

# Natural Language Processing for Human Resources: A Survey

Naoki Otani   Nikita Bhutani   Estevam Hruschka

Megagon Labs

{naoki, nikita, estevam}@megagon.ai

## Abstract

Advances in Natural Language Processing (NLP) have the potential to transform HR processes, from recruitment to employee management. While recent breakthroughs in NLP have generated significant interest in its industrial applications, a comprehensive overview of how NLP can be applied across HR activities is still lacking. This paper discovers opportunities for researchers and practitioners to harness NLP’s transformative potential in this domain. We analyze key fundamental tasks such as information extraction and text classification, and their roles in downstream applications like recommendation and language generation, while also discussing ethical concerns. Additionally, we identify gaps in current research and encourage future work to explore holistic approaches for achieving broader objectives in this field.

## 1 Introduction

Human Resources (HR) is a vital component of any organization, responsible for managing its most valuable resource—people. Over the years, computational tools have transformed HR processes like hiring, training, and administration, reshaping the labor market and workplace. At the same time, concerns about the accuracy and fairness of automated systems have also garnered significant attention, paving the way for ongoing and future research. Advancements in Natural Language Processing (NLP), especially with large language models (LLMs), have spurred interest in applying language technologies to a broad range of real-world problems, and the HR domain is no exception. However, this domain remains relatively underrepresented in the NLP research community.<sup>1</sup>

As breakthroughs in LLMs continue to advance various aspects of NLP, key challenges in the HR

<sup>1</sup>Despite the development of many innovative applications in the industry (Barth, 2024), major conferences such as ACL, NAACL, EMNLP, EACL, and COLING featured only three papers with “job” or “human resources” in their titles in 2024.



Figure 1: **Concept of this survey paper.** We review and categorize HR-related problems through the lens of core NLP research areas.

domain, such as the complexity of processing heterogeneous data, and the scarcity of publicly available data resources, may be alleviated in the coming years. Therefore, the HR domain holds substantial potential for growth and also presents unique challenges that can drive NLP research forward. To facilitate this transformation, it is essential to develop a comprehensive overview of key HR activities from an NLP perspective and examine how upstream tasks, such as skill extraction, contribute to downstream applications like job matching.

In this paper, we analyze HR activities through the lens of NLP research, categorizing them into key areas and examining how NLP techniques have been applied, along with remaining challenges (Figure 1).<sup>2</sup> We explore fundamental tasks like information extraction and text classification (§3), and their role in supporting core applications such as recommendation, language generation, and interaction (§4). Finally, we highlight underrepresented areas in NLP to guide future research (§5). By organizing the discussion around NLP research topics, our goal is to provide insights for two audiences: (1) NLP researchers seeking impactful problems in the HR domain, and (2) those exploring how NLP can address HR challenges.

<sup>2</sup>NLP research is relevant to various HR activities. However, most existing studies focus primarily on talent acquisition, which is why this topic receives greater emphasis in the paper.

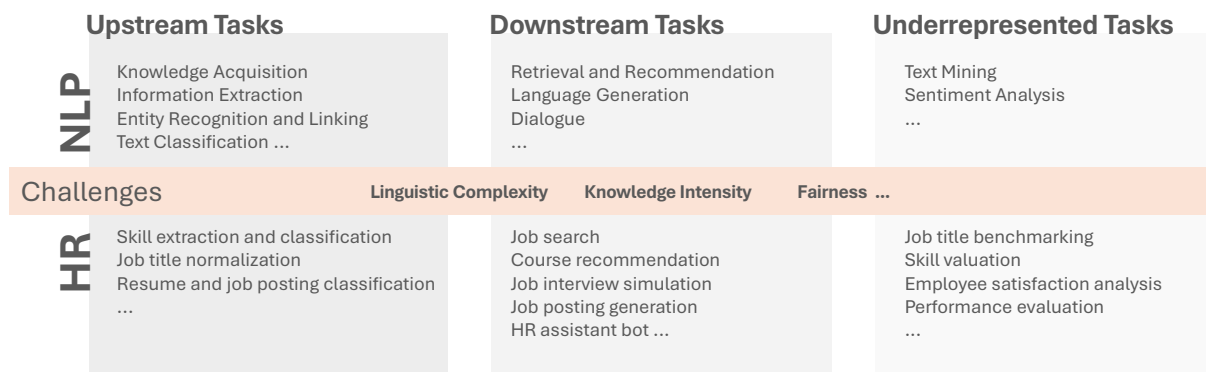


Figure 2: Landscape of NLP applications within the HR domain.

Previous surveys on this topic have typically focused on specific HR tasks and applications, such as information extraction from job postings (Khaouja et al., 2021; Senger et al., 2024), market analysis (Rahhal et al., 2024), job recommendation (Balog et al., 2012; de Ruijt and Bhulai, 2021; Freire and de Castro, 2021; Mashayekhi et al., 2024), conversational agents (Laumer and Morana, 2022), and fairness (Hunkenschroer and Luetge, 2022; Kumar et al., 2023; Fabris et al., 2024). While general literature reviews in this field provide a broad overview of relevant computational research (Budhwar et al., 2022; Sharma, 2021; Qin et al., 2024; Khan, 2024), they do not specifically explore insights into language technologies. In contrast, we focus on core NLP methodologies, such as information extraction, text classification, retrieval, and language generation, and discuss their evolving role in various HR applications.<sup>3</sup>

This paper provides a structured NLP-centric perspective that systematically maps NLP tasks to HR challenges, making it easier for readers with an NLP background to identify relevant research opportunities and for HR practitioners to connect with relevant methods.<sup>4</sup> We highlight how specialized tasks contribute to broader goals, such as job title understanding for skill extraction and skill extraction for job matching, and encourage future work to explore holistic approaches for achieving these objectives. To further advance this field, we recommend the development and use of real or real(istic) datasets to enhance the relevance and impact of research outcomes.

<sup>3</sup>While a position paper by Leidner and Stevenson (2024) also explores NLP applications in this field, it does not provide a comprehensive literature review.

<sup>4</sup>We describe our literature survey methodology in Appendix A.

## 2 What is HR Concerned with?

This section briefly describes HR activities and their links to NLP. Broadly, these activities can be categorized into pre-hiring and post-hiring tasks.

**Pre-hiring:** The pre-hiring process for recruiters includes drafting job postings, selecting candidates, conducting interviews, and extending offers. For job seekers, it involves exploring market trends, pursuing necessary training, preparing resumes, applying for jobs, preparing for interviews, and negotiating offers. These tasks rely heavily on nuanced domain-specific knowledge and are closely related to language generation (e.g., writing job postings and resumes, text-based communication) and specialized dialogue (e.g., interviews).

**Post-hiring:** Key HR functions include setting role requirements aligned with organizational goals, evaluating performance, optimizing team dynamics, and maintaining positive work environments. These tasks are complex, demanding occupation-specific insights and integration of diverse data sources like employee records, organizational network, and textual communications.

The application of NLP techniques for these activities faces several challenges: (1) **Diverse entities and language expressions** in HR data, such as the concise, bullet-pointed style of resumes or performance feedback, which vary across industries (e.g., software development vs. culinary arts). (2) The need for deep understanding of **domain-specific knowledge**, which is often not readily available in raw text corpora. (3) **Biases in data-driven systems**, reflecting stereotypes, proxies for sensitive attributes, and external barriers (Calanca et al., 2019; Glazko et al., 2024).

The following sections review existing research on HR activities, organized by NLP topics, with

a focus on upstream tasks (§3) and downstream tasks (§4). We then discuss underrepresented HR activities (§5) that could benefit from recent NLP advancements (Figure 2).

### 3 Upstream Tasks

Upstream HR tasks aim to enrich raw text corpora through information extraction and classification to facilitate knowledge-intensive downstream tasks.

#### 3.1 Taxonomy Creation

Significant efforts have been made to acquire domain-specific knowledge and develop HR-related taxonomies to organize information on occupations, industries, skills, education, and certifications. This has led to the creation of large-scale resources such as the European Skills, Competences, Qualifications and Occupations (ESCO; le Vrang et al., 2014) and others (Lau and Sure, 2002; International Labour Organization, 2012; Bastian et al., 2014; National Center for O\*NET Development). Expert-driven taxonomy creation can yield high-quality resources, but maintaining them is challenging. To reduce the costs, some studies have used Wikipedia (Kivimäki et al., 2013; Zhao et al., 2015) and the consolidation of web resources (Gugnani and Misra, 2020). However, taxonomy creation remains highly complex due to cultural and regional variations (Tu and Cannon, 2022).

