



CorpWiki: A self-regulating wiki to promote corporate collective intelligence through expert peer matching

Ioanna Lykourantzou^{a,*}, Katerina Papadaki^a, Dimitrios J. Vergados^a, Despina Polemi^b, Vassili Loumos^a

^aSchool of Electrical and Computer Engineering, National Technical University of Athens, Zographou Campus, 15773 Athens, Greece

^bSchool of Informatics, University of Piraeus, Karaoli and Dimitriou 80, 18532 Piraeus, Greece

ARTICLE INFO

Article history:

Received 1 October 2008

Received in revised form 14 July 2009

Accepted 4 August 2009

Keywords:

Web 2.0

Wiki

Collective intelligence

Expert peer matching

Feed-forward neural networks

ABSTRACT

One of the main challenges that organizations face nowadays, is the efficient use of individual employee intelligence, through machine-facilitated understanding of the collected corporate knowledge, to develop their collective intelligence. Web 2.0 technologies, like wikis, can be used to address the above issue. Nevertheless, their application in corporate environments is limited, mainly due to their inability to ensure knowledge creation and assessment in a timely and reliable manner. In this study we propose CorpWiki, a self-regulating wiki system for effective acquisition of high-quality knowledge content. Inserted articles undergo a quality assessment control by a large number of corporate peer employees. In case the quality is inadequate, CorpWiki uses a novel expert peer matching algorithm (EPM), based on feed-forward neural networks, that searches the human network of the organization to select the most appropriate peer employee who will improve the quality of the article. Performance evaluation results, obtained through simulation modeling, indicate that CorpWiki improves the final quality levels of the inserted articles as well as the time and effort required to reach them. The proposed system, combining machine-learning intelligence with the individual intelligence of peer employees, aims to create new inferences regarding corporate issues, thus promoting the collective organizational intelligence.

© 2009 Elsevier Inc. All rights reserved.

1. Introduction

The corporate collective knowledge, being a prerequisite for informed decision making, is considered to be a corporate asset and has become a strategic priority for organizations. Corporate knowledge, as a fluid mix of ideas, experience, intuition and lessons learned, does not exist only in information repositories but mainly resides in the minds of individuals [22]. It is the harnessing of this knowledge and the creation of collective intelligence [43], which really provides the enterprise with the ability to innovate and address the challenges of competition. In order to address the above issue of creating innovation and competing for excellence, organizations adopt various coordination approaches regarding collective knowledge creation, dissemination and usage.

In the hierarchy-based coordination approach [55], knowledge is considered to be concentrated in specialized units consisting of selected experts of the organization. The expertise of these individuals is a valuable asset for the organization, especially in cases of dealing with tasks that require specific predefined expertise [33]. However, the most challenging tasks,

* Corresponding author.

E-mail addresses: ioanna@medialab.ntua.gr (I. Lykourantzou), apapadak@mail.ntua.gr (K. Papadaki), djvergad@telecom.ntua.gr (D.J. Vergados), dpolemi@unipi.gr (D. Polemi), loumos@cs.ntua.gr (V. Loumos).

mainly related to the strategic plans affecting the whole organization, require the combination of various kinds of knowledge and broader expertise. In these cases, the hierarchically-coordinated expert consultancy presents specific limitations. First, since expertise is inescapably narrow, the experts of a specialized unit can focus only on a few aspects of knowledge and not on something as broad as the whole organizational information. Furthermore, expert consultancy may be accompanied by high latency in knowledge acquisition [75]. This is due to the fact that in hierarchical-based coordination, experts perform either isolated or in the limits of their specialized functional units and the necessary combination of knowledge is performed by moving up the hierarchy [33]. This practice inserts a considerable delay between the creation of knowledge and the moment it becomes available. Therefore, in cases of innovative and creative tasks where a broader expertise is required, the hierarchically-coordinated expert consultancy approach may limit the innovation capabilities and responsiveness of the organization.

As an alternative to the hierarchically-coordinated expert consultancy, organizations may attempt to adopt a more flat and open coordination approach, to harness the knowledge of their employees and make decisions based on the idea of the “wisdom of the crowds” [69]. According to this idea, the decisions of a diverse and large enough group of people will, over time, be intellectually superior to those of isolated expert individuals. There are four conditions to a wise crowd [69]. The first refers to diversity, which means that having a group of people with different backgrounds and perspectives allows the group to conceptualize problems in novel ways [78]. The second condition is independence, which means that if people in a group have relative freedom from each others’ influence it is far more likely to reach to a good decision. The third condition is decentralization, which enables individuals to make better decisions based on their own local and specific knowledge rather than on an omniscient or far reaching planner. Finally, the fourth condition of a “wise crowd” is aggregation, which helps turn individual judgments into a collective decision.

The aforementioned conditions have become more feasible by the advent of Web 2.0 systems, like wikis, which enable large groups of users to achieve impressive results on collective knowledge creation and sharing [71]. The openness, simplicity and support offered by these systems at a low cost, has led many organizations to increasingly invest on their adoption [11]. Nevertheless, the approach of collectively created content also presents certain shortcomings and critical voices exist that question the quality of the created information in wikis and other Web 2.0 technologies [44]. Especially as far as wikis are concerned, the lack of trust in the quality of the content as well as the uncertainty in its production time [29,30] are probably the main reasons why, after an initial period of promise and trial, companies are not so satisfied with their adoption [12]. A system is therefore needed to actively and continuously seek to ensure the quality of its material, by selecting, within the whole employee network of the organization, the most appropriate expert peers who will evaluate and contribute to the inserted information.

In this paper we propose CorpWiki, a self-regulating wiki-based system to collect corporate knowledge and combine it with the individual intelligence of the organizational members. Each article inserted in CorpWiki undergoes as many contribution and review processes as necessary to reach high-quality levels, within minimum time. To achieve this, an expert peer matching (EPM) algorithm is applied, which uses feed-forward neural networks to search the corporate employee network and select the appropriate in-house experts who will provide knowledge of high quality. The system combines machine-learning intelligence with the individual intelligence of peer users to create new inferences regarding corporate issues, thus using and promoting the collective organizational intelligence. CorpWiki manages to overcome the problem of questionable content quality which usually exists in the flat-based coordination approach, while, by coordinating contributions and reviews, it limits the unsystematic character that this approach usually has. In addition, contrary to the hierarchically-coordinated expert consultancy approach, the combination of expert knowledge in CorpWiki is not performed by moving up the hierarchy of the organization. Hence, the latency in knowledge acquisition is reduced and the responsiveness of the organization is increased. Finally, by combining more diverse user knowledge, CorpWiki manages to overcome the issue of narrow expertise related to the hierarchically-coordinated expert consultancy approach.

The rest of this paper is structured as follows: Section 2 presents the results and limitations of current research literature. In Section 3 the proposed CorpWiki system is presented including an analytical description of the EPM algorithm. In addition to presenting the basic CorpWiki functionality, this section also includes a potential use case of the system on risk management, a business process of strategic importance for organizations. The evaluation of CorpWiki is presented in Section 4. Finally, Section 5 concludes with the main findings and future work of this study.

2. Relevant literature

Various studies have been devoted to the development, application and evaluation of collective intelligence systems, which aim to gather the expertise of a group, rather than an individual, to make decisions. These systems may include groupware and group decision support systems [18,21,23,31], synchronous and asynchronous collaborative software [8,66], case based reasoning systems [6] and prediction markets [80]. The aforementioned systems facilitate the collection of opinions, the organizing of collaborative discussions towards the consensual prioritization of goals, as well as the support of group decisions and thus they are ideal for supporting task forces or project teams with predefined goals [33]. However, they present certain shortcomings regarding their complexity and cost, as well as their capability to efficiently represent the majority of the organizational knowledge [34].

An alternative approach, is the adoption of wikis, a popular Web 2.0 technology, which seeks to benefit from collective intelligence in a simple, cost-efficient and easy-to-update manner. The success of Wikipedia has promoted the application

of wiki technologies to various domains including enterprises where wikis offer a promising environment for collaboration and organizational learning, but also raise certain concerns. The nature of these concerns may be managerial, such as the loss of traditional organizational hierarchy and centralized control. It may also be social since wikis do not recognize authorship, and the quality of their articles is questionable. Finally legal concerns may also appear, posing intellectual property and liability for libel issues. Finally, an important issue is the employees' ability and willingness to use and contribute to the wiki. Studies that identify the above issues propose certain solutions including quality assurance through the volunteering contributions of qualified peers, the provision of corporate incentives to ensure participation and the development of corporate rules for the protection of intellectual property [29,30]. Ding et al. [24] implement a wiki to support corporate projects and identify two main issues of concern. The first issue refers to accessibility problems, since users cannot easily retrieve the information they search for. The second issue pertains to maintenance issues, since a lot of the corporate wiki information tends to be outdated. Majchrzak et al. [53] in their surveys on a diverse population of 168 corporate wiki users conclude that corporate wikis are a sustainable choice for an organization, especially when they are used for retrieving novel solutions and the contributions are made by credible users. The study of Wagner and Bolloju [76], in comparison to other Web 2.0 technologies, the characteristics of wikis render them beneficial for best practice communities.

The above studies conclude that the limiting factors to the success of corporate wikis are the lack of efficient information organization, the lack of authorship recognition which may lead to low credit for a user's contributions, the decreased participation and the low quality assurance that current wiki systems present. The lack of information organization inside a wiki is addressed by current research with semantic technologies that improve navigation and search [10], article classification [74], and content consistency and knowledge reusability [48]. Semantic wikis supporting domain specific needs have also been created [50,51]. To address user anonymity inside a wiki, two wiki-based systems, namely WikiGenes [35] and Knol [56] have been proposed. These systems combine the collaborative options provided by wiki technologies with explicit authorship in order to provide users with appropriate credit for their contributions. The issue of participation is also addressed by various studies proposing that incentives to participate should be given to corporate employees, to motivate them to contribute and maintain the enterprise wiki [30,56]. These incentives can be extrinsic or intrinsic and should be carefully selected in order to avoid negative effects such as the crowding-out effect, i.e., employees participating only in expectance of the reward [58]. On the issue of quality, a number of studies focus on the estimation of article quality by extracting article features and using various machine-learning techniques [7,39,68]. Other qualitative studies also indicate that article quality increases as cooperation levels increase [79], especially when contributions stem from experienced users [67].

Although the above approaches present ways to calculate quality, they do not actually assure article quality. To this end, Bughin [12] suggests that in order to improve the quality of their internal wikis, companies should have quality assurance practices which rely on appointed or self-appointed guardians to police quality issues. Other studies also propose the use of quality assurance techniques which usually wiki-gardeners, experience and trusted individuals, that have the task of ensuring and maintaining information quality [20,46]. To the best of our knowledge, no attempt has yet been made towards a wiki system that self-improves its article quality by selecting the most appropriate peer users to contribute. However, since most of the aforementioned studies have identified peer expert contribution as a significant factor of article quality assurance, we will summarize the relative literature in expert finding and peer matching techniques. Expert finding search systems use the documents created by a person, to form the textual evidence profile for this person. Then, using the similarity that a user's profile has to a specific query, the expertise of each user is calculated. Various studies have researched expert finding in enterprises using different types of documents and techniques which include language models [2–4,25,60,65], semantic analysis [81] or graph-based methods [52]. Matching someone's knowledge to another person's question, in peer social networks, is reported by a number of studies in the e-learning domain [45,59], as well as in the collaborative knowledge engineering domain [62].

Summarizing the above, one may conclude, that although wikis have the potential of promoting organizational learning and offering a promising collaborative environment, current wiki technologies present structural and maintenance problems, do not recognize authorship, and do not assure article quality. The proposed CorpWiki system, aims to address the aforementioned issues by following a structured information organization and by selecting the most appropriate expert users to contribute in order to assure content quality. Finally, although reviews and contributions are made anonymously to assure objectivity, authorship is also recognized on high-quality articles, thus motivating users to continue contributing.

