

Skillrank: Towards a hybrid method to assess quality and confidence of professional skills in social networks

José María Álvarez-Rodríguez^{1,2}, Ricardo Colomo-Palacios³, Vladimir Stantchev⁴

¹Universidad Carlos III de Madrid

Av. Universidad 30, 28911 Leganés, Madrid, Spain

0034 91 624 5958

josemaria.alvarez @uc3m.es

²Wroclaw University of Technology, Wyspianskiego 27, 50-370 Wroclaw, Poland

³Østfold University College, B R A Veien 4, 1783 Halden, Norway,

ricardo.colomo-palacios@hiof.no

⁴SRH University Berlin, Ernst-Reuter-Platz 10, 10587 Berlin, Germany

vladimir.stantchev@srh-hochschule-berlin.de

Abstract

The present paper introduces a hybrid technique to measure the expertise of users by analyzing their profiles and activities in social networks. The main focus lies on the ability of assessing the quality and confidence of professional skills after spreading endorsements and skills in a social web environment. Currently, both job seekers and talent hunters are looking for new and innovative techniques to filter jobs and candidates as well as candidates are trying to improve and make more attractive their profiles. In this sense, the Skillrank approach is based on the conjunction of existing and well-known information and expertise retrieval techniques that perfectly fits to the existing web and social media environment to deliver an intelligent component to integrate the user context in the analysis of skills confidence. A major outcome of this approach is that it actually takes advantage of existing data and information available on the web to perform both: a ranked list of experts in a field and a confidence value for every professional skill. Thus, expertise and experts can be detected, verified and ranked using a suited trust metric. The paper also presents a validation of Skillrank performance in terms of precision and recall by means of a sound qualitative and quantitative approach based on the opinions of a panel of experts on two different datasets: 1) ad-hoc created using real data from a professional social network and 2) real data extracted from the LinkedIn API. Finally, some discussion, conclusions and future work are also outlined.

Keywords: expertise retrieval, expert finder, hybrid technique, quality assessment, information retrieval, social networks

1. Introduction

In recent years, social network research has been carried out using data collected from online interactions and from explicit relationship links in online social network platforms like, for instance, Facebook and LinkedIn [1]. Among these tasks, expert and people search is one of the most challenging tasks that one can try in social networks [2] [3].

Expertise represents the skill of answering some questions or conducting some activities [4]. Thus, the focus of expertise location is finding an answer, a solution, or a person with whom details of a problem can be discussed [5] or a task can be performed [6]. In other words, expert finding addresses the task of identifying the right person with the appropriate skills and knowledge. Effective management of expertise can benefit both organizations and individuals by easing the access to knowledge, as well as sharing and applying knowledge [7].

In this light, expert finding involves two main aspects including expertise identification (“*Who are the experts on Topic X?*”) and expertise selection (“*What does Expert Y know?*”) [8]. In the later topic, expert profiling turns the expert-finding task around and asks the following: *What topic(s) does a person know about?* [9]. Topics such as expertise relevance and authority within a community have been pointed out as some of the factors to assess expert’s competence [10], [11]. Given that complete and accurate expert profiles enable people and search engines to effectively and efficiently locate the most appropriate experts for an information need [9], this paper presents an expert profiling approach to analyze expert’s skills confidence by means of hybrid soft computing techniques. One of the main advantages of Skillrank is the use of LinkedIn as a source of expertise. LinkedIn is likely the most notable example of business-oriented social networking site. The company was founded in December 2002 and launched six months later. LinkedIn reports by December 2014 more than 330 million users in 200 countries and territories. In this professional social networking site, users are allowed to track and publish their career paths, skills and past experiences, the size and tenure of the teams with whom they’ve worked, and the roles they played on each team [12]. LinkedIn users self-report their expertise and ask members of their social network to provide positive references or recommendations for them [13]. Although LinkedIn has been used in the literature for expertise search [14], [15], to the best of authors knowledge, there is not a study devoted to the application of self-disclosure and social network integrators to assess the quality and confidence of professional skills in this network.

On the other hand, a good number of techniques have been designed to exploit the information available in social networks and, in general, to address problems that contain an implicit graph. The well-known algorithm PageRank [16] by Google Inc. was developed to assign a measure of importance to each web page. This algorithm works by counting the number and quality of links to a page to determine a rough estimate of how important a website is. The underlying assumption is that more important websites are likely to receive more links from other websites. In the same way, the HITS (Hyperlink-Induced Topic Search) algorithm [17] also known as “hubs and authorities” is a kind of analysis technique that also rates web pages. It was designed by J. Kleinberg from the Department of Computer Science at Cornell and the idea behind hubs and authorities stemmed from a particular insight into the creation of web pages when the Internet was originally forming; that is, some web pages are known as hubs and serve as hubs that compile large directories of web pages. These directories are not actually authoritative in some topic but a good hub represents a page that points to many other pages and a good authority page is expected to be linked to many hubs. The main restriction of the HITS algorithm lies in its applicability since it only operates in a small sub graph. This sub graph is considered to be query dependent, whenever the search contains a different query phrase, the seed changes as well as the HITS algorithm ranks the seed nodes according to their authority and hub weights. The SPEAR (Spamming-resistant Expertise Analysis and Ranking) algorithm [18] is another tool for ranking users in social networks by their expertise and influence within a community. It is also a graph-based technique to measure the expertise of users by analyzing their activities and interaction.

The main idea behind this technique lies on the ability of users to find new and high-quality information on the Internet. This algorithm is an extension of the aforementioned HITS algorithm including two main elements: 1) Mutual reinforcement of user expertise and document quality and 2) Discoverers vs. followers. The combination of both elements has been demonstrated to reward quality over quantity of user activities and that is why it has been also applied to detect spam attacks [19]. Although graph analysis techniques [20] have been widely used to study social networks (e.g. trend detection, opinion mining, sentiment analysis, information retrieval, etc.) and, in most of cases, the PageRank algorithm can be seen as a precursor of this kind of approach there is still lack of techniques to deal with quality over quantity. In this sense, the SPEAR algorithm offers us a technique that can be applied to a rather wide of area of domains such as assessment of skills quality. In the context of graph-based algorithms for expertise ranking, the ExpertRank [10] algorithm proposes a novel technique to evaluate the expertise of users based on both document-based relevance and one's authority in this or her knowledge community. Authors modified the PageRank algorithm to evaluate one's authority so that it reduces the effect of certain biasing communication behavior in online communities. As an important cornerstone and relevant to this work, they explored three different expert ranking strategies that combine document-based relevance and authority: linear combination, cascade ranking, and multiplication scaling. This evaluation has been done using a popular online knowledge community showing that the proposed algorithm achieves the best performance when both document-based relevance and authority are considered

In this paper a re-interpretation and extension of the SPEAR algorithm, called Skillrank, is presented. Furthermore, the evaluation of the presented approach is carried out by comparing existing approaches for expertise ranking such as the HITS and SPEAR algorithms to the proposed technique when tests are executed on top of two datasets extracted from the LinkedIn API. To do so, a panel of experts has established a set of expected results and values that are compared to the real results provided by each algorithm with the aim of obtaining measures of precision and recall.

The remainder of this paper is organized as follows. Section 2 presents the related literature. The proposed approach for skill ranking is illustrated in Section 3. In Section 4, experiment evaluations are conducted to compare our approach with other methods. Section 5 presents main conclusions and future research directions.

2. State of the art

Expert and credibility finding is not a new issue in literature. As a result of this, literature has vastly reported works on the topic and even produced relevant surveys on the topic e.g. [21], [22]. Methodologies of expert finding can be divided into three categories [4], [7]: Content-Based Approach, Network-Based Approach and Hybrid Approach. On the other hand, other works [10], propose a different taxonomy of existing expert finding systems. These authors indicate that these systems are based on four kinds of expertise indicators: self-disclosed information, authored documents, social network analysis and hybrid techniques [23]. An analysis performed by these authors reveal that hybrid techniques are not combining self-disclosure indicators with social network analysis or document-based indicators. In other words, authors underline that self-disclosure indicators can be seen as isolated indicators that need deeper analysis.