#### 3.2 Information Extraction

The extraction of HR-related information, particularly job-related skills, has been extensively studied in the research community (Khaouja et al., 2021; Senger et al., 2024). Skills include a range of competencies, such as technical expertise, knowledge, and the ability to learn and apply new concepts<sup>5</sup>. Other studies have also focused on extracting information like work experience and education (De Sitter and Daelemans, 2003; Finn and Kushmerick, 2004; Green et al., 2022).

This challenge is often framed as a sequential labeling problem with models trained on in-domain corpora (Sayfullina et al., 2018; Green et al., 2022; Zhang et al., 2022). Recent studies have explored multi-task and transfer learning (Fang et al., 2023; Zhang et al., 2023, 2024a) to address the diversity and long-tail nature of job-related information. For

<sup>5</sup>Some literature differentiates skills from knowledge, competencies, and qualifications, but for simplicity, we consider skills to encompass all types of proficiency.

extraction from resumes, the use of layout information has proven useful. Early work by Yu et al. (2005) introduced a two-pass model that segments and labels resume sections before identifying specific details. A similar approach is adopted by Yao et al. (2023) for extracting information from resumes in PDF format.

#### 3.3 Classification and Entity Linking

**Classification of job-related documents** plays a crucial role, especially in hiring, by organizing the large volumes of content generated by job seekers and recruiters. Previous research has focused on classification tasks such as categorizing resumes by job type (Inoubli and Brun, 2022) and sorting job postings into occupation categories (Lake, 2022). Text classification within documents—such as detecting section types (Wang et al., 2022) or analyzing work experience details (Li et al., 2020a)—can also be useful for downstream applications like job recommendation. Automated text classification methods have already been widely used in society as part of Applicant Tracking Systems (ATS), which has also drawn attention to their potential bias issues (§4.4).

**Job title normalization** involves consolidating job titles expressed into a finite set of occupation categories. Prior work has addressed this by incorporating skill information (Decorte et al., 2021) and behavioral data into computational modeling (Liu et al., 2020a; Ha et al., 2020; *inter-alia*). For example, Zhang et al. (2019) integrated a job transition graph to model compositional meaning of job titles, while Zhu and Hudelot (2022) enhanced this graph by further adding edges from component words. Recent studies have demonstrated the effectiveness of Transformer-based text encoders (Yamashita et al., 2023; Laosaengpha et al., 2024), yet this task remains challenging due to issues such as the length of documents, the presence of irrelevant information (e.g., location), and the reliance on domain knowledge.

A similar task is **skill classification**, which involves mapping texts to a pre-defined taxonomy like ESCO (le Vrang et al., 2014). Some studies have employed methods based on similarity matching, while others have formulated the task as a multi-label classification problem (Senger et al., 2024). A notable challenge in this task is handling diverse skill labels.<sup>6</sup> Zhang et al. (2024b)

<sup>6</sup>For example, ESCO v1.2.0 contains 13,939 skills.




Figure 3: The problem of **job recommendation** (§4.1) is a two-sided process relying on multiple facets of information, such as expertise and requirements. Even if a job seeker prefers a particular job, the candidate may not necessarily be the best fit for the position.

demonstrated that entity linking models trained on Wikipedia data can be effectively adapted to the HR domain. Other studies have explored implicit relationships between occupations and skills to improve skill identification. [Bhola et al. \(2020\)](#) used a bootstrapping technique leveraging skill co-occurrence, while [Goyal et al. \(2023\)](#) used a job-skill graph to capture implicit relationships between skills. To collect training data efficiently, [Decorte et al. \(2022\)](#) proposed distant supervision, and recent studies have used LLMs to synthesize annotated texts ([Decorte et al., 2023](#); [Clavié and Soulié, 2023](#); [Magron et al., 2024](#)).

### 3.4 Summary

Upstream HR tasks face challenges such as language complexity and diversity, varying types of data, and insufficient labeled data for training. While existing research has introduced innovative approaches to address these issues, some challenges remain underexplored. These include handling implicit information (e.g., inferring job requirements like a “driver’s license” for truck drivers) and scaling extraction methods to accommodate emerging jobs and skills.

## 4 Downstream Tasks

Downstream HR applications broadly leverage NLP techniques across retrieval and recommendation, language generation, and dialogue systems. This section delves into these areas, followed by a discussion on the challenges of fairness and bias within these tasks.

### 4.1 Retrieval and Recommendation

**Job recommendation** (or Person-Job fit) is typically framed as a text matching problem between job descriptions and resumes, addressed by various encoding methods such as word/document

vectors ([Elsafty et al., 2018](#); [Zhu et al., 2018a](#); [Mogenet et al., 2019](#)) and Transformers ([Lavi et al., 2021](#); [Kaya and Bogers, 2023](#)).<sup>7</sup> The task is inherently two-sided, requiring consideration of the multifaceted preferences of both recruiters and job seekers (Figure 3). To address this problem, previous work has extracted and integrated fine-grained factors like skills ([Dave et al., 2018](#); [Li et al., 2020b](#); [Yao et al., 2022](#); *inter-alia*), experience levels ([Li et al., 2020a](#)), and more ([Ha-Thuc et al., 2016](#); [Luo et al., 2019](#); [Gutiérrez et al., 2019](#); [Lai et al., 2024](#)) into matching models. Leveraging the linguistic capability of LLMs is an emerging research area, with studies exploring how LLMs can refine documents to alleviate the challenge of linguistic complexity ([Zheng et al., 2023a](#); [Du et al., 2024](#)) and integrating structured knowledge to improve accuracy and interpretability ([Wu et al., 2024](#)).

**Course recommendation** aims to help people bridge skill gaps by matching them with relevant courses from various data sources. Existing methods identify underlying factors using Transformer encoders ([Hao et al., 2021](#)), Bayesian variational networks ([Wang et al., 2021](#)), and generative adversarial networks ([Zheng et al., 2023b](#)). Recently, LLM-based systems have emerged with modular components for upstream tasks like skill extraction, entity linking, and matching ([Frej et al., 2024](#)).

Retrieval and recommendation tasks in the HR domain are highly knowledge-intensive and often involve challenges associated with the heterogeneity of data sources such as documents and behavioral data. Although existing approaches have developed sophisticated methods to tackle these challenges, there remains substantial potential for in-

<sup>7</sup>For more comprehensive review of this field, refer to specialized survey papers ([Balog et al., 2012](#); [de Ruijt and Bhulai, 2021](#); [Freire and de Castro, 2021](#); [Mashayekhi et al., 2024](#)).



tegrating pre-trained language models to improve language comprehension (Zhu et al., 2024).

## 4.2 Language Generation

**Generating job postings and resumes** is an impactful real-world application<sup>8</sup> that requires a nuanced understanding of job-specific skills across diverse work environments. Creating accurate job requirements, in particular, heavily relies on domain knowledge. Liu et al. (2020b) represented the relationships between skills, company size, and job titles using graphs, employing graph neural networks to generate job requirements. Similarly, (Qin et al., 2023) used a topic model to incorporate skill information into a job requirements generator. Other work has addressed job posting generation as a data-to-text task using a rule-based system (Somers et al., 1997) and a fine-tuned language model (Lorincz et al., 2022), with a focus on the fluency and adequacy of the generated texts.

**Generating interview questions** is also a knowledge-intensive task in the HR domain that can streamline the time-consuming candidate screening process. Automated systems have shown promise in generating questions based on the key requirements of a job position (Shi et al., 2020). Beyond this, NLP technologies can assist in crafting personalized questions by leveraging contextual information (Inoue et al., 2020; Rao S B et al., 2020), structured knowledge (Su et al., 2019), or web search (Qin et al., 2019, 2023).

Language generation tasks in the HR domain present several characteristic challenges. For instance, these tasks often involve generating output based on lengthy inputs with mixed topics (e.g., job postings). Existing work has typically focused on simplified problem settings (e.g., inputs that have already been parsed into skill tags). Accurately and efficiently processing such complex inputs remains an open problem.

