Report on the 1st Workshop on Recommender Systems for Human Resources (RecSys in HR 2021) at RecSys 2021

David Graus
Randstad Groep Nederland
david.graus@randstadgroep.nl

Toine Bogers
Aalborg University Copenhagen
toine@hum.aau.dk

Mesut Kaya
Aalborg University Copenhagen
mkaya@hum.aau.dk

Francisco Gutiérrez
KU Leuven
francisco.gutierrez@kuleuven.be

Katrien Verbert
KU Leuven
katrien.verbert@kuleuven.be

Abstract

Recommender systems are increasingly used in more high risk application domains, including in the domain of Human Resources (HR). These recommender systems help end-users find relevant vacancies out of an abundant overload of available vacancies, but also support other important objectives such as job mobility. Despite the use in industry applications, there are several research challenges associated to such objectives that have not yet been addressed in detail in this context, such as supporting end-users to steer the recommendation process with input and feedback and increasing diversity of recommendations. The goal of our workshop is to build a strong research capacity around recommender systems for HR to address these challenges. This paper describes the goal and activities of the Workshop on Recommender Systems for Human Resources (RecSys in HR 2021), organized at the 15th ACM Conference on Recommender Systems.

Date: 10–14 September, 2021.

Website: https://foo2021.net.

1 Introduction and Motivation for the Workshop

The field of Human Resources (HR) is at the forefront of adopting AI technologies. According to PWC over 40% of HR-functions of international companies use AI-applications [Charlier and Kloppenburg, 2017]. This so-called HR Technology (HR Tech) aims to replace or support Human

Resource functions such as talent acquisition and management, employee compensation, workforce analytics, and performance management.

Recommender Systems, broadly defined as systems that aim to support users in decision making by suggesting and offering relevant content, play an integral role in the rapid rise of HR Tech. Their applications range from assisting the talent acquisition process through matching [Kenthapadi et al., 2017], analyzing resumes or other user representations for candidate screening [Wilson et al., 2021] and automated assessment [Naim et al., 2018; Liem et al., 2018], to broader tasks such as recommendations for upskilling [Umemoto et al., 2020].

The use of AI applications in the recruitment process, such as recommender systems, is considered high-risk by the European Commission [Zuiderveen Borgesius, 2018], as automation here can directly impact the (working) lives of people. In this light, the rise of AI-assisted hiring and screening is met with caution, and is a widely-used example application area in AI ethics and fairness literature [Raghavan et al., 2020; Deshpande et al., 2020; Köchling and Wehner, 2020; Sánchez-Monedero et al., 2020; Mujtaba and Mahapatra, 2019]. At the same time, there is a rising commercial interest around these technologies from companies and startups alike Raghavan et al. [2020]. We feel the prevalence and rise of recommender system technology in HR calls for a central forum where researchers and practitioners alike can study and discuss the domain-specific aspects, challenges, and opportunities of RecSys and other HR Tech.

Past editions of the RecSys conference have seen a steady number of research contributions on automating and (more commonly) supporting job recommendation [Saini et al., 2019; Frid-Nielsen, 2019; Gutiérrez et al., 2019; Paparrizos et al., 2011; Bastian et al., 2014; Kleinerman et al., 2018; Lacic et al., 2019], all of which have focused on the core HR task of recruitment through the development of automatic job recommendation algorithms. In addition to this research, the RecSys Challenges of 2016 [rec, 2016] and 2017 [rec, 2017] both focused on the task of job recommendation, with Xing, a social network for business operating mostly in German-speaking countries, providing the training data. The task proved popular with 119 and 103 participating teams in 2016 and 2017 respectively, which shows that there is a large potential audience at RecSys for research on recommender systems in an HR setting.

However, despite a handful of industry-focused events focused on HR tech[1,2,3], to the best of our knowledge there have not been academic workshops on job recommendation or HR tech in general at RecSys or related conferences. By gathering experts, interested researchers and practitioners from academia and industry at a single event, we hoped to provide the research area of recommender systems in HR–and HR tech in general–a more concentrated push forward.

2 Workshop Goals & Format

Our aim was to make the **RecSys in HR** workshop an inclusive, interactive, and inspiring event. Given the fact this was the first workshop on the topic, we made an effort to attract and invite participants from academia, industry and government with an interest in HR tasks and the technology to support them, hoping as such to gather a diverse range of perspectives at the

[1] https://events.cipd.co.uk/events/people-analytics/
[2] https://hrfutureconference.com/programme
[3] https://www.mihranalyticsconference.com/

workshop, which should make for more stimulating and engaged discussion.

