5-2018

Industry Classification Based on Labor Mobility Network Mining

Xiaoying Xu
South China University of Technology, bmxyxu@scut.edu.cn

Hanlin Qian
South China University of Technology, 201620123470@mail.scut.edu.cn

Chunmian Ge
South China University of Technology, bmgecm@scut.edu.cn

INDUSTRY CLASSIFICATION BASED ON LABOR MOBILITY NETWORK MINING

Xiaoying Xu
South China University of Technology
bmxyxu@scut.edu.cn

Hanlin Qian
South China University of Technology
201620123470@mail.scut.edu.cn

Chunmian Ge
South China University of Technology
bmgecm@scut.edu.cn

Abstract:

Industry classification is important for industry analysis and competitive intelligence. However, existing schemes and methods are limited by the small number of industry categories and the lagged information of firms' business. In this paper, we propose a novel industry classification method by constructing the labor mobility network from the LinkedIn profiles. We also propose a hierarchical extension of the community detection algorithm to better discover latent industry clusters on the constructed network. The evaluation conducted on real-world datasets shows that our method outperforms the best existing industry classification scheme and the state-of-the-art method and improves their explanatory power by 8.31% and 3.97% respectively. Moreover, our method is effective in earlier revealing firms' action of entering new industries.

Keywords:

Industry classification; Labor mobility network; LinkedIn; Community detection

1. Introduction

In the fields of industry analysis and competitive intelligence, industry classification of firms is a fatal step before investigating the target firm's competitive environment and growth opportunities(Fang, Dutta, & Datta, 2013). For example, when Google decided to enter the field of self-driving, it might want to know its potential competitors in this industry. To ensure the effectiveness of related analyses, a precise and timely industry classification approach is urgently needed by both researchers and practitioners in the market, as well as other stakeholders such as asset managers, investors and credit analysts.

There are some existing industry classification schemes, such as the Standard Industrial Classification (SIC) and the North American Industry Classification System (NAICS). However, the number of industry categories defined by these schemes are too limited to reveal the detailed business of firms. For example, using the existing schemes, it is difficult to find companies attributed to self-driving, which is a new and small industry. In addition, these existing schemes are not frequently updated so that they usually lag behind the firms' quick actions of entering new industries.

Recently, researchers have begun seeking solutions to supplement the industry classification schemes. For example, Hoberg and Phillips (2013) and Fang et al. (2013) performed industry classification by applying text mining to the 10-K forms. The solutions based on annual reports like 10-K forms have shown some advantages, but they still cannot timely reflect the dynamics of the firms' business due to the low update frequency and the lagged information. For example, firms like Google would update their new business on self-driving in the annual reports long after they have entered this field, but not at the early stage.

To target these problems, this paper aims at seeking for other available sources that can be utilized for industry classification. Guided by the theoretical work of Farjoun (1994) that contributes to resource-based view by characterizing industry groups from a human-resource-related perspective, we propose a brand-new industry classification method by constructing and analyzing the labor mobility network. In detail, we extract the turnover records of each individual from their public profiles on LinkedIn[1], a famous professional social network, with which we construct the labor mobility network among different firms. In order to discover the clusters of firms that consist of the latent industry classes, we adapt and extend a community detection algorithm to a hierarchical version, and then apply it on the constructed network. Our evaluation based on real-world datasets shows that the proposed method outperforms other commonly used schemes and state-of-the-art methods in terms of explaining stock return variation, and can reveal firm action of entering new industry. To summarize, this paper has the following contributions:

(1) Our work is among the first to consider the aspect of human resource in automatic industry classification. Theoretically, our work contributes to resource-based view by empirically examining the effectiveness of human-resource-based industry classification.

(2) Our work has revealed research potentials for professional social networks like LinkedIn. Most of the resume information on such social network sites is up to date. By analyzing the turnover records, a firm's action of entering new industry can be earlier detected when large scale of recruitment in related fields is observed, therefore addressing the lagging problem. Our work has shown that LinkedIn is an ideal and reliable source for industry analysis.

