicians, such as Andrew Yang and Bernie Sanders, are the most dominant in the dataset, partly explained by the English-speaking audience and many Reddit users living in the USA.




Fig. 4. Meta-topics and their combined emotions.

The redditers perceive that AI development hugely depends on major technology figures and organizations, such as Elon Musk, Bill Gates, and OpenAI, which are also mentioned among the topics, especially in the LLM impact topic. Notably, not only politicians and people in the business sector but also Stephen Hawking, a member of the scientific community, have often been mentioned, especially his vision and understanding of AI. Thus, members of the scientific community who can overcome the barrier of popularity can also substantially influence general perception. We also included Skynet (Topic 22) in this category. Besides real-life influencers, science fiction also plays a substantial role in forums. We found evidence of films, such as Terminator, Star Trek, and Matrix, in the dataset. Previously, when AI was the primary subject of science fiction, it was the only source of common AI perception (Cave and Dihal, 2019), and the literature was rarely utopian about the future of humans and AI. We usually observed representations of malevolent, hostile AI that turns against its creator. Only in some cases do they serve humanity to gain prosperity and peace. Today, when AI approaches reality, people in the Skynet topic are trying to compare what can already be seen with what has been expected through popular culture.

3.4.4. Meta-topic: job transformation

This meta-topic deals with job transformation within certain professions, perceived to be the most affected by AI. Some talk about how different jobs, especially intellectual ones, will change, such as doctors (Topic 3), software developers (Topic 6), therapists (Topic 12), and lawyers (Topic 5). This is the topic for which ChatGPT is the most dominant. Another topic (Topic 6) has discussions on generative AI in software coding. This study takes another dive into the topic of humans not needing more human therapists, as generative AI can provide sufficient help and social interaction. One topic is about how ChatGPT is used as a work aid, although we carefully removed all forums that only dealt with the practical workplace implications of ChatGPT from our database. Next, a topic (Topic 18) is on the future of translators. A job that many considered very human, as the understanding of culture and context is associated with it, can now be accomplished through generative AI. Interestingly, only a few topics concerned job replacement in poorly educated jobs. One topic (Topic 1) briefly discusses replacing drivers (trucks or taxis), which is strongly linked to the more prominent

topic of the development of self-driving cars. Another mentions that bank tellers and cashiers (self-checkout) are being automated. This category is dominated by jobs requiring high levels of education, which can either be explained by the demographic characteristics of redditers or by the fact that these are jobs that are traditionally considered challenging to replace, and any technology attempting to do so often gets more publicity. Further, one topic (see Topic 7 in Table A1 in the Appendix) is about a well-known analogy of human workforce replacement: the story of horse replacement over the last century. Horses have been used in all aspects of life for several hundreds of years. They were essential partners of humans in war, transportation, everyday manual work, and leisure activities. Nowadays, their role in human society is reduced to racing, leisure activities, and occasional labor. In France, for example, the total horse population since the 19th century fell to a third (Rzekeć et al., 2020) due to the changes in utilization. Forumers believe that humans can end up like horses, being entirely replaced in the labor market by intelligent technological innovations, such as AI, thus exhibiting a rather negative future.

3.5. Chronological alterations in sentiments and emotions

Next, we present the sentiment analysis and emotion detection findings to answer RQ2: “What are the general sentiments and emotions of forumers towards the topic of AI and work?” Fig. 5 shows the RoBERTa sentiment analysis results. The overall attitude of the 37 forums was negative.

Among emotionally driven topics, two changes in the negative trend are observed: First, from 2019 onwards, when the COVID-19 pandemic hit the world (decrease); and second, from 2022, from the release of ChatGPT (slow increase). ChatGPT’s impact can also be observed in the sharp rise in the number of comments that appeared in 2023, immediately after its launch to the public.

The proportion of positive attitudes remains relatively stable over time. Both negative and positive emotional intensities peaked following the launch of ChatGPT in 2023 (Fig. 6). During the COVID-19 period, negative emotional intensity remained stable, positive emotional intensity increased. However, the launch of ChatGPT impacted both. A RoBERTa emotion analysis was also conducted to understand the




Fig. 5. RoBERTa sentiment analysis results. (a) The number of posts (comments) and their sentiments chronologically; (b) the overall sentiments of all posts (comments), shown chronologically.

emotions underlying these sentiments and draw more accurate conclusions. Fig. 7 shows the results over the years grouped by negative and positive sentiments (see Fig. A2 in the Appendix). This analysis does not consider comments with neutral sentiments. Our two focal points, namely, the COVID-19 pandemic and the release of ChatGPT, triggered significant changes in particular emotions.

3.5.1. The COVID-19 impact on emotions

We observed that 2019 marked a drop in “fear”, a sentiment that only started to increase again after the COVID-19 pandemic. We also observed a gradual increase in positive emotions. “Surprise” began slowly rising, while no significant change in “joy” happened during the pandemic. Simultaneously, “sadness” and “anger” began increasing. These results may be explained by the acceleration of digitalization in several domains, such as education, e-commerce (Gryaznov, 2022), and workplace management (Kniffin et al., 2021) during the pandemic. It promoted remote work, and helped in re-evaluating the role of humans and technology in the labor market. Many workplaces were forced to accelerate the deployment of digital tools, which served as both work aids and tools of surveillance (Aloisi and De Stefano, 2022). There were high hopes that advanced technology would help humanity find the cure faster with technologies such as AI. Simultaneously, digital tools were also used to monitor the propagation of the virus and home office activities to impose quarantine measures. This, according to some, harmed the liberty of human (Chettri et al., 2020; Ting et al., 2020).

