Expertise Assessment with Multi-Cue Semantic Information

Jun Wang, Kush R. Varshney, Aleksandra Mojsilović, Dongping Fang, and John H. Bauer
Business Analytics and Mathematical Sciences
IBM Thomas J. Watson Research Center
1101 Kitchawan Road, Yorktown Heights, New York 10598
Email: {wangjun, krvarshn, aleksand, dfang, jhbauer}@us.ibm.edu

***Abstract*—Assessing and managing the expertise of employees in knowledge and service industries is critical because human capital is the key differentiator among companies. Moreover, professional social networks are becoming increasingly popular. Besides the well-known public professional social network site LinkedIn, enterprise social networks are also now being widely used inside corporations and companies. In this paper, we address the critical workforce analytics problem of automatically assessing employees' skills by mining multiple cues found in enterprise and social data. In particular, we treat the assessment of employees' expertise as a matrix completion problem, where the rows represent individual employees and the columns represent individual skills. The multiple cues about employee expertise come from data we integrate on the existing skill assessment process within the company, the social networking and social media activity of the employees, and the semantic similarity of skills. Assessment results are evaluated as a binary classification recommendation. Extensive empirical study using a real-world data set from a large multinational Fortune 500 corporation corroborates the effectiveness of multi-cue analytics to improve the coverage and accuracy of skill assessment.**

I. Introduction

Globalization, recognized as the process of extending social relations worldwide, has driven economic growth in the past decades. In addition, globalization has pushed many large multinational corporations to transform from collections of several domestically-based organizations to globally-integrated and cross-cultural enterprises [1]. In particular, advances in transportation and telecommunications infrastructures served as the first major impetus for such a globalization trend. More recently, the rise of social networking on the Internet has made social interaction across world-space easier and more frequent. For example, the world's largest social network, Facebook, consists of more than 1 billion users with over 1 trillion connections[1] and LinkedIn, the world's largest professional social network, reached 200 million members in January 2013.[2]

Social networks provide efficient ways for people to communicate and collaborate, and help businesses be more competitive and successful. Therefore, there is also an increasing use and need of social network and social media technologies inside every enterprise and corporation, which leads to a more connected, interactive, and enabled workforce [2], [3], [4].

[1]http://en.wikipedia.org/wiki/Facebook
[2]http://blog.linkedin.com/2013/01/09/linkedin-200-million

Traditionally, governments are the employers with the largest workforces. For example, the United States Department of Defense is the world's largest employer with around 3.2 million employees.[3] However, globalized business drives workforce growth of non-governmental employers since companies tend to require more employees to work directly with local clients [5]. Table I lists the 20 largest employers from the Fortune Global 500 by number of employees, where non-corporate public employers are excluded.[4] It is no surprise that companies from traditional industries, such as retail and energy, tend to have large workforces since they are labor-intensive. In addition, technology companies, like International Business Machines and Hon Hai Precision Industry, also have considerable workforce size due to their growth in the global market.

There are many emerging challenges for human resource (HR) management and workforce analytics and optimization in modern multinational corporations. To address these challenges, the fast growth of data sources like enterprise social networks can be utilized to help corporations manage, transform, engage, and plan for their workforces. In particular, social networking naturally captures information about the activities, interactions, and knowledge of employees in a digital form that can be mined for insight and business process improvements [6], [7], [8], [9].

Knowledgeable employees are not interchangeable because they each have specialized expertise and skills; this has an important role in a corporation's business, especially in service industries with much client interaction. Hence, it is critical to capture and understand the individual specialities of employees for successful human capital management and operation in large enterprises. For instance, if the expertise of each employee can be comprehensively catalogued, the assignment of projects can be extremely efficient and accurate since the desired experts can be called upon to meet clients' needs. From a more strategic perspective, an accurate understanding of employees' skill information can be used to plan for an enterprise's long term business goals.

However, existing HR management tools are insufficient in terms of handling such challenging issues in both scalability and depth [9]. First, although employees can be easily categorized by their organization charts and reporting chains, it

[3]http://www.bbc.co.uk/news/magazine-17429786
[4]http://money.cnn.com/magazines/fortune/global500/2012/performers/companies/biggest

TABLE I. THE WORLD'S 20 LARGEST EMPLOYERS FROM THE FORTUNE GLOBAL 500 BY NUMBER OF EMPLOYEES.

