Expediting Expertise: Supporting Informal Social Learning in the Enterprise

Jennifer Lai, Jie Lu, Shimei Pan, Danny Soroker, Mercan Topkara, Justin Weisz, Jeff Boston, Jason Crawford
IBM Thomas J. Watson Research Center
1101 Kitchawan Road, Yorktown Heights, NY 10598
{jlai, jielu, shimei, soroker, mtopkara, jweisz, daddyb, ccjason}@us.ibm.com

ABSTRACT
In this paper, we present Expediting Expertise, a system designed to provide structured support to the otherwise informal process of social learning in the enterprise. It employs a data-driven approach where online content is automatically analyzed and categorized into relevant topics, topic-specific user expertise is calculated by comparing the models of individual users against those of the experts, and personalized recommendation of learning activities is created accordingly to facilitate expertise development. The system's UI is designed to provide users with ongoing feedback of current expertise, progress, and comparison with others. Learning recommendation is visualized with an interactive treemap which presents estimated return on investment and distance to current expertise for each recommended learning activity. Evaluation of the system showed very positive results.

Author Keywords
Expertise, learning, personalized, social, informal, assessment, recommendation.

ACM Classification Keywords
H.5.m [Information Interfaces and Presentation]: Miscellaneous.

General Terms
Algorithms; Design; Experimentation.

INTRODUCTION
Expertise development is critical to knowledge workers who continually improve their knowledge and skills to stay current with changing technologies. It is typical for knowledge workers to learn on the job through informal learning as much as (and sometimes more than) from formal learning [31]. Informal learning through knowledge sharing and interaction with social software (e.g. online communities, wikis and file sharing) has emerged as an important form of learning to help knowledge workers increase their expertise [4, 18, 22]. Informal social learning has the advantage of providing a more personalized, open, dynamic, and knowledge-pull model of learning, as opposed to the one-size-fits-all, centralized, static, top-down and knowledge-push model of formal curriculum-based learning [11]. However, the effectiveness of informal learning in an online social environment is often hindered by two problems. First, because voluntary contribution accounts for a large amount of content in social software, the available content often lacks organization and quality control [23]. Without guidance, it is challenging for users to identify suitable learning materials. Second, no feedback is given to learners regarding their current expertise level and distance to their target level of expertise. As a result, learners may question their progress and stop before reaching their goal.

To address these problems, we have built the Expediting Expertise (EE) system that aims to provide structured support to informal social learning in the enterprise. The approach taken by the system is data-driven and grounded in Social Learning theory, which suggests that people learn in a social context by modeling and observing others (in our case, experts) [6, 8, 16]. With this approach, online content is automatically analyzed and categorized into relevant topics to provide users with an overview of the main subjects contained in the target domain. Individuals' expertise scores in these topics are computed by comparing activity and word vector based models built from users' online activities with the same models for experts to help users identify expertise gaps and give them ongoing feedback during expertise development. Learning recommendation is created by taking into consideration content relevancy along with social factors (e.g., popularity, ratings, and related expert activities), and further personalized based on individual user models.

The benefits of this data-driven and social-oriented approach are multifold. First, it requires little effort for domain customization, making the system easily applicable to a wide range of subjects and skills. Second, it allows the natural evolution of domain knowledge (e.g. emerging topics, expanding expertise) to be automatically captured and seamlessly handled, which provides adaptability and flexibility. Third, it creates a concrete social measure of

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.
IUI '14, February 24–27, 2014, Haifa, Israel.
Copyright 2014 ACM 978-1-4503-2184-6/14/02…$15.00.
http://dx.doi.org/10.1145/2557500.2557539

expertise by modeling and quantifying social participation and contribution. This not only encourages learners to actively contribute to the growth of their expertise, but also motivates experts to engage more in knowledge sharing and increase their digital eminence.

Our ultimate goal is to create an adaptive, engaging, and rewarding social learning environment that provides personalized learning feedback and guidance with an intuitive interface. For the work reported in this paper, we focus on the following aspects: dynamically modeling users' social learning activities; computing and presenting expertise scores to provide continuous feedback; and creating a personalized learning roadmap that is easy to follow and interact with.

Our work has three main contributions:

1. We developed a system built on top of existing enterprise social software to provide structured support to informal social learning. It is the first of this kind known to us.
2. We proposed a data-driven approach to automatically collect and organize learning materials and objectively assess users' expertise to provide instant and continuous feedback. It enables a learning solution that is more adaptive, personalized, and less costly to build and maintain than traditional e-learning systems.
3. We designed a web UI that employs interactive and intuitive graphical visualizations to present expertise assessment and learning recommendation to facilitate effective system-user interaction.

RELATED WORK

Traditional E-Learning Systems

Learning management system (LMS) is the traditional approach to e-learning. Learning in LMS is organized as courses (a.k.a. curriculum). Traditional LMS is mainly used for formal learning and is considered incompetent for supporting e-learning in the new era when learning is viewed as a self-governed, problem-based and collaborative social process [12].

Intelligent tutoring system (ITS) [17, 25, 26] is developed to support learning of a particular subject or skill. ITS dynamically builds a student model as the learning process advances and compares it against a pre-determined domain model, which serves as a source of expert knowledge and a standard for evaluating the student's performance. Based on the comparison result, the system provides timely customized instruction or feedback to the learner. ITS is expensive to develop and implement in general [2, 24, 25, 28]. It is particularly difficult and time-consuming to develop the domain model (i.e., the concepts, rules, and problem-solving strategies of the domain to be learned), which requires structured and well-defined domain knowledge [24, 25]. Only recently Educational Data Mining emerged as a research field that explores the application of data mining and machine learning techniques for domain modeling [14]. Similar to ITS, our system dynamically builds a user model for each learner and compares it against a domain model to provide personalized feedback and recommendation. The data-driven approach adopted by our system makes it much easier and less costly to create and dynamically update the domain model.