3.5.2. The ChatGPT impact

ChatGPT was released by OpenAI at the end of 2022 and has significantly affected AI-related discourse. This was the first publicly available, robust generative AI model. These trends were also reflected in the subreddits. ChatGPT related forums, though they can only represent less than 2 out of 10 years of this study’s scope, are the second most significant source of comments (see Table 1 and Fig. 3) and the only central forum where positive emotions (“surprise”) dominate. In 2023, the number of comments on AI and work-related forums significantly grew (see Fig. 5) due to the appearance of ChatGPT. In the sentiment and emotional analyses, the arrival of ChatGPT entailed an immense increase in “fear” and “surprise”, and a trend change in “anger”, which slowly decreased after the pandemic and began increasing again in 2022.

Given that Reddit users are usually young and have at least a college education, we can conclude that this is the first AI technology that could shake the world of this population. Previous automation technologies usually affected the work of the population with lower education, such as taxi or camion drivers (self-driving cars), manufacturing line workers (robotization), and shop assistants (self-check-out). This led the conversation to a very theoretical and bird’s-eye-view level, mostly discussing political, economic, and social impacts. With the arrival of ChatGPT, the jobs of the highly educated population were endangered for the first time, which changed the perspectives of redditors. Moreover, while neutral comments’ proportion remained stable during the




Fig. 6. Changes in sentiment intensity over time. (a) The change in the mean values categorized as negative sentiment using RoBERTa over ten years; (b) the change in the mean values classified as positive sentiment using RoBERTa over ten years.




Fig. 7. BERT emotion analysis results. (a) The number of posts (comments) and their emotions chronologically; (b) the overall emotions of all posts (comments) chronologically.

Table 1

Reddit forums with the highest number of comments and their emotions.

Name of the Reddit forum	Number of comments	The most dominant non-neutral emotion
Futurology	85,068	Anger
ChatGPT	10,820	Surprise
Worldnews	4,960	Disgust
Technology	4,223	Disgust
Capitalism vs. socialism	1,113	Anger

years, a considerable increase was observed from 2022. Thus, the conversation moved to more “fact-based”, given that they no longer had to guess or predict the impact of AI on work. Rather, they were experiencing it first-hand.

4. Discussion

Individuals’ work and private selves cannot be distinguished. Their attitudes, perceptions, fears, and expectations outside work strongly influence their adaptability and perceptions of AI at work. AI and the future of work are widely discussed on Reddit. Subreddits like “Futurology”, which is one of the primary used sources in AI-related forums, is the 34th⁹ biggest community on Reddit. Thus, a general curiosity regarding this topic exists. ChatGPT’s emergence has put the future of work even more in the spotlight, demonstrated by the rapid increase in ChatGPT-related forums in the past year, and sharp increase in the number of comments related to AI and job market since 2023. ChatGPT-related forums are the second most significant data source in our dataset, although they only represent just a little more than one year of the study period.

We have shown that the general perception of AI on job-related forums is mostly negative towards the future of work. Emotionally, we observe a present fear and anger towards AI and the future of work. A textual analysis of the comments showed that people are not only afraid of the possible labor market outcomes, potentially losing their jobs, income, or place in society, but also fear and vividly discuss the broader implications of AI, as demonstrated by the topic extraction. The impact of job transformation on the economy, politics, and society is discussed in detail. These were also the most emotionally driven topics. General

emotions are neutral and negative, mostly indicating disgust and fear. Moreover, modern-day influencers highly affect perception: politicians (primarily US-based), popular culture (Star Trek, Terminator, etc.), and, to some extent, scientists. Consequently, the media strongly influences people’s attitudes (Nader et al., 2022). Additionally, the chronological analysis of topics revealed that certain “breaking news”, such as the first autonomous car crossing the US, inflated the number of comments. Conversely, although sometimes not fact-checked, Internet resources also increasingly influence general perceptions (Nader et al., 2022). We also found traces of social media and YouTube channels in the forums. Undoubtedly, forums such as the one we studied, Reddit, also influence perceptions. When an in-depth understanding of a technology, such as AI, is low, it is easy to make people believe anything, good, or bad, which can create impossible expectations or unrealistic fears.

While the general AI perception of the future of work is negative, the current ChatGPT perception is quite positive primarily based on the general surprise at the capabilities of generative AI. ChatGPT was democratized quickly, allowed people to gain first-hand experience with technology and its capabilities, and created realistic expectations about workplace implications (Bain and McCay, 2024). Accordingly, analyzing the topics revealed that the most dominant subjects are possible task replacements, with almost no mention of complete technological unemployment, as opposed to the general AI topics mentioned before. Forumers believe that generative AI will make their jobs easier and lower career entry barriers. AI will always need supervision and human interaction to work correctly, which will not make people susceptible to automation, but rather transform how they work, potentially in a good way. These findings align with those of recent studies indicating that, while humans typically regard AI as an advanced technology capable of societal progress, apprehensions around job displacement and ethical considerations are prevalent among forum participants (Ocal and Crowston, 2024; Savela et al., 2024).

Interestingly, redditers discussed the potential job transformation or replacement of highly educated jobs: doctors, lawyers, software developers, educators/teachers and professors, translators, and therapists. The latter’s perceptions have already been examined in detail (Aktan et al., 2022). This dominance can be attributed to several factors. First, the population of the forumers: As mentioned above, they usually belong to more highly educated groups, meaning they are more concerned with the risks to their jobs than others. Second, these were traditionally deemed irreplaceable jobs, as they required human characteristics. Doctors and therapists need social skills and empathy; lawyers need to master persuasion and reasoning; coders need to understand humans to create software for them; translators need to understand humor, culture,

⁹ <https://www.reddit.com/best/communities/1/>.

and context; and finally, teachers need the most human skills among all of them to educate other humans from a very young age. For years, we have witnessed the automation of low-educated jobs such as manufacturing jobs, robotics, and self-service. However, with recent generative or general AI development, the highly educated jobs' intense "Ivory Towers" were penetrated. Still, they can still only be automated partially. This is consistent with the concerns raised in other studies, where respondents anticipated that AI could lead to worldwide job displacement, especially in regular or repetitive work (Paran et al., 2024).