3. Description of the proposed CorpWiki system

In an effort to overcome the limitations of existing techniques towards assuring the quality of collective intelligence, this study proposes a self-regulating wiki system, namely CorpWiki. The proposed system implements a peer review process among corporate employees to assess the quality of the inserted articles in order to decide whether an article needs improvement or not. In case improvement is required, CorpWiki automatically selects the most appropriate user to enhance the quality of the article, using an intelligent expert peer matching algorithm. The system thus receives contributed articles and reaches them to satisfactory quality levels, without the need for human coordination over the process of selecting the expert peers.

The rationale behind the proposed system is that the quality of an article can be objectively assessed by using the opinions of a large number of peer reviewers and significantly improved by selecting the most appropriate expert peers for its

revision. In addition, the system provides the corporate management with the flexibility to adopt different policies regarding its functionality.

3.1. CorpWiki terminology

Before CorpWiki starts processing an article, it first needs to be categorized into the knowledge domain it most appropriately belongs to, hereby referred to as “article domain”. The article domains can be determined by the management according to the specific business process on which CorpWiki will be applied. Then, for the categorization of incoming articles, a variety of document classification methods found in recent literature may be used. Such methods include document classification though the k-Nearest-Neighbor technique [5,70], neural networks [73], support vector machines [16,28], methods which are based on the vector space model [52], as well as methods which make use of the semantic relationships between the document words [49,82]. The aforementioned methods aim at classifying documents to knowledge domains based on the similarity that each document presents with the specific domain and have been successfully applied on large document collections, often of enterprise nature.

Each peer participating in the system may serve either as a “contributor”, when making a contribution or improvement, or as a “reviewer”, when commenting and grading an article. Therefore, the system stores two distinct profiles for each user, namely the “contributor profile” and the “reviewer profile”. The “article quality” is calculated by averaging the grades that a large number of reviewers have assigned to each article. Similarly, “user expertise” is directly derived from the quality level of the contributions that a user has made on a specific domain. In addition, a significant factor that is expected to affect the success of CorpWiki is the willingness of peer users to make a contribution or a review upon system request. Therefore, the “acceptance ratio” for each user is also computed as the ratio of the contributions that this user has made upon system request over the total number of requested contributions that the user has received. The acceptance ratio regarding the reviews of each user is calculated similarly. Thus, the contributor and reviewer profiles of a peer user each consist of the expertise on a specific domain, the acceptance ratio and the knowledge domain. The profile of each article comprises the current quality that the article has reached and the knowledge domain it belongs to. The system stores the profile of each peer as well as that of every article and updates them accordingly.

3.2. CorpWiki functionality

System functionality can be described as follows: As soon as a new article is inserted into the system, or a new contribution is made on an existing article, it undergoes a quality assessment control performed through the peer review process. During this procedure, the system selects a group of reviewers to assess the quality of this article. The most appropriate reviewers are those that demonstrate both high expertise and a high acceptance ratio in the knowledge domain of the article. After an adequate number of reviewers has accepted and performed their reviews, the system updates the expertise field on the contributor’s profile and the acceptance ratio field on the profile of each reviewer who received a system request. The number of reviewers assigned to each article can be either static, determined by the management according to the human resources that it wishes to commit, or it may be dynamic, based on whether the reviewers have determined with a specific amount of certainty that the article is of adequate quality or if it needs further revisions.

Next, if the article is found to be of inadequate quality, meaning that it has not exceeded a predefined quality threshold, the contributor selection process takes place. This process is controlled by the expert peer matching algorithm, which seeks to identify the most appropriate contributor for the specific article. The functionality of the expert peer matching algorithm is analytically described later on this section. As soon as a peer user, selected by the algorithm, contributes to the article, it undergoes re-evaluation in the same manner described earlier. The contribution, review and expert peer matching processes can be repeated several times until the inserted article has surpassed the quality threshold.

Articles that have acquired a satisfactory quality level are marked by the system as “stable versions” while articles which are under the process of review or contribution are marked as “working versions”. Users may optionally choose to be notified by the system as to the current status of a “working version” of an article – for instance that the article is under revision or that it is under improvement. Moreover, once an article has reached a stable version its contributors are displayed, to recognize authorship and motivate users to continue participating. Finally, CorpWiki preserves its wiki character by allowing users to voluntarily interact with an article at any moment, even after the article has reached a stable version. Fig. 1 presents an overview of the proposed system.

3.3. Expert peer matching using feed-forward neural networks

The expert peer matching algorithm (EPM) is used for selecting the most appropriate peer to contribute and improve a specific article. The algorithm uses a popular machine-learning technique, namely feed-forward neural networks (FFNNs) [32], to estimate the contribution quality that each peer would present if he was assigned with the specific article. More specifically, the EPM algorithm implements a number of FFNNs to approximate the function which maps the present quality of an article and the profile of a specific peer user, to the quality that this article will have upon the contribution of that specific peer. The number of the FFNNs used by the aforementioned algorithm is equal to the number of knowledge domains that exist in the CorpWiki system. Each one of these FFNNs approximates the following function:

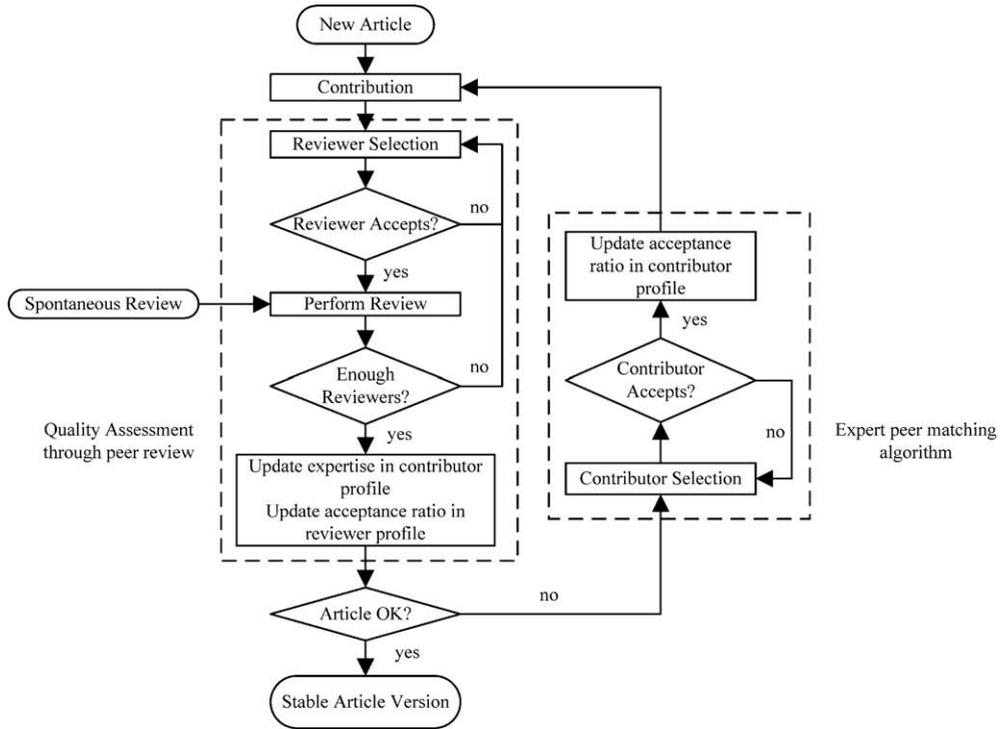


Fig. 1. CorpWiki functionality overview.

$$RAQ = f(CAQ, PEX, ACR), \quad (1)$$

where the network output RAQ stands for the Resulting Article Quality, that is the quality of the article after the contribution of the specific peer, and the input data of the FFNN consist of the current quality level of the article received, CAQ (Current Article Quality), the expertise of a specific peer user, PEX (Peer EXpertise) and the acceptance ratio of this peer, ACR (ACceptance Ratio). The CAQ element is calculated as the average of the grades that the current version of the article has received through the peer review process:

$$CAQ = \frac{\sum_{j=1}^n AG_j}{n}, \quad (2)$$

where AG is the Article Grade and n is the number of the reviewers who have graded the current article. The PEX and ACR elements both refer to the knowledge domain that the article belongs to. The PEX is calculated as the average grade that the peer's contributions have received in the knowledge domain of the specific article:

$$PEX = \frac{\sum_{j=1}^m PCG_j}{m}, \quad (3)$$

where PCG is the Peer Contribution Grade, and m is the number of contributions made by the specific peer user. The initialization of the PEX element is described at the end of this section, where the handling of new users is presented. The ACR is estimated as the ratio of the number of the submitted contributions that the specific peer has made upon system request, ACRQ (ACcepted ReQuests) to the total number of contribution requests that he has received, TRQ (Total ReQuests) on the knowledge field that the article belongs to:

$$ACR = \frac{ACRQ}{TRQ}. \quad (4)$$

A FFNN has the ability to learn from a set of examples, called the training set, and generalize its findings to an unseen population. Usually a FFNN is trained using the popular error back-propagation algorithm, established in [63,64]. The goal of this algorithm is to minimize a cost function, typically defined as the mean square error between the actual and target outputs of the network, by adjusting the synaptic weights and neuron biases. To this end, the network is presented with the training set, which consists of examples of an input vector and the corresponding output vector. Next, the information is passed forwardly from the input nodes, through the hidden layers, to the output nodes and the error between the desired and the actual response of the network is calculated. This error signal is then propagated backwards to the input neurons adjusting the network weights and biases. This process is repeated for each example in the training set and when the entire training set is presented to the network an epoch has elapsed. The network training phase may consist of several epochs.

A popular approach used to optimize the performance of the error back-propagation algorithm is the Levenberg–Marquardt algorithm [27] which has been found to increase the speed convergence and effectiveness of network training. Thus, the EPM algorithm uses the latter algorithm to implement the used FFNNs. The training set of each network consists of several pair examples of actual user contributions. More specifically, the network input information consists of the profiles of peer users and the quality of the article that they contributed to, while the network output consists of the quality rates that the article received after their contribution.

However, during its training, a FFNN may end up over-adjusting its weights and biases to represent only the training data and thus lose its generalization ability. This phenomenon is called overfitting and can be avoided using a separate data set called the validation set. The FFNN parameters are estimated only using the training set and the performance of the network is evaluated by computing the mean square error on the validation set. If the network performance is found to deteriorate, meaning that overfitting has probably occurred, training stops and the state of parameters from the previous network epoch is stored. Thus, to avoid overtraining the expert peer matching algorithm uses a separate validation set consisting of data that have not been used during training. The network training phase is terminated in case that one of the following conditions is met: if the cost function has reached to a minimum, meaning that the neural network cannot further reduce the mean square error between its actual and target outputs. Additionally, network training is terminated in case that the performance goal of the cost function, usually set to a very low numeric value – in the paper the performance goal was set to 10^{-10} – is met, meaning that the neural network has achieved to reach the desired mean square error. Finally training stops by detecting that the validation set produces increasing mean square error, a fact which threatens the generalization capabilities of the network, as described earlier in this section.

As soon as network training finishes, the network test phase takes place. During this phase the trained FFNN is presented with unseen data to evaluate its performance. These data compose the test set, a data set different to both the training and the validation sets. To evaluate the performance of the networks, the expert peer matching algorithm uses a test set that has not been used during the training or validation phases.

The first time the expert peer matching algorithm is applied, it randomly assigns contributors to articles, until a satisfactory amount of data have been gathered. This data is required for both training the FFNNs and for creating the initial user profiles which will be used by the EPM algorithm to select the most appropriate contributor. Then, the algorithm splits the gathered data set into three separate sets, the training, validation and test set. Next, it feeds the networks with the test set, and uses the validation set to avoid overfitting. When the networks are trained, it tests their performance on the test set.

Networks that have achieved satisfactory performance results on the test set are then used by the algorithm for the estimation of the contribution quality of new users. More specifically, as soon as an article arrives, the expert peer matching algorithm uses the FFNN, trained for the specific knowledge domain, to determine the article quality which would result if the article is given to a specific contributor. To this end, it feeds the trained network with the current quality of the article and the profile of each peer user. Then, based on the estimation that the FFNN has made for each possible contributor, it sorts users in descending estimated quality levels.