Rank	Company	Region	Employees(2012)	Fortune 500 Rank
1	Walmart Stores	North America	2.2 million	3
2	China National Petroleum Corporation	Asia	1.7 million	6
3	State Grid Corporation of China	Asia	1.6 million	7
4	Sinopec	Asia	1 million	5
5	Hon Hai Precision Industry	Asia	961 thousand	43
6	China Post Group	Asia	889 thousand	258
7	U.S. Postal Service	North America	602 thousand	135
8	Volkswagen	Europe	502 thousand	12
9	China Telecommunications	Asia	491 thousand	221
10	Aviation Industry Corp. of China	Asia	480 thousand	250
11	Compass Group	Europe	471 thousand	432
12	Agricultural Bank of China	Asia	447 thousand	84
13	International Business Machines	North America	433 thousand	57
14	Deutsche Post	Europe	424 thousand	98
15	McDonald's	North America	420 thousand	410
16	Carrefour	Europe	412 thousand	39
17	Industrial & Commercial Bank of China	Asia	409 thousand	54
18	Tesco	Europe	406 thousand	59
19	Gazprom	Europe	401 thousand	15
20	Sodexo	Europe	391 thousand	495

is difficult to characterize the employees by specific functions they carry out or knowledge they have. An expertise or skill taxonomy within a company is one structure for representing the various functional abilities or knowledge that employees may have [10]. Such taxonomies can be used for various business processes [11], [12].

Although it is important to capture skills information in a structured format, it is fairly time-consuming since constructing a skill taxonomy and assessing employees against it relies heavily on manual processing. Especially for a modern multinational corporation, since the number of employees can be hundreds of thousands and even several million (as shown in Table I), and the skills carried by the workforce could cover a very wide range of fields, it is extremely challenging to develop such a complex skill taxonomy with such a scale. Second, even if such a taxonomy could be built and populated to organize and index employees by their expertise, search and retrieval of experts using the skill taxonomy is not straightforward. This is mainly because the enterprise's skill taxonomy often consists of specific technical terms and there exists a clear semantic gap between those terms and natural search keywords. So it is necessary to map the technical skill terms to a common set of concepts.

Realizing all of these challenges and opportunities, in this paper, we focus our study on automatic expertise assessment using various types of semantic sources that each reveal a different cue of skill, knowledge, and expertise. In particular, we treat the expertise assessment as a recommendation problem, where we utilize several prediction techniques for solving such a problem. First, we conduct a basic matrix completion task using a set of incomplete observed skill assessments. Then we explore the social context and skill semantics to perform collaborative filtering, content filtering and a hybrid approach to predict the skills of employees. Empirical study using a real world employee data set from a multinational Fortune 500 corporation clearly demonstrates the strengths and weaknesses of each method. The results indicate that a further study for such a problem should be directed towards combined models with multi-cue semantic information.

The remainder of this paper is organized as follows. Section 2 briefly reviews the problem of expertise assessment and the corresponding use cases. Section 3 presents the methods used for predicting experts' skills. Section 4 provides experimental validations and comparative studies, and, finally, Section 5 concludes the paper and discusses further work.

II. Skill Assessment and Use Cases

In this section, we set forth the concrete workforce analytics problem that we are considering and discuss the use cases that a solution to the problem would enable.

A. *Problem Formulation*

In this work, we consider a corporation with an existing expertise taxonomy of hundreds or thousands of fine-grained skills with detailed textual descriptions that has been created and curated manually. Within this structure of skills, employees have assessment scores that indicate their level of mastery of the skill, ranging from no expertise at all, to having acquired the skill, to having mastered the skill, to being so skilled that the employee is viewed as a thought leader. However, due to the multitude of skills and the semantic gap discussed in Section I, most employees are not assessed on most skills. If we think of employees as rows and skills as columns in a table, with skill assessment values as entries of table cells, then a large fraction of the cells are empty.

The problem we tackle is to fill in the empty cells through analytics-based means in order to get a complete picture about the expertise within the corporation. The basic premise for such analytics is that similar employees have similar assessment levels on similar skills. Moreover, we have access to multiple cues to define these similarities. The already-entered skill assessments are one cue, the textual descriptions of the skills are another cue, and the data about employee knowledge and interaction captured through enterprise social networking technology is a third cue. The overall idea is illustrated in Fig. 1. We discuss the precise mathematical formulation for filling in the missing skill assessments in Section III.