Education has been moving towards increasingly digital solutions since the COVID-19 pandemic (Gryaznov, 2022). Robust generative AI algorithms can already perform translators' jobs. The diagnostic work of doctors can be performed more effectively and precisely with AI (Dvijotham et al., 2023). With a basic programming language understanding and prompt writing, with some exaggeration, many can create programs today using generative AI (Liu et al., 2023). This can replace software developers. These changes in the status quo certainly interest many forumers, regardless of whether they are affected, and are often represented in media outlets. This does not mean that less-educated jobs were not mentioned. As presented in the Results section, topics such as self-driving cars (driver automation) and mention of bank tellers or cashiers (self-check-out) were replaced by machines. Additionally, we could often find instances of discussion about robotization in manufacturing. However, their focus was on intellectual job automation.

When examining sentiments and emotions, we found two main events that changed the dynamics of emotions: the COVID-19 pandemic and release of ChatGPT. The pandemic showed that humans cannot be entirely replaced. In addition, some professions, such as those of essential workers, are more in need than before. Concurrently, it was also a moment of social deprivation and, to some extent, surveillance imposed by technology, which negatively influenced AI perceptions. The pandemic has brought about negative and positive changes in sentiments and emotions: AI has become both the savior and oppressive power, changing from the beginning (hopeful) to the end (deprived) of the pandemic. ChatGPT has also brought about ambiguity in emotions, but on a much broader scale. A massive increase in positive emotions and sentiments was dominated by surprise and simultaneously increased the fear brought about by job replacement/transformation. Our findings are comparable with those of recent research (Ocal and Crowston, 2024), which found that general attitudes towards the impact of AI on work are slightly positive, with curiosity being the most expressed emotion.

Implementing an AI solution in the workplace is a complex challenge from a technological perspective, let alone when we add a human component to the equation. Perceptions are brought from "home" from the external environment and injected into the work environment. Company employees who experience general fear and anger with themselves are likely to resist AI. We could observe that people fear losing their jobs, their humanity being exploited, and sinking into an economic crisis. A detailed investigation of the forums revealed concerns about housing crises, the increasing retirement age, overpopulation, and environmental issues. Companies need to be clear about the goals of using AI applications: Which jobs will be replaced and which ones will be augmented. They need to prepare plans for handling employees who need to change the required job scope skills, participate in further re-education, or be laid off. People need to voice their need for skills, knowledge, and infrastructure to governmental organizations and companies to create policies to protect the human workforce, and avoid negative impacts on society and the economy. However, excessive trust in AI can lead to negative consequences. People are generally surprised by the capabilities of generative AI technology, such as ChatGPT, and have high expectations. However, moderating these expectations is essential. Many technologies are as innovative as their users. Correct prompt writing skills are needed to effectively utilize these tools. People must be trained in data privacy, copyright, and data reliability issues (e.

g., when the tool was trained the last time, or the danger of AI's hallucinations). Additionally, which tasks can be automated by AI, to what extent, and who is responsible for validating the AI work need to be clarified. Employees should not accept anything from generative AI as the ultimate truth. While some forumers believe that generative AI is the end of human critical thinking and creativity, we believe it is the opposite. It has never been more important to approach technology from a critical perspective.

5. Conclusions

We examined a large dataset of 114,377 top-level comments from 37 Reddit forums where users talked about how AI may change the labor market. We examined sentiments and emotions towards the AI-induced labor market changes and performed a topic analysis using BERT. We found four dominant meta-topics: job transformation by AI, general AI topics about the future of work, economy, and politics, and influencers (RQ1). Forumers equally like to discuss how AI is transforming professions in general and individually (e.g., doctors and software developers), and AI's long-term economic and social implications. US politicians, technological company owners (e.g., Bill Gates and Elon Musk), popular culture (e.g., science fiction), and, to some extent, widely popular scientists (e.g., Stephen Hawking) strongly influence the discussions. ChatGPT has drastically transformed the discussion of topics. Forums related to generative AI have a very high volume of comments. Moreover, their associated emotions show intense ambiguity: strong anticipation and surprise but also increasing fear.

Next, we observed generally negative sentiments towards the upcoming and already existing labor market changes driven by the development of AI technologies in the past ten years (RQ2). To better understand sentiments, we conducted emotional detection. Disgust, anger, and surprise (primarily associated with ChatGPT) were the most dominant emotions. At the emotional level, we found two game-changers in the forums: the COVID-19 pandemic and release of ChatGPT. In the topic analysis, the most critical game-changer was the arrival of ChatGPT. While representing only one year of the 10-year study, the topic was the second most significant source of comments (RQ3).

6. Limitations and future research perspectives

First, we only used Reddit as the data source. Future research can consider other social media platforms to overcome the limitations imposed by Reddit's demographics. The results can also be compared with a survey to reveal to what extent social media distorts public perception. Second, we only examined top-level comments, encompassing the ones that received the most appreciation from the community. The platform did not allow us to scrape unmoderated content (comments deleted by the user or by admins). Considering all comments can potentially provide more depth to the analysis. However, for this, the researcher must actively follow and scrape certain daily redits. Finally, although we have tried to gather all relevant redits in the past almost ten years, some were undoubtedly missed or needed to be ignored due to the few comments they received. Social media is a formidable tool for mining public perceptions. However, it also has a "hype effect": only topics deemed interesting by the end users or topics about everyone having an opinion can rank first and receive enough contributions to be worthwhile examining. This can narrow down the perspective of studies.

