Placement Support for Learners in Learning Networks

(Plaatsingsondersteuning voor lerenden in leernetwerken)

"Technologies are never just tools, they are evocative objects. They cause us to see ourselves, and our world, differently"
Sherry Turkle, 2002

Placement Support for Learners in Learning Networks

Proefschrift

ter verkrijging van de graad van doctor
aan de Open Universiteit Nederland
op gezag van de rector magnificus
prof. dr. ir. F. Mulder
ten overstaan van een door het
College voor promoties ingestelde commissie
in het openbaar te verdedigen

op vrijdag 16 oktober 2009 te Heerlen
om 16.00 uur precies

door

Marco Kalz

geboren op 7 december 1975 te Aken, Duitsland

Promotor:
Prof. dr. E.J.R. Koper
Open Universiteit Nederland

Co-promotor:
Dr. J.M. van Bruggen
Open Universiteit Nederland

Overige leden van de beoordelingscommissie:
Prof. dr. P. Baumgartner
Danube University Krems, Austria

Prof. dr. Ph. Dessus
University of Grenoble, France

Dr. ir. F. J. M. Mofers
Open Universiteit Nederland

Prof. dr. E. O. Postma
University of Tilburg, The Netherlands

Prof. dr. P. B. Sloep
Open Universiteit Nederland

The research reported in this thesis was carried out at the Open University of the Netherlands in the

Centre for Learning Sciences and Technologies
celstec.org

in the context of the research school

SIKS

(Dutch Research School for Information and Knowledge Systems)
and as a part of the TENCompetence project.

TEN Competence
Building The European Network for Lifelong Competence Development

The TENCompetence Integrated Project is funded by the European Commission's 6th Framework Programme, priority IST/Technology Enhanced Learning. Contract 027087 (www.tencompetence.org).

SIKS Dissertation Series No. 2009-36
The research reported in this thesis has been carried out under the auspices of SIKS, the Dutch Research School for Information and Knowledge Systems.

Cover design: Stephan Kochs, EinGrafikbüro, Aachen & Jeroen Berkhout
Printed by: Datawyse/Universitaire Pers Maastricht
ISBN 978-90-79447-30-5

Placement Support for Learners in Learning Networks

Marco Kalz

Synopsis

Accreditation of prior learning (APL, in the Netherlands referred to as EVC) enables institutions to offer personalized learning arrangements which take into account the prior learning of individuals. Unfortunately, this placement process of learners is time-and cost-intensive and cannot be scaled up easily to meet higher demands. Much work needs to be undertaken even before the potential learner makes a commitment to enter a study. Not only the assessment of an APL portfolio is time-consuming; literature shows that learners are unsure what to put in their portfolios for the APL procedure. Technological support may help reduce time and costs involved in this process and to improve the quality of the APL process. In this thesis we explore the application of advanced text mining and statistical natural language processing to offer a solution for these procedures which can be used in traditional APL procedures and informal learning networks. Based on the identification of three different situations which depend on the availability of data from learners we focus on the most complicated case where no (meta)data about the learners and the target learning content exists.

A prototype model of a placement web-service for lifelong learning is proposed that is able to function as a supporting service by pre-analyzing documents in the APL procedure and by assisting in deciding whether these documents are relevant or irrelevant for the target course or study programme. This service employs in its core a reduced vector space model similar to Latent Semantic Analysis (LSA). While LSA uses very large corpora, in this thesis we evaluate the use of dimensionality reduction methods with smaller domain specific corpora. For this purpose we empirically evaluate the use of dimensionality reduction on the basis of two exemplary small corpora. We demonstrate that a combination of filtering strategies, the use of multiple criteria and dimension reduction that takes into account the variance accounted can help maximize performance. Based

on these finding we have conducted a study with real learner data. Data is collected in a psychology course of the Open University of the Netherlands. The results of this study show that the application of dimensionality reduction techniques for APL procedures is a promising supporting method to decide about relevancy of learner portfolios. Based on these results we introduce two technological artifacts that have been developed during the project and which allow us to evaluate the approach in several different environments and to extend our approach on the different placement support situations introduced at the beginning of the thesis.

Table of Contents

Index of Figures

Index of Tables

Section 1: Introduction[1]

[1] Partially based on Van Bruggen, J., Kalz, M., & Joosten-Ten Brinke, D. (2009). Placement Services for Learning Networks. In Koper, R. (Ed). Learning Network Services for Professional Development. Berlin: Springer.

Introduction

Mobility of learners is a concept that is discussed on several national and international levels. Especially on the level of the European Union there are efforts to make competences and performance of individuals comparable across borders. One important example of such unification is the European Qualification framework (EQF). The EQF serves as a reference framework allowing national bodies to align their qualifications systems and certificates to an international standard to reach an open area for work and learning across borders in Europe.

This target is not limited to cross-national transferability but it exists at the same time within a national qualification system if people want to change their careers or study programme. Even within an educational institution this problem can occur. The accreditation or recognition of prior learning (APL) is a procedure to allow exemptions for students and to offer an individualized curriculum. Within this process domain experts study the portfolio of potential students and provide exemption decisions based on the prior knowledge they have identified in the portfolio of the learner. Toyne (1979) describes APL as an "essential process whereby qualifications, part-qualifications and learning experiences are given appropriate recognition (or credit), to enable students to progress in their studies without unnecessarily having to repeat material or levels of study" .

But this procedure is expensive and time-consuming for the domain experts. Minton (2008) mentions two main problems of the current APL procedures. On the one hand the current APL system cannot be easily scaled up to meet the increased demand on various levels, on the other hand it is problematic that much work needs to be undertaken before the potential student makes a commitment to study and pays any fees. To decrease the costs involved in the APL procedure technological support is needed. In the UK this need has been recognized and the e-APEL project proposes two supporting services for the APL procedure. An "Estimator" service should guide students through the APL process in which they have to decide which documents they can submit to show a proof or prior learning. This service should approximate the (knowledge) level and scope of the documents submitted by the students. In addition an "Advisor" service should support the domain experts to make decisions about exemptions (Haldane, Meijer, Newman, & Wallace, 2007).

Besides the importance of technological solutions to support the APL process in traditional educational offerings there is at the same time a requirement to implement this process in technology-enhanced learning for personalization of learning offerings. Nordeng, Lavik, & Meloy (2004) reformulate this problem in the following way:

> "How can the students themselves be able to assess their position relative to a future learning environments consisting of a diverse set of learning activities from which learners somehow may take their pick? The learner's history and goals define an entry position relative to the learning activities. A different entry position is likely to result in a different partition of the set of available activities in activities to skip and to complete".

In traditional e-learning contexts much knowledge and existing (meta)data about learners and (course) content are available. This allows the application of learner modeling approaches to approximate prior knowledge of learners (Brusilovsky, Kobsa, & Nejdl, 2007). We refer to these approaches as "top-down" since the availability of this information is assumed and personalized learning can be offered solely on this information. However in the context of informal learning scarce or no knowledge about the learner and the learning content is available. Therefore bottom-up techniques are needed to approximate the prior knowledge of learners and their position in relation to their current competence development goal. Bottom-up means in this context, that no static knowledge about learners and learning content and learning activities is available since these can constantly change. In addition there is no agreed-upon vocabulary available to describe competences and learning activities in informal learning.

Learning networks are a promising approach to realize the idea of flexible (in time, place and space), personalized (optimal suited to the learner) learning environments for lifelong learning. A learning network is an ensemble of actors, institutions and learning resources which are mutually connected through and supported by information and communication technologies in such a way that the network self-organizes (Koper, Rusman, & Sloep, 2005). All actors in the network strive for the development or improvement of their competences. A learning network should serve as an environment that can bridge the worlds of informal and formal learning based on the approach of competence-based learning (Sloep, 2008).

Placement in Learning Networks, or 'positioning' as it was called elsewhere (Van Bruggen et al., 2004), is a process by which we try to put a learner on the most appropriate place in one of the alternative paths leading to a specific competence taking into account his goals and preferences, as well as the history. In informal learning networks two independent but connected supporting services are needed. A *placement support service* seeks to identify an efficient path towards a goal by identifying those learning activities in one or more paths that the learner might skip, taking into account the history. In addition, based on the result of the placement support service a *navigation service* assists the student in making a decision on which learning activity to engage in next (Drachsler, Hummel, & Koper, 2008). These two services can make use of a on a *learning path* that is able to describe a sequence of learning activities in informal learning (Janssen, Berlanga, Vogten, & Koper, 2008).

Due to the open nature of the environment offered by Learning Networks and in particular its implication that all members of a learning network may freely contribute learning activities or even complete curricula to the network the placement process is complicated. A diverse set of learning activities might lead to the same competence development goal. In addition, there may be offerings ranging from accredited curricula provided by educational institutions that lead to a formal, full degree, to collections of materials of non-formal nature. A road towards a specified competence goal may be predefined by an institution in form of a curriculum or it might be more loosely defined e.g. by an individual learner. It is likely that such a learning network will expand to offer various roads leading to a competence outcome their contributors claim to be the same or much alike.

Next to complete learning paths towards a predefined competence goal, what we called a curriculum, there may be a vast amounts of resources, not bound together in a curriculum, that are contributed by individual actors. Today's Internet already offers a plethora of lessons, examples, transcriptions et cetera that are made available by large numbers of amateurs and professionals. In a learning network there will be overlaps in the learning activities contained in these predefined roads: several universities will offer master programmes in Psychology; different actors, including institutes for formal music education, will offer roads leading to advanced guitar playing etc. And, as already witnessed in the World Wide Web today, there will be vast amounts of informal learning resources that are offered, on an as-is basis by various individual or collective actors.

All this makes placement a process in which extended, but not necessarily complicated, search through alternatives is necessary. The real complicating factor is the absence of a common vocabulary. A learning network does not enforce its users to adopt a common controlled vocabulary to describe the learning activities and its learning outcomes. Note that the latter also has implications for the portfolio of members of a learning network. Although, as emphasized above, no controlled vocabulary to describe competences and learning activities is enforced in Learning Networks, this does not preclude that areas with a common vocabulary exist. Metadata on curricula offered by European institutes may conform to the Dublin descriptors or to the European Qualification Framework. Even in these cases it can be challenging to recognize that metadata are compliant with a controlled vocabulary and to identify that vocabulary. It seems likely that compliant parts of the learning network are associated to providers offering formal education. The real challenge for placement then is to use *whatever data is available* to locate material and to map that onto the learner data available.

This thesis explores an alternative approach to approximate prior knowledge of learners via the texts they have written in their prior education and work context. This approach requires a method to map the competence level represented in the texts written by the learners to a target domain in which the learner wants to improve his/her competences. For this purpose we have employed in this thesis a text vector space model in which the text documents of the learner and the domain are represented as vectors (Salton, Wong, & Yang, 1975). The approach we have chosen is very similar to the application of Latent Semantic Analysis (LSA) (Deerwester, Dumais, Furnas, Landauer, & Harshman, 1990) which has been applied to educational problems like automated essay scoring or the selection of educational material. In the thesis we use the term "LSA" loosely because we apply the same steps like in a Latent Semantic Analysis. The only difference is that we do not use a large general language corpus since it was not available in Dutch at the time of our studies.

The problems we are addressing in this thesis are the following:

1. The assessment of prior knowledge of learners is a very time-consuming process. To improve the quality of this process and to reduce the costs technological support is needed. For this technological support a model of a placement support service is needed that is able to support the APL procedure and the

approximation of prior knowledge. The development of such a model is complicated by the fact that a method to approximate prior knowledge of learners based on their writing must be generic enough to be implemented in different contexts but at the same time flexible enough to be changed and aligned to different local standards and thresholds. To evaluate such a model an evaluation strategy is needed that is based on the aspects of *scalability*, *sensitivity*, *reliability* and *validity* and the *fit into APL procedures*.

2. While the Latent Semantic Analysis method is based on the availability of a large general language corpus to allow the algorithm to "learn" the underlying structure of the application language in our studies we could not follow such an approach. On the one hand there was no large corpus available in Dutch at the time of our studies; on the other hand we cannot assume that in learning networks such a training corpus is available at all. The same problem applies for multilingual communities where no lingua franca is available (Vuorikari & Ochoa, 2009). For this reasons we have decided to focus in our studies on the application of the text vector space model and dimensionality reduction techniques to smaller domain specific corpora.

3. For this application the question needs to be addressed how a placement support service operating under the restrictions above can reach its best performance. One important aspect to consider here is the *sensitivity* of such a service. The service must demonstrate that it can differentiate between similar and dissimilar documents. This aspect is complicated by the fact that the target learning units in learning networks might share a high amount of domain specific concepts.

4. The support of a placement support service will have the form of a *classification* since it should help to decide about relevant and irrelevant documents for a specific target domain. This classification model should be *valid* in regard to the model that human experts apply. Besides the validity such a placement support service should demonstrate ideally the same reliability as human experts.

5. The placement support service should be *flexible* and *scalable*. It should be flexible enough to be implemented and trained in different contexts and it should be scalable in regard to the future use of larger corpora.

6. For the implementation of our model of a placement support service a *prototype* is needed that can implement the model proposed in this thesis. In addition for future research a technological basis is needed to conduct experiments which go beyond the approach evaluated in regard of corpus construction, data used and algorithms applied.

In section 2 of this thesis we present the theoretical foundations for our research. These foundations have a technological and educational part. In chapter 2.1 we define a long-term research agenda for placement support. This long-term research agenda focuses on the availability of the data in learning networks. In this chapter we propose a long-term data-driven research agenda and identify three different situations for placement support. For these situations we describe the state-of-the art and discuss a strategy for this long-term research agenda. In chapter 2.2 we focus on the educational foundations of our research project and discuss methods and approaches from research on eAssessment, Accreditation of Prior Learning (APL) and (electronic) Portfolios. Based on a discussion of these perspectives on the placement problem we focus on the content-based placement situation and describe Latent Semantic Analysis (LSA) as a promising approach to support this situation. We develop and discuss a model of a placement support service.

In section 3 we present our empirical work. In chapter 3.1 we present a study dealing with the *sensitivity* aspect of our approach. For this purpose we have developed a method that optimizes performance in regards to the *discriminatory power* and the *reduction of noise* in the data. The discriminatory power should prove that our approach is able to recognize similar and dissimilar documents while the noise reduction aspect is realized with the identification of an ideal number of dimensions to retain. To optimize on these criteria we had to develop filtering strategies that filter away high-frequency general language terms while keeping important domain-specific terms. In small scale corpora we cannot make use of weighting functions like inverse document frequency (IDF) or entropy because these would also filter away important domain specific terms that contribute to the discrimination between the target learning units. For the dimensionality reduction we have developed guidelines to identify a bandwidth of dimensions to retain and reduce error in the data. This approach is tested with two corpora and we can demonstrate sufficient sensitivity of our approach. In chapter 3.2 we present and discuss a validation study in which we apply our method within a real life context. Based on data collected from students at the psychology faculty of the Open University of the Netherlands we evaluate the method developed in the first empirical study in regards of *reliability* and *validity*. In this study we confirm the method evaluated in the first study on a psychology corpus and we present the results of the model performance assessment.

In section 4 we present the technology development part of the thesis. In chapter 4.1 we describe the role of a *placement support service* within the TENCompetence infrastructure and discuss its impact on the wayfinding support for open educational practices. Then we introduce two technological artifacts that have been developed within the PhD project. In chapter 4.2 the placement web-service prototype is discussed. We introduce the architecture of this web-service and present results of a first technical evaluation of this web-service. In chapter 4.2 we introduce the Semantic Weblog Monitoring Framework (SWeMoF). SWeMoF is a rapid text mining infrastructure for the social web. In this prototype we have implemented methods to construct corpora easily from sources of the social web. This prototype addresses the problem of availability and construction of corpora on the one hand and on the other hand it goes beyond the content-based approach presented in the thesis. The framework offers the possibility to work with other data and to extend the application of the text vector space model and dimensionality reduction with methods from data mining, specifically classification and clustering.

In section 5 we discuss the research project as a whole. We review the results achieved in the project and discuss its scope and limitations. Furthermore we discuss the practical implications and provide an outlook on future research perspectives.

Section 2: Theoretical Foundations

Section Introduction

In this section we introduce the theoretical foundations of our research project. The theoretical foundations are based on the one hand on a data-driven research agenda for placement support and on the other hand on the educational perspective on learner placement. To cover both perspectives we develop in chapter 2.1 a long-term research agenda for placement support. This long-term research agenda describes three different situations for placement support taking into account the data available about the learner and the learning content. The situations are ranging from unstructured information over metadata to ontological information. For each of these situations we summarize the state of the art and discuss several existing standards that can be used in placement support for situation where structured information or metadata are available. Hence, in this thesis we focus on the most complicated situation where no metadata are available.

In chapter 2.2 we develop a model for a placement support service that can operate in this specific situation and support traditional APL procedures and the placement problem in learning networks. Based on a review on existing approaches for eportfolio assessment and on the basis of the assessment triangle of Pellegrino, Chudowsky, & Glaser (2001) the contribution of a placement support service is discussed within APL procedures.

Chapter 2.1: Positioning of Learners in Learning Networks with Content, Metadata and Ontologies[2]

[2] Based on Kalz, M, Van Bruggen, J., Rusman, E., Giesbers, B., & Koper, R. (2007). Positioning of Learners in Learning Networks with content, metadata and ontologies. Interactive Learning Environments, 15, 191-200.

Abstract

Placement in learning networks is a process that assists learners in finding a starting point and an efficient route through the network that will foster competence building. In this chapter we discuss the placement problem in learning networks and introduce a long-term research agenda for placement in learning networks. We discuss several cases and give an outlook on the development of a placement support service for learning networks.

Introduction

Technology enhanced lifelong learning promises learners the possibility to learn and build competences in every context and every phase of their life. To meet this promise the individual should stand in the centre of every effort in lifelong learning instead of institutions and organizations. The concept of learning networks offers a framework to bridge the different distributed parts of current technology enhanced lifelong learning (Koper, Rusman, & Sloep, 2005). A learning network connects actors, humans as well as agents, institutions and learning resources which are organized in competence development programs. Information and communication technologies are used in such a way that the network self-organizes. The actors in the learning network share one common goal: furthering the development of competence by learners. A common approach to overcome the limitations of institutional dependencies in is the concept of Accreditation or Recognition of Prior Learning (APL/RPL) (Merrifield, McIntyre, & Osaigbov, 2000).

APL offers methods and techniques to identify prior learning experiences from formal and informal education. This procedure is especially important if a person crosses the boundaries between work and learning or between academic disciplines. Most of the methods for APL rely on experts who study the learners' profiles and decide which parts of educational programs could be exempted and which ones are best suited as starting point for the students. However, this way of analyzing prior learning experiences is a very time-consuming and expensive method. We propose therefore as an alternative the usage of computational approaches to address this problem for lifelong learning in learning networks. Previous work at the Open University of the Netherlands focused on content-based approaches to address this problem in learning networks (Van Bruggen, Rusman, Giesbers

& Koper 2006). This article widens the focus to metadata and ontologies and provides a research agenda for an ongoing project that aims at the research and development of a web service for placement in learning networks.

Placement in learning networks

Placement in learning networks has to take into account various forms of learning, including non-formal learning. No matter if the competence development programs are formal or informal, learners engage in series of learning activities that may take a long time to complete. In learning networks for lifelong learning, prolonged interruptions of such series of learning activities are likely to occur. Moreover, learners may engage intermittently in different types of learning. Whenever such a learner enters or returns to a learning network we are faced with what we call the 'placement problem': Taking into account the goals and the history of the learner, what route or routes of learning activities through the learning network can we advise and what is the best place for the learner to start (Van Bruggen et al., 2004)?

Placement means to compare the already acquired (levels of) competences of a learner to the (levels of) competences that result from a particular competence development program in the current learning network. Assume that this learning network contains pre-arranged routes towards particular goals and that every route is a competence development program. Then, the placement problem is one of determining which learning activities in the routes need to be completed and which ones can be skipped, because they do not add to the competences, skills and knowledge that the learner has acquired in the past. How exactly the competences of the learner and his history can be mapped onto the learning outcomes of activities in the learning network, depends to a great extent on the given data. Learner data may result from formal, accredited learning as well as from experience gained in informal learning situations. Description of competences may range from completely absent to being based on an ontology or at least a controlled vocabulary. To address the placement problem in learning networks we assume that learners will enter a learning network with a variety of different data stored in learner profiles or electronic portfolios.

On the other hand the learning network itself can consist of a loosely coupled collection of material or a very well structured collection of learning activities where we have a connection between them and competences or competence levels. We surmise that alternative approaches to placement need to be based on the type of competence descriptions (of learners as well as programs) that are available. We seek computational approaches to placement that ultimately fulfil the criteria of reliability (the same situation leads to the same recommendation) as well as validity (the recommendation matches that of experts). A reliable placement support service has to bring always the same result from the same given data, while the validity can only be compared to human performance for the placement problem. The different situations and data for placement are shown in the Placement Situations-Matrix in figure 2.1:

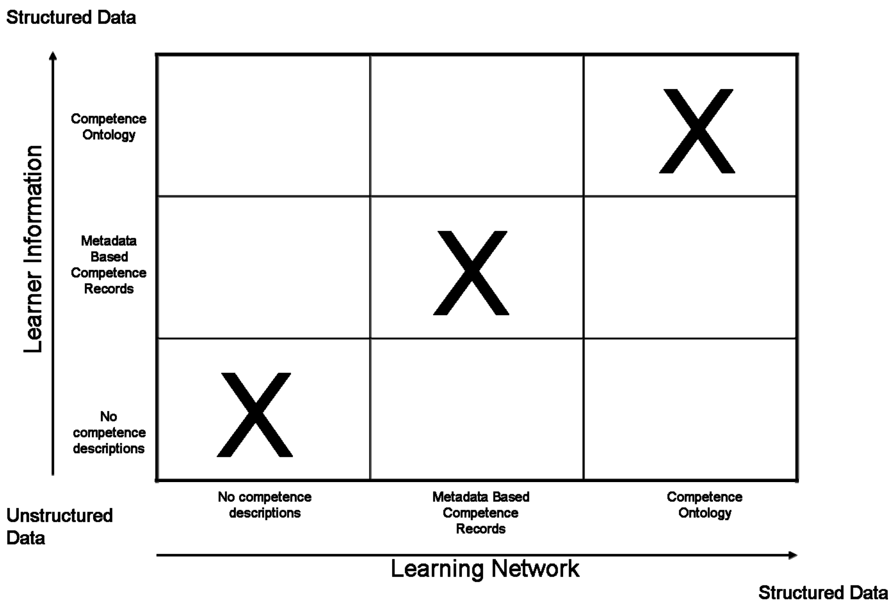

Figure 2.1: The Placement Situations Matrix

The three cases discussed here represent the "symmetrical placement" where similar data are compared. The more complicated placement would be the "asymmetrical positioning" where for example a competence ontology in the learner profile should be mapped to the content of a competence development program. To cover all these different situations and to ensure the best achievable position inside a learning network we compare different situations and approaches. In this chapter we limit the discussion to three cases of symmetrical positioning:

Case 1: Informal descriptions
The learner enters an educational environment without any explicit competence descriptions. The competence development program is highly informal without information about the resulting competence of learning activities. An example: A learner wants to update his competences in accounting. He enters a learning network that deals with the finance and accounting domain. His competence development goal will be reached through an informal collection of learning activities. His electronic portfolio contains only some documents he produced in his former education. Here, a content based approach, as discussed in the next section is best suited for placement.

Case 2: Metadata based placement
If a learner enters with a standards-compliant ePortfolio the situation would be different for a placement support service. To take the same example as in case one the learner enters with a standards-based description of his competences and the activities in his chosen learning network have detailed information about requirements or competence result for a learning activity. In section four we review and discuss standards and the way they can support the placement process.

Case 3: Ontology-based placement
If there are competence ontologies inside the learner profiles and the competence development program the placement problem can be based on mappings between the ontologies. The same learner as in case one and two enters a learning network with a very detailed competence-ontology or competence map that shows his already acquired competences. The learning network contains an agreed upon domain-ontology where all aspects of the domain are modelled. Additionally a competence-ontology has relations to the domain ontology. This highly structured description in the learning networks allows a direct comparison between the competence-ontology in the learner profile and the competence-ontology in the current learning network.

A Content-Based Approach to the Placement Problem

The rationale and the research agenda for a content-based approach to placement was described in (Van Bruggen et al., 2004). The approach rests on the following assumptions. Because it would require extensive assessment it does not aim to directly demonstrate that the learner has already acquired knowledge, skills and competences that are equivalent to the outcomes of learning activities within the routes considered. The core assumption is that equivalence of outcomes will be reflected in, or can be approximated by, the similarity of the contents of (learning) materials studied or produced by the student (source material) and the material contained in the learning activities in the learning network (target). If a placement support service determines that the content of source and target materials overlap substantially, the target activity is exempted. In our content-based placement support service document similarity is computed using latent semantic analysis (LSA) (Deerwester, Dumais, Furnas, Landauer, & Harshman, 1990).