(3) We design and implement a brand-new solution for automatic industry classification. Our innovative approach of constructing the labor mobility network has implications for both researchers and practitioners in related fields. In addition, our proposed hierarchical extension of the community detection algorithm has technical innovativeness.

The remainder of this paper is organized as follows: first we review related research. Then we depict the intuition behind our proposed method and provide an overview followed by a more detailed elaboration. Besides, we evaluate the method and present the evaluation results. Lastly, we summarize our work and discuss the possibilities for future work.

[1] https://www.linkedin.com/

2. Related Work

In the following, we will review related work on industry classification, including the latest extensions. In addition, we will discuss related research and applications on professional social network. Methods for community detection will also be introduced, which will be adapted in our approach for industry classification.

2.1. Industry Classification

There are three traditional industry classification schemes: SIC, NAICS and GICS. American government first developed the Standard Industrial Classification (SIC) in 1937 and later replaced it with the North American Industry Classification System (NAICS) in consort with Canadian and Mexican governments. Besides, Standard & Poor's and Morgen Stanley Capital International (2002) established the Global Industry Classification Standard (GICS) that is based on the judgment of financial analysts to determine which firms are financially comparable. Many researchers compare these schemes in terms of reliability and precision, and they find that GICS outperforms the other two in most situations (Bhojraj, Lee Charles, & Oler Derek, 2003; Chan, Lakonishok, & Swaminathan, 2007; Hrazdil, Trottier, & Zhang, 2013; Hrazdil & Zhang, 2012).

Although widely used, the existing schemes are limited by some well-known drawbacks. Firstly, due to the small number of industry categories, they may fail to reflect real industrial structure and therefore lead to misclassification (Dalziel, 2007). Furthermore, the low update frequency would prevent them from capturing the evolvement of firms' business and the change of industry structure (Chong & Zhu, 2012). Lastly, Fang et al. (2013) point out that these schemes assume binary relationship and cannot measure the degree of similarity.

Noticing these limitations, some researchers have shifted their focus to classify industries on different aspects. For example, Dalziel (2007) defines a sector on the basis of similarity in needs to which firms collectively respond, that which can better reflect industry structure. Hoberg and Phillips (2013) make it easier to find peer firms by using the product descriptions in form 10-K to generate a new set of industry categories. Similarly, Fang et al. (2013) propose to extract topic features from business descriptions of form 10-K for industry classification. And Chong and Zhu (2012) introduce a firm clustering method that employs XBRL-based financial information. These new approaches work well by revealing more dimensions of the industry structure. But they still lack timeliness caused by the lagged information from the data they use.

The work of Farjoun (1994) proposes the Resource-Related Industry Groups (Resource-RIGs) identified by the similarities in the requirements for human expertise, which lays the theoretical foundation for our work. To bridge the research gap, our proposed method considers the human resource aspect, and makes use of the data from professional social network that will be introduced below.

2.2. Professional Social Network

Professional social network, such as LinkedIn, Viadeo and Open Science Lab, is a type of social network focusing exclusively on business relationships and interactions (Ahmed Ben, Ahlem, & Faïez, 2013). Many studies have been conducted on LinkedIn profiles. For example, Ge, Huang, and Png (2016) use LinkedIn data to investigate the effect of human capital on mobility of engineers and scientists, and find that LinkedIn provides more accurate career histories than patents through an inventor survey. Antoine, Cécile, and Hilda (2017) demonstrate that LinkedIn allows decomposing firms' value chains and permits to develop interlocking networks dedicated to firms' divisions. Wang, Li, and Zhou (2016) implement personal profile summarization by leveraging both textual information and social connection information in LinkedIn. All the existing studies have highlighted that LinkedIn is a valuable and reliable source for investigating the human-resource-related problems. However, no existing work has been found to utilize the LinkedIn data for industry analysis, which is the focus of our paper.