Social Learning

Web 2.0 and collaboration technologies (e.g., social networking, wiki, blog, forum) have gained increasing popularity in social learning and brought profound impact to education [19, 29]. Existing collaboration and social software tools primarily focus on basic content sharing features such as file sharing, rating, commenting, tagging and searching, which are not tailored specifically to support learning. The work presented in [32] pointed out the needs for new social learning technologies to help learners find the right content, connect with the right people, and to motivate participation. Vassileva et al. described different approaches to motivate users to participate in social computing applications [33], and explored recommendation techniques for e-learning focused discussion forums [1]. Despotakis et al. conducted automatic semantic analysis to capture viewpoints and enable visual exploration of the conceptual spaces of user-generated content [13]. In comparison, the focus of our work is on supporting informal social learning with automatic and dynamic expertise assessment and progress checking, and creation and visualization of personalized learning roadmaps.

Expert Finding

Expert finding systems have been developed to answer questions like "who are the experts on topic X?" Two main approaches are employed to automatically identify experts [5, 27, 30]. The first approach models one's expertise based on all the materials s/he is associated with, and then uses information retrieval methods to rank individual expertise models by their relevancy to the topic. The second approach first locates materials highly relevant to the topic, and then builds and ranks expertise models only for those who are associated with these materials. Despite the differences, both approaches have demonstrated the effectiveness of using content (text and metadata) associated with a person to assess this person's expertise on different topics. In comparison, our system is capable of identifying experts for topics derived from available content based on user expertise models created from activities associated with this content. However, the main purpose of this functionality is to facilitate social interactions in learning, not to find experts for specific Q&A.

Recommender Systems

The field of recommender systems is very rich and over years, many algorithms (e.g., collaborative filtering [9, 20], content [3, 21] and social based [10] recommendation) have been developed to support applications ranging from movie, music and book recommendation to news filtering. In this

work, we primarily use content and social based strategies in learning recommendation. Unlike typical recommender systems that suggest a list of ranked items, our system characterizes the utility of each recommended learning item along two dimensions: return on investment and distance to the user's current expertise. It then employs an interactive treemap visualization to present both dimensions simultaneously.

MODEL REPRESENTATIONS

Topic Model

Topics provide a form of guidance to help users navigate through the learning space. The system employs a three-step process to create a model representation that associates individual documents with topics. In the first step, the system uses unsupervised statistical text analytics to infer topic distribution from the collection of documents in the target domain. This automatic topic analysis step reduces the manual effort needed to organize content. Topics created at this step are referred to as *system topics*. In the second step, domain leaders map system topics to a set of *labeled topics.* Each labeled topic corresponds to one or more system topics. This human validation step ensures the quality of the topic model and produces easier to understand topic labels and descriptions. In the third step, the system creates a topic-specific word vector for every <document, labeled topic> pair. This representation serves as the basis for creating topic-specific user and domain models. Since both expertise assessment and learning recommendation are conducted at the labeled topic level, labeled topics are simply referred to as *topics* in the later sections of this paper unless otherwise noted.

Specifically, in the first step, Latent Dirichlet Allocation (LDA) [7] is conducted over the collection of documents to infer topic assignment for every word token in each document. Based on the topic assignment, two types of information are derived: topic-specific word distribution Φ and document-specific topic distribution Θ.

In the second step, domain leaders inspect top-ranked words and documents associated with each of the system topics (derived from Φ and Θ respectively), and label them to produce a set of labeled topics. Multiple system topics may be given the same label, essentially merging them to a single topic. Incoherent system topics may be removed.

In the third step, a word vector representation $V_{d,t}$ is created for every document d on every labeled topic t. The i^{th} dimension of $V_{d,t}$ corresponds to a unique word w_i in the vocabulary of the document collection. Let k_j denote the j^{th} word token in document d, $T_s(k_j)$ denote the system topic assigned to k_j in the first step, $T(T_s(k_j))$ denote the labeled topic associated with $T_s(k_j)$ in the second step, then:

$$V_{d,t}(i) = \sum_j (I(w_i = k_j) \times I(T(T_s(k_j)) = t)) \quad (1)$$

where $I(x = y)$ denotes the indicator function that returns 1 when x is equal to y and 0 otherwise. In other words, $V_{d,t}(i)$ is the count of all the words in d that are equal to w_i and have a topic assignment equal to t. The overall topic strength $S_{d,t}$ of document d for labeled topic t is the L1-norm (sum) of $V_{d,t}$:

$$S_{d,t} = \left\| V_{d,t} \right\|_1 \quad (2)$$

The system uses the above formula instead of Θ to compute the topic strength in order to give equal weight to each topic-related word, regardless of the length of the document that contains the word.

User Model

User modeling infers users' expertise based on the content associated with their knowledge sharing (e.g. author, comment) and acquisition (e.g. read, download) activities in the enterprise social environment. Even though models built this way may not fully capture all aspects of individuals' expertise, previous success of expert finding systems [5, 27, 30] show that they provide more reliable and granular representations of expertise than the expertise profiles created based on manual/self assessment, and we believe that they are sufficient for the purpose of providing users feedback and guidance in social learning.