LSA is based on word (co)-occurrences in documents, thus all order (syntax) of words or semantics in the original documents is ignored. All analyses are performed on a Term-by-Document matrix with word frequencies in the cells. The dimensions of this matrix are computed and the largest dimensions found (the semantic factors) are retained to reproduce the original matrix (Landauer & Dumais, 1997). In the reproduced matrix each document is represented as a vector. The smaller the angle between two document vectors the higher they are correlated, that is, they are expected to contain materials that have substantial overlap. Learners are represented by one or more documents that they have produced or studied. If one or more of these learner document vectors demonstrate a high correlation with learning material vectors, then the learning material may be considered redundant. Although the content-based approach has modest requirements on the way data are expressed, there are several limitations and assumptions that we need to consider like the amount and quality of available material in the learner profiles for content-based placement.

Metadata-Approaches for the Placement Problem

Metadata are used to describe learning resources as well as learner profiles. Several efforts from standardization bodies and working groups aim at unifying competence descriptions and competence levels. The IMS Reusable Definition of Competency or Educational Objective (RDCEO) specification aims at a standard description of competences and educational objectives for online and distributed learning. RDCEO is expected to promote common understanding of competences that can be used in competency development (learning and career development) or in specifying learning pre-requisites or learning outcomes (IMS Globals Learning Consortium (IMS), 2002). The RDCEO offers a unique identifier to assign an unstructured competency description to an object for example in a Unit-of-Learning (UoL). Based on the RDCEO a draft standard for Reusable Competency Definitions (RCD) is being defined in the IEEE. Although RCD does not intent to offer a solution to the aggregation of competences from sub-competences the data-model allows the integration of relational information or competence ontologies through embedding additional metadata (IEEE LTSC, 2006). For portfolios two specifications are of interest. The IMS Learner Information Package Specification (LIP) is designed to package learner information for the exchange of data (IMS, 2001). The IMS ePortfolio specification builds on the LIP specification to ensure portability and exchange of ePortfolio records for learners (IMS, 2005). The specification is addressing different usage possibilities (assessment, planning of learning) and it can store produced artifacts form the learner and formal achievement records like references. A slightly different approach comes from the HR-XML Consortium. The consortium develops a standard suite of XML-specifications to allow the exchange of Human-Resource-related data, such as a competency schema for a variety of business contexts that is applicable in recruitment processes (HR XML, 2004). The model allows the evaluation, rating and ranking of competences which are an important issue in recruiting processes.

While these metadata were all related to the learner profiles different standards in the learning network are also important for the placement problem. The IMS Learning Object Metadata (LOM) is used to assign metadata to learning objects. For the placement support service it is important that there is no element in the LOM standard to store competence related information at the moment (Ng, Hatala, & Gasevic, 2007). They could be stored in the educational segment of the metadata as proposed in (Sánchez-Alonso & Sicilia, 2005) but this does not seem to be a widely adopted solution to the problem. On the authoring level, IMS Learning Design (IMS LD) can also be used to take into account prior knowledge (IMS, 2003). One of the use cases states that IMS LD can be used to reduce the content in a learning path to reduce the time required to reach learning objectives. Through conditions a learning activity can be skipped if a learner already knows enough about the specific subject of the learning activity.

For the placement support service not only the specifications and standards are important but also the way they can be compared to each other. Since the underlying data models differ in the above presented standards the interoperability of competence related metadata is a problem. (Chan & Zeng, 2006) present different options for comparing and mapping metadata. One solution for this problem is the development of a crosswalk for competence-related metadata. A crosswalk is a specification for mapping one metadata standard to another (St.Pierre & LaPlant, 1998). While crosswalking works well when the number of involved schemas is small it can become a very complicated matter when a high number of schemas should be compared. Especially the exponential increase of relationships for the different schemas leads to a complexity problem. Therefore Chan & Zeng (2006) introduce the method of metadata switching. Instead of using a many-to-many relationship the switching method uses one schema as the central relation for all other schemas. Other options for interoperability are a metadata framework or a metadata registry. As a metadata framework tries to integrate all solutions in a common architecture a metadata registry collects information about different schemas in an environment and allows crosslinking and mapping.

One can imagine that it could still make sense to combine these approaches with one that is based on the content of learner profiles and the learning network. The specifications discussed here allow the integration of external competence models. They make (meta-)data available for a placement support service, and may serve the purpose of opening more data for content-based placement. The standardization activities alone, however, have a limited usefulness for competence mapping and the formalized description of complex competence relationships. The interoperability standards discussed above serve the purpose of sharing data. They themselves do not ensure the semantics of the data, i.e. there are still different ways to describe the same learning outcomes, such as competences.

Placement with Competence Ontologies

The missing link between the standards and the competency mapping may emerge from the use of competence ontologies and semantic web technology (Koper, 2004). Ontologies are metadata schemas providing a controlled vocabulary of concepts and they can be useful to share common understanding in a domain in a machine-readable way. For competence development ontologies or taxonomies can be used to define competences related to learning activities. Competence ontologies could be either added to the learner profiles (Dolog & Schafer, 2005), learning objects (Ng, Hatala, & Gasevic, 2007) or the competence development programs (Woelk, 2002). But the design and implementation of competence ontologies is still a very complex and time consuming task. Su (2002) presents three different situations for ontology matching: the single ontology approach where all information sources are related to a unified global ontology, a multiple ontology approach where every information source has its own ontology without a shared vocabulary and a hybrid approach where all information sources have their own ontology but they use a unified shared vocabulary. In an ideal situation every learning network could share a common understanding of the competences needed for successful running through a competence development program based on ontologies. In this case placement inside a learning network can be based on the relations between a domain ontology and the competence ontology (Posea & Harzallah, 2004). The process of adding competences in the learner profile could be derived from successfully finished learning activities in the learning network. Parts of the competence ontology in the learning network could be added to the learner profiles step-by-step after they have successful passed the related assignments. For the multiple ontology-approach, ontology similarity is the key factor for successful placement (Maedche & Staab, 2002). In the next part of the chapter we will discuss the presented approaches and try to give an outlook for our research on placement in the future.

Discussion

Placement of a learner in a learning network for lifelong learning is a complex task by itself and this is exacerbated by conditions that prevent any simple mapping of learner profiles and competency descriptions onto the educational resources. The two most extreme situations that we considered are the clearest: (1) no competency descriptions inside the learner profiles and the learning network and (2) competence ontologies in the learner profile and the learning network. In the first case a content-based approach is the one to take. The content-based approach to the placement problem has the advantage that it can be used for placement right now, where most learners do not have a detailed profile with explicit competence descriptions. The drawback of the approach is that it is only related to the produced content of the learner and not to his earned competences. So the success is dependent on the amount of text the learner can provide in relation to his educational history. If he can for example only add content from several parts of his educational background, the placement recommendation will be biased. Additionally, the concentration on content may effectively limit the approach to domains with a strong verbal character. For the same reason, domains with psycho-motor content, for example practical skills, may not be adequately represented. In the second case a mapping of ontologies could be a feasible technique to reach the ideal position for the learner.

For the placement problem all the data models can be useful because having machine-readable information about the competences of the learner simplifies the placement task if we have also competence descriptions inside the chosen learning network. But, there are drawbacks to this approach and those related to it: all the presented metadata-based initiatives offer a way to ensure a standardized description of competence related data. The models differ from openness (in terms of the possibility to embed ontologies or taxonomies) and intention (packaging focus or description focus). But the biggest drawback with metadata- and ontology-based approaches is the economical side of the medal. A huge amount of work has to be invested to enrich learning resources and learner profiles with metadata and competence ontologies. Another problem is that metadata and ontologies are always arbitrary models to a knowledge domain and that objective ontologies do not exist (Shirky, 2003). Besides it is very expensive to let domain experts guarantee the quality of the metadata used. Several experiences from repositories have shown that it is not an advisable idea to pass the burden of metadata-enrichment to the users. The quality of user created metadata cannot be compared to the quality of experts (Barton, Currier, & Hey, 2003).

Outlook

Our research will focus in the future on computational approaches to address the three presented cases for placement of learners in learning networks. While there are already several individual experiences in all of the presented cases a combination of them is a new and previously untried approach to address the placement problem. Because of feasibility reasons we will concentrate for the development of a placement support service on the most complicated situation first. All experiments will be done in an introductory psychology learning network of the Open University of the Netherlands. While the focus of the project is on the comparison of similar data it is still an open question how an asymmetrical positioning could be addressed.

Chapter 2.2: A Model for New Linkages for Prior Learning Assessment[3]

[3] Based on Kalz, M., Van Bruggen, J., Giesbers, B., Eshuis, J., Waterink, W., & Koper, R. (2008). A Model for New Linkages for Prior Learning Assessment. Campus Wide Information Systems, 25(4), 233-243.

Abstract

The purpose of the chapter is twofold: First the chapter should sketch the theoretical basis for the use of electronic portfolios for prior learning assessment. Second it should introduce Latent Semantic Analysis as a powerful method for the computation of semantic similarity between texts and a basis for a new observation link for prior learning assessment. A short literature review about e-assessment was conducted with the result that none of the reviews included new and innovative methods for the assessment of open responses and narrative of learners. On a theoretical basis the connection between ePortfolio research and research about prior learning assessment is explained based on existing literature. After that, Latent Semantic Analysis (LSA) is introduced and several examples from similar educational applications are provided. A model for prior learning assessment on the basis of LSA is presented. A case study at the Open University of the Netherlands is presented and preliminary results are discussed. Some limitations of the approach are discussed in the chapter that will be the basis for future research with similar purpose (e.g. the modeling and assessing cognitive growth in a domain based on narrative of learners).

Introduction

Although technology may have lead to educational innovations in some institutions most assessment practices of today are still the same as 10 years ago. McDonald, Boud, Francis, & Gonzci (1995) argue that students can escape bad teaching bad not bad assessment. Assessment is always embedded into a social context and it influences behavior of students because it transports a message about what is appreciated in a given learning context and what is not. Sluijsmans, Prins, & Martens (2006) point to the fact that current technology-enhanced assessment practice still focuses more on testing than assessment. Additionally in most higher education institutions assessment is still done completely without the use of technology. This leads to a "bizzare practice" „where students use ICT tools such as word processors and graphic calculators as an integral part of learning, and are then restricted to paper and pencil when their "knowledge" is assessed" (Ridgway, McCusker, & Pead, 2004).
For the use of computers in testing and assessment different concepts like computer-assisted assessment (CAA) or eAssessment are used. Conole & Warburton (2005) present a review of computer-assisted assessment.

According to them computer-assisted assessment includes also optical mark reading to analyze paper-and-pencil tests and the use of portfolios to collect learning products. Computer-based assessment (CBA) is – according to them – the use of computers to "mark answers that were entered directly into a computer" and they differentiate between web-based, networked and standalone CBA. Ridgway, McCusker, & Pead (2004) conducted another literature review on e-assessment with a similar perspective. In conclusion they define an agenda for the future of technology-enhanced assessment that includes the assessment of metacognition, the analysis and assessment of cognitive processes and the support of reflection and critical thinking skills.

Apparently all of the above mentioned reviews of the field of technology-enhanced assessment do not mention several new approaches to analyze and score open responses or narrative text from learners. This chapter introduces a new method and technique to assess students' prior learning through the use of electronic portfolios in combination with a content analysis technique called latent semantic analysis (LSA). In the next section we will provide context for our assessment approach and present an assessment framework. Next we introduce the electronic portfolio as an important technological advancement for assessment practice and define its role in prior learning assessment. Third we introduce a model for prior learning assessment with Latent Semantic Analysis as and (electronic) portfolios. Fourth we report about a case study we conducted in the framework of the European integrated project TENCompetence, and finally discuss preliminary results and give an outlook on future research.

New Linkages for Prior Learning Assessment

While traditional assessment is focused on the comparison of learners in competence based educational programs assessment judgements should be based on comparisons between individual performance and performance requirements set in a standard or learning target description. Competence-based assessment is not a traditional examination but a process in order to collect evidence about the performance and knowledge of a person with respect to such a competence standard. Joosten – ten Brinke et al. provide an overview about the traditional and new assessment methods and they point to the difference between performance assessment and competence assessment (Joosten-ten Brinke et al., 2007). While performance assessment is focused only on an isolated part of a "performance" of a learner competence assessment is much broader and can include several test and assessment types like self-assessment, peer-assessment or portfolio assessment.

A competence assessment process can use several sources to judge about the competence level of learners. These sources can stem from tests, a monitoring of behaviour or documents that were written by the learner. In the literature authors often differentiate between formative and summative assessment. While formative assessment is given during learning as a kind of feedback summative assessment is more a judgment at the end of a performance mostly connected to grading. Many students think of summative assessment when it comes to assessment situations because this is the dominant practice in higher education institutions. But especially formative assessment is a powerful tool to support students to reach high-order skills (Sadler, 1989). No matter what kind of assessment is used every assessment situation consists of several elements. Pellegrino, Chudowsky, & Glaser (2001) have developed a framework for assessment called the 'assessment triangle'. According to this framework any assessment consists of the following elements that should be made explicit.

Every assessment has an underlying model of cognition and cognitive growth in a domain. This model should be clear to assess and differentiate between low-level concepts and high-level concepts in a domain. The observation part consists of a "set of beliefs about the kinds of observations…that provide evidence of students' competences" (Pellegrino, Chudowsky, & Glaser, 2001). These observations are based on tasks or a performance that demonstrates their knowledge or skills. The interpretation part is about making sense of this evidence. New assessment methods can provide new linkages between the aspects of this framework.

In our project we focus on providing a new linkage from observation to interpretation for the assessment of prior learning. In some European countries and in Canada this issue is addresses by a procedure called APL/RPL (Accreditation/Recognition of Prior Learning) or PLAR (Prior Learning Assessment and Recognition). PLAR is used in the admission phase of educational programs to assess possible prior learning experiences and to allow exemptions in the study program chosen (Merrifield, McIntyre, & Osaigbov, 2000). The decisions for exemptions are based on prior output of learners. In a typical case the students send in material they have written in their former education or work context. Domain experts of the institution have to decide about possible exemptions after analyzing this material. The result of the time-and cost-intensive procedure is an individualized curriculum.

For technology-enhanced learning Nordeng, Lavik, & Meloy (2004) reformulate this problem in the following way: "How can the students themselves be able to assess their position relative to a future learning environments consisting of a diverse set of learning activities from which learners somehow may take their pick? The learner's history and goals define an entry position relative to the learning activities. A different entry position is likely to result in a different partition of the set of available activities in activities to skip and to complete".

Later on we will present Latent Semantic Analysis as a new linkage for the assessment of prior learning as introduced in (Van Bruggen et al., 2004). But first we will discuss the role of the electronic portfolio in assessment and accreditation of prior learning.

ePortfolios in APL

The implementation and use of electronic portfolios (eportfolios) has been recently discussed intensively although the targets of the electronic portfolio roadmap to equip every citizen of Europe with an ePortfolio until 2010 were too courageous. Barker (2006) states that "the word "ePortfolio" has almost become a code word for a variety of important concepts … an ePortfolio can be one of many different things depending on audience perspective and purpose". We see electronic portfolios as digital collections of what a person has learned or produced over time. This includes the products as well as the process to these products. Reformative educationalists like Freinet introduced the use of portfolios in his classrooms already in the 1920ies of the last century. Although the technical progress has changed tremendously since then the targets for using portfolios in education have stayed nearly the same. Documentation and self-reflection of the learning process are the main reasons to use portfolios in learning and competence development (Tillema, 2001).
Electronic portfolios can serve several roles in competence development. Smith & Tillema (2003) introduce different types of portfolios to clarify the many interpretations of this instrument: The dossier portfolio, the training portfolio, the reflective portfolio and the personal development portfolio. A dossier portfolio is a collection of performance proofs for entry to a

profession or programme. A training portfolio is an exhibit of learning during a programme, which focuses on products or competences build from the time the learners participate in the programme. A reflective portfolio is a composed collection of evidence of a specific competence requirement consisting of best-practices in combination with a self-appraisal. A personal development portfolio is a documentation of professional growth of an individual over a longer time that might also include discussions with peers with similar interest.

Although all types of electronic portfolios are important for the lifelong learning perspective for our focus the dossier-type electronic portfolio is the most important one. In the process of prior learning assessment the electronic portfolio is at the same time a means and an outcome of the assessment situation. Barker points to the conjunction between (electronic) portfolios and prior learning assessment. The PLAR procedure is often the starting point for an electronic portfolio. Learners pick products from their prior education and enrich them with additional more structured information. But the authors see much more potential for the use of electronic portfolios if they are used continuously: "The idea of developing an ELR [electronic learning record] in advance of choosing a training option or seeking career advancement is not unconventional, however, it is made more by the application of assessment techniques and principles inherent in good PLAR prior to choosing a training option or seeking career advancement, to help make those decisions, rather than after making decisions and seeking, e.g., advanced placement in a course or program" (Barker, 1999).

The electronic portfolio can serve indeed as a good tool to support these advanced placements decisions. But the electronic portfolio alone is not enough because it can only help to support the observation part of the above presented framework because it offers learners a place for documentation and reflection. To provide computer-support also in the assessment linkage between observation and interpretation we introduce Latent Semantic Analysis in the next part of the chapter as a method to assess the prior learning of students and to support these placement decisions.

A model for Prior Learning Assessment with Latent Semantic Analysis

Latent Semantic Analysis (LSA), in the past sometimes referred to as Latent Semantic Indexing (LSI), is a theory and method for extracting and representing the contextual-usage meaning of words by statistical computations (Landauer, Foltz, & Laham, 1998). It provides a method to calculate the similarity of text or parts of textual information. The whole process of this analysis consists of several steps like the pre-processing of the text, some weighting and normalizing mechanisms, the construction of a term-document matrix and a mathematical function called singular-value decomposition (SVD), which is similar to factor-analysis. The end result of this process is a latent semantic space, in which the main concepts or documents of the input are represented as vectors. Concepts and documents in this space are similar if they appeared in the same context and so their vectors are close together in the space providing a measurement for the similarity of text. LSA is applied in several research fields like informatics, psychology or medicine.

For technology-enhanced learning the application of Latent Semantic Analysis can help to solve some basic problems like increased tutor load or formative feedback during learning. Since LSA is only a general "theory of meaning" as one of the inventors of the technique, Tom Landauer, stated it recently, there are several applications of LSA in technology-enhanced learning (Wild, Kalz, Van Bruggen, & Koper, 2007; Landauer, 2007). The most prominent example for the use of LSA in an educational environment is the assessment and feedback of free text in intelligent tutoring systems. Some examples of these applications are the Intelligent Essay Assessor (Foltz, Laham, & Landauer, 1999), Summary Street (Steinhart, 2001) and Select-a-Kibitzer (Wiemer-Hastings & Graesser, 2000) to mention only a few. Some researchers have used LSA to provide students with text that is appropriate to their current knowledge (Wolfe et al., 1998; Dessus, 2004).

Our application of LSA is similar but has a different motivation and context. In the framework of the European Integrated project TENCompetence we are currently aiming at the development of an infrastructure for lifelong competence development (Koper & Specht, 2008). We are using LSA to assess prior knowledge of learners for placement or positioning decisions and finally the construction of personalized learning paths or individualized curriculum through a learning network. The model for the application is presented in figure 2.2.

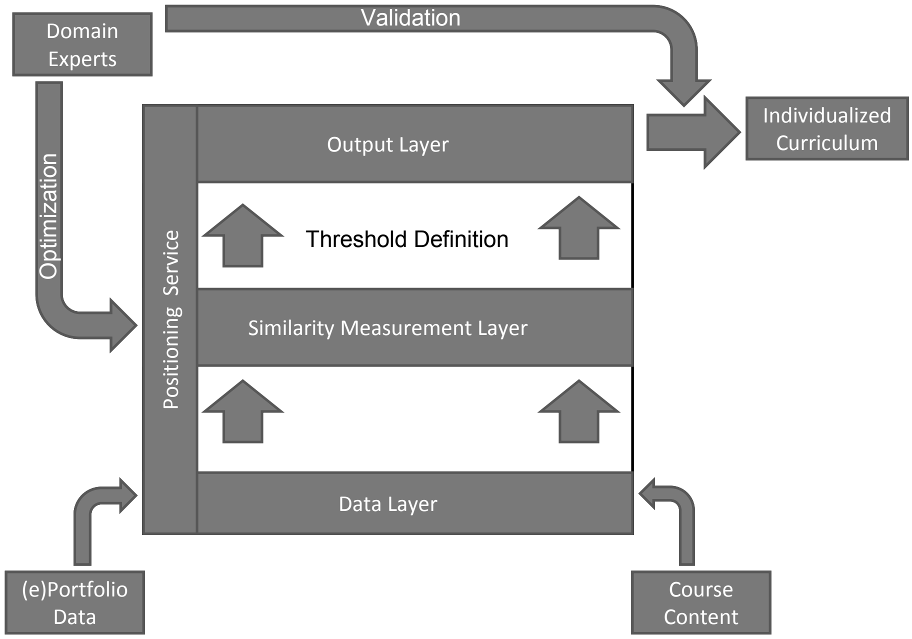

Figure 2.2: Positioning Service Model

The content of courses and data in (e)portfolios of students is compared regarding their similarity based on the assumption that the similarity of concepts in a domain and a personal portfolio will give an indication about the student's prior knowledge for this domain. Domain experts are used to validate the model and to help in optimising the results from the positioning service. The result of these analyses should be taken into account for the creation of a personalized learning path/individualized curriculum. Some learning activities on the way to the target competences a learner wants to achieve may be exempted because of the results of this prior learning analysis. In the next part of the chapter we present a case study about this application of Latent Semantic Analysis.

Prior Learning Assessment Case Study

To test our model and the usefulness of LSA for prior learning assessment we conducted a case study in an introductory psychology course at the

Open University of the Netherlands. The course was an online course consisting of 18 learning activities based on a textbook. Every chapter covers a subtopic of the psychology domain. Students were asked in advance to build a dossier-type portfolio of products they produced in their past education or work context. Since we could not expect that students knew exactly which topics would be presented in the chapters they have been asked again after every learning activity, how much of the presented material was new for them.

We used Latent Semantic Analysis to analyze the similarity between the students' documents and the content in the learning activities of the course. The basic corpus to build the semantic space consisted of other psychology books, texts from the Dutch Wikipedia and the content of the course. All student documents were "projected" into this latent semantic space and we calculated the cosine similarity measure between the student's documents and the learning activities of the course.

Depending on the policies of the current environment the learners could get exemptions for learning activities with high similarity measure. To evaluate these results we are currently conducting an expert validation. Domain experts were asked to rate the similarity of documents and to decide about exemptions based on this similarity. Another measure we are interested in is the time that experts spend to come to a decision because one of our main reasons to research technology-enhanced assessment for prior learning is the increase of the efficiency of today's assessment practice.

Preliminary results

The results of the analysis are promising. A first inspection of the results shows us that the similarity measurement that are produced by the system can differentiate between learners who sent in different material and between the learning activities and chapters. While the material of some students who sent in non-scientific psychological content produced very low values a bachelor thesis in psychology that has been collected from a colleague produced high values to the learning activities that show a topical similarity to the thesis. Table one shows a (cosine) similarity measure table between learning activities and documents in an electronic portfolio. While some documents in this portfolio show low values there are several very high results.

Table 2.1: Cosine similarity measure matrix as a result from LSA analysis of eportfolios/course content

Learning Activitiy/ Student Documents	Learner Document 1	Learner Document 2	Learner Document 3	Learner Document 4
Learning Activity 1	0.34	0.38	0.44	0.51
Learning Activity 2	0.26	0.31	0.28	0.35
Learning Activity 3	0.24	0.41	0.20	0.29
Learning Activity 4	0.33	0.46	0.23	0.29
Learning Activity 5	0.94	0.90	0.24	0.39
Learning Activity 6	0.26	0.50	0.30	0.53
Learning Activity 7	0.51	0.28	0.89	0.33

In the TENCompetence project a so called "placement support service" delivers these results to a navigation service so that learning activities with a very high correlation can be exempted for the recommendation of the next best learning activity and in the future for the construction of a personalized learning path. Another possible application is the support of the traditional PLAR procedure. LSA can support the domain experts to analyze student's material. In the next part of the chapter we discuss some limitations of the presented approach and give an outlook on future research.

Discussion and Outlook

Although the results of the presented approach are encouraging we have to keep in mind that an assessment situation has more elements according to the framework presented above. While we provide here a new linkage between the observation and interpretation part the results of the analysis still need interpretation. In addition, it has to be clear which model of cognitive growth is the basis for the assessment. Especially in domains where a high level performance cannot be measured through textual expression the presented approach will not be of much help.

But there are more limitations of the presented approach. Some limitations are connected to the use of electronic portfolios in general and some limitations stem from the use of Latent Semantic Analysis to analyze prior learning.

A general problem of electronic portfolios – especially in the context of lifelong learning – is an issue like portability of the electronic portfolio as a whole and the collected artefacts (Carroll & Calvo, 2005). Since there are several technical standards like the IMS ePortfolio standard (IMS, 2005) or the IMS LIP (IMS, 2001) we believe that this problem is merely an implementation and development issue. Every electronic portfolio system should be based on such standards to guarantee the portability. Another more general issue of the use of electronic portfolios is the validation and verification of evidence submitted. Especially in times where plagiarism in higher education is increasing the origin of artifacts is an important issue that involves also ethical implications and trust issues (Barker, 1999). Is the presented work really done by the owner of the portfolio?