B. *Use Cases*

Based on predicted values for employee skill assessments, there are several business processes that can be improved

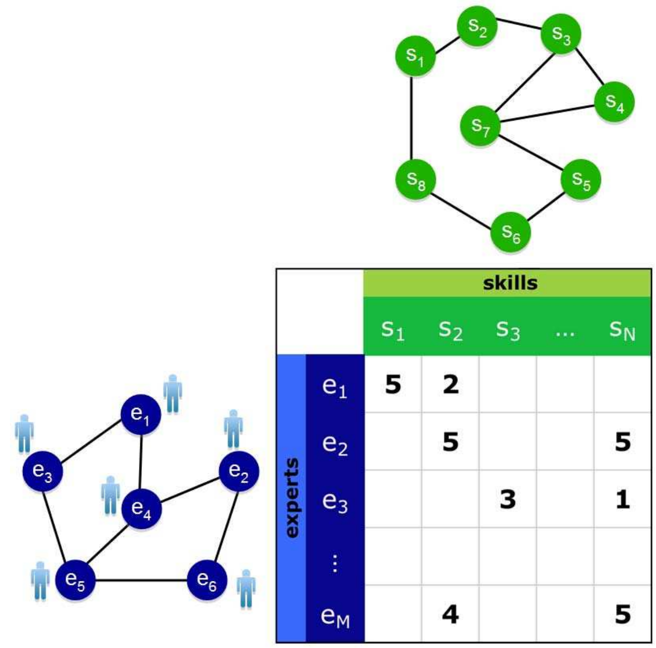

Fig. 1. Illustration of skill prediction problem formulation using multi-cue semantic information.

and uses cases satisfied [9]. First, the predicted values for employees can be recommended to them in a user-friendly manner in which they can confirm or reject the prediction. Second, a similar interface can be used for employees to endorse the skills of their peers. For planning and management purposes, the predictions themselves can be used as proxy values for characterizing expertise in the organization. Another use case is using analytics to construct new or emerging skills in the taxonomy. Also, the predicted values can be used in locating experts within the company for various reasons.

III. METHODOLOGY

We first give a brief introduction of the notations used in this paper. Assume we have a set of employees $\{\mathbf{e}_i\}_{i=1}^{M}$ and that there are a total of N skills for assessment. The skill level of an employee $\mathbf{e}_i$ for the j-th skill is denoted by m_{ij}, hence the employee-skill matrix $\mathbf{M} = \{m_{ij}\} \in \mathbb{R}^{M \times N}$ indicates all of the skill assessments. Given a partial observation of the skill assessment as $\tilde{\mathbf{M}}$, the goal for skill assessment is to complete the underlying matrix $\mathbf{M}$.

A. Matrix Completion for Skill Assessment

Without considering any context, the problem of skill assessment can be treated as a matrix completion problem, where the objective is to complete the missing values in the employee-skill matrix $\mathbf{M} \in \mathbb{R}^{M \times N}$ under certain structural assumptions on $\mathbf{M}$. Here we describe two related methods: matrix factorization and low-rank matrix estimation.

Matrix factorization-based approaches assume that the target matrix can be decomposed into two matrices $\mathbf{W} \in \mathbb{R}^{L \times M}$ and $\mathbf{H} \in \mathbb{R}^{L \times N}$ as $\mathbf{M} = \mathbf{W}^\top \mathbf{H}$ [13]. Given a partially observed matrix $\tilde{\mathbf{M}} = \{\tilde{m}_{ij}\}$ with $\mathcal{O} = \{(i,j)\}$ indicating the set of matrix cells with observed values, a basic formulation of matrix factorization can be written in a square loss form as:

$$\arg\min_{\mathbf{W},\mathbf{H}} \sum_{(i,j)\in\mathcal{O}} \|\tilde{m}_{ij} - \mathbf{w}_i^\top \mathbf{h}_j\|^2 \tag{1}$$

where $\mathbf{w}_i$ and $\mathbf{h}_j$ are the i-th and j-th column vectors of the hidden matrices $\mathbf{W}$ and $\mathbf{H}$, respectively.