To ensure a wider range of perspectives, we hosted a panel on the current challenges in HR tech, for which we invited both academic researchers and industry practitioners with HR and technical backgrounds. Some of these panel members would be unlikely to normally attend RecSys, which allowed the workshop to provide a complementary experience to the main conference.

In order to ensure the interactivity of the workshop, we planned enough time for Q&A sessions after each presentation and the keynotes, in addition to organizing break-out sessions in which all attendees had the opportunity to actively engage and discuss topically relevant issues and challenges.

RecSys in HR was a full-day, hybrid workshop. The first half took place physically at the ACM RecSys 2021 Conference venue (and was also streamed via Zoom to remote workshop attendees), and the second half took place virtually via Zoom (and was also streamed from and followed by the in-person attendees at the venue). The workshop attracted around 80 participants, of which 20 attended physically and around 60 online. Three out of five of the workshop organizers were present at the workshop venue in Amsterdam.

3 Workshop Activities

Below we list the different workshop activities for the full-day workshop.

3.1 Physical morning slot

3.1.1 Opening keynote

Our first keynote was given by Quirine Eijkman, Deputy President at the National Human Rights Institute of the Netherlands, and co-author of their recently published report[4] "*Research into algorithms and discrimination in recruitment and selection*". Quirine started by defining discrimination, both from a Dutch and an EU perspective, after which she highlighted several possible areas of friction with recruitment algorithms, such the right to non-discrimination, the right to privacy, and the right to an effective remedy and to a fair trial. Quirine added that the recruitment process without algorithms is not perfect either, as there have been plenty of examples of discriminatory biases in human recruitment behavior. Legally, it does not matter whether a computer or a human discriminates and that intent is irrelevant, as it is the discriminatory impact that counts. Some of the particular human rights challenges for recruitment that Quirine identified included transparency and explainability—how can an applicant, employer or judge know that there was discriminatory treatment by an algorithm—and the general challenges with applying decades-old laws to modern technology. To conclude her keynote, Quirine advocated for designing algorithm for non-discrimination from the start through increased cooperation between lawyers and algorithm designers as well as frequent verification of recruitment algorithms to ensure they meet the relevant standards.

[4] https://mensenrechten.nl/en/netherlands-institute-human-rights

3.1.2 Paper presentation sessions

In total, 8 papers were submitted for peer-review to this workshop. Out of these, 8 papers were accepted and published in the CEUR proceedings: 7 as regular papers, 1 as short paper. Below we summarize the accepted contributions.

In our first paper session, Mashayekhi et al. [2021] presented their work on studying the mismatch between skill supply and demand in the job market as a network imbalance problem. The authors propose (a) a novel method for quantifying imbalance in a network between two sets of nodes based on network embeddings, and (b) a Graph Balancing algorithm (`GraB`). They evaluate their algorithm on different datasets, including a job market network, and find that it outperforms baselines in reducing network imbalance.

Next, Bogers and Kaya [2021] presented an exploratory user study of how recruiters search for candidates at Scandinavia's largest job portal. They applied the Contextual Inquiry (CI) methodology for gaining understanding and gathering information on how recruiters seek candidates. Their findings include the importance of separating primary from secondary qualifications, and the relative importance of location, salary, and experience in the relevance assessment of candidates.

In the second paper session, de Groot et al. [2021] presented their work on constructing a skills and occupation knowledge graph (KG), enriched with skills and occupations extracted from a large vacancy dataset. They present different downstream tasks that can be applied with the resulting KG, including career path-finding and identifying most relevant skills per occupation group.

The next two papers in this session addressed the use of textual embeddings for matching resumes or job seekers to job postings. First, Kaya and Bogers [2021] compare embedding job titles (from both resumes and job postings) to using textual embeddings based on the full-text resumes and job postings. Their findings include that job title embeddings usually outperform full text embeddings, and that customly trained embeddings outperform pre-trained embeddings.

The latter finding is in line with the work presented by Lavi et al. [2021]. In their paper, the authors compare pre-trained vectorization methods (TF-IDF weighted vectors and BERT embeddings) to vectorization fine-tuned on a custom binary classification dataset of over 270,000 hand-labeled matches between job postings and job seekers' resumes. Their presented model, `conSultantBERT`, significantly outperforms all other baselines on the same matching task.