In [9] authors make in-depth review of benchmarking techniques and components that constitute a test collection with special emphasis on error analysis. They also give an overview of different test collections for expert profiling and expert finding. In [24] author reviews more recent examinations of the validity of a test collection approach and evaluation measures as well as he outlines trends in current research exploiting query logs and live labs to finally show that despite its age, this long-standing evaluation method is still a highly valued tool for retrieval research. Furthermore, the Text REtrieval Conference (TREC) or the Yandex Personalized Web

Search Challenge also falls in this area of methods for assessment ad-hoc datasets through train and test processes in different topics. For instance, in 2008, the main topic of the TREC conference was focused on expert finding where a dataset from the Tilburg University¹ was used as input of a competition to find and rank experts. Usually these evaluation methodologies are in charge of asserting the results of an information retrieval process by comparing expected results to actual results and taking into account measures [25] of precision, recall, sensitivity or stability. On the other hand, relevance assessment methods are usually created by a panel of experts, a good number of collections can be found in different domains such as web search, movie/tourism recommendation, medical diagnosis, etc. Finally statistical significance tests are used to estimate the average of system performance according to a set of queries that can be generalized. In this sense, the Wilcoxon test, Student's t test and the Fisher pairwise comparison are common techniques to assess a p-value under a certain degrees of freedom. All these techniques that have been reviewed in [9] and they are relevant to this work due to the fact that a validation of the expected results must be done to assess if the Skillrank algorithm is properly working. To do so the validation section introduces the evaluation method that is a combination of an ad-hoc collection built by a panel expert with measures of precision and recall.

Regarding online skills evaluation, in [26] a technique is introduced to establish a credibility rank (also known as Skillrank) to online profiles based on user's confirmations and six requirements for online skills evaluation. According to these six requirements, a credibility model is defined and populated from on-line profiles. Afterwards, a pilot is implemented to show the functional architecture that supports the on-line evaluation of skills. Nevertheless, the real evaluation of skills and online profiles is still an open issue and, only the architecture is presented. Authors also comment some of the limitations of their approach: 1) skills are evaluated as they are and a scale will be necessary to establish an order of expertise and 2) spread of experience and skills are also under evaluation. On the other hand, they also raise some relevant questions in the evaluation of on-line profiles: *What kind of information can be incorporated to enrich the skill evaluation model?* and *How half-time jobs, activities or tasks can also be included in the evaluation model?* This work is very closely related to the approach presented in this paper. However, they have been focused in the definition of a skill model and a pilot architecture instead of comparing different algorithms working on the same datasets. Other recent works [27] can also be found applying gamification techniques to build online personal skills and boost the learning process. Thus, it is possible to ensure the acquisition of skills from the early stages of learning by analyzing the online behavior and interactions [28] between peers. Finally, other works are also paying attention to the evaluation of online profiles with different purposes such as digital inclusion. As an example, authors in [29] theorize how people's online social networking skills may condition their uses of various digital media for communication.

On the other hand, reputation management systems have emerged to understand the influence of individuals or groups in a certain group. Different metrics with a particular level of effectiveness [30] are applied to assess the reputation in a social network, mailing list or any other collaborative site. For instance the *Stackoverflow* system, a question and answer system [31], uses a simple formula to establish a level of "karma" for each individual depending on their participation in the system. The Research Gate site, a social network for science and research, also establishes a score depending on publications and contributions that you have added to the site. In this case, reputation is passed from researcher to researcher, allowing us to build and leverage our reputation based on anything we choose to contribute. Interactions or activities in this social network will determine our score by looking in our activities (how our peers receive them) but also at who these peers are. Higher scores will be reached as much as higher scores peers interact with our activities, it can be seen as an application of the Spreading Activation technique [32] that has been widely used in information/document retrieval and recommending systems [33].

¹ <http://ilk.uvt.nl/uvv-expert-collection/documentation/documentation.html>

The idea behind of all these reputation management systems lies in a set of internal metrics (they are usually private to avoid fraudulent profiles) that are collected in just one value to create a rank of users by tracking their activity: asking and responding questions, using online feedback of other users (How much a response is better than another?), etc. These systems are also relevant to professional social network sites such as LinkedIn, ResearchGate or Xing in which users try to complete, at the most, their profiles adding own education, professional experience, rewards, publications, etc. as well as feedback from their connections to improve and enrich their profiles. In this sense, talent hunters have got a new way of detecting specialists in a topic by performing advanced search through tools in these websites. Nevertheless the access to this valuable information is commonly restricted and only quantitative information can be found. In this context some works have emerged for topic extraction systems and online reputation management [34] which they use a set of evaluation metrics based on handmade metadata annotation to assess the quality of different factors. Thus trust and provenance analysis is becoming a major challenge to avoid vandalism, fraud, etc. in public profiles, more specifically in user and company profiles. Existing works e.g. [35] are then focused on applying techniques to characterize profiles in reputation management systems to demonstrate through algorithms such as Eigen Trust, TNA-SL or distributed approaches [8] are secure enough to effectively manage trust in communities.

In the field of information retrieval, expertise retrieval [36] and expert finding [37] systems have been widely studied to provide a new approach to tackle the discovery of experts in some area [15]. Currently practices such competitions [38] or challenges are common techniques to retrieve experts based on their performance in a certain topic. The main objective of expertise retrieval [36] lies in the application of existing information retrieval techniques as building blocks to design advanced algorithms that can serve to create content-based links between topics and people. Nevertheless authors have also outlined both applications to other domains such as entity retrieval as well as some conjectures on what the future may hold for expertise retrieval research. For instance last times new novel approaches [39] are emerging to include time and evolution of skills over time as variables in expertise retrieval processes.

As a closely related field to expertise retrieval, community detection [40] in social networks is a widely active research area to segment large communities by a certain criteria. In the specific case of expertise in some topic, these algorithms can be applied to narrow down the search of experts. These techniques can be roughly divided into two groups: 1) global and 2) local approaches. The first ones assume knowledge of the entire network while local ones only assume knowledge of sub-communities featured by some attributes such as location. Global detection algorithms were firstly proposed by Girvan and Newman by iteratively removing edges until the social graph is partitioned (each partition can be consider as a community). The key point of these techniques lies in the selection of the edge to be removed and, in general, some metrics such as betweenness centrality are calculated for each edge. Thus a large social network is divided into high-dense connected communities that share some features and, therefore, can be considered sub-communities. The main drawback of global approaches lies in the necessity of knowing the full graph (it is usually expensive in terms of time and size). In order to decrease the complexity of handling a large graph, local approaches aim for detecting communities in a more scalable and applicable way starting with a set of seed nodes to detect implicit communities. For instance the Clauset's algorithm uses intra-community and inter-community measures to iteratively establish or remove the connection between two nodes. Thus from a starting node communities dynamically emerge. Although these approaches are completely correct to detect underlying communities the use or inferring of dynamic attributes of nodes (users in most of cases) is still an open issue that has been studied in some works such as [41] and it allows a better and more accurate community partition. As a possible application of these aforementioned techniques, the detection of violent communities, hostility or rivalry is currently under study [42] since the internet is understood to be a social space conducive to increased hostility, greater disinhibition and increased social freedom. Moreover, these authors see a link between virtual hostility and actual violence. In social networks research, [43] predicts user personality by mining social

interactions including Aggression-Hostility traits and [44] modeled online social interactions incorporating the effects of hostile interactions.

Finally the relevance of expertise ranking in social networks and Internet has been presented in some works to understand and exploit enterprise know-how [45], find competence gaps and learning needs inside corporations [46], improve Scrum processes [47], improve human to human interactions [48] or tackling information asymmetries in electronic marketplaces [49] to name a few.