As a classical latent factor model, the above minimization problem can be interpreted as the representation learning of both user and item features in an unknown feature space $\mathbb{R}^L$ with L as the latent dimension. A similar formulation has been widely used for designing modern recommendation systems [14], [15], where a very successful application is recommending movies to the customers of the on-demand Internet streaming media company Netflix [16]. In our particular application of skill assessment, the employees are the users and the skills are the items. Imposing nonnegativity constraints on both the $\mathbf{W}$ and $\mathbf{H}$ matrices leads to another well-know variant: nonnegative matrix factorization (NMF) [17], [18]. In addition, probabilistic algorithms have been applied to matrix factorization modeling; the resulting probabilistic matrix factorization (PMF) methods have been shown effective in handling large-scale applications as they scale linearly with the number of observations [19], [20].

The low-rank matrix approximation approach assumes that the underlying rank of the target matrix is low. Intuitively, the users and items are clustered in a small number of groups. Hence, given a partial observation, it is possible to recover the target matrix directly. Although matrix factorization and matrix rank minimization methods have different explicit formulations, they are closely related and sometime equivalent [21], [22]. For instance, Ma et al. propose an iterative approach, namely fixed point continuation with approximate singular value decomposition to solve the following matrix rank minimization problem [22]:

$$\begin{aligned} &\min_{\mathbf{M}} \operatorname{rank}(\mathbf{M}) \\ &\text{s.t. } m_{ij} = \tilde{m}_{ij} \quad \text{for} \quad (i,j) \in \mathcal{O}. \end{aligned} \tag{2}$$

Matrix factorization-based solutions are also regarded as model-based collaborative filtering approaches because the resulting predictions are based on a certain form of statistical modeling [23]. Though such model-based approaches show promising results in many applications, one of their disadvantages is that the training cost is usually fairly high when handling large-scale data sets.

B. Collaborative Filtering Using Multi-Cue Semantics

Memory based approaches are another popular category for performing such predictions due implementation ease and computational efficiency. Typically, one can first use the observed skill assessment data to compute similarity between employees and skills. Then the nearest neighborhood mechanism can be applied to perform either item-based or user-based top-N prediction [23]. For example, the similarity between two employees $\mathbf{e}_i$ and $\mathbf{e}_j$ can be computed using Pearson correlation as

$$sim(\mathbf{e}_i, \mathbf{e}_j) = \frac{(\tilde{\mathbf{m}}_i - \bar{m}_i)^\top (\tilde{\mathbf{m}}_j - \bar{m}_j)}{\sqrt{\|\tilde{\mathbf{m}}_i - \bar{m}_i\|^2 \|\tilde{\mathbf{m}}_j - \bar{m}_j\|^2}}, \tag{3}$$

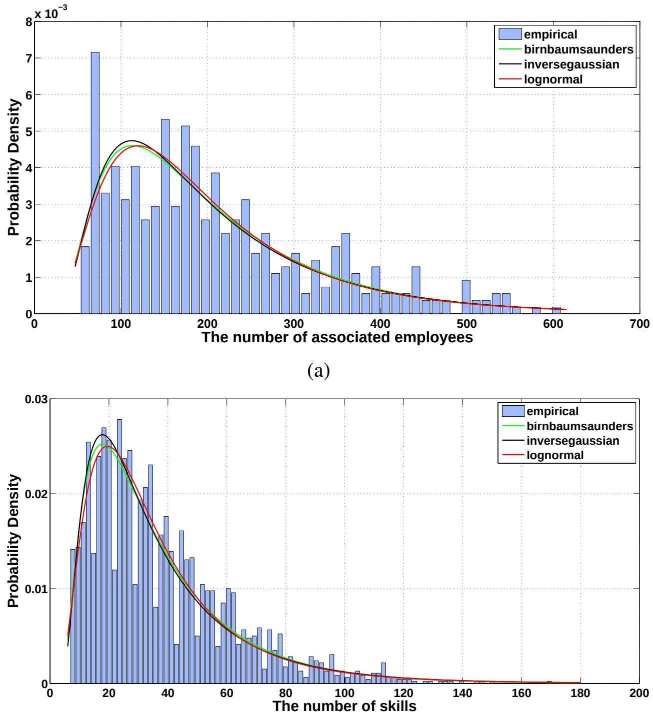

Fig. 2. The distributions of a) the number of skills per employee, and b) the number of associated employees per skill.

where $\bar{m}_i$ and $\bar{m}_j$ represent the average skill rating for the i-th and j-th employee across all skills. Here $\tilde{\mathbf{m}}_i$ and $\tilde{\mathbf{m}}_j$ are the observed skill assessment for the i-th and j-th employee. Similarly the similarity between skills can also be calculated using the partially observed matrix $\tilde{\mathbf{M}}$. However, such neighborhood-based algorithms rely on the sufficient completeness of the observed assessments. In realistic scenarios, usually only a very small portion of employee-skill cells are filled with assessed values; thus, the estimation of the similarity could be extremely unstable.