3.1.3 Breakout Sessions

Our breakout sessions aimed to provide an opportunity to discuss relevant issues and challenges of RecSys in HR between all the workshop attendees. The break-out sessions' topics were seeded from a short survey we sent out prior to the workshop to conference attendees that expressed interest in joining our workshop. In this survey, we requested the most pressing challenges and interesting topics in HR Tech according to our prospective participants' perspectives. We split the attendees into three offline groups, one for each topic and each corresponding to an online Zoom breakout rooms so online participamts could also participate. The topics were:

1. Challenges and opportunities in fairness, accountability and transparency in job recommendations.

2. How to design more inclusive and fair HR technology?
3. How to design better recruiter-in-the-loop tools to assist recruiters more intelligently?

After the breakout sessions, a representative of each of the breakout groups summarized their discussions and insights.

In the first group, participants discussed topics ranging from model-agnostic explainability methods, identifying relevant stakeholders in an organization, and more generally accountability (e.g., identifying who is responsible for overseeing whether HR Technology tools are working as expected), challenges in disclosing what (type of) data is being used for what purpose, e.g., how training data is constructed, and effective methods of explaining and communicating complex information and data with uncertainty. One observation shared by an attendee is that the group organically seemed to focus on challenges of fair, accountable, and transparent recommendations, omitting opportunities — which as the attendee reflected may be a signal in itself: perhaps challenges are more urgent, important, and/or top-of-mind? Finally, another attendee stressed that explainability is not merely a data challenge: transparency may mean very different things to different stakeholders, e.g., an applicant may be happy with a trivial *"why am I rejected?"*-explanation, whereas for other stakeholders having to explain complex data pipelines (e.g., how is training data constructed?) and inner workings of algorithms (e.g., how is the model trained?) may be important.

The second group discussed the design of fairer HR technology, and focused on the algorithmic perspective. The first step that came to mind was to not explicitly model protected attributes (e.g., gender or age) in designing algorithms for HR, but because of proxy variables and the existence of bias correlated with these attributes, even without explicitly modeling these features, it is likely the solution may still learn this through correlations and proxy variables. The behavior of recruiters was also discussed, which may turn out to be different between one person to the next; the attendees discussed how this relates to the (un)fairness of the HR tech they use. The final topic discussed in the summary was the impact of the rise of remote working that came with the Covid pandemic. The attendees discussed whether this perhaps could resolve existing unfairness in hiring to some extent (when anyone can be considered), or whether it may cause completely new types of biases?

Finally, the third group discussed human-in-the-loop systems. Their reflection started with the importance of not automating hiring decisions, but rather building "power tools" for helping people make those decisions. This reflection led to the importance of explainability, and, much like the first group, on the diverse requirements or contexts of explanation goals. One characteristic of the recruitment task is the time costliness of the task, and differences between the "decisiveness" of different recruiters. A proper HR tech-power tool should ideally be able to make recruiters aware of biases in the results or underlying data, e.g., through explicitly showing when a result list contains a highly biased or non-diverse list of candidates, and perhaps even take it one step further by pro-actively offering profiles that differ highly from those shown in the results. Finally, the group discussed the challenge of creating trust with the users in these systems they may use, and the challenges of adoption of HR technology in organizations.

3.2 Virtual afternoon slot

3.2.1 Industry keynote

Our industry keynote was given by Kaare Danielsen, CEO of JobIndex, Scandinavia's largest job portal. The invited talk was given live over Zoom, and consisted of a 30 minute keynote with 20 minutes for discussion and Q&A. In his keynote, Kaare presented the history of Jobindex and how they currently work with recruiters and job seekers. He then discussed the dilemma of what to optimize for when matching candidate CVs to job ads and highlighted the benefits and drawbacks of four different optimization targets: (1) clicks from the candidates, (2) ad spending by companies, (3) best content-based match, and (4) value to society. Kaare's keynote was then followed by a lively Q&A session.

3.2.2 Virtual paper presentation session (1)

Our virtual afternoon slot consisted of pre-recorded videos of paper presentations with live Q&A sessions with the presenting authors via Zoom. First in the session, Burke et al. [2021] presented their approach to the Society for Industrial and Organizational Psychology (SIOP) ML Competition, where the task was to create an accurate and fair method to identify potential hires from a collection of new candidates, given as input assessment test data and and historic (binary) labels that corresponded to the (historic) hire's performance, retention, and membership of a(n unspecified) protected group (which was used in the fairness metric). The authors presented their solution to this task, which includes applying spatial partitioning search to re-rank members from the protected group for increased fairness.

Next, Lakhani [2021] presented his idea on leveraging sentiment, intent analysis, and engagement scores from (SMS) chats between candidates and recruiters for candidate recommendations. The presented work uses a custom hand-labeled SMS dataset with a 5-level sentiment score and a multi-class intent classification scheme. The author propose a combination of sentiment scores with engagement scores based on response rates and times from candidates to surface candidates to recruiters.

