This is an electronic reprint of the original article.
This reprint may differ from the original in pagination and typographic detail.

Sayfullina, Luiza; Malmi, Eric; Liao, Yiping; Jung, Alexander

Domain Adaptation for Resume Classification Using Convolutional Neural Networks

Published in:

Analysis of Images, Social Networks and Texts: 6th International Conference, AIST 2017, Moscow, Russia, July 27--29, 2017, Revised Selected Papers

DOI:

[10.1007/978-3-319-73013-4_8](https://doi.org/10.1007/978-3-319-73013-4_8)

Published: 01/01/2018

Document Version

Peer reviewed version

Please cite the original version:

Sayfullina, L., Malmi, E., Liao, Y., & Jung, A. (2018). Domain Adaptation for Resume Classification Using Convolutional Neural Networks. In W. M. P. van der Aalst, D. I. Ignatov, M. Khachay, S. O. Kuznetsov, V. Lempitsky, I. A. Lomazova, N. Loukachevitch, A. Napoli, A. Panchenko, P. M. Pardalos, A. V. Savchenko, & S. Wasserman (Eds.), *Analysis of Images, Social Networks and Texts: 6th International Conference, AIST 2017, Moscow, Russia, July 27--29, 2017, Revised Selected Papers* (pp. 82-93). (Lecture Notes in Computer Science; Vol. 10716). SPRINGER. https://doi.org/10.1007/978-3-319-73013-4_8

This material is protected by copyright and other intellectual property rights, and duplication or sale of all or part of any of the repository collections is not permitted, except that material may be duplicated by you for your research use or educational purposes in electronic or print form. You must obtain permission for any other use. Electronic or print copies may not be offered, whether for sale or otherwise to anyone who is not an authorised user.

Domain Adaptation for Resume Classification Using Convolutional Neural Networks

Luiza Sayfullina, Eric Malmi, Yiping Liao and Alexander Jung

Department of Computer Science, Aalto University, Finland

Abstract. We propose a novel method for classifying resume data of job applicants into 27 different job categories using convolutional neural networks. Since resume data is costly and hard to obtain due to its sensitive nature, we use domain adaptation. In particular, we train a classifier on a large number of freely available job description snippets and then use it to classify resume data. We empirically verify a reasonable classification performance of our approach despite having only a small amount of labeled resume data available.

1 Introduction

The fast paced online job-market industry requires recruiters to screen through vast amounts of resume data in order to evaluate applicants fast and reliable. The design of accurate automatic classification systems typically requires the availability of labeled resume data which can be used to train the classifier.

Due to its sensitivity, resume data is difficult and costly to obtain. In contrast, data about job descriptions can be obtained much easier. However, the two domains constituted by resume data and job description data are intrinsically related. Indeed, both data domains are related to the same job recommendation task, which is to match applicant to suitable job offers. Moreover, the resumes of applications have semantic similarities with job descriptions which belong to the same job category. For instance, they both can contain skills, education, duties as well as personal characteristics of the desired candidate.

So far, there are two main flavors along with their hybrids [9], of job recommendation systems. One class, referred to as content-based recommendation systems, is based mainly on the available job descriptions. A second class, referred to as collaborative filtering recommendation systems, is mainly based on the preferences of users who are interested in similar jobs. Content-Based recommendation system suggests to a user textually similar jobs to what he/she viewed or liked previously [2].

It seems therefore reasonable to use transfer learning in order to implement a domain adaptation in order to leverage the information contained in vast amounts of labeled job description data in order to classify resume data. Since resume and job summaries belong to similar domains, we expect features extracted by a convolutional neural network for job classification to be highly relevant for resume summaries as well.

The theory of learning in different domains was theoretically approached in [3,4]. The authors provided generalization bound for domain adaptation using $\mathcal{H}$ -divergence. It consists of two components and tries to find a trade-off between source-target similarity and source training error. Based on that assumption several researchers [1,6,7] came up with the domain-adversarial approach, where high-level representations from neural network are optimized to minimize the loss on the source domain and maximize the loss on the domain classifier. [13] proposed another approach based on convolutional neural network special architecture. First three layers are domain-invariant, next two layers are fined-tuned and fully connected layers aim to fit specific tasks, but regularized by multiple kernel variant of maximum mean discrepancy that enforces distributions to be similar. The proposed network is optimized for the image domain, being more specific.