However, in the application of expertise assessment, besides the partially observed expert-skill matrix, rich semantic information about the employees and skills can be acquired. For instance, as mentioned earlier, the usage of enterprise social networking creates a platform for employees to discuss their projects and research topics. Mining such social media data can help derive the social proximity between employees. For example, the employees' social activity on technical communities and microblogs can be extracted to form semantic information, which can be used to estimate the their skill background. Assume we can extract the semantic representation $\mathbf{e}_i$ for the i-th employees. Then the social proximity between any two employees can be estimated using such representations. Simply the vector cosine-based similarity can be computed as

$$sim(\mathbf{e}_i, \mathbf{e}_j) = \frac{\mathbf{e}_i^\top \mathbf{e}_j}{\sqrt{\|\mathbf{e}_i\|^2 \|\mathbf{e}_j\|^2}}.$$

Then a user-based collaborative filtering algorithm can be used to estimate an unknown skill assessment m_{ik} as

$$m_{ik} = \frac{\sum_{(j,k)\in\mathcal{O}} sim(\mathbf{e}_i, \mathbf{e}_j)\tilde{m}_{jk}}{\sum_{(j,k)\in\mathcal{O}} sim(\mathbf{e}_i, \mathbf{e}_j)}. \quad (4)$$

Intuitively, the skill assessment for the k-th skill on the i-th employee is the weighted sum of the skill assessment of all the similar employees on the same skill. The above prediction is essentially a weighted nearest-neighbor method in the employee space if we truncate the similarity between $\mathbf{e}_i$ and his or her similar employees and only choose the most similar ones.

In addition, the skill can be further analyzed by using its descriptions and definitions. Such analysis can help derive semantic features for each skill. Then we can compute the semantic similarity between skills. Let us assume the skills are represented by feature vectors as $\{\mathbf{s}_i\}_{i=1}^N$. Similarly, the semantic similarity $sim(\mathbf{s}_i, \mathbf{s}_j)$ between two skills $\mathbf{s}_i$ and $\mathbf{s}_i$ can be computed accordingly. Hence, the skill assessment can also be estimated using the skill similarity as

$$m_{ik} = \frac{\sum_{(i,j)\in\mathcal{O}} sim(\mathbf{s}_j, \mathbf{s}_k)\tilde{m}_{ij}}{\sum_{(i,j)\in\mathcal{O}} sim(\mathbf{s}_j, \mathbf{s}_k)}. \quad (5)$$

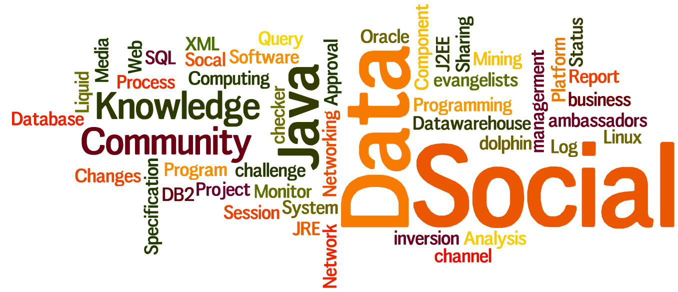

Fig. 3. An example of the employee's semantic information on the enterprise social network represented by keywords.

Finally, since both employee and skill can be represented by semantic features, it is straightforward to combine both pieces of information to perform two-way prediction as

$$\begin{aligned} m_{ik} &= (1-\mu)\frac{\sum_{(j,k)\in\mathcal{O}} sim(\mathbf{e}_i, \mathbf{e}_j)\tilde{m}_{jk}}{\sum_{(j,k)\in\mathcal{O}} sim(\mathbf{e}_i, \mathbf{e}_j)} \\ &+ \mu\frac{\sum_{(i,j)\in\mathcal{O}} sim(\mathbf{s}_j, \mathbf{s}_k)\tilde{m}_{ij}}{\sum_{(i,j)\in\mathcal{O}} sim(\mathbf{s}_j, \mathbf{s}_k)}, \end{aligned} \quad (6)$$

where μ is a parameter that indicates the relative weights given to each type of semantic information.