Finally, Vogiatzis and Kyriakidou [2021] propose a framework for responsible data management for HR, motivated by the prevalence of biased data from direct and indirect sources (proxies) in the HR domain. The framework is based on several data management, organizational, and ethical and legal requirements. The authors proceed to describe a data ecosystem consisting of job applications, job postings, and an integrated knowledge graph to illustrate their responsible data management framework.

3.2.3 Panel

We hosted a panel discussion to bring together practitioners and researchers from different backgrounds to discuss challenges of automatic job recommendation and related tasks.

The panel consisted of Kaare Danielsen (JobIndex), Christiaan Duijst (The Netherlands Institute for Human Rights), Gerd Goetschalckx (VDAB), and Krishnaram Kentapadi (Amazon AWS AI, formerly at LinkedIn), and reflected a wide variety of backgrounds spanning academic HR expertise, industry HR, (semi-)governmental, and with general knowledge on fairness, bias, and transparency of recommender systems.

The panel, chaired by Toine Bogers, discussed the following three topics:

Topic 1: EU framework on AI The European Commission is proposing the first-ever legal framework on AI,[5] which aims to provide AI developers, deployers, and users with precise requirements and obligations regarding specific uses of AI. This legal framework proposes a risk-based approach based on different levels of risk—minimal, limited, high, and unacceptable—in which HR is marked as a high-risk domain.

Kaare expresses his appreciation for the risk-based approach taken, i.e., the distinction between high-impact decisions and those with lesser impact (e.g., algorithms for job seekers vs. personalized advertising, respectively). He mentions that as opposed to the GDPR, which is a more "reactive" regulation, the new AI directive approval needs to be sought before taking an AI system into production, which may create a lot of work but is generally considered to be in the right direction. Krishnaram reflects that when GDPR took effect, technology companies were scrambling to improve their processes and meet requirements (e.g., for the right to be forgotten). However, engineers and others in tech companies may not have initially liked the amount of additional work these regulations created; it was widely supported that these regulations are necessary.

Krishnaram raises one issue the regulations solve; that of different companies being left to decide what is acceptable and what is not. Gerd agrees to this point and stresses how regulation is required to push organizations and tech partners to a more professional level of working with HR data. She stresses that leaving regulated aspects open to interpretations can keep existing biases in place.

One concern raised by Krishnaram is the speed at which our field evolves; he asks whether regulations can anticipate how things will evolve? And, perhaps inadvertently, can these regulations turn into a barrier to innovation? In response to this, Christiaan mentions how going from GDPR to new regulations for algorithms underlines the necessity of continuously adapting to new technological realities. With respect to the concern of regulations "hampering innovation," Christiaan explains there is a provision for so-called "regulatory sandboxes," that allow people to experiment and try things that "technically are not allowed," and facilitate innovation. Even if it does not mean innovation will not be hindered in any way, it shows how regulators do explicitly consider it.

Krishnaram further raises a concern around the notion of fairness and bias, which in the literature mostly comes from western (e.g., US and Europe-based) conceptions of trustworthiness or fairness. What could these corresponding notions of fairness be in other contexts? He illustrates his point with the example of the Indian company Ola recently opening a giant e-scooter factory run entirely by women, what may be an acceptable notion of affirmative action in some contexts, may not be considered acceptable according to US/European laws.

A final observation from Krishnaram is how these regulations highlight the need for communities such as ours to come together and invite people with different backgrounds to address these challenges.

[5] https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai

Topic 2: Fair and inclusive design of HR tech How can we design HR tech in such a manner that will be inclusive of all relevant stakeholders and that its predictions and decision are fair to all people affected by such systems? How can we make HR tech more fair? How do technological and legal issues and 'company values' intertwine in the design and use of such systems?

Gerd states that AI brings a great potential but does not deliver on it yet. She explains how old patterns — in which professionals tell job seekers what they should do — are repeated. There are many different aspects related to choosing a job: e.g., experiencing interest in a job, knowing whether you can or want to move to another area, knowing whether you can or want to find housing there, or a school for your children. These are elements that may not be in the vacancy, which are currently hard to be picked up by AI systems. Gerd summarizes that current AI systems are too limited and traditional, and need further development to move away from the recruiter perspectives to that of job seekers and citizens.