2.3. Community Detection

Communities, also called clusters or modules, are groups of vertices that probably share common properties and/or play similar roles within the graph (Newman, 2004). The methods for finding communities in a certain network are called community detection. Compared to traditional clustering methods, community detection is more suitable for large-scale and rapid-change situation (Fortunato, 2010), making it appropriate for us to discover firm clusters. Community detection algorithms has been widely used in many areas, such as social communities (Moody & White, 2003), protein-protein interaction networks (Chen & Yuan, 2006) and food webs (Krause, Frank, Mason, Ulanowicz, & Taylor, 2003), and has been proved to have good quality in discovering clusters in networks. However, scant attention has been drawn to applying community detection to industry analysis. Therefore, our method adapts and extends the Louvain algorithm (Blondel, Guillaume, Lambiotte, & Lefebvre, 2008), which has been found to have low computational cost and high quality, for detecting firm clusters.

3. Intuition and Overview

We are interested in helping market practitioners and researchers with a better industry classification approach that is able to reveal more aspects of firms' business and response quickly to market changes. Our proposed method is guided by the theory of resource-based view, and is motivated by a simple fact that, when a firm starts new business, it must recruit people in relevant areas. Hence, from the human resource the firm owns or needs, we may infer its business. Furthermore, people with expertise in certain area would tend to stay among the firms related to the same area. For instance, a person who works in Google is more likely to take an offer of Apple rather than McDonalds. Knowing this nature, it is reasonable to expect that, in a network representing the labor mobility among different firms, those firms focusing on similar business would be well connected with each other due to the high mobility of labors between them. We call such network labor mobility network, in which a node indicates a firm, and a weighted directed edge between two nodes indicates the number of people who leave one firm and join another. Based on the network structure, our method is designed to classify the firms in the labor mobility network into subgroups that consist of the latent industry classes.

Since the industry classes are self-generated rather than a small number of pre-defined categories, our method would be able to cover more aspects of firms' business and industry structure, contributing to more precise industry classification. Moreover, we make use of the human resource aspect of a firm to capture its changes in business, which can quickly response to market changes, addressing the lagging problem and therefore improving the precision of industry classification.

4. Solution Details

In this section, we will first describe the architecture of our proposed solution, followed by the details of each component in the architecture. The system architecture of our approach is depicted in Figure 1. There are three major components in our architecture: Profile Crawler, Network Constructor and Industry Classifier. More details of each component will be discussed in the ensuing sub-sections.