In conclusion, a list of methods and techniques for benchmarking has been introduced with the aim of comparing existing approaches to assess personal skills quality. Furthermore, existing works related to reputation management systems and, more specifically, trust and provenance analysis can be also applied to the Skillrank technique since methods to evaluate public profiles are an emerging topic due to the current use of the web. That is why the Skillrank technique seeks for providing an innovative method to assess user profiles from a qualitative point of view through an agnostic technique that can help both talent hunters and managers to exactly know where a capability or skill can be found in their connections or employees with a certain degree of trust and provenance.

3. Skillrank: reinterpreting the SPEAR algorithm to assess skills quality in professional social networks

3.1. Summary of the HITS and SPEAR algorithms

As the previous section has introduced, the HITS algorithm [50] identifies good authorities and hubs for a certain topic by assigning two numbers to a page: an authority and a hub weight where weights are recursively defined. A higher authority weight occurs if the page is pointed to by pages with high hub weights. A higher hub weight occurs if the page points to many pages with high authority weights. More specifically in the context of web search, the HITS algorithm first collects a base document set for each query. After that it recursively calculates the hub and authority values for each document. In order to gather the base document set I , first, a root set R matching the query is fetched from the search engine. Once this root set is configured for each document $r \in R$, a set of documents that point to r another set of documents L' that are pointed to by r are added to the set I as R 's neighborhood. Then for each document $i \in I$, let a_i and h_i be the authority and hub values respectively that are initialized to 1. While the values have not converged, the algorithm iteratively proceeds as follows:

1. For all $i' \in I$ which points to i

$$a_i = \sum_{i'} h_{i'}$$

2. For all $i' \in I$ which is pointed to i

$$h_i = \sum_{i'} a_{i'}$$

3. Normalize a_i and h_i values so that $\sum_i a_i = \sum_i h_i = 1$

A good hub increases the authority weight of the pages it points to. A good authority increases the hub weight of the pages that point to it. The idea is then to apply the two operations above alternatively until equilibrium values for the hub and authority weights are reached. The author also demonstrated that the algorithm will likely converge but the bound on the number of

iterations is unknown (in practice the algorithm converges quickly). New improved versions of this algorithm have emerged such as BHITS by giving a document a default authority weight of $\frac{1}{k}$ if the document is in a group of k documents on a first host which link to a single document on a second host, and a default hub weight of $\frac{1}{l}$ if there are l links from the document on a first host to a set of documents on a second host. Nevertheless and according to its authors, this new version of the algorithm generated bad results when a root link has few in-links but a large number of out-links that are not relevant to the query.

On the other hand the SPEAR algorithm [18] [19] makes use of the HITS definition to introduce the concept of expert, someone with a high level of knowledge, technique or skills in a particular domain. This implies that experts are reliable sources of relevant resources and information but with two main assumptions:

1. Mutual reinforcement of user expertise and document quality. The expertise of an user in a particular domain will depend on the quality of the documents he/she has found. In the same way, quality of documents will depend on the expertise of the user who has found them. This is an issue that has been studied in Psychology and it states that expertise involves the ability of selecting best and relevant information in a certain context. The SPEAR algorithm is based on this assumption and an expert should be someone who selects by quality instead of quantity.
2. Discoverers vs. followers. The second assumption of the SPEAR algorithm lies in the definition of a *discoverer* (expert user that find high-quality and relevant information) vs a *follower* (an user that annotates a document after a *discoverer* does).

Under the aforementioned assumptions the SPEAR algorithm produces a ranking of users with regard to a set of one or more tags. It assumes that a topic of interest is represented by a tag t . The algorithm works as follows [18] [19]:

- Firstly the set of tags R_t is extracted from an underlying folksonomy in a certain social network. Each tag is represented by the tuple $r = (u, t, d, c)$ where u is the user, c is the time when the tag t was assigned to the document d and $c_1 < c_2$ if c_1 refers to an earlier time than c_2 .
- Then the next vectors are defined:
 - a vector $\vec{E} = (e_1, e_2, \dots, e_M)$ containing the expertise scores of users where $M = |U_t|$ is the number of unique users in R_t
 - a vector $\vec{Q} = (q_1, q_2, \dots, q_N)$ containing the quality scores of documents where $N = |D_t|$ is the number of unique documents in R_t
- According to the first assumption, mutual reinforcement refers to the idea that the expertise score of an user depends on the quality scores of the documents to which he tags with t , and the quality score of a document depends on the expertise score of the users who assign tag t to it. Authors define an adjacency matrix A of size $M \times N$ where $A_{i,j} = 1$ if user i has annotated with the tag t the document j , and $A_{i,j} = 0$ otherwise. Based on this matrix, the calculation of expertise and quality scores is an iterative process similar to that of the HITS algorithm: $\vec{E} = \vec{Q} \times A^T$ and $\vec{Q} = \vec{E} \times A$
- On the other hand, the second assumption is implemented by changing the definition of the aforementioned adjacency matrix. Instead of assigning either 0 or 1 (like the HITS algorithm) the following equation is used to populate the initial values of the matrix A :
 - $A_{i,j} = |\{u \mid (u, t, d_j, c), (u_i, t, d_j, c_i) \in R_t \wedge c_i < c\}| + 1$

- Thus, the cell $A_{i,j}$ is equal to 1 plus the number of users who have assigned tag t to document d_j after user u_i . Hence, if u_i is the first to assign t to d_j , $A_{i,j}$ will be equal to the total number of users who have assigned t to d_j . If u_i is the most recent user to assign t to d_j , $A_{i,j}$ will be equal to 1. The effect of such initialization is that the matrix A represents a sorted timeline of any users who tagged a given document d_j .
- The last step is to assign a proper credit score to users by applying a credit scoring function C to each element $A_{i,j}$. According to the authors three different functions could be applied to the matrix A :
 - A linear credit score $C(x) = x$. This function was initially discarded by the authors due to discoverers of a popular document would receive a comparatively higher expertise score although they might have not contributed in any other document thereafter.
 - An increasing function but with a decreasing first derivative to retain the ordering of the scores in A . Authors demonstrated that this kind of function enables the possibility of keeping discoverers score higher than followers but differences between higher scores will be reduced to avoid the undesirable effect of assigning high expertise scores to users who were the first in tagging a few set of popular documents but without further contribution in high-quality documents thereafter. Finally authors selected the function $C(x) = \sqrt{x}$ as credit score for their experiments.

3.2. Skillrank in online-communities

A simplistic definition of an on-line community or social network is a set of $C = \{U, F, DF, R\}$ where U is the set of users that interact with each other, F is a set of static features or attributes, DF is a set of dynamic or inferred attributes that define the community and R is the set of all resources generated by users. More specifically, an user $u_i \in U$ is also described by a set of attributes $u_i = \{S, D\}$ where S is the set of static attributes that describe the user profile and they are usually defined by the own user. On the other hand, D represents a dynamic set of attributes that can be inferred or predicted by tracking the user's activity and interaction in the context of social network C . Furthermore, any social network can be divided into different sub-communities (sub graphs) C_k that are also communities, and by extension, a social network can be also be defined as the union of several sub-communities $C = \bigcup_1^k \{C_k\}$. Formally, let K be an index set, and for each $k \in K$ then the family of sets $\{C_k: k \in K\}$ is the union set that represents an online-community,

$$C = \bigcup_1^k \{ \{U_1, F_1, DF_1, R_1\}, \{U_2, F_2, DF_2, R_2\}, \dots, \{U_k, F_k, DF_k, R_k\} \}.$$

Commonly, the set of users U_k are not disjoint sets, so an user u_i can be a member of several sub-communities. Nevertheless, the set of features F_k , dynamic features DF_k and resources R_k could be shared among sub-communities but they could be also disjoint sets depending of the characteristics of the social network.

Following these definitions, we can describe a social network such as *LinkedIn* containing a sub-community “*MyLinkedIn*” that can be also partitioned in several sub-communities such as “*MyUniversity*” or “*MyWork*”. According to the theoretical model:

- $C_{LinkedIn} = \{U_{LinkedIn}, F_{LinkedIn}, DF_{LinkedIn}, R_{LinkedIn}\}$ where
- $U_{LinkedIn}$ is the set of all registered users.
- $F_{LinkedIn} = \{id = 1, type =$

"professional social network", name="LinkedIn", descriptors = $\{d_1, d_2, \dots, d_k\}$...} is a set of key-value pairs .