Individual user models are built at the topic level to represent topic-specific user expertise. The model of user u for topic t has two parts. The first part captures the strength of various user activities in relation to the topic. It is represented with a multi-dimensional activity vector $A_{u,t}$ where the m^{th} dimension corresponds to an activity type a_m (e.g. create-wiki, read-blog), and its value measures to what degree the user's activities of type a_m is about topic t:

$$A_{u,t}(m) = L(\sum_{d \in D(u,a_m)} S_{d,t}) \times \omega(a_m) \quad (3)$$

$$L(x) = \frac{1}{1 + a \times e^{-b \times x}} \quad (4)$$

where $D(u, a_m)$ denotes the set of documents associated with user u through activity type a_m, $S_{d,t}$ is the topic strength of document d for topic t (see the previous section), L is a logistic function to control the scale of the value, and $\omega(a_m)$ is the weight of activity type a_m. Activity type weight reflects the different levels of expertise required/implied by different types of activities. Its value is the product of empirically determined action type weight and content type weight. For instance, the action type of "create" has a higher weight (e.g. 1.0) than "read" (e.g. 0.4), and the content type of "wiki" has a higher weight (e.g. 0.8) than "forum" (e.g. 0.6).

The second part of the topic-specific user model captures the user's familiarity with the terminologies of the topic. It is represented with a word vector $W_{u,t}$ where the i^{th} dimension corresponds to a unique word w_i:

$$W_{u,t}(i) = \sum_m (\omega(a_m) \times \sum_{d \in D(u,a_m)} V_{d,t}(i)) \quad (5)$$

In other words, $W_{u,t}$ is the weighted sum of the topic-specific word vectors for all the documents associated with the user.

Topic-Specific Domain Model

A topic-specific domain model is an amalgam of experts' activities and knowledge related to a specific topic in the target domain. It is created by aggregating individual topic-specific user models of those who are identified as experts for the topic. Specifically, for each topic *t*, the activity and word vectors $A_{u,t}$ and $W_{u,t}$ of the experts for this topic are averaged to create the vector representations $A_{e,t}$ and $W_{e,t}$, which are stored in the topic-specific domain model for *t*.

The identification of experts for a topic is a combination of automatic and manual effort. First, the system ranks all users based on the L1-norm values of their activity vectors for the topic and outputs a list of top-ranked users as candidate experts. Second, domain leaders select one or more candidates from the list and designate them as experts. This procedure ensures that the experts (whom all others are compared against) are both well established and have a prominent digital presence in the domain.

EXPERTISE ASSESSMENT

The goal of expertise assessment is to create a quantified measure of expertise to help users keep track of their current status, progress, and distance to target expertise when they engage in social learning activities to improve their expertise. We chose to use a relative measure of expertise as it is easier to interpret and understand in the social context. Specifically, expertise assessment of a user for a topic takes the form of an expertise score with its value ranging from 0% to 100% to indicate the user's relative expertise level when compared against the aggregate expertise of the experts for the topic (a.k.a. topic-specific domain model). Below we describe the method to calculate expertise scores and the research work around the presentation of expertise scores at the user interface.

Creation

The system evaluates a user's expertise for a topic from three aspects. The first aspect assesses the user's overall activity level for the topic in relative to the aggregate expert activity level:

$$S_1 = \min\{1, \frac{\|A_{u,t}\|_1}{\|A_{e,t}\|_1}\} \quad (6)$$

where $A_{u,t}$ denotes the activity vector from the user's model for topic *t*, $A_{e,t}$ denotes the aggregate expert activity vector from the topic-specific domain model, and $\|\bullet\|_1$ denotes the L1-norm (sum) of the vector.

The second aspect compares the user and the experts at a finer granularity by measuring how closely the user's activities resemble the aggregate expert activities. It is calculated as:

$$S_2 = 1 - \sqrt{\frac{\sum_m (\max\{0, A_{e,t}(m) - A_{u,t}(m)\})^2}{\sum_m A_{e,t}(m)^2}} \quad (7)$$

The third aspect estimates the user's familiarity with the topic by calculating the similarity between the user's topic-specific word vector $W_{u,t}$ and the aggregate expert's topic-specific word vector $W_{e,t}$:

$$S_3 = 1 - \sqrt{\frac{\sum_i (\max\{0, W_{e,t}(i) - W_{u,t}(i)\})^2}{\sum_i W_{e,t}(i)^2}} \quad (8)$$

The final expertise score *S* of a user *u* for a topic *t* is the average of S_1, S_2, and S_3.

New user activities and/or expert activities cause expertise scores to change. Scores can not only increase, but also decrease. For example, when there are new expert activities but no new activities from a user, the user gets into the situation of "if you are not getting better, you are getting worse."

Presentation

The user interface design for displaying expertise scores aims to find a graphical presentation that maximizes the users' trust in the system's expertise assessment and their perceived utility of the system. We explored the effects of different graphical presentations of expertise scores on trust and perceived system utility with a controlled experiment (N=60). Our results indicated that the users' trust in the system's assessment of their expertise increased when the scores were presented in an exact percentage format (e.g. 35%) instead of a range format (e.g. 30-40%). We also tested the effects of providing social comparison, and observed that the users found more utility in the system when they were provided with an average score next to their scores. Whether the average score was above or below their scores did not have a significant effect on their response. These results on increased perception in the utility of the system align with the literature on social comparison [8].

Based on the results of this experiment, our design decision was to present the expertise score of a user for a topic in an exact percentage format, together with the average score calculated based on the expertise scores of all users with non-zero scores for the topic.