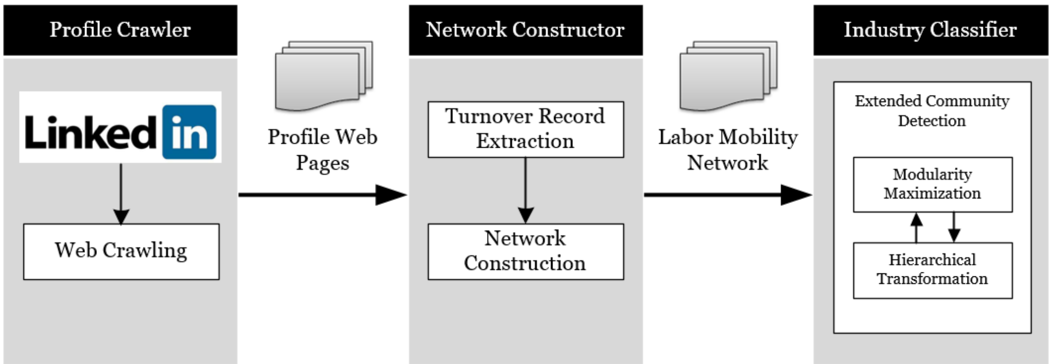

Figure 1. System Architecture

4.1. Profile Crawler

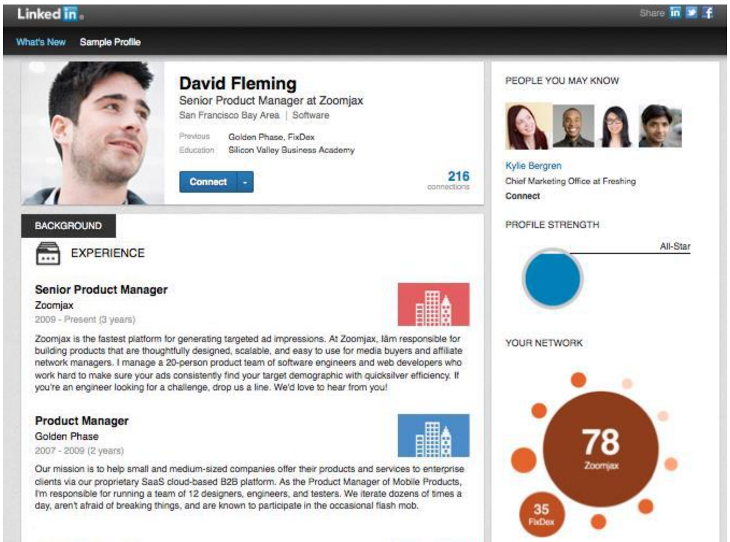

Figure 2. Sample Profile Webpage on LinkedIn

The main task of the Profile Crawler is to collect LinkedIn profiles for further analysis. The crawler integrated with Google API is able to systematically and automatically browse the LinkedIn website and download public profile webpages. Figure 2 shows a sample of LinkedIn profiles with working experience and other personal information. During the process of data collection, all non-job-related content like the advertisements is filtered and excluded. The collected data is then passed to the next component.

4.2. Network Constructor

This component is designed to construct the labor mobility network, an important element in our implementation. There are two main steps to implement this component:

4.2.1. Turnover Record Extraction.

The turnover records are embedded in the working experience section in each LinkedIn profile. They need to be extracted and reorganized. For example, in the profile shown in Figure 2, we can see that David left Golden Phase and joined Zoomjax in 2009, so a turnover record "Golden Phase -> Zoomjax, 2009" can be extracted from this profile.

4.2.2. Network Construction.

With the extracted turnover records at individual level, this step is to aggregate these records at firm level and to construct the labor mobility network that represents the mobility of labors between different firms. Let $G_t = (V, E)$ be a directed graph with a set of vertices V and a set of directed edges E during time period t. A vertex V_i denotes a company i, and a directed edge $E_{i,j}(\rho)$ denotes that there were ρ employees who left company i and joined company j. Figure 3 shows a simple example of a labor mobility network.

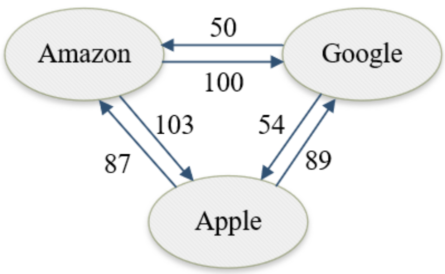

Figure 3. Example of Labor Mobility Network

4.3. Industry Classifier

The industry classifier is a core component to discover the latent industry classes and perform industry classification. To fulfill this task, we extend an algorithm called Louvain (Blondel et al., 2008), which is acknowledged as one of the most excellent and fastest community detection algorithms, under our context. This process consists of two main steps:

4.3.1. Modularity Maximization.

The objective of the Louvain algorithm is maximizing the modularity Q that is defined as:

$$Q = \frac{1}{2M} \sum_{c} \left[I - \frac{T^2}{2M} \right]. \tag{1}$$

where M denotes the total number of labors moving on this network; I denotes the total number of labors moving inside community c that is a specific cluster of firms; and T denotes the total number of labors moving between community c and the outside nodes. Generally, the original algorithm is implemented in the following steps. Initially, we consider every firm as an independent community. Then we assign firm I to each other community, and calculate the changes of modularity in each trial. Firm I is allocated to the target community k where the largest modularity gain is obtained. We iterate this community assignment process for other firms until the algorithm converges, which means there is no significant change in modularity. The details of the original algorithms can be found in the work of Blondel et al. (2008).

