JPLink: On Linking Jobs to Vocational Interest Types

Amila Silva¹(✉), Pei-Chi Lo², and Ee-Peng Lim²

¹ School of Computing and Information Systems, The University of Melbourne,
Melbourne, Australia

amila.silva@student.unimelb.edu.au

² School of Information System, Singapore Management University,
Singapore, Singapore

pclo.2017@phids.smu.edu.sg, eplim@smu.edu.sg

Abstract. Linking job seekers with relevant jobs requires matching based on not only skills, but also personality types. Although the Holland Code also known as RIASEC has frequently been used to group people by their suitability for six different categories of occupations, the RIASEC category labels of individual jobs are often not found in job posts. This is attributed to significant manual efforts required for assigning job posts with RIASEC labels. To cope with assigning massive number of jobs with RIASEC labels, we propose JPLINK, a machine learning approach using the text content in job titles and job descriptions. JPLINK exploits domain knowledge available in an occupation-specific knowledge base known as O*NET to improve feature representation of job posts. To incorporate relative ranking of RIASEC labels of each job, JPLINK proposes a listwise loss function inspired by learning to rank. Both our quantitative and qualitative evaluations show that JPLINK outperforms conventional baselines. We conduct an error analysis on JPLINK's predictions to show that it can uncover label errors in existing job posts.

Keywords: Job profiling · Representation learning · Learning to rank

1 Introduction

Motivation. Job profiling refers to uncovering important characteristics of jobs for generating useful insights about job trends and for matching jobs with talents. According to Holland's theory [5], each occupation (or applicant) can be assigned 1 to 3 out of 6 personality types characterizing different personality types. These personality types are assigned RIASEC labels: REALISTIC (R), INVESTIGATIVE (I), ARTISTIC (A), SOCIAL (S), ENTERPRISING (E) and CONVENTIONAL (C) (see Fig. 1a). Ideally, one should match people with jobs based on personality

This research was supported by the National Research Foundation, Prime Minister's Office, Singapore under its International Research Centres in Singapore Funding Initiative.

Fig. 1. (a) The Hexagonal Model of Holland’s vocational interest types (source: [9]); and (b) a few examples for jobs and occupations (source: O*NET)

types, assuming that all other job criteria (e.g., skills, abilities, knowledge, etc.) are already met. For example, doctor, researcher, and lawyer are jobs ideal for people with INVESTIGATIVE personality type, while photographer, musician, and architect are jobs ideal for people with ARTISTIC personality type.

Nevertheless, RIASEC labels are usually not found in the job descriptions. Experts in the past attempted to focus on manually performing personality type annotations at the occupation level as each occupation represents a set of jobs involving similar job tasks (see Fig. 1b for examples). Such an approach assumes that jobs of the same occupation share the same personality types. This assumption does not always work well when people are often expected to be recommended specific jobs instead of occupations. Also, the manual approach does not timely profile new occupations which are expected to emerge at faster pace due to recent technology disruptions.

Objectives. We aim to determine the personality types of a large collection of job posts. In this task, one has to address a few research challenges: (a) limited labeled data; (b) noisy word semantics; and (c) ranked personality types.

First, there are very limited labeled data available for training and evaluation. One can certainly find pre-existing occupations labeled with personality types but not jobs. To the best of our knowledge, the number of occupations is usually at the scale of less than 1500, much smaller than the millions of job posts available. In this work, we use the labeled data available in O*NET at the occupation level, which consists of 974 different occupations with RIASEC labels. Without labeled jobs as ground truth, we have to use labeled occupations in some distant supervision solution approach.

Second, the description for the same job can be vastly different due to different word choices, job scopes and requirements. One has to accommodate these differences in developing an accurate prediction model. Although there are

previous attempts to use pre-trained word embedding to profile job posts [16], domain-specific words are not accurately represented in pre-trained word embeddings as they often do not carry their domain-specific meanings in a general corpus. For instance, the word *spark* in a software developer job mostly refers to a cluster computing framework, although it means “an emission of fire or electricity” in a general corpus.