To further apply ideas from social learning, the design of the user interface includes a list of experts for each topic. Names and expertise scores of these experts are displayed under their pictures (hyperlinked to their profiles in an enterprise social network), informing users of who the experts are and how far his/her expertise score for a topic is from those of the experts. This design decision was based on several social learning concepts such as being able to find models (experts), as well as providing a way to easily reach out to experts to ask questions.

Figure 1 illustrates how we implemented the "My Interests" page at the user interface based on these design decisions. The page displays four types of information: 1) the topics

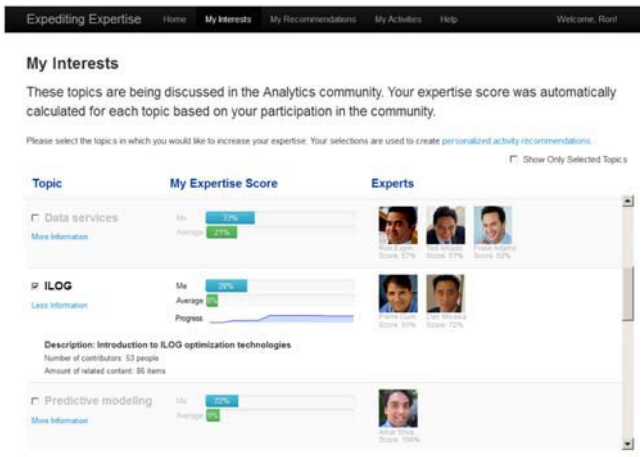


Figure 1. EE user interface – "My Interests"

available for learning, 2) the user selection of the topics s/he is interested in learning (which can be changed at any time during learning), 3) topic-specific expertise scores, and 4) the experts for each topic. All the information is organized into sections, where each section corresponds to one topic, listed in alphabetical order based on topic labels.

Clicking on "More Information" below the label of a topic expands the section for this topic to display a description of the topic, the number of documents related to the topic, and the number of content contributors for the topic.

The charts in the central portion of each topic section show the user's expertise score for the topic, along with the average expertise score as a comparison. The interface provides a sparkline of the progress of the user, which gives the user an idea about the changes to her/his expertise score over time. More detail on the exact value and time of those score changes can be obtained by hovering over different points on the progress chart. As the user progresses through learning, the scores are dynamically updated.

The "My Activities" page (not shown here due to space limitation) lists all of the user's activities recorded by the system. This helps the user track his/her progress and gives an idea of what are included in the user model used to compute his/her expertise scores. The user can choose to have the list categorized by action type (e.g. read, author), or content type (e.g. wiki, blog, forum, file).

LEARNING RECOMMENDATION

The goal of learning recommendation is to create personalized learning roadmaps that can provide users guidance on what to learn and help them get the most learning benefits out of the time and effort spent. The recommendation is on a topic basis and takes the form of learning activities. Each recommended learning activity is associated with two different dimensions of information: *ROI* (**r**eturn **o**n **i**nvestment) to provide an estimate of "the bang for the buck" provided by this activity, and *DUCE* (**d**istance to the **u**ser's **c**urrent **e**xpertise) to give some idea about how familiar the user is with this activity. Both *ROI* and *DUCE* values can change over time as a result of the user's evolving expertise and the related social learning activities of other users. Below we describe the algorithm for creating the recommended learning activities, and the design for presenting them at the user interface.

Creation

To create personalized learning recommendation for a topic, the system first generates two candidate activity lists, one for reading activities and the other for creation activities. While reading activities aim to help users acquire knowledge, the purpose of creation activities is mainly to encourage users to demonstrate and share expertise for others to benefit from social learning. The system ranks the activities from both lists together based on their *ROI* values and recommends the top-ranked activities (up to a value or rank threshold). Below we describe separately the creation of each list and the calculation of the *ROI* and *DUCE* values for each recommended activity.

Reading Activity Recommendation

The *ROI* value of a reading activity for a topic is a combination of the following factors related to the recommended document to read:

- Relevance *R*: the topic strength $S_{d,t}$ of the document
- Importance *I*: the weighted count of expert activities associated with the document, with each count weighted by the action type (e.g., read, create)
- Popularity *P*: the count of user activities associated with the document by all users
- Average rating *A*: the product of the average relevance rating and the average helpfulness rating of the document
- Type *T*: the weight (empirically determined) of the content type (e.g. wiki, blog)
- Size *S*: the length (word count) of the document

The value of each factor is scaled to be between 0 and 1 using Formula 4, where the normalizing parameters *a* and *b* are set to different values for different factors.

The *ROI* value is calculated as:

$$ROI = \frac{R \times I \times P \times A \times T}{S} \qquad (9)$$

To measure the familiarity of a user *u* with a recommended document *d*, we compute *DUCE* (distance to the user's current expertise) as the distance between the user's topic-specific word vector $W_{u,t}$ and the topic-specific word vector $V_{d,t}$ of the recommended document:

$$DUCE = \sqrt{\frac{\sum_i (\max\{0, V_{d,t}(i) - W_{u,t}(i)\})^2}{\sum_i V_{d,t}(i)^2}} \qquad (10)$$