Next, it selects the most appropriate peer to contribute using the aforementioned sorted list. The selection process may undergo different policies discussed in detail in Section 3.4. Peers that receive a system request to contribute can either accept or reject it and their acceptance ratio is updated accordingly. It should be noted that, the same contributor is not assigned the same article because in this case no significant improvements are expected to occur, whereas another person would probably contribute new knowledge and perspective to the article. The algorithm implements this by preserving a list of peers that have already contributed to the article. The contributors in this list are not included in the list of possible contributors evaluated by the expert peer matching algorithm during the peer selection process for the particular article. The aforementioned process is repeated until a peer user accepts to make a contribution. An overview of the expert peer matching algorithm is presented in Fig. 2.

The profiles of the peers that have not yet contributed nor received a request for contribution to the specific domain are initialized as soon as they make their first contribution. Finally, the networks are periodically retrained to include new data that have been received since the last training. This process is performed to enable the system to use all the available information, including recently arrived pairs of examples, in order to make more accurate approximations on the expected result. The retraining period of the network depends on the rate that the dataset increases, as well as on how much the approximated model changes over time. A possible method for determining this time interval is by monitoring degradations in system performance, in terms of its ability to generalize well over the new incoming data.

3.3.1. Strengths and limitations of using FFNNs on the specific problem

The specific problem tackled by the EPM algorithm of CorpWiki requires an efficient approximation of the function that relates user and information characteristics to the predicted user contribution quality. This problem will involve a variety of complex real-world user data patterns and thus it is expected to be nonlinear, while the relationship among its data is expected to be difficult to describe analytically. In addition, in this problem, the only available information is past data. Taking the aforementioned into account, the prediction of future user performance can either be made by assuming the data model or by approaching it using a machine-learning technique, such as feed-forward neural networks. FFNNs present various strengths that make them suitable for the regression analysis requirements of the expert peer matching algorithm. Firstly, being universal function approximators [19,26,36–38] FFNNs have the ability to efficiently map nonlinear relationships between their input and output. Secondly, they can generalize, meaning that they can correctly predict the output of unseen

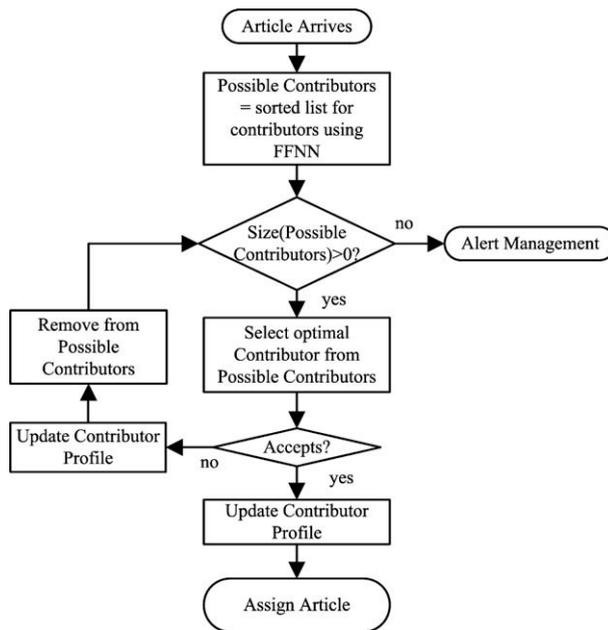


Fig. 2. Expert peer matching algorithm.

data, even if their training set contains noisy information [72]. Additionally, FFNNs are data-driven instead of model-driven, which means that they do not a priori assume an explicit relationship model among the data. Finally, trained neural networks are able to quickly make accurate predictions on unseen data, which makes them suitable for real-world problems where training needs to be made sporadically but predictions should be made in real time.

Nevertheless, neural networks also present certain shortcomings. Firstly, they usually require more time for training than linear methods, due to the number of iterations needed to achieve their optimal prediction. Secondly, FFNNs depend on the size and quality of the data used for their training [32]. Finally, neural networks are black-box methods and as such, they cannot be easily analyzed in great detail like linear models [1]. Since the success of the whole CorpWiki system greatly lies on the efficient approximation of the function that estimates user contribution quality, and given the fact that training will be made only sporadically to include new data patterns, the FFNN limitations – training time, dataset dependency, black-box method – are acceptable.

3.4. CorpWiki policies

An important element that the proposed system features is that it can be used along with different policies determined by the organization, regarding the number of the reviewers assigned to each article version, as well as the article quality threshold that is considered to be adequate. Moreover, different policies can be related to the formula used to calculate the article quality. For instance, this might be the simple average of the reviewer grades on the article or a weighted average based on the expertise of each reviewer. Furthermore, the activation of the peer matching algorithm can be subject to different policies. For example, one policy may indicate that the quality of an article is the only factor determining the need for revision, and a second policy may indicate that quality also depends on the number of the contributions that the article has already received.

The expert matching algorithm may also adopt various selection policies, ranging from most-appropriate-expert based, which aims to expedite the processes by selecting the contributor with the highest estimated contribution quality, fairness-based, which considers both the workload that each peer has and the peer's estimated contribution quality, or a combination of the above. To measure the fairness in the distribution of the workload, the Jain's fairness index [41] is used:

$$f(x_1, \dots, x_n) = \frac{(\sum_{i=1}^n x_i)^2}{n \sum_{i=1}^n x_i^2}, \quad (5)$$

where x_i is the number of articles assigned to the person i and n is the total number of persons involved. If user k makes an additional contribution, then the workload of this user will become x_{k+1} and the fairness index will become:

$$\begin{aligned} f(x_1, \dots, x_n) &= \frac{(\sum_{i=1}^n x_i + 1)^2}{n \left(\sum_{i=1}^{k-1} x_i^2 + (x_k + 1)^2 + \sum_{i=k+1}^n x_i^2 \right)} = \frac{(\sum_{i=1}^n x_i + 1)^2}{n \left(\sum_{i=1}^{k-1} x_i^2 + x_k^2 + 2x_k + 1 + \sum_{i=k+1}^n x_i^2 \right)} \\ &= \frac{(\sum_{i=1}^n x_i + 1)^2}{n \left(\sum_{i=1}^n x_i^2 + 2x_k + 1 \right)}. \end{aligned} \quad (6)$$

Since $x_k \geq 0$, the fairness f is maximized when x_k has the smallest value. Thus the user that will maximize the fairness if selected is the one with the minimum x_k . In addition, consider the following formula for the fitness of user k :

$$\text{fitness}_k = \text{RAQ}_k^p \left(\frac{1}{x_k} \right)^{1-p}, \quad (7)$$

where RAQ_k is the estimated resulting article quality if user k is selected and p is a constant in the $[0, 1]$ interval.

For $p = 0$ the fitness of a user decreases as his workload x_k increases (Eq. (7)). Thus, users with lower workload x_k will have larger fitness values and their selection will result in higher fairness due to Eq. (6). Similarly, users with high assigned workload will have small fitness values indicating that their selection will lead to decreased fairness. Thus $p = 0$ leads to the highest degree of fairness and the fairness-based policy. For $p = 1$, the fitness of a user is equal to the user's RAQ value and does not depend on their assigned workload. Thus, in this case the user selection does not consider fairness and selects users based only on their estimated resulting article quality, thus adopting the most-appropriate-expert-based policy. Intermediate values of p result in more balanced policies where both the fairness and the RAQ factors are taken into consideration for the user selection.

3.5. CorpWiki use case on risk management

In the previous section the basic functionality of the CorpWiki system has been presented. Nevertheless, in order to create value for the organization, CorpWiki should be integrated into the business fields that the company wants to support. In this section we present a practical use case of the application of CorpWiki on risk management, a business process of strategic importance for the organization. Risk management refers to the identification of risks that can affect an organization, as well as their respective countermeasures. Typically, risk management in each organization is performed using a limited number of experts who identify and assess possible risks and propose the necessary solutions [15]. However, the demanding environment – due to globalization, mergers and acquisitions, outsourcing – in which contemporary enterprises operate, makes the nature of the risks too complex and challenging to be efficiently addressed using only limited expert resources. In addition, risk management is a creative rather than a routine organization task and as such it would greatly benefit from the combination of the expertise and ideas from all the layers of the organization. Finally, since decision making based on the results of risk management significantly affects the viability of the organization, risk information quality assurance is also a major requirement. The capabilities of CorpWiki to identify expertise through the whole corporate network and to assure the quality of the information gathered, as well as the importance that risk management has for the organization, make this process an ideal use case for demonstrating the CorpWiki potential in creating value for the enterprise.

In this use case, CorpWiki receives risk information contributions from all the layers of the corporation, a fact that is expected to result to more complete and concrete coverage of enterprise risks. Then, the quality of the inserted risk information is improved internally, through the intelligent selection of risk experts, identified as such by their peers. Two objectives are thus met: risk identifications and assessments are made by experts on the field, while the expert selection is not restricted to a small group of individuals, but includes potentially all corporate employees. Therefore, the collective intelligence of the organization is promoted and used for arriving to more informed decisions regarding the risks that the organization faces, as well as preventing and minimizing the risk consequences more efficiently.

3.5.1. Description of the CorpWiki risk management application

As with every system that is applied on a specific business process of the organization, the basic CorpWiki system may require certain adaptations to best suit the needs of the particular risk management process. Furthermore, CorpWiki can cooperate with additional application-specific modules depending on the specific business requirement of the organization that it needs to be applied on.

In the general CorpWiki approach, each article comprises of a textual description and a numerical part corresponding to the rating of its quality. In the specific application of CorpWiki on the risk management field the textual information of each article, hereby referred to as risk article, consists of the risk name, the risk description and the solution proposed to address it. To measure the quality of the risk article two numerical ratings are used which assess the description adequacy and the solution effectiveness, respectively. Thus, the risk article is an instantiation of the basic wiki page adapted to the specific risk management application requirements. Moreover, in each risk article two additional risk management specific elements are used: the expected impact and the likelihood of occurrence of the risk described in the article. Impact refers to the significance of the potential consequences of the risk, while likelihood of occurrence expresses how likely the consequences of the risk are. Each of these two elements corresponds to a numerical value in the scale 1–3, where 1 represents a low and 3 represents a high impact or likelihood of occurrence.

As soon as a certain risk document is inserted into the CorpWiki, the system takes all the necessary actions to reach it to sufficient quality levels, as described earlier in Section 3.2. More specifically, each inserted risk document undergoes a peer review process during which the reviewers rate the quality of the risk article in terms of description adequacy and solution effectiveness. If, at the end of the peer review process, the average of either one of these two quality metrics falls below a predefined threshold, then the expert peer matching algorithm selects the appropriate contributor to enhance the article. The process is repeated until the risk document reaches satisfactory quality levels in both its description and its proposed

solution. Throughout the aforementioned process both the reviewers and the contributors, can rate the impact and likelihood of occurrence of the risk. This allows the system to gather as many employee assessments as possible and use this information to reach to more accurate estimations. It should also be noted that users do not review or see the ratings that other users have assigned the article with, as this would be a violation of the second condition of the wisdom of the crowds, referring to independence. According to this condition people inside a group should provide their judgments independently of each others' influence, in order to ensure that their opinions will not be biased, their mistakes will not be correlated and that new information will be brought into the group [69]. The aggregation of the impact and likelihood of occurrence as well as its use are described in more detail later on this section.

Thus the CorpWiki, applied on the risk management field, can be used to harness the collective capabilities of the employees in two ways: firstly by improving the quality of each risk article and secondly, by obtaining more accurate estimations on its impact and likelihood of occurrence, in order to lead to an overall more effective risk management process. A number of modules can cooperate with the CorpWiki system, to better serve the specific purposes of the risk management:

3.5.1.1. Classification module. This module is responsible for the correct classification of risks into their relative knowledge domain, according to the risk taxonomy which usually exists in the risk management framework of the organization [15]. The classification module guides users into inserting each new risk to the exact knowledge domain that it logically relates to, using a top-down question approach through the risk management taxonomy. To ensure that each risk is categorized correctly, one of the popular categorization algorithms, referred to in Section 3.1, can be activated periodically as well as upon insertion of a new risk to detect possible risk information duplication or misclassification.

