Using Interaction Signals for Job Recommendations

Benjamin Kille[1], Fabian Abel[2], Balázs Hidasi[3], Sahin Albayrak[1]

[1] Berlin Institute of Technology,
benjamin.kille@tu-berlin.de
sahin.albayrak@dai-labor.de
[2] XING AG,
fabian.abel@xing.com
[3] Gravity Research,
hidasi.balazs@gravityrd.com

Abstract. Job recommender systems depend on accurate feedback to improve their suggestions. Implicit feedback arises in terms of clicks, bookmarks and replies. We present results from a member inquiry conducted on a large-scale job portal. We analyse correlations between ratings and implicit signals to detect situations where members liked their suggestions. Results show that replies and bookmarks reflect preferences much better than clicks.

Key words: job recommendation, interactions, reciprocity, survey, ratings

1 Introduction

Online job portals are becoming more and more popular among professionals as well as recruiters. They facilitate exchanging information. Professionals gain instant access to newly added job offers. Recruiters can spread their job advertisements to a larger base of recipients. Increasingly simple access to larger collections of data induces an information overload problem. Professionals struggle to discover interesting job offers. Recruiters struggle to discover suited candidates. The portals seek to support their members to overcome these issues. They incorporate mechanisms matching professionals and job offers based on available data. The vast set of professionals reduces to few candidates. The overwhelming collection of job offers reduces to few positions. The quality of the reduction depends on how well the system estimates recipients' preferences. Professionals expect the system to display relevant positions. Recruiters expect the system to display candidates best suited for the respective position. We refer to systems automatically learning users' preferences as *recommender systems*. Research on recommender systems has introduced a plethora of algorithms. Collaborative Filtering (CF) (cf. Koren and Bell [2]) and Content-based Filtering (CBF) (cf. Lops et al. [4]) represent two wide-spread paradigms. CF takes a collection of preferences and infers unknown (user, item)-pairs thereof. CBF additionally considers features describing items. Both techniques require data expressing the preferences between user and items. Professionals include information on their general preferences in their profiles. Still, job portals lack preferences directed toward specific job offers. They have to infer such information from implicit signals. Whether professionals appreciate recommended job offers depends on a variety of aspects. For instance, they may like the job description but be unwilling to move to another location. We collaborate with the job portal XING[1]. XING is the leading job portal in the German speaking world with millions of members. There we observe professionals interacting with job postings in different ways. Types of interactions include clicking, bookmarking, and replying. Interactions have a time stamp assigned. This gives rise to temporal analysis. We look for patterns indicating situations in which professionals show interest in a job posting. Alternatively, we search for activities indicating interest toward a job posting. Our contributions are three-fold:

– we analyse interaction signals on a job portal,
– we conducted an inquiry providing ratings for job recommendations,
– we investigate the correlation between interaction signals and ratings.

Frist, we present previous works on job recommendation in Section 2. Section 3 illustrates the data we observed on a large-scale job portal. Section 4 describes the notion of relevance and how we learn correlations with interaction signals. We conclude and discuss future research in Section 6.

2 Related Work

Malinowski et al. [5] investigated how to automatically match professionals and open positions. Matching professionals and jobs requires a bilateral perspective. They argue that combining two recommenders promises high matching quality. On the one hand, we ought to determine relevant positions for a given professional. On the other hand, we should find the most promising candidates for an open position. Combining both yields matches satisfying both professionals and recruiters needs. Malinkowski et al. [5] conduct a user study to determine how well their recommender matches professionals and open positions. However accurate results such a study yields, a large-scale system will lack resources to apply similar evaluation protocols. Mine et al. [6] extend the idea as they consider interactions between professionals and recruiters. Their approach iteratively updates matching lists of recruiters and professionals as they exchange messages. The evaluation considers the time taken to establish matchings along with matching quality. Mine et al. [6] simulate the actions of professionals and recruiters. Behaviours of professionals and recruiters are prone to change over time. Hence, we expect to observe a more representative picture of interactions in a real, large-scale system. Paparrizos et al. [7] modell the problem as a supervised machine learning task. They take a data set of job transitions. Their

[1] www.xing.com

method predicts which position a given professional will accept next. The data comprise a selection of large companies. They measure the classification accuracy. Contrarily, we consider a more comprehensive view on the labour markets. Our ability to represent the problem as classification task depends on the cardinality of potential class labels. Paparrizos et al. [7] obtain a manageable set of labels focusing on few companies. A large-scale job portal offers thousands of companies. Therefore, the classification problem would become too complex to manage. Recommender systems matching people to one another are referred to as *reciprocal recommender systems*. Online dating websites matching partners have been subject to many studies (cf. Akehurst et al. [1], Kunegis et al. [3], Pizzato et al. [8], and Zhao et al. [9]). Online dating matches exactly two persons. On the other hand, recruiters might be looking for several candidates at once to fill several positions. Thus, we face a slightly different problem. We collect observations from a large-scale job portal. Thereby, we expect to discover new insights enlightening interactions between professionals and job offers. In particular, we investigate correlations between explicitly stated preferences and implicitly observed signals.