Finally, up to three personality types can be associated with a job post and they are ranked. We thus require the prediction model to recover the ranking.

Contributions. To address these challenges, we propose JPLINK, a framework to profile jobs with their personality types, which: (1) jointly learns domain-specific word and occupation representations using knowledge available in O*NET. To the best of knowledge, this is the first work to learn word representations specific to occupations and jobs; (2) incorporates a novel supervised approach to assign RIASEC labels to occupations and job posts using their text context, which considers the inter-correlations among RIASEC dimensions. JPLINK outperforms the conventional baselines by 4.86% and it yields a high ranking accuracy measured by NDCG (=0.949); and (3) predicts the personality type labels for a set of job posts, extracted from Singapore’s JobsBank¹, for which weak labels are assigned using a distant supervision approach. Our error analysis shows JPLINK can effectively overcome the imperfections caused by the limitation of assuming all jobs of the same occupation share the identical personality type.

2 Related Work

Personality types have been extensively studied using empirical evidence. The study in [10] provides empirical evidence of the personality type dimensions underlying the hexagon of Hollands’ theory and it shows that individuals’ personality types characterize the tasks they prefer to perform in jobs. [11] introduces a procedure to associate occupations with Holland’s hexagon (*Hexagon Congruence Index*) using the personality type scores and empirically shows that *Hexagon Congruence Index* provides a basis for a new index of congruence (e.g., person-occupation, occupation-occupation), by which similar occupations and user profiles can be identified.

Nevertheless, there are very little efforts on assigning personality types to jobs as opposed to occupations. To the best of our knowledge, there are three types of methods for occupational level personality type determination: (a) *Incumbent Method*, in which the personality type of an occupation is the average of the personality types of a representative sample of workers taking jobs of that occupation (based on the Holland’s idea that people sharing the same occupation represent the occupation). This method requires the personality types of many individuals and is therefore very costly (i.e., less practical); (b) *Empirical Method*, which uses occupational analysis data (collected using *Incumbent*

¹ A central repository of job posts in Singapore <https://www.jobsgov.sg/>.

Table 1. Descriptive Statistics of O*NET knowledge base

	Actual correlations						% of Proportions
	C	E	I	S	R	A	
Conventional (C)	1.00	0.21	-0.28	-0.34	-0.05	-0.50	23.76%
Enterprising (E)		1.00	-0.46	0.15	-0.51	-0.06	14.75%
Investigative (I)			1.00	-0.09	-0.17	0.11	16.01%
Social (S)				1.00	-0.60	0.24	11.16%
Realistic (R)					1.00	-0.39	28.18%
Artistic (A)						1.00	6.13%

Method) to develop classifiers to predict the personality type of occupations; and (c) *Judgment Method*, which involves trained experts making direct personality type assignment to occupations. An example for such an approach is O*NET interest profiler [12]. It constructs initial personality type of occupations based on discriminant functions derived from the ratings of Occupational Units [12]. Judgment method is then performed to fine-tune the constructed personality profiles using expert knowledge. Usually, human judges are trained to determine occupations’ personality profiles after looking at attributes of occupations such as title, description, and job tasks. Hence, such an approach requires a considerable human effort. It is also difficult to extend this kind of approach for new emerging occupations and jobs. Almost all these previous efforts share the assumption of all job posts of the same occupation share the same personality type. This differentiates our research from them.