3.5.1.2. User interaction module. This module controls the user interactions, namely contributions or reviews, with a risk article. These interactions are either initiated by a user or requested by the system. This module is responsible for maintaining information homogeneousness on each new or updated risk article by guiding users, through a series of steps, to the completion of all the elements of the risk article. Additionally, this module is responsible for updating each risk article with new information regarding its quality.

3.5.1.3. Risk level calculation module. This module receives all the assessments made by the users on the impact and likelihood of occurrence for each risk, and aggregates them, through majority voting, to calculate the overall impact and the overall likelihood of occurrence for that specific risk. Then, by combining the overall impact and overall likelihood, the module calculates the risk level [40]. Each time a new assessment is made on the impact or likelihood of a specific risk, it is aggregated with the respective previous values and the level of the risk is updated accordingly.

3.5.1.4. Recommendation module. Risk reporting and recommendation are controlled by the recommendation module, which allows management to monitor possible risks and filter them based on their risk level, risk category or solution effectiveness. Being capable to easily examine a variety of risk categories, the management is facilitated towards a more effective realization of risk management tasks and better decision making.

3.5.1.5. Cooperation measuring module. The cooperation measuring module, measures the level of engagement of employees with the system. Since cooperation is an important element of the proposed system, as well as of all Web 2.0-based systems, this module aims at being used to recognize the efforts that employees make and further motivate them to participate. At individual level, the cooperation module measures the participation that individual users show in the CorpWiki system. This information can then be used to demonstrate the top participating users, a social recognition reward which has been found to increase individual incentives to share one's knowledge [13,77]. Low participation information is not to be demonstrated by the cooperation module since it would act in a demotivating way, which is not the purpose of the module. At community level the participation information produced by the cooperation module, can demonstrate the overall usage level of CorpWiki. Then, in case this level is low, management could provide more incentives in the form of group rewards [13], linked to the overall success of the knowledge sharing efforts of the community. As an alternative means of increasing participation, management could raise the team spirit [13,58] of the participating employees, through, for instance, the creation of knowledge sharing teams, specialized in the various knowledge domains covered by CorpWiki.

3.5.1.6. Alert generation module. This module is responsible for generating alerts notifying management about high-level risks which require immediate actions. It also generates alerts in case that the number of topics in a certain risk category has grown significantly, thus indicating the need for their division in sub-categories.

4. Performance evaluation

Since the proposed CorpWiki system relies on the cooperation of a large number of users for the development and evaluation of high-quality articles, such a system is difficult to be evaluated prior to its operational deployment. However, some of the advantages of the proposed system may be foreseen through simulation modeling. Thus, the performance evaluation

of the proposed system was completed in two distinct phases: the design of the system models to be used during simulation and the simulation phase itself.

The design phase of the performance evaluation includes the design of two elements to be used during simulation, namely the *Real-World Generator* and the *CorpWiki System Model*. The Real-World Generator is used to generate the peer users and their characteristics, as well as the interaction that each of these users has with the proposed system. The CorpWiki System Model contains a key-subset of all the necessary data and logic that form the proposed system, including the expert peer matching algorithm. During the simulation, the users, generated by the Real-World Generator, interact with the CorpWiki System Model, composing and reviewing articles and generally responding to the system requests. Therefore, the simulation consists of the parameterization and execution of the CorpWiki system, as well as the performance logging and presentation.

It should be noted here that the aforementioned modeling does not intend to achieve an accurate prediction of the studied system in operational use, since it relies on a simplified real-world model. On the contrary, it intends to illustrate the way in which the proposed algorithm will improve system performance in a qualitative manner. These results are intended for fine-tuning the system and for justifying its performance prior to its testing in a large organization.

4.1. Simulation model design

4.1.1. Real-World Generator

The Real-World Generator is an abstraction of the anticipated behavior of the environment that interacts with the proposed system. It is used to generate the individuals and the articles that will interact with each other during system simulation. It is important to note here that the real-world model is not known nor taken into account by the simulated system, at any step of the process. For the purpose of our evaluation, we have made the following assumptions:

4.1.1.1. Person. Each person generated by the real-world model can act both as a reviewer and as a contributor depending on the request each time received by the system. Every person is characterized by an expertise in each domain, which indicates the knowledge that this individual has on the domain. The knowledge of a person on a specific domain is considered to be constant and it is assigned independently for each domain. The person expertise is initialized according to two different distributions: (1) uniformly in the $[0, 1]$ interval and (2) according to the beta distribution, which receives values in the range of $[0, 1]$ and its probability density function is maximized on the 0.5 value and minimized at the interval boundaries. In addition, a person's probability to accept to contribute to, or review a specific article is correlated to the expertise that the person has on the domain of the article, being higher for domains where his expertise is higher.

4.1.1.2. Article. An article which is created by a person of the real-world model generator, is firstly considered to belong to a domain. Moreover, each article is assumed to have an internal quality. The internal article quality is set irrespective of the individual subjective review opinions that it may receive. The internal quality of the article is neither known nor used by the CorpWiki System Model. On the contrary, it is only used by the Real-World Generator to generate the reviewers' grades on the article, and to produce the internal quality of a future revised version.

More specifically, the reviewing function that is used by the real-world generator to calculate the rating of an article by a reviewer is random, and follows a beta distribution. The reason the beta distribution is selected, is because it is double-bounded. The probability density function is

$$f(x; a, b) = \frac{x^{a-1}(1-x)^{b-1}}{\int_0^1 u^{a-1}(1-u)^{b-1} du}, \quad (8)$$

where the parameters a and b , depend on the internal article quality (q) and on the knowledge of the reviewer (e). It should be noted that beta distribution is bounded in the $(0, 1)$ interval. Since the article qualities in the model range from $(0, 10)$, the internal article quality is divided by 10 before calculating the parameters a and b , and after the rating has been calculated, the resulting value is multiplied by 10. For $q' = \frac{q}{10}$, the mean rating every reviewer assigns, regardless of his knowledge, is modeled to be equal to the internal quality of the article, thus $\frac{a}{a+b} = q'$.

On the contrary, the variance (v) of the ratings depends on the reviewers' knowledge. Assuming a linear relation between the variance and the knowledge, and that an ideal reviewer ($e \rightarrow 1$) would provide an accurate rating ($v \rightarrow 0$) we have that $v = q'(1-q')(1-e)/c$, where c is a scaling factor. The term $q'(1-q')$ is used for reducing the variance near the values 0 and 1. This is required because otherwise no beta distribution would be able to provide a mean equal to the internal quality.

Since the variance of the beta distribution is $v = \frac{ab}{(a+b)^2(a+b+1)}$, the parameters a and b may be calculated by solving the above equations. Thus $a = \frac{q^2(1-q')}{v} - q'$ and $b = \frac{1-q'}{q'} a$.

For the scaling factor c , the value 3 is selected, so that a reviewer with $e = 0$ (the worst possible reviewer) rates an article with quality $q = 0.5$ following a uniform distribution.

In addition, the contribution function, used to produce the internal quality of a future revised article version, is the following:

$$q_{t+1} = 10e + 0.85q_t - 0.85eq_t, \quad (9)$$

where q_t is the quality prior to the contribution, e is the contributor's knowledge, and q_{t+1} is the quality after the contribution. This function has the following properties:

An ideal contributor ($e \rightarrow 1$) produces perfect articles.

A contributor with $e \rightarrow 0$ (the worst possible contributor) will reduce the article quality by 15%. The initial version of the article ($q_t = 0$) is proportional to the contributor's knowledge.

It should be noted that these assumptions may have an impact on the quantitative behavior of the system, but we believe that this simplified and intuitive model can produce results that will have a qualitative resemblance to the actual deployment.

4.1.2. CorpWiki system model

The development of the CorpWiki system model intends to focus on the performance evaluation of the novelties introduced by the study, more specifically the self-regulating wiki using the expert peer matching algorithm (EPM). Thus, the other features of the system, which are based on well-accepted algorithms, such as the categorization of the articles in knowledge domains, will be considered as already evaluated by relative literature. Therefore, the CorpWiki system model consists of the following elements:

4.1.2.1. Person profile. The person profile contains the specific variables that will be used by the EPM algorithm, for the selection of the most appropriate expert. These variables include the average of the reviewer ratings that the contributions of a specific person have received on each domain, and the number of articles that this person has accepted or rejected to review or contribute to.

4.1.2.2. Article profile. The article profile comprises the variables that are used by the system to determine whether an article has reached sufficient quality or further improvement, through additional contributions, is required. These variables include the score that each reviewer has assigned to the article since its last revision, the last person who edited the article, the number of revisions so far, and the average score of the previous article version.

4.1.2.3. Expert peer matching algorithm. This algorithm uses a feed-forward neural network, as described in Section 3.3, to estimate the resulting article quality that the contribution of a person will have.

4.1.2.4. Contributor selection policy. The contributor selection policy is used by the expert peer matching algorithm to select the person who will be requested to improve an article that has been found to require further improvement. The three policies that were implemented and used during simulations are the following: (1) the “*most-appropriate-expert*” policy, where the only criterion for selection is the expected result of each peer contribution, (2) the “*fairness-based*” policy, where the selection is made based on a fair workload distribution and (3) the “*random-selection*” policy, where users are randomly selected. The latter is used as the reference basis.

4.1.2.5. Article acceptance policy. This policy is used to determine whether an article meets the desired quality standards or not. The article acceptance policy implemented in this performance evaluation is either “*threshold-based*”, with articles being revised until their average reviewers' scores exceed the predefined threshold, or “*revision-based*”, in which articles are revised a certain number of times.

4.2. Simulation results

After designing the real-world generator and the CorpWiki system model, the simulation phase took place and various scenarios, referring to possible corporation requirements, were examined. These scenarios include the necessity for improved article quality, prompt article production and the examination of different levels of fairness in workload distribution.

4.2.1. Simulation configuration

Before presenting the results of the system evaluation, the simulation configuration features are presented. Fifty individuals were generated and used during the simulation and the domains that the articles are categorized into three. Then, the neural networks to be used by the EPM algorithm were implemented using the Mathworks MATLAB R2007b platform and trained according to the training process and the Levenberg–Marquardt algorithm described in detail in Section 3.3. The precise training process specifications are as follows: a separate FFNN, consisting of two hidden layers of eight and five neurons respectively, was trained for each knowledge domain. From the available training data, 85% was used as the training set and the remaining 15% as the validation set, to prevent overfitting.

To determine the adequate size of training data for the specific problem addressed, different sizes of training sets were examined and the mean absolute error between the trained network outputs and the respective actual outputs was measured in relation to the size of the training set. More specifically during these tests, a large data set of randomly assigned

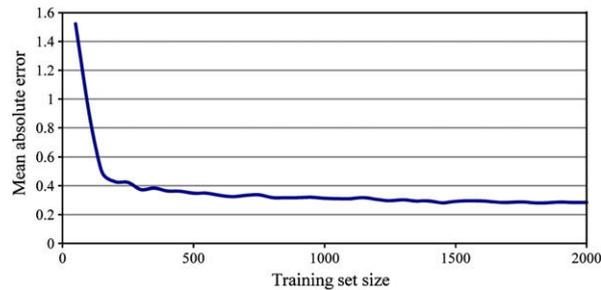


Fig. 3. FFNN training set size in relation to the mean absolute error.

Table 1

Simulation parameters.

Users	50
Domains	3
Articles per simulation run	1000 (unless otherwise specified)
Reviews per article (static)	20
Minimum reviews per article (dynamic)	3
Maximum reviews per article (dynamic)	20

4000 contributions was created. Then, this dataset was split in two parts. The first part was used as the training set and the second part as the test set. The experiments were conducted for increasing training set sizes which ranged from 50 to 2000 contributions. Fig. 3 corresponds to the results of this comparison for the first knowledge domain. The results of the networks in the other two domains were similar. It may be observed that the mean absolute error decreases as the training set size increases. It may also be observed that a mean absolute error less than 0.5 in a scale of 10 can be achieved if 150 contributions, have been made. Thus, for the specific problem encountered in the simulations, using 150 or more contributions for each domain may be considered a satisfactory amount of training data.