3 Data Description

In this section, we describe the observable signals. Our data capture professionals, job offers, and interactions between both. Professionals create profiles describing themselves and their careers. These profiles include demographics, education, and interests. The job portal can track activities by identifying members. Recruiters create job offers. These offers outline required skills and portray a candidate's characteristics. In addition, they introduce the hiring company. We can track interactions with such job offers as they can be identified. Job portals let their members interact with offered jobs in different ways. Professionals may click, bookmark, or reply to suggested job postings. Additionally, they can query an integrated search engine. Details refer to random sample of more than one million members of XING. The sample has been taken over a period of several weeks in early 2015. Note that clicks, bookmarks, and replies may occur in any order. For instance, members may bookmark job postings without clicking. We highlight actions from the perspectives of professionals, job offers, and their combinations.

3.1 User Activity

We observe varying levels of activity among members. Professionals can either click, bookmark, or reply to a recommended job offer. Replies comprise three actions. First, professionals can access the job offering company's website to apply for the position. Second, professionals can request additional information about the offered position. Third, professionals can message the recruiter. Clicks require fairly low levels of cognitive effort. Conversely, replies force professionals to carefully study job offers. Our data reflect this aspect. We observe on average 95.7 % clicks, 3.3 % replies, and 1.0 % bookmarks. Our data cover a range of 3 months. Therein, professionals click on average 6.5 (standard deviation $\sigma = 15.6$) job postings. They bookmark on average 2.0 ($\sigma = 3.4$) job postings. They reply on average to 3.1 ($\sigma = 6.9$) job offers.

3.2 Item Activity

Job postings attract varying numbers of members. On average job offers obtain 87.6 ($\sigma = 219.5$) clicks. The large standard deviation indicates a popularity bias. Some job offers appear to receive much more clicks than others. Similarly, job positions obtain on average 4.2 ($\sigma = 7.2$) bookmarks. Finally, job offers attract on average 6.3 ($\sigma = 11.6$) replies.

3.3 Interaction Activity

Professionals spread limited spans of attention across the collection of job postings. Similarly, job offers target varying subsets of professionals according to their skills. Both phenomena combined cause a sparse pattern of interactions between professionals and offered jobs. We compute the density for each type of interaction. Let I refer to the number of interactions. Further, let U and J refer to the numbers of members and open positions. The density is defined as $\rho = {}^{I}/_{UJ}$. We observe densities of $\rho_{\text{clicks}} = 0.01\,\%$ for clicks, $\rho_{\text{bookmarks}} = 0.01\,\%$ for bookmarks, and $\rho_{\text{replies}} = 0.01\,\%$ for replies. In other words, if we were to pick a pair of professional and job posting at random, chances of observing any type of interactions are ≈ 1 in 10 000. This implies that the job portal will struggle to estimate relevance of most combinations of members and open positions.

We notice that we observe sparse signals. We do not expect to obtain a clear picture of preferences from professionals by observing their interactions with job postings. Job portals depend on preferences to filter the most promising positions. Thus, we have to avoid misinterpreting interactions. Members might click on suggestions due to curiosity. Further, professionals might like some aspects of the position. Simultaneously, they may object to other aspects. For instance, a given position's description sounds interesting. At the same time, the position is located far from the current residence. Consequently, the position as a whole becomes irrelevant.

4 Relevance Prediction

Previous research determined the relevance of suggested jobs by simulation and small scale user studies (see Section 2). Large-scale job portals cannot rely on few individuals. We cannot expect preferences of thousands of members to conform to pre-defined simulation parameter. Hence, we propose to estimate preferences from signals derived from interactions between professionals and job postings. As a first step, we validate the information conveyed by these signals. We analyse the correlation between observed signals and ratings obtained via member inquiry. We asked professionals how they perceive job recommendations in form of an online survey. Professionals assigned stars from 1 to 5 or selected a 0 indicating they were unable to quantify their preference. Subsequently, we analysed how the responses correlate with professionals' interactions.

4.1 Clicks ∼ Ratings

First, we relate clicks and ratings. We filtered all clicks of suggested job offers which professionals had rated. Figure 1 illustrates the relation. We encode proportions in form of colours. We observe the largest proportion referring to few clicks. Additionally, we observe that ratings vary throughout the scale with a limited number of clicks. We conjecture that few clicks fail to indicate a high level of relevance. Higher ratings dominate the range of 5 to 10 clicks. We conclude that as professionals increasingly click on suggested job offers, they become more relevant.