- $DF_{LinkedIn} = \{trends = \{t1, t2 \dots t3\}, posts = \{(p1, u1), \dots, (pk, uk)\}, time\}$ is also a set of dynamic key-value pairs in a certain moment $time$.
- On the other hand we can also define this social network by the union of several disjoint sub-communities. Thus $C_{LinkedIn} = \{C_{MyLinkedIn} \cup C_k\}$ where $C_{MyLinkedIn} = \{C_{MyUniversity} \cup C_{MyWork}\}$
- Finally, users will be members of some community so u_{me} represents a *LinkedIn* user that creates the sub-community “*MyLinkedIn*”. This user and its sub-community generate resources such as “posts”, “connections”, “endorsements”, etc. as an example $r_{u_k}^1 = \{user = u_{me}, type = "endorsement", time = timestamp, tag = "skill"\}$ describes the resource in the community “*MyLinkedIn*” that was created by the user u_k using a “endorsement” in a certain moment $time$ on the user u_{me} .

Although communities, users and resources can be described through different static attributes there is still a set of dynamic or behavioral features that must be inferred to make a better description of foreknown communities and to be able to create new inter-community relationships. Since communities, user endorsements, etc. are evolving characteristics, it is necessary to analyze emerging or implicit user’s behaviors [41] [51]. In this sense, the aforementioned community detection algorithms follow a similar approach but studying the structure of the social graph instead of analyzing contents.

Here, we propose the adaptation of the SPEAR algorithm to support the quality assessment of endorsements generated by a sub-community; more specifically the following contexts can be identified:

- Community $C_{MyLinkedIn}$ (Local context). Figure 1 shows that a user u_1 endorses another user u_{me} with the skill “java” in the time t_1 . After that another user u_2 belonging to the same sub-community generated by user u_{me} also uses the same endorsement but at time t_2 where $t_1 < t_2$. The assumption behind this behavior is that after seeing the new endorsement (made by u_1) u_2 also realizes that this endorsement is correct and adds again the same endorsement to the user u_{me} . This situation implies that the first post (see *discoverer* in the SPEAR algorithm) has activated new annotations (see *follower* in the SPEAR algorithm) reinforcing both: 1) the skill “java” in user u_{me} and 2) the initial annotation of the user u_1 . Similarly, if a user u_2 notices that a user u_1 has endorsed another user u_{me} at time t_1 this can lead to an endorsement of user u_1 for the same skill by user u_2 at time t_2 where $t_1 < t_2$ (Figure 2). Finally Figure 3 depicts the situation in which a user is activated by some activity but instead of applying the same tag she uses another tag to annotate knowledge of user u_1 (the one that started the interaction).

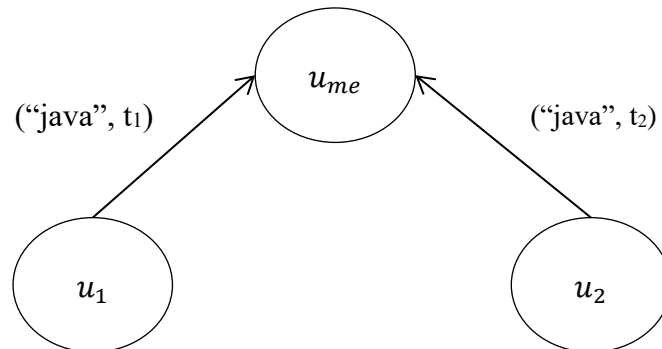


Figure 1 Correlated-endorsements to the same user u_{me} in a sub-community.

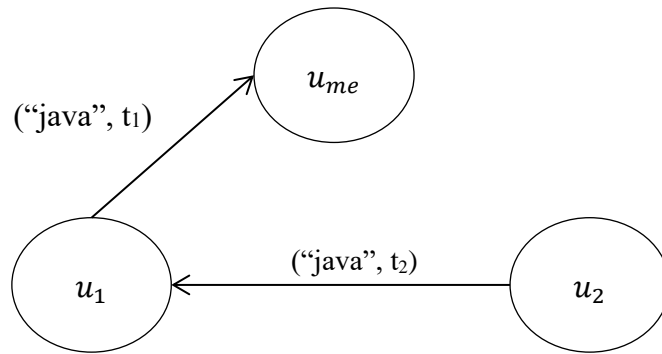


Figure 2 Correlated-endorsements to different users in a sub-community.

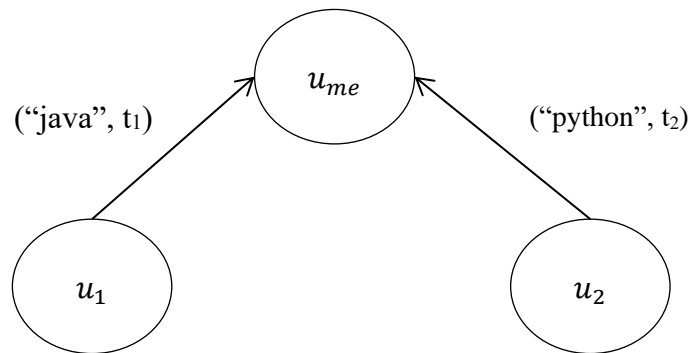


Figure 3 Independent endorsements to the same user u_{me} in a sub-community.

- Community $C_{LinkedIn}$ (Global context). Figure 4 shows that an user u_1 endorses another user u_{me} with the skill “java” in the time t_1 . After that another user u_2 , outside of the sub-community (represented by a dashed circle) also uses the same endorsement but in time t_2 , where $t_1 < t_2$, to endorse user u_1 . The idea behind this behavior is that after seeing the new endorsement (made by u_1) u_2 also realizes that this endorsement can be applied to u_1 . This situation implies that the first post (see *discoverer* in the SPEAR algorithm) has activated new annotations (see *follower* in the SPEAR algorithm) reinforcing the skill of the user u_1 .

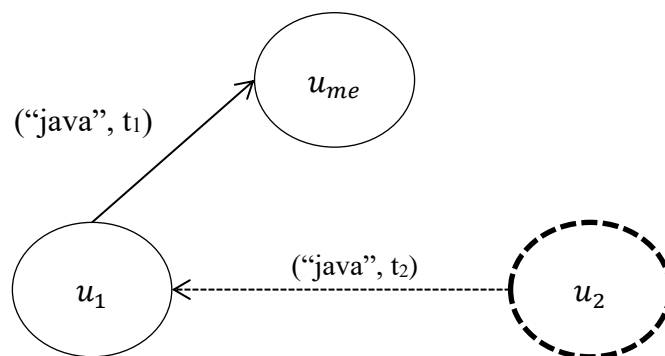


Figure 4 Correlated-endorsements to different users in a community.

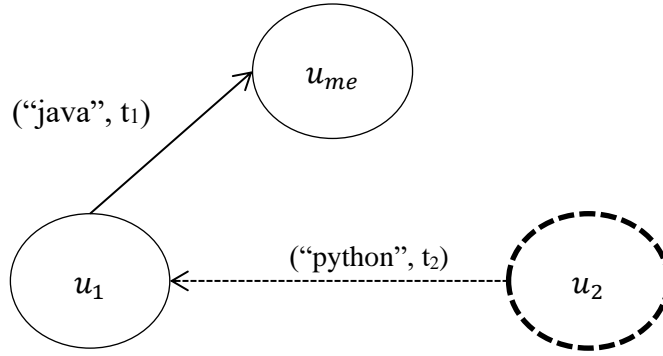


Figure 5 Independent endorsements to the same user u_{me} in a community

On the other hand, Figure 3 and Figure 5 depict a situation in which an user u_1 assigns a skill to user u_{me} in time t_1 but although other user u_2 is activated and assigns another skill, different from the one assigned by u_1 , there is not actually a correlation between them and both assignments can be interpreted as independent endorsements. The Skillrank technique covers the aforementioned scenarios to take advantage of the data delivered by tracking user activities.