Convolutional neural networks (CNN) have been successfully applied to not only image, but also text classification [12,16], provided that enough training data is available. We propose a domain adaptation approach [5,8] where we train a CNN based classifier on 85,000 job description snippets which are labeled using 27 industrial job classes. After the classifier has been trained, we apply it to classify unlabeled resume data.

The paper is organized as follows. First, in Section 2 we describe job, resume and children dream job datasets, used for classification. Then in Section 3, we describe the **fastText** baseline model and the CNN for short text classification model. Experimental results are provided in Section 4, where classification accuracies are reported along with t-SNE visualization built on latent CNN representations. Finally, we present conclusions in Section 5.

2 Datasets

We study three different datasets: *job descriptions*, which are used for training models, *resume summaries*, which are our main target domain used for testing the models. *Children’s dream job* descriptions is rather a toy data lacking enough samples for fair evaluation, but these job descriptions significantly differ and thus are interesting to experiment with.

2.1 Job Descriptions

We collected 90,000 job description snippets using the Indeed Job Search API¹, that enables access to short job summaries given a key word. As key words, we used 27 different industrial job categories listed in Table 1.

Here is an example of a job summary from the category *Accountant*: “Entering journal entries, posting cash, and account reconciliations/supporting schedules. This position is responsible for supporting the daily operations and ...”.

Note that the snippets provided by Indeed are generated based on the full description of the job postings, thereby they encapsulate only the condensed

¹ <https://www.indeed.com/publisher>

information regarding the job. Furthermore, since the descriptions are unstructured text snippets, the contents provided by different companies for similar positions may be inconsistent. For example, some job snippets or summaries do not include informative sentences or keywords related to job titles or categories. However, since the descriptions are not limited by a predefined structure, they may provide richer and more detailed information about the jobs in varying industries.

Table 1. 27 industrial job categories from <https://www.indeed.com/find-jobs.jsp>.

1. Accounting/Finance	10. Banking/Loans	19. Education/Training
2. Healthcare	11. Human Resources	20. Legal
3. Non-Profit/Volunteering	12. Restaurant/Food service	21. Telecommunications
4. Administrative	13. Construction/Facilities	22. Engineering/Architecture
5. Computer/Internet	14. Insurance	23. Manufacturing/Mechanical
6. Pharmaceutical/Bio-tech	15. Retail	24. Transportation/Logistics
7. Arts/Entertainment/Publishing	16. Customer Service	25. Government/Military
8. Hospitality/Travel	17. Law Enforcement/Security	26. Marketing/Advertising/PR
9. Real Estate	18. Sales	27. Upper Management/Consulting

2.2 Resume Summaries

We collected 523 anonymous resume data samples, each sample labeled with one of the 27 categories based on the type of a job the candidate is looking for. The distribution of the categories is shown in Figure 1.

Here is an example of a resume self-description summary:

“Experienced analyst with an excellent academic profile and having several years of invaluable experience in domestic and international consultancy and management. Highly focused with a comprehensive knowledge and understanding of project management, technical issues and financial practices. Good at meeting the deadlines. Consider myself to be sociable person and good team worker.”

2.3 Children’s Dream Jobs

Children, unlike grown-ups, can express their dream jobs more emotionally, without being attached to skills, but rather following their interests. So in addition to resumes, we decided to use children dream job descriptions that were categorized manually into the same 27 job categories. The data set contains 98 children’s short essays on their dream job parsed from ². Below is an example essay:

² http://www.valleymorningstar.com/sie/what_you_think/article_692e1ac9-bae5-5705-8005-c22dac04ebf6.html

“As far as I can remember I have always wanted to become a medical doctor. More specifically, a cardiologist. I love the thought of saving a person’s life. The road to becoming a doctor is a long process, but worth it in the end. Having the feeling of accomplishment and knowing that I have made an impact on a family’s life, would be the greatest satisfaction for me.”

Fig. 1. The comparison of class distribution in job description and resume datasets.

2.4 Comparison of Job Descriptions and Resumes

Since our aim is to leverage easily available job description data to train a model for classifying resume summary snippets, it is important to understand how these two domains differ from each other. In order to compare the two, we study word frequencies to see whether certain terms are over-represented in one domain compared to the other.