IV. Empirical Study

To perform an empirical study using the aforementioned prediction approaches, we acquired a set of employee data from a multinational Fortune 500 corporation. Below we start by introducing the data and experimental settings, followed by the evaluation results.

Fig. 4. An example of the skill's semantic information represented by keywords.

A. Data and Overview

We collected data for a total of 2618 employees who assessed 471 unique skills. Hence, we have an employee-skill matrix with the size 2618×471. The number of assessed skills for each employee ranges from 6 to 180 with a average number of skills as 40. The popular skills have up to 615 experts and the most rare skill only has 47 assigned experts. Figure 2(a) shows the distribution of the number of associated employees per skill and Figure 2(b) shows the distribution of the number of skills per employee, where Birnbaum-Saunders, Inverse Gaussian, and Log-normal distributions all give good approximation of the empirical histograms.

For employee's semantics, we crawl the content from enterprise social networks, including research/project communities and microblogs, and the associated textual tags. Based on the frequency of the online social activities, a relevance score is estimated for each textual tag for the employee. Figure 3 and 4 demonstrate examples of the extracted employee and skill semantic information in the form of a weighted keyword representation, where the size of the word indicates the strength of the relevance for that keyword.[5]

Clearly, such a semantic representation provides an informative description of the employee. In this case, this particular employee very likely has database, data analysis, and social networking-related skills since his or her online activities cover related topics. In our study, we constructed a semantic dictionary with a total of 5701 keywords to represent each employee's semantic information. Based on such a representation, we can compute the online social similarity, as described above. Meanwhile, for each skill, we build similar semantic representation with a 2514-keyword dictionary and compute the skill similarity.

B. Evaluations and Results

To show the effectiveness of the computed semantic similarity of employees and skills, we plot the curve of semantic similarity between employees $sim(\mathbf{e}_i, \mathbf{e}_j)$ versus the values of $sim(\mathbf{m}_i, \mathbf{m}_j)$ in Figure 5(a), and the skill semantic similarity $sim(\mathbf{s}_i, \mathbf{s}_j)$ versus the values of $sim(\mathbf{m}_{\cdot i}, \mathbf{m}_{\cdot j})$ in Figure 5(b). Here, $\mathbf{m}_{\cdot i}$ and $\mathbf{m}_{\cdot j}$ are the i-th and j-th column vectors, representing the assessment level of th i-th and j-th skills across

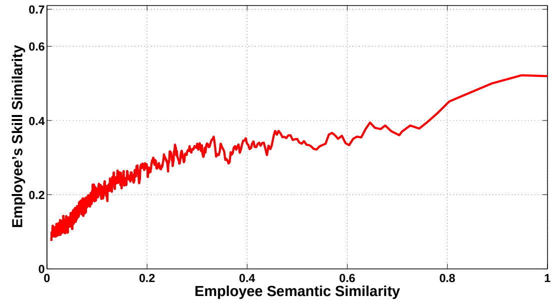

(a)

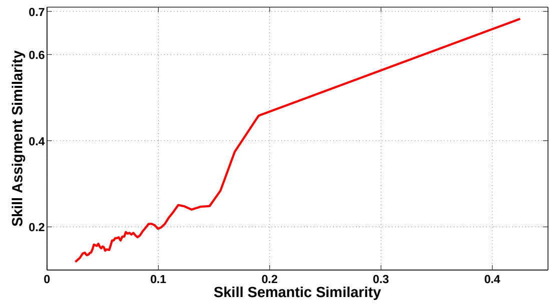

(b)

Fig. 5. Evaluation of the effectiveness of (a) employee semantic similarity and (b) skill semantic similarity for the prediction of skill assessment.

the employee population. It is straightforward to conclude that the computed semantic similarities of employees and skills are correlated to the skill assessment. In other words, similar employees tend to share a similar set of skills and that similar skills tend to be associated to similar employee populations.