## 4.3 Dialogue Systems

**Job interviews** present significant NLP research opportunities. Researchers have developed automated interviewing systems for communication skills training, which provide feedback through visualizations of user behavior (Hoque et al., 2013; Rao S. B et al., 2017) and adapt their interactions

<sup>8</sup>For instance, on Indeed’s platform, more than 750,000 employers have used an automated job posting generation system for approximately 2 million jobs as of July 2024 (Batty, 2024).

based on emotional states (Anderson et al., 2013; Hartholt et al., 2019; Kawahara, 2019). Additionally, techniques for post-interview assessment have been proposed, combining various visual and audio features linearly (Nguyen et al., 2014; Rao S. B et al., 2017; Naim et al., 2018) or with advanced neural networks (Hemamou et al., 2019). These studies have advanced the state of the art in processing multi-modal information, such as facial expressions, gestures, and speech. The rapid development of multi-modal LLMs could lead to new advancements in job interview systems. However, simply applying LLMs without domain-specific tuning can be ineffective, as a deep understanding of specialized knowledge is crucial for conducting meaningful conversations (Li et al., 2023).

Interactive systems can also be used for **managing HR-related inquiry**. A case study by Malik et al. (2022) showed positive effects of chatbots on employee experiences in HR activities. Collecting interactive data in specialized domains is challenging, but Xu et al. (2024) demonstrated the effectiveness of LLMs to simulate interactions for post-hiring HR transactions.

## 4.4 Ethics, Bias, and Fairness

Fairness concerns in algorithmic hiring have been widely studied in various research fields (Hunkenschroer and Luetge, 2022; Kumar et al., 2023; Fabris et al., 2024), with bias mitigation techniques focusing on reducing disparities in algorithmic outcomes across sensitive groups. These techniques span multiple stages of system development and evaluation (Quiñonero-Candela et al., 2023), including biased keyword removal from input text (De-Arteaga et al., 2019), balanced data sampling, internal representation adjustments (Hauzenberger et al., 2023; Masoudian et al., 2024), and post-processing methods (Geyik et al., 2019).

The association between occupations and sensitive attributes has also been a significant focus in text representation and generation. Studies have shown that word embeddings link gender pronouns with specific job titles, such as “she” with “nurse” and “he” with “physician” (Sun et al., 2019). Similar gender biases are found in system-generated texts (Sheng et al., 2019; Borchers et al., 2022). For example, Wan et al. (2023) found that person names, which can serve as proxies for sensitive attributes, influence LLM-generated reference letters. An et al. (2024) and Nghiem et al. (2024) also report name-related biases in LLM-based hiring

decisions, highlighting the need for careful consideration in these applications.

Blodgett et al. (2020) conducted a literature review and argued the importance of carefully conceptualizing “bias” and grounding it in theories established outside of NLP. In the HR domain, fairness and bias have been extensively studied for decades (Bertrand and Mullainathan, 2004). This rich theoretical and empirical foundation could offer valuable insights to NLP research. A notable example is the bias evaluation framework by Wang et al. (2024). This framework is informed by insights from labor economics, legal principles, and existing benchmarks, enabling a comprehensive and theoretically grounded evaluation of hiring decisions generated by LLMs.

#### 4.5 Summary

Downstream HR tasks are highly knowledge-intensive and also necessitate ethical and safety considerations. Researchers have addressed these with advanced modeling techniques that leverage detailed information such as extracted skills. Looking ahead, the contextual understanding, and reasoning capabilities of modern LLMs present an opportunity to develop holistic approaches that integrate specialized modules to address overarching goals in downstream HR tasks.

### 5 Underrepresented Tasks

Finally, we discuss HR activities that have been underrepresented in NLP research. Some of these tasks have received attention in broader research communities, but significant opportunities remain to leverage language resources for advancing computational methods.

#### 5.1 Data Analytics

Analyzing the labor market (Rahhal et al., 2024) can greatly benefit from data/text mining techniques. The insights gained can be valuable for policymakers, educators, and businesses.

**Job title benchmarking** involves matching job titles with equivalent expertise levels across different companies. Similarly, **job mobility analysis** focuses on identifying transferability between jobs while accounting for their specialties and work environments. These tasks are similar to the task of job title normalization (§3.3) but require a deeper analysis of individual roles and organizations. For example, a company’s industry and size often in-

fluence an employee’s next career move. Therefore, previous work has developed methods to integrate diverse information linked to career trajectories with LSTMs (Li et al., 2017), multi-view learning (Zhang et al., 2019) and graph neural networks (Zhang et al., 2021; Zha et al., 2024).

**The assessment of skill demand and value** is important not only for hiring but also for economic research (Zhu et al., 2018b; Cao et al., 2021) and education (Hao et al., 2021; Patacsil and Acosta, 2021). While this area has not yet gained much attention within the NLP community, a variety of techniques have been explored in the broader research field. For instance, Sun et al. (2021) introduced a neural model to break down job positions into required skills and assess their market value through salary prediction. Chao et al. (2024) proposed a graph encoder over a skill co-occurrence graph to capture demand-supply patterns in skill evolution. More recently, Chen et al. (2024) developed a large-scale dataset for forecasting job-skill demand, which opens avenues for future research. Although these studies effectively utilize structured data, skills are often described by simple phrases that may not fully convey their true functions. For example, “communication skills” can differ significantly based on the context (e.g., schools vs. consulting firms). Future research could focus on extracting rich contextual information from textual data such as job postings to enhance the depth of analysis.

#### 5.2 Sentiment Analysis and Opinion Mining

Sentiment about jobs and organizations can be collected through questionnaires or reviews from platforms like Glassdoor.<sup>9</sup> This information has the potential to help organizations create work environments, boost productivity, and improve business outcomes (Harter et al., 2002).

**Employee satisfaction (job satisfaction)** analysis focuses on evaluating work environments and identifying areas for improvement based on employee feedback. Moniz and de Jong (2014) applied topic modeling to online employee reviews to uncover key themes related to the organization’s future. Rink et al. (2024) approached this as an aspect-based sentiment analysis task, creating annotated datasets and fine-tuning transformer-based classifiers. While these studies highlight valuable use cases of sentiment analysis, addressing the di-

<sup>9</sup><https://www.glassdoor.com/>

versity of job categories remains an open challenge.

**Company profiling** focuses on identifying the key characteristics of a company. Early work relied mainly on numerical data, but recent studies have successfully incorporated textual data for deeper insights (Bajpai et al., 2018; Lin et al., 2020). For example, Lin et al. (2020) proposed a model-based topic approach that integrates review texts with numerical data to perform both qualitative opinion analysis and quantitative salary benchmarking.

### 5.3 Summary

This section highlighted several HR activities that offer significant opportunities to explore NLP techniques with heterogeneous data. In a similar vein, other core HR tasks, such as employee performance evaluation (Ye et al., 2019; Yu et al., 2023) and turnover analysis (Teng et al., 2019; Gamba et al., 2024), also provide interesting challenges. Future efforts should focus on constructing publicly accessible datasets to drive advancements in this area. Applying LLMs to synthesize data or de-identify Personally Identifiable Information (PII) in real-world datasets could offer a promising solution to the problem of data scarcity. However, they should be used with caution, as issues such as amplifying biases (§4.2) and exposing sensitive information from training data (Carlini et al., 2021) remain.

## 6 Conclusion and Future Directions

In this paper, we have categorized critical research challenges within the HR domain and identified significant opportunities for future exploration. To inspire future research in this domain and the broader NLP community, we provide a list of papers and public data resources on GitHub,<sup>10</sup> which we plan to update regularly.

**Toward Broader Goals:** The HR domain encompasses a variety of specialized problems where NLP techniques have been successfully applied (e.g., skill extraction). These problems are often tied to broader goals, such as matching talent with appropriate job opportunities and optimizing employee productivity. For example, accurate skill extraction can significantly improve job recommender systems. To accurately extract this skill information, it is useful to perform semantic analysis of documents to identify relevant sections and understand job titles. Intermediate tasks like these

can improve system performance in downstream applications and provide detailed information that can improve the fairness and transparency of final outcomes. The orchestration of specialized NLP tools to perform complex tasks is increasingly gaining the interest of the research community (e.g., Schick et al., 2023). The HR domain would benefit from exploring holistic approaches, which could also provide research opportunities to push the boundaries of language technologies.

**Knowledge Transfer:** Some successful research in the HR domain has introduced techniques and knowledge transferable to problems in other applications or domains. This trend is particularly evident in studies on job recommendation and bias mitigation, where the HR domain has established a strong position within the research community. We can also see similar knowledge transfer in some other specialized domains. For instance, the e-commerce domain has been one of the key drivers of multiple core NLP areas such as information extraction, sentiment analysis, and summarization. Promoting knowledge transfer to other domains will be key to conducting impactful NLP research in HR in the future.