4.3.2. Hierarchical Transformation.

The original community detection algorithm tends to discover small communities, resulting in too many firm clusters that reduce the usefulness of the results. For example, in a network with 100 thousand firms, the Louvain algorithm may detect about 10 thousand communities. In order to get desirable results with appropriate scale, we propose a hierarchical extension to the original Louvain algorithm. The extended algorithm combines the nodes in the same community given by the previous step into new nodes, and uses the combined nodes as the input to run the Modularity Maximization process. We iterate this loop until the number of communities is scalable. With the extended algorithm, as is shown in Figure 4, we are able to obtain the hierarchical results of community detection, which contributes to improving the usefulness of the original algorithm by allowing industry classification at different scales.

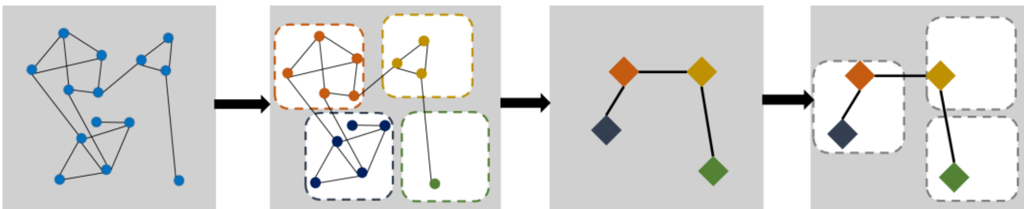

Figure 4. Hierarchical Community Detection

5. Evaluation

In this section, we report the results of the experiment we conduct to evaluate the effectiveness of our proposed method. First, we describe the dataset we use for the evaluation. Then we introduce our evaluation metrics, followed by the evaluation results on network analysis and comparison with state-of-the-arts, as well as a case of application.

5.1. Data Collection

We collect 1.52 million LinkedIn profiles containing 5.74 million working records and 2.83 million turnover records for constructing the labor mobility network. In order to compare our method with other existing methods, following Hoberg and Phillips (2013), we choose the constituents of the Standard and Poor's 1500 (S&P 1500), which is a combination of the S&P 500, the S&P MidCap 400, and the S&P SmallCap 600 and covers approximately 90% of market capitalization in the U.S., as our firm sample, and collect the corresponding capitalization information from Yahoo Finance.

5.2. Evaluation Metrics

5.2.1. Metrics for Network Analysis.

The purpose of network analysis is to take a general view of the network structure and to justify whether the network is suitable for further analysis. We choose three typical indicators to analyze our network: *Modularity*, *Clustering Coefficient* and *Density*. *Modularity* defined in Equation 1 measures the strength of division of a network into modules. *Clustering Coefficient* is a measure of the degree to which nodes in a network tend to cluster together:

$$C = \frac{3\tau_\Delta}{\tau_3}. \tag{2}$$

where τ_Δ is the number of triangles and τ_3 is the number of connected triplets of vertices in the network. *Density* of the network is defined as the number of connections a network has, divided by the total possible connections the network could have:

$$D = \frac{2L}{N(N-1)}. \tag{3}$$

where L and N denote the total number of edges and nodes respectively in a network.

5.2.2. Metrics for Industry Classification.

Following Bhojraj et al. (2003), Hoberg and Phillips (2013) and Fang et al. (2013), we compare our proposed industry classification method with other methods and commonly-used industry classification schemes by measuring how they can help to explain the movements in target firms' stock returns. We will estimate the regression specification defined in the following equation:

$$R_{i,t} = \alpha_t + \beta_t R_{p,t} + \varepsilon_{i,t} \tag{4}$$

where $R_{i,t}$ denotes the quarterly return for firm *i*, and $R_{p,t}$ denotes the average quarterly portfolio return of the industry *p* which firm *i* belongs to. Following Hoberg and Phillips (2013), we pick 10 firms in a portfolio, and run cross-sectional regressions between 2011 and 2013 on a monthly basis. Based on 36 regressions, we can obtain an average adjusted R^2, which can be used as a measurement for the effectiveness of industry classification: the higher value of the average adjusted R^2, the more effective the industry classification scheme.