Taking into account the inputs required by the SPEAR algorithm, the following vectors are re-defined and the pseudo-code of the algorithm is also presented in **Error! Reference source not found.** :

- The set of skills S_t is extracted from the activities generated by a sub-community C_k in a certain social network. Each endorsement is also represented by the tuple $r = (u_s, s_k, u_t, c)$ where u_s is the source user that endorses with the skill s_k to a target user u_t at time c .
- Then the next vectors are also re-defined:
 - a vector $\vec{E} = (e_1, e_2, \dots, e_M)$ containing the expertise scores of users where $M = |U_t|$ is the number of unique users in C_k
 - a vector $\vec{Q} = (q_1, q_2, \dots, q_N)$ containing the quality scores of skills where $N = |S_t|$ is the number of unique skills in C_k

```

Algorithm 1 Skillrank: Re-interpreting the SPEAR algorithm [18] [19]
Input: Number of users  $M$ 
Input: Number of skills  $N$ 
Input: A set of skills  $S_t \in C_k = \{(u_s, s_k, u_t, c)\}$ 
Input: Credit scoring function  $C$  (the same as in the standard SPEAR)
Input: Number of iterations  $k$ 
Output: A list  $L$  of users.
1: Set  $\vec{E}$  to be the vector  $(1, 1, \dots, 1) \in Q^M$ 
2: Set  $\vec{Q}$  to be the vector  $(1, 1, \dots, 1) \in Q^N$ 
3:  $A \leftarrow$  Generate Adjacency Matrix( $S_t, C$ )
4: for  $i = 1$  to  $k$  do
5:    $\vec{E} \leftarrow \vec{Q} \times A^t$ 
6:    $\vec{Q} \leftarrow \vec{E} \times A$ 
7:   Normalize  $\vec{E}$ 
8:   Normalize  $\vec{Q}$ 
9: end for
10:  $L \leftarrow$  Sort users by their expertise scores in  $\vec{E}$  and quality skills
    scores in  $\vec{Q}$ 
11: return  $L$ 

```

Listing 1 Skillrank pseudo-code. Re-interpreting the SPEAR algorithm [18] [19].

According to these definitions the unique difference between the original version of the SPEAR algorithm and this new version seems to be the naming of elements (“document” by “skill”). Nevertheless, the Skillrank facilitates a two-step process to run the SPEAR algorithm in C before C_k with the aim of populating both vectors \vec{E} and \vec{Q} with real values. Hence, a new interpretation of the adjacency matrix can be done. Instead of considering the adjacency matrix for the whole social network or folksonomy, each user will generate an adjacency matrix in which rows represent connections and columns skills respectively (see Table 1). The interpretation of this table is as follows: 0 represents that the user u_k have not yet endorsed using the skill s_k while another value such 2 in cell (u_1, s_2) and 1 in cell (u_2, s_1) represents that user u_1 used the skill s_2 before user u_2 .

	s_1	s_2	...	s_N
u_1	0	2	0	2
u_2	2	1	3	1
...	3	5	0	4
u_M	4	2	1	0

Table 1 Example of generated adjacency matrix for an user u .

On the other hand, an analysis of the temporal and spatial complexity of the algorithm can be carried out to show the computational complexity of the technique depicted in **Error! Reference source not found.** Firstly and regarding the spatial complexity, the algorithm makes use of a set of skills S_t which contains all skill endorsements registered for a community C_k , the aforementioned vectors \vec{E} and \vec{Q} and a matrix A of dimensions: $M = |U_t|$, number of users in a community C_k and $N = |S_t|$, number of unique skills in a community C_k . Secondly and regarding the temporal complexity, the standard SPEAR algorithm performs k iterations containing a main operation (matrix multiplication) two times, an operation to transpose a matrix and two vector normalizations. Assuming that the adjacency matrix, A , has dimension $M \times N$, the computation complexity of the algorithm can be expressed through the next expression using the *Big-O* notation: $O(k [2 * (M * N * N) + (M * N) + M + N]) = O(k [M * N^2]) = O(M * N^2)$.

4. The case study

To illustrate the performance in terms of precision and recall of the presented algorithms, HITS and SPEAR, with regards to the adaptation designed in Skillrank a case study using different datasets is provided. Here, the evaluation of performance is not a mere question since there is a lack of real datasets containing the required information of users, connections, skills and time. To mitigate this problem, a synthetic dataset, as the basis of the experiment, has been designed after collecting real data from the LinkedIn API (currently, this API provides access to valuable but incomplete information that must be fixed by the own users). Thus, simulated communities, users, skills and endorsements are generated to study the behavior of the different approaches. To carry out both experiments the following steps have been carried out:

1. Select and prepare dataset. For every dataset to be evaluated a set of tuples in the form $r = (u_s, s_k, u_t, c)$ must be provided.
2. Create a dataset for unit testing purposes. To do so a panel of experts has established a category, using an official competence scale [53], for every user and skill.
3. Definition of precision and recall. In order to calculate both measures, next definitions are also required (given an user):
 - True positives (tp): “number of skills that were expected to reach a certain level of quality”

- False positives (fp): “number of skills that have reached a different level of quality”
- True negative (tn): “number of skills that were not expected to reach a certain level of quality”.
- False negative (fn) “number of skills that have not reached a different level of quality”

Once we have the aforementioned definitions, precision and recall can be defined and calculated as follows:

- Precision is defined as “*the number of user skills that have reached the proper level of quality established by the panel of experts*”.

$$Precision = P = \frac{tp}{tp + fp}$$

- Recall is defined as “*the number of user skills that have not reached the proper level of quality established by the panel of experts*”.

$$Recall = R = \frac{tp}{tp + fn}$$

- F_1 score is then defined as:

$$F_1 = 2 * \frac{Precision * Recall}{Precision + Recall}$$

4. Inclusion of the frequency as a basic technique for each user and skill. Given an user u_k and a skill s_k the quality of the skill is calculated as follows:

$$F_{u_k}^{s_k} = \frac{\text{number of tuples } (u_s, s_k, u_k, c)}{\text{number of connections of } u_k}$$

5. Run the experiment for every dataset and technique.
6. Configure all techniques with default parameters (credit score function, etc.) as previous section has presented.
7. Extract measures of precision, recall and F_1 by comparing expected results to real results.

4.1. Design of the experiment

The first step to run the experiments lies in the proper creation of a synthetic dataset inspired by real data extracted from the LinkedIn API. To do so a community $C_{synthetic}$ must be modeled including the required input parameters for the target algorithms. Thus a set of users, $U_{synthetic}$ containing 10 different profiles has been designed including an average between 30-50 connections per user (these values have been inferred from the real data). On the other hand, a set of skills S_t , see Table 2, must be also designed according to next features: 1) technical, professional and management skills must be available for each user and 2) all skills must be, at least, in one profile but no all profiles contain all skills. Finally, a set of endorsements in the form $r = (u_s, s_k, u_t, c)$ are generated from each user being u_s the user/connection that assigns the skill s_k to the user u_t in a certain time c .

Once the input dataset is designed, it is necessary to create a dataset containing the expected results with the aim of performing automatic unit testing. To do so, a panel experts that has already participated easing the access to their profiles in LinkedIn, has also established for their real connections a level of expertise for each skill in S_t . Thus, this dataset contains a set of tuples in the form $r_k = (u_k, s_k, l_k)$ where u_k is an user with a level of competence l_k on the skill s_k . The different levels of competence have been taken from [53] in which authors present “The Individual Competency Index (ICI)”, see Table 3.

Id	Skill
s_1	Java
s_2	Python
s_3	Data mining
s_4	UML
s_5	MySQL
s_6	CMMI
s_7	Sales Management
s_8	Negotiation
s_9	Technical Management
s_{10}	Business Management

Table 2 Set of selected skills S_t .