To validate the performance of the aforementioned methods, here we simply view skill assessment as a binary prediction problem, where the goal is to predict whether an employee has certain skills. Since the employee-skill matrix is very sparse with only around $3\% - 4\%$ nonzero elements, here we measure the prediction error by the average of false positive rate and false negative rate as:

$$\alpha = \frac{FalsePositive}{FalsePositive + TrueNegative}$$
$$\beta = \frac{FalseNegative}{TruePositive + FalseNegative}$$
$$\text{error} = \frac{1}{2}(\alpha + \beta) \tag{7}$$

where α and β are the false positive and false negative rates, respectively. In addition, we randomly split the cells in the $\mathcal{M}$ matrix into 6 folds and use cross-validation to compute the average errors.

We evaluate the following methods: 1) basic matrix completion (BMC) using the low rank approximation [22], 2) basic matrix completion combined with skill semantic information (BMC-SSI), 3) basic matrix completion combined with employee semantic information (BMC-ESI), and 4) basic matrix completion combined with both types semantic information (BMC-BSI). Since the BMC method gives the prediction in the continuous values, we need to binarize the predicted values

[5]The visualization of the word representation is generated using the online tool: http://www.wordle.net.

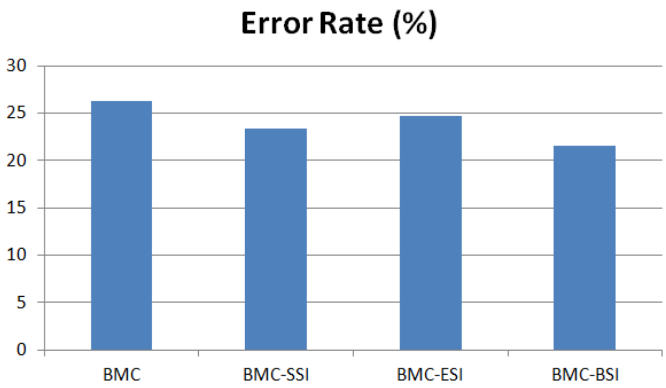

Fig. 6. Prediction performance measured by the average error percentage for different approaches.

to generate the results. In particular, we use the training data to estimate the ratio of positive and negative values in the expert-skill matrix. Such estimates are then been used to compute an adaptive threshold value to generate the binary predictions for the test subset, where the positive/negative ratio is the same as the training subset. To combine the predictions from basic matrix completion with semantic information for BMC-SSI, BMC-ESI, and BMC-BSI, we use a simple after-fusion approach by applying a weighting and normalization scheme [24].

Figure 6 demonstrates the prediction performance measured by the error rates. Clearly, both semantic information of employees and skills help to improve the performance of the basic matrix completion methods, and using both semantic information can further boost the performance. In addition, the performance gain by using the skills' semantic information is higher than that using employees' semantic information, which is also consistent with the observation in Figure 5.

V. CONCLUSIONS

In this paper, we address the problem of automatically assessing employees' expertise, which plays an important role in human capital management and workforce analytics. In particular, we treat the assessment of employees' expertise as a matrix completion problem with the rows and columns representing individual employees and skills. Besides the well-known matrix completion techniques, we propose to explore the professional social network to retrieve multi-cue semantic information for both employees and skills. Empirical study using real world human resource data corroborates that the extracted multi-cue semantic information can help boost the performance of skill assessment. One of our future directions of research is to design a generic prediction model to combine the matrix structure information and the semantic information for robust skill assessment.

REFERENCES

[1] S. J. Palmisano, "The globally integrated enterprise," *Foreign Aff.*, vol. 85, no. 3, pp. 127–136, May/Jun. 2006.

[2] P. Raghavan, "Social networks: From the web to the enterprise," *IEEE Internet Comput.*, vol. 6, no. 1, pp. 91–94, Jan.–Feb. 2002.

[3] A. P. McAfee, "Enterprise 2.0: The dawn of emergent collaboration," *MIT Sloan Manage. Rev.*, vol. 47, no. 3, pp. 21–28, Spr. 2006.

[4] A. Wu, J. M. DiMicco, and D. R. Millen, "Detecting professional versus personal closeness using an enterprise social network site," in *Proc. SIGCHI Conf. Hum. Fact. Comput. Syst.*, Atlanta, GA, Apr. 2010, pp. 1955–1964.

[5] C. K. Prahalad and V. Ramaswamy, *The Future of Competition: Co-Creating Unique Value with Customers*. Boston, MA: Harvard Business School Press, 2004.