The three networks which were used during the simulations were trained using 1000 articles. The trained networks were then used during the performance evaluation for a given real-world model parameter. In case such a parameter changed, the networks were retrained. The set of articles used during the training phase was disjoint to the set of articles used during the performance evaluation.

During the simulation a number of different scenarios were examined. Each scenario tested the system performance according to one independent variable, namely the article quality threshold, the number of article revisions and the p -value. The p -value corresponds to a constant in the $[0, 1]$ interval and it is used to calibrate the system to the desired fairness level, with $p = 1$ leading to the most appropriate expert policy, $p = 0$ leading to the highest degree of fairness and intermediate values resulting in more balanced policies. For each value of the independent variable, 10 different simulation rounds were executed and their results were averaged. One thousands articles were created and improved in each simulation round. The results of the FFNN accuracy in predicting the author contribution quality, measured in terms of the mean absolute error between the network output and the actual value, were 0.2765, 0.2659 and 0.3073 for the 1st, 2nd and 3rd domains, respectively. The results of the FFNN accuracy in predicting the author contribution quality, measured in terms of the mean absolute error between the network output and the actual value, were 0.2765, 0.2659 and 0.3073 for the 1st, 2nd and 3rd domains, respectively. The simulation parameters are depicted in Table 1.

4.2.2. Examined scenarios

A number of different scenarios were examined to evaluate the system performance under different conditions and performance metrics. In the following we present the rationale for examining each one of these scenarios, their description and the respective system performance.

4.2.2.1. Scenario 1. Independent variable: article quality threshold. In the first scenario we examine the improvement that the proposed system can provide the organization with, regarding the number of revisions needed to surpass a certain article quality threshold, and the average article quality achieved. To examine system performance on different quality threshold requirements, the threshold-based article selection policy is used. The most-appropriate-expert-based contributor selection policy and the random-selection policy are compared in terms of the average quality achieved at each specific threshold and the average revisions required for each policy to reach this threshold. Ten different article threshold levels were examined, ranging from 0.5 to 9.5. Each policy was tested 10 times per threshold level and the results it yielded at each level were averaged. Figs. 4 and 5 present the aforementioned simulations in case that the expertise of individuals follows a beta and a uniform distribution respectively. As it can be observed in Fig. 4, the EPM algorithm results in higher average article quality and

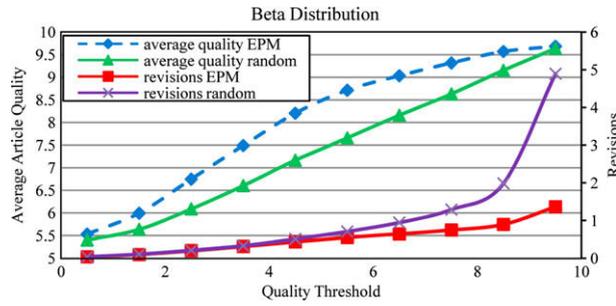


Fig. 4. Average article quality and number of required revisions in relation to the quality threshold – beta distribution.

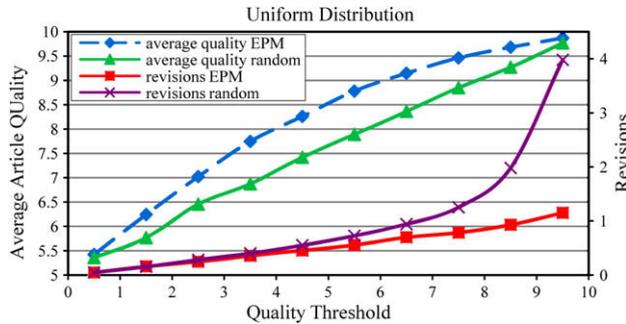


Fig. 5. Average article quality and number of required revisions in relation to the quality threshold – uniform distribution.

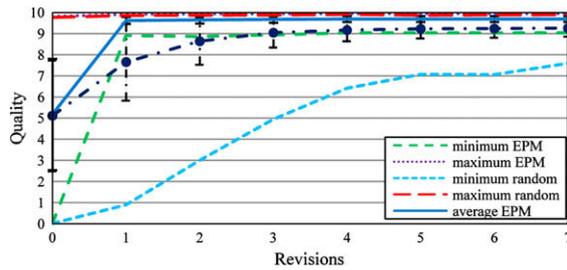


Fig. 6. Article quality in relation to the number of revisions.

fewer average article revisions in all threshold levels. Especially at the very high-quality thresholds of 8.5 and 9.5, the number of revisions needed by the random-selection policy is 2.2 and 3.6 times larger, respectively, compared to the number of revisions needed by the expert peer matching algorithm.

As depicted in Fig. 5, where the distribution of user expertise is uniform, the performance of the proposed algorithm presents similar characteristics irrespective of the real-world model used.

4.2.2.2. Scenario 2. Independent variable: number of revisions. In case the corporate management requires that articles should be finished within a specific time period, the number of revisions that each article will receive is predetermined. The second evaluation scenario refers to the improvement that the proposed system can achieve on the article quality according to a specified number of revisions. The revision-based article selection policy is used to examine the system performance on eight different revision levels, ranging from 0 to 7 revisions. For each revision level, the most-appropriate-expert contributor selection policy and the random-selection policy were tested 10 times each in terms of the average, the maximum, the minimum and the standard deviation of the article quality they yield. These results were then averaged for each revision level and policy.

As one may observe in Fig. 6, the expert peer matching algorithm performs better in comparison to the random-selection policy. More specifically, the EPM algorithm ensures that articles will have an average quality of at least 9.62 in a scale of 10, even after only one revision. On the other hand, the average quality achieved by the random-selection of authors is considerably low when the number of revisions is restricted to 1 or 2, and at least 3 revisions of the article are needed to reach

quality level 9. Moreover, the average quality achieved by the random-selection of the authors never manages to surpass the respective quality of the proposed algorithm, even after a considerable number of revisions.

The maximum quality achieved by both methods is the same, since a person with a high-level of expertise might be selected even with the random-selection policy. However, the minimum quality levels achieved by the proposed method are usually around 8.9, while the minimum quality levels achieved by the random-selection policy are substantially lower, especially when the number of revisions is limited. It is interesting to observe that, the minimum quality achieved by the proposed expert peer matching algorithm at low revision levels, corresponding to 1 or 2 revisions, is so high that it exceeds even the average random-selection quality level.

4.2.2.3. Scenario 3. Independent variable: *p*-value. The third scenario examines the tradeoff between the system performance and the workload assigned on system users. Therefore the contributor selection policy used is fairness-based. To this end, different *p*-values corresponding to different levels of the fairness index are tested in relation to the maximum number of contributions assigned to a single user, to the average article quality and to the average number of revisions necessary to reach a predetermined quality threshold.

Since the management is more likely to use the proposed algorithm to achieve high-quality articles, this scenario is tested on two high-quality thresholds, namely 8.5 and 9.5. The *p*-value levels examined range from 1, where the algorithm always selects the most appropriate expert, to 0, where the selection targets towards the maximum level of fairness in workload distribution. The algorithm was tested 10 times for each *p*-value and the results were then averaged.

As illustrated in Fig. 7a, as the *p*-value increases, the average article quality corresponding to the 8.5 threshold increases as well. On the contrary, the average article quality does not present significant changes as far as the 9.5 threshold is concerned, since the minimum acceptable quality value is initially high. In addition, according to the results presented in Fig. 7b, as the *p*-value decreases, the average number of revisions needed to achieve the quality thresholds increases. This is due to fact that, in case the corporation increases its fairness requirements, the article contributions may be assigned to less suitable users, who improve the article quality only marginally, resulting in more revisions. On the contrary, if the fairness requirements are loosened and more suitable users are selected, significant improvements are added to the article quality, resulting in fewer revisions. However, in this case the maximum workload assigned to a single user increases.

Fig. 7a and b indicates that a tradeoff does exist, between the need to reduce the maximum workload, and the need to increase the article quality as well as to expedite article production. Therefore, the corporate management has to decide on the policy which best fits the organizational requirements.

4.2.2.4. Scenario 4. Considering the quality of the reviewers. In the previous scenarios, which focused on improving the article quality and production time, the quality of an article was calculated as the simple average of the grades that the reviewers assigned the article with, according to Eq. (2). However, since all the reviewers do not necessarily possess the same expertise, and taking into account that expert reviewers are expected to be more accurate in rating an article than non-expert ones, in this scenario we estimate the current article quality (CAQ) as the weighted average of the ratings that the reviewers assign to an article, with a weighting factor that depends on their expertise.

The peer expertise (PEX), i.e., the average grade the user has received, in each domain increases for users that have contributed high-quality articles in the specific domain. Thus, it is used for calculating the *RevRank* value for every reviewer *j*:

$$RevRank_j = \frac{1}{1 - \frac{PEX_j}{10}} \tag{10}$$

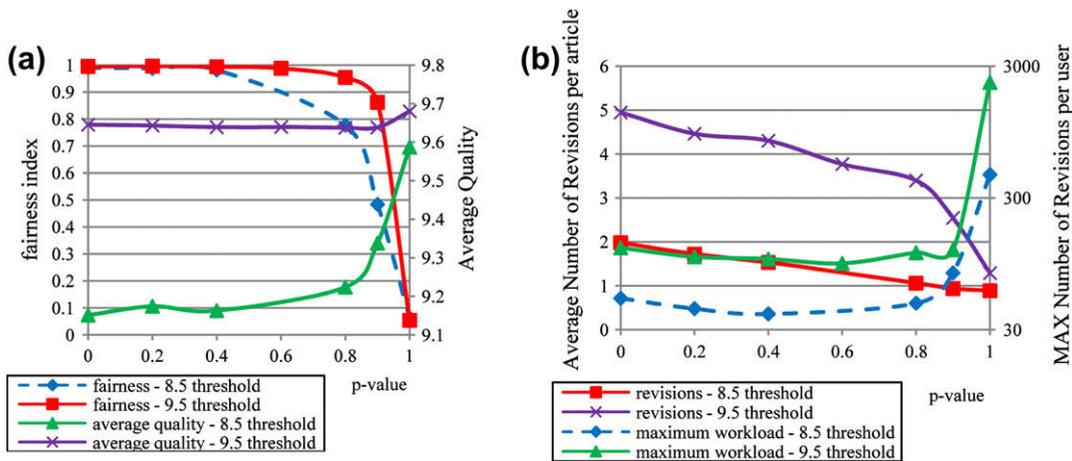


Fig. 7. (a) Fairness index and average article quality in relation to the *p*-value, thresholds 8.5 and 9.5 and (b) average number of revisions per article and maximum number of revisions per person in relation to the *p*-value, thresholds 8.5 and 9.5.

This value is used as a weighting factor to calculate the current article quality more accurately, as follows:

$$CAQ_{RevRank} = \frac{\sum_{j=1}^n RevRank_j * AG_j}{\sum_{j=1}^n RevRank_j}, \quad (11)$$

where n is the number of reviews on the specific article, and AG_j is the article review grade given by reviewer j . The above equation is inspired by PageRank [9], the popular algorithm used by Google that considers the quality of the votes. More specifically, in the PageRank algorithm the more “strong” a web page is, the more “strength” its vote has. Similarly, in the RevRank version of Eq. (2), the more expert a reviewer is, the more influential his opinion will be to the final rating of the article. In order to evaluate the accuracy of this RevRank grading method, a simulation was carried out using 1000 articles, with a threshold of 9. After every review is received, the article quality is calculated using the RevRank and the simple average methods. These values are compared to the internal article quality, and the absolute error in the estimation is calculated. Finally, the mean absolute error is calculated for varying numbers of reviewers for both algorithms, and the results are depicted in Fig. 8.