Id	Conceptual knowledge	Description
l_0	None	Level 0 denotes a lack of competence in a specific area or topic.
l_1	Basic	Level 1 denotes an understanding of fundamentals and some initial practical application.
l_2	Intermediate	Level 2 denotes a solid conceptual understanding and some practical application.
l_3	Advanced	Level 3 denotes significant conceptual knowledge and practical experience in performing a competency to a consistently high standard.
l_4	Expert	Level 4 denotes extensive knowledge, refined skill and prolonged experience in performing a defined competency at the highest standard.

Table 3 The Individual Competency Index (ICI).

After the creation of the input and test datasets, the algorithms and unit tests can be executed to finally extract the measures of precision, recall and F_1 and compare the different techniques. Last step involves the creation of a function to convert numerical values into a level of expertise. To do so, a percentile rank for every level of expertise is defined. The aforementioned steps have been also followed to perform the same experiment on the LinkedIn dataset. As a final remark it is relevant to discuss some research limitations that have emerged during the creation of both datasets.

- The use of the LinkedIn API is restricted and it is not possible to access all information that is available through the public website. Thus some relevant information with regards to the skills is missing such as who has endorsed someone. To overcome this issue our panel of experts and collaborators were asked to complete this information.
- Another issue in the use of the LinkedIn API lies in the lack of time for each endorsement. This is a critical point since algorithms are based in this assumption. To overcome this issue we have follow two strategies: 1) ask the panel of experts and collaborators to estimate a date in which the endorsements were created and 2) estimate the time of the endorsement by using the join data in the social network.
- The LinkedIn API also provides a level of proficiency for each skill. Nevertheless these features cannot be used since it is not available in all skills and it is based on a particular taxonomy.

- Finally, in order to access all required information, an URL to query the official LinkedIn REST API² was designed, see Listing 1. Nevertheless and due to privacy setting of the API, every participant in the experiments was asked to execute this request through the APIgee service³ using their own OAuth credentials and to send us the request's output as XML. Through this request the user will be asked to grant access to their full profile. Then, all public personal information, a profile containing: first name, last name, headline, industry, location, number of connections, summary, specialties, positions, associations, honors, interests, publications, patents, languages, skills (id, proficiency and years), certifications, education, courses, volunteer, three-current-positions, number of recommenders, connections (and their full profile) will be gathered using the LinkedIn REST API.

[https://api.linkedin.com/v1/people/~connections:\(id,first-name,last-name,formatted-name,email-address,headline,industry,location,num-connections,summary,specialties,positions,site-standard-profile-request,public-profile-url,api-standard-profile-request,proposal-comments,associations,honors,interests,publications,patents,languages,skills:\(id,skill,proficiency,years\),certifications,educations,courses,volunteer,three-current-positions,num-recommenders,following,job-bookmarks,date-of-birth,member-url-resources,connections\)](https://api.linkedin.com/v1/people/~connections:(id,first-name,last-name,formatted-name,email-address,headline,industry,location,num-connections,summary,specialties,positions,site-standard-profile-request,public-profile-url,api-standard-profile-request,proposal-comments,associations,honors,interests,publications,patents,languages,skills:(id,skill,proficiency,years),certifications,educations,courses,volunteer,three-current-positions,num-recommenders,following,job-bookmarks,date-of-birth,member-url-resources,connections))

Listing 1 URL to extract user data from the LinkedIn API.

4.2. Results and discussion

After the execution of the different techniques the averaged measures (for all skills) and for every user in dataset $C_{synthetic}$ are presented in Table 4. Obviously, the first technique based on the number of times a user has been endorsed is not actually relevant in terms of quality as results show. On the other hand, the HITS algorithm provides better results in terms of precision but the drawbacks of this algorithm (not considering time as a relevant variable - see Section 3.1) implies a low-precision in some users with a behavior close to the frequency-based technique. As an improvement or more accurate version of the HITS algorithm, the SPEAR technique seems to get better results that are closer to the expert's opinion. Here, it is clear that the assumption of time as a key-variable to assess quality is a determinant to detect the level of expertise. Finally, the Skillrank technique that previously configures the level of expertise of every user before making endorsement seems to have a similar behavior to the SPEAR algorithm. Although in some cases there is a relevant gain, the truth is that values in both techniques are very similar and to actually assert that Skillrank is better than the simple version of the SPEAR technique more data should be used.

User	$F_{u_k}^{Sk}$			HITS			SPEAR			Skillrank		
	P	R	F_1	P	R	F_1	P	R	F_1	P	R	F_1
u_1	0.43	0.80	0.56	0.67	0.79	0.73	0.73	0.87	0.79	0.80	0.81	0.80
u_2	0.25	0.89	0.39	0.65	0.78	0.71	0.90	0.71	0.79	0.65	0.72	0.68
u_3	0.44	0.83	0.58	0.63	0.82	0.71	0.82	0.80	0.81	0.85	0.86	0.85
u_4	0.52	0.74	0.61	0.63	0.86	0.73	0.83	0.86	0.84	0.86	0.77	0.81
u_5	0.48	0.84	0.61	0.56	0.79	0.66	0.67	0.84	0.75	0.77	0.77	0.77
u_6	0.43	0.79	0.56	0.53	0.78	0.63	0.89	0.72	0.80	0.84	0.89	0.86
u_7	0.35	0.71	0.47	0.53	0.89	0.66	0.82	0.82	0.82	0.86	0.85	0.85
u_8	0.46	0.84	0.59	0.74	0.88	0.80	0.80	0.86	0.83	0.80	0.84	0.82
u_9	0.45	0.73	0.56	0.74	0.87	0.80	0.82	0.84	0.83	0.66	0.90	0.76
u_{10}	0.29	0.77	0.42	0.68	0.89	0.77	0.69	0.90	0.78	0.78	0.78	0.78

² <https://developer.linkedin.com/apis>

³ <https://apigee.com/console/linkedin>

Average	0,41	0,79	0,54	0,64	0,84	0,72	0,80	0,82	0,80	0,79	0,82	0,80
----------------	-------------	-------------	-------------	-------------	-------------	-------------	-------------	-------------	-------------	-------------	-------------	-------------

Table 4 Aggregated measures S_t in dataset $C_{synthetic}$.

Following this discussion, Table 5 shows the results of the different techniques using real data. In general, a decrease of precision can be found in this table with regards to previous results. This can be explained due to the fact that this dataset is not customized and real behavior of users and skills is found implying, in general, worse results.

As final remark, a change in the parameters such as the set of users and skills could lead us to get better results. Nevertheless, this initial effort will be used as a baseline to compare further improvements. Regarding similar approaches that have been implemented, as the Related work section has outlined, the presented approach is closely related to [26] since the same problem is being addressed. The main difference is that they have also outlined a functional architecture while we focus here in addressing some of the existing open issues: alignment of the skills quality to an existing competency index. On the other hand, we have tried to re-use at the most existing techniques. That is why, we have made use of the well-known techniques such as the HITS and SPEAR algorithms that have been demonstrated to detect experts under certain characteristics of a graph. The Skillrank technique gets inspiration of these techniques to adapt the underlying concepts and execution steps to the problem of quality assessment of skills available in online profiles.

User	$F_{u_k}^{S_k}$			HITS			SPEAR			Skillrank		
	<i>P</i>	<i>R</i>	<i>F₁</i>	<i>P</i>	<i>R</i>	<i>F₁</i>	<i>P</i>	<i>R</i>	<i>F₁</i>	<i>P</i>	<i>R</i>	<i>F₁</i>
u_1	0.39	0.68	0.50	0.53	0.68	0.60	0.66	0.71	0.68	0.79	0.79	0.79
u_2	0.22	0.80	0.35	0.59	0.69	0.64	0.56	0.83	0.67	0.77	0.75	0.76
u_3	0.40	0.60	0.48	0.40	0.70	0.51	0.58	0.87	0.70	0.78	0.89	0.83
u_4	0.23	0.80	0.36	0.55	0.60	0.57	0.75	0.70	0.72	0.67	0.71	0.69
u_5	0.40	0.60	0.48	0.43	0.76	0.55	0.74	0.82	0.78	0.84	0.89	0.86
u_6	0.33	0.69	0.45	0.63	0.68	0.65	0.53	0.77	0.63	0.67	0.89	0.76
u_7	0.50	0.78	0.61	0.67	0.62	0.64	0.64	0.79	0.71	0.71	0.76	0.73
u_8	0.27	0.60	0.37	0.60	0.73	0.66	0.52	0.88	0.65	0.75	0.87	0.81
u_9	0.30	0.67	0.41	0.66	0.69	0.67	0.81	0.80	0.80	0.81	0.84	0.82
u_{10}	0.45	0.80	0.58	0.63	0.76	0.69	0.72	0.89	0.80	0.73	0.81	0.77
Average	0,35	0,70	0,46	0,57	0,69	0,62	0,65	0,81	0,71	0,75	0,82	0,78

Table 5 Aggregated measures S_t in dataset $C_{LinkedIn}$.