Credit author statement

Szabina Fodor: Writing – original draft, Visualization, Validation, Supervision, Software, Project administration, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Bettina Boncz:**

Writing – original draft, Visualization, Software, Methodology, Investigation, Formal analysis, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A. Supplementary data

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.dsm.2025.11.001>.

References

- AbuMusab, S., 2024. Generative AI and human labor: who is replaceable? *AI Soc.* 39 (6), 3051–3053.
- Acemoglu, D., Restrepo, P., 2018. Artificial Intelligence, Automation and Work. National Bureau of Economic Research. Working paper 24196. <http://www.nber.org/papers/w24196>. (Accessed 22 October 2025).
- Adoma, A.F., Henry, N.M., Chen, W., 2020. Comparative analyses of Bert, Roberta, Distilbert, and XLNet for text-based emotion recognition. In: 17th International Computer Conference on Wavelet Active Media Technology and Information Processing (ICCWAMTIP). IEEE, pp. 117–121.
- Ahlqvist, T., 2005. From information society to biosociety? On societal waves, developing key technologies, and new professions. *Technol. Forecast. Soc. Change* 72 (5), 501–519.
- Aktan, M.E., Turhan, Z., Dolu, İ., 2022. Attitudes and perspectives towards the preferences for artificial intelligence in psychotherapy. *Comput. Hum. Behav.* 133 (Aug.), 107273.
- Al-Emran, M., Abu-Hijleh, B., Alsewari, A.A., 2025. Examining the impact of generative AI on social sustainability by integrating the information system success model and technology-environmental, economic, and social sustainability theory. *Educ. Inf. Technol.* 30 (7), 9405–9426.
- Albarrán Lozano, I., Molina, J.M., Gijón, C., 2021. Perception of artificial intelligence in Spain. *Telematics Inf.* 63 (Oct.), 101672.
- Aldous, K.K., An, J., Jansen, B.J., 2023. What really matters? Characterising and predicting user engagement of news postings using multiple platforms, sentiments and topics. *Behav. Inf. Technol.* 42 (5), 545–568.
- Alekseeva, L., Azar, J., Giné, M., et al., 2021. The demand for AI skills in the labor market. *Labour Econ.* 71 (Aug.), 102002.
- Allen, R.C., 2017. Lessons from history for the future of work. *Nature* 550 (7676), 321–324.
- Aloisi, A., De Stefano, V., 2022. Essential jobs, remote work and digital surveillance: addressing the COVID-19 pandemic panopticon. *Int. Labour Rev.* 161 (2), 289–314.
- Aoki, N., 2021. The importance of the assurance that “humans are still in the decision loop” for public trust in artificial intelligence: evidence from an online experiment. *Comput. Hum. Behav.* 114 (Jan.), 106572.
- Arntz, M., Gregory, T., Zierahn, U., 2016. The Risk of Automation for Jobs in OECD Countries: a Comparative Analysis. OECD Social, Employment and Migration Working Papers. OECD Publishing, Paris. No. 189.
- Atalay, E., Phongthientham, P., Sotelo, S., et al., 2018. New technologies and the labor market. *J. Monetary Econ.* 97 (Aug.), 48–67.
- Bain, M., McCay, A., 2024. The neural democratisation of AI. *AI Soc.* 39 (5), 2589–2591.
- Bárány, Z.L., Siegel, C., 2020. Biased technological change and employment reallocation. *Labour Econ.* 67 (Dec.), 101930.
- Barbieri, F., Camacho-Collados, J., Neves, L., et al., 2020. TweetEval: unified benchmark and comparative evaluation for tweet classification. [arXiv:2010.12421](https://arxiv.org/abs/2010.12421).
- Barbieri, F., Anke, L.E., Camacho-Collados, J., 2021. Multilingual language models in Twitter for sentiment analysis and beyond. [arxiv.2104.12250](https://arxiv.org/abs/2104.12250).
- Battisti, S., Agarwal, N., Brem, A., 2022. Creating new tech entrepreneurs with digital platforms: meta-organizations for shared value in data-driven retail ecosystems. *Technol. Forecast. Soc. Change* 175 (Feb.), 121392.
- Benvenuti, M., Cangelosi, A., Weinberger, A., et al., 2023. Artificial intelligence and human behavioral development: a perspective on new skills and competences acquisition for the educational context. *Comput. Hum. Behav.* 148 (Nov.), 107903.
- Betlen, A., 2023. llama-cpp-python. <https://github.com/abetlen/llama-cpp-python>. (Accessed 22 October 2025).
- Bhargava, A., Bester, M., Bolton, L., 2021. Employees' perceptions of the implementation of robotics, artificial intelligence, and automation (RAIA) on job satisfaction, job security, and employability. *J. Technol. Behav. Sci.* 6 (1), 106–113.
- Blei, D.M., Ng, A.Y., Jordan, M.I., 2003. Latent dirichlet allocation. *J. Mach. Learn. Res.* 3 (Mar.), 993–1022.
- Bond, M., Khosravi, H., De Laat, M., et al., 2024. A meta systematic review of artificial intelligence in higher education: a call for increased ethics, collaboration, and rigour. *Int. J. Educ. Technol. High. Educ.* 21 (1), 4.
- Botzer, N., Gu, S., Weninger, T., 2023. Analysis of moral judgment on Reddit. *IEEE Trans. Comput. Soc. Syst.* 10 (3), 947–957.
- Brett, E.I., Stevens, E.M., Wagener, T.L., et al., 2019. A content analysis of JUUL discussions on social media: using Reddit to understand patterns and perceptions of JUUL use. *Drug Alcohol Depend.* 194 (Jan.), 358–362.
- Bubeck, S., Chandrasekaran, V., Eldan, R., et al., 2023. Sparks of artificial general intelligence: early experiments with GPT-4. [arXiv:2303.12712](https://arxiv.org/abs/2303.12712).