As the number of reviews increases, the mean absolute error decreases for both algorithms. The RevRank method however results in smaller mean absolute errors than the average method. Thus, fewer reviewers are required to reach a certain level of accuracy.

4.2.2.5. Scenario 5. Dynamic reviewer number calculation. In the previous scenarios, which focused on the effectiveness of the system to ensure timely creation of qualitative content, the number of reviewers assigned to each article in each review round was static and relatively large, since we wanted to obtain as accurate article quality ratings as possible. However, since the review process consumes a part of the reviewer’s productive time, assigning an increased number of reviewers to an article may in turn increase the organizational costs. On the other hand assigning too few reviewers may not result in accurate article quality rating. Thus, determining the appropriate number of reviewers is essential for the practical usage of the system in real-world organizations. In this scenario, a dynamic reviewer number estimation method has been incorporated to the expert peer matching algorithm. The method calculates the number of reviewers as follows:

After each review is received, the system estimates the quality of the article as the average of the reviewers’ ratings (CAQ). It also estimates the standard error for these ratings, as: $s = \frac{\sigma}{\sqrt{n-1}}$, where n is the number of ratings and σ is the standard deviation. At this point, if $|CAQ - Th| \geq \rho s$, where Th is the threshold and ρ is a constant related to the confidence of the decision, then the system may decide with a high-level of confidence that the article can be correctly classified as stable or as needing revision. Thus, in this case, no more reviews are needed. On the contrary, if $|CAQ - Th| < \rho s$, then there is uncertainty on whether the article quality is actually higher than the threshold, and for this reason more reviews are requested. It should be noted that the number of reviewers per contribution is at least 3, and at most 20, in order to avoid extreme numbers of reviewers when the quality is very close to the threshold. The number of reviewers needed for deciding whether the article is above or below the quality threshold is different for the RevRank and the simple average methods of calculating CAQ, since both the mean, and the standard deviation are calculated differently.

In order to evaluate the performance of the above reviewing policies, we compared it to the static number of reviewers allocation method, for various thresholds, using both the random contributor selection policy, and the EPM algorithm (Fig. 9). It may be observed that the use of the dynamic reviewer number calculation method decreases the number of reviews per article, without significantly affecting the system performance, in terms of misclassifications and average article quality achieved. In addition, the use of the RevRank method of calculating CAQ further reduces the number of reviews required.

An additional observation may be made on Fig. 9c. More specifically, one may observe that in this figure, the functions related to the random-selection policy are monotonically increasing while the respective functions related to the EPM selection policy are not. This may be explained as follows:

For the random policies, the classification error monotonically increases with the increasing threshold. This is due to the fact that at low thresholds it is more likely to exceed the threshold significantly after each revision and thus the reviewers

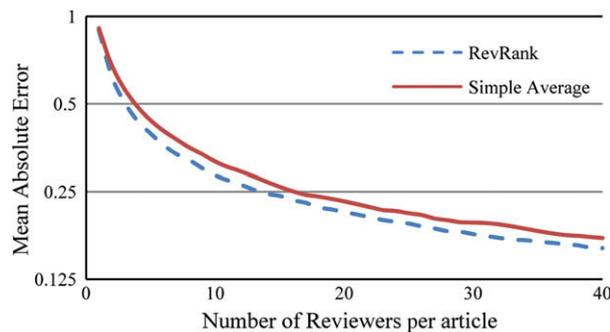


Fig. 8. Mean absolute error of article quality in relation to the number of reviewers, using the simple average and the RevRank methods of calculating CAQ.

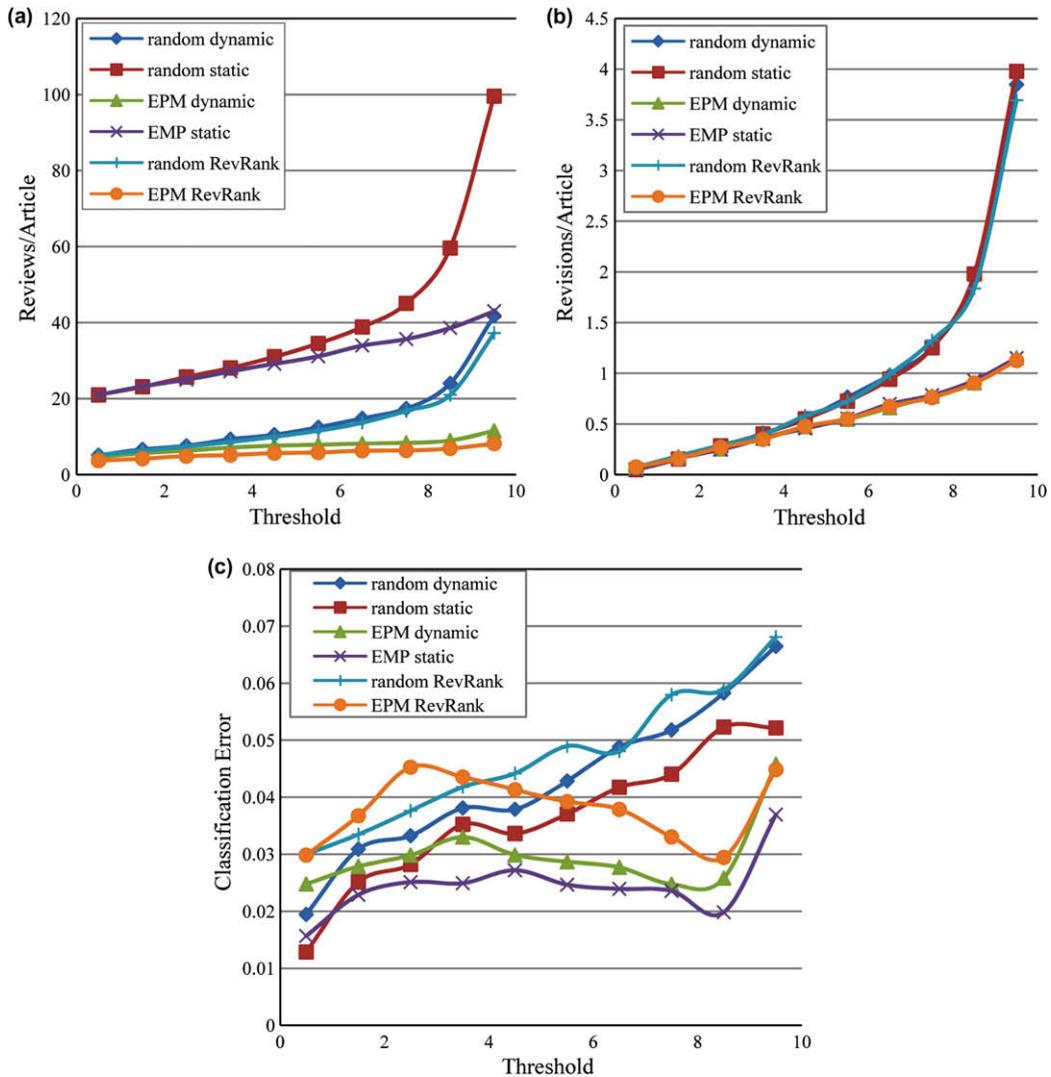


Fig. 9. Average reviews/article (a), average revisions/article (b), average misclassifications (c) on various threshold levels for the dynamic and static reviewer number calculation methods.

are more likely to classify the article correctly. On the contrary, when the threshold is high, a revision with the random-selection will probably just marginally exceed the threshold, leading to more misclassifications of the article by the reviewers and therefore leading to higher classification errors.

For the EPM policy, the classification error does not increase monotonically with the increasing threshold. This is explained as follows: In this policy the only case that the internal article quality may be close to the threshold is the first time that the article is created, since the initial contributors may have any expertise. If a revision is decided then it will be performed by an expert and thus the resulting article quality will be much higher than the threshold. Thus, when the EPM policy is used, the reviewers are more likely to misclassify the article on the initial contribution, except for the case of the very high threshold of 9.5. In addition, as the threshold increases fewer articles exceed it as initial contributions, and thus the percentage of articles with a high risk of misclassification decreases. The classification error is higher for the dynamic and RevRank policies, compared to the static policy, due to the smaller number of reviewers, which leads to more misclassifications. Finally at the high threshold value of 9.5 the selected experts, despite being the best possible, produce articles that surpass the quality threshold only marginally, thus increasing the classification error.

The dynamic reviewer number calculation method is expected to decrease the number of employees that are required to review an article and as a result to lead to decreased costs for the organization. However, this method has the drawback that it is slower compared to assigning the article directly to a large number of reviewers and as such it is more suitable for organizations that are more concerned about the organizational cost and less about the turn-around time of the process. In addition to using the dynamic reviewer calculation method, another option to further decrease the organizational costs would be

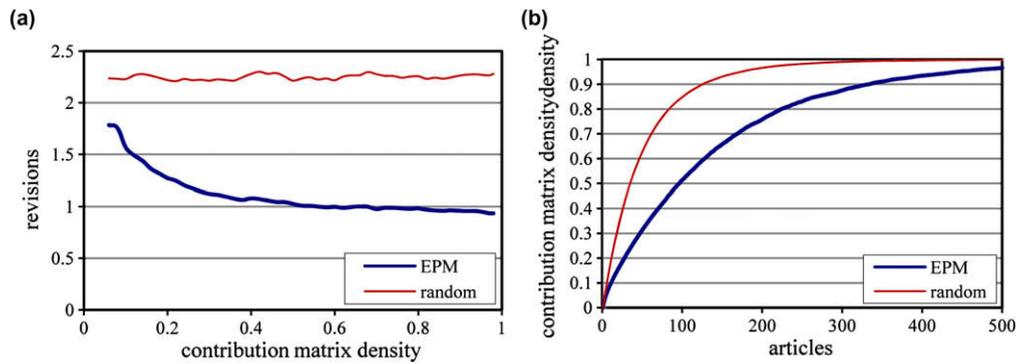


Fig. 10. (a) Average number of revisions per article in relation to the density of the contribution matrix and (b) density of the contribution matrix in relation to the number of articles.

the review system to highlight the changes made by the contributor, so that the reviewers would easily distinguish them, thus performing more robust and less time-consuming reviews.

4.2.2.6. Scenario 6. Density of the contribution matrix. In this scenario we examine the performance of the expert peer matching algorithm in relation to the user data that are each time available to the system.

Even for a well trained network the expert peer matching (EPM) algorithm can only select a contributor among those users for whom an available profile exists. In other words, knowledgeable users may exist that are not considered for selection by the expert peer matching algorithm because they have not yet contributed anything. Therefore system performance when using this algorithm depends also on how many user profiles are available at the moment. Moreover, the use of the EPM algorithm limits the set of future non-spontaneous contributors to the current profiles. On the other hand the random-selection policy requests contributions also from users that may have not yet contributed.

Therefore, the performance of the expert peer matching algorithm depends on the density of the contribution matrix, i.e., the fraction of the users that have contributed to an article per domain. Specifically, if C is the contribution matrix and $C_{ij} = 1$ denotes the case when author i has contributed at least one article to the domain j , then the density of the contribution matrix d is: $d = \frac{\sum_{i,j} C_{ij}}{U \cdot D}$, where U is the total number of users and D is the number of domains.

During the system initialization, the density is low. As users gradually contribute to the knowledge domains, the system gathers increasingly more profile data and the density increases.

In this scenario we examine the performance of the EPM algorithm for various densities and compare it to the respective performance achieved by the random contributor selection policy.

Fig. 10a illustrates the average number of revisions needed, for various levels of density and a quality threshold equal to 9.0. It may be observed that the performance of the random policy does not depend on the density, whereas the performance of the expert peer matching algorithm improves as the density increases.