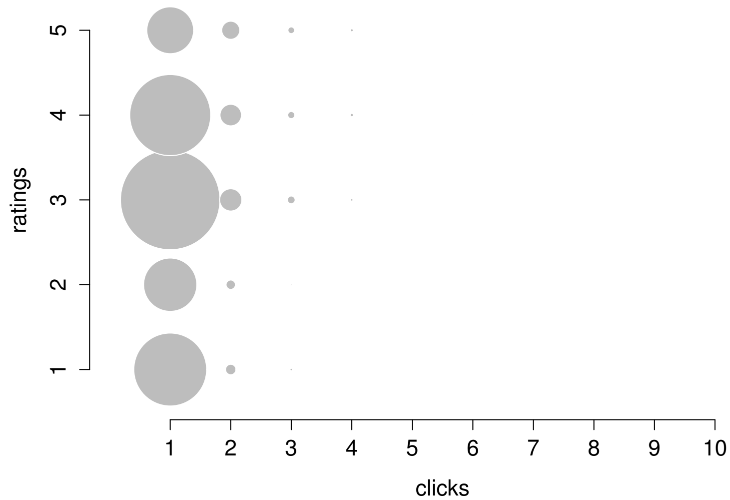

Fig. 1. Relation between clicks and ratings. The size of the circles encodes the proportion of data points. We observe that few clicks fail to indicate high relevancy. Professionals clicking once express both dislike and like toward recommended jobs. Professionals tend to click more often on relevant positions.

4.2 Bookmarks ∼ Ratings

Why would professionals bookmark job postings? We suppose that bookmarking indicates a higher level of relevance than clicking. Figure 2 confirms our intuition. We filtered ratings assigned to bookmarked job offers. We observe that these job postings obtained on average 3.6 out of 5 stars. The distribution is skewed in favour of higher ratings. Hence, we conclude that bookmarks better reflect relevance than clicks do.

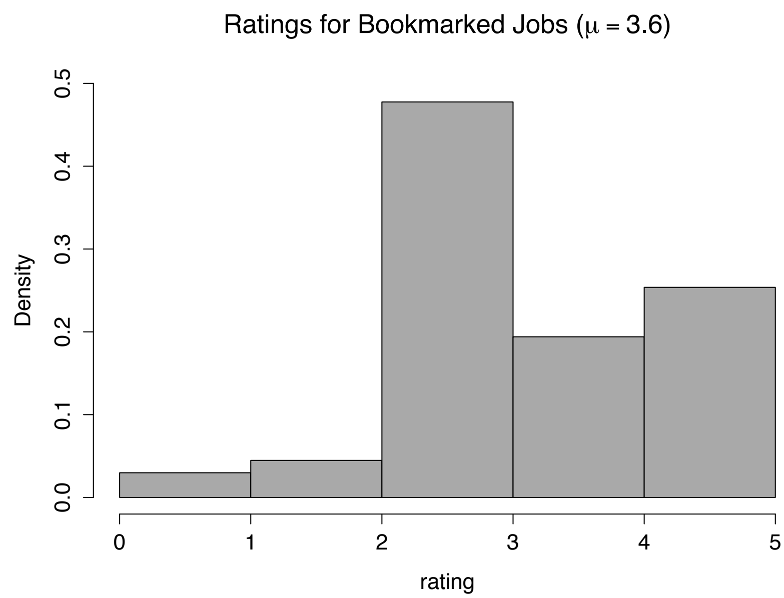

Fig. 2. Distribution of Ratings for bookmarked Job Offers. We observe that bookmarked job offers obtain rather positive ratings. Professionals rated bookmarked job offers on average with 3.6 out of 5 starts.

4.3 Replies ∼ Ratings

Replying refers to three possible actions. Professionals can either message the recruiter, request additional information, or access an external application form for the position. The systems allows multiple actions for an offered position. Figure 3 relates the number of replies with the ratings. We observe that few replies with high ratings collect the largest proportion. Thereby, we conclude that replies reliably indicate relevance.