Other issues stem from the use of Latent Semantic Analysis. LSA results depend on several corpus factors and pre-processing procedures that cannot be described here into detail. An important issue for successful analysis is the size of the basic corpus that is used as a query basis for the Latent Semantic Space. In the future we will address this issue to collect experiences about the trade-off between the size of the corpus and the reliability of the results of LSA for prior learning assessment. Another disadvantage of using LSA for assessment is the limitation to highly textual domains. Competence assessment that takes into account a physical performance cannot be analyzed with the presented method. In addition LSA can only find a similarity when the concepts used by the learners are represented in the semantic space. But there are several special presentation types (forms, descriptions of experimental designs etc.) that show an inherent higher prior learning than the purely textual content can show. In this case domain experts can deduct this but LSA cannot. A real advantage of using LSA for prior learning assessment is that students do not have to think about the design of their portfolios because it is only based on textual information and it does not rely on the format, structure or design.

While we worked with dossier portfolios at this time, for lifelong learning the personal development portfolio has several implications for a prior learning assessment that does not only take into account products of prior learning but also the reflection about these products. A really continuously updated electronic portfolio could help the learner not only on a course level but for the lifelong learning perspective without the need to collect material every time when entering a new educational context again. Currently we are dealing in this project only with a content-based approach to analyze prior learning of learners. In the future we will address also more structured data like metadata and ontologies for prior learning assessment (Kalz, Van Bruggen, Rusman, Giesbers, & Koper, 2007).

Section 3: Empirical Work

Section introduction[4]

In this section we present our empirical work to evaluate our model of a placement support service. This work is based on a validation strategy that is meant to test the viability of a placement support service to reduce the work load of assessors by filtering or annotating the contents of a portfolio. In learning networks we aim for a placement service that, at least in principle, can prepare a decision on exemptions for one or more target learning activities, comparable to the assessment phase in the APL process. It depends on the policies in the learning network and authority delegated by the provider of the learning activities whether or not the system can do the actual accreditation. Before such a complex service can be offered, interim solutions seem more realistic. We can, for example consider a service that assists the new member of the learning network in the preparation of the portfolio by signaling relevant and irrelevant documents. Of course, such a service – assuming that it will reduce the content of the portfolio - will also reduce the load of human assessors in traditional APL procedures.

The validation scenario we are using demands that the placement service meets the following criteria:

1. *Scalability*: Currently, there still is no large Dutch corpus readily available on which one could build a domain corpus that would allow training an LSA-based service (Louwerse & Van Peer, 2006). LSA is often based on large scale training corpora for a domain of 10000 and more documents. Opposed to that, in learning networks we will have to deal with corpora that are considerably smaller. One may question whether the size of the corpus is a problem, since domains such as Veterinary Medicine or Psychology are potentially very large. However, the learning network will often contain specific subsets of such domains, such as “Diagnostics” or ‘Psychology on a Bachelor level”. Moreover, the content encountered in learning networks is often limited to learning

[4] Based on Kalz, M., van Bruggen, J., Giesbers, B., Rusman, E., Eshuis, J., Waterink, W., & Koper, R. (2009). A Validation Scenario for a Placement Service in Learning Networks. In Koper, R. (Ed.) Learning Network Services for Professional Development. ISBN: 978-3-642-00977-8. Berlin:Springer, August 2009.

2. material, rather than having complete sets of reference works and background documents included. In general the service should offer a roadmap to a placement service that can use large corpora in the future.

3. *Sensitivity*: The placement service is capable to identify material dealing with the same topic, whilst discriminating material that is dealing with different topics.

4. *Reliability and validity*: A placement service needs to be *reliable*, that is treating similar cases in an equal manner. In practice this means that the scores or ratings assigned to a document by the placement service need to match, within reasonable boundaries, those of a human observer. There is an upper limit here; it is highly unlikely that two human assessors would always rate student's input exactly the same; their inter-rater reliability will be less than 1.00. In practice, one may be quite satisfied with an observed inter-rater reliability of .75 and higher. The observed inter-rater reliability is the maximum that one may expect from the ratings of a placement support service. The *validity* of a placement support service is the extent to which the service actually measures what it is supposed to measure. As the reader will have noticed, this a crucial criterion because the placement service is based on the assumption that content can be used as a proxy to learning outcomes. Measuring similarities between texts is not an issue, various standard solutions to achieve that are available, but whether these similarities correspond with similar learning outcomes is the validity question.

5. *Fit in APL procedures*: The placement service needs to fit in regular APL procedures. This implies that, first we need to establish how the service can collaborate with students and assessors in such a way that the work load, in particular of assessors, is reduced. This has implications beyond considerations of reliability: the service needs to be optimized so as to avoid errors that are costly in terms of assessor time. Second, the placement service should use the results of the APL process in future developments to improve its performance. Thus, ratings of assessors of a particular case can be fed back into the service by using the material of the case as a positive or negative example. In our LSA-based placement service this type of case-based reasoning will be achieved by adding these

examples as new positive or negative standards to which student data can be compared.

In chapter 3.1 we focus on the sensitivity aspect of our validation strategy. We address specific issues related to the application of the text vector space model and dimensionality reduction on small scale corpora. In this context we cannot rely on traditional filtering methods applied in LSA research based on global weights such as Inverse Document Frequency (IDF) or entropy since these methods would filter away important domain specific terms that contribute to the discriminatory power in small scale corpora. To address this problem we evaluate the use of different stopword strategies in combination with a method to keep important domain terms. We develop guidelines how to maximize performance on two indicators: The discrimination between different document sets and high similarity within document sets. In addition we develop guidelines to reduce error in the data. To decide about a bandwidth of dimensions retained we propose and evaluate a method that takes into account the variance accounted for.

In chapter 3.2 we present the evaluation of applying our approach to classify document as relevant or irrelevant for the target learning activities in APL procedures. For this purpose we present results form a study that has applied the method to real learner data and exemptions decisions in a psychology course of the Open University of the Netherlands.

Chapter 3.1: Latent Semantic Analysis of Small-Scale Corpora for Placement in Learning Networks[5]

[5] Kalz, M., Van Bruggen, J., Rusman, E., Giesbers, B., & Koper, R. (2009). Latent Semantic Analysis of Small Scale Corpora for Placement in Learning Networks. Manuscript submitted for publication.

Abstract

Placement in learning networks is the process of establishing a starting point and an efficient route along which a learner may build competences. We investigate computational approaches to services such as placement that avoid or reduce labor-intense procedures and that are based on the contents of the learning network and the behavior of those participating in it, rather than in predefined procedures and (meta-) data. In this chapter we consider some preliminary questions related to the use of Latent Semantic Analysis as a computational approach to placement. In learning networks LSA will often be used on small-scale corpora of learning materials. In this article we develop guidelines for the application of LSA in this type of corpora and we present empirical evidence that substantiates these guidelines. Although the results reported are encouraging we discuss some limitations that need addressing in subsequent work.

Introduction

Self-organizing learning networks are seen as a solution to the problem of realizing flexible (in time, place and space), personalized (optimal suited to the learner) learning environments for lifelong learning. A learning network is an ensemble of actors, institutions and learning resources which are mutually connected through and supported by information and communication technologies in such a way that the network self-organizes (Koper, Rusman, & Sloep, 2005). In a learning network all actors are involved in furthering the development of competence, which implies that they not only act as learners but will engage in other roles as well, such as consulting expert, teacher, or librarian. Consider, for example, a learning network within the field of veterinary medicine. The learning network contains descriptions of the competences of veterinarians and it offers learning activities and knowledge resources with which learners can develop these competences. Experienced veterinarians may participate to refresh their knowledge and also offer support to students.

When a learner (re-)enters a learning network we have to face the 'placement' problem: where shall the learner be placed in this network and what route through the learning network is the most efficient, considering the competences that the learner wants to develop and the history of the learner. This boils down to identifying the learning activities that need to be completed and those that may be considered redundant considering the needs and prior experience of the learner. Accreditation of prior learning using portfolio assessment and formal assessment of prior knowledge are methods to solve the placement problem. In learning networks, however, we try to avoid these laborious routines wherever possible. Our first efforts are therefore aimed at defining computational solutions that are based on the data encountered in the network and in the learner data.

Elsewhere we described a computational approach to learner positioning in learning networks where latent semantic analysis (LSA) is used to analyze and match the content found in the learner data with the content available in the learning network (Van Bruggen et al., 2004).

LSA is a method for extracting and representing the contextual-usage meaning of words by statistical computations. It has been shown in a number of studies that LSA can produce high similarity ratings between texts which discuss similar subjects even if they do not share any words. These similarity ratings correlate significantly with those of human raters. To appreciate the analysis presented in this chapter, we recommend that those unfamiliar with LSA first read (Landauer, Foltz, & Laham, 1998; Landauer, 2007; Yu, Cuadrado, Ceglowski, & Payne, 2002).

The semantic space of the domain is created using several documents, including, but not limited to, those that the learner may have to study. These documents may be as small as individual sentences or as large as essays, articles or web pages. The collection of documents (corpus) is represented in a term-document matrix with term frequencies in the cells. This matrix is decomposed using singular value decomposition. This will yield the number of dimensions needed for an exact reproduction of the data. Latent Semantic Analysis (LSA) gives an approximate reproduction of the data by dropping the smallest dimensions and using less dimensions which is thought to better capture the semantics in the data (Deerwester, Dumais, Furnas, Landauer, & Harshman, 1990; Landauer & Dumais, 1997). In the reproduced data each document is represented as a vector. If two document vectors are highly correlated, they represent materials that have substantial overlap in their contents (and thus may be considered redundant). Learners are represented by one or more documents that they have produced or studied. If one or more of these learner document vectors demonstrate a high correlation with learning material vectors, then the learning material may be considered redundant.

This application of LSA to placement is new, yet is not very remote from well-documented educational applications as reported by (Wolfe et al., 1998), (Zampa & Lemaire, 2002) and (Kintsch, 2002) who applied the LSA-technique to select learning material appropriate to prior knowledge. We do not expect that our computational approach will completely replace the current practice of accreditation of prior learning. We intend to offer pre-processing that will help learners and tutors to identify content that is a suitable candidate for accreditation of prior learning.

Before we may hope to implement this approach to placement, we have to find answers to several questions related to the use of LSA in learning networks. In particular, we have to address how LSA can be applied to small-scale corpora. LSA is often used for document retrieval from very large document bases, containing ten thousands of documents and an input of 5000 documents to LSA is quite common. The most important question to answer is the amount of training material needed for LSA to work properly. Ideally, a large general language corpus is combined with a smaller domain specific corpus consisting of background material of the domain. One may question whether the size of the corpus is a problem, since domains such as Veterinary Medicine or Psychology are potentially very large. But in learning networks we often have only subsets of a domain as corpus material. Moreover, the content encountered in learning networks is often limited to learning material, rather than having complete sets of reference works and background documents included. Therefore, we cannot avoid catering for small-scale corpora in learning networks.

The number of documents may not be the main issue however: the corpora used in learning networks are specific to particular sub-domains. It is obvious that in these corpora the dimensionality in the data is less than in general language corpora such as those build on the basis of a complete encyclopedia. There is some empirical evidence that LSA can be robust when applied to small-scale corpora. Wiemer-Hastings & Graesser, (2000) reduced the size of their text corpus from 2.3 MB to a minimum of 15 %, the performance of LSA in terms of correspondence with human raters decreased 12%. These results are indications that LSA can perform reasonably well in small scale corpora.

There are several topics to consider when creating, pre-processing and analyzing the corpus, including document size, filtering the corpus, deciding on frequency and association measures and in particular the number of singular values to be used in reproducing the data (Haley, Thomas, Roeck, & Petre, 2005). In this article we study the effects on the performance of LSA of (a) filtering noise words from the corpus and (b) varying the number of singular values used to reproduce the data. Moreover we offer guidelines for the determination of the number of singular values to retain by taking into account the explained variance and we propose a method to filter small domain specific corpora.

Filtering the data can be done by stemming and by stopping. Stemming refers to the parsing of tokens to their semantic stem. Thus, tokens such as "hypothesis", "hypotheses" and "hypothesized" would all be stemmed to a semantic root "hypothes". Stemming has the potential of raising the semantic relevance of the results. For example, the "mountain gorilla" is called the "berglandgorilla" in Dutch, which is parsed as a single, completely different token than "gorilla" Stopping refers to the filtering of noise words, such as "the" and "not" that occur frequently and indiscriminately in the corpus. Noise terms appear throughout the corpus and do not contribute to the discrimination of documents; from a measurement point of view these terms only add error variance to the corpus. In large corpora that reflect broad domains, such as those based on an encyclopedia or a collection of web sites, the terms to be stopped may be determined by consulting a general list of word frequencies in the written language. In the small-scale corpora that we are considering, this approach is not feasible, because relevant terms may appear in higher frequencies than they do in large, general corpora. Filtering these terms would exclude them from any query, which is not a desirable option. We therefore had to apply a different stopping method that will be presented in the method section of this article.

The most important decision in any LSA is the selection of the number of singular values (i.e. the number of dimensions) that will be used to reproduce the data. Dimension reduction itself is not a goal of LSA and it has become accepted practice that for corpora with the size of 5000 documents and above, at least 300 dimensions are used. Larger numbers are not uncommon. In smaller corpora, however, we need other heuristics to determine the number of singular values.

One method that might be used to decide on the number of singular values is a Scree test (Cattell, 1966). In this test one visually inspects the data to find the place where a sudden drop occurs in the size of the singular values. This point is used as a cut-off point, beyond which additional singular values are believed to add little more than error to the data.

A second approach is to normalize the document vectors to unitary length, and retain only the singular values that have a length greater than one. This corresponds to the practice in factor analysis of retaining only the eigenvalues larger than one. However, to routinely apply normalization to the data, would loose all information related to the length of documents. One might even argue that in such a case the difference between Singular Value Decomposition and Principal Components Analysis is obfuscated.

The third approach, as proposed by Wiemer-Hastings & Graesser (2000) is to empirically determine the number of dimensions, for example, by selecting the number of singular values that optimizes a performance criterion, such as the correlation with ratings by human raters (Wiemer-Hastings, 1999). For placement we would seek the number of singular values at which there is (a) a maximum correlation between documents that have overlapping content, and (b) a minimum correlation between documents that share little or no content.

To this raw performance criterion we add the constraint that the number of singular values for the selection is chosen from a range that corresponds to a minimum and maximum amount of explained variance in the data. This additional constraint will counter, as we expect, the effects that a low number of singular values may lead to inflated correlations, whereas a large number of singular values will attenuate all correlations.

To calculate this constraint, we exploit the connection between singular values and the variance they account for in the Term-Document matrices analyzed by LSA. Term document-matrices are sparse, that is, they contain many empty or zero-filled cells. Consequently, the mean of the cell frequencies is close to zero and the variance in the cell frequencies is small as well. For sparse matrices, singular values are closely related to the variance in the matrix. Consider the formula for variance:

$$\frac{\sum x^2}{n} - M^2$$

, where x are the cell frequencies, M is the mean frequency in the matrix and n is equal to the number of observations (that is the number of cells in the matrix). In sparse matrices the mean is close to zero and n is a (large) constant. The variance is therefore mainly determined by the sum of squares of the cell frequencies. Since the sum of the squared singular values is equal to the sum of the squared cell frequencies, the squared singular values, in sparse matrixes, can be used to estimate the proportion of variance that is accounted for by this singular value (or rather dimension). Thus, one may select a minimum and a maximum number of singular values that correspond with a bandwidth of variance accounted for.

In an ideal world, one would retain the number of singular values that corresponded to the reliability of the data. This would ensure that the maximum systematic variance is retained, whilst the maximum of error is being removed. Not knowing these reliabilities we can only make an attempt at reducing error in the data. We do so by filtering the data before the analysis (stopping) and by limiting the number of singular values that are used to reproduce the data after the analysis. Filtering the data is done by excluding terms from the Term-Document matrix that do not contribute to the systematic variance. The primary candidates are high frequency terms that occur in most documents of the corpus. “Stopping” these terms can reduce error from the data before analysis. However, the effect is uncertain and needs empirical testing on which we report in the next sections

In order to test the effects of filtering the data before analysis and applying our heuristics to limit the number of singular values we construed two test corpora. In the next section we report on the characteristics of these corpora and on the (separate and joined) effects on performance of filtering and the selection of the number of singular values.

Method

Haley, Thomas, Roeck, & Petre (2005) urge researchers to describe their analyses and data in more detail so that research results can be better compared and understanding of LSA and its further development and refinement are being fostered. We will therefore discuss our two corpora and the way in which we analyzed them in detail.

Preparation of the test corpora

We decided to prepare two corpora in Dutch because LSA should become the basis for several learning technology applications and semantic web services which should be implemented and evaluated at the Open University of the Netherlands. To create test corpora with characteristics that we expect in small scale domain specific corpora, we first decided to create two corpora in domains where the material would be primarily factual, so as to facilitate text comparison. Second, we decided to use sources that would be mainly textual in order to minimize pre-processing. Third, we decided, more or less random, to build two corpora: One corpus (corpus 1) around 'apes' (the Dutch term 'aap' refers to apes as well as monkeys, we will use 'apes' in this broad sense) and the other corpus (corpus 2) around the concept of "geheugen" (memory). In the following we will report details of both corpora.

The first corpus is based on a collection of documents obtained from the Dutch version of Encarta encyclopedia (Microsoft, 2005) and from a Web search of documents in Dutch on 'apes' and 'monkeys'. This resulted in a series of documents, including a number of overview articles, in which several species of apes were described. The first version of the text corpus resulted from filtering and splitting the documents from the queries. A few duplicate documents were removed and the remaining ones were manually split by the authors according to topics treated. Wherever possible, splitting was based on headers and paragraphs in the text, thus mainly creating documents of a paragraph in length. If a split would result in a document with less than 20 words, no splitting was done. All splitting was discussed and agreed between the authors. The final mean document size was 311 words.

The second corpus is based on a selection of documents taken from the Dutch psychology textbook "Psychologie" by Academia Press which is used at the Open University of the Netherlands for an introductory psychology course. We have applied the same filtering and splitting methods as with the first corpus resulting in a mean document size of 158 words.

Running statistics such as word frequencies revealed a number of spelling errors in the two corpora. These errors were corrected only when a document contained a systematic misspelling as, e.g. 'oran utan'. The final corpus contains 287 (560) documents and has a size of 611216 bytes (607007 bytes). TextStat (Hüning, 2005) reported 10736 (9661) different tokens of which 4376 (5000) occur in at least 2 documents. The sum of the term frequencies is 89401 (93114). The matrix is sparse, many cells are empty, which is reflected in a mean frequency of 0.071 (0.026) and a variance of the cell frequencies of 0.711 (0.187).

Applying stopping to the corpus

Since we could not find a satisfactory stemming application for Dutch (the language of our test corpora) we refrained from stemming. In order to prevent domain specific terms being filtered away solely on the basis of their frequencies, we compared the rank orders of the terms within the corpus to those within a more general corpus, based on several volumes of a Dutch newspaper (Bouma & Klein, 2001). Terms with the highest ranking in the corpus were selected for stopping, unless the difference with the rankings in the general corpus exceeded 50 ranks (an arbitrarily chosen number). Thus domain specific terms that are "overrepresented" in the corpus will be retained. For example, in our test corpus the keyword "aap" (ape or monkey) would be removed from the corpus if we only considered general word frequencies, but with our method it was retained. For analysis purposes a number of variants were made that would filter 0, 30, 40 or 50 percent of the occurrences of terms in the corpus.

The main characteristics of the corpus under different stopping strategies are presented in Table 3.1. As a first stopping method we exclude all terms that occur in one document only. This reduces the number of terms with about 60% (53%) to 4376 (5000) terms. Further stopping removes high frequency terms according to the rules for retaining domain specific terms. Since we are reducing high frequency terms, filtering 20 (17) terms (for the 30% stop list) reduces the overall term frequency with about 30000 (26500). The 40 and 50 pct stop list filtered 46 (39) and 135 (99) terms respectively. In the 40 pct list three terms (no term) were retained because their ranking in the corpus was substantially higher than in the general corpus. For the 50 pct list 68 (6) terms were retained for the same reason. As Table 3.1 shows, there is little variance in the data, especially after applying a stopping strategy.

Table 3.1: Statistics for terms in the two corpora under different stopping strategies

stopping strategy corpus 1	0%	30%	40%	50%
Number of terms	4376	4356	4330	4241
Sum of term occurrences	89401	59393	49736	40126
Mean frequency	0.071	0.048	0.040	0.081
Variance	0.711	0.146	0.105	0.001

stopping strategy corpus 2	0%	30%	40%	50%
Number of terms	5000	4983	4961	4001
Sum of term occurrences	88453	62016	53014	43778
Mean frequency	0.026	0.022	0.020	0.017
Variance	0.187	0.064	0.049	0.04

Performance measures: correlations

As stated earlier the placement problem boils down to the crucial question whether we can obtain high correlations among documents that share contents, and simultaneously obtain low correlations with documents with other contents. With the experiments conducted we aim at identifying an ideal interval of singular values to use depending on a bandwidth of minimum and maximum variance explained. To find this bandwidth we worked with two different sets of documents in the corpora. In corpus 1, for example, we expect that documents about gorillas are highly correlated. We expect the same documents to demonstrate low or moderate correlations with 'non-gorilla' documents. To test this performance question, we created a set of documents dealing with a particular species of apes. We refer to this collection as an H-set (H for homogeneous). For each H-set two types of correlations can be obtained: the correlations within the H-set (an indication of how well documents about the same species correlate) and the correlations of the documents of the H-set with the other documents in the corpus (we refer to those as the N-Set). To summarize these correlations we used the mean correlation within the H-Set (a measure that is biased downward) and the mean of the absolute values of the correlations with the N-set (a measure that is biased upward). These results were corrected for autocorrelation. We first performed analyses under two conditions: a zero and a 30 percent stopping strategy. Within each of these conditions a series of LSAs using different numbers of singular values was performed. For each

LSA summary correlations were computed. Finally, we replicated the results using a different H-Set for each corpus.

Parameters for the analysis

All singular value decompositions and LSA analyses were done with the General Text Parser (GTP) program (Giles, Wo, & Berry, 2003). GTP has several parameters that dictate the parsing and processing of the documents as well as the resulting LSA analyses. The analysis was based on tokens (terms) of length between 2 and 25 characters that occurred in at least two documents. Tokens that appeared on a stop list were filtered out. The singular value decomposition used raw frequencies. Other types of frequency measures, such as inverse frequencies or entropy based measures have the drawback that they will filter away domain-specific terms with relative high frequency. Document vectors were not normalized, but most of the documents in the corpora are of similar length anyway. The parameters that control the computing of the singular value decomposition were set to maximize the number of iterations as well as the number of singular values to be extracted. Cosine values were used as association measures.

Results

Singular value decompositions of the Term-Document matrix were performed under a 0 %, 30 % and 50 % stopping strategy. For each strategy the cumulative proportion of variance accounted for by the singular values were estimated on the basis of the squared singular values. Figures 3.1 & 3.2 summarizes the results of this analysis and it demonstrates that stopping leads to a less steep curve in the proportion of variance accounted for. The difference between a 30% and a 50% stopping strategy are marginal for both corpora. Recall that the 50% stopping strategy forced us to manually retain several high frequency domain specific terms in both corpora. To define the number of singular values to chose for LSA we can see that in a bandwidth between 35 and 55 singular values between 70 % and 80% of the data are represented in case of the 50 % stopping strategy (see figure 3.1).

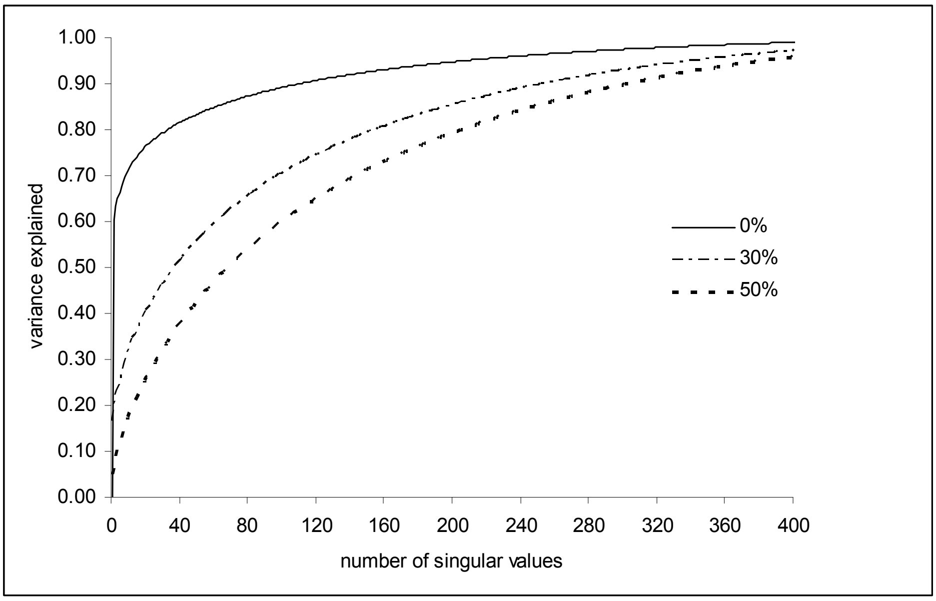

Figure 3.1. Number of singular values and variance explained under different stopping strategies for corpus 1

In corpus 2 this bandwidth is between 150 and 200 singular values which would explain 70 % to 80 % of the data (see figure 3.2).