**Data Challenge:** The availability of real or realistic datasets is a critical factor for advancing NLP research in the HR domain. Many types of HR documents involve privacy concerns that make them unsuitable for public release. However, approaches such as shared tasks with restricted data licenses, data donation,<sup>11</sup> anonymization, and data synthesis could provide valuable resources to the research community. Moreover, working with real-world datasets would also help researchers identify system constraints and requirements in practical scenarios such as latency requirements, increasing the social impact of research artifacts.

**The Application of LLMs:** The application of LLMs has gained popularity in the HR domain. While the collection and annotation of HR documents pose significant challenges, some studies have demonstrated the potential of LLMs to alleviate these issues. Furthermore, LLMs may introduce a new paradigm for many problems, offering substantial opportunities for researchers to generate innovative ideas that benefit both the HR domain and the broader research community.

<sup>10</sup><https://github.com/megagonlabs/nlp4hr-survey>

<sup>11</sup>FINDHR collected more than 1,100 CVs through donations (<https://findhr.eu/datadonation/>).



## Limitations

Due to space constraints, this paper aims to provide a focused literature review to offer readers a concise yet effective overview of HR applications. For those interested in a broader collection of NLP research in HR, we provide a list of papers and language resources on GitHub,<sup>10</sup> which we plan to update regularly. While there are numerous other NLP challenges in HR, such as linguistic and societal analysis (e.g., demographic, language, and cultural differences), we did not extensively cover these topics due to space limitations. As a result, the majority of papers discussed focus on widely spoken languages like English and Chinese. Lastly, while many companies are adopting modern NLP solutions in HR tasks, we have only reviewed techniques published in academic conferences.

## Acknowledgments

We thank the anonymous reviewers, our colleagues at Megagon Labs, and the participants of the NLP4HR 2024 workshop for their valuable feedback and discussions.

## References

- Haozhe An, Christabel Acquaye, Colin Wang, Zongxia Li, and Rachel Rudinger. 2024. [Do Large Language Models Discriminate in Hiring Decisions on the Basis of Race, Ethnicity, and Gender?](#) In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics*, pages 386–397, Bangkok, Thailand. Association for Computational Linguistics.
- Keith Anderson, Elisabeth André, T. Baur, Sara Bernardini, M. Chollet, E. Chryssafidou, I. Damian, C. Ennis, A. Egges, P. Gebhard, H. Jones, M. Ochs, C. Pelachaud, Kaška Porayska-Pomsta, P. Rizzo, and Nicolas Sabouret. 2013. [The TARDIS Framework: Intelligent Virtual Agents for Social Coaching in Job Interviews](#). In *Advances in Computer Entertainment*, pages 476–491, Cham. Springer International Publishing.
- Rajiv Bajpai, Devamanyu Hazarika, Kunal Singh, Sruthi Gorantla, Erik Cambria, and Roger Zimmerman. 2018. [Aspect-Sentiment Embeddings for Company Profiling and Employee Opinion Mining](#). In *Proceedings of the 19th International Conference on Computational Linguistics and Intelligent Text Processing*, Hanoi, Vietnam. Springer.
- Krisztian Balog, Yi Fang, Maarten de Rijke, Pavel Serdyukov, and Luo Si. 2012. [Expertise Retrieval](#). Now Foundations and Trends.
- Jill Barth. 2024. [Unveiling the Winners: 2024 Top HR Tech Products of the Year](#). Accessed: 21 February, 2025.
- Mathieu Bastian, Matthew Hayes, William Vaughan, Sam Shah, Peter Skomoroch, Hyungjin Kim, Sal Uryasev, and Christopher Lloyd. 2014. [LinkedIn Skills: Large-Scale Topic Extraction and Inference](#). In *Proceedings of the 8th ACM Conference on Recommender systems*, pages 1–8, New York, NY, USA. Association for Computing Machinery.
- Ryan Batty. 2024. [Attract Great Candidates with Indeed’s New AI-Powered Job Description Tool](#). Accessed: 21 February, 2025.
- Marianne Bertrand and Sendhil Mullainathan. 2004. [Are Emily and Greg More Employable Than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination](#). *American Economic Review*, 94(4):991–1013.
- Akshay Bhola, Kishaloy Halder, Animesh Prasad, and Min-Yen Kan. 2020. [Retrieving Skills from Job Descriptions: A Language Model Based Extreme Multi-label Classification Framework](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 5832–5842, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Su Lin Blodgett, Solon Barocas, Hal Daumé III, and Hanna Wallach. 2020. [Language \(Technology\) is Power: A Critical Survey of “Bias” in NLP](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5454–5476, Online. Association for Computational Linguistics.
- Conrad Borchers, Dalia Gala, Benjamin Gilbert, Eduard Oravkin, Wilfried Bounsi, Yuki M Asano, and Hannah Kirk. 2022. [Looking for a Handsome Carpenter! Debiasing GPT-3 Job Advertisements](#). In *Proceedings of the 4th Workshop on Gender Bias in Natural Language Processing*, pages 212–224, Seattle, Washington. Association for Computational Linguistics.
- Pawan Budhwar, Ashish Malik, M. T. Thedushika De Silva, and Praveena Thevisuthan. 2022. [Artificial Intelligence – Challenges and Opportunities for International HRM: A Review and Research Agenda](#). *The International Journal of Human Resource Management*, 33(6):1065–1097. Routledge.
- Federica Calanca, Luiza Sayfullina, Lara Minkus, Claudia Wagner, and Eric Malmi. 2019. [Responsible Team Players Wanted: An Analysis of Soft Skill Requirements in Job Advertisements](#). *EPJ Data Science*, 8(1):1–20.
- Lina Cao, Jian Zhang, Xinquan Ge, and Jindong Chen. 2021. [Occupational profiling driven by online job advertisements: Taking the data analysis and processing engineering technicians as an example](#). *PLOS ONE*, 16(6):e0253308.