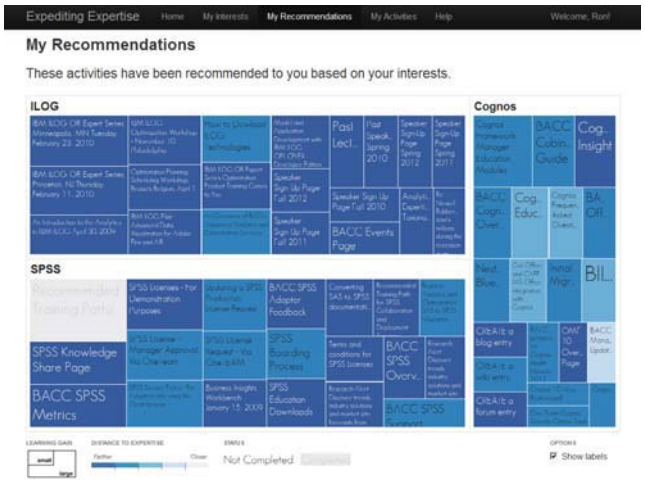


Figure 2. EE user interface – "My Recommendations"

Creation Activity Recommendation

The system recommends four types of creation activity: create a blog entry, create a wiki page, create a new forum post, and create a reply to an existing forum post. For the first three types of creation activity, because there is no content associated with them, the *ROI* value is calculated as the product of the user's expertise score for the topic and the popularity of the data source (i.e. total number of user activities) where the content is to be created. The idea is to encourage (via a high *ROI* value) users with high expertise scores to share their knowledge by creating content in the data source of high popularity and potentially high impact. The *DUCE* value is the inverse of the normalized count of the same type of activities by the user on the topic, assuming that the user is more familiar with a creation activity if s/he has conducted similar activities for the topic previously.

The last type of creation activity (i.e. create a reply to an existing forum post) is recommended to users with expertise scores above a threshold for the topic. Content from the existing forum post is used to calculate the *ROI* and *DUCE* values using the methods described in the previous section.

Presentation

The user interface design for presenting learning recommendation has two main goals. First, it aims to provide an easy-to-follow visual guidance to help the users to prioritize the learning activities for improving expertise. Second, the presentation aims to give the users an overview of their expertise strengths and gaps across topics, so they can easily assess current status and track progress as the presentation gets updated over time.

With the above goals in mind, we chose to present learning recommendation in the form of an interactive treemap for two reasons. First, compared with traditional list views, a treemap provides a more compact and visually appealing display which gives the users more information at a glance. Second, it enables simultaneous presentation of multiple data dimensions with size, color, and grouping, which can be utilized in a straightforward manner to display the *ROI*, *DUCE*, and topic information.

Figure 2 illustrates how we implemented the "My Recommendations" page at the user interface. The recommendation treemap groups learning activities by topic and renders only activities for the user-selected topics. Each recommended learning activity is depicted as a box in the treemap, labeled with document title if it is a reading activity or activity type such as "create a wiki" if it is a creation activity. The *ROI* and *DUCE* values of the activity are mapped to the size and color dimensions of the corresponding box respectively. Bigger boxes correspond to greater *ROI* values, and darker boxes correspond to greater *DUCE* values (different *DUCE* values are binned with each bin mapped to a different shade of the same color, which is blue in our case). Completed activities are marked gray to differentiate them from those that had not yet been completed. As a result, when a user completes more activities and the distance between the recommended activities and the user's expertise becomes smaller due to increased user expertise, the color distribution of the treemap gets lighter, which resembles a cleaner plate. This makes it easy for the user to quickly assess his/her progress. Our design choice for the color distribution also enables the system to avoid any restrictions for color-blind users.

Clicking on a recommendation box opens up an overlay modal window which provides an overview of the recommended learning activity. It displays a snippet from the content (for reading activity), content type (e.g., wiki, blog, forum), degree of estimated learning gain (return on investment), popularity, and an estimation of the time (in minutes) needed to complete the activity if it is a reading activity (calculated based on the word count of the associated document). The overview is intended to let the user know more about the recommended activity to help him/her decide whether to follow the recommendation to complete it.

For a reading activity, there is a "Read Now" button which when clicked on opens a new browser tab to display the content to read. For a creation activity, the window shows a "Create Now" button instead, which takes the user to the online data source where s/he can create new content. The user can spend as much time as s/he wants/needs to complete the activity in the new browser tab.

After the user is finished with the activity, s/he can simply close the new tab and return to the "My Recommendations" page, where the same overlay window is updated to enable user ratings of the activity in terms of its relevancy to the topic and helpfulness to the user's learning. The ratings will be incorporated into the calculation when the system creates new recommendations for this user and others. Another update to the overlay window is its buttons. For a reading

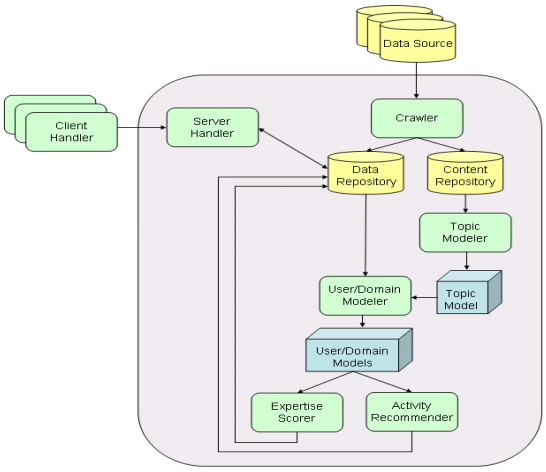


Figure 3. EE system architecture and flow

activity, the user is required to explicitly mark the activity as completed by clicking on the "Mark as Read and Close" button (which replaces the "Read Now" button) to close the overlay window. The treemap is automatically updated to color the corresponding box gray, indicating that the activity has been completed. For a creation activity, the user is not required to mark it as completed as creation activities are automatically detected and tracked by the system. The user can simply click on the "Close" button to close the overlay window. The color of the corresponding box will only be updated after the system recalculates the *DUCE* value of the recommended creation activity based on the updated user model which includes the new creation activity. It turns lighter over time as the user creates more content, but never gray (which signals "completed") in order to encourage the users to keep contributing to content.