5.3. Evaluation Results

In this section, we first present the results on network analysis, and then report the evaluation results on industry classification compared with other methods. Lastly, we use a case to show how our method can reveal potential company actions of entering a new field.

5.3.1. Results on Network Analysis

Although we can construct the labor mobility network in any time span, to make our method comparable with other existing methods, we construct the network on a yearly basis. The descriptive results on the network analysis is shown in Table 1.

Year	**2011**	**2012**	**2013**
Vertex scale	194530	189435	199893
Modularity	0.582	0.684	0.600
Density	0.312	0.359	0.318
Clustering Coefficient	0.647	0.702	0.680
Central Companies	Google Amazon Apple Disney	Apple Google Disney Microsoft	Apple Google Pixar Intel Adobe Intuit

Table 1. Results on Network Analysis

From the indicator of *Density*, we notice that the network is dense enough. The *Modularity* and *Clustering Coefficient* indicate that there are potential clusters, and the nodes inside each potential cluster may be tightly connected with each other. We can also see that there are some changes on the central companies during the 3 years, which means the network is dynamic but not stable. This is consistent with our expectation that the labor mobility network could be aligned with the market changes. Overall, the results on network analysis have shown that it is feasible to proceed to perform community detection.

To show the effectives of our method based on labor mobility network (denoted as LMN), we compare the average adjusted R^2 of the proposed LMN method with the commonly used schemes, i.e. SIC, NACIS and GICS. In addition, we also compare the proposed LMN method with the HP method (Hoberg & Phillips, 2013) and the FAK method (Fang et al., 2013), both of which are based on the textual descriptions of 10K forms, representing the state-of-the-arts. To make the results comparable, the number of communities is controlled to 20, which is close to the number of industries of other methods. The results are given in Figure 5, and the percentage of improvement of our method on other methods are shown in Table 2.

5.3.2. Results on Comparison with State-of-The-Arts

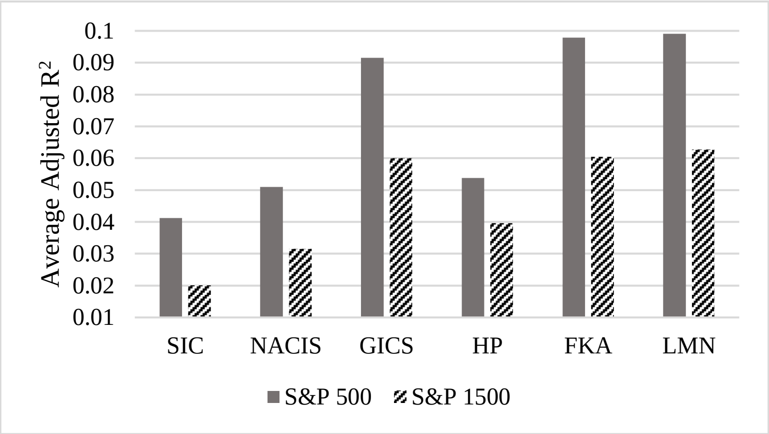

Figure 5. Comparison on Average Adjusted R^2

From the result, we can see that the average adjusted R^2 values of our LMN method is higher than all the other methods. Specifically, compared with the traditional schemes, our LMN method improves the average adjusted R^2 of the best one, i.e. GICS, by 8.31% and 4.67% on the S&P 500 and S&P 1500 samples respectively. When it comes to the state-of-the-arts, our method outperforms the FKA method which is the better one by improving its average adjusted R^2 by 1.23% and 3.97% on the two samples respectively. We notice that the magnitude of improvement is relatively small. Possible explanation is that the stock returns may be affected by many factors, and the industrial effect is only a very small part. When the explanatory power of the industrial effect is close to the upper bound, it is hard to be significantly improved.