[6] K. Ehrlich, C.-Y. Lin, and V. Griffiths-Fisher, "Searching for experts in the enterprise: Combining text and social network analysis," in *Proc. ACM Conf. Supporting Group Work*, Sanibel Island, FL, Nov. 2007, pp. 117–126.

[7] C. Chelmis, V. Sorathia, and V. K. Prasanna, "Enterprise wisdom captured socially," in *Proc. IEEE/ACM Int. Conf. Adv. Soc. Netw. Anal. Min.*, Istanbul, Turkey, Aug. 2012, pp. 1228–1235.

[8] C.-Y. Lin, L. Wu, Z. Wen, H. Tong, V. Griffiths-Fisher, L. Shi, and D. Lubensky, "Social network analysis in enterprise," *Proc. IEEE*, vol. 100, no. 9, pp. 2759–2776, Sep. 2012.

[9] K. R. Varshney, J. Wang, A. Mojsilović, D. Fang, and J. H. Bauer, "Predicting and recommending skills in the social enterprise," in *Proc. Int. AAAI Conf. Weblogs Soc. Med.*, Cambridge, MA, Jul. 2013.

[10] D. R. Ilgen and J. R. Hollenbeck, "The structure of work: Job design and roles," in *Handbook of Industrial and Organizational Psychology*, M. D. Dunnette and L. M. Hough, Eds. Palo Alto, CA: Consulting Psychologists Press, 1991, pp. 165–207.

[11] J. Hu, B. K. Ray, and M. Singh, "Statistical methods for automated generation of service engagement staffing plans," *IBM J. Res. Dev.*, vol. 51, no. 3/4, pp. 281–293, May/Jul. 2007.

[12] Y. Naveh, Y. Richter, Y. Altshuler, D. L. Gresh, and D. P. Connors, "Workforce optimization: Identification and assignment of professional workers using constraint programming," *IBM J. Res. Dev.*, vol. 51, no. 3/4, pp. 263–279, May/Jul. 2007.

[13] Y. Koren, R. Bell, and C. Volinsky, "Matrix factorization techniques for recommender systems," *IEEE Computer*, vol. 42, no. 8, pp. 30–37, Aug. 2009.

[14] P. Resnick and H. R. Varian, "Recommender systems," *Comm. ACM*, vol. 40, no. 3, pp. 56–58, Mar. 1997.

[15] F. Ricci, L. Rokach, and B. Shapira, "Introduction to recommender systems handbook," in *Recommender Systems Handbook*, F. Ricci, L. Rokach, B. Shapira, and P. B. Kantor, Eds. New York, NY: Springer, 2011, pp. 1–35.

[16] J. Bennett and S. Lanning, "The Netflix prize," in *Proc. KDD Cup Workshop*, San Jose, CA, Aug. 2007, pp. 3–6.

[17] D. D. Lee and H. S. Seung, "Algorithms for non-negative matrix factorization," in *Adv. Neural Inf. Process. Syst. 13*. Cambridge, MA: MIT Press, 2001, pp. 556–562.

[18] ——, "Learning the parts of objects by non-negative matrix factorization," *Nature*, vol. 401, no. 6755, pp. 788–791, Oct. 1999.

[19] R. Salakhutdinov and A. Mnih, "Probabilistic matrix factorization," in *Adv. Neural Inf. Process. Syst. 20*. Cambridge, MA: MIT Press, 2008, pp. 1257–1264.

[20] ——, "Bayesian probabilistic matrix factorization using Markov chain Monte Carlo," in *Proc. Int. Conf. Mach. Learn.*, Helsinki, Finland, Jul. 2008, pp. 880–887.

[21] B. Recht, M. Fazel, and P. A. Parrilo, "Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization," *SIAM Rev.*, vol. 52, no. 3, pp. 471–501, Aug. 2010.

[22] S. Ma, D. Goldfarb, and L. Chen, "Fixed point and Bregman iterative methods for matrix rank minimization," *Math. Program.*, vol. 128, no. 1–2, pp. 321–353, Jun. 2011.

[23] X. Su and T. M. Khoshgoftaar, "A survey of collaborative filtering techniques," *Adv. Artif. Intell.*, vol. 2009, p. 421425, Aug. 2009.

[24] A. Jain, K. Nandakumar, and A. Ross, "Score normalization in multimodal biometric systems," *Pattern Recogn.*, vol. 38, no. 12, pp. 2270–2285, Dec. 2005.