3 Dataset Construction

3.1 Occupation-Specific Knowledge Base Extraction

For the purpose of learning domain-specific word embedding as well as for evaluation of interest profile prediction, we crawled the O*NET occupation knowledge base which covers 1110 different occupations and their RIASEC profiles, which is publicly available at <https://www.onetonline.org/>. Each occupation has a profile consisting of six numerical scores in the range [0, 100], one for each RIASEC dimension. The dimensions with scores more than 50 are known as the *interest codes* of the occupation. For example, an occupation with profile of (R = 80, I = 40, A = 20, S = 60, E = 10, C = 70) will be assigned the Holland codes R, C and S. The rightmost column in Table 1 shows the RIASEC distribution for the 1110 occupations found in O*NET. Table 1 depicts the Spearman’s correlation between personality type dimensions. We observe that the correlations among dimension is consistent with Holland’s Hexagonal Model (Fig. 1a). The opposite and consecutive dimensions in Holland’s Hexagon have significant negative correlations (e.g., Realistic vs Social, Enterprising vs Investigative, and Artistic vs Conventional) and positive correlations (e.g., Artistic vs Social) respectively.

In other words, jobs for realistic people may not suit social people, jobs for enterprising people may not suit investigative people, etc. On the other hand, artistic people may be able to take on social jobs.

O*NET also identifies similar occupations for each occupation. We thus construct a network of similar occupations known as *occupation network* and measure homophily of RIASEC labels in *occupation network* using *Affinity* [7] measure, which is defined as the ratio between the observed fraction of links between interest dimension sharing occupations in the network, and the expected fraction of links between interest dimension sharing occupations. Here, we assign each occupation in the network with its highest-scored RIASEC dimension. The *occupation network* reports an *Affinity* of 2.82 ($\gg 1$), implying that occupations connected by a link in the network are 2.82 times more likely to share the similar dimension than that between any two random occupations. Such a strong homophily property demonstrates the importance of taking advantage of O*NET occupation network information for RIASEC prediction.

3.2 Extraction of Job Posts from Singapore’s JobsBank

To quantitatively evaluate JPLINK, we crawled a set of 217,874 job posts from Singapore’s Jobsbank posted during the period from September 2017 to December 2018. Each job post consists of fields such as job title, skill description, and SSOC (Singapore Standard Occupational Classification)² category. These job posts however do not come with any personality type profile.

Weakly Labelling Personality Type Profile of Job Posts. As mentioned in Sect. 2, personality type profiles are only available at the occupation level, not at the job level. While each job post has an SSOC code corresponding to some occupation, the code is not associated with RIASEC profiles. To derive the latter, we propose a distant supervised approach to map SSOC occupations to most similar O*NET occupations (for which personality types are assigned). Since there is no direct mapping between SSOC and SOC, we first use a mapping table³ matching SSOC occupation codes with ESCO occupation codes which are the occupation codes standardized in the EU region. Subsequently, we determine the most similar SOC occupations for each ESCO occupation using another mapping table from Bureau of Labor Statistics of USA⁴. Finally, the RIASEC profile of an SSOC occupation is defined as the average of its similar SOC occupations’ profiles. Following the assumption that jobs have RIASEC profiles similar to their occupations, we assign weak interest profile labels to the job posts. In this way, we are able to map 96.71% of SSOC occupations to 96.57% of ESCO occupations, and finally 75.23% of SOC occupations. This amounts to 171,946 job posts assigned with weak interest profile labels, which are used in Sect. 6 to train JPLINK and to evaluate the prediction accuracy.

² <https://www.singstat.gov.sg/standards/standards-and-classifications/ssoc>.

³ <https://www.singstat.gov.sg/ssoc2015-v2018-isco-08-correspondence.xlsx>.

⁴ https://www.bls.gov/soc/ISCO_SOC_Crosswalk.xls

4 RIASEC Profile Prediction Problem

We define the RIASEC profile prediction problem to consist of (a) learning domain-specific representations for words and occupations; and (b) prediction of personality profiles (RIASEC dimensions) for occupations and job posts.