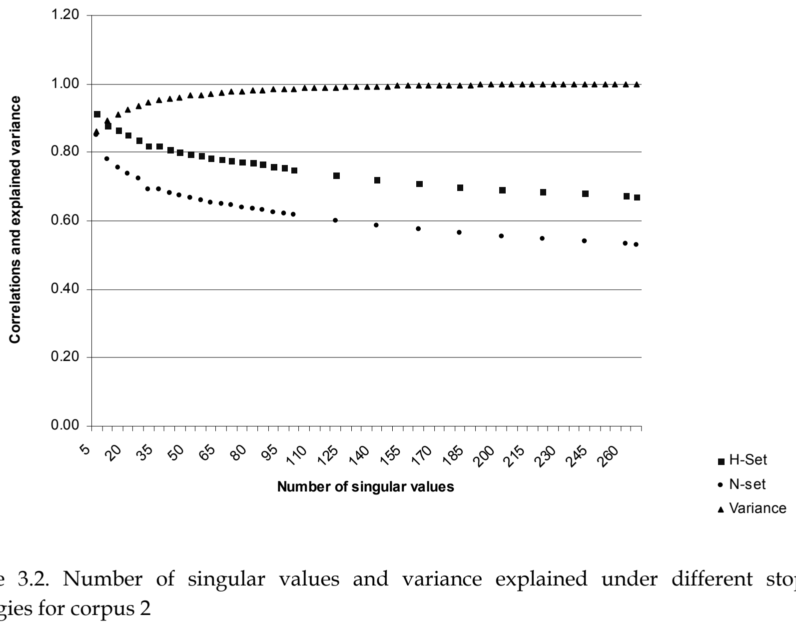

Figure 3.2. Number of singular values and variance explained under different stopping strategies for corpus 2

For each corpus separate analyses were run for the two query sets. In each stopping condition several LSAs were performed using different numbers of singular values to obtain correlations and summary measures for the H-set and N-sets.

Figures 3.3 – 3.6 summarize the results for the analyses under a 0% stopping strategy. As these figures show about 10%-20% of the number of singular values extracted account for a bandwidth of 70 % - 80 % of the variance. As was already shown in figure 3.1 and 3.2, especially the first, large singular values account for much of the variance explained.

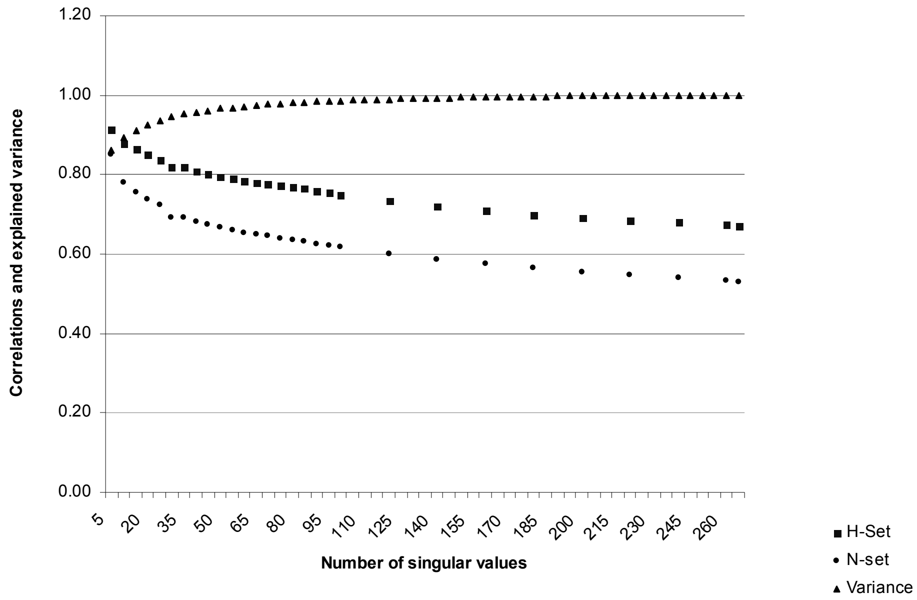

Figure 3.3. Query set 1 correlations and explained variance under 0% stopping for corpus 1

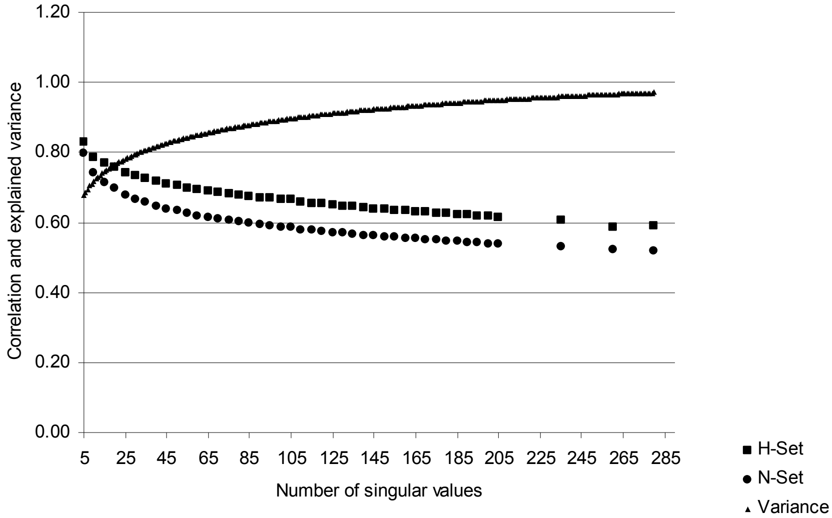

Figure 3.4. Query set 1 correlations and explained variance under 0% stopping for corpus 2

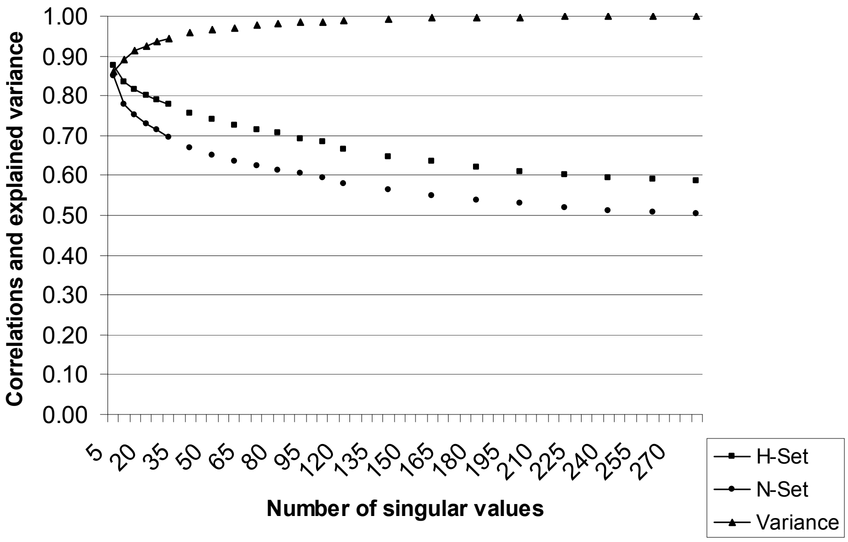

Figure 3.5. Query set 2 correlations and explained variance under 0% stopping for corpus 2

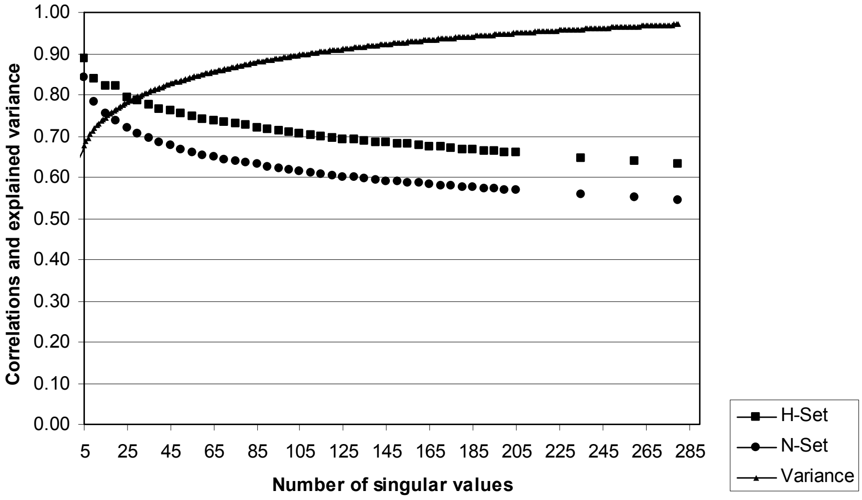

Figure 3.6. Query set 2 correlations and explained variance under 0% stopping for both corpora

Second, as expected, the correlations show a gradual decrease when the number of singular values is increased. Third, the correlations obtained from the reproduced data are considerable. In both cases, all the H-set correlations are above .50. The same applies to the correlations with the N-set unfortunately, thus making it different to differentiate between the sets. The results for the analysis under a 30% stopping strategy are summarized in figures 3.7 – 3.10.

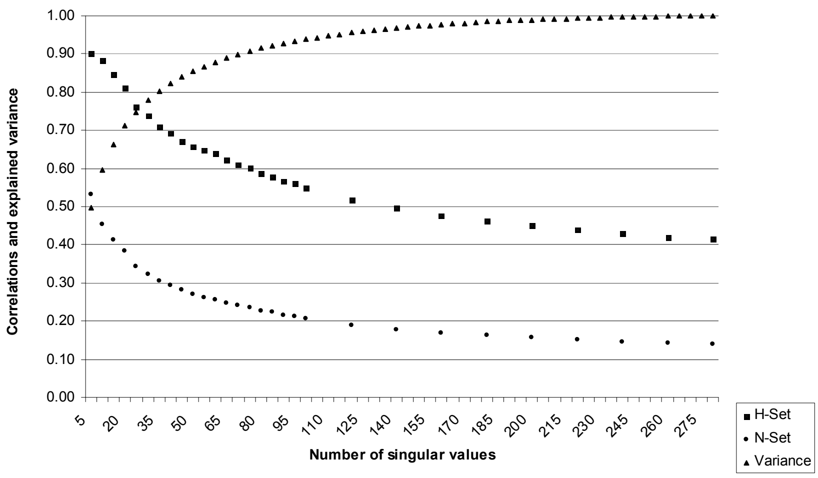

Figure 3.7. Query set 1 correlations and explained variance under 30% stopping for corpus 1

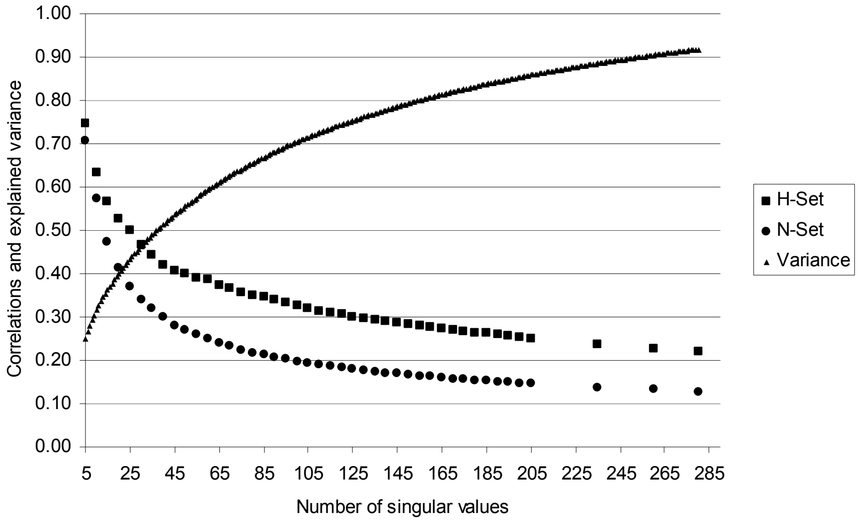

Figure 3.8. Query set 1 correlations and explained variance under 30% stopping for corpus 2

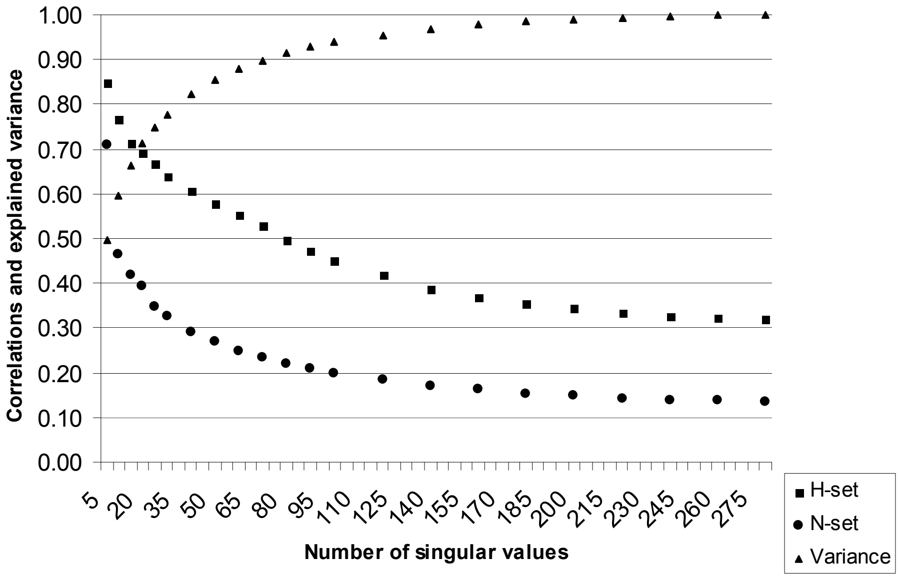

Figure 3.9 Query set 1 correlations and explained variance under 30% stopping for corpus 1

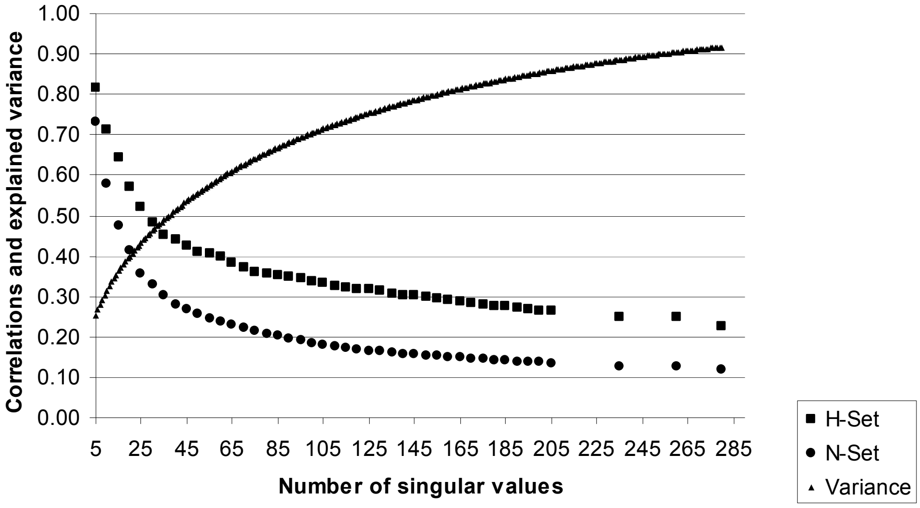

Figure 3.10 Query set 2 correlations and explained variance under 30% stopping for corpus 2

Figures 3.7 – 3.10 also illustrate that an increase in the number of singular values is associated with a gradual decrease in the size of the correlations obtained from the reproduced data. There are a number of differences as well. First, and not surprisingly, more singular values are needed to reproduce the same proportion of variance. Second, the correlations are now lower than those obtained under a zero stopping strategy, but the differences between the H-Sets and N-Set are larger. This effect is clearer for corpus 1 than for corpus 2. This means that we can discriminate better between the H-Sets and N-Set and select the number of singular values that correspond to an optimum discrimination within the range of minimum and maximum percentage of variance explained.

Since we wanted to control if the discrimination for the second corpus can be increased we ran additional LSAs for the second corpus with a 50% stopping strategy. The results of these runs are shown in figures 3.11 and 3.12.

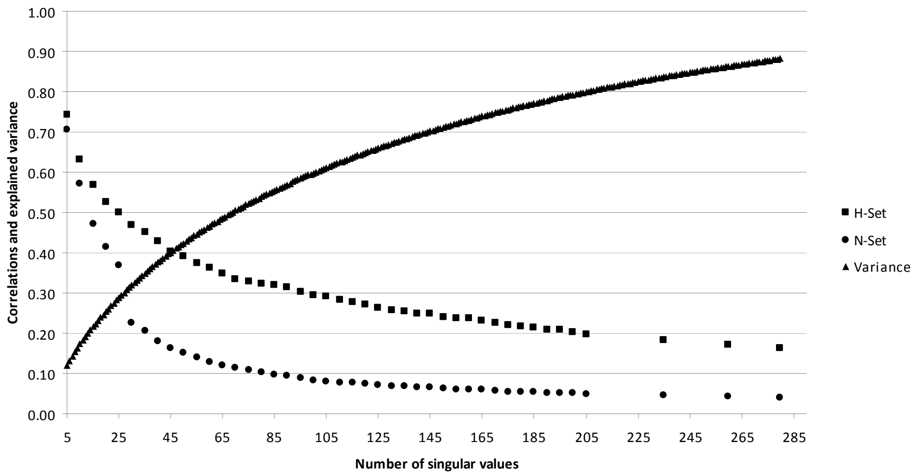

Figure 3.11 Query set 1 correlations and explained variance under 50% stopping for corpus 2

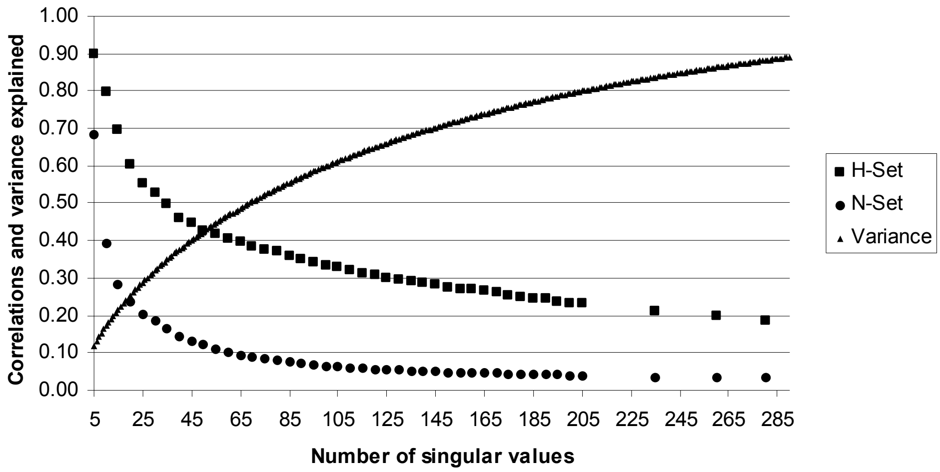

Figure 3.12 Query set 2 correlations and explained variance under 50% stopping for corpus 2

Again the correlations decreased and the number of singular values needed to reproduce the same proportion of variance increased. The increase from a 30% to a 50% stopping strategy also improved the discrimination between the H-Set and the N-Set comparably to corpus 1.

Discussion

The results presented here indicate that LSA can be applied to the type of small-scale, domain-specific corpora that we expect to encounter in the implementation of learning networks. The results also indicate that in doing so we have to consider a number of trade-offs that occur because of the omnipresence of error in the data. Beyond a certain (unknown) threshold, using additional singular values in the reproduction of the data will introduce more and more error. Clear indications of this effect are the diminishing size of the correlations that we found in all analyses. Limiting the number of singular values may help to prevent this effect. An additional measure is to filter the corpus before performing analyses. Our data clearly show that reducing high-frequency terms improves results. It was only when we applied stopping strategies that we could obtain correlations that were high within homogenous sets and low between non-related sets simultaneously. This is also an indication that optimizing for a single criterion, such as a high correlation with a target variable, may be insufficient if we also want to obtain sufficient discriminatory power. According to our results (figure 3.4 to 3.12) optimizing on H-set correlations alone would lead to selecting no more than five singular values. Optimizing for both correlations would lead to a choice of about ten singular values for corpus 1 and around 30 singular values for corpus 2, which would only account of 50% of the variance. Using joint criteria, such as to account for a bandwidth between 70% and 80% of the variance; optimize homogeneity measures and maximize discriminatory power, one may pinpoint the number of singular values that maximizes performance and explanatory power. In our corpora this will point to a number of singular values between a bandwidth of 30 and 50 singular values for corpus 1 and a bandwidth between 145 and 200 singular values for corpus 2 with higher dimensionality. In our analyses stopping was crucial to obtain the type of performance that we tried to maximize. In case of the carefully constructed corpus 1 a 30% stopping strategy delivered sufficient discrimination while in case of corpus 2 a 50% stopping strategy was needed to reach comparable discrimination between the two sets.

There are a number of limitations that need to be considered however. In the first place we have not demonstrated that LSA is able to match learners and learning materials in a sufficient manner. In a follow-up we will compare LSA (dis)similarity ratings of learning materials in Introductory Psychology to those of human experts. Moreover, we will then examine to what extent these ratings are valid approximations or predictors of prior knowledge and learning outcomes.

In the second place the approach is limited to textual materials, and one may even argue that it is limited to textual materials where syntax or temporal order plays a minor role. This limitation will obviously restrict applications to 'textual domains' and may exclude lots of skills training. For these and other reasons we intent to broaden the scope of the analysis by, on the one hand, including metadata of the materials and, on the other, integrating the computational approach in methods such as (e)portfolio assessment.

Conclusion

We described a computational approach to placement in Learning Networks that defines the problem as one of matching the contents of the learner's portfolio to the contents of the learning materials that are made available in the Learning Network. We explored some of the major methodological issues in the application of LSA and reported our results in detail to avoid some of the pitfalls mentioned by Haley, Thomas, Roeck, & Petre (2005). We demonstrated that we were able to detect correlations within homogenous content sets whilst discriminating between different sets.

We presented a heuristic to maintain domain-specific terms that have high frequencies. Although there are no mathematical restrictions to apply LSA in small scale corpora, one may question the stability of the results. We have obtained comparable results for two different corpora with similar size but different dimensionality and we have used two different H-Sets within these corpora. We expect that the approach to the selection of singular values proposed here increases the stability of results because it combines performance measures with constraints of performance optimization by the amount of variance accounted for.

Thus, returning to our initial example, there is a promise that this computational approach may detect sufficient overlap between the history of a learner, such as our veterinarian and the materials in the learning network so as to avoid recommending redundant materials to the learner.

Chapter 3.2: Where am I? – An Empirical Study on Learner Placement based on Semantic Similarity[6]

[6] Kalz, M., van Bruggen, J., Giesbers, B., Waterink, W., Eshuis, J. & Koper, R. (2009). Where am I? – An Empirical Study about Learner Placement based on Semantic Similarity. Manuscript submitted for publication.

Abstract

The Accreditation of Prior Learning (APL) is a procedure to offer learners an individualized curriculum based on their prior experiences and knowledge. The placement decisions in this process are based on analyses of student material by domain experts, which makes it a time-consuming and expensive process. In order to reduce the workload of these domain experts we are seeking for ways in which the preprocessing and selection of student submitted material can be done with technological support. In this study we have evaluated the use a method similar to Latent Semantic Analysis (LSA) to support the process of learner placement decisions. The study was conducted in a psychology course of the Open University of the Netherlands. The results of the study confirm our findings from an earlier study regarding the identification of the ideal number of dimensions and the use of stopwords for small scale corpora. Furthermore the study indicates that the application of the vector space model and dimensionality reduction produces a well performing classification model for deciding about relevant documents for APL procedures. Along with the results of this study we discuss methodological issues and limitations of our study and provide an outlook on future research.

Introduction

Some institutions for higher education allow for the accreditation of prior learning (APL) (Merrifield, McIntyre, & Osaigbov, 2000). A typical APL procedure consists of four main phases (Van Bruggen, Kalz, & Joosten-ten Brinke, 2009):

1. In a *profiling phase* the institution collects information about the learner's needs and personal background.
2. In the second phase learners *collect and present evidence* about their qualifications and experience. This evidence should support a claim for credit for the new qualification they are seeking.
3. In the *assessment phase* the evidence submitted by the learner is analyzed and reviewed confirming to the local assessment standards. The result of this phase in the procedure is an answer to the question whether the student should be granted recognition.
4. In the *accreditation phase* the results are verified by the department responsible for awarding the credit or recognizing the outcome of the assessment.

The procedures of APL are cost-and time-consuming because they involve domain experts to assess the contents of the portfolios submitted by the students. Two different ways of accreditation in higher education institutions can be compared. On the one hand there is a generalized accreditation procedure which is based for example on certificates from vocational education which are expected to be equivalent to local certificates. On the other hand there is an individual accreditation procedure which also takes into account prior learning from non-formal and informal contexts. For the second type of accreditation technological support is needed to approximate the prior knowledge of learners (Joosten-ten Brinke, Brand-Gruwel, Sluijsmans, & Jochems, 2008). This support may range from a form of pre-advising the experts which documents are relevant for the target course or study programme or it could help students fill their portfolios only with material that is relevant to possible exemption decisions.

