HIDDEN PATTERNS IN LEARNER FEEDBACK
Generalizing from Noisy Self-assessment during Self-directed Learning

Thomas Markus and Eline Westerhout
UiL-OTS Institute, Utrecht University, Trans 10, 3512 JK Utrecht, The Netherlands

Keywords: Topic modeling, Ontology, Informal learning, Computer-aided assessment.

Abstract: We propose a method which uses high-level learner feedback to recommend learning materials that match the knowledge level of a specific learner. Machine learning and topic inference techniques will be applied to documents that were rated by the learner to infer information on the learner's conceptual development. The inferred topics will be linked to a domain ontology, allowing us to offer the learner knowledge-rich feedback regarding his level of understanding. In addition, appropriate learning materials can be recommended on the basis of the learner's computational model. The proposed method is especially useful in lifelong learning contexts, in which tutor support is often not available.

1 INTRODUCTION

In a Lifelong Learning context, learners access and process information in an autonomous way. They often rely on informal learning materials, that is, on (non-)textual content available through the web which is uploaded and accepted by a community of learners and not necessarily by an institution. These learners, however, do not have the support of tutors or teachers when trying to comprehend these learning materials. The educational practice of building learning support systems is shifting from pedagogically orientated approaches that focus on acquiring a fixed curriculum to just-in-time (JIT) and life-long learning (LLL) approaches (Collis and Moonen, 2002). JIT and LLL both rely on a large body of accessible learning materials that target a specific area of interest or skill. Both can be accessed using social networks and social bookmarking services (Marlow et al., 2006) or regular search engines.

Social networks and collaborative bookmarking systems are a natural fit for undirected informal learning since they allow an almost unprecedented amount of personalization. Current solutions aim to suggest relevant documents tailored to a specific task or a person's interests. However, from a learning perspective, the personalization should also take a learner's background knowledge and learning goals into account (Ley et al., 2010). The goal would be to provide learning objects that extend and build on familiar knowledge and while doing so to continuously improve the level of understanding of the subjects of interest. Taking this into account would allow the learner to be presented with learning objects that support his or her development on established subjects. This doesn't pre-suppose that learners actively search for such resources nor requires an appropriate level of lexical competence for composing effective search queries. Providing such resources to a learner pre-supposes a pedagogical model of the learner that captures not only his interests, but also his level of understanding of different subjects. Such a pedagogical model is further complicated by the LLL-environment, where the presence of a dedicated tutor cannot be assumed. This necessitates a great level of automation for such a model to be applicable. The EU FP7 "Language Technologies for Lifelong Learning"-project (LTfLL)[1] has developed pedagogic approaches and software which leverage NLP-tools and techniques, ontologies and social media for tutor support and self directed lifelong learning.

This paper will describe a methodology that builds on the LTfLL models and tools and embeds itself in current web practices. The methodology results in a learner model based on self-directed learning that can support lifelong learners by providing appropriate feedback. We will employ knowledge rich resources such as domain ontologies to visualize this model in order to make it understandable and to reinforce and acquire domain concepts and their relations to one another. Section 2 will shortly summarize its theoreti-

[1] http://www.ltfll-project.org/

Markus T. and Westerhout E..
HIDDEN PATTERNS IN LEARNER FEEDBACK - Generalizing from Noisy Self-assessment during Self-directed Learning.
DOI: 10.5220/0003343302850290
In *Proceedings of the 3rd International Conference on Computer Supported Education* (CSEDU-2011), pages 285-290
ISBN: 978-989-8425-49-2

cal assumptions, followed by section 3 which will address the advantages and challenges in using subjective ratings provided by individual learners. Section 4 will describe the process of determining the conceptual contents of documents and how these are to be linked to the subjective ratings. Section 5 will provide details of knowledge rich representations for knowledge feedback in order to make the acquired learner model understandable and accessible.