5. Conclusions and future work

The present paper has introduced different techniques to assess quality in graph-based structures. The well-known algorithms HITS and SPEAR have been also presented as inspiration for the Skillrank technique. This approach reinterprets the notions and underlying concepts of the SPEAR algorithm to apply them to the context of skills quality assessment in professional social networks. On the other hand two main experiments have been conducted using synthetic and real data to evaluate the behavior of the aforementioned techniques in terms of precision and recall. Both approaches – the SPEAR and Skillrank algorithms – have shown similar results in the test datasets implying that these techniques can be meaningfully applied to assess quality of skills. From another perspective the quality assessment of user profiles and more specifically user skills is an active research area that ranges from applying expertise retrieval techniques to expertise profiling, topic extraction, etc. In this sense there are still some open issues that must be tackled

in order to provide automatic methods for user profiling, talent hunter or expert finding processes. Currently a lot of professional social networks are emerging but the problem of creating groups of users by a certain topic is becoming a major challenge since it is necessary to improve trust and provenance of information or user's activities. The relevance of this work is that it can serve to manage enterprise know-how and to detect experts inside organizations. Their competitiveness can be increased by better human resources management processes that could take advantage of exploiting existing data generated by tracking user's activities. From a technical point of view Skillrank is a first substantial effort (including a good number of "artisan" tasks) and new capabilities such as including new data sources to assess skills quality, use of more advanced ordered weighted averaging (OWA) operators and adaptation of other datasets for experimenting purposes should be added in the future as well as new variables in the core of the algorithm. Finally, we also plan to release the information of our experiments using some existing standard such as nanopublications to ease the reuse and comparison with new techniques.

Acknowledgments

This work has been partially supported by the European Commission (programme LifeLong Learning – action Leonardo da Vinci – Transfer of Innovation) through project "ECQA Certified Social Media Networker Skills" (2011-1-ES1_LEO05-35930), by the Spanish Ministry of Economy and Competitiveness through INNFACTO project Post-Via 2.0 (IPT-2011-0973-410000), by the German Federal Ministry of Education and Research through project "OpSIT - Optimaler Einsatz von Smart-Items-Technologien in der Stationären Pflege" (16SV6048), and by the German Federal Ministry of Economics and Technology through project "PreventAB" (KF3144902DB3).

References

- [1] F. Bonchi, C. Castillo, A. Gionis, and A. Jaimes, "Social Network Analysis and Mining for Business Applications," *ACM Trans Intell Syst Technol*, vol. 2, no. 3, pp. 22:1–22:37, May 2011.
- [2] M. Neshati, D. Hiemstra, E. Asgari, and H. Beigy, "Integration of scientific and social networks," *World Wide Web*, vol. 17, no. 5, pp. 1051–1079, Sep. 2014.
- [3] K. Musial and P. Kazienko, "Social networks on the Internet," *World Wide Web*, vol. 16, no. 1, pp. 31–72, 2013.
- [4] X. Tang and C. C. Yang, "Ranking User Influence in Healthcare Social Media," *ACM Trans Intell Syst Technol*, vol. 3, no. 4, pp. 73:1–73:21, Sep. 2012.
- [5] T. Schleyer, B. S. Butler, M. Song, and H. Spallek, "Conceptualizing and advancing research networking systems," *ACM Trans Comput-Hum Interact*, vol. 19, no. 1, pp. 2:1–2:26, May 2012.
- [6] R. Colomo-Palacios, E. Tovar-Caro, Á. García-Crespo, and J. M. Gómez-Berbís, "Identifying technical competences of IT Professionals: the case of software engineers," *Int. J. Hum. Cap. Inf. Technol. Prof.*, vol. 1, no. 1, pp. 31–43, 2010.
- [7] Y. Xu, X. Guo, J. Hao, J. Ma, R. Y. K. Lau, and W. Xu, "Combining social network and semantic concept analysis for personalized academic researcher recommendation," *Decis. Support Syst.*, vol. 54, no. 1, pp. 564–573, Dec. 2012.
- [8] D. W. McDonald and M. S. Ackerman, "Expertise recommender: a flexible recommendation system and architecture," in *Proceedings of the 2000 ACM conference on Computer supported cooperative work*, New York, NY, USA, 2000, pp. 231–240.
- [9] R. Berendsen, M. de Rijke, K. Balog, T. Bogers, and A. van den Bosch, "On the assessment of expertise profiles," *J. Am. Soc. Inf. Sci. Technol.*, vol. 64, no. 10, pp. 2024–2044, 2013.

- [10] G. A. Wang, J. Jiao, A. S. Abrahams, W. Fan, and Z. Zhang, "ExpertRank: A topic-aware expert finding algorithm for online knowledge communities," *Decis. Support Syst.*, vol. 54, no. 3, pp. 1442–1451, Feb. 2013.
- [11] R. Colomo-Palacios, I. González-Carrasco, J. L. López-Cuadrado, A. Trigo, and J. E. Varajao, "I-Competere: Using applied intelligence in search of competency gaps in software project managers," *Inf. Syst. Front.*, vol. 16, no. 4, pp. 607–625, Sep. 2014.
- [12] A. Capiluppi, A. Serebrenik, and L. Singer, "Assessing Technical Candidates on the Social Web," *IEEE Softw.*, vol. 30, no. 1, pp. 45–51, 2013.
- [13] M. Taylor and D. Richards, "Finding and validating expertise.," presented at the 19th European Conference on Information Systems, ECIS 2011, Helsinki, Finland, 2011.
- [14] E. Ben Ahmed, A. Nabli, and F. Gargouri, "Group extraction from professional social network using a new semi-supervised hierarchical clustering," *Knowl. Inf. Syst.*, vol. 40, no. 1, pp. 29–47, 2014.
- [15] V. Boeva, M. Krusheva, and E. Tsiporkova, "Measuring expertise similarity in expert networks," in *Intelligent Systems (IS), 2012 6th IEEE International Conference*, 2012, pp. 053–057.
- [16] L. Page, S. Brin, R. Motwani, and T. Winograd, "The PageRank Citation Ranking: Bringing Order to the Web," Stanford Digital Library Technologies Project, 1998.
- [17] J. M. Kleinberg, "Hubs, authorities, and communities," *ACM Comput. Surv.*, vol. 31, no. 4es, p. 5–es, Dec. 1999.
- [18] C. A. Yeung, M. G. Noll, N. Gibbins, C. Meinel, and N. Shadbolt, "SPEAR: SPAMMING-RESISTANT EXPERTISE ANALYSIS AND RANKING IN COLLABORATIVE TAGGING SYSTEMS," *Comput. Intell.*, vol. 27, no. 3, pp. 458–488, Aug. 2011.
- [19] M. G. Noll, C. Au Yeung, N. Gibbins, C. Meinel, and N. Shadbolt, "Telling experts from spammers: expertise ranking in folksonomies," 2009, p. 612.
- [20] B. Hoppe and C. Reinelt, "Social network analysis and the evaluation of leadership networks," *Leadersh. Q.*, vol. 21, no. 4, pp. 600 – 619, 2010.
- [21] T. Lappas, K. Liu, and E. Terzi, "A Survey of Algorithms and Systems for Expert Location in Social Networks," in *Social Network Data Analytics*, C. C. Aggarwal, Ed. Springer US, 2011, pp. 215–241.
- [22] X. Liu, G. A. Wang, A. Johri, M. Zhou, and W. Fan, "Harnessing global expertise: A comparative study of expertise profiling methods for online communities," *Inf. Syst. Front.*, vol. 16, no. 4, pp. 715–727, Sep. 2014.
- [23] J. J. Jung and P. Kazienko, "Advances on Social Network Applications," *J UCS*, vol. 18, no. 4, pp. 454–456, 2012.
- [24] M. Sanderson, "Test Collection Based Evaluation of Information Retrieval Systems," *Found. Trends® Inf. Retr.*, vol. 4, no. 4, pp. 247–375, 2010.
- [25] K. Hofmann, S. Whiteson, and M. D. Rijke, "Fidelity, Soundness, and Efficiency of Interleaved Comparison Methods," *ACM Trans. Inf. Syst.*, vol. 31, no. 4, pp. 1–43, Nov. 2013.
- [26] T. Haselmann, A. Winkelmann, and G. Vossen, "Towards a Conceptual Model for Trustworthy Skills Profiles in Online Social Networks," in *Information Systems Development*, J. Pokorny, V. Repa, K. Richta, W. Wojtkowski, H. Linger, C. Barry, and M. Lang, Eds. New York, NY: Springer New York, 2011, pp. 285–296.
- [27] M.-E. Del-Moral Pérez and A.-P. Guzmán-Duque, "CityVille: collaborative game play, communication and skill development in social networks," *J. New Approaches Educ. Res.*, vol. 3, no. 1, pp. 11–19, Jan. 2014.
- [28] A. Cabrales, A. Calvó-Armengol, and Y. Zenou, "Social interactions and spillovers," *Games Econ. Behav.*, vol. 72, no. 2, pp. 339–360, 2011.
- [29] Y. P. Hsieh, "Online social networking skills: The social affordances approach to digital inequality," *First Monday*, vol. 17, no. 4, 2012.
- [30] G. E. Bolton, E. Katok, and A. Ockenfels, "How effective are electronic reputation mechanisms? An experimental investigation," *Manag. Sci.*, vol. 50, no. 11, pp. 1587–1602, 2004.