- Nicholas Carlini, Florian Tramèr, Eric Wallace, Matthew Jagielski, Ariel Herbert-Voss, Katherine Lee, Adam Roberts, Tom Brown, Dawn Song, Úlfar Erlingsson, Alina Oprea, and Colin Raffel. 2021. [Extracting Training Data from Large Language Models](#). In *Proceedings of the 30th USENIX Security Symposium*, pages 2633–2650, Online. USENIX Association.
- Wenshuo Chao, Zhaopeng Qiu, Likang Wu, Zhuoning Guo, Zhi Zheng, Hengshu Zhu, and Hao Liu. 2024. [A Cross-View Hierarchical Graph Learning Hypernetwork for Skill Demand-Supply Joint Prediction](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 19813–19822. AAAI Press.
- Xi Chen, Chuan Qin, Chuyu Fang, Chao Wang, Chen Zhu, Fuzhen Zhuang, Hengshu Zhu, and Hui Xiong. 2024. [Job-SDF: A Multi-Granularity Dataset for Job Skill Demand Forecasting and Benchmarking](#). *Advances in Neural Information Processing Systems*, 37:129329–129356.
- Benjamin Clavié and Guillaume Soulié. 2023. [Large Language Models as Batteries-Included Zero-Shot ESCO Skills Matchers](#). In *Proceedings of the 3rd Workshop on Recommender Systems for Human Resources*, Singapore.
- Vachik S. Dave, Baichuan Zhang, Mohammad Al Hasan, Khalifeh AlJadda, and Mohammed Korayem. 2018. [A Combined Representation Learning Approach for Better Job and Skill Recommendation](#). In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*, pages 1997–2005, New York, NY, USA. Association for Computing Machinery.
- Maria De-Arteaga, Alexey Romanov, Hanna Wallach, Jennifer Chayes, Christian Borgs, Alexandra Chouldechova, Sahin Geyik, Krishnamurthy Kenthapadi, and Adam Tauman Kalai. 2019. [Bias in Bios: A Case Study of Semantic Representation Bias in a High-Stakes Setting](#). In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, pages 120–128, New York, NY, USA. Association for Computing Machinery.
- Corné de Ruijt and Sandjai Bhulai. 2021. [Job Recommender Systems: A Review](#). *arXiv*.
- An De Sitter and Walter Daelemans. 2003. Information Extraction via Double Classification. In *Proceedings of the International Workshop on Adaptive Text Extraction and Mining*, pages 66–73, Cavtat, Dubrovnik, Croatia.
- Jens-Joris Decorte, Jeroen Van Haute, Johannes Deleu, Chris Develder, and Thomas Demeester. 2022. [Design of Negative Sampling Strategies for Distantly Supervised Skill Extraction](#). In *Proceedings of the 2nd Workshop on Recommender Systems for Human Resources*, volume abs/2209.05987, Seattle, WA, USA.
- Jens-Joris Decorte, Jeroen Van Haute, Thomas Demeester, and Chris Develder. 2021. [JobBERT: Understanding Job Titles through Skills](#). *FEAST: International Workshop on Fair, Effective, and Sustainable Talent Management using Data Science*.
- Jens-Joris Decorte, Severine Verlinden, Jeroen Van Haute, Johannes Deleu, Chris Develder, and Thomas Demeester. 2023. [Extreme Multi-Label Skill Extraction Training using Large Language Models](#). In *Proceedings of the 4th International workshop on AI for Human Resources and Public Employment Services*, Turin, Italy.
- Yingpeng Du, Di Luo, Rui Yan, Xiaopei Wang, Hongzhi Liu, Hengshu Zhu, Yang Song, and Jie Zhang. 2024. [Enhancing Job Recommendation through LLM-Based Generative Adversarial Networks](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 8363–8371, Vancouver, BC, Canada.
- Ahmed Elsafty, Martin Riedl, and Chris Biemann. 2018. [Document-based Recommender System for Job Postings using Dense Representations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 216–224, New Orleans, LA, USA. Association for Computational Linguistics.
- Alessandro Fabris, Nina Baranowska, Matthew J. Dennis, David Graus, Philipp Hacker, Jorge Saldivar, Frederik Zuiderveen Borgesius, and Asia J. Biega. 2024. [Fairness and Bias in Algorithmic Hiring: a Multidisciplinary Survey](#). *arXiv*.
- Chuyu Fang, Chuan Qin, Qi Zhang, Kaichun Yao, Jingshuai Zhang, Hengshu Zhu, Fuzhen Zhuang, and Hui Xiong. 2023. [RecruitPro: A Pretrained Language Model with Skill-Aware Prompt Learning for Intelligent Recruitment](#). In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 3991–4002, New York, NY, USA. Association for Computing Machinery.
- Aidan Finn and Nicholas Kushmerick. 2004. [Multi-level Boundary Classification for Information Extraction](#). In *Proceedings of the 15th European Conference on Machine Learning*, pages 111–122, Pisa, Italy. Springer Berlin Heidelberg.
- Mauricio Noris Freire and Leandro Nunes de Castro. 2021. [e-Recruitment recommender systems: a systematic review](#). *Knowledge and Information Systems*, 63(1):1–20.
- Jibril Frej, Anna Dai, Syrielle Montariol, Antoine Bosselet, and Tanja Käser. 2024. [Course Recommender Systems Need to Consider the Job Market](#). In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 522–532, New York, NY, USA. Association for Computing Machinery.

- David Gamba, Yulin Yu, Yuan Yuan, Grant Schoenebeck, and Daniel M. Romero. 2024. [Exit Ripple Effects: Understanding the Disruption of Socialization Networks Following Employee Departures](#). In *Proceedings of the ACM Web Conference 2024*, pages 211–222, New York, NY, USA. Association for Computing Machinery.
- Sahin Cem Geyik, Stuart Ambler, and Krishnaram Kenthapadi. 2019. [Fairness-Aware Ranking in Search & Recommendation Systems with Application to LinkedIn Talent Search](#). In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2221–2231, New York, NY, USA. Association for Computing Machinery.
- Kate Glazko, Yusuf Mohammed, Ben Kosa, Venkatesh Potluri, and Jennifer Mankoff. 2024. [Identifying and Improving Disability Bias in GPT-Based Resume Screening](#). In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency*, pages 687–700, New York, NY, USA. Association for Computing Machinery.
- Nidhi Goyal, Jushaan Kalra, Charu Sharma, Raghava Mutharaju, Niharika Sachdeva, and Ponnurangam Kumaraguru. 2023. [JobXMLC: EXTreme Multi-Label Classification of Job Skills with Graph Neural Networks](#). In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 2181–2191, Dubrovnik, Croatia. Association for Computational Linguistics.
- Thomas Green, Diana Maynard, and Chenghua Lin. 2022. [Development of a Benchmark Corpus to Support Entity Recognition in Job Descriptions](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 1201–1208, Marseille, France. European Language Resources Association.
- Akshay Gugnani and Hemant Misra. 2020. [Implicit Skills Extraction Using Document Embedding and Its Use in Job Recommendation](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(08):13286–13293. Number: 08.
- Francisco Gutiérrez, Sven Charleer, Robin De Croon, Nyi Nyi Htun, Gerd Goetschalckx, and Katrien Verbert. 2019. [Explaining and Exploring Job Recommendations: A User-driven Approach for Interacting with Knowledge-based Job Recommender Systems](#). In *Proceedings of the 13th ACM Conference on Recommender Systems*, pages 60–68, New York, NY, USA. Association for Computing Machinery.
- Phong Ha, Shanshan Zhang, Nemanja Djuric, and Slobodan Vucetic. 2020. [Improving Word Embeddings through Iterative Refinement of Word- and Character-level Models](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 1204–1213, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Viet Ha-Thuc, Ye Xu, Satya Pradeep Kanduri, Xianren Wu, Vijay Dialani, Yan Yan, Abhishek Gupta, and Shakti Sinha. 2016. [Search by Ideal Candidates: Next Generation of Talent Search at LinkedIn](#). In *Proceedings of the 25th International Conference Companion on World Wide Web*, pages 195–198, Republic and Canton of Geneva, CHE. International World Wide Web Conferences Steering Committee.
- Bowen Hao, Jing Zhang, Cuiping Li, Hong Chen, and Hongzhi Yin. 2021. [Recommending Courses in MOOCs for Jobs: An Auto Weak Supervision Approach](#). *Machine Learning and Knowledge Discovery in Databases: Applied Data Science Track (ECML PKDD 2020)*, pages 36–51.
- James K. Harter, Frank L. Schmidt, and Theodore L. Hayes. 2002. [Business-Unit-Level Relationship Between Employee Satisfaction, Employee Engagement, and Business Outcomes: A Meta-Analysis](#). *The Journal of Applied Psychology*, 87(2):268–279.
- Arno Hartholt, Sharon Mozgai, and Albert "Skip" Rizzo. 2019. [Virtual Job Interviewing Practice for High-Anxiety Populations](#). In *Proceedings of the 19th ACM International Conference on Intelligent Virtual Agents*, pages 238–240, New York, NY, USA. Association for Computing Machinery.
- Lukas Hauzenberger, Shahed Masoudian, Deepak Kumar, Markus Schedl, and Navid Rekasaz. 2023. [Modular and On-demand Bias Mitigation with Attribute-Removal Subnetworks](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 6192–6214, Toronto, Canada. Association for Computational Linguistics.
- Léo Hemamou, Ghazi Felhi, Vincent Vandembussche, Jean-Claude Martin, and Chloé Clavel. 2019. [HireNet: A Hierarchical Attention Model for the Automatic Analysis of Asynchronous Video Job Interviews](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 573–581. Number: 01.
- Mohammed (Ehsan) Hoque, Matthieu Courgeon, Jean-Claude Martin, Bilge Mutlu, and Rosalind W. Picard. 2013. [MACH: My Automated Conversation Coach](#). In *Proceedings of the 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, pages 697–706, New York, NY, USA. Association for Computing Machinery.
- Anna Lena Hunkenschroer and Christoph Luetge. 2022. [Ethics of AI-Enabled Recruiting and Selection: A Review and Research Agenda](#). *Journal of Business Ethics*, 178(4):977–1007.
- Wissem Inoubli and Armelle Brun. 2022. [DGL4C: a Deep Semi-supervised Graph Representation Learning Model for Resume Classification](#). In *Proceedings of the 2nd Workshop on Recommender Systems for Human Resources*, Seattle, WA, USA.
- Koji Inoue, Kohei Hara, Divesh Lala, Kenta Yamamoto, Shizuka Nakamura, Katsuya Takanashi, and Tatsuya