SYSTEM ARCHITECTURE AND FLOW

Figure 3 illustrates the architecture and flow of the system integrating all the components described above to support informal social learning. First we describe the back-end process. It starts with the *crawler* gathering data it has not seen before from available data sources in the target domain and extracting documents (e.g. wikis, blogs, forum posts, files) as well as time-stamped user activities associated with each document. Document metadata and activities are stored in the *data repository*, while document text is stored in the *content repository*. The *topic modeler* retrieves data from the content repository and conducts text analytics to create the *topic model* if there isn't one already, or to fold new documents into the existing topic model otherwise. The *user/domain modeler* accesses the data repository and the topic model to build topic-specific *user models*. Individual topic-specific models of the users who are identified as experts for a topic are aggregated to create the topic-specific *domain model* of the topic. The *expertise scorer* computes topic-specific expertise scores for each user. The *activity recommender* creates a list of recommended learning activities for each <user, topic>. Each learning activity has the attributes of activity type, link to the content (for reading) or to the data source (for creation), and the *ROI* and *DUCE* values. Both expertise scores and recommended learning activities are recorded in the data repository. Periodically (e.g. overnight) the back-end process is re-executed to gather new content/activities, and update models, expertise scores, and recommendations.

On the front-end, each of the three main pages has its own *client hander* to provide the UI functionality for the page. Each client handler communicates with the centralized *server handler* using a set of REST APIs and JSON-formatted data to retrieve the data to be displayed and report user interaction to the back-end for processing.

USER STUDY

Methodology

We conducted a controlled laboratory study to collect data on the user experience with the Expediting Expertise (EE) system. The user study was designed to evaluate to which degree the system meets the original design goals, and to gather feedback on how the system could be improved.

We crawled data (documents and activities) from two internally public communities dedicated to data analytics products and services in an enterprise social platform. We chose these communities based on the high level of activities. The data set contains 2,533 documents (including wiki pages, blog posts, forum entries, and files), and 10,755 activities (including create, update, comment, and download) from 4,748 community members. We set the number of system topics (a configurable parameter) to 30 for the automatic topic modeling step. A domain expert mapped these system topics to 13 labeled topics by merging related topics together during the human validation step.

Due to the geographically dispersed nature of the community members, it turned out to be difficult to run a laboratory study with the actual members. Therefore, we chose data from three community members with various degrees of activities ("active", "moderate", "inactive") and mapped them to three hypothetical users of the EE system. The topics of interest for learning were pre-selected prior to the evaluation to keep the setup same for all participants.

There were a total of three exercise sessions (one per each hypothetical user) in the study. For each exercise, the study participants were asked to log in to the system as one of the hypothetical users and perform three tasks. Each task required the participants to explore one page of the system (i.e., "My Activities", "My Interests", and "My Recommendations") and answer 9 questions using the information available on that page (e.g., number of documents authored, topic with the highest expertise score, recommended learning activity with the highest *ROI*/*DUCE* value and its content type, popularity, or estimated completion time). A full list of these questions is omitted here due to space limitation. The error rate of the answers to

these questions provides objective measures of the participants' understanding of EE's main features. Each task also included 3 questions asking for subjective evaluation of the expertise scores and recommendation for the associated hypothetical user on a 5-point Likert scale (e.g. "the recommendation map gives the user an idea of his/her overall expertise strengths and gaps for the selected topics"). The order of the exercises (hypothetical users) was randomly selected for each participant.

For the main study, we recruited 10 participants from a large international corporation, 7 of who were male and 3 were female. Before performing the exercises, each user was given a tutorial on EE, during which s/he was instructed to log in to the system using a test account and visit all of the main pages. The users had time to freely explore all the features at the UI. For each exercise in the study, the participants read the descriptions of the tasks and the questions they needed to answer in print form. After the participants completed all three exercises, we asked them to fill out an exit survey of 10 questions that aimed to collect feedback on the overall system. The questions were structured to ask about degrees of agreement with a number of statements regarding the expertise scores, the recommendation map and the overall system, measured using a 5-point Likert scale (from "strongly disagree" to "strongly agree").

Results

In this section we report on objective and subjective data collected from the lab study.

The training time varied from 5 minutes to 18 minutes across the users, with an average of 11 minutes. In spite of this very brief training time, all participants were able to complete their exercises with few difficulties. The total exercise completion time varied from 6 minutes to 27 minutes, with an average of 14 minutes. There was no significant variance among the completion time of different exercises by a single participant. Almost all participants answered all task questions correctly. Only one participant made a mistake on answering one question in his last exercise, where he was asked to identify the recommended learning activity that provided the most learning gain (return on investment) for a topic. This same participant had correctly answered the same question for other topics in his two earlier exercises. Overall, the short exercise completion time and the low error rate indicate that the users could easily understand and interact with the system.