Figure 2 shows the normalized frequencies of all words appearing at least five times in both datasets. The two frequencies are correlated ($\rho = 0.59$), but we can see, that for some words, the frequencies differ considerably. In Table 2, we list the words for which the relative difference is the largest.³ The results

³ The relative difference is measured by dividing the two normalized frequencies. For low frequencies, this measure will be noisy but we ignore this since the purpose of the experiment is to merely gain an overview of the differences between the datasets.

Fig. 2. Normalized frequencies of all the words appearing at least five times in both datasets. Each word corresponds to a dot whose x and y coordinates denote the frequencies among job descriptions and CVs (resumes), respectively.

show that in resumes people are much more likely to use adjectives describing themselves, such as *adaptable* and *polite*, whereas job descriptions mention more often roles, such as *director* and *coordinator*.

3 Industrial Category Classification Methods

The objective of industrial category classification is to classify user profiles, represented as text snippets, into 27 industrial categories shown in Table 1. We apply a CNN based methods to this task because they have shown state-of-art performance in text classification [11]. As a baseline method, we employ the **fastText** classifier [10] which is presented next.

3.1 Fast text Classifier

The **fastText** method has been proposed recently by Joulin et al. [10] to efficiently classify text data. The method is based on learning word embeddings, averages them, and feeds the resulting vector into a linear classifier. The method also supports learning word embeddings for n-grams which allows capturing word order information.

Supported by a few algorithmic and implementation improvements, **fastText** is able to train and test extremely fast without access to GPUs. We have chosen **fastText**, since it was shown by Joulin et al. [10] to be a competitive baseline for deep learning models, outperforming models like CNN, char-CNN and slightly (1%) underperforming LSTN-GRNN models.

Table 2. Words which normalized frequency differs the most between job descriptions (f_{Job}) and resume summaries (f_{CV}). Difference is measured by dividing the frequencies.

$f_{\text{CV}}/f_{\text{Job}}$ Word			$f_{\text{Job}}/f_{\text{CV}}$ Word		
1.	284.9	uk	15.3		program
2.	242.2	gained	10.4		assist
3.	239.3	adaptable	9.0		director
4.	95.0	polite	8.4		medical
5.	82.1	keen	6.7		provides
6.	76.0	bsc	6.7		coordinator
7.	73.3	trustworthy	6.6		accounting
8.	73.3	ambition	6.5		executive
9.	63.3	licence	6.2		representative
10.	59.6	confident	5.9		assistant
11.	59.5	adapt	5.8		report
12.	57.0	versatile	5.7		food
13.	57.0	consultancy	5.7		perform
14.	52.9	approachable	5.7		equipment
15.	48.8	punctuality	5.6		related

3.2 Convolutional Neural Networks for Sentence Classification

Word2vec model [15] is a widely used method for learning vector representations of words, so that semantically similar words are close to each other in the vector space. Based on the word vectors, contextual information can be extracted to learn the semantic similarity between words and sentences.

Convolutional neural networks trained on the top of pre-trained word2vec representations proposed by [11] showed state-of-the-art performance on several datasets, including sentiment analysis. In this model, words in the sentences are embedded word2vec representation of the same length. Then vectors of words are concatenated by rows thus forming a matrix, to which CNN is applied.

Let us introduce some notations. First, $\mathbf{x}_i \in \mathbb{R}^k$ is the i -th word embedded into a vector of k dimensions. A sentence which consists of n words is represented as $\mathbf{x}_{1:n} = \mathbf{x}_1 \oplus \mathbf{x}_2 \cdots \oplus \mathbf{x}_n$, where $\oplus$ is the concatenation operator. The convolutional filter with window size of h words is denoted as $\mathbf{w} \in \mathbb{R}^{hk}$. Then the new feature map is generated by:

$$c_i = f(\mathbf{w} \cdot x_{i:i+h-1} + b), \quad (1)$$

where $b \in \mathbb{R}$ is a bias term and f is a hyperbolic tangent. Following the convolution operation, the max-pooling operation is applied to capture the most important feature and the output is forwarded to a fully connected soft-max layer whose output is a probability distribution over classes. The regularization of the network is done by applying dropout to prevent co-adaptation and re-scaling weights to prevent large (and possibly noisy) gradient updates during training. At the testing phase, the learned weight vectors are scaled by $\mathbf{w} \leftarrow p\mathbf{w}$.