- Casal-Otero, L., Catala, A., Fernández-Morante, C., et al., 2023. AI literacy in K-12: a systematic literature review. *Int. J. STEM Educ.* 10 (1), 29.
- Cavaglia, C., Etheridge, B., 2020. Job polarization and the declining quality of knowledge workers: evidence from the UK and Germany. *Labour Econ.* 66 (Oct.), 101884.
- Cave, S., Dihal, K., 2019. Hopes and fears for intelligent machines in fiction and reality. *Nat. Mach. Intell.* 1 (2), 74–78.
- Chang, W., 2022. The effectiveness of AI salesperson vs. human salesperson across the buyer-seller relationship stages. *J. Bus. Res.* 148 (Sep.), 241–251.
- Chen, Y., Zhang, H., Liu, R., et al., 2019. Experimental explorations on short text topic mining between LDA and NMF based Schemes. *Knowl. Base Syst.* 163 (Jan.), 1–13.
- Cheng, B., Lin, H., Kong, Y., 2023. Challenge or hindrance? How and when organizational artificial intelligence adoption influences employee job crafting. *J. Bus. Res.* 164 (Sep.), 113987.
- Chettri, S., Deb Nath, D., Devi, P., 2020. Leveraging digital tools and technologies to alleviate COVID-19 pandemic. *SSRN Electron. J.* <https://doi.org/10.2139/ssrn.3626092>.
- Chigbu, B.I., Makapela, S.L., 2025. AI in education, sustainability, and the future of work: an integrative review of industry 5.0, education 5.0, and work 5.0. *J. Open Innov. Technol. Mark. Complex.* 11 (4), 100645.
- Chiu, Y.T., Zhu, Y.Q., Corbett, J., 2021. In the hearts and minds of employees: a model of pre-adoptive appraisal toward artificial intelligence in organizations. *Int. J. Inf. Manag.* 60 (Oct.), 102379.
- Choi, T., Leigh, N.G., 2024. Artificial intelligence's creation and displacement of labor demand. *Technol. Forecast. Soc. Change* 209 (Dec.), 123824.
- Coates, J.F., 1998. Ready children for the future. *Technol. Forecast. Soc. Change* 57 (1/2), 157–165.
- Cooper, R.G., 2024. The AI transformation of product innovation. *Ind. Mark. Manag.* 119 (May), 62–74.
- Crawford, K., Calo, R., 2016. There is a blind spot in AI research. *Nature* 538 (7625), 311–313.
- Damioli, G., Van Roy, V., Vértessy, D., et al., 2024. Drivers of employment dynamics of AI innovators. *Technol. Forecast. Soc. Change* 201 (Apr.), 123249.
- David, B., 2017. Computer technology and probable job destructions in Japan: an evaluation. *J. Jpn. Int. Econ.* 43 (Mar.), 77–87.
- De Obesso Arias, M., Pérez Rivero, C.A., Carrero Márquez, O., 2023. Artificial intelligence to manage workplace bullying. *J. Bus. Res.* 160 (May), 113813.
- Dennehy, D., Griva, A., Pouloudi, N., et al., 2023. Artificial intelligence for decision-making and the future of work. *Int. J. Inf. Manag.* 69 (Apr.), 102574.
- Devlin, J., Chang, M.W., Lee, K., et al., 2019. BERT: pre-training of deep bidirectional transformers for language understanding. In: Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies. ACL, pp. 4171–4186.
- Dirican, C., 2015. The impacts of robotics, artificial intelligence on business and economics. *Procedia Soc. Behav. Sci.* 195 (Jul.), 564–573.
- Downey, M., 2021. Partial automation and the technology-enabled deskilling of routine jobs. *Labour Econ.* 69 (Apr.), 101973.
- Duch-Brown, N., Gomez-Herrera, E., Mueller-Langer, F., et al., 2022. Market power and artificial intelligence work on online labour markets. *Res. Pol.* 51 (3), 104446.
- Dvijotham, K.D., Winkens, J., Barsbey, M., et al., 2023. Enhancing the reliability and accuracy of AI-enabled diagnosis via complementarity-driven deferral to clinicians. *Nat. Med.* 29 (7), 1814–1820.
- Egger, R., 2022. Topic modelling. In: Egger, R. (Ed.), *Applied Data Science in Tourism*. Tourism on the Verge. Springer, Cham, pp. 375–403.
- Eldakruri, T., Senyurek, E., 2025. Evaluating the economic feasibility of labor replacement through robotics and automation in Qatar. [arXiv:2509.10152](https://arxiv.org/abs/2509.10152).
- Fast-Berglund, Å., Salunkhe, O., Åkerman, M., 2020. Low-cost automation—changing the traditional view on automation strategies using collaborative applications. *IFAC-PapersOnLine* 53 (2), 10285–10290.
- Felten, E.W., Raj, M., Seamans, R., 2023. Occupational heterogeneity in exposure to generative AI. *SSRN Electron. J.* <https://doi.org/10.2139/ssrn.4414065>.
- Flathmann, C., Schelble, B.G., Rosopa, P.J., et al., 2023. Examining the impact of varying levels of AI teammate influence on human-AI teams. *Int. J. Hum. Comput. Stud.* 177 (Sep.), 103061.
- Focacci, C.N., 2021. Technological unemployment, robotisation, and green deal: a story of unstable spillovers in China and South Korea (2008–2018). *Technol. Soc.* 64 (Feb.), 101504.
- Frey, C.B., Osborne, M.A., 2017. The future of employment: how susceptible are jobs to computerisation? *Technol. Forecast. Soc. Change* 114 (Jan.), 254–280.
- Frey, C.B., Osborne, M., 2023. Generative AI and the future of work: a reappraisal. *Brown J. World Aff.* 30 (1), 1–17.
- Fui-Hoon Nah, F., Zheng, R., Cai, J., et al., 2023. Generative AI and ChatGPT: applications, challenges, and AI-human collaboration. *J. Inf. Technol. Case Appl. Res.* 25 (3), 277–304.
- Furendal, M., Jebbari, K., 2023. The future of work: augmentation or stunting? *Philos. Technol.* 36 (2), 36.
- Gama, F., Magistretti, S., 2025. Artificial intelligence in innovation management: a review of innovation capabilities and a taxonomy of AI applications. *J. Prod. Innovat. Manag.* 42 (1), 76–111.