	SIC	**NACIS**	**GICS**	**HP**	**FKA**
S&P 500	140.53%	94.31%	8.31%	84.20%	1.23%
S&P 1500	212.43%	98.73%	4.67%	58.59%	3.97%

Table 2. Percentage of Improvement on Average Adjusted R^2

5.4. A Case of Application

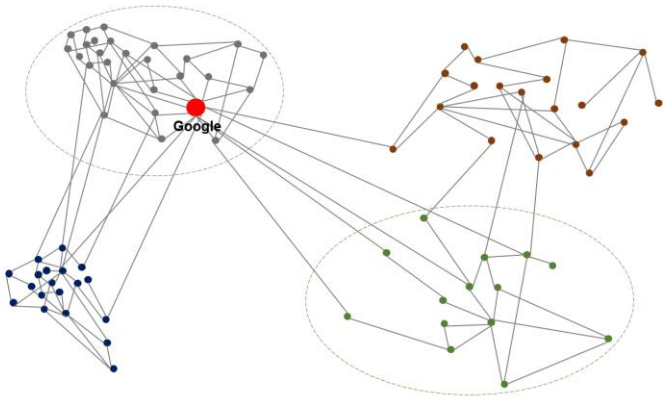

Figure 6. Partial Network in 2011

Besides improving the explanatory power of the industrial effect, our method is designed in a hierarchical manner that allows to reveal the industry structure at different scales, which is not supported by other methods. In addition, the proposed method investigates the industry structure from the perspective of human resource that may provide useful information ahead of the annual reports and other industry classification schemes.

We use a case to demonstrate the possible practical application of our method. A partial labor mobility network in 2011 constructed from the LinkedIn profiles is depicted in Figure 6. Nodes with different colors represent different communities detected by our method. The cluster in grey on the top-left is related to the industry "Internet", and the cluster in green on the bottom-right denotes the industry related to "Automobile".

It is reasonable to see that Google (denoted as a red node in Figure 6) is allocated among the Internet firms. It is interesting to find that there are some connections between Google and the "Automobile" community, from which we may infer that Google were entering the area of self-driving. In fact, Google's self-driving project began around the year 2010, and beta tests were conducted in 2012. Only after that, most of the information about the self-driving project was disclosed in its annual report. Hence, this case shows that our method is able to reveal a firm's action of entering a new field earlier, which provides great practical potentials.

6. Conclusion

In this paper, we propose a novel industry classification method with the aim of addressing the problems caused by the limited number of industry categories and the lagged information of firms' business. To achieve this goal, we employ the resource-based-view as our theoretical foundation, and propose to classify firms from the perspective of human resource. We design and implement an automatic industry classification solution. Our method extracts the turnover records from LinkedIn profiles for constructing the labor mobility network. We also extend the community detection algorithm to a hierarchical version, making it more suitable for our purpose of discovering latent firm clusters. The evaluation conducted on real-world datasets shows that the constructed labor mobility network is able to support the community detection process. In addition, our proposed method outperforms both the best industry classification scheme and the best state-of-the-art method and improves their explanatory power by 8.31% and 3.97% respectively. An application case also shows the advance of our method in revealing firms' action of entering new industries.

Our work contributes to resource-based-view by empirically testing the effectiveness of using human resource in industry analysis. Our novel approach of constructing labor mobility network and the proposed extension of hierarchical community detection have both theoretical and practical implications for related fields.

There are a few limitations in our research. First, the LinkedIn users are not evenly distributed among all industries. Generally, the IT industry tends to have more users. This sample bias problem could be addressed in future work by integrating other data sources. Second, the timeliness of our method in reflecting market changes has not been quantitatively tested. It is possible to come up with suitable metrics for quantifying and evaluating the timeliness in future work. And lastly, in future work, it would be interesting to apply our method in investigating the impact of intra-industry labor mobility and inter-industry labor mobility on the performance of firms, which has not been discussed in this work.