Kaare posits that he is not worried about creating fair systems, but more worried about so-called "unusual candidates," or candidates that may not directly meet all requirements but still be suitable for a job. He explains how in China it is not unusual to receive thousands of applications for an open position; turning automated hriing systems into an absolute requirement, instead of just a nice-to-have tool to support the hiring process. If we would then use the same systems to hire all employees, that could make for a very homogeneous workforce with no diversity in backgrounds. AI systems have no issues finding matching qualifications, and it is easy to argue selecting applicants with the best qualifications is fair, the question is what then happens to these candidates with unusual backgrounds, who may still be qualified for the job?

Christiaan mentions the work by Hangartner et al. [2021] as a great opportunity of using AI with recruiters, where machine learning was applied to monitor the behavior of recruiters, and showed a 7% penalty for women in male professions. He also raises the issue of deciding how regulators should be able to trust the AI recruitment systems to be fair.

Krishnaram adds another dimension of fairness, by explaining how fully automated hiring is not typical for higher-paying jobs but more common for blue-collar or low-wage jobs with many thousands of applications, which offer significant challenges in scaling recruiting. Simple economics of these jobs can make it difficult to justify processing all of the applications by humans, and even if it would economically be possible, it would not be clear whether a human would be able to effectively process such large numbers of applicants without their own biases, and considering time constraints.

Topic 3: Activating hidden workers A recent report by Harvard Business School and Accenture titled "Hidden Workers: Untapped Talent" ([Fuller et al., 2021]) proposed several recommendations for how to hire more so-called 'hidden' workers: part-time workers who would like to work full-time, long-time unemployed who are seeking employment, and people missing from the workforce are willing and able to work under the right circumstances. One of the reasons for these hidden workers to stay invisible is inflexibly configured in HR tech. The panel was asked for the most significant challenges from the report and ideas on addressing them.

Krishnaram mentions how the problems surfaced in the report are broader than HR tech alone. In his experience, many technology companies see bringing those who have been out of the workforce for a number of years back as a huge opportunity. This awareness is an important start towards fixing the underlying challenges. One direction for a technical solution suggested by

Krishnaram is to look more closely at evaluating machine learning models. Krishnaram proposes to move beyond assessing performance on entire populations towards measuring how the model may perform on different subpopulations (for example, previously working mothers who have had a break before going back to the labor market).

In the same light, Gerd mentions that even before HR technology appears in the process, "hidden workers" may result from cultural or psychological aspects hindering people from even engaging with any HR services or technology, e.g., people over the age of 55 may face expectations they will not be able to find employment, leading them to not participate at all.

Kaare explains how the current Danish job market with a shortage of employees helps people with "unusual backgrounds" to get into the job market. Employers may be open to hiring candidates even if they do not meet all the qualifications needed. But even then, automation can still make it difficult to find employment when not all qualifications are met.

Christiaan reiterates to stress the report's main findings; the researchers estimate a total of 27 million hidden workers in the US, partly due to automated hiring. They expect similar numbers in the UK and Germany. In light of these numbers, it concerns him that the Netherlands Institute for Human Rights has never received a single complaint concerning recruitment algorithms. Putting these observations together suggests it may be challenging for people to go to a judge and get their rights. Transparency, or the lack thereof, must play an important role here.

Krishnaram responds to Christiaan's question on the plausibility of the report that, even if the number of hidden workers is not 27 million, it is likely to be in the ballpark in terms of order of magnitude. He also mentions that workers can get displaced due to external factors not mentioned in the report, e.g., globalization and competition from different countries. He illustrates the rust belt area in the US, where high-paying middle-class jobs disappeared as trends changed. In the light of this kind of development, companies and organizations such as LinkedIn, JobIndex, and governmental agencies have an important role, Krishnaram says. They have holistic views of the job market, with which these organizations have access to an early indication of which skills or roles are becoming more or less important. According to Krishnaram, these signals are underused but could effectively guide potential workers in upskilling.

Take-homes Finally, Toine asks each panelist to summarize their main findings and take-homes.

First, Christiaan stresses the importance of the human touch in a domain in which decisions affect people's lives beyond simply the salary they receive, towards giving meaning and purpose. In addition, he says that regulators are waking up to this problem. In the Netherlands new laws are being prepared, and algorithm watchdogs are also emerging. He stresses that lawmakers and algorithm makers should work together; even when speaking different languages, we need to solve these challenges together.

Kaare mentions we discussed bias in AI systems, but there is lots of bias in humans, both recruiters and job seekers. In AI systems, he says, it is easier to measure bias, and if we use AI in the right way, there will be opportunities to improve the way the job market works.