- Garcia-Murillo, M., MacInnes, I., Bauer, J.M., 2018. Techno-unemployment: a framework for assessing the effects of information and communication technologies on work. *Teleomatics Inf.* 35 (7), 1863–1876.
- Gasteiger, E., Prettner, K., 2022. Automation, stagnation, and the implications of a robot tax. *Macroecon. Dynam.* 26 (1), 218–249.
- Gignac, G.E., Szodorai, E.T., 2024. Defining intelligence: bridging the gap between human and artificial perspectives. *Intelligence* 104 (May), 101832.
- 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.
- Goos, M., Manning, A., Salomons, A., 2014. Explaining job polarization: routine-biased technological change and offshoring. *Am. Econ. Rev.* 104 (8), 2509–2526.
- Grootendorst, M., 2022. BERTopic: neural topic modeling with a class-based TF-IDF procedure. *arXiv:2203.05794*.
- Gryaznov, S.A., 2022. How digital technologies are changing business education. In: *Digital Technologies in the New Socio-Economic Reality*. Springer, Cham, pp. 801–807. *ISCDE* 2021.
- Harari, Y.N., 2017. Reboot for the AI revolution. *Nature* 550 (7676), 324–327.
- He, L., Yin, M., Shi, Y., 2020. Love, hate thy neighbour? Or just don't care much about them: a sentiment analysis of China-related posts and comments on Reddit. *Com. China Rep* 56 (2), 204–220.
- Hofmann, T., 2001. Unsupervised learning by probabilistic latent semantic analysis. *Mach. Learn.* 42 (1), 177–196.
- Hughes, J., 2014. A strategic opening for a basic income guarantee in the global crisis being created by AI, robots, desktop manufacturing and BioMedicine. *J. Eth. Emerg. Tech.* 24 (1), 45–61.
- Hui, X., Reshef, O., Zhou, L., 2024. The short-term effects of generative artificial intelligence on employment: evidence from an online labor market. *Organ. Sci.* 35 (6), 1977–1989.
- Ilić, D., Gignac, G.E., 2024. Evidence of interrelated cognitive-like capabilities in large language models: indications of artificial general intelligence or achievement? *Intelligence* 106 (Sep.), 101858.
- Jaradat, S., Matskin, M., 2019. On dynamic topic models for mining social media. In: *Emerging Research Challenges and Opportunities in Computational Social Network Analysis and Mining*. Lecture Notes in Social Networks. Springer, Cham, pp. 209–230.
- Jazdauskaitė, J., Privarova, M., Baranskaite, E., et al., 2021. Evaluation of the impact of science and technology on the labour market. *Mark. Manag. Innov.* 5 (4), 153–167.
- Jeon, E., Yoon, N., Sohn, S.Y., 2023. Exploring new digital therapeutics technologies for psychiatric disorders using BERTopic and PatentSBERTa. *Technol. Forecast. Soc. Change* 186 (Jan.), 122130.
- Joržik, P., Klein, S.P., Kanbach, D.K., et al., 2024. AI-driven business model innovation: a systematic review and research agenda. *J. Bus. Res.* 182 (Sep.), 114764.
- Kherwa, P., Bansal, P., 2018. Topic modeling: a comprehensive review. *ICST Trans. Scalable Inf. Syst.* (Jul.), 159623.
- Khogali, H.O., Mekid, S., 2023. The blended future of automation and AI: examining some long-term societal and ethical impact features. *Technol. Soc.* 73 (May), 102232.
- Kim, J., Im, I., 2023. Anthropomorphic response: understanding interactions between humans and artificial intelligence agents. *Comput. Hum. Behav.* 139 (Feb.), 107512.
- Kim, S., Lee, K., 2023. The paradigm shift of mass customisation research. *Int. J. Prod. Res.* 61 (10), 3350–3376.
- Kim, Y.J., Kim, K., Lee, S., 2017. The rise of technological unemployment and its implications on the future macroeconomic landscape. *Futures* 87 (Mar.), 1–9.
- Kniffin, K.M., Narayanan, J., Anseel, F., et al., 2021. COVID-19 and the workplace: implications, issues, and insights for future research and action. *Am. Psychol.* 76 (1), 63–77.
- Kolade, O., Owoseni, A., 2022. Employment 5.0: the work of the future and the future of work. *Technol. Soc.* 71 (Nov.), 102086.
- Korinek, A., Juefils, M., 2024. Preparing for the (Non-Existent?) Future of Work. In: Bullock, J.B., Chen, Y.-C., Himmelreich, J., Hudson, V.M., Korinek, A., et al. (Eds.), *The Oxford Handbook of AI Governance*. Oxford University Press, Oxford.
- Kshetri, N., Dwivedi, Y.K., Davenport, T.H., et al., 2024. Generative artificial intelligence in marketing: applications, opportunities, challenges, and research agenda. *Int. J. Inf. Manag.* 75 (Apr.), 102716.
- Kutyauripo, I., Rushambwa, M., Chivazi, L., 2023. Artificial intelligence applications in the agrifood sectors. *J. Agric. Food Res.* 11 (Mar.), 100502.
- Lau, J.H., Newman, D., Baldwin, T., 2014. Machine reading tea leaves: automatically evaluating topic coherence and topic model quality. In: *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics*. ACL, pp. 530–539.
- Lee, D., Seung, H.S., 2000. Algorithms for non-negative matrix factorization. *Adv. Neural Inf. Process. Syst.* 13. <https://www.mlpack.org/papers/nmf.pdf>. (Accessed 22 October 2025).
- Leichtmann, B., Humer, C., Hinterreiter, A., et al., 2023. Effects of explainable artificial intelligence on trust and human behavior in a high-risk decision task. *Comput. Hum. Behav.* 139 (Feb.), 107539.
- Lima, Y., Barbosa, C.E., dos Santos, H.S., et al., 2021. Understanding technological unemployment: a review of causes, consequences, and solutions. *Societies* 11 (2), 50.
- Liu, L., Tang, L., Dong, W., et al., 2016. An overview of topic modeling and its current applications in bioinformatics. *SpringerPlus* 5 (1), 1608.
- Liu, Y., Ott, M., Goyal, N., et al., 2019. RoBERTa: a robustly optimized BERT pretraining approach. *arXiv:1907.11692*.
- Liu, J., Xia, C.S., Wang, Y., et al., 2023. Is your code generated by ChatGPT really correct? Rigorous evaluation of large language models for code generation. In: *Proceedings of the 37th International Conference on Neural Information Processing Systems*. ACM, pp. 21558–21572.