Acknowledgments

We gratefully acknowledge the funding support from the National Natural Science Foundation of China (Grant 71601081 and 71503084), the Guangdong Natural Science Foundation (Grant 2016A030310426) and the Fundamental Research Funds for the Central Universities (Grant 2017BQ048 and 2017BQ049). Dr. Chunmian Ge is the corresponding author.

References

Ahmed Ben, E., Ahlem, N., & Faïez, G. (2013). Group extraction from professional social network using a new semi-supervised hierarchical clustering. *Knowledge and Information Systems, 40*(1), 29-47.

Antoine, S., Cécile, S., & Hilda, G. (2017). Advanced logistics in Italy: A city network analysis. *Tijdschrift voor economische en sociale geografie, 108*(6), 753-767.

Bhojraj, S., Lee Charles, M. C., & Oler Derek, K. (2003). What's my line? A comparison of industry classification schemes for capital market research. *Journal of Accounting Research, 41*(5), 745-774.

Blondel, V. D., Guillaume, J.-L., Lambiotte, R., & Lefebvre, E. (2008). Fast unfolding of communities in large networks. *Journal of Statistical Mechanics Theory & Experiment, 2008*(10), 155-168.

Chan, L., K. C. , Lakonishok, J., & Swaminathan, B. (2007). Industry classifications and return comovement. *Financial Analysts Journal, 63*(6), 56-70.

Chen, J., & Yuan, B. (2006). Detecting functional modules in the yeast protein-protein interaction network. *Bioinformatics, 22*(18), 2283-2290.

Chong, D., & Zhu, H. (2012). Firm clustering based on financial statements. *Social Science Electronic Publishing*.

Dalziel, M. (2007). A systems-based approach to industry classification. *Research Policy, 36*(10), 1559-1574.

Fang, F., Dutta, K., & Datta, A. (2013). *LDA-based industry classification*. Paper presented at the Thirty Fourth International Conference on Information Systems, Milan.

Farjoun, M. (1994). Beyond industry boundaries: Human expertise, diversification and resource-related industry groups. *Organization Science, 5*(2), 185-199.

Fortunato, S. (2010). Community detection in graphs. *Physics Reports*, *486*(3-5), 75-174.

Ge, C., Huang, K.-W., & Png, I. P. L. (2016). Engineer/scientist careers: Patents, online profiles, and misclassification bias. *Strategic Management Journal, 37*(1), 232-253.

Hoberg, G., & Phillips, G. (2013). Text-based network industries and endogenous product differentiation. *Social Science Electronic Publishing, 124*(5), 1423-1465.

Hrazdil, K., Trottier, K., & Zhang, R. (2013). A comparison of industry classification schemes: A large sample study. *Economics Letters, 118*(1), 77-80.

Hrazdil, K., & Zhang, R. (2012). The importance of industry classification in estimating concentration ratios. *Economics Letters, 114*(2), 224-227.

Krause, A., E. , Frank, K., A. , Mason, D., M. , Ulanowicz, R., E., & Taylor, W., W. . (2003). Compartments revealed in food-web structure. *Nature, 426*(6964), 282-285.

Moody, J., & White, D., R. . (2003). Structural cohesion and embeddedness: A hierarchical concept for social groups. *American Sociological Association, 68*(1), 103-127.

Newman, M. E. J. (2004). Detecting community structure in networks. *The European Physical Journal B - Condensed Matter, 38*(2), 321-330.

Standard & Poor's and Morgen Stanley Capital International. (2002). *Global industry classification stardard-a guide to the GICS methodology*.

Wang, Z., Li, S., & Zhou, G. (2016). Personal summarization from profile networks. *Frontiers of Computer Science, 11*(6), 1085-1097.