Fig. 2. JPLINK framework

Formally, let $O = \{o_1, o_2, \dots, o_n\}$ is the 1110 ($= n$) occupations available in O*NET. Each occupation $o \in O$ is a tuple $\langle W_o, E_o, y_o \rangle$, where W_o is a sequence of words describing the tasks, job activities, and other aspects of occupation o , and E_o is the set of similar occupations to o . The personality profile of o is denoted as y_o a vector with elements $y_o^d \in \mathbb{R}^{[0,100]}$ which is the score of the d^{th} dimension, where $d \in \{R, I, A, S, E, C\}$ represents the personality type dimensions.

Learning Domain-Specific Word and Occupation Representations. The goal of this task is to learn mapping function $f : w|o \rightarrow \mathbb{R}^{1 \times k}$ for $o \in O, w \in W$, where $W = \bigcup_o W_o$ and k is the dimension of the embedding space. W_o denotes the set of words in occupation o .

Prediction of RIASEC/Personality Type Profiles for Occupations and Job Posts. Here, our goal is to predict personality type profile $\hat{y}_o$ for a given $o \in O$ or $\hat{y}_j$ for a given job post $j \in J$, where J is the set of job posts. A job post j is a tuple $\langle W_j^t, W_j^d \rangle$, where both W_j^t and W_j^d are word sequences to represent j 's title and description respectively.

5 JPLink

Our proposed JPLINK framework consists of three parts as depicted in Fig. 2: (1) learning job-domain specific embedding for words and occupations (JPLINK-EMB); (2) prediction of personality type profiles for occupations and jobs (JPLINK-PRED); and (3) representation construction for new occupations or job posts.

5.1 Learning Job-Domain Specific Embedding for Words and Occupations (JPLink-Emb)

JPLINK-EMB is motivated by skip-gram model [6], a popular language model in NLP. For a given set of sentences, the skip-gram model loops on the words in each sentence and uses the current word to predict its neighbors. Formally, for a given sentence $w_1, w_2, w_3, \dots, w_T$, the skip-gram model maximizes:

$\frac{1}{T} \sum_{t=1}^T \sum_{-c \leq j \leq c} \log p(w_{t+j}|w_t)$ where c is the window size around the center word w_t . By doing so, representations of the words with similar context will be closed to each other in the embedding space (second order proximity is preserved). Generally, the skip-gram model is trained on a large general text corpora (i.e., Wikipedia and Google News) which do not carry any occupation-specific knowledge. The trivial way of constructing representations for multi-word entities is taking the mean representation of the words in the particular entity (e.g., the representation of “computer programmer” is the mean of the representations of “computer” and “programmer”). We refer this approach as PRETRAINED EMB for the rest of this paper.

Knowledge Graph Construction. In our model, a knowledge graph G is constructed, in which both $w \in W$ and $o \in O$ are nodes. We create edges between two occupations by connecting each $o \in O$ to the similar occupations in E_o . For each $o \in O$, a set of discriminative words are selected from W_o and these are connected to the o ’s node. This derives a graph G with useful first-order context to represent each node (word or occupation).

We perform random walks around each node in G and then the skip-gram model is used to learn representations for nodes, considering random walk sequences are analogous to sentences in a language. Due to the importance of first-order context, we perform breadth-first random walks with random restarts (inspired by node2vec [4]). We observe that such a procedure provides the relevant context for each node and outperforms conventional network representation learning techniques (e.g., deepwalk [8] and LINE [14]) for this particular task.

In this work, the discriminative words are determined empirically by their normalized document frequency $< 10\%$. We leave the exploration of other threshold settings and the embedding schemes for knowledge graph to future work.

5.2 Predictions of Personality Type Profiles for Occupations and Jobs (JPLink-Pred)

We formulate the predicted personality type profile ($\hat{y}_o \in \mathbb{R}^{1 \times 6}$) of an occupation o as a linear projection A of input features followed by a bias offset b and a softmax activation function. Formally, $\hat{y}_o = \text{softmax}(A * f(o)^T + b)$, where $f(o)$ denotes the embedding of o while $A_{6 \times k}$ and $b_{1 \times 6}$ are the trainable linear projection matrix and bias offsets respectively.