2 THEORETICAL BACKGROUND

The social learning support system is based on the theoretical framework proposed by Stahl (Stahl, 2006), who views the knowledge building process as a mutual construction of the individual and the social knowledge building, striking a balance between the Acquisition (individual) and the Participation (social) Metaphors. In this model knowledge is a socially mediated product. Individuals develop personal representations and beliefs from their own perspectives, socio-cultural knowledge building, shared language and external representations. These are further extended and corrected through social interaction, communication, discussion, clarification and negotiation. Learners build knowledge collaboratively and then internalize it in a personal knowledge building process. Learners can then decide to try and become skilled members of a Community of Practice (Lave and Wenger, 1991), mastering a domain speech genre (Bakhtin et al., 1986).

The process of mastering a domain speech genre is expressed through the consumption and generation of certain language artefacts. Large parts of the social mediation of knowledge currently takes place in social networks which are used to discuss and share learning resources. An increased amount of lexical competence in a domain is evidence for improved understanding and integration in its corresponding Community of Practice. Our proposed exploitation of the hidden patterns in learner self-assessments goes beyond the recommendation of topics of interest, because it models a learner's current level of understanding and can therefore provide added value.

3 SELF-ASSESSMENTS

Learners frequently perform self-assessments on potential learning material in order to estimate and adjust their subjective level of understanding of a topic. On the fly self-assessments drive a lot of exploratory search requests which target comprehensive learning materials in contrast to short-lived fact-finding queries. Naturally, the learning material needs to be relevant and should contain a decent amount of well presented information. There is however, an important self-assessment phase during this selection process that should not be overlooked. Learning resources that, although relevant, are above or below a learner's current level of understanding will be discarded as potential learning materials. Consider for example a search for the mathematical procedure of 'integration' by a 15 year old which is interested in next-week's subject. Although the Wikipedia article is somewhat helpful, a step-by-step tutorial is much more suited to that specific learner's level of understanding.

A search using a social bookmarking service where the learner has established a suitable social network structure or provided tagged resources will likely result in suitable resources. The learner would then use either the number of users that bookmarked the resource or the average rating of the resource to decide on whether to explore the resource or not. Finally the remaining resources are inspected for suitability by the learner through a fast self-assessment of the material. The learner ideally selects resources that provide additional information that is neither too difficult nor too trivial, but this process is slow and error-prone for new topics of interest. Neither the rating nor the popularity of a resource are reliable indicators for the utility of the resource. The learner is thus forced to manually decide whether a description or tag attributed to a resource is indicative of the appropriateness of the resource.

The success of this process however depends on the assumption that learners can perform adequate self-assessment of their current level of knowledge about a topic. (Baker, 1989) argues that learners are rather bad at assessing their comprehension of both texts within and outside of their domain of expertise. Surprisingly, domain experts were shown to overestimate their text comprehension on texts from their own domain when compared to novices, whose self-assessments were actually closer to their true level of comprehension. Assuming that learners' self-assessments are quite noisy and in some cases over-estimates, does this mean that these are useless? Most of the studies in (Baker, 1989) were conducted some time ago and most were based on a small set of texts where comprehension was measured to some previously determined gold standard. Good performance on self-assessments mostly correlated with better overall reading skills. It thus seems likely that improved meta-cognitive skills (evaluation of your own cognitive performance) lead to better manual selection of appropriate learning resources. Although (Baker, 1989) argues that the self-assessments are skewed, the assessments still capture an overall trend.

The search queries for retrieving learning objects are primarily constructed by the learners themselves in search of new information. This means in practice that learners will frequently default to formulating search queries that yield simplified tutorial-style resources. This is understandable considering the information overload and the temptation of ‘sticking to what you know’, but it does create challenges for self-directed learning approaches. Ideally the quality and difficulty of resources will improve as a learner’s level of understanding increases, short-lived fact-finding queries are both effective and easy which may keep the learner contained within a community of beginners instead of slowly migrating towards a community of experts. The amount of effort required by learners to construct search queries for high quality resources which support self-directed learning may prove to be too cumbersome to maintain in the long term.