For technology-enhanced learning APL procedures are related to approaches to model prior knowledge from learners with methods from adaptive hypermedia research like scalar models, overlay models, perturbation models or genetic models (Brusilovsky & Millan, 2007). But it is rather a closed world in which these models operate; systems can only take account of what they represent/know about the learner in the current learning environment. All "system experiences" are lost after changing the learning environment and cannot be reused in another system. This problem has been recognized as the ‘open corpus’ problem in adaptive hypermedia research (Brusilovsky & Henze, 2007).

While these models have their value in traditional e-learning processes in learning networks they are not applicable at all (Koper, Rusman, & Sloep, 2005). Alternative bottom-up approaches are needed to offer personalized learning paths which do not need extensive sets of metadata to reason about prior knowledge of learners. The basic assumption of our research is that prior knowledge of learners can be approximated by the content of the learner portfolio and therefore overlap between the documents in the portfolio and the courses of the plan/curriculum is a proxy to give exemptions and provide a personalized curriculum. This similarity calculation is done in our study with the help of a dimensionality reduction technique similar to Latent Semantic Analysis (LSA).

LSA for Prior Knowledge Approximation

Latent Semantic Analysis (LSA) is a theory and method for extracting and representing the contextual-usage meaning of words by statistical computations (Landauer, 2007). The latent semantic analysis process consists of several steps: In the *indexing phase* all words and documents construct a so called Term-Document-Matrix (TDM), in which all terms are listed in the columns and all documents in the rows. After the counting of the occurrence of each word in all documents several weighting and normalization options are possible. In the *dimensionality reduction phase* a mathematical function called singular-value decomposition (SVD) is applied which is similar to factor-analysis. The end result of this process is a latent semantic space, in which the input documents are represented as vectors. Documents in this space are similar if they contain words which appear in the same context and so their vectors are close together in the space providing a measurement for the similarity of text. In the *retrieval or query phase* the query text is projected into the space and the distance to the document vectors is calculated via the cosine or the Euclidian mean.

Latent Semantic Analysis has been applied to several problems in the domain of Technology-Enhanced Learning like peer tutoring (Van Rosmalen, 2008), provision of feedback (Graesser, Chipman, Haynes, & Olney, 2005), automated essay scoring (Foltz, Laham, & Landauer, 1999) and the selection of educational material (Graesser, Wiemer-Hastings, Wiemer-Hastings, Harter, & Person, 2000). Zampa & Lemaire (2002) applied LSA to the user modeling problem. In their model learners learn a domain by acquiring the most important “lexemes” of a domain. These most important concepts are identified beforehand and the end result is a recommendation the next best item to look at based on the theory of proximal development (Vygotskii & Cole, 1978). (Wolfe et al., 1998) focused on the use of LSA for the selection of educational material. Since we cannot elaborate on LSA more into detail in this chapter we recommend the papers mentioned above as an introduction to LSA.

Our application of LSA is similar to those presented here but differs in a number of important aspects. In contrast to the approach of Zampa & Lemaire (2002) we do not aim to model the student background and the learning resources beforehand. Since we use the text of students as a proxy to their prior knowledge we have a more dynamical model which could build on learning progress documented e.g. in an electronic portfolio. Our application of LSA also differs from a "simple" essay scoring scenario where a student text is compared to one ore more clearly defined gold standard texts which provide the basis for judging about the quality of an essay. In our case the topical range of the student documents can be potentially very broad while the target documents or course units can be very similar to each other. This makes it very important to find an ideal dimensionality which results in the best discrimination between target documents.

Study

In this chapter we present a study that has been conducted at a psychology course of the Open University of the Netherlands. The foundations for this study are described in (Van Bruggen et al., 2004), and the long term research agenda for it has been presented in (Kalz, Van Bruggen, Rusman, Giesbers, & Koper, 2007). The study employed a text vector space model and dimensionality reduction technique inspired by Latent Semantic Analysis (LSA) to calculate a similarity between documents in the learner portfolios and the content of the course units. To test the validity of the results an expert validation with domain experts is performed and the performance of using LSA as a classifier for relevant or irrelevant documents is evaluated with a Receiver Operating Characteristic also called ROC curve. ROC curves are a method from signal detection theory and they are often applied in binary classification experiments and model performance evaluation.
Our hypothesis tested with the study is the following:

H1: LSA can classify documents as relevant/irrelevant in a manner comparable to human experts.

Our target of this study is to minimize the false-negative and the false-positive cases under a threshold level of 10%. Too many false-negative cases would hinder students from exemptions while too many false-positive cases would results in unnecessary work for the domain experts. In addition we evaluate results we have reported in a prior study about the identification of an ideal number of dimension for small corpora (Kalz, Van Bruggen, Rusman, Giesbers, & Koper, 2009a).

It is important to mention that this study was not intended to evaluate the use of automated exemption rules for study programmes but is intended to evaluate the applicability of the text vector space model and dimensionality reduction techniques for learner placement in general. The exemption rules and accompanied problems of these (trust issues, validity of thresholds etc.) are not a specific problem of the method we are evaluating but of the general exemption procedure. This means that exemptions standards can differ between institutions and that they have to set their own thresholds for exemptions.

Method

Participants

In this chapter we report on a study in the context of an introductory psychology course of the Open University of the Netherlands. From the 244 students that enrolled in the course around 6% submitted material to apply for the accreditation of prior learning (14 students). To have a broader range of different content to be examined we asked several colleagues to submit additional material. Overall we had a total number of 18 participants providing 28 documents which could be compared to the units of the target course. Thus the total number of similarity ratings was 504 (28 documents x 18 course units).

Procedures

The introductory psychology course consisted of 18 units each of them dedicated to a subtopic of psychology. The course was offered in an online environment. Before students could enter the course they had to read an introduction about the content of the units of the course. Then the students filled out a questionnaire on any prior knowledge for the course or parts of the course. To substantiate claims on prior knowledge students were invited to submit materials they had produced in their prior education or working environments. Documents submitted were work reports, (bachelor / master) theses, technical reports, essays, reference lists and presentations.

We have chosen the ideal dimensionality for the study based on two performance criteria. On the one hand we have employed the connection between singular values and the variance they account for in the Term-Document matrices analyzed by LSA. Our target was to reduce the variance accounted for under a threshold of 70 – 80 % of variance represented in the data. We have successfully evaluated this method in a prior experiment (Kalz, Van Bruggen, Rusman, Giesbers, & Koper, 2009a.). Then we have used the corpus documents as queries to control the discrimination between the chapters. These queries were tested for self-correlations and discrimination to the other chapters under different conditions. We have varied the dimensions used between 5 and 1000 and we have used different stopword settings (no stopword list, 30% stopword, 50% stopword list). Stopwords are removed because they appear very seldom or very often in the corpus. In addition we have also evaluated the use of local and global weighting options. In this process we have followed a method by (Rosmalen et al., 2006) to calibrate and test several LSA parameters.

To evaluate the results of the study two domain experts independently rated the similarity between the student documents and the domain documents on a 5-point Likert-scale. For each student document the experts evaluated the semantic similarity and marked if they would give an exemption. In addition they wrote down how much time it took to review the material. We interviewed one of the domain experts qualitatively.

We have calculated a raw overall percentage of agreement between the two raters. We have calculated the interrater agreement according to the consensus and the constistency of the ratings by the two judges (Stemler, 2004). The consensus of the ratings by the two domain experts was calculated using Cohen's Kappa while we have used the Spearman rank coefficient to calculate the consistency among the ratings. For the performance assessment we have recoded the Likert scale into a binary scale. Here we have used the most optimistic rule that a document is seen as relevant when at least one of the raters has rated the document with a value higher then 1.

The model performance assessment of our method as a classifier for relevant and irrelevant documents for APL procedures was analyzed via a confusion matrix and ROC-curve (Fawcett, 2004). ROC-Curves (Receiver Operating Characteristics) are a method from signal detection theory that has been applied to evaluate model performance assessment of classification models. In ROC curves the true positive rate (tpr) of a classifier is plotted against the false positive rate (fpr) while varying the thresholds used for the classification. One of the main advantages of the application of ROC-curves is that they are not sensitive to class skew (Hamel, 2008). With this approach we have also compared the effect of applying different weighting functions reported in the literature to contribute positively to performance in small scale corpora (Nakov, Popova, & Mateev, 2001; Wild, Stahl, Stermsek, &

Neumann, 2005). We have compared the use of a logarithmic weighting function for a local weighting and the use of entropy and inverse document frequency for the global weighting and the combination of logarithmic and entropy weighting.

In addition we have calculated raw success scores on the document level (How many documents would have been recommended right?) and the person level (How many learners would have been exempted right if LSA would have decided about exemptions?).

The corpus for the experiment consisted of the content of the 18 psychology learning units. This corpus had 28165 terms with 490431 occurences. The corpus size was 3,1 MB. 13283 terms only appeared once in the corpus. After keeping only terms that appeared more than once the corpus was reduced to 14882 terms with 477147 occurences. All corpus and learner documents have been manually cut into paragraphs. The paragraph length was between 250 and 500 words. The corpus consisted in the end of 2246 paragraphs. In this study all analysis was done using the Text to Matrix Generator (TMG) which is a Matlab implementation of Latent Semantic Analysis and other techniques (Zeimpekis & Gallopoulos, 2006). For the experiment a script was written which calculates the mean of cosines for all paragraph to paragraph comparisons and writes down the mean correlation to all 18 chapters in a spreadsheet file.

Results

Dimensionality reduction and Sensitivity

For the estimation of the ideal dimensionality we could reproduce results obtained in a prior study (Kalz, Van Bruggen, Rusman, Giesbers, & Koper, 2009). Figure 3.13 shows our research corpus with different stopping strategies applied and with different numbers of singular values.

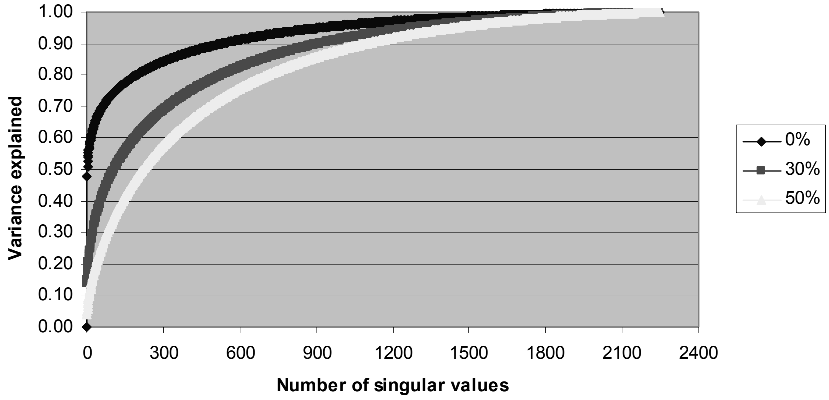

Figure 3.13: Variance explained for different numbers of singular values and 3 different stopping strategies

In this figure we can see that the variance zone of at least 80% variance accounted for starts with different numbers of singular values depending on the stopping strategy used. For no stopping this zone starts at a reduction to 194 singular values. For the 30% stopping strategy this zone starts at 530 singular values while it starts for the 50% stopping strategy at 713 singular values.

The second performance criterion was tested with the discrimination between the target learning units. Figures 3.14 – 3.16 illustrate the auto-correlation and the correlation to the other learning units for the learning unit 1. Note that we did not analyze the full range of singular values but only a selection since the figures are only used for illustrative purposes.

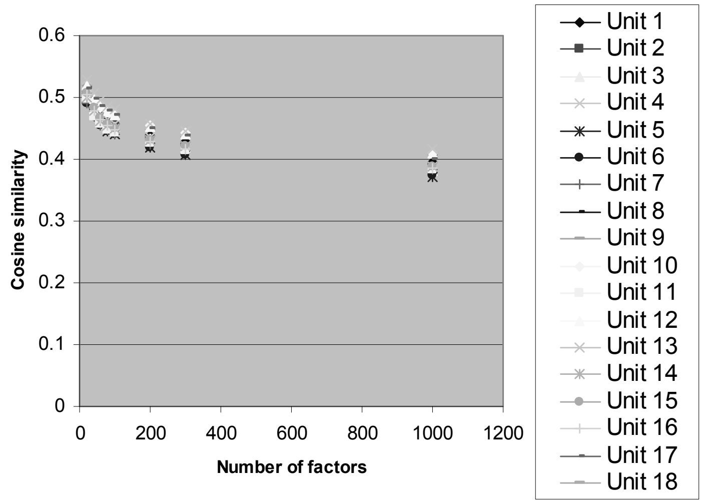

Figure 3.14: Correlations and auto-correlation for learning unit 1 with no stopping.

As we can see in figure 3.14 it is not possible to discriminate between unit 1 and the other learning units of the course. In figure 3.15 we have increased our stopword strategy to 30 %.

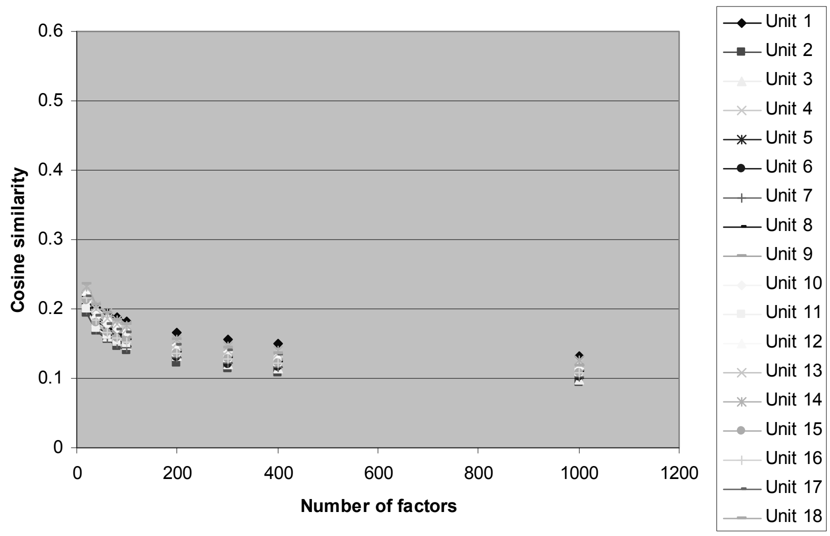

Figure 3.15: Correlations and auto-correlation for learning unit 1 with 30% stopping.

As we can see in figure 3.15 the cosine values drop and discrimination between the chapters improves with a 30% stopping strategy. But still it is not easy to clearly discriminate between the chapters. In figure 3.16 we increased our stopping strategy to 50% stopwords.

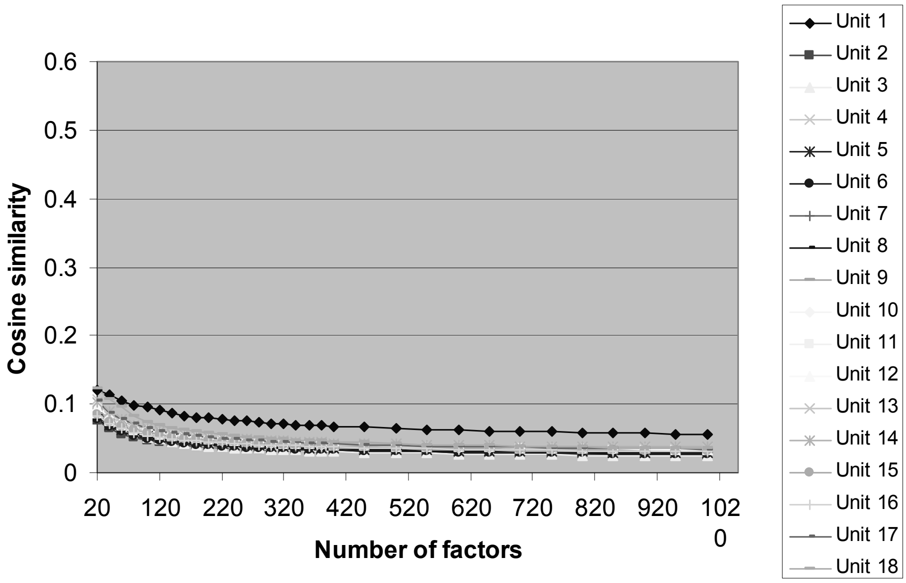

Figure 3.16: Correlations and self correlation for learning unit 1 with 50% stopping

Overall we can see that the cosine values drop the more stopwords we apply, but at the same time the discrimination between the target documents in the corpus improves sufficiently to discriminate between the chapters. This effect could be replicated for all 18 units. After having defined the ideal number of factors and stopwords for the analysis we have used these parameters (800 factors with a 50% stopword list) for querying the student documents and comparing the cosine values to the 18 learning units.

Model Performance Assessment

The raw agreement between the two domain experts was 95%. This high agreement between raters was mainly based on the high number of submissions rated as not-relevant for the APL procedure. The interrater reliability for the raters was found to be Kappa = 0.77 (p <.0.001). According to Spearmans rho there was a high consistency between the ratings (rho=0.75, p<0.01). Again these high statistics are a biased upwards because of the skewness of the data. After the recoding into a binary classification the rater agreement was Kappa = 0.74 (p<0.001). Only 8 % of all cases (40 cases) were evaluated as relevant for APL. Because of this the data were negatively skewed. The expert data confirmed our basic assumption that content similarity is related to exemptions in APL procedures. If we take into account only cases with content similarity higher 2 on the Likert scale then 83 % of the cases have been proposed for an exemption in the mean of both raters. Seven learners would have been given exemptions based on the decisions of experts which equals 38 % of all participants who submitted material for the study. In our evaluation of the impact of different weighting functions we have compared five different parameters shown in figure 3.17.

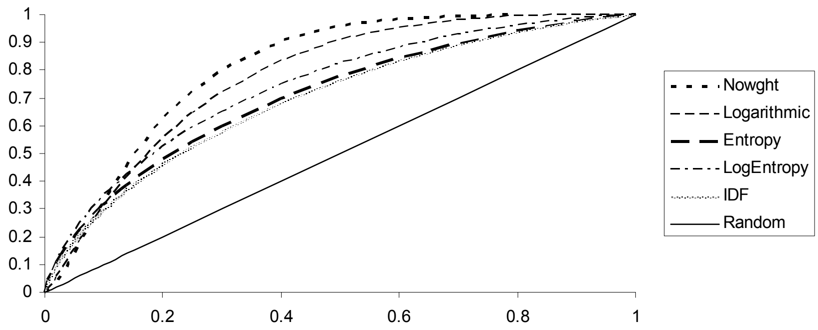

Figure 3.17: Receiver operating characteristics curve (ROC curve) for LSA with different weighting parameters

There are several important aspects to be reported. In general, we can report that without taking into account the different weighting functions our method can be clearly distinguished from a random classifier that would rate documents by chance into each of the categories. Overall the use of the weighting functions reported in the literature has slightly decreased the performance of our model and the best performance could be reached without any local and global weighting. To compare this performance more detailed we provide in table 3.2 an overview about the size of the Area und the ROC curve (AUC) which can be used to evaluate the performance of a classification model.

Table 3.2: Influence of weighting to performance measured with the size of the Area under the ROC-Curve and the standard error

	AUC	**Std. Error.**
No wght.	0.811	0.0262
Logarithmic	0.777	0.0313
Entropy	0.708	0.0413
Log./Entropy	0.741	0.0386
IDF	0.697	0.0417

We can see that the inverse document frequency function (idf) performs worst for our model and that the logarithmic function was the second best option. The bad performance of the global weighting methods was on the other hand not surprising since we have filtered the corpus already before we have applied the weighting functions. If we also take into account the confidence intervals of the ROC curves we have to summarize that the weighting functions decrease the performance but not significantly because the confidence intervals of all settings overlap.

We have decided to use the 800 dimensions, a 50% stopword list, no weighting and a cutoff point of 0.3 for separating between relevant and irrelevant documents in our model. For this settings human raters and LSA results show only a fair agreement of Kappa = 0.33 ($p < .0.001$). The mean of the ratings was 0.11 (Std. error = 0.21) while the mean variance was 0.23 and the mean standard deviation 0.47. To discuss the performance of our model more into detail we provide an overview of the results with a confusion matrix (table 3.3). We can see that LSA could successfully classify 88 % of all cases right. The false positive cases are 8 % of all cases while 2 % of the cases fall into the false negative category.

Table 3.3: Classification results Human ratings vs. LSA ratings (n=504)

		Human rating	
		Relevant	Irrelevant
LSA rating	Relevant	28	40
	Irrelevant	12	424

The *accuracy* of our classification model is 0.89. The *sensitivity* is 0.41 while the *specificity* is 0.91. The false positive rate is 0.08. Overall the *negative predictive value* of our model is 0.91 while the *false discovery rate* is 0.58. The Area Under the ROC-Curve (AUC) for our model was 0.81 (95% CI, Std Err. = 0.0262). Overall this is a sign of a good predictive classification model.

On the document level we have analyzed how many documents of the 36 would have been passed to the human experts if our LSA model would have been implemented to evaluate relevant documents. Here we have compared two different methods: The mean of all cosine values to each unit for every learner document and the maximum cosine value of the comparison between learner document and all 18 units. The raw percentage of right classified documents of the mean method was 85 % while it was only 43 % with the maximum method. On the level of the learner our model would have recognized 7 of 12 given course unit exemptions of the human raters but it would have added 40 false positive exemptions on top for these 7 learners.

Discussion

Based on a method that combines an approach of filtering, multiple criteria and variance-based selection we were able to achieve sufficient discrimination between the target units. With these results we confirm our findings from an earlier study using small scale corpora. After testing different weighting options we could show that additional weighting does not improve the performance of our model. Overall we could reach a satisfactory classification model because of two reasons. First we could reach our self-set target of less then 10% false positive and false negative cases. Second, an AUC value of 0.81 is seen as a good performance indicator for a classification model. The relative low sensitivity of the classification model is not too problematic since we still can reach a significant number of cases that would have not been analyzed by the domain experts.

But there are several limitations to the study presented here. First of all the low number of relevant documents is problematic in case of generalizability of the results. The documents we collected for an APL procedure resulted in negatively skewed data. After discussing these findings with an APL expert we could get a confirmation that students are often not aware what to submit for exemptions. This means that collecting more student data would likely lead to a similar skewed dataset. In fact, it is a part of the problem we are trying to solve with technology.

Some of the false negative cases reveal that human experts decide about exemptions on more factors then just the semantic similarity between learner documents and target documents. Especially one case was interesting in this regard: One of the learners submitted a very detailed description about an experiment conducted. This description did not contain sufficient semantic concepts which had a relation to the target documents. But human raters deducted from the document that this learner must have specific prior knowledge in psychology to be able to write such a document. For this purpose other techniques and approaches to approximate prior knowledge are needed which go beyond semantic similarity.

The qualitative interview showed us that the analysis model of the human experts is based on semantic similarity of documents but the cognitive process of the experts is more complicated. One domain expert described the analysis process with different steps involving keyword analysis, semantic analysis and quality ranking of the student documents.

While our study confirms that the application of dimensionality reduction techniques like LSA for the support of APL procedures and the approximation of prior knowledge is a promising research and development direction the approach needs to be validated in several different contexts. In this regard we expect that a training phase is needed in all implementations to align the approach to local thresholds and local decision boundaries.

Section 4: Technological Artifacts

Section Introduction

In this section we present the technological artifacts developed within the project. In chapter 4.1 we present the technological foundations and embedding of placement support in the TENCompetence infrastructure. We present its relation to the navigation service and discuss the importance of wayfinding approaches to educational scenarios which make extensive use of open educational resources.

In chapter 4.2 we present the placement web service prototype that we have proposed in the theoretical foundations. We present the architecture of the web-service and discuss results of a technical evaluation. A perspective for future development is provided.

In chapter 4.3 we present a framework that extends the approach of the web-service. The Semantic Weblog Monitoring Framework (SWeMoF) addresses two problems we had to tackle within this study. On the one hand SWeMoF allows the construction of corpora on the basis of sources from the social web. This should enable us to construct larger corpora and domain specific corpora more easily. In addition this framework enables us to extent our focus on other data formats and other methods that we have defined in the long-term research agenda for placement support.

Chapter 4.1: Wayfinding Services for Open Educational Practices[7]

[7] Kalz, M., Drachsler, H., Van Bruggen, J., Hummel, H. G. K., & Koper, R. (2008). Wayfinding Services for Open Educational Practices. International Journal of Emerging Technologies in Learning, 3(2), 24-28.

[8] Kalz, M., Drachsler, H., Van Bruggen, J., Hummel, H. G. K., & Koper, R. (2008). Wayfinding Services for Open Educational Practices. International Journal of Emerging Technologies in Learning, 3(2), 24-28.

Introduction

The role of content for technology-enhanced learning has been attributed a lower priority from several recently stressed theories and models of learning like socio-constructivist theories (Duffy & Jonassen, 1992), situated cognition (Lave & Wenger, 1991) or constructionism (Gergen, 1999). In addition the role of user-generated content is currently intensively discussed with its impact and importance on learning and competence development (NMC, 2007). Nonetheless, learning content is still an important factor for technology-enhanced learning and a huge amount of learning content is critical to allow a wide-scale diffusion of self-directed lifelong-learning for the individual.