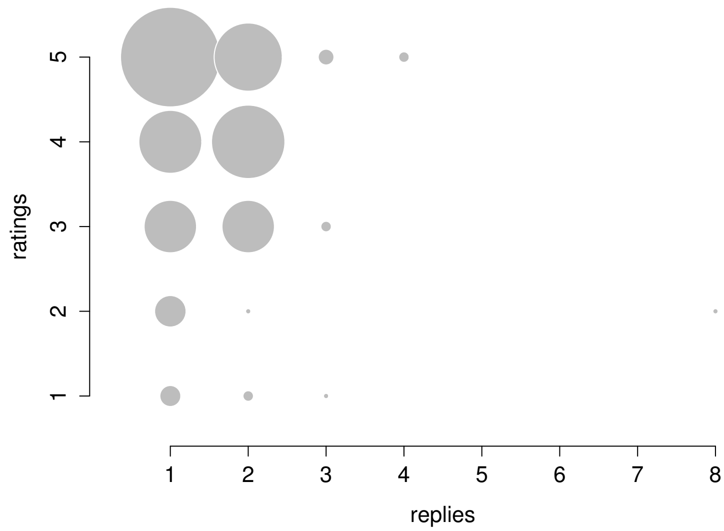

Fig. 3. Relation between Replies and Ratings. The figure encodes proportions by the size of the circle. We observe the largest proportion for a single reply and the rating 5. In addition, the next largest proportions gather around high ratings and few replies.

5 Discussion

Recommending jobs is subject to vastly different requirements than recommending movies, music, or products. Recruiters require open positions to be filled within a specified time. In addition, positions are frequently open to at most one professional. In contrast, arbitrarily many users can consume movies, music, and products. Some items allow re-consumption, for instance, watching a movie several times. Users interact with recommender systems with varying motivation. Some interactions are due to information needs or decision support. Other interactions occur as users satisfy their curiosity in exploratory fashion. Movie recommender systems tend toward explorative use. On the other hand, job recommender systems deal much more with decision support. As a result, they face users with interchanging periods of activity and absence thereof. Looking for a new position, professionals intensively interact with the system. Satisfied with their working environment, professionals scarcely interact with the system. Conversely, movie recommender systems face a more stable condition with users regularly visiting. Furthermore, movie recommender systems observe the effect of their recommendations. They can track how long their users watch a movie. Job recommender systems cannot monitor the interactions between employer and candidate outside their platform.

6 Conclusion and Future Work

Determining situations where professionals deem suggested job offers relevant is crucial to improve job recommender systems. We analysed observable signals incurring as professionals interact with job postings. We showed results of a member inquiry providing explicit preferences. We investigated how observable signals and responses from the inquiry correlate. Although, clicks are more common than bookmarks and replies, they convey less information concerning the relevance of a suggestion. Replying gives the best indication on whether a professional deemed a recommended job offer relevant.

This knowledge will support our future research. We will conduct experiments with a variety of existing recommendation algorithms. The selection includes collaborative filtering, content-based filtering, location-aware recommendation, and context-aware filtering. We will investigate which algorithm or combination of algorithm maximises replies and bookmarks.

7 Acknowledgements

The research leading to these results has received funding from European Union Seventh Framework Programme (FP7/2007–2013) under Grant Agreement № 610594.

References

1. Joshua Akehurst, Irena Koprinska, Kalina Yacef, Luiz Pizzato, Judy Kay, and Tomasz Rej. CCR - A content-collaborative reciprocal recommender for online dating. In *IJCAI*, pages 2199–2204, 2011.
2. Yehuda Koren and Robert Bell. *Recommender Systems Handbook*, chapter Advances in Collaborative Filtering, pages 145–184. Springer, 2011.
3. Jérôme Kunegis, Gerd Gröner, and Thomas Gottron. Online dating recommender systems: the split-complex number approach. In *ACM RecSys workshop on Recommender systems and the social web*, pages 37–44, 2012.
4. Pasquale Lops, Marco de Gemmis, and Giovanni Semeraro. *Recommender Systems Handbook*, chapter Content-based Recommender Systems: State of the ARt and Trends, pages 73–100. Springer, 2011.
5. Jochen Malinowski, Oliver Wendt, Tobias Keim, and Tim Weitzel. Matching people and jobs: A bilateral recommendation approach. *Proceedings of the Annual Hawaii International Conference on System Sciences*, 6(C):1–9, 2006.
6. Tsunenori Mine, Tomoyuki Kakuta, and Akira Ono. Reciprocal Recommendation for Job Matching with Bidirectional Feedback. (2009):39–44, 2013.
7. Ioannis Paparrizos, Barla Cambazoglu, and Aristides Gionis. Machine Learned Job Recommendation. In *RecSys*, pages 325–328, 2011.
8. Luiz Pizzato, Tomek Rej, Thomas Chung, Irena Koprinska, and Judy Kay. RECON: a reciprocal recommender for online dating. In *ACM RecSys*, pages 207–214, 2010.
9. Kang Zhao, Xi Wang, Mo Yu, and Bo Gao. User recommendations in reciprocal and bipartite social networks-An online dating case study. *IEEE Intelligent Systems*, 29:27–35, 2014.