We would therefore like to automatically steer learners towards resources that are both relevant and slightly challenging such that they go beyond fact-finding and move towards increased understanding of the domain. This approach however requires an accurate and up-to-date model of the subjects of interests of a learner and an estimate of the current level of understanding of each subject. It is likely that the learner will be unable to provide much detail on the conceptual decomposition of the difficulties he or she encountered when trying to understand certain learning objects. Moreover, requesting too much additional information from a learner is likely to disrupt the existing workflow which in turn creates additional boundaries for adoption of this approach. Luckily present day interaction using social networks and search engines allows us to acquire a huge number of simple learner self-assessments. Each individual self-assessment by itself may be skewed or wrong, but generalizing from a larger collection will yield stable trends. Naturally these trends will change over time as the learner progresses which means that older self-assessments should be properly discounted.

The aggregation of self-assessments needs integrate well within a learner’s existing workflow and should be simple and easy to use. A suitable candidate would be the 5-star rating process that is already familiar to learners on the Internet which can be repurposed to capture a simplistic summary of a learner’s assessment of a learning object. The advantage of using this type of simple and unspecific feedback is that it takes very little effort on the learner’s side, which increases the chances of the learner actually providing enough feedback. The feedback could for example be a simple likert scale which ranges from: 1 (too easy), 3 (just right), to 5 (too difficult).

The approach assumes that a learner is able to judge whether a specific learning resource is too complicated, but is unable to explain why. Only a maximally simplistic self-assessment is required from the learner that can be provided with a single mouse-click for each resource. Taking such a minimalist approach with respect to the feedback provided by learners minimizes the amount of additional effort required from learners which increases the likelihood of learners providing a large number of such resource feedbacks.

A computer-based machine learning approach allows us to analyze large amounts of data from each learner without much effort. Machine learning can be employed to automatically find complex patterns in that data collection. Machine learning allows us to build a model that links topics of interest to subjective levels of understanding. The system can then use this model to predict the most likely self-assessment for a new resource for a new particular learner. This model, which can be automatically learned from the self-assessments, can provide feedback which supports learners in their search for appropriate learning materials or can be used to recommend new resources. The approach is largely data-driven and only relies on the assumption that there is some level of consistency in the learner provided feedback.

The rating of a resource as provided by the learner says something about the two things that the document is composed of: (1) The way the information in the document is presented and structured (length of sentences, clarity of the language, ...) and (2) The information in the document itself; a number of topics. At present we are not addressing (1) which, although important, is about readability measures (Crossley et al., 2007). Incorporating a readability measure will allow the system to differentiate between text readability and conceptual complexity.

4 DECOMPOSING LEARNING OBJECTS

The learners provide feedback at the document level, and not separately for each of the individual subjects covered in a particular document. In order to determine a learner’s current level of understanding, it is necessary to identify which subjects (topics) are present in each document and what their relative proportion is. Latent Drichilet Allocation (Blei et al., 2003) (LDA) can be used to infer the distribution of topics for any particular document. In LDA, a topic is a set of words where the presence of those words in each other's context is evidence for the presence of the topic in question. An LDA-based topic inferencer is first trained on a generic document corpus that spans multiple subjects in order to determine the most likely topic composition of the corpus and the words that each topic consists of. Such a corpus could for example be an encyclopedia like Wikipedia which covers a wide range of subjects. Increasing the total number of topics will make each topic more specific, but the data that it is based on decreases. Proper sampling and inference can generate a probabilistic distribution of the topics present in any document. It is important to note that topics themselves have no name, but the most prominent words of a topic usually give a good impression of the semantically related subject(s) that the topic covers.

Reducing a document to a set of topics with their proportions will allow us to identify the subjects that the document covers. This information can then be used to identify the relation between the feedback provided by the learner at the level of the document and the individual topics that make up the document. The overall process is depicted in figure 1 which also gives a succinct example of the presence of three topics in a hypothetical document.