As we defined in Sect. 4, the objective of this prediction model is to produce correct ranking of the personality type dimensions for a given o . This motivates us to propose an approach to predict the personality type profiles by learning to rank [1, 2]. We introduce a listwise loss function for learning, which consider the relative ranking information of all RIASEC dimensions together. We adopt ListNet [2], a listwise learning to rank method, to compute the target top one probability of dimension d for a given occupation o as,

$$P_d(y_o) = \frac{\exp y_o^d}{\sum_{d' \in \{R, I, A, S, E, C\}} \exp y_o^{d'}} \quad (1)$$

$\hat{y}_o$ is used as the corresponding modeled top one probability. Then, JPLINK-PRED optimizes the following loss function using SGD,

$$L(y_o, \hat{y}_o) = \sum_{d \in \{R, I, A, S, E, C\}} P_d(y_o) * \log(\hat{y}_o^d) + (1 - P_d(y_o)) * \log(1 - \hat{y}_o^d) \quad (2)$$

5.3 Construct Representations of New Occupations or Jobs

Once the prediction model is learned, we can apply the model to predict personality type profiles of jobs. Before that, we construct the representations of a new job. Given a job j , we use words appearing in the job title W_j^t and job description W_j^d to generate representation $f(j)$ as shown in Eq. 3. The representation for a new occupation can be constructed in the same way.

$$f(j) = \beta * f(W_j^t) + (1 - \beta) * f(W_j^d) \quad (3)$$

where $f(W_j^t) = \sum_{w \in W_j^t} f(w) / \sum_{w \in W_j^t} 1$ and $f(W_j^d) = \sum_{w \in W_j^d} f(w) / \sum_{w \in W_j^d} 1$ denote the representations constructed for the job (or occupation) title and job (or occupation) description respectively. The parameter β controls the importance given to title and description in the final representation. The optimal value of β is determined using a grid search (see [13] for detailed results).

6 Evaluation

Baselines. We compare JPLINK-EMB with two baselines: (1) PRE-TRAINED EMB, which adopts the pretrained word embeddings⁵ as introduced

⁵ <https://github.com/mmihaltz/word2vec-GoogleNews-vectors>.

in Sect. 5.1; and (2) WIKIPEDIA2VEC, which adopts the model proposed in [15] to jointly learn embedding for words and multi-word entities, considering occupations as entities.

We compare JPLINK-PRED with two well known baselines: (1) POINT, which adopts Logistic Regression classifier with binary cross entropy loss. POINT ignores the relative ranking of the RIASEC dimensions; and (2) PAIR, which adopts the pairwise learning to rank loss function proposed in RANKNET [1]. RANKNET considers the relative ranking information between pairs of output dimensions.

Measuring Performance. As claimed by [3], occupations should be characterized by a variable size set of interest dimensions. Hence the prediction engine should be capable to produce the correct ranking of RIASEC dimension for a given occupation o or job j . We thus measure the performance of RIASEC dimension prediction using Normalized Discounted Cumulative Gain (NDCG), which is commonly used to evaluate the performance of rankers in a multi-graded relevance setting. To calculate NDCG, each RIASEC dimension ($d \in \{R, I, A, S, E, C\}$) is assigned a relevance score $R_d(o)$ using the ground truth score (y_o^d) taken from O*NET for o as $R_d(o) = \lfloor y_o^d/20 \rfloor$, where $\lfloor \cdot \rfloor$ denotes the conventional floor operator. Suppose a method ranks the interest dimensions of an occupation o as $\hat{\sigma}_o$ such that $\hat{\sigma}_o(l)$ ($1 \leq l \leq 6$) returns the l^{th} ranked interest dimension. We then define NDCG to be:

$$NDCG(\sigma_o) = \frac{DCG(\sigma_o)}{\max_{\sigma'_o} DCG(\sigma'_o)} \quad DCG(\sigma_o) = \sum_{l=1}^6 \frac{2^{R_{\sigma_o(l)}(o)} - 1}{\log_2(1 + l)} \quad (4)$$

NDCG for job j can be similarly defined by replacing o by j , and σ_o by σ_j .