Krishnaram agrees and underlines it is at least as important to understand the bias in broader processes too. For algorithmic components, it is possible to articulate what bias or fairness metrics we want, and once we can articulate that, we can embed them into broader mechanisms. When humans make the same decisions, it can be harder to articulate those decision-making criteria. Hence, we should not dismiss algorithms with humans in the loop altogether, as they may result in better outcomes in some settings. Finally, Krishnaram calls attention to one area that can benefit from collaboration across different disciplines: the bias and behavioral differences in job seekers (e.g., how to take into account the different ways in which men or women may apply for jobs or describe their profiles).

Gerd, finally, mentions there is still a lot of education that needs to be done in society, employers, job seekers, and citizens in creating awareness and understanding of the labor market. This needs to be adapted in HR services too, where there is lots of knowledge and expertise. She also suggests more close collaboration between AI researchers and the human sciences. Finally, Gerd states that it is up to us to keep pushing local, national, and European governments for an inclusive labor market and society.

4 Acknowledgments

We would like to thank all the workshop attendees, our keynote speakers Quirine Eijkman and Kaare Danielsen, our panelists Christiaan Duijst, Gerd Goetschalckx, and Krishnaram Kentapadi, all authors of accepted papers and presenting authors, the RecSys Workshop chairs Jen Golbeck, Marijn Koolen, and Denis Parra. And finally, we would like to thank our program committee: Himan Abdollahpouri, Emma Beauxis-Aussalet, Asia J. Biega, Ludovico Boratto, David Brazier, Robin Burke, Carlos Castillo, Robin De Croon, Michael Ekstrand, Snorre Frid-Nielsen, Sahin Geyik, Nyi-Nyi Htun, Bo Kang, Sepideh Mesbah, Bamshad Mobasher, Manish Raghavan, Nava Tintarev, and Christo Wilson.

The workshop materials, including the program, list of accepted papers and the full recording of our workshop can be found on the dedicated workshop website: https://recsyshr2021.aau.dk/. The workshop proceedings have been published in the CEUR-WS Proceedings: http://ceur-ws.org/Vol-2967/. Finally, the recording of the entire workshop is available on the RecSys YouTube channel at https://www.youtube.com/watch?v=UTsbRBLNuek.

References

RecSys Challenge '16: Proceedings of the Recommender Systems Challenge, New York, NY, USA, 2016. Association for Computing Machinery. ISBN 9781450348010.

RecSys Challenge '17: Proceedings of the Recommender Systems Challenge 2017, New York, NY, USA, 2017. Association for Computing Machinery. ISBN 9781450353915.

Mathieu Bastian, Matthew Hayes, William Vaughan, Sam Shah, Peter Skomoroch, Hyungjin Kim, Sal Uryasev, and Christopher Lloyd. Linkedin skills: Large-scale topic extraction and inference. In *Proceedings of the 8th ACM Conference on Recommender Systems*, RecSys '14, page 1–8, New York, NY, USA, 2014. Association for Computing Machinery. ISBN 9781450326681. doi: 10.1145/2645710.2645729. URL https://doi.org/10.1145/2645710.2645729.

Toine Bogers and Mesut Kaya. An exploration of the information seeking behavior of recruiters. In *Proceedings of the First Workshop on Recommender Systems for Human Resources (RecSys in*

HR 2021) co-located with the 15th ACM Conference on Recommender Systems (RecSys 2021), volume 2967. CEUR-WS, 2021.

Ian Burke, Robin Burke, and Goran Kuljanin. Fair candidate ranking with spatial partitioning: Lessons from the siop ml competition. In *Proceedings of the First Workshop on Recommender Systems for Human Resources (RecSys in HR 2021) co-located with the 15th ACM Conference on Recommender Systems (RecSys 2021)*, volume 2967. CEUR-WS, 2021.

Robert Charlier and Sander Kloppenburg. *Artificial Intelligence in HR: a No-brainer*. PwC, 2017.

Maurits de Groot, Jelle Schutte, and David Graus. Job posting-enriched knowledge graph for skills-based matching. In *Proceedings of the First Workshop on Recommender Systems for Human Resources (RecSys in HR 2021) co-located with the 15th ACM Conference on Recommender Systems (RecSys 2021)*, volume 2967. CEUR-WS, 2021.

Ketki V. Deshpande, Shimei Pan, and James R. Foulds. Mitigating demographic bias in ai-based resume filtering. In *Adjunct Publication of the 28th ACM Conference on User Modeling, Adaptation and Personalization*, UMAP '20 Adjunct, page 268–275, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450379502. doi: 10.1145/3386392.3399569. URL https://doi.org/10.1145/3386392.3399569.