When a learner gives feedback about a resource on the Internet, topic inference can link the overall document content (topics) to the learner's rating of the document. A machine learning approach using neural network-based classifiers is used to learn the relation between topic distributions and ratings. This will result in a neural network classifier for each individual learner. Each document for which the learner has provided feedback is decomposed using LDA into a set of topic probabilities. These topic probabilities are then used to train the neural network with each corresponding rating as the desired output value.

Each of the neural-network classifiers realizes a learner model that is able to predict the most likely rating that the learner would give to new resources. This information can then be used to re-order results from other search engines tailored to a learner's model. For example resources which are likely to be classified as "just right" could replace earlier search results predicted to be classified as "too easy". The model can also provide the overall patterns in the ratings provided by the learner. Such generalizations are of the form: "There is a statistically significant chance that when Mary encounters a document that is about topics A, F and D, she will judge it as too difficult. However, documents only about topic A will be judged as easy". The classifier not only learns the examples by heart, but also builds a model of the underlying generalizations.

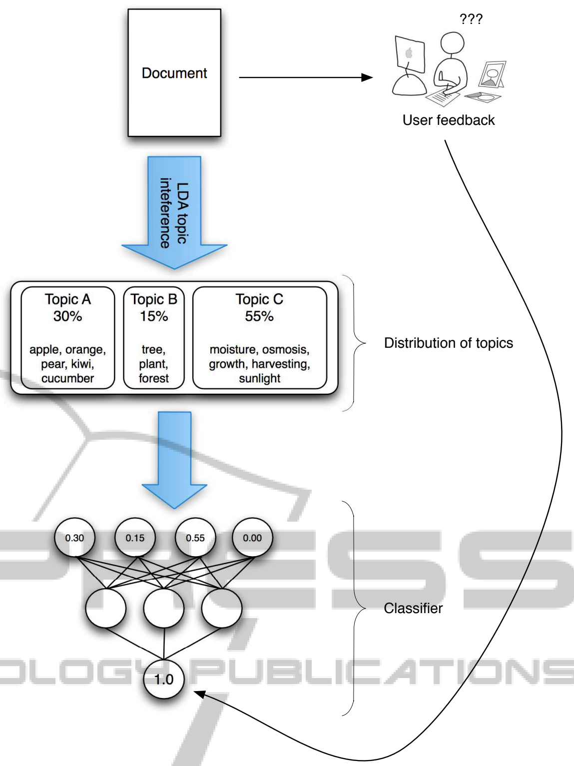

Figure 1: Overview of the topic inference process and the training of a classifier that represents an individual learner model.

These can be visualised by, for example, showing the most important topic terms that have been correlated with a particular rating.

5 KNOWLEDGE RICH FEEDBACK

A purely term based approach still provides challenges for learners when trying to internalize the feedback. Consider for example that we provide the learner with a list of terms that are representative of the topics that he consistently classified as being too difficult. Since the terms and their interrelations may be unknown or not apparent, because they were indirectly classified as difficult, the learner may experience great difficulty in understanding such raw results.

The language artefacts which the classifier generates are still relevant and useful, but it should not be assumed without question that the learner is able to gain this from the available textual feedback. Ad-

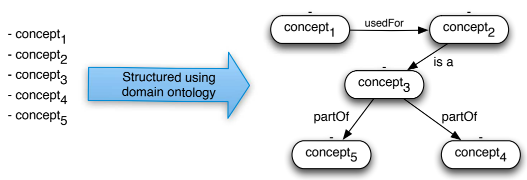

Figure 2: Structuring terms using a domain ontology.

ditional effort is needed to impose structure which should make it easier to internalize and relate the topics to the learners current conceptualization of the domain. It is because of this reason that we employ knowledge rich models of a domain to structure and enhance the results. This allows us to provide well structured feedback that makes the structure between terms and concepts explicit.