Thus, it is beneficial to rapidly increase the density when it is low, in order to optimize the results of the EPM algorithm. Fig. 10b illustrates the number of articles required to achieve various levels of density. As one may observe, when the expert peer matching algorithm is used, the density increases slowly over time, since the contributors are selected based on their previous contributions. On the contrary, the density increases more rapidly when the random-selection is applied, since more new users are requested to contribute articles. Thus, the system performance may increase if the contributors are selected randomly during the initialization of the system, in order to rapidly increase the density of the contribution matrix and discover possibly better contributors than the ones that are so far known. This process will in turn result in an increased efficiency of the expert peer matching algorithm when it is eventually used. For the above scenario, a density equal to $d = 0.2$ produces an average number of revisions which is only 30% worse than the one produced at $d = 1$ for the expert peer matching algorithm.

5. Issues for further consideration

The simulations present a number of simplified scenarios of the proposed system. However, in a realistic setting other issues may also occur. These issues are discussed in the following.

5.1. Aging of articles

As time elapses, the information included in the articles of acceptable quality may age and no longer be up-to-date. To address this, and protect the gathered knowledge against becoming outdated, an article updating process may periodically take place. That is, as soon as a specific time interval has elapsed, each article is re-inserted to the expert peer matching

algorithm in order to undergo the process of improvement, accompanied by a message that the information of the article may need updating. The time intervals between article updates can initially be the same for all domains but, since different knowledge domains may require different updating frequency, these initial time intervals can be dynamically adjusted over time. Take the case of two knowledge domains, biotechnology and history. At first, the initial time interval for the re-evaluation of the articles of both domains is six months. After the re-evaluation, a significant drop may be observed in the quality of the biotechnology articles which had initially passed the quality threshold, mainly due to the fact that they contain outdated information, while the respective articles in the domain of history, have not met significant quality deterioration. Therefore the updating time interval should be changed accordingly, to lengthen the re-evaluation period for the domain of history and shorten the respective period for the domain of biotechnology. This dynamic change will allow the system to address the issue of article aging while, at the same time, excessive use of valuable peer resources for the unnecessary re-evaluation of the articles is avoided.

5.2. Profiles of new users

In this paper, the profiles of new users are initialized as soon as they make their first contribution. However, as an alternative, the management may set the initial expertise values based on the assumed knowledge and expertise that employees of each department have on each domain, according to their everyday work responsibilities. For instance, instead of initializing all new users to the same expertise value, the management might choose to define that users from the IT department have higher expertise value on the domain of technical computer issues than the users from the law department. Whichever strategy is chosen to define the initial values, the expert peer matching algorithm will, after a few runs, define the level of expertise for each user based on peer ratings and will act accordingly.

5.3. Topic classification

While the system is being used, the domains into which the information is organized may need to be split, deleted or merged. As far as domain splitting is concerned, a first indication that this may be necessary is that the domain includes an increased number of articles. Another indication may be when a significant diversion on the grades of the same contributors is observed, in other words when the same users receive high ratings for a group of articles but low ratings for another on the same domain. A third indication of domain splitting is when the contextual information of some groups of articles, as extracted by the applied categorization algorithms, differs significantly from the contextual information of other groups of articles. In case one or more of the aforementioned conditions are present, then the system should generate an alert requesting the attention of the management on the possibility of splitting the domain. If the categorization algorithms find that the contextual information of the articles of two or more domains is very similar, then the system should prompt the management to consider merging the domains and deleting the remaining ones.

5.4. Related domains

In the simulation scenarios examined in Section 4.2.2, the domains used were independent from each other. Nevertheless, in some cases, the knowledge domains of CorpWiki may be related, and thus the expertise of a person in a specific domain might be related to the person's expertise in another. Therefore, as future work, the expert peer matching algorithm could be extended to initialize a person's expertise in a specific domain not on a predefined value, but on the basis of the expertise that the person has in related domains. This is expected to result in the need for less profile information from each user in order to achieve efficient estimations.

5.5. Enhancing CorpWiki with semantic capabilities

Another issue that could be considered as a further design option of CorpWiki, is enhancing the proposed system with semantic capabilities, along the lines of already successful semantic wiki systems, for example the Semantic MediaWiki [47]. That is, instead of containing only text, CorpWiki could include semantic annotations which will provide structured access to its content, support information reuse and enable users to organize and process the knowledge contained more efficiently [57].

5.6. Global vs. local wiki quality

The system proposed in this paper focuses on the issue of quality assurance for collective content that consists of individual article pages. However, apart from assuring the aforementioned local quality, the assurance of the global wiki quality is also an issue which should be further considered. This issue could for instance include ensuring the quality of CorpWiki in terms of interlinking, missing pages or consistency among the articles.

5.7. Negative review changes

In some instances, there is the possibility that a negative change in the review may exist, either due to a bad contribution or due to randomness. In case that this negative change is large, CorpWiki could be enhanced to inform the management in order to check upon this issue.

5.8. Motivating users to participate

One of the most important considerations for a collective intelligence system, such as the one proposed in this study, is the fact that the system can only become useful to the corporate community if it is accepted and used by the employees. Thus, for the organization to actually benefit from the advantages of a wiki-based system, like CorpWiki, a non-hierarchical and bottom-up knowledge sharing culture, where employees are willing to share their knowledge, should be embraced by both the management and the employees prior to system launch [75]. This type of culture, however, is not present in most top-down corporate environments and it cannot be changed solely by a system, no matter how well designed. Therefore in order to successfully integrate CorpWiki in the organization, the management should promote an open information sharing culture among the employees of the organization by providing them with the appropriate incentives to participate and share their knowledge.

In their recent study Malone et al. [54] suggest that although no systematic studies on the motivational issue are known to exist, however, the incorrect identification of the factors that motivate users to participate is probably the most important cause for the launch failure of a new collective intelligence system. In the aforementioned study, three main motivating factors are identified, namely the financial, the self-fulfillment and the recognition incentive. Additionally, if CorpWiki is viewed in the broader field of knowledge management solutions, the existing knowledge sharing literature can be applied to identify possible incentives to motivate participation. These incentives can be extrinsic or intrinsic motivating factors.

More specifically, extrinsic motivation could be provided in the form of tangible rewards, e.g., monetary compensation, promotion incentives or better work assignments. This type of reward has the advantage of “providing satisfaction independent of the actual activity itself” [14]. However, this type of motivation is often perceived by the employees as a controlling aspect, and as such it has been reported to produce the so-called “crowding-out effect”, i.e., the employees do not perform the task at hand unless they receive the reward [42,58]. In addition, if these rewards are provided based on individual participation, there is the risk that users may artificially inflate their contributions at the expense of quality. Thus, it has been proposed that such incentives should not be provided individually, but as group rewards, linked to the success of a group-wide knowledge sharing program [13]. Apart from increasing the employees’ benefits, the organization could reduce the individual costs of knowledge sharing by providing, for example, proper training to use the system as well as the necessary time and resources for knowledge contribution. Additionally, the management should consider expanding the concept of employee performance to include not only business results but also contributions to the knowledge of the organization [13].

Furthermore the organization should make efforts to increase the intrinsic motivation of the employees to share their knowledge. Intrinsic motivation is more effective than extrinsic, in cases of innovative work that entails the generation and transfer of knowledge, while it also appears to be self-sustained [58]. More specifically, individual non-financial rewards in the form of social recognition – e.g., top participants – have been found to increase the perceived self-efficacy of the employees and their perceived benefit from sharing their knowledge [13,42,58,77]. In addition, feedback that their contribution has been helpful to other employees is also an intrinsic motive that is positively related to increased participation [13,42]. Another means of motivating employees to participate is to promote team spirit [61], which can be achieved by creating a group identity, for example through communities of practice [17]. This factor has been found to increase the social responsibility for one’s actions and the commitment to common welfare, while it is also expected to prevent “free riding” [13,58]. It should be noted however, that intrinsic motivation is more difficult to inspire and its outcome is more uncertain.

Based on the above, and taking into consideration the advantages and disadvantages that each approach presents, an organization that wishes to integrate CorpWiki may select, according to its needs, the most appropriate extrinsic and intrinsic motivators, that will promote employee participation and use of the system.

6. Conclusion

Corporate collective intelligence is an asset, which, if harnessed and properly used, has the potential to create significant competitive advantage for an organization. This paper presents CorpWiki, an innovative self-regulating wiki-based system that aims to assure timely creation of high-quality collective intelligence content, in corporate environments. To perform this, the system searches through the human resources network of the organization and identifies the most appropriate peer employees who will assess and improve the created content. The core functionality of CorpWiki is a novel expert peer matching algorithm, based on feed-forward neural networks. Comparative evaluation results, obtained through simulation modeling, between the expert peer matching algorithm and a random peer selection policy indicate that the proposed system achieves significant improvement on the average article quality, on the number of revisions and on the time needed in order to reach the predefined quality threshold.

Concluding, the potential that CorpWiki presents in using the collective intelligence of the organization to ensure timely creation of quality content, is expected to lead to more complete corporate knowledge, to more innovative solutions and thus to an increased value for the enterprise.