Additionally, an L2-norm constraint is applied to rescale $\mathbf{w}$ to have $\|\mathbf{w}\|_2 = s$ when $\|\mathbf{w}\|_2 > s$ after gradient decent step.

4 Experimental results

We trained our models by fixing training (80,000 samples) and validation data (5,000 samples) consisting of job summaries and used all available samples from resume and children’s dream job data for testing. We also used 5,000 job summary samples for testing a classifier on purely job data. All selected data samples were trimmed to 100 words.

Our CNN model was based on the implementation by Kim ⁴. In order to avoid strong overfitting, we increased the L2-norm constant up to 10 and set the width of CNN filters to be [2,3,4] instead of [2,3,4,5]. We tried the filters of size [50, 100, 200] per each type and have chosen 50 based on validation set from job description data. The dropout rate was set to 0.5 and we found it useful for model regularization. A non-static setting of CNN was chosen, where Google-News pretrained word vectors are fine-tuned while training.

For the **fastText** model, we optimized the lengths of the n-grams and the learning rate hyperparameters using the validation data, obtaining the values 4 and 0.25, respectively. These values were kept fixed for all three test datasets.

The overall prediction accuracies are shown in Table 3. When moving from the source domain to the target domain, the accuracy drops from 74.88% to 40.15%. CNN outperforms **fastText** for each dataset and particularly for resume and dream job data, which shows that the CNN model generalizes better to new domains.

Table 3. Job category prediction accuracies (%) for the fastText method and CNN for short text classification.

Dataset	fastText	CNN
Job description	71.99	74.88
Resume	33.40	40.15
Children’s dream job	28.5	51.02

The confusion matrices for job description and resume summary classification are shown in Figure 3.

From Figure 3 we can see that the hardest categories to classify in job description dataset are Management, Administrative, Sales, Customer Service and Manufacturing. Probably, it happens due to semantic closeness of some job categories, like Management and Administrative, Sales and Retail, since even humans can have trouble clearly distinguishing between them. Manufacturing category samples were classified with Construction, Engineering and Transporta-

⁴ https://github.com/yoonkim/CNN_sentence

Fig. 3. Confusion matrix of resume (top) and job (bottom) classification results. In both datasets Management, Administrative, Customer Service, Retail and Manufacturing categories have a low recall. We assume that this happens due to the semantic closeness of these categories, since even a human can not always correctly make a clear distinction between them.

Fig. 4. t-SNE visualization using the CNN first layer outputs on job and resume data. We used all job (5,000) and resume (523) test data for fitting t-SNE. By visualizing vectors on 2D space, we check how useful are the representations learned by CNN to distinguish between the classes in familiar and new domains. There are 5000 test samples from job data, marked with circles, and 523 samples from resume data, marked with crosses, in total. We can observe the presence of category clusters formed by job samples, however they are not perfectly separable. Since the resume data has differences in underlying distribution, some resume clusters are neighbours with corresponding job clusters, e.g from Non-Profit, Computer/Internet, Arts, Retail and Engineering categories. In fact, resume classes from these classes form neighbouring clusters or intersect with corresponding job clusters.

tion/Logistics labels. The highest recall belongs to Legal, Real Estate, Arts, Law and Non-Profit categories.

Resume dataset has a small number of samples per class, so we can not make general conclusions from confusion matrix showed in Figure 3. Still, the results on our resume data have common trends with job data. For example, similarly to job confusion matrix, Management, Administrative, Customer Service, Retail and Manufacturing categories have a low recall. Legal, Government, Arts, Healthcare and Pharmacy show the highest recall. Management category, consisting of 9 samples, was not detected at all, probably since this position can be quite general and related to various job fields.

One of the ways to achieve better generalization is building latent representations. In our case, the concatenated outputs c_i of the first layer of the CNN model form latent space representations. Therefore, we visualized those outputs both for job and resume test data using t-SNE [14] projection and show the results in Figure 4.