- Long, C., Lucey, B.M., Yarovaya, L., 2021. I just like the stock” versus “fear and loathing on main street”: the role of Reddit sentiment in the GameStop short squeeze. *SSRN Electron. J.* <https://doi.org/10.2139/SSRN.3822315>.
- Mäkelä, E., Stephany, F., 2024. Complement or substitute? How AI increases the demand for human skills. *arXiv:2412.19754*.
- Melton, C.A., Olusanya, O.A., Ammar, N., et al., 2021. Public sentiment analysis and topic modeling regarding COVID-19 vaccines on the Reddit social media platform: a call to action for strengthening vaccine confidence. *J. Infect. Public Health* 14 (10), 1505–1512.
- Mirbabaie, M., Brünker, F., Möllmann Frick, N.R.J., et al., 2022. The rise of artificial intelligence—understanding the AI identity threat at the workplace. *Electron. Mark.* 32 (1), 73–99.
- Mitchell, T., Brynjolfsson, E., 2017. Track how technology is transforming work. *Nature* 544 (7650), 290–292.
- Mutasu, M., 2021. Artificial intelligence and unemployment: new insights. *Econ. Anal. Pol.* 69 (Mar.), 653–667.
- Grootendorst, M., 2020. KeyBERT: Minimal Keyword Extraction with BERT. <https://zenodo.org/records/4461265>. (Accessed 22 October 2025)
- Nader, K., Toprac, P., Scott, S., et al., 2022. Public understanding of artificial intelligence through entertainment media. *AI Soc.* 39 (Apr.), 1–14.
- Nazareno, L., Schiff, D.S., 2021. The impact of automation and artificial intelligence on worker well-being. *Technol. Soc.* 67 (Nov.), 101679.
- Neri, H., Cozman, F., 2020. The role of experts in the public perception of risk of artificial intelligence. *AI Soc.* 35 (3), 663–673.
- Noy, S., Zhang, W., 2023. Experimental evidence on the productivity effects of generative artificial intelligence. *Science* 381 (6654), 187–192.
- Obadimu, A., Mead, E., Agarwal, N., 2019. Identifying latent toxic features on YouTube using non-negative matrix factorization. In: *SOTICS 2019: the Ninth International Conference on Social Media Technologies, Communication, and Informatics*. IARIA, pp. 25–31.
- Ocal, A., Crowston, K., 2024. Framing and feelings on social media: the futures of work and intelligent machines. *Inf. Technol. People* 37 (7), 2462–2488.
- OECD, 2021. What Happened to Jobs at High Risk of Automation? Organisation for Economic Co-operation and Development. https://www.oecd.org/en/publication/s/what-happened-to-jobs-at-high-risk-of-automation_10bc97f4-en.html. (Accessed 22 October 2025).
- Orchard, T., Tasiemski, L., 2023. The rise of generative AI and possible effects on the economy. *Econ. Bus. Rev.* 9 (2), 9–26.
- Paran, L.R.L., Maleptey, J.C., De Leon, J.L., et al., 2024. Discovering the future and impact of artificial intelligence: a quantitative exploration. *Cogniz. J. Multidiscip. Stud.* 4 (11), 148–157.
- Pianta, M., Vivarelli, M., 2000. Unemployment, Structural Change and Globalization. *International Labour Organization*. https://training.ilo.org/actrav_cdrom1/eng/sh/global/art/3.htm. (Accessed 22 October 2025).
- Raisch, S., Krakowski, S., 2020. Artificial intelligence and management: the automation-augmentation paradox. *Acad. Manag. Rev.* 46 (1), 192–210.
- Ray, D., Mookherjee, D., 2022. Growth, automation, and the long-run share of labor. *Rev. Econ. Dynam.* 46 (Oct.), 1–26.
- Revilla-León, M., Gómez-Polo, M., Vyas, S., et al., 2023. Artificial intelligence applications in implant dentistry: a systematic review. *J. Prosthet. Dent* 129 (2), 293–300.
- Röder, M., Both, A., Hinneburg, A., 2015. Exploring the space of topic coherence measures. In: *Proceedings of the Eighth ACM International Conference on Web Search and Data Mining*. ACM, pp. 399–408.
- Rodríguez-Ibáñez, M., Casáñez-Ventura, A., Castejón-Mateos, F., et al., 2023. A review on sentiment analysis from social media platforms. *Expert Syst. Appl.* 223 (Aug.), 119862.
- Rosner, F., Hinneburg, A., Röder, M., et al., 2014. Evaluating topic coherence measures. *arXiv:1403.6397*.
- Rzekeć, A., Vial, C., Bigot, G., 2020. Green assets of equines in the European context of the ecological transition of agriculture. *Animals* 10 (1), 106.
- Saeed, W., Omlin, C., 2023. Explainable AI (XAI): a systematic meta-survey of current challenges and future opportunities. *Knowl. Base Syst.* 263 (Mar.), 110273.
- Sammet, J., Krestel, R., 2021. Domain-specific keyword extraction using BERT. In: *International Conference on Language, Data, and Knowledge*. ACL, pp. 659–665.
- Sánchez-Franco, M.J., Rey-Moreno, M., 2022. Do travelers' reviews depend on the destination? An analysis in coastal and urban peer-to-peer lodgings. *Psychol. Market.* 39 (2), 441–459.
- Sanfilippo, F., Hamza Zafar, M., Wiley, T., et al., 2025. From caged robots to high-fives in robotics: exploring the paradigm shift from human-robot interaction to human-robot teaming in human-machine interfaces. *J. Manuf. Syst.* 78 (Feb.), 1–25.
- Savelle, N., Pellert, M., Latikka, R., et al., 2024. Affective, cognitive, and contextual cues in Reddit posts on artificial intelligence. *J. Comput. Soc. Sci.* 8 (1), 6.
- Segal, M., 2018. How automation is changing work. *Nature* 563 (7733), 132–135.
- Sheikh, H., Prins, C., Schrijvers, E., 2023. Artificial intelligence: definition and background. In: *Mission, A.I. (Ed.), Research for Policy*. Springer, Cham, pp. 15–41.
- Singh, A., Glińska-Neweś, A., 2022. Modeling the public attitude towards organic foods: a big data and text mining approach. *J. Big Data* 9 (1), 2.
- Sinha, N., Singh, P., Gupta, M., et al., 2020. Robotics at workplace: an integrated Twitter analytics-SEM based approach for behavioral intention to accept. *Int. J. Inf. Manag.* 55 (Dec.), 102210.
- Sowa, K., Przegalińska, A., Ciechanowski, L., 2021. Cobots in knowledge work human-AI collaboration in managerial professions. *J. Bus. Res.* 125 (Mar.), 135–142.
- Su, C.W., Yuan, X., Umar, M., et al., 2022. Does technological innovation bring destruction or creation to the labor market? *Technol. Soc.* 68 (Feb.), 101905.