6.1 Evaluation of the Job-Domain Specific Embeddings

In this section, we compare different embedding techniques (i.e., JPLINK-EMB, PRETRAINED EMB and WIKIPEDIA2VEC). The size of the embedding is set to 300 for all three methods. We train both JPLINK-EMB and WIKIPEDIA2VEC using O*NET, which are initialized using PRETRAINED EMB.

Qualitative Evaluation. Our first evaluation is qualitative: we randomly select a few target occupations and manually inspect the 5 most similar occupations and words to the target occupations (by cosine similarity) using the three occupation/word embedding representations. Due to the space limitations, we direct readers to [13] for the detailed results of the experiment.

We first discuss the most similar occupations. For almost all the target occupations, PRETRAINED EMB returns similar occupations which have overlapping words with the target occupations. For example, the target occupation “financial managers” is similar to several other occupations of “manager” role. The target occupation “technical writers” is similar to occupations containing “technical” and/or “writers” in their titles. On the other hand, JPLINK-EMB

Table 2. RAISEC profile prediction result - NDCG@6

	Prediction methods		
Embedding methods	POINT	PAIR	JPLINK-PRED
PRETRAINED EMB	0.905	0.912	0.928
WIKIPEDIA2VEC	0.921	0.924	0.941
JPLINK-EMB	0.928	0.934	0.949

Table 3. Correlations among the dimensions of the linear transformation matrix and the bias values in JPLINK

	Correlations						Bias
	C	E	I	S	R	A	
Conventional (C)	1.00	0.09	-0.11	-0.29	0.08	-0.52	0.60
Enterprising (E)		1.00	-0.55	0.12	-0.31	-0.23	0.12
Investigative (I)			1.00	-0.40	0.10	-0.07	0.15
Social (S)				1.00	-0.52	-0.01	-0.63
Realistic (R)					1.00	-0.34	0.65
Artistic (A)						1.00	-0.89

produces similar occupations that may not have overlapping words. For example, JPLINK-EMB shows “financial managers” and “loan officers” are similar, “health educators”, “rehabilitation counselors” and “recreational therapists” are similar, which are reasonable with respect to their domains. In contrast, WIKIPEDIA2VEC produces some of the similar occupations that are less obvious, e.g., an occupation similar to “technical writers” is “broadcast news analysts”. Also, when we compute the proportion of the 20 most similar occupations returned by each embedding scheme appearing as similar occupations of the target occupations in O*NET. We found that only 20% of similar occupations by PRETRAINED EMB appears in O*NET for the target occupations. In contrast, JPLINK-EMB and WIKIPEDIA2VEC show 55% and 40% of their 20 most similar occupations appearing in O*NET respectively. This shows the importance of incorporating O*NET knowledge base into the learning of occupation representations.

We next analyse the most similar words to the target occupations. PRETRAINED EMB returns either target occupation title words or their synonyms. E.g., “writers” and “technical” are the most similar words for technical writers, while “investors” and “bankers” are for financial managers. JPLINK-EMB and WIKIPEDIA2VEC, on the other hand, returns domain-specific words associated with the occupational tasks. E.g., “lifestyle”, “smoking”, “vaccines” and “nutritionally” are words relevant to the tasks of the target occupation health educators. These words are useful to understand the requirement of health educator occupation. Because usually words overlapped with the title and its synonyms

do not appear in job descriptions. Instead of that domain-specific information (tasks and skills) related to the jobs are appeared in the description.