In the past several initiatives have been started to offer learning resources on a wide-scale on the internet for free. One of the first and most successful initiatives was the Opencourseware project from the Massachusetts Institute of Technology in which the content from 1550 courses has been made publicly available. Several initiatives followed and initiated a new discussion about openness and access to learning resources in education. The UNESCO summarizes these new development 2002 in their Forum on the Impact of Open Courseware for Higher Education in Developing Countries as "the open provision of educational resources, enabled by information and communication technologies, for consultation, use and adaptation by a community of users for non-commercial purposes" (UNESCO, 2002).

Although the availability of open educational resources is currently increasing there is a lack between the mere availability of these resources and the educational use of the available material. To fill this gap it is important "how educational repositories of Open Educational Resources, which often want to grow based on user contributions and sharing among users, will manage to become more useful for communities of practice" (Geser, 2007). The increase of open access and the publication of open educational resources do not imply the creative use of these resources for learning.

This chapter presents two services that are currently developed in the framework of the European Integrated project TENCompetence. These services deal with a similar problem that users of open educational resources are facing. The next part of the chapter deals with user requirements and existing technological solutions to improve the competence and learning related search and use of open educational resources. Then we present positioning and navigation as two independent but connected services that can help learners to decide about which learning activity or resource to choose as next step in their personal competence development. Finally, we discuss the (dis)advantages and give an outlook about future research.

Supporting learners and learning designers to find suitable OER

The main users for open educational resources are self-directed learners on the one hand and learning designers on the other hand. Both groups have specific requirements for using open educational resources for learning or developing learning opportunities. First of all, learners need to have orientation to choose the best suited learning activities from the vast amount of available resources. The label "best-suited" implies several options regarding the choice of material. In general the best suited resources for learners are the ones that help them to reach the "zone of proximal development" (Vygotskii & Cole, 1978) regarding their competence development goals. This zone can be identified through an analysis of the learners' prior knowledge, his topical interest and/or a comparison to the next steps similar learners have taken.

For learning designers it is important to know which resources can be combined to produce a sound competence development program for learners. Again, there are several aspects for learning designers to decide about the appropriateness of open educational resources for constructing learning activities and courses. They have to address their target group based on specific characteristics like prior knowledge, learning goal, study time or preferred study style. Therefore, the most important problem for both user groups is an individualized search facility tailored to their needs and competence development targets. This search-and-find problem can be addressed on several levels: On the level of the learning objects, on the level of the technology for storing the objects (the repository level) and on the user level:

Learning Objects Level: To unify the description of learning resources the IEEE LOM standard has been used in many repositories to describe the contained resources (LTSC, 2002). But the IEEE LOM standard has been criticized because of its limited possibilities to enrich learning objects with educational meaningful information (Foroughi, 2003). In addition, research has shown that it is not recommendable to let authors enrich learning objects with metadata because this does not lead to sufficient quality of the metadata (Barton, Currier, & Hey, 2003). To ensure high qualitative metadata domain experts are needed who tag the resources with an agreed upon taxonomy of keywords. As an alternative to IEEE LOM several repositories use an extended set of the Dublin Core Standard (Mason & Sutton, 2005). This extended set offers more flexibility to enrich learning resources with educational and competence development related information but in essence the expert problem remains.

Repository Level: On the level of the learning object repositories the Open Archive Initiative Protocol for Metadata Harvesting (OAI PMH) was a first step towards pooling of resources from different origins (OAI, 2005). Currently, the work on the Simple Query Interface (SQI) has enabled interoperability for search between multiple repositories (Van Assche et al., 2006). This intra-repository specification allows users to find learning objects in several distributed repositories.

But solutions on the level of the object and the repository do not address the problem that the success of the search is still dependent to a large extent on the quality of the metadata attached to the learning objects. Additional approaches on the user layer are needed to address the above described problem description for users of open educational resources. Especially when the search functions should address and support competence development the contribution of users is needed to add this information to the learning resources in open content repositories. These distributed resources are not only used in “formal learning environments” but also used by distributed self-directed learners who use them in a more informal way. Therefore, we imagine a mixed use of different OER repositories with additional user driven contents in future. Such an emerging OER environment requires highly automated services to support the learners and limit their level of maintenance as far as possible.

In the European Integrated Project TENCompetence we are currently researching ways to personalize distributed learning resources, units of learning and competence development programs in Learning Networks (LN) (Koper & Specht, 2008). In LNs learners, team, or institution are free to add any kind of content they want to. Thus, LNs could consist of different mixed OER, formal learning offers, or separated learner contributions at the same time.

Two wayfinding services on the user level are responsible to offer individualized competence development programs in LNs. Wayfinding services should tell learners where they currently stand in a chosen competence development program and which way they should choose to foster competence building in a learning network. The placement support service analyzes the prior learning of learners through a content analysis method while the navigation service recommends the next best learning activity for a student. In this contribution we introduce these services and discuss their potential as a bridge from distributed open educational resources to open educational practices.

TENCompetence wayfinding approach

To provide learners with orientation for their competence development our research focus is to answer the following questions: Where do I stand in the "curriculum" and which step should be next to reach my personal learning goal? To answer this question we have conducted research how to support this orientation process with technology. Two services haven been recently implemented and tested in the framework of the TENCompetence project: A placement support service uses language technology for prior learning assessment while a navigation service applies research from recommender systems to help learners to find orientation.

Placement

Especially from the perspective on lifelong learning it is an important question for learners where they should start their competence development on the basis of what they already know and what they want to achieve as competence development goal. In traditional educational settings this problem is addressed through the Accreditation of prior learning (APL) (Merrifield, McIntyre, & Osaigbov, 2000). This process – which is most of the times carried out the submission phase of study programs in Higher Education – relies on domain experts who study the portfolios of learners and decide about exemptions for them. Due to the fact that this is impossible to follow this approach when people change their domains and institutions quite often during their life, we use language technology to support this process.

Our current project uses Latent Semantic Analysis (LSA) to support the APL process for technology enhanced learning. Latent Semantic Analysis (LSA), in the past sometimes referred to as Latent Semantic Indexing (LSI), is used to calculate a similarity between the learner documents and the learning resources. LSA is a theory and method for extracting and representing the contextual-usage meaning of words by statistical computations (Landauer, Foltz, & Laham, 1998). The whole process of this analysis consists of several steps like the pre-processing of the text, some weighting and normalizing mechanisms, the construction of a term-document matrix and a mathematical function called singular-value decomposition (SVD), which is similar to factor-analysis. The end result of this process is a latent semantic space, in which the main concepts (or types) of the input are represented as vectors. Concepts in this space are similar if they appeared in the same context and so their vectors are close together in the space providing a measurement for the similarity of text. LSA is successfully applied in several research fields like informatics, psychology or medicine.

LSA has been used in an educational environment for assessment and feedback of free text in intelligent tutoring systems. Some examples of these applications are the Intelligent Essay Assessor (Foltz, Laham, & Landauer, 1999), Summary Street (Steinhart, 2001) and Select-a-Kibitzer (Wiemer-Hastings & Graesser, 2000) to mention only a few. Wolfe et al. (1998) and Dessus (2004) have used LSA to provide students with appropriate texts that fit to their current knowledge
Our application of LSA is similar but it is applied in a different way and on the basis of a different motivation. We are using LSA to assess prior knowledge of learners for placement or positioning decisions and finally the construction of personalized learning paths through a learning network. A high correlation between documents in the portfolio and learning resources leads to an exemption of this specific learning activity. The result of these analyses should be taken into account for the creation of a personalized learning path. Some learning activities on the way to the target competence a learner wants to achieve may be exempted because of the results of this prior learning analysis. We conducted an expert validation of the positioning service and compared the results of LSA to results that experts have given. The first results of the service look promising.

Navigation

Navigational support is necessary for providing learners with appropriate learning resources when there is not a clear curriculum. We have recently designed a navigation service as a personal recommender system (PRS) for learning resources. The general concept of the PRS is in line with hybrid recommender systems in other domains. Hybrid recommender systems combine different kind of recommendation techniques to achieve a higher accuracy in their recommendation (Good et al., 1999; Melville, Mooney, & Nagarajan, 2002; Soboro & Nicholas, 1999). Every single recommendation technique has its own advantages and disadvantages. Thus, a combination of techniques in a hybrid recommender system can balance the advantages and disadvantages of single techniques through switching or cascading them in a recommendation strategy (Drachsler, Hummel, & Koper, 2007) to increase the overall accuracy of the recommendations. Figure 4.1 shows an example of a possible recommendation strategy that combines a stereotype filtering technique with an ontology based recommendation.

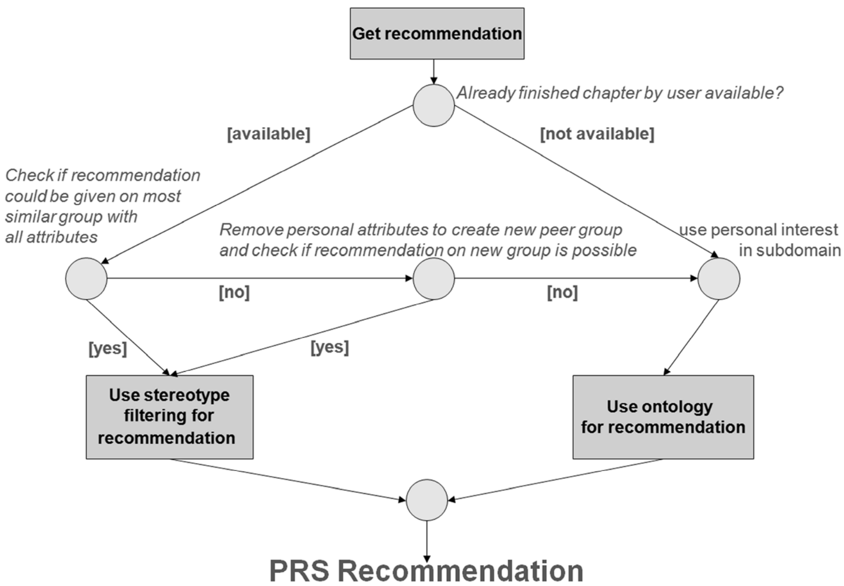

Figure 4.1: Combination of stereotype filtering and item based filtering in a recommendation strategy

Recommendation strategies decide which specific recommendation technique provides the highest accuracy for the current user based on specific history information about users and items.

The following recommendation techniques are promising for recommendations in OER in order to have less maintenance and highest guidance for the learners 1: User-based collaborative filtering 2. Stereotype filtering and 3. Item-based filtering.

1. Collaborative filtering techniques (or social-based approaches) use the collective behaviour of all learners or learning resources. Both user- and item-based collaborative filtering use the same mechanism of correlation for different objects. To underline the differences between these two techniques we now describe them together. User-based filtering correlates users by mining their (shared) ratings, and then recommend new items that were preferred by similar users (see Figure 4.2). Similar users are calculated through shared preferences for learning resources.

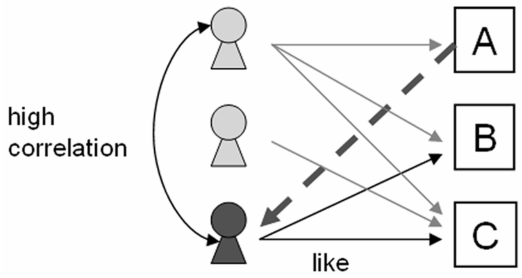

Figure 4.2: User-based collaborative filtering

2. Item-based techniques correlate items by mining (similar) ratings they own, and then recommend new, unknown but similar items to the learners (see Figure 4.3).

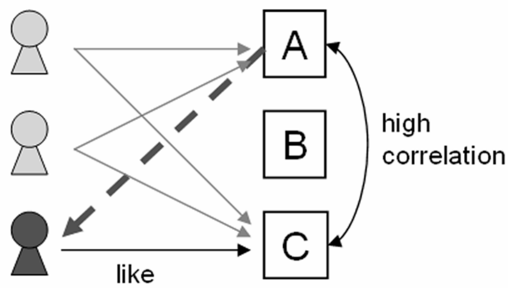

Figure 4.3: Item-based collaborative filtering

Main advantages of both techniques are that they use information provided bottom-up by user rating, that they are domain independent and require no content analysis, and that the quality of the recommendation increases over time (Herlocker, Konstan, Terveen, & Riedl, 2004). User- and item-based techniques are useful for networks which are dealing with different topics like OER. They do not have to be adjusted for specific topics, which is important because we expect many LN for different topics. CF techniques can identify learning resources with high quality, allow learners to benefit from experiences of other, successful learners. CF techniques can be based on pedagogic rules that are part of the recommendation strategy. Characteristics of the current learner could be taken into account to allocate learners to groups (e.g., based on similar ratings) and to identify most suitable learning activities. For instance, suitable learning activities can be filtered by the entrance level that is required to study the learning activity. The prior knowledge level of the current learner would than be taken into account to identify the most suitable learning activity.

A disadvantage of both techniques is the so called 'cold-start' problem. It is due to the fact that CF techniques depend on sufficient historical user behavior data. Even when such techniques have been running for a while, adding new users or new items will suffer this problem. New users first will have to give a sufficient amount of ratings to items in order to get accurate recommendations based on user-based CF (new user problem). New items have to be used or rated from a adequate amount of users to be recommended (new item problem). To solve these disadvantages, user- and item-based CF have to be combined with other techniques, like Stereotype filtering, in a recommendation strategy.

3. Stereotype filtering techniques can recommended preferred items to similar users based on their mutual attributes (see Figure 4.4). They are also domain independent, but they do not require that much history data to provide recommendations. Therefore, it is useful to solve the 'cold-start' problem of User- and Item-based techniques. Additionally, stereotype filtering offers exploration through 'serendipity' what is less the case for user- or item-based filtering and an advantage for a combined recommendation strategy.

Figure 4.4: Stereotype filtering

The stereotype recommendation technique is an accurate way to allocate learners to groups if no behavior data is available. In combination with techniques that suffer from the 'cold-start' problem, stereotypes complement a recommendation strategy, enabling valuable recommendations from the very beginning.

Beside the general issue of selecting the most appropriated techniques for an environment, the domain of learning demands to be addressed by specific characteristics. For PRS in lifelong learning context it is not possible to simply take or adjust an existing PRS for consumer products (like in amazon.com). PRS for lifelong learning should support the efficient use of available resources to improve the competence development, taking into account the specific characteristics of learning. PRS have to be driven by pedagogical rules, which could be part of the recommendation strategy. The recommendation strategy looks for available data to decide on which technique(s) to select for which situation. The same situation is given when users are dealing with open educational resources for their personal competence development.

Our approach of navigation support was evaluated in a psychology pilot implementing the ISIS recommender system. This recommender system used an attribute-based recommendation technique in combination with an ontology based recommendation technique in a hybrid recommender system.

Services for open educational practices

In the TENCompetence project we are developing a Personal Competence Manager (PCM) that will support individuals in building competences. One important feature of this application is the underlying theoretical approach of LNs. LNs should enable learners to develop their competences together with peer-learners who have a similar competence development goal. Learners in LNs are able to develop their own learning paths including the use of openly available resources and learning activities. Since all users in the TENCompetence environment are able to share their learning paths and learning activities and resources they have used for competence building the environment should enable users to collect competence related information about open educational resources and educational/contextual metadata.

In addition, the above described placement support service can give a valuable contribution to help learners finding their way through open educational resources. Instead of using this service only for exemptions the similarity rate between a learner's portfolio and documents in repositories can provide an individual "interestingness factor" for open educational resources. A high correlation between these resources and a portfolio can show that the learner already knows most of the concepts represented in these resources while a very low correlation would mean that these resources are completely out of the learners' context. While the positioning service takes only into account individual information of learners the navigation services uses also information by other learners to provide a recommendation.

For OER several situations would lead to different recommendation strategies. Without user information only a recommendation strategy based on topics or ontologies is possible. If a network of learners who use the content from several distributed repositories can be established the user behavior can be taken into account to recommend best suited learning resources. Since the learner groups of open educational resources are already available there is a need for a technology that connects these distributed learners and helps them with their competence development.

Discussion and Outlook

This chapter has introduced placement and navigation as two wayfinding services that haven been developed in the framework of the TENCompetence project. We believe that the combination of prior knowledge analysis and a personal recommender system has a high potential to bridge the gap between the distributed resources and distributed self-directed learners who have the burden to choose suited learning activities and resources. Both services have been recently analyzed in user studies and first results of these studies are promising. But this approach has also some issues: The use of Latent Semantic Analysis is limited to highly textual domains. In addition LSA can only find a similarity when the concepts used by the learners are represented in the semantic space. But there are several special presentation types (forms, descriptions of experimental designs etc.) that show an inherent higher prior learning than the purely textual content can show. In this case domain experts can deduct this but LSA cannot. For the navigation service the use social based approaches like collaborative filtering techniques is limited by a number of disadvantages. New users first will have to give a sufficient amount of ratings to items in order to get accurate recommendations based on user-based CF (new user problem). New items have to be rated from a sufficient amount of users to be recommended (new item problem). Another disadvantage for CF techniques is the sparsity of past user actions in a network. Since these techniques are dealing with community driven information, they support popular taste stronger than unpopular. Learners with unusual taste may get less qualitative recommendations, and others are unlikely to be recommended unpopular items (of high quality).

In the future we will implement and test our services in several settings where open educational resources are used for competence development, specifically in the OpenER project of the Open University of the Netherlands where learning resources from the Open University are published under an open content license. We are planning to setup a pilot with OER in a community space. Additionally, we will integrated the positioning and the navigation service to evaluated the support of these services for lifelong learners. The positioning service is planned to be used every time a new learner enters the community or a learner decided to study for a new learning goal. The placement support service uses the profile and historical information of learners to determine which learning resources are beneficial and which are not beneficial for their competence development.

The navigation service uses the outcome this calculation to recommend the most suitable learning resource to the learners. Therefore, it is based on a recommendation strategy that combines different kind of bottom-up collaborative filtering techniques in order to priorities the available learning resources for the individual competence development.

Chapter 4.2: A Placement Web Service for Lifelong Learners[9]

[9] Kalz, M., Drachsler, H., Van der Vegt, W., Van Bruggen, J., Glahn, C., & Koper, R. (2009c). A Placement Web-Service for Lifelong Learners. In K. Tochtermann & H. Maurer (Eds.), Proceedings of the 9th International Conference on Knowledge Management and Knowledge Technologies (pp. 289-298). September, 2-4, 2009, Graz, Austria: Verlag der Technischen Universität Graz.

Introduction

Aside from traditional educational offerings by institutional providers of technology-enhanced learning more and more focus is given to learning scenarios outside these institutions. Especially in professional education the importance of learning networks (Koper, Rusman, & Sloep, 2005) for non-formal education is rising (Kamtsiou et al., 2008). This is, besides other factors, also provoked through the use of social software for self-directed learning and competence development (Klamma et al., 2007). In these learning networks people get together who share one or more competence development goals. Several providers can contribute learning activities to this network. Based on these contributed resources, learning activities and the behaviour of the participants emerging effects can be used as the basis for several service offers.

In the TENCompetence project one focus is the development of models and tools which help to bridge the worlds of formal education and non-formal approaches all subsumed under the concept of competence development programmes (Koper & Specht, 2008). In the TENCompetence Integrated Project two independent but intertwined Web services have been developed to provide learners with orientation and feedback in learning networks. The placement Web service provides learners with a starting position in learning networks while the navigation web-service (Kalz et al., 2008a) leads the learners through the learning network based on the position provided by the placement service. These services have been integrated in one recommender system for learning resources which can use and reason with information from several sources and backgrounds. While data from formal providers of competence development programs can most of the time offer well structured information about the programmes and e.g. dependencies between learning activities in non-formal learning scenarios this can not be assumed.

For this purpose the hybrid personalizer has been developed (Herder & Kärger, 2008). The hybrid personalizer tries to overcome the problems of both learning contexts (formal & non-formal) through a collection of "top-down" and "bottom-up' Web Services which contribute to a recommendation of learning activities. The recommendation is computed based on information available about the learner, the learning activities, and the behaviour of other (successful) learners. The recommendation provided by the system is two-fold: on the one hand, it ranks learning activities based on how close they are to the learner's current knowledge level; that is, learning activities that are still too advanced are scheduled at a later point in the initial recommended visualization of the curriculum (or rather, competence development program).

On the other hand, as a second dimension, the system ranks learning activities based on to what extent they match the learners' preferences—as explicitly indicated in their profiles and as estimated from the behaviour of similar users. The placement service was used here mainly to compute the similarity between the content of the learner's portfolio and the content attached to the learning activities from non-formal sources without an extensive set of metadata. To test the reliability and performance of the placement web-service we have conducted a technical evaluation for further development.

In this chapter we focus on the placement web-service for lifelong learning. Based on the discussion of learners needs in non-formal professional learning we introduce the problem of placement in learning networks. The architecture of the placement web-service is discussed and a first technical evaluation and empirical results are presented. Several limitations are discussed and a research and development outlook is provided.

Learner needs in informal technology enhanced learning

Lifelong learners which make use of Technology-Enhanced-Learning have different needs and different expectations of technology-support for learning then learners in formal educational settings. Both learners share the wish to get personalized educational offers which fit to their prior knowledge and their learning goal. Most approaches from the traditional Adaptive Hypermedia (AH) literature assume for this problem the existence of a top-down-design for learning and a clear structure with pre-requisites defined and adaptation based on these predefined structures. In research on adaptive sequencing of learning resources for example a domain expert is assumed that models the learning goals and the domain concept ontology (Karampiperis & Sampson, 2005).

In non-formal education these assumptions often do not hold, because the sources and routes through the learning networks are not clearly defined and structured. This is also recognized in recent literature from adaptive hypermedia and referred to as the “open corpus problem” (Brusilovsky & Henze, 2007).

While these traditional models calculate the prior knowledge of learners only internal in their systems through a logic which is based on which pages a learner has visited inside the system and which pages contain similar content, for learning networks for non-formal professional development other techniques to take into account prior knowledge of learners are needed due to the open and dynamic structure of these networks. One approach to solve this problem is to ask contributors of learning networks to define the related learning goals to every learning activity added to the network and retrieve goal-related learning activities (Lindstaedt et al., 2008). While this approach may work well in more hierarchically organized learning and working environments its success depends to a very large extend on the willingness and ability of contributors to add these information. In learning networks other techniques have to be applied which can handle the dynamic nature of the learner profiles and the learning content in the network as well. Because of this dynamic nature we have chose an approach which takes the similarity between the content in the learning network and the content in the learner portfolio as a proxy for prior knowledge analysis (Kalz et al., 2008b). This is what we refer to as the “placement problem” in learning networks which we discuss in the next section in detail.

In a nutshell the placement problems boils down to the question which learning activities a learner should take in the learning network and which can be omitted taking into account his prior knowledge. In traditional formal education this problem is mirrored in a process called Accreditation of Prior Learning (APL). In this process learners apply with their portfolio for a study programme. Domain experts study these portfolios and decide based on the submitted material which parts of the study programme can be exempted and which personalized curriculum the learner gets. This process is time-consuming and expensive so a need is to support this process with technology.
In the UK this process has been recently covered in a JISC funded project which focused on the development of services for students and assessor in this procedure. An “estimator” service should serve in the pre-entry phase of the APL process to estimate the “importance” of a specific part of the

learner portfolio for a potential claim (Haldane, Meijer, Newman, & Wallace, 2007). In the context of the TENCompetence project we have evaluated and developed a solution for the same purpose which is based on the extensive use of Language Technology. In the next part we describe the architectural design of this application and present the context in which this service is embedded.

Placement Web-Service Prototype

Latent Semantic Analysis

The core technology of the Placement Web Service Prototype (PWSP) is Latent Semantic Analysis (LSA) (Landauer, 2007). LSA is a method for extracting and representing the contextual-usage meaning of words by statistical computations. The whole process of this analysis consists of several steps like the pre-processing of the text involving weighting and normalizing mechanisms, the construction of a term-document matrix, a mathematical function called singular-value decomposition (SVD) similar to factor-analysis, the rank reduction of the Term Document matrix and finally the projection of a query vector into the latent semantic space. In this latent semantic space the main concepts of the input are represented as vectors. Concepts or documents containing these concepts in this space are similar if their vector representations are close together in the space which provides a measurement for the similarity of text and concepts. The power of LSA in comparison to classical keyword techniques is the ability to match similarity even if documents do not contain any joint terms. LSA has been applied in several fields like medicine, psychology or computer science and it has shown good performance in essay scoring (Foltz, Laham, & Landauer, 1999) and tutoring systems (Graesser, Chipman, Haynes, & Olney, 2005).

Based on these experiences we have evaluated the use of LSA for approximating prior knowledge of learners based on their writings and results of this evaluation are promising. For the TENCompetence project a web-service prototype has been developed whose architecture is explained next.