Snorre S. Frid-Nielsen. Find my next job: labor market recommendations using administrative big data. In Toine Bogers, Alan Said, Peter Brusilovsky, and Domonkos Tikk, editors, *Proceedings of the 13th ACM Conference on Recommender Systems, RecSys 2019, Copenhagen, Denmark, September 16-20, 2019*, pages 408–412. ACM, 2019. doi: 10.1145/3298689.3346992. URL https://doi.org/10.1145/3298689.3346992.

Joseph B. Fuller, Manjari Raman, Eva Sage-Gavin, and Kristen Hines. *Hidden Workers: Untapped Talent*. Harvard Business School Project on Managing the Future of Work and Accenture, September 2021. URL https://www.hbs.edu/managing-the-future-of-work/research/Pages/hidden-workers-untapped-talent.aspx.

Francisco Gutiérrez, Sven Charleer, Robin De Croon, Nyi Nyi Htun, Gerd Goetschalckx, and Katrien Verbert. Explaining and exploring job recommendations: a user-driven approach for interacting with knowledge-based job recommender systems. In Toine Bogers, Alan Said, Peter Brusilovsky, and Domonkos Tikk, editors, *Proceedings of the 13th ACM Conference on Recommender Systems, RecSys 2019, Copenhagen, Denmark, September 16-20, 2019*, pages 60–68. ACM, 2019. doi: 10.1145/3298689.3347001. URL https://doi.org/10.1145/3298689.3347001.

Dominik Hangartner, Daniel Kopp, and Michael Siegenthaler. Monitoring hiring discrimination through online recruitment platforms. *Nature*, 589(7843):572–576, 2021. doi: 10.1038/s41586-020-03136-0. URL https://doi.org/10.1038/s41586-020-03136-0.

Mesut Kaya and Toine Bogers. Effectiveness of job title-based embeddings on résumé-to-job ad recommendation. In *Proceedings of the First Workshop on Recommender Systems for Human Resources (RecSys in HR 2021) co-located with the 15th ACM Conference on Recommender Systems (RecSys 2021)*, volume 2967. CEUR-WS, 2021.

Krishnaram Kenthapadi, Benjamin Le, and Ganesh Venkataraman. Personalized job recommendation system at linkedin: Practical challenges and lessons learned. In *Proceedings of the Eleventh ACM Conference on Recommender Systems*, RecSys '17, page 346–347, New York, NY, USA, 2017. Association for Computing Machinery. ISBN 9781450346528. doi: 10.1145/3109859.3109921. URL https://doi.org/10.1145/3109859.3109921.

Akiva Kleinerman, Ariel Rosenfeld, and Sarit Kraus. Providing explanations for recommendations in reciprocal environments. In *Proceedings of the 12th ACM Conference on Recommender Systems*, RecSys '18, page 22–30, New York, NY, USA, 2018. Association for Computing Machinery. ISBN 9781450359016. doi: 10.1145/3240323.3240362. URL https://doi.org/10.1145/3240323.3240362.

Alina Köchling and Marius Claus Wehner. Discriminated by an algorithm: a systematic review of discrimination and fairness by algorithmic decision-making in the context of hr recruitment and hr development. *Business Research*, 13(3):795–848, 2020. doi: 10.1007/s40685-020-00134-w. URL https://doi.org/10.1007/s40685-020-00134-w.

Emanuel Lacic, Markus Reiter-Haas, Tomislav Duricic, Valentin Slawicek, and Elisabeth Lex. Should we embed? a study on the online performance of utilizing embeddings for real-time job recommendations. In *Proceedings of the 13th ACM Conference on Recommender Systems*, RecSys '19, page 496–500, New York, NY, USA, 2019. Association for Computing Machinery. ISBN 9781450362436. doi: 10.1145/3298689.3346989. URL https://doi.org/10.1145/3298689.3346989.

Ashish Lakhani. Recommendations for recruiters with sentiment detection. In *Proceedings of the First Workshop on Recommender Systems for Human Resources (RecSys in HR 2021) co-located with the 15th ACM Conference on Recommender Systems (RecSys 2021)*, volume 2967. CEUR-WS, 2021.

Dor Lavi, Volodymyr Medentsiy, and David Graus. consultantbert: Fine-tuned siamese sentence-bert for matching jobs and job seekers. In *Proceedings of the First Workshop on Recommender Systems for Human Resources (RecSys in HR 2021) co-located with the 15th ACM Conference on Recommender Systems (RecSys 2021)*, volume 2967. CEUR-WS, 2021.