Quantitative Evaluation. To quantitatively evaluate our end-to-end framework, we report the results for the task of predicting the RIASEC labels for occupations. In this experiment, we trained the prediction models (POINT, PAIR and JPLINK-PRED) using 50% of 974 O*NET occupations and then the model is evaluated using the predicted RIASEC labels for the rest 50%, using different occupation representations. As shown in Table 2, JPLINK-EMB outperforms PRETRAINED EMB and WIKIPEDIA2VEC for all the three prediction methods. WIKIPEDIA2VEC is slightly better than PRETRAINED EMB in this prediction task. These results are consistent with our qualitative analysis results.

Among the three prediction methods, POINT shows the worst performance as it does not consider the relative ranking of RIASEC dimensions. And JPLINK-PRED consistently outperforms PAIR which only considers the relative ranking of pairs of RIASEC dimensions. This observation shows the importance of having ranking information of RIASEC dimensions in the training phase to capture significant relationships among RIASEC dimensions in Table 1. To illustrate this fact further, we further analyze the learned parameters of the JPLINK-PRED model. As depicted in Table 3, in JPLINK-PRED, each row in the A projection matrix corresponds to the linear mapping learned for a RIASEC dimension. As shown in Table 3, the correlations among the rows of A reflect the actual correlations between the corresponding RIASEC dimensions. These correlations are consistent with those seen in Table 1. Moreover, we also found the bias values (b) of our model are consistent with the actual distribution of RIASEC dimensions in O*NET. Hence, we can say that JPLINK is capable of capturing the inherent patterns of the interest dimensions.

Table 4. A few job posts predicted with “wrong” RIASEC profile by JPLINK

Assigned SSOC Occupation	Job Title	Key words in Job Description	Actual (y) and Predicted ($\hat{y}$) RIASEC Ranking (first has the highest rank)
Graphic and Multimedia Designers and Artists	lead full stack developer	widgets saas javascript html css community	$y : \{A, R, E, I, C, S\}$ $\hat{y} : \{C, I, R, E, S, A\}$
Graphic and Multimedia Designers and Artists	singapore researcher	subject matter experts project managers executive leadership travel presentations oral independent	$y : \{A, R, E, I, C, S\}$ $\hat{y} : \{I, C, R, E, S, A\}$
Manufacturing Labourers and Related Workers	on executive	online market place inventory management customer service stocks filing enquires commerce	$y : \{R, C, E, I, A, S\}$ $\hat{y} : \{C, E, I, S, R, A\}$

6.2 Prediction of Personality Type Profiles for Job Posts

Here, we analyze the predictive power of our model trained using occupation data to predict the personality type profiles of 171,946 job posts with weak RIASEC profiles (see Sect. 3). To derive the representation of these job posts, we use the approach introduced in Eq. 3. We observe that JPLINK-PRED model with JPLINK-EMB embeddings consistently gives the best performance for all different β settings. The β values in the range of $[0.6, 0.7]$ give the best performance for JPLINK (see [13] for detailed results). This means that job title should be given more weight (60%) when constructing the representation of job posts, which is consistent with the results in [16].

Error Analysis. In the above results, we use the weak labels derived by a distant supervision approach for the evaluation. There might be some imperfections in these labels. In addition, we use the SSOC occupation label assigned to job posts from Singapore Jobsbank in the weakly labeling approach (see Sect. 3), which might not be perfectly accurate too. To identify whether these imperfections adversely affect our evaluation, we manually inspect a few job posts that are wrongly predicted from our model. Table 4 lists a few cases where JPLINK wrongly predicts labels. In these example cases, job titles and descriptions do not tally with the corresponding occupations (assigned by the current system). For instance, the job title *lead full stack developer* is assigned to the occupation category, “Graphics and Multimedia Designers and Artists”, instead of an occupation related to software development (e.g., Software Developer). JPLINK actually predicts a personality type profile appropriate to the occupations like Software Developer ($\{I, R, C, E, A, S\}$). Similarly, the predicted interest profiles for other example cases are reasonable with respect to their descriptions and titles. These results further signify the potential of this work to replace the current labor-intensive manual approach to profile job posts.