- [31] D.-R. Liu, Y.-H. Chen, W.-C. Kao, and H.-W. Wang, "Integrating expert profile, reputation and link analysis for expert finding in question-answering websites," *Inf. Process. Manag.*, vol. 49, no. 1, pp. 312–329, 2013.
- [32] J. M. Alvarez-Rodríguez, J. E. L. Gayo, and P. O. de Pablos, "An Extensible Framework to Sort out Nodes in Graph-Based Structures Powered by the Spreading Activation Technique: The ONTOSPREAD Approach," *IJKSR*, vol. 3, no. 4, pp. 57–71, 2012.
- [33] J. M. Alvarez-Rodríguez, L. Polo, W. Jimenez, P. Abella, and J. E. L. Gayo, "Application of the spreading activation technique for recommending concepts of well-known ontologies in medical systems," in *Proceedings of Bioinformatics, Computational Biology and Biomedicine (BCB)*, 2011, pp. 626–635.
- [34] E. Amigó, D. Spina, B. Beotas, and J. Gonzalo, "Towards an Automatic Evaluation for Topic Extraction Systems for Online Reputation Management," in *NyNaK*, 2010.
- [35] A. G. West, A. J. Aviv, J. Chang, V. S. Prabhu, M. Blaze, S. Kannan, I. Lee, J. M. Smith, and O. Sokolsky, "QuanTM: a quantitative trust management system," in *EUROSEC*, 2009, pp. 28–35.
- [36] K. Balog, Y. Fang, M. de Rijke, P. Serdyukov, and L. Si, "Expertise Retrieval," *Found. Trends® Inf. Retr.*, vol. 6, no. 2–3, pp. 127–256, 2012.
- [37] A. Bozzon, M. Brambilla, S. Ceri, M. Silvestri, and G. Vesci, "Choosing the right crowd: expert finding in social networks," in *EDBT*, 2013, pp. 637–648.
- [38] Ç. Aslay, N. O'Hare, L. M. Aiello, and A. Jaimes, "Competition-based Networks for Expert Finding," in *Proceedings of the 36th International ACM SIGIR Conference on Research and Development in Information Retrieval*, New York, NY, USA, 2013, pp. 1033–1036.
- [39] J. Rybak, K. Balog, and K. Nørnvåg, "Temporal Expertise Profiling," in *Proceedings of the 36th European conference on Advances in Information Retrieval*, 2014.
- [40] S. Fortunato and A. Lancichinetti, "Community detection algorithms: a comparative analysis: invited presentation, extended abstract," in *VALUETOOLS*, 2009, p. 27.
- [41] A. Mislove, B. Viswanath, P. K. Gummadi, and P. Druschel, "You are who you know: inferring user profiles in online social networks," in *WSDM*, 2010, pp. 251–260.
- [42] D. Upton Patton, R. D. Eschmann, and D. A. Butler, "Internet banging: New trends in social media, gang violence, masculinity and hip hop," *Comput. Hum. Behav.*, vol. 29, no. 5, pp. A54–A59, Sep. 2013.
- [43] A. Ortigosa, R. M. Carro, and J. I. Quiroga, "Predicting user personality by mining social interactions in Facebook," *J. Comput. Syst. Sci.*, vol. 80, no. 1, pp. 57–71, Feb. 2014.
- [44] C.-C. Musat, B. Faltings, and P. Roussille, "A Model of Online Social Interactions Based on Sentiment Analysis and Content Similarity," *HUMAN*, vol. 2, no. 1, pp. pp. 55–66, Sep. 2013.
- [45] R. Colomo-Palacios, C. Casado-Lumbreras, P. Soto-Acosta, F. J. G. Peñalvo, and E. T. Caro, "Competence gaps in software personnel: A multi-organizational study," *Comput. Hum. Behav.*, vol. 29, no. 2, pp. 456–461, 2013.
- [46] F. J. G. Peñalvo, R. Colomo-Palacios, and M. D. Lytras, "Informal learning in work environments: training with the Social Web in the workplace," *Behav. IT*, vol. 31, no. 8, pp. 753–755, 2012.
- [47] R. Colomo-Palacios, I. Gonzalez-Carrasco, J. L. López-Cuadrado, and Á. García-Crespo, "ReSySTER: A hybrid recommender system for Scrum team roles based on fuzzy and rough sets," *Appl. Math. Comput. Sci.*, vol. 22, no. 4, pp. 801–816, 2012.
- [48] P. Kazienko, N. Szozda, T. Filipowski, and W. Blysz, "New business client acquisition using social networking sites," *Electron. Mark.*, vol. 23, no. 2, pp. 93–103, 2013.
- [49] V. Stantchev and G. Tamm, "Reducing Information Asymmetry in Cloud Marketplaces," *Int. J. Hum. Cap. Inf. Technol. Prof. IJHCITP*, vol. 3, no. 4, pp. 1–10, 2012.
- [50] L. Li, Y. Shang, and W. Zhang, "Improvement of HITS-based Algorithms on Web Documents," in *Proceedings of the 11th International Conference on World Wide Web*, New York, NY, USA, 2002, pp. 527–535.
- [51] A. Hannak, P. Sapiezynski, A. M. Kakhki, B. Krishnamurthy, D. Lazer, A. Mislove, and C. Wilson, "Measuring personalization of web search," in *WWW*, 2013, pp. 527–538.

- [52] A. Mislove, B. Viswanath, P. K. Gummadi, and P. Druschel, "You are who you know: inferring user profiles in online social networks," in *WSDM*, 2010, pp. 251–260.
- [53] B. Succar, W. Sher, and A. Williams, "An integrated approach to BIM competency assessment, acquisition and application," *Autom. Constr.*, vol. 35, pp. 174–189, Nov. 2013.