Cynthia C. S. Liem, Markus Langer, Andrew Demetriou, Annemarie M. F. Hiemstra, Achmadnoer Sukma Wicaksana, Marise Ph. Born, and Cornelius J. König. *Psychology Meets Machine Learning: Interdisciplinary Perspectives on Algorithmic Job Candidate Screening*, pages 197–253. Springer International Publishing, Cham, 2018. ISBN 978-3-319-98131-4. doi: 10.1007/978-3-319-98131-4_9. URL https://doi.org/10.1007/978-3-319-98131-4_9.

Yoosof Mashayekhi, Bo Kang, Jefrey Lijffijt, and Tijl De Bie. Quantifying and reducing imbalance in networks. In *Proceedings of the First Workshop on Recommender Systems for Human Resources (RecSys in HR 2021) co-located with the 15th ACM Conference on Recommender Systems (RecSys 2021)*, volume 2967. CEUR-WS, 2021.

Dena F Mujtaba and Nihar R Mahapatra. Ethical considerations in ai-based recruitment. In *2019 IEEE International Symposium on Technology and Society (ISTAS)*, pages 1–7. IEEE, 2019.

I. Naim, M. I. Tanveer, D. Gildea, and M. E. Hoque. Automated analysis and prediction of job interview performance. *IEEE Transactions on Affective Computing*, 9(2):191–204, 2018. doi: 10.1109/TAFFC.2016.2614299.

Ioannis Paparrizos, B. Barla Cambazoglu, and Aristides Gionis. Machine learned job recommendation. In *Proceedings of the Fifth ACM Conference on Recommender Systems*, RecSys '11, page 325–328, New York, NY, USA, 2011. Association for Computing Machinery. ISBN 9781450306836. doi: 10.1145/2043932.2043994. URL https://doi.org/10.1145/2043932.2043994.

Manish Raghavan, Solon Barocas, Jon Kleinberg, and Karen Levy. Mitigating bias in algorithmic hiring: Evaluating claims and practices. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, FAT* '20, page 469–481, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450369367. doi: 10.1145/3351095.3372828. URL https://doi.org/10.1145/3351095.3372828.

Amar Saini, Florin Rusu, and Andrew Johnston. Privatejobmatch: a privacy-oriented deferred multi-match recommender system for stable employment. In Toine Bogers, Alan Said, Peter Brusilovsky, and Domonkos Tikk, editors, *Proceedings of the 13th ACM Conference on Recommender Systems, RecSys 2019, Copenhagen, Denmark, September 16-20, 2019*, pages 87–95. ACM, 2019. doi: 10.1145/3298689.3346983. URL https://doi.org/10.1145/3298689.3346983.

Javier Sánchez-Monedero, Lina Dencik, and Lilian Edwards. What does it mean to 'solve' the problem of discrimination in hiring?: social, technical and legal perspectives from the UK on automated hiring systems. In Mireille Hildebrandt, Carlos Castillo, Elisa Celis, Salvatore Ruggieri, Linnet Taylor, and Gabriela Zanfir-Fortuna, editors, *FAT* '20: Conference on Fairness, Accountability, and Transparency, Barcelona, Spain, January 27-30, 2020*, pages 458–468. ACM, 2020. doi: 10.1145/3351095.3372849. URL https://doi.org/10.1145/3351095.3372849.

K. Umemoto, T. Milo, and M. Kitsuregawa. Toward recommendation for upskilling: Modeling skill improvement and item difficulty in action sequences. In *2020 IEEE 36th International Conference on Data Engineering (ICDE)*, pages 169–180, 2020. doi: 10.1109/ICDE48307.2020.00022.

Dimitrios Vogiatzis and Olivia Kyriakidou. Responsible data management for human resources. In *Proceedings of the First Workshop on Recommender Systems for Human Resources (RecSys in HR 2021) co-located with the 15th ACM Conference on Recommender Systems (RecSys 2021)*, volume 2967. CEUR-WS, 2021.

Christo Wilson, Avijit Ghosh, Shan Jiang, Alan Mislove, Lewis Baker, Janelle Szary, Kelly Trindel, and Frida Polli. Building and auditing fair algorithms: A case study in candidate screening. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '21, page 666–677, New York, NY, USA, 2021. Association for Computing Machinery. ISBN 9781450383097. doi: 10.1145/3442188.3445928. URL https://doi.org/10.1145/3442188.3445928.

F. Zuiderveen Borgesius. *Discrimination, artificial intelligence, and algorithmic decision-making*. Strasbourg: Council of Europe, Directorate General of Democracy, 2018.