Domain ontologies will be used to structure the domain terms and to place them in an expert-approved relational structure (Gruber, 1993). Domain ontologies serve as approved reference conceptualizations of domains. The personal knowledge building process is supported by the clear and explicit structure of a domain ontology which improves the internalization. The salient terms extracted from the relevant topics extracted from the learner model can be linked to concepts from a domain ontology using a word sense disambiguation (WSD) algorithm. Such an algorithm can determine the appropriate meaning (word sense) for terms that are ambiguous or have only a single interpretation. Each meaning that the WSD-algorithm yields is represented by a concept from a domain ontology.

A graph-based visualization of a set of concepts can be generated given a domain ontology which serves as a user friendly method to access the ontology's conceptual structure (Westerhout et al., 2010). Such a domain ontology not only provides the concepts themselves, but also shows which relations they have to other concepts in the domain. This visualization can be enhanced to also show concepts already acquired by the learner and in the way in which they are connected to as of yet unacquired concepts.

Figure 2 illustrates the difference between an ontology based representation and a term-based representation of feedback. The added value of the rich relational structure of the ontology reduces the effort required from the learner to interpret and internalize the representation.

We can thus convert a list of terms, as provided by a learner model, to a list of concepts from an ontology. This list of concepts can then be used to generate an ontology fragment tailored to a particular learner. The relational, expert approved, structure of the domain ontology supports the learner in interpreting and exploring the trained model. It provides a frame of reference starting from known concepts to new unknown concepts which allows learners to start from already acquired domain concepts and to explore new subjects and relations between subjects which allow learners to gradually expand their knowledge. Learners are motivated to explore new subjects, because the domain ontology shows how these subjects relate to what they are already familiar with.

6 CONCLUSIONS

We propose a learner support system that employs the knowledge available in social networks to recommend relevant learning materials tailored to the conceptual level of the learner. The system aggregates a large number of learner provided non-textual feedbacks instead of using learner provided text in order to minimize the disruption of the normal workflow. The resources for which the learner provides feedback are decomposed in topics. This allows us to identify the differences between the topics already understood by the learner and those that are not. To this end, a personalized model of each learner is created from the data which is used to predict the level of conceptual competence for new resources and to provide an overview of unacquired concepts through the use of domain ontologies. The conceptual structure provided by the ontologies facilitates the acquisition and reinforcement of domain concepts.

REFERENCES

Baker, L. (1989). Metacognition, comprehension monitoring, and the adult reader. *Educational Psychology Review*, 1(1):3–38.

Bakhtin, M., Holquist, M., and Emerson, C. (1986). *Speech genres and other late essays*. Univ of Texas Pr.

Blei, D., Ng, A., and Jordan, M. (2003). Latent dirichlet allocation. *The Journal of Machine Learning Research*, 3:993–1022.

Collis, B. and Moonen, J. (2002). Flexible learning in a digital world. *Open Learning: The Journal of Open and Distance Learning*, 17(3):217–230.

Crossley, S., Dufty, D., McCarthy, P., and McNamara, D. (2007). Toward a new readability: A mixed model approach. In *Proceedings of the 29th Annual Conference of the Cognitive Science Society*, pages 197–202.

Gruber, T. (1993). What is an Ontology? *Knowledge Acquisition*, 5(2):199–220.

Lave, J. and Wenger, E. (1991). *Situated learning: Legitimate peripheral participation*. Cambridge university press.

Ley, T., Kump, B., and Gerdenitsch, C. (2010). Scaffolding Self-directed Learning with Personalized Learning Goal Recommendations. *User Modeling, Adaptation, and Personalization*, pages 75–86.

Marlow, C., Naaman, M., Boyd, D., and Davis, M. (2006). Position paper, tagging, taxonomy, flickr, article, toread. In *In Collaborative Web Tagging Workshop at WWW06*. Citeseer.

Stahl, G. (2006). *Group cognition: Computer support for building collaborative knowledge*. MIT Press.

Westerhout, E., Monachesi, P., Markus, T., and Posea, V. (2010). Enhancing the Learning Process: Qualitative Validation of an Informal Learning Support System Consisting of a Knowledge Discovery and a Social Learning Component. *Sustaining TEL: From Innovation to Learning and Practice*, pages 374–389.