One can observe the presence of category clusters formed by job samples, although some of them are not perfectly separable. However, for 27 classes this is relatively good separation. If resume samples, represented by crosses, were semantically close to job descriptions, they could belong to the same job clusters. However, since the resume data has differences in the underlying distribution, some of its clusters are at least neighbours with corresponding job clusters, e.g., for Non-Profit, Computer/Internet, Arts, Retail and Engineering categories. We can not make any general conclusions about resume clusters due to the lack of data, but we can clearly find clusters for some categories, that are sometimes distant from corresponding job clusters. This suggests that the learned CNN representations are useful for resume classification as well, since clusters can be found using them.

5 Conclusion

We have devised a resume classification method which is able to exploit the information contained in vast amounts labeled job description data in order to achieve higher accuracy. Since resumes are more sensitive data and difficult to obtain, compared to job summaries, we trained the proposed model only on job summaries and tested its performance on resume data with the same job category labels. A convolutional neural network for short text classification using word embeddings was trained and validated on 85,000 short job summaries mined from Indeed. Then this network was used to classify a set of 523 candidate resumes and compared with a simple but effective **fastText** model. Our method achieved 74.88% accuracy on job classification task and 40.15 % on resume classification, thereby outperforming the existing **fastText** model by more than 6% on resume classification task and 3% on the job description task. Moreover, we applied our method to a small imbalanced dataset consisting of 98 children dream job descriptions. In this task CNN outperformed **fastText** by 22%.

Given the fact that no labels were used from resume data for training or validation, we consider CNN for short classification to be useful in a domain adaptation scenario. An interesting direction for future work would be to study whether the results can be improved by leveraging a small number of labeled resume samples to fine-tune the CNN model.

References

1. Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., Marchand, M.: Domain-adversarial neural networks. arXiv preprint arXiv:1412.4446 (2014)
2. Al-Otaibi, S.T., Ykhlef, M.: A survey of job recommender systems. *International Journal of Physical Sciences* 7(29), 5127–5142 (2012)
3. Ben-David, S., Blitzer, J., Crammer, K., Kulesza, A., Pereira, F., Vaughan, J.W.: A theory of learning from different domains. *Machine learning* 79(1), 151–175 (2010)
4. Ben-David, S., Blitzer, J., Crammer, K., Pereira, F., et al.: Analysis of representations for domain adaptation. *Advances in neural information processing systems* 19, 137 (2007)
5. Daume III, H., Marcu, D.: Domain adaptation for statistical classifiers. *Journal of Artificial Intelligence Research* 26, 101–126 (2006)
6. Ganin, Y., Lempitsky, V.: Unsupervised domain adaptation by backpropagation pp. 1180–1189 (2015)
7. Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., Marchand, M., Lempitsky, V.: Domain-adversarial training of neural networks. *Journal of Machine Learning Research* 17(59), 1–35 (2016)
8. Glorot, X., Bordes, A., Bengio, Y.: Domain adaptation for large-scale sentiment classification: A deep learning approach. In: *International Conference on Machine Learning*. pp. 513–520 (2011)
9. Hong, W., Zheng, S., Wang, H., Shi, J.: A job recommender system based on user clustering. *Journal of Computers* 8(8), 1960–1967 (2013)
10. Joulin, A., Grave, E., Bojanowski, P., Mikolov, T.: Bag of tricks for efficient text classification. arXiv preprint arXiv:1607.01759 (2016)
11. Kim, Y.: Convolutional neural networks for sentence classification pp. 1746–1751 (2014)
12. Kim, Y., Jernite, Y., Sontag, D., Rush, A.M.: Character-aware neural language models. In: *Thirtieth AAAI Conference on Artificial Intelligence*. pp. 2741–2749 (2016)
13. Long, M., Cao, Y., Wang, J., Jordan, M.: Learning transferable features with deep adaptation networks. In: *International Conference on Machine Learning*. pp. 97–105 (2015)
14. Maaten, L.v.d., Hinton, G.: Visualizing data using t-sne. *Journal of Machine Learning Research* 9(Nov), 2579–2605 (2008)
15. Mikolov, T., Sutskever, I., Chen, K., Corrado, G.S., Dean, J.: Distributed representations of words and phrases and their compositionality. In: *Advances in neural information processing systems*. pp. 3111–3119 (2013)
16. Zhang, X., Zhao, J., LeCun, Y.: Character-level convolutional networks for text classification. In: *Advances in neural information processing systems*. pp. 649–657 (2015)