9 out of 10 participants indicated that they understood that the expertise scores were a relative measure calculated based on a comparison of the user's activities against the activities of others in the communities. All participants agreed/strongly agreed that the expertise scores would be helpful in keeping track of the user's current status, progress, and target in expertise building. However, their opinions varied regarding whether 'the expertise scores are an accurate representation of the hypothetical user's expertise as reflected in his/her activities listed on the "My Activities" page' (our actual question text). For both the "active" and "inactive" hypothetical users, 8 participants agreed/strongly agreed with this statement, while only 2 participants strongly agreed with this statement for the "moderate" hypothetical user, 7 remained neutral, 1 disagreed. The average rating was 3.90 out of 5. We believe that it was easier for the participants to judge the relation between the expertise score and the activity level for the "active" and "inactive" users than for the "moderate" user. Overall, we were encouraged to see that the majority of the responses about the expertise scores were positive, and felt that by asking participants to evaluate the hypothetical users' expertise scores given the available data, these responses were relatively objective and free of any bias that may be introduced (if we asked them to evaluate their own expertise scores) due to people's desire to be seen in the best possible light (e.g. social desirability bias, selective self-presentation effect).

All participants agreed that it would be useful to have the recommendation map personalized based on the user's background, and the meanings of the color and size of the boxes in the recommendation map were easy to understand. 9 out of 10 participants agreed that it was helpful to have *ROI* and *DUCE* information for each recommended learning activity, and it was easy to prioritize learning activities (i.e. determining which to focus on first) based on such information. All except one user (who remained neutral) agreed that they would prefer the recommendation map to the list style display of recommended learning activities. One user commented: *"List is dull."* Another user said *"I like the map view but lists are sometimes useful too. Both views should be available."* 9 out of 10 participants agreed that the recommended learning activities were relevant to the topics. Only one user found it difficult to judge relevancy mentioning that he was not familiar with the domain chosen for the study. We saw a slight drop in the participants' agreement with the statement of "the recommendation map would give the user an idea of his/her overall expertise strengths and gaps for the topics" when the hypothetical user's activity level was lower. While all agreed or strongly agreed with this statement in the "active" user's case, that number was 8 for the "moderate" user, and 7 for the "inactive" user.

For the overall assessment of the system, all participants agreed that they found the user interface intuitive and easy to use, and that they would use the tool for their own expertise building if available, and recommend it to others.

For all the study results, we performed one-sample t-tests comparing means to the midpoint of the Likert scales (3) and found them all statistically significant ($p<.005$).

DISCUSSION

Content

Because the primary goal of the Expediting Expertise system is to support informal social learning, we have largely focused on content organically created in social environments (i.e., content created by and shared with peers and colleagues). We use topic modeling to automatically group content into topics and assess content quality based on a number of factors such as relevancy to topic, popularity, and user ratings. This effort helps structuralizing content and making it more suitable for learning, but it is only the start. One of our user study participants pointed out that even though the recommendation map provided an intuitive roadmap for learning about a topic, the lack of information about how different pieces of content about this topic are related to each other could still make learners feel at loss, especially if they are new to the topic and want to start with the basics. This suggests a need for granular structures on content within topic boundaries. However, it is a challenge to obtain such structures automatically without extensive domain knowledge.

One way to address this issue is that, instead of focusing purely on informal learning, we could explore technologies to associate organically created content with curated content (including formal training materials) wherever available to synergistically blend formal and informal learning. Because information such as dependency and difficulty is more easily available for curated content, utilizing content and metadata created for formal learning could help to better access the quality and suitableness of particular informal learning content, which would result in more effective learning. An additional benefit of blending formal and informal learning is that it allows users to easily access different types of learning materials that are complementary to each other in one setting without having to go to different places.

One should note that although there has been work on integrating social learning into existing learning management systems, organically created content is still largely viewed as supplementary to formal learning content. As a result, it is not systematically analyzed and its use is limited. We would argue that organically created content should be put on an equal footing with formal learning content in terms of the attention it should receive, so users can benefit more from knowledge and experience of others, especially for fast growing areas where formal learning content is scarce.

Assessment

One common feedback from our users was that the usefulness of expertise scores and personalization of learning recommendation mostly depend on how accurate a user model is in representing one's expertise. Since the system infers a user's expertise based on what s/he does, it is important to build user models based on comprehensive sources of user activity information. A feature that allows users to provide information about their activities and associate content (e.g. a publication authored) through user interface would enable the system to build more accurate user models by incorporating information from sources that are otherwise unknown or ignored.

A concern frequently raised about expertise assessment was whether it could have a negative impact on employee performance evaluation. Users tended to be more critical of the accuracy of expertise scores and less willing to make their scores visible to others when they perceived that the scores would be associated with the evaluation of their job performance. Therefore, it is important to emphasize the learning support purpose of expertise assessment for the system to be successfully adopted.

Recommendation

Several users called for a more granular personalization of learning recommendation than what *ROI* (size) and *DUCE* (color) values could provide in the recommendation map. They commented that a more useful recommendation would require an understanding of the user's goal and preferences in learning. Their feedback suggests that learning recommendation should be created based not only on what the user knows and doesn't know, but also on what s/he needs/wants to know. To deal with the cold start or data sparsity problems, the recommendation could be based on activities and preferences from people who are "similar" to the target user (e.g. people in the same job role or from the user's personal social network) by adapting collaborative filtering techniques.

Other features suggested for learning recommendation include providing a more dynamic and adaptive Amazon style recommendation based on the specific learning activity the user is currently focusing on, and enabling users to bookmark/pin a particular recommendation for easy access at a later time.

CONCLUSION

In this paper, we present Expediting Expertise, a system that facilitates continuous expertise building in the enterprise by providing a social environment to support personalized informal learning. It employs a data-driven approach to automatically organize learning materials, concretely measure users' current expertise levels in their topics of interest and create personalized recommendations of learning activities. This approach requires minimal domain customization and automatically captures the evolution of domain knowledge. The system's web UI presents expertise assessment and learning recommendation using graphical visualizations to facilitate effective system-user interaction. Results from our user study demonstrate the usefulness of our main design features and the effectiveness of the overall system.