References

- [1] R. Andrews, J. Diederich, A.B. Tickle, Survey and critique of techniques for extracting rules from trained artificial neural networks, *Knowledge-Based Systems* 8 (6) (1995) 373–389.
- [2] K. Balog, L. Azzopardi, M. de Rijke, Formal models for expert finding in enterprise corpora, in: *Proceedings of the 29th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2006, pp. 43–50.
- [3] K. Balog, M. de Rijke, *Associating People and Documents*, Advances in Information Retrieval, Springer, Berlin, Heidelberg, 2008.
- [4] K. Balog, M. de Rijke, Finding experts and their details in e-mail corpora, in: *Proceedings of the 15th International Conference on World Wide Web*, 2006, pp. 1035–1036.
- [5] S.L. Bang, J.D. Yang, H.J. Yang, Hierarchical document categorization with k-NN and concept-based thesauri, *Information Processing and Management* 42 (2) (2006) 387–406.
- [6] R. Bergmann, A. Freßmann, K. Maximini, R. Maximini, T. Sauer, *Case-Based Support for Collaborative Business*, Advances in Case-Based Reasoning, Springer, Berlin, Heidelberg, 2006.
- [7] J.E. Blumenstock, Size matters: word count as a measure of quality on wikipedia, in: *Proceedings of the 17th International Conference on World Wide Web*, 2008, pp. 1095–1096.
- [8] S. Bowen, F. Maurer, Process support and knowledge management for virtual teams doing agile software development, in: *Proceedings of the 26th Annual International Computer Software and Applications Conference*, 2002, pp. 1118–1120.
- [9] S. Brin, L. Page, The anatomy of a large-scale hypertextual Web search engine, *Computer Networks and ISDN Systems* 30 (1–7) (1998) 107–117.
- [10] M. Buffa, F. Gandon, SweetWiki: semantic web enabled technologies in Wiki, in: *Proceedings of the 2006 International Symposium on Wikis*, 2006, pp. 69–78.
- [11] J. Bughin, J. Manyika, How businesses are using Web 2.0: a McKinsey global survey, *The McKinsey Quarterly* (2007).
- [12] J.R. Bughin, How companies can make the most of user-generated content, *The McKinsey Quarterly* (2007).
- [13] A. Cabrera, E.F. Cabrera, Knowledge-sharing dilemmas, *Organization Studies* 23 (5) (2002) 687–710.
- [14] B.J. Calder, B.M. Staw, Self-perception of intrinsic and extrinsic motivation, *Journal of Personality and Social Psychology* 31 (1975) 599–605.
- [15] R.J. Chapman, *Simple Tools and Techniques for Enterprise Risk Management*, John Wiley & Sons, 2006.
- [16] R.C. Chen, C.H. Hsieh, Web page classification based on a support vector machine using a weighted vote schema, *Expert Systems with Applications* 31 (2) (2006) 427–435.
- [17] R.S. Chen, C.H. Hsiang, A study on the critical success factors for corporations embarking on knowledge community-based e-learning, *Information Sciences* 177 (2) (2007) 570–586.
- [18] I. Cil, O. Alpturk, H.R. Yazgan, A new collaborative system framework based on a multiple perspective approach: IntelTeam, *Decision Support Systems* 39 (4) (2005) 619–641.
- [19] G. Cybenko, Approximation by superpositions of a sigmoidal function, *Mathematics of Control Signals and Systems (MCSS)* 2 (4) (1989) 303–314.
- [20] E. Da Lio, L. Fraboni, T. Leo, TWiki-based facilitation in a newly formed academic community of practice, in: *Proceedings of the 2005 International Symposium on Wikis*, 2005, pp. 85–111.
- [21] S. Damart, L.C. Dias, V. Mousseau, Supporting groups in sorting decisions: methodology and use of a multi-criteria aggregation/disaggregation DSS, *Decision Support Systems* 43 (4) (2007) 1464–1475.
- [22] T. Davenport, L. Prusak, *Working Knowledge: How Organizations Manage What They Know*, Harvard Business School Press, 2000.
- [23] J.S. Dhaliwal, L.L. Tung, Using group support systems for developing a knowledge-based explanation facility, *International Journal of Information Management* 20 (2) (2000) 131–149.
- [24] X. Ding, C. Danis, T. Erickson, W.A. Kellogg, Visualizing an enterprise Wiki, in: *Proceedings of the Extended Abstracts on Human factors in Computing Systems*, 2007m pp. 2189–2194.
- [25] H. Fang, C. Zhai, Probabilistic Models for Expert Finding, *Advances in Information Retrieval*, Springer, Berlin, Heidelberg, 2007.
- [26] K. Funahashi, On the approximate realization of continuous mappings by neural networks, *Neural Networks* 2 (3) (1989) 183–192.
- [27] M.T. Hagan, M.B. Menhaj, Training feedforward networks with the Marquardt algorithm, *IEEE Transactions on Neural Networks* 5 (6) (1994) 989–993.
- [28] P.Y. Hao, J.H. Chiang, Y.K. Tu, Hierarchically SVM classification based on support vector clustering method and its application to document categorization, *Expert Systems with Applications* 33 (3) (2007) 627–635.
- [29] H. Hasan, C.C. Pfaff, Emergent conversational technologies that are democratising information systems in organisations: the case of the corporate Wiki, in: *Proceedings of the Information Systems Foundations (ISF): Theory, Representation and Reality conference*, 2006.
- [30] H. Hasan, C.C. Pfaff, The Wiki: an environment to revolutionise employees' interaction with corporate knowledge, in: *Proceedings of the 20th Conference of the Computer–Human Interaction Special Interest Group (CHISIG) of Australia on Computer–Human Interaction: Design, Activities, Artefacts and Environments*, 2006, pp. 377–380.
- [31] W.D. Haseman, D.L. Nazareth, S. Paul, Implementation of a group decision support system utilizing collective memory, *Information and Management* 42 (4) (2005) 591–605.
- [32] S. Haykin, *Neural Networks: A Comprehensive Foundation*, Prentice-Hall, 1999.
- [33] C. Heckscher, P.S. Adler, *The Firm as a Collaborative Community: Reconstructing Trust in the Knowledge Economy*, Oxford University Press, 2007.
- [34] A. Hester, Innovating with organizational wikis: factors facilitating adoption and diffusion of an effective collaborative knowledge management system, in: *Proceedings of the 2008 ACM SIGMIS CPR Conference on Computer Personnel Doctoral Consortium and Research*, 2008, pp. 161–163.
- [35] R. Hoffmann, A wiki for the life sciences where authorship matters, *Nature Genetics* 40 (9) (2008) 1047–1051.
- [36] K. Hornik, Approximation capabilities of multilayer feedforward networks, *Neural Networks* 4 (2) (1991) 251–257.
- [37] K. Hornik, Some new results on neural network approximation, *Neural Networks* 6 (9) (1993) 1069–1072.
- [38] K. Hornik, M. Stinchcombe, H. White, Multilayer feedforward networks are universal approximators, *Neural Networks* 2 (5) (1989) 359–366.
- [39] M. Hu, E.-P. Lim, A. Sun, H.W. Lauw, B.-Q. Vuong, On improving wikipedia search using article quality, in: *Proceedings of the 9th Annual ACM International Workshop on Web Information and Data Management*, 2007, pp. 145–152.
- [40] IRM, *A Risk Management Standard*, The Institute of Risk Management, London, 2002.
- [41] R. Jain, D. Chiu, W. Haw, A quantitative measure of fairness and discrimination for resource allocation in shared systems, DEC Technical Report TR-301, Digital Equipment Corporation, 1984.
- [42] A. Kankanhalli, B.C. Tan, K.K. Wei, Contributing knowledge to electronic knowledge repositories: an empirical investigation, *MIS Quarterly* 29 (1) (2005) 113–143.
- [43] E. Kapetanios, Quo Vadis computer science: from turing to personal computer, personal content and collective intelligence, *Data and Knowledge Engineering* 67 (2) (2008) 286–292.
- [44] A. Keen, *The Cult of the Amateur: How Today's Internet is Killing Our Culture*, Doubleday Business, 2007.
- [45] L. Kester, P.v. Rosmalen, P. Sloep, F. Brouns, M. Kone, R. Koper, Matchmaking in learning networks: bringing learners together for knowledge sharing, *Interactive Learning Environments* 5 (2007) 117–126.

- [46] V. Klabbers, N. Kruidierink, Users and tools: the art of matchmaking. Challenges in choosing appropriate online collaboration tools for development professionals and practitioners, *Knowledge Management for Development Journal* 3 (1) (2007) 5–25.
- [47] M. Krötzsch, D. Vrandečić, M. Völkel, Semantic MediaWiki, *The Semantic Web – ISWC 2006*, 2006.
- [48] M. Krötzsch, D. Vrandečić, M. Völkel, H. Haller, R. Studer, Semantic wikipedia, web semantics: science, Services and Agents on the World Wide Web 5 (4) (2007) 251–261.
- [49] L.F. Lai, A knowledge engineering approach to knowledge management, *Information Sciences* 177 (19) (2007) 4072–4094.
- [50] C. Lange, Towards scientific collaboration in a semantic wiki, in: *Proceedings of the European Semantic Web Conference 2007, Workshop “Bridging the Gap between Semantic Web and Web 2.0”*, 2007, pp. 119–126.
- [51] C. Lange, T. Hastrup, S. Corlosquet, Improving mathematical knowledge items by acting on issue-based community feedback, in: *Proceedings of the 2nd Workshop on Scientific Communities of Practice*, 2008.
- [52] C. Lin, H. Zhou, Z. Huang, W. Wang, REC: a novel model to rank experts in communities, in: *Proceedings of the 9th International Conference on Web-Age Information Management*, 2008, pp. 301–308.
- [53] A. Majchrzak, C. Wagner, D. Yates, Corporate wiki users: results of a survey, in: *Proceedings of the 2006 International Symposium on Wikis*, 2006, pp. 99–104.
- [54] T.W. Malone, R. Laubacher, C. Dellarocas, Harnessing crowds: mapping the genome of collective intelligence, CCI Working Paper, 2009-001, MIT Center for Collective Intelligence, 2009.
- [55] J.A. Nickerson, T.R. Zenger, A knowledge-based theory of the firm – the problem-solving perspective, *Organization Science* 15 (6) (2004) 617–632.
- [56] D.E. O’Leary, Wikis: from each according to his knowledge, *Computer* 41 (2) (2008) 34–41.
- [57] E. Oren, J.G. Breslin, S. Decker, How semantics make better wikis, in: *Proceedings of the 15th International Conference on World Wide Web, Conference*, 2006.
- [58] M. Osterloh, B.S. Frey, Motivation, knowledge transfer and organizational forms, *Organization Science* 11 (5) (2000) 538–550.
- [59] P.B.S.P. Van Rosmalen, F. Brouns, L. Kester, A. Berlanga, M. Bitter, R. Koper, A model for online learner support based on selecting appropriate peer tutors, *Journal of Computer Assisted Learning* 24 (6) (2008) 483–493.
- [60] D. Petkova, W.B. Croft, Hierarchical language models for expert finding in enterprise corpora, in: *Proceedings of the 18th IEEE International Conference on Tools with Artificial Intelligence*, 2006, pp. 599–608.
- [61] S. Petter, V. Vaishnavi, Facilitating experience reuse among software project managers, *Information Sciences* 178 (7) (2008) 1783–1802.
- [62] D. Richards, A social software/Web 2.0 approach to collaborative knowledge engineering, *Information Sciences* 179 (15) (2009) 2515–2523.
- [63] D.E. Rumelhart, G.E. Hinton, R.J. Williams, Learning internal representations by error, propagation, Parallel distributed processing: explorations in the microstructure of cognition, *Foundations*, vol. 1, MIT Press, Cambridge, MA, 1986.
- [64] D.E. Rumelhart, G.E. Hinton, R.J. Williams, Learning representations by back-propagating errors, *Nature* 323 (1) (1986) 533–536.
- [65] P. Serdyukov, D. Hiemstra, Modeling Documents as Mixtures of Persons for Expert Finding, *Advances in Information Retrieval*, Springer, Berlin, Heidelberg, 2008.
- [66] V. Sinha, B. Sengupta, S. Chandra, Enabling collaboration in distributed requirements management, *Software IEEE* 23 (5) (2006) 52–61.
- [67] K. Stein, C. Hess, Does it matter who contributes: a study on featured articles in the german wikipedia, in: *Proceedings of the 18th Conference on Hypertext and Hypermedia*, 2007, pp. 171–174.
- [68] B. Stvilia, M. Twidale, L. Smith, L. Gasser, Assessing information quality of a community-based encyclopedia, in: *Proceedings of the International Conference on Information Quality*, 2005, pp. 442–454.
- [69] J. Surowiecki, *The wisdom of the crowds: why the many are smarter than the few*, Abacus, 2004.
- [70] V. Tam, A. Santoso, R. Setiono, A comparative study of centroid-based, neighborhood-based and statistical approaches for effective document categorization, in: *Proceedings of the 16th International Conference on Pattern Recognition*, 2002, pp. 235–238.
- [71] D. Tapscott, A. Williams, *Wikinomics: How Mass Collaboration Changes Everything*, Atlantic Books, 2008.
- [72] S.B. Thrun, Extracting provably correct rules from artificial neural networks, Technical Report IAI-TR-93-5, Institute für Informatik III, Universität Bonn, 1993.
- [73] A.J.C. Trappey, F.-C. Hsu, C.V. Trappey, C.-I. Lin, Development of a patent document classification and search platform using a back-propagation network, *Expert Systems with Applications* 31 (4) (2006) 755–765.
- [74] D. Vrandečić, M. Krötzsch, Reusing ontological background knowledge in semantic wikis, in: *Proceedings of the First Workshop on Semantic Wikis – From Wikis to Semantics*, 2006, pp. 16–30.
- [75] C. Wagner, Breaking the knowledge acquisition bottleneck through conversational knowledge management, *Information Resources Management Journal* 19 (1) (2006) 70–83.
- [76] C. Wagner, N. Bolloju, Supporting knowledge management in organisations with conversational technologies: discussion forums weblogs and wikis editorial preface, *Journal of Database Management* 16 (2) (2005) 1–8.
- [77] M.M. Wasko, S. Faraj, Why should I share? Examining social capital and knowledge contribution in electronic networks of practice.
- [78] D. West, S. Dellana, Diversity of ability and cognitive style for group decision processes, *Information Sciences* 179 (5) (2009) 542–558.
- [79] D.M. Wilkinson, B.A. Huberman, Cooperation and quality in wikipedia, in: *Proceedings of the 2007 International Symposium on Wikis*, 2007, pp. 157–164.
- [80] J. Wolfers, E. Zitzewitz, Prediction markets, *Journal of Economic Perspectives* 18 (2) (2004) 107–126.
- [81] J. Zhang, J. Tang, L. Liu, J. Li, A Mixture Model for Expert Finding, *Advances in Knowledge Discovery and Data Mining*, Springer, Berlin, Heidelberg, 2008.
- [82] H.T. Zheng, B.Y. Kang, H.G. Kim, Exploiting noun phrases and semantic relationships for text document clustering, *Information Sciences* 179 (13) (2009) 2249–2262.