- Kawahara. 2020. [Job Interviewer Android with Elaborate Follow-up Question Generation](#). In *Proceedings of the 2020 International Conference on Multimodal Interaction*, pages 324–332, New York, NY, USA. Association for Computing Machinery.
- International Labour Organization. 2012. *International Standard Classification of Occupations 2008 (ISCO-08): Structure, group definitions and correspondence tables*. International Labour Organization, Geneva, Switzerland.
- Tatsuya Kawahara. 2019. Spoken Dialogue System for a Human-like Conversational Robot ERICA. In *9th International Workshop on Spoken Dialogue System Technology*, pages 65–75, Singapore. Springer Singapore.
- Mesut Kaya and Toine Bogers. 2023. [An Exploration of Sentence-Pair Classification for Algorithmic Recruiting](#). In *Proceedings of the 17th ACM Conference on Recommender Systems*, pages 1175–1179, New York, NY, USA. Association for Computing Machinery.
- Mohammad Rashed Khan. 2024. [Application of Artificial Intelligence for Talent Management: Challenges and Opportunities](#). In *the 7th International Conference on Intelligent Human Systems Integration: Integrating People and Intelligent Systems*, volume 119, pages 234–329, Palermo, Italy. AHFE Open Acces. ISSN: 27710718 Issue: 119.
- Imane Khaouja, Ismail Kassou, and Mounir Ghogho. 2021. [A Survey on Skill Identification From Online Job Ads](#). *IEEE Access*, 9:118134–118153.
- Ilkka Kivimäki, Alexander Panchenko, Adrien Dessy, Dries Verdegem, Pascal Francq, Hugues Bersini, and Marco Saerens. 2013. [A Graph-Based Approach to Skill Extraction from Text](#). In *Proceedings of TextGraphs-8 Graph-based Methods for Natural Language Processing*, pages 79–87, Seattle, Washington, USA. Association for Computational Linguistics.
- Deepak Kumar, Tessa Grosz, Navid Rekabsaz, Elisabeth Greif, and Markus Schedl. 2023. [Fairness of Recommender Systems in the Recruitment Domain: An Analysis from Technical and Legal Perspectives](#). *Frontiers in Big Data*, 6.
- Kai-Huang Lai, Zhe-Rui Yang, Pei-Yuan Lai, Chang-Dong Wang, Mohsen Guizani, and Min Chen. 2024. [Knowledge-Aware Explainable Reciprocal Recommendation](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 8636–8644, Vancouver, BC, Canada. Number: 8.
- Thom Lake. 2022. [Flexible Job Classification with Zero-Shot Learning](#). In *Proceedings of the 2nd Workshop on Recommender Systems for Human Resources*, volume 3218, Seattle, WA, USA.
- Napat Laosaengpha, Thanit Tatavannarat, Chawan Piansaddhayanon, Attapol Rutherford, and Ekapol Chuangsuwanich. 2024. [Learning Job Title Representation from Job Description Aggregation Network](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 1319–1329, Bangkok, Thailand. Association for Computational Linguistics.
- Thorsten Lau and York Sure. 2002. [Introducing Ontology-based Skills Management at a Large Insurance Company](#). In *Modellierung 2002, Modellierung in der Praxis – Modellierung für die Praxis*, pages 123–134, Tutzing, Germany. Gesellschaft für Informatik e.V.
- Sven Laumer and Stefan Morana. 2022. [Chapter 12: HR Natural Language Processing - Conceptual Overview and State of the Art on Conversational Agents in Human Resources Management](#). In *Handbook of Research on Artificial Intelligence in Human Resource Management*. Edward Elgar Publishing, Cheltenham, UK.
- Dor Lavi, Volodymyr Medentsiy, and David Graus. 2021. [conSultantBERT: Fine-tuned Siamese Sentence-BERT for Matching Jobs and Job Seekers](#). In *Proceedings of the Workshop on Recommender Systems for Human Resources co-located with the 15th ACM Conference on Recommender Systems*, Amsterdam, The Netherlands. CEUR Workshop Proceedings.
- Martin le Vrang, Agis Papantoniou, Erika Pauwels, Pieter Fannes, Dominique Vandestein, and Johan De Smedt. 2014. [ESCO: Boosting Job Matching in Europe with Semantic Interoperability](#). *Computer*, 47(10):57–64. Conference Name: Computer.
- Jochen L. Leidner and Mark Stevenson. 2024. [Challenges and Opportunities of NLP for HR Applications: A Discussion Paper](#). *arXiv*.
- Changmao Li, Elaine Fisher, Rebecca Thomas, Steve Pittard, Vicki Hertzberg, and Jinho D. Choi. 2020a. [Competence-Level Prediction and Resume & Job Description Matching Using Context-Aware Transformer Models](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*, Online. Association for Computational Linguistics.
- Liangyue Li, How Jing, Hanghang Tong, Jaewon Yang, Qi He, and Bee-Chung Chen. 2017. [NEMO: Next Career Move Prediction with Contextual Embedding](#). In *Proceedings of the 26th International Conference on World Wide Web Companion*, pages 505–513, Republic and Canton of Geneva, CHE. International World Wide Web Conferences Steering Committee.
- Mingzhe Li, Xiuying Chen, Weiheng Liao, Yang Song, Tao Zhang, Dongyan Zhao, and Rui Yan. 2023. [EZ-Interviewer: To Improve Job Interview Performance with Mock Interview Generator](#). In *Proceedings of the Sixteenth ACM International Conference on Web Search and Data Mining*, page 1102–1110, New York, NY, USA. Association for Computing Machinery.



- Shan Li, Baoxu Shi, Jaewon Yang, Ji Yan, Shuai Wang, Fei Chen, and Qi He. 2020b. [Deep Job Understanding at LinkedIn](#). In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2145–2148, New York, NY, USA. Association for Computing Machinery.
- Hao Lin, Hengshu Zhu, Junjie Wu, Yuan Zuo, Chen Zhu, and Hui Xiong. 2020. [Enhancing Employer Brand Evaluation with Collaborative Topic Regression Models](#). *ACM Transactions on Information Systems*, 38(4):32:1–32:33.
- Junhua Liu, Yung Chuen Ng, Kristin L. Wood, and Kwan Hui Lim. 2020a. [IPOD: A Large-scale Industrial and Professional Occupation Dataset](#). In *Companion Publication of the 2020 Conference on Computer Supported Cooperative Work and Social Computing*, pages 323–328, New York, NY, USA. Association for Computing Machinery.
- Liting Liu, Jie Liu, Wenzheng Zhang, Ziming Chi, Wenxuan Shi, and Yalou Huang. 2020b. [Hiring Now: A Skill-Aware Multi-Attention Model for Job Posting Generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3096–3104, Online. Association for Computational Linguistics.
- Anna Lorincz, David Graus, Dor Lavi, and Joao Lebre Magalhaes Pereira. 2022. [Transfer Learning for Multilingual Vacancy Text Generation](#). In *Proceedings of the 2nd Workshop on Natural Language Generation, Evaluation, and Metrics (GEM)*, pages 207–222, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Yong Luo, Huaizheng Zhang, Yonggang Wen, and Xinwen Zhang. 2019. [ResumeGAN: An Optimized Deep Representation Learning Framework for Talent-Job Fit via Adversarial Learning](#). In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, pages 1101–1110, New York, NY, USA. Association for Computing Machinery.
- Antoine Magron, Anna Dai, Mike Zhang, Syrielle Montariol, and Antoine Bosselut. 2024. [JobSkape: A Framework for Generating Synthetic Job Postings to Enhance Skill Matching](#). In *Proceedings of the First Workshop on Natural Language Processing for Human Resources*, pages 43–58, St. Julian’s, Malta. Association for Computational Linguistics.
- Ashish Malik, Pawan Budhwar, Charmi Patel, and N. R. Srikanth. 2022. [May the Bots Be with You! Delivering HR Cost-effectiveness and Individualised Employee Experiences in an MNE](#). *The International Journal of Human Resource Management*, 33(6):1148–1178. Publisher: Routledge. eprint: <https://doi.org/10.1080/09585192.2020.1859582>.
- Yoosof Mashayekhi, Nan Li, Bo Kang, Jeffrey Lijffijt, and Tijl De Bie. 2024. [A Challenge-based Survey of E-recruitment Recommendation Systems](#). *ACM Comput. Surv.*, 56(10):252:1–252:33.
- Shahed Masoudian, Cornelia Volaucnik, Markus Schedl, and Navid Rekabsaz. 2024. [Effective Controllable Bias Mitigation for Classification and Retrieval using Gate Adapters](#). In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics*, pages 2434–2453, St. Julian’s, Malta. Association for Computational Linguistics.
- Adrien Mogenet, Tuan Anh Nguyen Pham, Masahiro Kazama, and Jialin Kong. 2019. [Predicting On-line Performance of Job Recommender Systems with Offline Evaluation](#). In *Proceedings of the 13th ACM Conference on Recommender Systems*, pages 477–480, New York, NY, USA. Association for Computing Machinery.
- Andy Moniz and Franciska de Jong. 2014. [Sentiment Analysis and the Impact of Employee Satisfaction on Firm Earnings](#). In *Proceedings of the 36th European Conference on IR Research on Advances in Information Retrieval*, pages 519–527, Cham. Springer International Publishing.
- Iftekhar Naim, Md. Iftekhar Tanveer, Daniel Gildea, and Mohammed Ehsan Hoque. 2018. [Automated Analysis and Prediction of Job Interview Performance](#). *IEEE Transactions on Affective Computing*, 9(2):191–204. Conference Name: IEEE Transactions on Affective Computing.
- National Center for O\*NET Development. [O\\*NET OnLine](#). Accessed: 21 February, 2025.
- Huy Nghiem, John Prindle, Jieyu Zhao, and Hal Daumé Iii. 2024. [“You Gotta be a Doctor, Lin”: An Investigation of Name-Based Bias of Large Language Models in Employment Recommendations](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 7268–7287, Miami, Florida, USA. Association for Computational Linguistics.
- Laurent Son Nguyen, Denise Frauendorfer, Marianne Schmid Mast, and Daniel Gatica-Perez. 2014. [Hire Me: Computational Inference of Hirability in Employment Interviews Based on Nonverbal Behavior](#). *IEEE Transactions on Multimedia*, 16(4):1018–1031. Conference Name: IEEE Transactions on Multimedia.
- Frederick F. Patacsil and Michael Acosta. 2021. [Analyzing the Relationship between Information Technology Jobs Advertised On-line and Skills Requirements Using Association Rules](#). *Bulletin of Electrical Engineering and Informatics*, 10(5):2771–2779. Number: 5.
- Chuan Qin, Le Zhang, Yihang Cheng, Rui Zha, Dazhong Shen, Qi Zhang, Xi Chen, Ying Sun, Chen Zhu, Hengshu Zhu, and Hui Xiong. 2024. [A Comprehensive Survey of Artificial Intelligence Techniques for Talent Analytics](#). *arXiv*.