REFERENCES

1. Abel, F., Bittencourt, I., Costa, E., Henze, N., Krause, D. and Vassileva J. (2010). Recommendations in online discussion forums for e-learning systems. *IEEE Transactions on Learning Technologies*, 3(2):165–176.
2. Anderson, J.R., Corbett, A.T., Koedinger, K.R. and Pelletier, R. (1995). Cognitive tutors: Lessons learned. *The Journal of the Learning Sciences*, 4:167–207.
3. Antonopoulus, N., Salter, J. (2006). Cinema screen recommender agent: Combining collaborative and content-based filtering. *IEEE Intelligent Systems*, 35–41.
4. Ballantyne, N.J. and Quinn, K. (2006). Informal learning and the social web. *http://informallearning.pbworks.com/w/page/19914104/FrontPage*.
5. Balog, K., Azzopardi, L. and de Rijke M. (2006). Formal models for expert finding in enterprise corpora. In *SIGIR '06*.
6. Bandura, A. (1977). Self-efficacy: Toward a unifying theory of behavioral change. *Psychological Review*, 84(2):191–215.
7. Blei, D., Ng, A. and Jordan, M (2003). Latent dirichlet allocation. *Journal of Machine Learning Research*, 3:993–1022.
8. Buunk, B.P., Collins, R.L., Taylor, S.E., Van Yperen N.W. and Dakof, G.A. (1990). The affective consequences of social comparison: Either direction has its ups and downs. *Journal of Personality and Social Psychology*, 59(6):1238–1249.
9. Candillier, L., Meyer, F., Boullé, M. (2007). Comparing state-of-the-art collaborative filtering systems. *Lecture Notes in Computer Sciece*, 4571: 548–562.
10. Carmagnola, F., Vernero, F., Grillo, P. (2009). SoNARS: A social networks-based algorithm for social recommender systems. In *Proceedings of the 17th International Conference on User Modeling, Adaptation, and Personalization*, 223–234.
11. Chatti, M.A. and Jarke, M. (2007). The future of e-learning: A shift to knowledge networking and social software. *International Journal of Knowledge and Learning*, 3(4/5):404–420.
12. Dalsgarrd, C. (2006). Social software: e-learning beyond learning management systems. *European Journal of Open, Distance, and E-Learning.*
13. Despotakis, D., Dimitrova, V., Lau, L., Thakker, D., Ascolese, A. And Pannese, L. (2013). ViewS in user generated content for enriching learning environments: A semantic sensing approach. *International Journal of Artificial Intelligence in Education*, (7926):121–130.
14. Educational Data Mining. *http://www.educationaldatamining.org.*
15. Enterprise Strategy Group. (2012). Research report: Social enterprise adoption trends. *http://www.esg-global.com/research-reports/social-enterprise-adoption-trends/*.
16. Festinger, L. (1954). A theory of social comparison processes. *Human Relations*, 7(2):117–140.
17. Freedman, R. (2000). What is an intelligent tutoring system? *Intelligence*, 11(3):15–16.
18. Gary, B. (2004). Informal learning in an online community of practice. *Journal of Distance Education*, 19(1):20–35.
19. Greenhow, C., Robelia, B., Hughes, J. (2009). Learning, teaching, and scholarship in a digital age web 2.0 and classroom research: What path should we take now? *Educ Res* 38(4):246–259.
20. Herlocker, J., Konstan, J., Borchers, A., Riedl, J. (1999). An algorithmic framework for performing collaborative filtering. In *SIGIR '99*, 230–237.
21. Lang, K. (1995). NewsWeeder: learning to filter netnews. In *Proceedings of the 12th International Conference on Machine Learning*, 331–339.
22. Mason, R. and Rennie, F. (2007). Using Web 2.0 for learning in the community. *The Internet and higher education*, 10(3):196–203.
23. Minocha, S. (2009). Role of social software tools in education: A literature review. *Education and Training*, 51(5/6):353–369.
24. Murray, T. (1999). Authoring intelligent tutoring systems: An analysis of the state of the art. *International Journal of Artificial Intelligence in Education (IJAIED)*, 10:98–129.
25. Nkambou, R., Mizoguchi, R. and Bourdeau, J. (2010). Advances in intelligent tutoring systems. Heidelberg: Springer.
26. Nwana, H.S. (1990). Intelligent tutoring systems: An overview. *Artificial Intelligence Review*, 4:251–277.
27. Petkova, D. and Croft W.B. (2006). Hierarchical language models for expert finding in enterprise corpora. In *ICTAI '06*.
28. Polson, M.C. and Richardson, J.J. (1988). Foundations of Intelligent Tutoring Systems. Lawrence Erlbaum.
29. Ravenscroft, A. (2009). Social software, web 2.0 and learning: status and implications of an evolving paradigm. *Journal of Computer Assisted Learning*, 25(1):1–5.
30. Seid, D.Y. and Kobsa, A. (2003). Expert finding systems for organizations: Problem and domain analysis and the DEMOIR approach. *Journal of Organizational Computing and Electronic Commerce*, 13(1):1–24.
31. Shank, P. (2012). Smart companies support informal learning. The eLearning Guild.
32. Vassileva, J. (2008). Toward social learning environments. *IEEE Transactions on Learning Technologies*, 1(4):199–214.
33. Vassileva, J. (2012). Motivating participation in social computing applications: a user modeling perspective. *User Modeling and User-Adapted Interaction*, 22(1–2):177–201.