Architecture of the Placement Web-Service Prototype

The Web-Service is built on top of an LSA implementation in PHP. The PHP implementation offers compiled methods for all basic LSA steps from importing documents, through cleaning up, building the Term-Document-Matrix (TDM) matrix, decomposing and reducing it and performing queries. This significantly reduces the code complexity without removing the option to implement custom scripted steps in between. It is possible to switch between a built-in and (two) external decomposition applications: As built in method we have used the ALGLIB library (Bochkanov & Bystritsky, 2008), while users can use General Text Parser (Giles, Wo, & Berry, 2003) or SvdLibC (Rohde, 2005) as well if they request them from the developers. No matter which engine is used, the output is always formatted in Harwell-Boeing format, making it easy to compare the various engines and to switch between them. The PHP implementation also allows various locations for input documents like ftp, (local) disk and nntp (usenet). Other document sources can be implemented easily.

It also has a wide range of textual cleaning functions built-in. The implementation can work directly with documents in txt-format or Word, PowerPoint and PDF. It also allows these type of documents as query. Because it saves all intermediate output in an easy readable non compressed format if the matrices are small enough, the implementation is suitable for research purposes. Figurer 4.5 shows a high level overview about the web-service prototype.

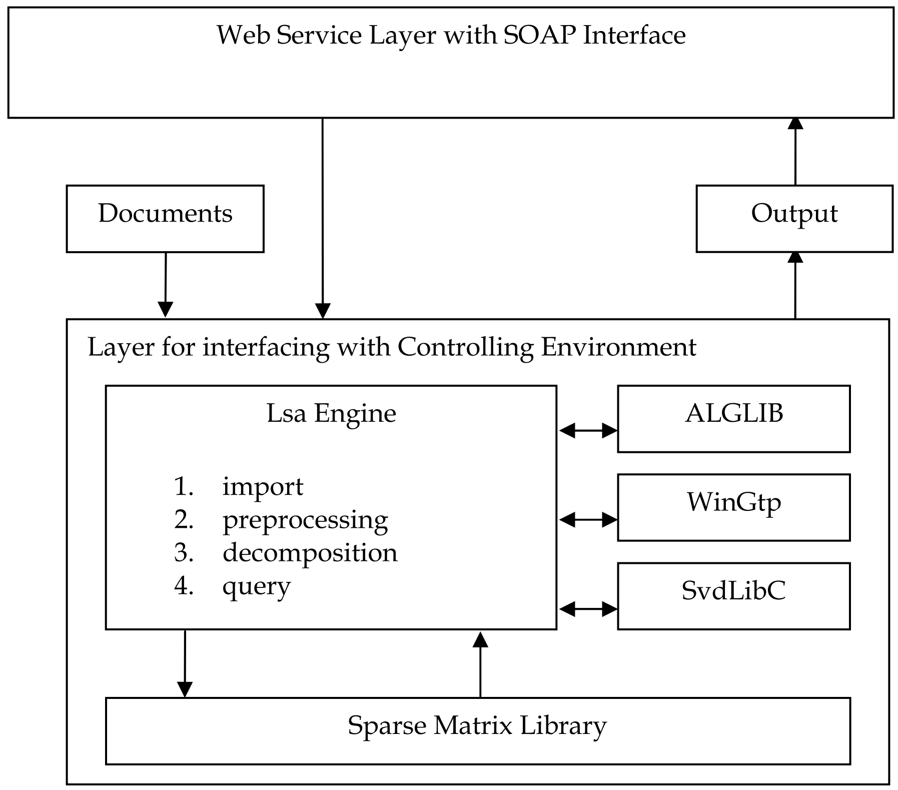

Figure 4.5: Architecture of the PWSP

On top of the Sparse Matrix Library the Lsa engine is built that performs the various actions necessary for Latent Semantic Analysis. These steps are:

1. importing documents
2. extraction of plain text from text files, Microsoft Word, Microsoft PowerPoint and Adobe Acrobat documents
3. pre-processing the imported text
4. constructing the Term x Document matrix (TDM)
5. performing the singular values decomposition and transforming the output of the decomposition
6. reducing the rank of the Term x Document matrix (TDM)
7. building a query vector
8. projecting a query vector onto the reduced Term x Document matrix
9. performing the query by calculating distances, dot products and cosines between the query vector and the reduced Term x Document space
10. allow retrieval of results

The Web-Service layer on top of the implementation offers a SOAP interface which can be called from external applications. As a result the web-service delivers a list of annotated documents with their similarity scores attached. Table 4.1 shows an overview of the API of the placement web-service prototype.

Table 4.1: Placement Web-Service Interface (API)

Placement Web-Service Interface (API)				
Name	**Method**	**Description**	**Input**	**Output**
getPositionValues	Get	Return a list of UoL annotated with cosine values	Iduser=xx	2-dimensional array of floats. Each UoL with its calculated cosine values.
Frequency		DATA Fields		Format
On request.	EventType	Iduser = Integer Learninggoal = Array (Strings)		

In the next part of the chapter we discuss the evaluation of the web-service prototype.

Evaluation

Performance Evaluation

In two prior studies we have evaluated the performance of using LSA for prior knowledge approximation. Since LSA was used in the past on the basis of large corpora from information retrieval we first had to evaluate its applicability on much smaller corpora which we expect to be available in learning networks. In this first study we could show the importance of applying stopwords and we have proposed a method to estimate the number of factors to be retained in LSA for small corpora (Kalz, Van Bruggen, Rusman, Giesbers, & Koper, 2009). In a second empirical study we have collected real data from students and compared LSA results with results from two domain experts who analyzed the documents according to semantic similarity and prior knowledge of students (Kalz et al., 2009a). In this study we could show that we can optimize the LSA procedure by applying the stopword and dimensionality method from the former study and we could reach a false negative and false positive rate under 10%. In this regard the correctness of the web-service results was satisfactory.

Technical Evaluation

There are several approaches how to evaluate semantic web-services and frameworks from a technical perspective. While we cannot discuss here extensively the several approaches for evaluation of semantic web-services we would like to present only the categories in which these web-services and their frameworks are evaluated. Küster, König-Ries, Petrie, & Klusch (2008) present the following aspects for such an evaluation: Performance & Scalability, Usability & Effort, Correctness, Decoupling and Scope & Automation. In Table 4.2 we show the evaluation results from a first evaluation of the PSWP.

Table 4.2: First Evaluation Results of the PWSP

Category	Result
Performance & Scalability	30 seconds for a dataset of around 800 documents (total size 7 MB).
Usability & Effort	Usability is high. Currently low effort, but might increase in future versions.
Correctness	Has been evaluated in an empirical study, but we expect that a "training phase" is needed for every implementation.
Decoupling	Not tested/Does not apply
Scope & Automation	Not tested/Does not apply

We have tested the performance and stability of the web-service with the same dataset which we have used for the empirical evaluation of LSA for placement purposes. For the dataset which had a size of 7MB in total the service needed around 30 seconds to respond and deliver the annotated list of learning activities (Core 2 Duo T7200 (2 GHz) PC with 1 GB of RAM). All documents used in the evaluation were text files so that no conversion into text format was needed. Scalability has been tested with datasets of different size and the time needed to get the result list increased on the one hand with the number of documents that have to be converted into textformat and on the other hand with the number and size of documents.

The usability of the current prototype is high, since we have coupled it with a WAMP (Windows, Apache, MySQL, PhP) package and so the setup time for the web-service is low. We could not test the aspects of decoupling and scope/automation since these aspect fit better to large frameworks than to a single web-service. In the next part of the chapter we discuss the results of the evaluation and provide a perspective for further development of the web-service.

Conclusions and Outlook

In this contribution we have introduced a placement web-service which has been developed in the context of the TENCompetence Integrated project. While we think that the work presented here is an important contribution for formal and non-formal contexts in technology-enhanced learning, there are several limitations of our approach. Since the services focuses on word usage in texts it can only work in domains where textual expression is important. Another limitation of using LSA for placement support is the inability to recognize more structured information like metadata. For this purpose other techniques have to be applied to allow for a matching like e.g. case-based reasoning approaches. The result of the first technical evaluation was that the web service performance was not suited for a productive system. This result was mainly due to scalability and effort issues. Response times higher than 30 seconds could be observed. Although this long reaction time depends to a large extent on the hardware of the Web Service machine we also analyzed improvements of this response time in the architecture. An important architectural aspect which leads to these response times was the combination of updating the Term-Document-Matrix (TDM) and singular value decomposition every time before a query is executed. Since an update of the matrix is only needed when sufficient new content has bee added to the learning network or the portfolio of the learner this process does not have to be executed before every query. For further development of the web-service we will work on the performance issues and we will split the process of updating the TDM and execution of the query. This is a well know problem in LSA research (Zha & Simon, 1999) and several solutions for this problem have been proposed which will be explored in the future.

Another important aspect of future work is related to the effort to set up the web-service to work correctly. Although we could evaluate a method to improve the web-service results sufficiently on a data set from psychology studies a new “training phase” will be needed for implementations in other context. For the further development of the web-service a rating system should be introduced that aligns the behaviour to the ratings of web-service users. For this purpose an adjustment layer should be implemented that can “learn” from the ratings of the results by users.

While we have focused in this project only on the application of Latent Semantic Analysis for prior knowledge approximation we combined text and data mining approaches in another recent project. The Semantic Weblog Monitoring Framework (SWeMoF) will enable users to conduct similarity analysis, classification and clustering experiments on corpora generated via RSS-feeds from weblogs (Kalz et al 09d). With this framework we will be able to improve the construction of test corpora on the one hand and we will extend our approach on other methods and approaches that go beyond the application of Latent Semantic Analysis alone.

Chapter 4.3: SWeMoF: A semantic framework to discover patterns in learning networks[10]

[10] Kalz, M., Beekman, N., Karsten, A., Oudshoorn, D., van Rosmalen, P., van Bruggen, J., & Koper, R. (2009d). SWeMoF: A semantic framework to discover patterns in learning networks. In Cress, U., Dimitrova, V., & Specht, M. (Eds.). Learning in the Synergy of Multiple Disciplines. Proceedings of the Fourth European Conference on Technology-Enhanced Learning. Nice, France. Lecture Notes in Computer Science Vol. 5794. pp. 160-165. Berlin:Springer-Verlag.

Abstract

In this chapter we introduce SWeMoF, a semantic framework to discover patterns in learning networks and the blogosphere. Based on a description of the state of the art in data mining, text mining and blog mining we discuss the architecture of the Semantic Weblog Monitoring Framework (SWeMoF) and provide an outlook and an evaluation perspective for future research and development.

Introduction

In the past we have concentrated on the evaluation of Latent Semantic Analysis (LSA) to approximate the prior knowledge of learners in learning networks. We could show that Latent Semantic Analysis (LSA) is a promising method to support this process (Kalz et al., 2009b). Several other examples show that semantic services and language technology have the potential to help to reduce tutor load and to increase efficiency in technology-enhanced learning (Van Rosmalen, 2008). We expect that the application of such approaches can help in personalization processes, the automatic generation of metadata and the discovery of structural patterns in learning networks. On the other hand we made the experience that the effort to develop and evaluate learner support services based on text- and data-mining methods is a very challenging task since a lot of different tools and sources are involved and manually processing of data is needed. Based on this aspect and the need to extend our research to other methods and approaches we have developed a prototypical solution that can help to find semantic patterns in learning networks. In this contribution we present the Semantic Weblog Monitoring Framework (SWeMoF). The prototypical framework that we discuss in this article employs feed parsing techniques and data- and text-mining algorithms for several types of experiments and prototyping scenarios.

A similar framework as proposed here has been described by Joshi & Belsare as BlogHarvest (Joshi & Belsare, 2006) and Chau et al. (Chau, Xu, Cao, Lam, & Shiu, 2009). The BlogHarvest framework is a conceptual framework for opinion and sentiment analysis that employs part-of-speech tagging, association rules and several miners for clustering and classification. The second proposal by Chau et al. consist of a blog spider to collect content, a blog parser to extract information, a blog analyzer and a blog visualizer. On the other hand this is a very general framework without any prototype or a detailed architecture.

The Semantic Weblog monitoring Framework (SWeMoF) enables researchers, course designers and learning technology developers to conduct several kinds of semantic experiments using different algorithms from natural language processing and data mining. We expect this framework to support the development of semantic technologies and web services to solve some basic problems in educational technology like formulated by Koper (Koper, 2004). Applying common data and text mining techniques for discovery, recommendation and similarity classification can help to overcome problems of efficiency and effectiveness of the learning process and the workload of tutors. In the next part of the paper we describe the state-of-the art in data mining, text mining and blog mining. Afterwards we introduce the architecture of SWeMoF and provide an evaluation outlook.

Data Mining, Blog Mining and Text Mining

Data Mining is a process to find patterns in large numbers of data (Witten & Frank, 2000). While data mining is applied most of the time to numerical data in large databases the application of techniques from data mining to textual data is called text mining. Inside the text mining research the application of text mining to weblogs is called blog mining.

The target of data mining is to discover meaning in a vast amount of data and to find patterns that are not recognizable by traditional statistical measurement and direct visual inspection. Witten and Frank (2000) refer to an increasing gap in today's society between the generation of data and the understanding of it [6]. In this sense data mining does not have the target to generate new data but to use existing data and to find structures which have not been explored before. Fayyad et al. describe the data mining process as an interactive and iterative process which involves several steps with different tasks (Fayyad, Piatetsky-Shapiro, Smyth, & Uthurusamy, 1996). A special focus on using data mining in educational settings and with educational data has been developed in the last years and applied to different educational problems (Romero & Ventura, 2007).

The same procedure as described above can also be applied to non-numeric data. If data mining is applied to text the process is called text data mining or text mining. Hearst (1999) defines text mining as a means of exploratory data analysis and he stresses the distinction between text mining and information retrieval. While information retrieval and information access are only about finding information which are hard to find because of a lot of similar information text mining in his opinion is a process which has the target to discover information that have never been encountered before. Several disciplines contribute to text mining research, the most important one computational linguistics/natural language processing (NLP) (Manning & Schutze, 2003). In addition several disciplines from literature studies to genetics and bio-informatics have applied text mining to solve some basic problems in their domain of research. The application of data mining techniques and text mining problems to weblogs is coined as Blog Mining. Blog mining is a very recent research direction. Barone (2007) provides a good overview of research done in this area until 2007. The framework proposed in this contribution will allow blog mining experiments with a special focus on discovery, classification and clustering. In the next part we describe the architecture of the framework.

Architecture of the Semantic Weblog Monitoring Framework

SWeMoF is an object-oriented, web-based application designed for semantic experiments on the basis of content produced from weblogs and other text-based applications which offer an RSS-feed. Within this framework several data mining/natural language processing experiments are possible. Every experiment takes the content of one or more weblogs as input, applies one or more algorithms/miners to the content and gives an output which can be downloaded. The level of input can be the whole content of a weblog (set level), content from a dedicated category in a weblog (category level) or even only dedicated postings (document level).

The prototype has implemented 5 example algorithms/miners for three different experiments: Semantic Similarity, Classification and Clustering. The prototype is written in Java and makes use of an integrated database and the Echo framework for the interface. The example algorithms are implemented using the Weka framework, but the SWeMoF framework does not depend on it. Both filters and text mining algorithms can be written from scratch or by using any available components and libraries. For the design of the system the following use cases have been defined:

- Corpus Creation

A corpus has to be defined before an experiment can be created. This corpus can be constructed from several RSS-Feeds and/or OPML files. Besides this functionality, the domain corpus can be combined with a general language corpus which has been discussed as an important option in several information retrieval scenarios. For classification experiments several examples need to be classified manually before an experiment can be executed. In the classification experiment these 'gold standard' examples are needed to allow a semantic comparison between the classified documents and the unclassified documents. This step can be done by inspecting the corpus directly or during the creation of an experiment.

- Experiment Creation

In the experiment creation phase the parameters for a text mining experiment can be configured. These parameters consist of a corpus, an optional general language corpus, filters and a text mining algorithm. Further, the level on which the experiment is conducted (set, category or document) must be configured. It is also possible to disable a part of the corpus on any level: set, category or document. After an experiment has been created it can be executed. This division between experimentation and execution allows for repeating experiments and comparing results with different settings.

- Result Presentation & Download

After the execution of an experiment the results are presented to the user and the user can download the results.

- Adding of additional miners

In the current prototype the following miners have been implemented: Naive Bayes Classifier, IB1 Classifier, EM Clusterer, Simple K-Means Clusterer and a similarity rater using LSA. In addition, LSA can be combined with the miners implemented. But it is easy to add additional miners into the system.

The SWeMoF Framework allows the user to either create new experiments or retrieve and execute older experiments that have been stored. The parameters of an experiment (corpus, general language corpus, filters, miner, mining level) are saved in an experiment configuration. The following figure shows how a text mining experiment is conducted with SWeMoF.

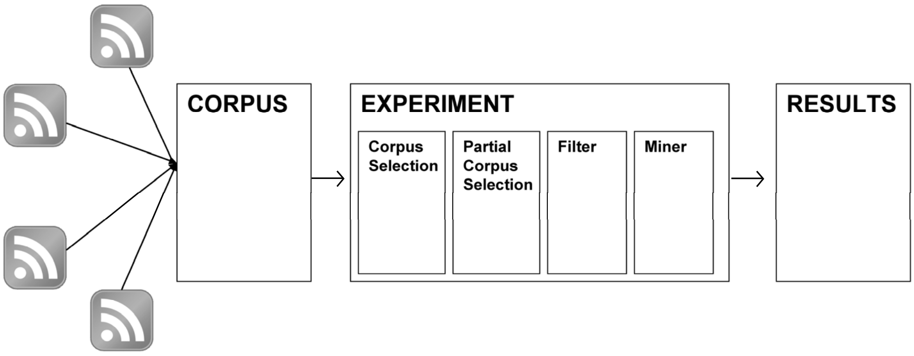

Figure 4.6 Overview of the components of the SWeMoF system

The *Input module* is responsible for the import of text. Single texts (e.g. single web posts) are organized in groups to create a hierarchy. Single text documents must be grouped in a document category, document categories must be grouped in document sets. Since SWeMoF's main focus is on web feeds, this design has been chosen to reflect the structure of these feeds. Even when only a single text document is imported it will have to be placed inside a document category, and the document category inside a document set. It is important to note that the corpus is not created by the Input module but by the Corpus module. For the feed parsing we have used the ROME library. ROME is a set of open source Java tools for parsing, generating and publishing RSS and Atom feeds. The *Corpus module* is responsible for the aggregation of documents generated from the input text by the input module. A corpus contains a collection of document sets. The structure within these sets is as described in the Input module section. After execution of an experiment the results are generated. The *View module* can display these results in different ways. In the prototype the View module is not implemented, instead a textual output is generated directly from the Result object. Three types of information can be stored through the *DAO module*: the corpus, the configuration parameters of the experiment, and the results of an experiment. Finally a *GUI* takes care of the interaction with the user, enabling him to create new experiments, retrieve old experiments, retrieve results of experiments, and set parameters of an experiment. The

GUI takes care of the interaction between users and the SWeMoF application. It is designed to let the user select text to convert (single document or web feeds); select filters (preprocessors) to generate the appropriate text corpus; select an experiment and choose a way to study the outcomes of the experiment. The GUI has been implemented with the Echo web framework.

The SWeMoF framework can be extended on several areas. The framework focuses on weblog monitoring and thus the focus for the prototype has been on implementing RSS and OPML as the document source. The input module however is designed in such a way that it can easily be extended with other input sources by implementing the appropriate interfaces. The second more important part where SWeMoF can be extended is in the filters and miners. To add a new filter or text mining algorithm all that needs to be done is implement the interface Filter or Miner and create a descriptor. The descriptor will tell the GUI what the Filter or Miner does and which options can be set. After this has been done, the descriptor can be added in to the registry. SWeMoF will then automatically make this filter or miner available to the end user.

Discussion, Outlook and Future Work

At the current stage of the development we could conduct several tests related to code functionality and result quality. After the components have been tested alone the integrated system has been tested to see if the system supports the use cases for which it was designed for. In addition we have compared the system results with the results of using Weka directly. The integration testing confirmed that the system is able to support the use cases and the comparison to Weka was successful as well. A real end-user and usability testing could not be conducted yet, but we are planning to present the system to researchers and learning technology developers with different levels of prior knowledge about data and text mining. For this purpose we are planning to combine traditional usability testing with the hedonic and pragmatic approach developed by Hassenzahl (2001). In this framework the "hedonic quality" aspect covers non-task-oriented quality aspects like innovativeness or originality and takes appealingness of a software system into account as well.

As a next step we will conduct an end-user testing with colleagues in the field. Based on the feedback of the potential end-users we will improve the system. The full code of the framework has been released under a GPL license (Kalz, 2009a) and a demonstration of the framework is available (Kalz, 2009b). Depending on the reaction of end-users of the system we might improve the storage and presentation of the results. In addition we are going to extend the system with more miners from Weka and use it as an evaluation instrument for the development of several semantic web-services in the future.

Section 5: General Discussion

Introduction

With the Bologna Declaration (Council of Europe, 1999) the European Union has set the target to establish a European Higher Education Area (EHEA). There are some core problems that need to be solved to reach this target. First of all a system of easily readable and comparable degrees needs to be established to allow learners to switch between educational institutions across borders. This has been addressed to some extent by the diploma supplement and the European Credit Transfer System (ECTS). Mobility of learners and trainers is the other core target that must be reached to establish the EHEA. But the aspect of comparable degrees and the mobility of learners cannot be reached only by top-down-agreements over the comparability of degrees. Especially the important perspective of informally acquired competences and their comparability and recognition is not easy to tackle with top-down approaches. The accreditation or prior learning (APL) is a process that takes into account formal accredited degrees and competences but also informally acquired competences of learners. But this process has many problems in terms of scalability, costs and quality.

At the same time one of the main targets of technology-enhanced learning is to allow for personalized learning that adapts according to preferences and characteristics of learners like competence level or prior knowledge. While this problem can be addressed in formal educational context based on the metadata and logic of the offering institution in non-formal settings like learning networks top-down approaches for personalizing content to the prior knowledge of learners are useless.

For this situation we have proposed a model for a placement support service that can take the similarity between content of learner portfolios and target learning units as a proxy for prior knowledge assessment. For this purpose we have employed a reduced text vector space model similar to Latent Semantic Analysis. As introduced at the beginning of this thesis this approach was complicated due to several constraints. The absence of large document collections to allow LSA to learn the language was identified as one important issue. Based on the specific situation in learning networks we have decided to evaluate our model on what we call small scale corpora. This approach has been evaluated in regard to sensitivity, reliability and validity.
We could show with two empirical studies that using the combined approach of filtering, multiple criteria and variance-based selection we were

able to achieve sufficient sensitivity. In addition we could show in the second study that the classification results by our model perform well and that we can minimize the false negative and false positive cases under 10%.

This method has been implemented in two software prototypes that allow us to integrate the approach in several placement support services. With the second prototype we made a step into the situations that go beyond the application of semantic similarity and dimensionality reduction defined in the long-term research agenda.

Review of the results

One important first result of our research project is that our general assumption can be confirmed that domain experts estimate an exemption of learners based on the semantic similarity between documents in the learner portfolio and the target documents. In 83 % of the cases in our second empirical study a semantic similarity higher then 2 (which was equal to some semantic similarity on the Likert scale) lead to an exemption by the domain experts. This basic assumption was also the foundation of our model for a placement support service that we have introduced based on a review of existing literature from learning technology standards and metadata, APL research and portfolio assessment in chapter 2.2 of this thesis.

Because of the unavailability of a large general language corpus in Dutch at the time of our studies and because of the learning networks context we have decided to apply the text vector space model and dimensionality reduction similar to Latent Semantic Analysis (LSA) for the research and development of a placement support service. This focus on what we call "small scale corpora" has led to specific issues to reach a sufficient sensitivity of such a service and to develop a method to identify the ideal numbers of dimensions to retain. These issues have been addressed in the empirical section of the thesis.

In chapter 3.1 we have addresses these issue. The specific situation of learning networks and APL procedures that requires to work with smaller corpora demanded specific approaches to optimize the method according to two criteria. The first criterion is the sensitivity of the method. We could show in this study that the application of a stopword strategy between 30% and 50% is an important step to reach a sufficient discrimination. These results could be confirmed for a manually constructed corpus that has been prepared with much pre-processing and filtering and a second corpus that was taken from course material from the Open University of the Netherlands. It was only when we applied stopping strategies that we could obtain correlations that were high within homogenous sets and low between non-related sets simultaneously in the study. This is also an indication that optimizing for a single criterion, such as a high correlation with a target variable, may be insufficient if we also want to obtain sufficient discriminatory power. The second criterion was the identification of the ideal number of dimension to use. In this study we have proposed a method to identify the ideal number of dimensions by taking into account the connection between singular values and the variance they account for in the Term-Document matrix analyzed by LSA. Since the sum of the squared singular values is equal to the sum of squared cell frequencies and the Term-Document matrices are sparse we can therefore estimate the proportion of variance accounted for by using the squared singular values only. Thus, one may select a minimum and a maximum number of singular values that corresponds with a bandwidth of variance accounted for. The ideal dimensionality is very important because beyond a certain (unknown) threshold, using additional singular values in the reproduction of the data will introduce more and more error. Clear indications of this effect are the diminishing size of the correlations that we found in all analyses in this study. Limiting the number of singular values helps to prevent this effect. This study has established the methodological foundations for the empirical validation of the second empirical study and it has contributed an important part to realize the model presented in chapter 2.2.