- Sun, Q., Wojcieszak, M., Davidson, S., 2021. Over-time trends in incivility on social media: evidence from political, non-political, and mixed sub-reddits over eleven years. *Front. Polit. Sci.* 3 (Nov.), 741605.
- Thompson, L., Mimno, D., 2020. Topic modeling with contextualized word representation clusters. *arXiv:2010.12626*.
- Ting, D.S.W., Carin, L., Dzau, V., et al., 2020. Digital technology and COVID-19. *Nat. Med.* 26 (4), 459–461.
- Țițan, E., Burciu, A., Manea, D., et al., 2014. From traditional to digital: the labour market demands and education expectations in an EU context. *Procedia Econ. Finance* 10, 269–274.
- Triguero, I., Molina, D., Poyatos, J., et al., 2024. General purpose artificial intelligence systems (GPAIS): properties, definition, taxonomy, societal implications and responsible governance. *Inf. Fusion* 103 (Mar.), 102135.
- Upadhyay, V., 2021. Can capitalism survive the high degree of automation? A comparison with Thomas Piketty's argument. In: Upadhyay, V., Singh, P. (Eds.), *Global Political Economy*. Routledge, London, pp. 179–191.
- Venugopalan, M., Gupta, D., 2022. An enhanced guided LDA model augmented with BERT based semantic strength for aspect term extraction in sentiment analysis. *Knowl. Base Syst.* 246 (Jun.), 108668.
- Wilson, H.J., Daugherty, P.R., Morini-Bianzino, N., 2017. The jobs that artificial intelligence will create. *MIT Sloan Manag. Rev.* 58 (Mar.), 14–16.
- Yadav, A., Vishwakarma, D.K., 2020. Sentiment analysis using deep learning architectures: a review. *Artif. Intell. Rev.* 53 (6), 4335–4385.
- Yin, B., Yuan, C.H., 2022. Detecting latent topics and trends in blended learning using LDA topic modeling. *Educ. Inf. Technol.* 27 (9), 12689–12712.
- Zhang, Y., Zhang, L., 2022. Movie recommendation algorithm based on sentiment analysis and LDA. *Procedia Comput. Sci.* 199, 871–878.
- Zhang, X., Qi, X., Teng, Z., 2025. Performance evaluation of Reddit comments using machine learning and natural language processing methods in sentiment analysis. In: *Computational and Experimental Simulations in Engineering. ICCES 2024*. Springer, Cham, pp. 14–24.
- Zhao, X., Wang, D., Zhao, Z., et al., 2021. A neural topic model with word vectors and entity vectors for short texts. *Inf. Process. Manag.* 58 (2), 102455.