- Chuan Qin, Hengshu Zhu, Dazhong Shen, Ying Sun, Kaichun Yao, Peng Wang, and Hui Xiong. 2023. [Automatic Skill-Oriented Question Generation and Recommendation for Intelligent Job Interviews](#). *ACM Transactions on Information Systems*, 42(1):27:1–27:32.
- Chuan Qin, Hengshu Zhu, Chen Zhu, Tong Xu, Fuzhen Zhuang, Chao Ma, Jingshuai Zhang, and Hui Xiong. 2019. [DuerQuiz: A Personalized Question Recommender System for Intelligent Job Interview](#). In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, page 2165–2173, New York, NY, USA. Association for Computing Machinery.
- Joaquin Quiñero-Candela, Yuwen Wu, Brian Hsu, Sakshi Jain, Jennifer Ramos, Jon Adams, Robert Hallman, and Kinjal Basu. 2023. [Disentangling and Operationalizing AI Fairness at LinkedIn](#). In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, pages 1213–1228, New York, NY, USA. Association for Computing Machinery.
- Ibrahim Rahhal, Ismail Kassou, and Mounir Ghogho. 2024. [Data Science for Job Market Analysis: A Survey on Applications and Techniques](#). *Expert Systems with Applications*, 251:124101.
- Pooja Rao S B, Manish Agnihotri, and Dinesh Babu Jayagopi. 2020. [Automatic Follow-up Question Generation for Asynchronous Interviews](#). In *Proceedings of the Workshop on Intelligent Information Processing and Natural Language Generation*, pages 10–20, Santiago de Compostela, Spain. Association for Computational Linguistics.
- Pooja Rao S. B, Sowmya Rasipuram, Rahul Das, and Dinesh Babu Jayagopi. 2017. [Automatic Assessment of Communication Skill in Non-conventional Interview Settings: A Comparative Study](#). In *Proceedings of the 19th ACM International Conference on Multimodal Interaction*, pages 221–229, New York, NY, USA. Association for Computing Machinery.
- Lois Rink, Job Meijdam, and David Graus. 2024. [Aspect-Based Sentiment Analysis for Open-Ended HR Survey Responses](#). In *Proceedings of the First Workshop on Natural Language Processing for Human Resources*, pages 16–26, St. Julian’s, Malta. Association for Computational Linguistics.
- Luiza Sayfullina, Eric Malmi, and Juho Kannala. 2018. [Learning Representations for Soft Skill Matching](#). *Analysis of Images, Social Networks and Texts*, pages 141–152.
- Timo Schick, Jane Dwivedi-Yu, Roberto Dessi, Roberta Raileanu, Maria Lomeli, Eric Hambro, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. 2023. [Toolformer: Language Models Can Teach Themselves to Use Tools](#). *Advances in Neural Information Processing Systems*, 36:68539–68551.
- Elena Senger, Mike Zhang, Rob van der Goot, and Barbara Plank. 2024. [Deep Learning-based Computational Job Market Analysis: A Survey on Skill Extraction and Classification from Job Postings](#). In *Proceedings of the First Workshop on Natural Language Processing for Human Resources*, pages 1–15, St. Julian’s, Malta. Association for Computational Linguistics.
- Gaurav Sharma. 2021. [A Literature Review on Application of Artificial Intelligence in Human Resource Management and Its Practices in Current Organizational Scenario](#). In *2021 Fifth International Conference on I-SMAC (IoT in Social, Mobile, Analytics and Cloud)*, pages 594–600. ISSN: 2768-0673.
- Emily Sheng, Kai-Wei Chang, Premkumar Natarajan, and Nanyun Peng. 2019. [The Woman Worked as a Babysitter: On Biases in Language Generation](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing*, pages 3407–3412, Hong Kong, China. Association for Computational Linguistics.
- Baoxu Shi, Shan Li, Jaewon Yang, Mustafa Emre Kazdagli, and Qi He. 2020. [Learning to Ask Screening Questions for Job Postings](#). In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 549–558, New York, NY, USA. Association for Computing Machinery.
- Harold Somers, Bill Black, Joakim Nivre, Torbjorn Lager, Annarosa Multari, Luca Gilardoni, Jeremy Ellman, and Alex Rogers. 1997. [Multilingual Generation and Summarization of Job Adverts: the TREE Project](#). In *Proceedings of the Fifth Conference on Applied Natural Language Processing*, pages 269–276, Washington, DC, USA. Association for Computational Linguistics.
- Ming-Hsiang Su, Chung-Hsien Wu, and Yi Chang. 2019. [Follow-Up Question Generation Using Neural Tensor Network-Based Domain Ontology Population in an Interview Coaching System](#). In *Proceedings of the Interspeech 2019*, pages 4185–4189.
- Tony Sun, Andrew Gaut, Shirlyn Tang, Yuxin Huang, Mai ElSherief, Jieyu Zhao, Diba Mirza, Elizabeth Belding, Kai-Wei Chang, and William Yang Wang. 2019. [Mitigating Gender Bias in Natural Language Processing: Literature Review](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1630–1640, Florence, Italy. Association for Computational Linguistics.
- Ying Sun, Fuzhen Zhuang, Hengshu Zhu, Qi Zhang, Qing He, and Hui Xiong. 2021. [Market-oriented Job Skill Valuation with Cooperative Composition Neural Network](#). *Nature Communications*, 12(1):1992. Number: 1 Publisher: Nature Publishing Group.
- Mingfei Teng, Hengshu Zhu, Chuanren Liu, Chen Zhu, and Hui Xiong. 2019. [Exploiting the Contagious Ef-](#)