7 Conclusion

In this paper, we proposed JPLINK, a framework to automate the profiling of jobs with their interest profiles. JPLINK: (1) explored the domain-specific knowledge available in O*NET to improve existing word and occupation representations, which are subsequently used as input features to assign corresponding interest profiles; and (2) proposed a novel loss function for the prediction of RIASEC profiles, which captures the interrelationship between RIASEC dimensions. Finally, we profiled a set of job posts using JPLINK and showed that our model managed to identify a type of imperfection, existed in the current profiling system.

There might be other imperfections in the current system (e.g., imperfections in mapping between different occupation taxonomies), which could be identified via a deep analysis of our predictions with the knowledge of domain experts. Also, our model can be used to understand individuals’ behaviors using the job posts that they prefer to apply. For instance, if individuals tend to apply for jobs with similar interest profiles, our model might be used to infer individuals’

interest profile, which is currently collected via surveys. Likewise, we believe that our effort will open the door for many promising research directions.

References

1. Burges, C., et al.: Learning to rank using gradient descent. In: Proceedings of ICML (2005)
2. Cao, Z., Qin, T., Liu, T.Y., Tsai, M.F., Li, H.: Learning to rank: from pairwise approach to listwise approach. In: Proceedings of ICML (2007)
3. Gati, I.: Using career-related aspects to elicit preferences and characterize occupations for a better person-environment fit. *J. Vocat. Behav.* **52**(3), 343–356 (1998)
4. Grover, A., Leskovec, J.: node2vec: scalable feature learning for networks. In: Proceedings of KDD (2016)
5. Holland, J.L.: Making Vocational Choices: A Theory of Vocational Personalities and Work Environments. Psychological Assessment Resources, Lutz (1997)
6. Mikolov, T., Sutskever, I., Chen, K., Corrado, G.S., Dean, J.: Distributed representations of words and phrases and their compositionality. In: Proceedings of NIPS (2013)
7. Mislove, A., Viswanath, B., Gummadi, K.P., Druschel, P.: You are who you know: inferring user profiles in online social networks. In: Proceedings of WSDM (2010)
8. Perozzi, B., Al-Rfou, R., Skiena, S.: Deepwalk: online learning of social representations. In: Proceedings of KDD (2014)
9. Prediger, D.J.: Mapping occupations and interests: a graphic aid for vocational guidance and research. *Vocat. Guidance Q.* **30**(1), 21–36 (1981)
10. Prediger, D.J.: Dimensions underlying Holland’s hexagon: missing link between interests and occupations? *J. Vocat. Behav.* **21**(3), 259–287 (1982)
11. Prediger, D.J., Vansickle, T.R.: Locating occupations on Holland’s hexagon: beyond RIASEC. *J. Vocat. Behav.* **40**(2), 111–128 (1992)
12. Rounds, J., Smith, T., Hubert, L., Lewis, P., Rivkin, D.: Development of occupational interest profiles for O*NET. Raleigh, NC: National Center for O*NET Development (1999)
13. Silva, A., Pei-Chi, L., Ee-Peng, L.: JPLink: on linking jobs to vocational interest types. In: [arXiv:2002.02557](https://arxiv.org/abs/2002.02557) (2020)
14. Tang, J., Qu, M., Wang, M., Zhang, M., Yan, J., Mei, Q.: Line: large-scale information network embedding. In: Proceedings of WWW (2015)
15. Yamada, I., Shindo, H., Takeda, H., Takefuji, Y.: Joint learning of the embedding of words and entities for named entity disambiguation. In: Proceedings of SIGNLL (2016)
16. Zhu, Y., Javed, F., Ozturk, O.: Document embedding strategies for job title classification. In: Proceedings of FLAIRS (2017)