- fect for Employee Turnover Prediction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 1166–1173. Number: 01.
- Suyi Tu and Olivia Cannon. 2022. **Beyond human-in-the-loop: scaling occupation taxonomy at Indeed**. In *Proceedings of the 2nd Workshop on Recommender Systems for Human Resources*, Seattle, WA, USA.
- Yixin Wan, George Pu, Jiao Sun, Aparna Garimella, Kai-Wei Chang, and Nanyun Peng. 2023. **“Kelly is a Warm Person, Joseph is a Role Model”: Gender Biases in LLM-Generated Reference Letters**. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 3730–3748, Singapore. Association for Computational Linguistics.
- Chao Wang, Hengshu Zhu, Peng Wang, Chen Zhu, Xi Zhang, Enhong Chen, and Hui Xiong. 2021. **Personalized and Explainable Employee Training Course Recommendations: A Bayesian Variational Approach**. *ACM Transactions on Information Systems*, 40(4):70:1–70:32.
- Jin Wang, Yuliang Li, Wataru Hirota, and Eser Kandogan. 2022. **Machop: an End-to-end Generalized Entity Matching Framework**. In *Proceedings of the Fifth International Workshop on Exploiting Artificial Intelligence Techniques for Data Management*, pages 1–10, New York, NY, USA. Association for Computing Machinery.
- Ze Wang, Zekun Wu, Xin Guan, Michael Thaler, Adriano Koshiyama, Skylar Lu, Sachin Beepath, Ediz Ertekin, and Maria Perez-Ortiz. 2024. **JobFair: A Framework for Benchmarking Gender Hiring Bias in Large Language Models**. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 3227–3246, Miami, Florida, USA. Association for Computational Linguistics.
- Likang Wu, Zhaopeng Qiu, Zhi Zheng, Hengshu Zhu, and Enhong Chen. 2024. **Exploring Large Language Model for Graph Data Understanding in Online Job Recommendations**. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 9178–9186, Vancouver, BC, Canada. Number: 8.
- Weijie Xu, Zicheng Huang, Wenxiang Hu, Xi Fang, Rajesh Cherukuri, Naumaan Nayyar, Lorenzo Malandri, and Srinivasan Sengamedu. 2024. **HR-MultiWOZ: A Task Oriented Dialogue (TOD) Dataset for HR LLM Agent**. In *Proceedings of the First Workshop on Natural Language Processing for Human Resources*, pages 59–72, St. Julian’s, Malta. Association for Computational Linguistics.
- Michiharu Yamashita, Jia Tracy Shen, Thanh Tran, Hamoon Ekhtiari, and Dongwon Lee. 2023. **JAMES: Normalizing Job Titles with Multi-Aspect Graph Embeddings and Reasoning**. In *2023 IEEE 10th International Conference on Data Science and Advanced Analytics (DSAA)*, pages 1–10.
- Kaichun Yao, Jingshuai Zhang, Chuan Qin, Xin Song, Peng Wang, Hengshu Zhu, and Hui Xiong. 2023. **Re-suFormer: Semantic Structure Understanding for Resumes via Multi-Modal Pre-training**. In *2023 IEEE 39th International Conference on Data Engineering (ICDE)*, pages 3154–3167. ISSN: 2375-026X.
- Kaichun Yao, Jingshuai Zhang, Chuan Qin, Peng Wang, Hengshu Zhu, and Hui Xiong. 2022. **Knowledge Enhanced Person-Job Fit for Talent Recruitment**. In *2022 IEEE 38th International Conference on Data Engineering*, pages 3467–3480. ISSN: 2375-026X.
- Yuyang Ye, Hengshu Zhu, Tong Xu, Fuzhen Zhuang, Runlong Yu, and Hui Xiong. 2019. **Identifying High Potential Talent: A Neural Network Based Dynamic Social Profiling Approach**. In *Proceedings of the 2019 IEEE International Conference on Data Mining*, pages 718–727. ISSN: 2374-8486.
- Kun Yu, Gang Guan, and Ming Zhou. 2005. **Resume Information Extraction with Cascaded Hybrid Model**. In *Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics*, pages 499–506, Ann Arbor, Michigan. Association for Computational Linguistics.
- Yulin Yu, Longqi Yang, Siân Lindley, and Mengting Wan. 2023. **Large-Scale Analysis of New Employee Network Dynamics**. In *Proceedings of the ACM Web Conference 2023*, pages 2719–2730, New York, NY, USA. Association for Computing Machinery.
- Rui Zha, Chuan Qin, Le Zhang, Dazhong Shen, Tong Xu, Hengshu Zhu, and Enhong Chen. 2024. **Career Mobility Analysis With Uncertainty-Aware Graph Autoencoders: A Job Title Transition Perspective**. *IEEE Transactions on Computational Social Systems*, 11(1):1205–1215. Conference Name: IEEE Transactions on Computational Social Systems.
- Denghui Zhang, Junming Liu, Hengshu Zhu, Yanchi Liu, Lichen Wang, Pengyang Wang, and Hui Xiong. 2019. **Job2Vec: Job Title Benchmarking with Collective Multi-View Representation Learning**. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, pages 2763–2771, New York, NY, USA. Association for Computing Machinery.
- Le Zhang, Ding Zhou, Hengshu Zhu, Tong Xu, Rui Zha, Enhong Chen, and Hui Xiong. 2021. **Attentive Heterogeneous Graph Embedding for Job Mobility Prediction**. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pages 2192–2201, New York, NY, USA. Association for Computing Machinery.
- Mike Zhang, Rob Goot, Min-Yen Kan, and Barbara Plank. 2024a. **NNOSE: Nearest Neighbor Occupational Skill Extraction**. In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 589–608, St. Julian’s, Malta. Association for Computational Linguistics.

Mike Zhang, Rob Goot, and Barbara Plank. 2024b. [Entity Linking in the Job Market Domain](#). In *Findings of the Association for Computational Linguistics: EACL 2024*, pages 410–419, St. Julian’s, Malta. Association for Computational Linguistics.

Mike Zhang, Kristian Jensen, Sif Sonniks, and Barbara Plank. 2022. [SkillSpan: Hard and Soft Skill Extraction from English Job Postings](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4962–4984, Seattle, WA, USA. Association for Computational Linguistics.

Mike Zhang, Rob van der Goot, and Barbara Plank. 2023. [ESCOXLM-R: Multilingual Taxonomy-driven Pre-training for the Job Market Domain](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics*, pages 11871–11890, Toronto, Canada. Association for Computational Linguistics.

Meng Zhao, Faizan Javed, Ferosh Jacob, and Matt McNair. 2015. [SKILL: A System for Skill Identification and Normalization](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 29, pages 4012–4017. Number: 2.

Zhi Zheng, Zhaopeng Qiu, Xiao Hu, Likang Wu, Hengshu Zhu, and Hui Xiong. 2023a. [Generative Job Recommendations with Large Language Model](#). *arXiv*.

Zhi Zheng, Ying Sun, Xin Song, Hengshu Zhu, and Hui Xiong. 2023b. [Generative Learning Plan Recommendation for Employees: A Performance-aware Reinforcement Learning Approach](#). In *Proceedings of the 17th ACM Conference on Recommender Systems*, pages 443–454, New York, NY, USA. Association for Computing Machinery.

Chen Zhu, Hengshu Zhu, Hui Xiong, Chao Ma, Fang Xie, Pengliang Ding, and Pan Li. 2018a. [Person-Job Fit: Adapting the Right Talent for the Right Job with Joint Representation Learning](#). *ACM Trans. Manage. Inf. Syst.*, 9(3):12:1–12:17.

Jun Zhu and Celine Hudelot. 2022. [Towards Job-Transition-Tag Graph for a Better Job Title Representation Learning](#). In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 2133–2140, Seattle, United States. Association for Computational Linguistics.

Tingting Juni Zhu, Alan Fritzier, and Jan Orlowski. 2018b. [World Bank Group-LinkedIn Data Insights : Jobs, Skills and Migration Trends Methodology and Validation Results](#). Technical report, World Bank Group, Washington D.C., USA.

Yaochen Zhu, Liang Wu, Binchi Zhang, Song Wang, Qi Guo, Liangjie Hong, Luke Simon, and Jundong Li. 2024. [Understanding and Modeling Job Marketplace with Pretrained Language Models](#). In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*, pages 5143–5150,

New York, NY, USA. Association for Computing Machinery.

## A Paper Collection

In this paper we aimed to offer a curated overview of key research challenges rather than a systematic and exhaustive literature review due to the page limit. Before curating papers, we employed the following approach to gather relevant papers. We begin by identifying recently published HR-related papers using keywords such as “job,” “occupation,” “hiring,” “recruit,” “resume,” “HR,” “company,” and “skill” from venues such as ACL conferences, KDD, CIKM, WWW, SIGIR, RecSys, AAAI, IJCAI, and relevant workshops. Additionally, we conduct keyword searches on Google Scholar and Semantic Scholar to collect non-computational papers. Subsequently, we employ snowball sampling from the citations of these papers to further gather relevant literature. We include peer-reviewed academic papers available as of December 2024 and exclude the others unless they are cited from multiple academic papers.