In the second study reported in chapter 3.2 we have applied the method developed in the first study in an APL procedure at the psychology faculty of the Open University of the Netherlands. We could confirm our findings of the first empirical study. The application of a 50% stopping strategy allowed us to differentiate between the target learning units of the study. Though, we could improve our results and reach our self-set target of a false positive and false negative rate under 10% of all cases. We could show based on a ROC-curve that our classification model delivers satisfying results.

To allow the implementation of the method applied in this thesis a technological solution was needed that could on the one hand be generic enough to be implemented in different contexts but at the same time flexible enough to be changed and aligned to different local standards and thresholds. For this purpose we have developed a web-service that allows the implementation of this method in other software-systems like e.g. environments that lead learners through the APL procedures. A prototype of such a web-service has been presented in chapter 4.2 of this thesis. A first technical evaluation has revealed some problems with the updating of the term-document matrix which is computationally an intensive process. For this purpose we have proposed several ideas to improve the web-service to decrease the computation time involved when contacting the web-service.

The development of the Semantic Weblog Monitoring Framework (SWeMoF) presented in chapter 4.3 of this thesis will allow us to solve some of the issues that have hindered us from realizing already a step towards the long-term research agenda defined in chapter 2.2 of this thesis. On the one hand, the framework can solve the problem of corpus construction. It allows the construction of corpora from sources of the social web. This means that we will be able to construct different domain corpora and larger corpora more effortless with freely available sources. In addition it allows also the use of other data formats like metadata. Furthermore we will be able to test several methods from text mining and data mining to study the best method or combination of methods applied to the placement problem. This possibility will also allow us to evaluate more easily several algorithms and miners before they are implemented in other semantic web-services developed in the future.

The scope and limitations of the research

The research project and its assumptions have several limitations which need to be discussed in the light of a generalizability of the results presented here. First of all, the method proposed in this thesis can only be applied in domain with a highly textual basis. This means that a placement support services will be able to provide satisfactory results in one domain while it will fail to capture the underlying structure of a domain like mathematics or chemistry which operate on a high abstraction level and use a lot of signs and symbols.

Another limitation of the method proposed here is that only parts of the APL procedure are covered. While a placement support service is an important tool for students to pick the right material to be presented it can help domain experts to presort relevant document submitted by learners. If the method proposed here can be optimized in a way that the actual decisions about exemptions can be automated is unclear and we tend to say that such an anonymous procedure would maybe reduce costs but would also decrease the effect of offering personalized learning.

While it is an important target to increase the efficiency of learning by increasing the ability of students to progress in their studies without unnecessarily having to repeat material or levels of study it can also be argued from an educational perspective that repetition of topics is recommended or needed.

The scope of our empirical research was the application of the method on a course level. In principal the same approach should work on a programme level as well but it needs to be empirically evaluated if we would have to deviate from the method shown here.

Both empirical studies have their own individual limitations. The first study was a methodological study involving two different Dutch test corpora. One corpus was carefully constructed and controlled while the other corpus was taken from parts of course material of the OUNL. This research was explorative in the sense that the use of LSA and similar dimensionality reduction techniques for Dutch corpora was very scarcely evaluated at the time we conducted the study. Only a few number of studies have applied the method to Dutch corpora and if it is possible to reach similar results at all with Dutch texts as with English texts is unclear.

The second study revealed that the application of LSA for prior knowledge approximation provides a well performing model for document classification in APL procedures. But the generalizability of the results from this study is limited since the data we have collected where skewed due to the low occurrence of relevant documents submitted by the learners. On the other hand, the collection of more data would have probably not improved the situation since this rate of relevant submission has been confirmed by an APL expert as a general pattern. This is a part of the problem we want to address with a placement support service.

Another limitation was identified by the second study. Some of the false negative cases have shown that the domain experts apply methods beyond simple text similarity. In one interesting case a document describing an experimental setup was used as a basis for exemptions although the semantic similarity was rated low. In an interview one of the domain experts described that the process of analyzing learner material according to prior knowledge takes semantic similarity only as one of the aspects that lead to exemptions. In this case the domain experts deducted that the material presented needs specific prior knowledge before a learner is able to describe an experiment on the level presented. This shows that the other dimensions of the long-term research agenda needs to be researched as well to come up with placement support-service that can work on the basis of different data sources and types from unstructured to more structured data. The reasons why we did not go a step further towards the direction defined in our long-term research agenda had to do mainly with the availability of datasets and tools for the evaluation. We could not easily lay our hands on a sufficient number of learner and course data that can be used for the evaluation of the more structured approaches. In addition after testing some of the available tools which we could have used in these analysis we have seen that much more work was needed before these tools could be applied. The SWeMoF framework addresses some of these problems identified in this analysis.

The practical implications

This project contributed to the field of technology-enhanced learning by its application of methods and approaches from the field of language technology and text mining to educational problems. The author beliefs that this combination offers a lot of potential to contribute new methods and approaches to the field and for the learner. In this regard we expect the appearance of probabilistic services which are not accurate like a personal tutor but which are able to support learners and their specific need for orientation and feedback and which contribute to the decrease of tutor load in online learning environments.

The method and technologies developed here should be integrated in further studies and implementations. We have not developed a reference implementation in an exemplary learning infrastructure but a web-service to allow on the one hand the implementation of the method in APL procedures in different institutions and on the other hand the evaluation of the method in technology-enhanced learning environments to construct individualized learning paths.

The web-service is based on a PHP implementation of LSA (Van der Vegt, 2009) and its code and some instructions how to set-up and use the web-service are available on the project page (Kalz, 2009f). With the software it is possible to set-up a placement web-service and to manipulate the underlying LSA settings and to evaluate its performance in a given learning environment. The challenge to set-up a placement service is to identify the thresholds that should be used to identify learning activities as relevant or not relevant. In this regard the interpretation aspect of the results provided by the web-service are important. The threshold depends on the local context of the implementation. As discussed in the last section in large networks a procedure to update the Term-Document-Matrix is needed.

The methodological approach to apply Latent Semantic Analysis (LSA) in small scale corpora might contribute to research with similar requirements like the ones presented here. We believe that the method to estimate the ideal number of singular values to retain is a promising step towards a more empirically grounded way of choosing the dimensions in LSA and LSA-like studies.

Further research

This research project should pave the way for bottom-up approaches to estimate the prior knowledge of learners in situation where no detailed and structured data about the learners and the learning content is available. Since at the time of our studies no large corpus was available in Dutch we decided to go for a small scale corpora solution. This situation has slightly improved since then and we can expect to lay our hands on a large general language corpus in the near future. Then we are able to study the influence of a large corpus on the performance of LSA for placement support. Furthermore our empirical results need to be evaluated on the basis of larger datasets and from several different domains. If the same results can be obtained in other languages is still an open question.

Our long-term research agenda already defines a direction for further research. In this thesis we have only concentrated on the most complicated situation where no clear metadata about learners and learning content is available. Solutions for the other two situations where metadata or ontologies are describing them are still missing. We started some work in this direction on two levels. On the first level we have concentrated on the development of the SWeMoF framework that will allow us to evaluate a combination of approaches for the placement problem. Specifically it needs to be addressed if traditional data mining methods can reach similar or better results. Furthermore we started to explore some ideas to compare representations of texts on a higher level like semantic networks or concept maps for the estimation of prior knowledge of learners (Kalz et al., 2009e; Berlanga et al., 2009).

From a technology-development perspective the performance of our web-service prototype needs to be improved. For this purpose we have already started to work on an efficient method to update the Term-Document Matrix when new texts get added to the learning network. The SWeMoF framework should be extended in the future with more miners and algorithms.

The real fit into APL procedures must be evaluated in several different contexts. The role of the placement support service in the APL process is that of a system that advises during the stage where a portfolio of evidence is put together on the relevance of documents submitted. As a preparation of the stage in which the portfolio is assessed the service annotates the documents as to their relevance for possible exemptions. To avoid introducing workload the system should be fined-tuned to avoid flagging documents as relevant where they are not.

An important option for aligning the service to local thresholds could be realized in the future via a feedback loop in which the system can also compare documents to documents that were previously classified by human raters. Thus documents that were accepted as evidence leading to an exemption, as well as documents that were rejected as evidence can be fed back to the system as standards to which new documents can be compared. The role of such a system would be left unchanged: it can scrutinize a portfolio and signal which documents bear no similarity to target material or prior positive cases. This would then result in a recommendation to remove the document from the portfolio. For assessors the service would annotate the documents in the portfolio with similar content and links to portfolio's with relevant (positive or negative) documents that were assessed earlier.

Although we conceived the service as assisting and not replacing actors, it has become evident that more attention should be given to automatic corpus creation. Of particular interest here, is the automatic segmentation of documents based on semantic principles, rather than on an arbitrary maximum size. A way forward here is shown by Gounon & Lemaire (2002) who used LSA to identify "topical paragraphs" in a text. It remains to be seen whether this will also handle the problem of rating similarity on general language characteristics. Beside adding more domain specific material and topical segmentation we are considering feeding the system with (several copies of) metadata and related material, such as keywords, index terms, summaries, and learning objectives.

References

Barker, K. (1999). The electronic learning record: Assessment and Management of Skills and Knowledge. A literature review, Vancouver. Canada:Literacy BC and the National Literacy Secretariat.

Barker, K. (2006). Foreword. In A. Jafari & C. Kaufman (Eds.). Handbook of research on ePortfolios (pp. xxvi-xxviii). London, UK: Hershey.

Barone, F. (2007). Current approaches to data mining blogs (pp. 1-7). Kent: University of Kent.

Barton, J., Currier, S. & Hey, J. M. N. (2003). Building quality assurance into metadata creation: an analysis based on the learning objects and e-prints communities of practice. In Sutton, S. and Greenberg, J. and Tennis, J. (Eds.). *Proceedings 2003 Dublin Core Conference: Supporting communities of discourse and practice - metadata research and applications*, 2003 Seattle, Washington (USA).

Berlanga, A., Brouns, F., Van Rosmalen, P., Rajagopal, K., Kalz, M., Stoyanov, S., et al. (2009). Making use of language technologies to provide formative feedback. In S. D. Craig & D. Dicheva (Eds.), *Proceedings of the workshop of natural language processing in support of learning: Metrics, feedback and connectivity. 14th Int. Conference on Artificial Intelligence in Education (AIED 2009)* (pp. 1-7). Brighton, UK.

Bochkanov, S., & Bystritsky, V. (2008). ALGLIB libraries [Computer software].

Bouma, G., & Klein, E. H. (2001). Volkskrant database and practicum [Data file].http://hagen.let.rug.nl/~hklein/Volkskrant/Volkskrant/
http://hagen.let.rug.nl/~hklein/pract_rainbow2.html

Brooks, C. H., & Montanez, N. (2006). Improved annotation of the blogosphere via autotagging and hierarchical clustering. In *Proceedings of the 15th international conference on World Wide Web* (pp. 625-632). Edinburgh, Scotland: ACM Press.

Brusilovsky, P., Kobsa, A., & Nejdl, W. (2007). The adaptive web: methods and strategies of web personalization (Lecture Notes in Computer Science / Information Systems and Applications, incl. Internet/Web, and HCI, 4321) (p. 770). Berlin: Springer.

Brusilovsky, P., & Millan, E. (2007). User models for adaptive hypermedia and adaptive educational systems. In P. Brusilovsky, A. Kobsa, & W. Nejdl (Eds.), The adaptive web: methods and strategies of web personalization (Lecture Notes in Computer Science / Information Systems and Applications, incl. Internet/Web, and HCI, 4321) (pp. 3 - 53). Berlin:Springer.

Brusilovsky, P., & Henze, N. (2007). Open corpus adaptive educational hypermedia. In P. Brusilovsky, A. Kobsa, & W. Nejdl (Eds.), The adaptive web: methods and strategies of web personalization (Lecture Notes in Computer Science / Information Systems and Applications, incl. Internet/Web, and HCI, 4321) (pp. 671-696). Berlin: Springer.

Carroll, N., & Calvo, R. (2005). Certified assessment artifacts for eportfolios. In *Proceedings of the third international conference on information technology and applications*, ICITA 2005. (p. 130 – 135). Palo Alto, USA: IEEE.

Cattell, R. (1966). The scree test for the number of factors. *Multivariate behavioral research*, 1 (2), 245-276.

Chan, L., & Zeng, M. (2006). Metadata interoperability and standardization–A study of methodology Part I. *D-Lib Magazine*, 12 (6). Retrieved 1st of July, 2006 from http://www.dlib.org/dlib/june06/chan/06chan.html

Chau, M., Xu, J., Cao, J., Lam, P., & Shiu, B. (2009). A blog mining framework. *IT Professional*, 11 (1), 36-41.

Conole, G., & Warburton, B. (2005). A review of computer-assisted assessment. *ALT-J, Research in Learning Technology*, 13 (1), 17-31.

Council of Europe (1999). The bologna declaration on the european space for higher education. Joint declaration of the European ministers of education. Bologna:European Commission.

Dolog, P., Schaefer, M. (2005). A framework for browsing and manipulating and maintaining interoperable learner profiles, *Proceedings of UM2005 - 10th International conference on user modeling*, July, 2005, Edinburgh, UK (Lecture Notes in Computer Science, 3538/2005). Berlin:Springer.

Deerwester, S., Dumais, S., Furnas, G., Landauer, T., & Harshman, R. (1990). Indexing by latent semantic analysis. Journal of the American society for information science, 41 (6), 391-407.

Dessus, P. (2004). Simulating student comprehension with latent semantic analysis to deliver course readings from the web. *Cognitive Systems*, 6 (2-3), 227 – 237.

Drachsler, H., Hummel, H., & Koper, R. (2007). Recommendations for learners are different: Applying memory-based recommender system techniques to lifelong learning. In *Proceedings of workshop on social information retrieval for technology-enhanced learning* (SIRTEL'07), 2nd European conference on technology enhanced learning (EC-TEL'07).Crete, Greece.

Drachsler, H., Hummel, H., & Koper, R. (2008). Personal recommender systems for learners in lifelong learning: requirements, techniques and model. *International Journal of Learning Technology*, 3 (4), 404-423.

Duffy, T., & Jonassen, D. (1992). Constructivism and the technology of instruction: A conversation. Hillsdale, New Jersey: Lawrence Erlbaum Associates.

Fawcett, T. (2004). ROC graphs: notes and practical considerations for researchers. HP Labs Tech Report. Palo Alto (CA), USA.

Fayyad, U., Piatetsky-Shapiro, G., Smyth, P., & Uthurusamy, R. (1996). Advances in knowledge discovery and data mining (p. 625). AAAI Press.

Foltz, P., Laham, D. & Landauer, T. (1999). Automated essay scoring: Applications to educational technology. In B. Collis & R. Oliver (Eds.), *Proceedings of world conference on educational multimedia, hypermedia and telecommunications* 1999 (pp. 939-944). Chesapeake, VA: AACE.

Foroughi, R. (2003). Proposing new elements for pedagogical descriptions in LOM. In V. Uskov (Ed.) *Proceedings of the 7th IASTED international conference on computers and advanced technology in education*. Kauai, Hawaii, USA: Acta Press.

Gergen, K. (1999). An invitation to social construction. London: Sage Publications.

Geser, G. (2007). Open educational practices and resources: OLCOS Roadmap 2012. Salzburg, Austria:Salzburg Research.

Giles, J., Wo, L., & Berry, M. (2003). GTP (General Text Parser) software for text mining. In H. Bozdogan (Ed.), Statistical data mining and knowledge discovery. Boca Raton, USA: CRC Press.

Good, N., Schafer, J., Konstan, J., Borchers, A., Sarwar, B., Herlocker, J., et al. (1999). Combining collaborative filtering with personal agents for better recommendations. In *Proceedings of the national conference on artificial intelligence* (pp. 439-446). John Wiley & Sons.

Gounon, P., & Lemaire, B. (2002). Semantic comparison of texts for learning environments. In F. Garijo, J. C. Riquelme Santos, & M. Toro (Eds.), Advances in artificial intelligence - IBERAMIA 2002, *Proceedings of the 8th Ibero-American conference on AI*, Seville, Spain, November 12-15, 2002 (pp. 724-733) (Lecture Notes in Computer Science, 2527), Berlin:Springer.

Graesser, A., Chipman, P., Haynes, B., & Olney, A. (2005). AutoTutor: An intelligent tutoring system with mixed-initiative dialogue. *IEEE Transactions on Education*, 48 (4), 612-618.

Haldane, A., Meijer, R., Newman, C., & Wallace, J. (2007). To what extent can the APEL process be facilitated through the use of technology? In D. Young & J. Garnett (Eds.), *Proceedings from the work-based learning futures conference* (pp. 151-161). Buxton, UK: University of Derby, Middlesex University.

Haley, D., Thomas, P., Roeck, A. D., & Petre, M. (2005). A research taxonomy for latent semantic analysis-based educational applications. In G. Angelova, K. Bontcheva, R. Mitkov, N. Nicolov, & N. Nikolov (Eds.), *Proceedings of the international conference on recent advances in natural language processing* (pp. 575-579). Borovets, Bulgaria: Bulgarian Academy of Sciences.

Hamel, L. (2008). Model assessment with ROC curves. In The encyclopedia of data warehousing and mining (2 ed.). Hershey:Idea Group Publishers.

Hassenzahl, M. (2001). The effect of perceived hedonic quality on product appealingness. *International Journal of Human-Computer Interaction,* 13 (4), 481-499.

Hearst, M. (1999). Untangling text data mining. In *Proceedings of the 37th annual meeting of the Association for Computational Linguistics* (pp. 3-10). Morristown, NJ, USA: Association for Computational Linguistics.

Herder, E., & Kärger, P. (2008). Hybrid personalization for recommendations. In *Proceedings of the 16th workshop on adaptivity and user modeling in interactive system* (pp. 20-25). Würzburg, Germany: University of Würzburg.

Herlocker, J., Konstan, J., Terveen, L., & Riedl, J. (2004). Evaluating collaborative filtering recommender systems. *ACM Transactions on Information Systems,* 22 (1), 5-53.

HR XML Consortium (2004). Competencies (neasurable characteristics). Retrieved 1st of July, 2006 from http://ns.hr-xml.org/2_3/HR-XML-2_3/CPO/Competencies.html

Hüning, M. (2005). TextStat simple text analysis tool [Computer software]. Berlin, Germany: Free University of Berlin.

IEEE Learning Technology Standards Committee (2006). Reusable competency definitions (P1484.20.1, Draft 3). Retrieved 1st of July, 2006 from http://www.ieeeltsc.org/wg20Comp/wg20rcdfolder/IEEE_1484.20.1.D3.pdf

IMS Global Learning Consortium (2001), IMS earner information packaging specification. Retrieved 1st of July, 2006 from http://www.imsglobal.org/profiles/

IMS Global Learning Consortium (2002). IMS reusable definition of competency or educational objective specification. Retrieved 1st of July, 2006 from http://www.imsglobal.org/competencies/

IMS Global Learning Consortium (2003). IMS learning design best practice and implementation guide.

IMS Global Learning Consortium (2005). IMS eportfolio specification. Retrieved 1st of July, 2006 from http://www.imsglobal.org/ep/index.html

Janssen, J., Berlanga, A., Vogten, H., & Koper, R. (2008). Towards a learning path specification. International *Journal of Continuing Engineering Education and Lifelong Learning*, 18(1), 77–97.

Joosten-ten Brinke, D., Van Bruggen, J., Hermans, H., Burgers, J., Giesbers, B., Koper, R., et al. (2007). Modeling assessment for re-use of traditional and new types of assessment. *Computers in Human Behavior*, 23 (6), 2721-2741.

Joosten-ten Brinke, D. J., Brand-Gruwel, S., Sluijsmans, D., & Jochems, W. (2008). The quality of procedures to assess and credit prior learning: Implications for design. *Educational Research Review*, 3 (1), 51-65.

Joshi, M., & Belsare, N. (2006). BlogHarvest: Blog mining and search framework. In L. V. Lakshmanan, P. Roy, & A. K. Tung (Eds.), *Proceedings of the 13th international conference on management of data* (COMAD). Delhi, India: Computer Society of India.

Kalz, M, Van Bruggen, J., Rusman, E., Giesbers, B., & Koper, R. (2007). Positioning of learners in learning networks with content, metadata and ontologies. *Interactive Learning Environments*, 15, 191-200.

Kalz, M., Drachsler, H., Van Bruggen, J., Hummel, H., & Koper, R. (2008a). Wayfinding services for open educational practices. *International Journal of Emerging Technologies in Learning*, 3 (2), 1-6.

Kalz, M., Van Bruggen, J., Giesbers, B., Waterink, W., Eshuis, J., Koper, R., et al. (2008b). A model for new linkages for prior learning assessment. *Campus-Wide Information Systems*, 25 (4), 233-243.

Kalz, M., Van Bruggen, J., Giesbers, B., Rusman, E., Eshuis, J., Waterink, W., & Koper, R. (2009). A validation scenario for a placement service in learning networks. In Koper, R. (Ed.) Learning Network Services for Professional Development. Berlin:Springer.

Kalz, M., Van Bruggen, J., Rusman, E., Giesbers, B., & Koper, R. (2009a). Latent semantic analysis of small scale corpora for placement in learning networks. Manuscript submitted for publication.

Kalz, M., van Bruggen, J., Giesbers, B., Waterink, W., Eshuis, J. & Koper, R. (2009b). Where am I? – An empirical study about learner placement based on semantic similarity. Manuscript submitted for publication.

Kalz, M., Drachsler, H., Van der Vegt, W., Van Bruggen, Glahn, C., J., & Koper, R. (2009c). A placement web-service for lifelong learners. In Tochtermann, K., Lindstaedt, S. (eds.). *Proceedings of the 9th international conference on knowledge management and knowledge technologies*. 2-4 September 2009, Graz, Austria.

Kalz, M., Beekman, N., Karsten, A., Oudshoorn, D., van Rosmalen, P., van Bruggen, J., & Koper, R. (2009d). SWeMoF: A semantic framework to discover patterns in learning networks. In Cress, U., Dimitrova, V., & Specht, M. (Eds.). Learning in the Synergy of Multiple Disciplines. *Proceedings of the Fourth European Conference on Technology-Enhanced Learning*. Nice, France. Lecture Notes in Computer Science Vol. 5794. pp. 160-165. Berlin:Springer-Verlag.

Kalz, M., Berlanga, A., Van Rosmalen, P., Stoyanov, S., Van Bruggen, J., Koper, R., et al. (2009e). Semantic networks as means for goal-directed formative feedback. In Kreativität und Innovationskompetenz im digitalen Netz - Creativity and innovation competencies in the web, Sammlung von ausgewählten Fach- und Praxisbeiträgen der 5. EduMedia Fachtagung 2009 (pp. 88-95). Salzburg, Austria: Salzburg Research.

Kalz, M.(2009a). Semantic Weblog Monitoring Framework project page, http://swemof.sf.net

Kalz, M. (2009b) Semantic Weblog Monitoring Framework demonstration page, http://www.swemof.org

Kamtsiou, V., Naeve, A., Kravcik, M., Burgos, D., Zimmermann, V., Klamma, R., et al. (2008). A roadmap for technology enhanced professional learning (TEPL). Prolearn Consortium. Deliverable D12.15.

Karampiperis, P., & Sampson, D. (2005). Adaptive learning resources sequencing in educational hypermedia systems. *Journal of Educational Technology & Society*, 8 (4), 128-147.

Kintsch, W. (2002). The potential of latent semantic analysis for machine grading of clinical case summaries. *Journal of biomedical informatics*, 35 (1), 3-7.

Klamma, R., Chatti, M. A., Duval, E., Hummel, H., Hvannberg, E. H., Kravcik, M., et al. (2007). Social software for life-long learning. *Journal of Educational Technology & Society*, 10 (3), 72-83.

Koper, R. (2004). Use of the semantic web to solve some basic problems in education: increase flexible, distributed lifelong learning, decrease teacher's workload. *Journal of Interactive Media in Education*, 2004 (6). Special Issue on the Educational Semantic Web.

Koper, R., Rusman, E., & Sloep, P. (2005). Effective learning networks. *Lifelong Learning in Europe*, 1 (1), 18-27.

Koper, R., & Specht, M. (2008). Ten-Competence: Life-long competence development and learning. In M.-A. Sicilia (Ed.), Competencies in Organizational e-learning: concepts and tools (pp. 234-252). Hershey (PA):Idea Group Publishing.

Küster, U., König-Ries, B., Petrie, C., & Klusch, M. (2008). On the evaluation of semantic web service frameworks. *International Journal on Semantic Web and Information Systems*, 4 (4).

Landauer, T., & Dumais, S. (1997). A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction and representation of knowledge. *Psychological Review*, 104 (2), 211-240.

Landauer, T., Foltz, P., & Laham, D. (1998). An introduction to latent semantic analysis. *Discourse Processes*, (25), 259-284.

Landauer, T. (2007). LSA as a theory of meaning. In T. Landauer, D. McNamara, S. Dennis, & W. Kintsch (Eds.), Handbook of latent semantic analysis (pp. 3-35). Lawrence Erlbaum Associates.

Lave, J., & Wenger, E. (1991). Situated learning: Legitimate peripheral participation. Cambridge:Cambridge University Press.

Lindstaedt, S. N., Scheir, P., Lokaiczyk, R., Kump, B., Beham, G., Pammer, V., et al. (2008). Knowledge services for work-integrated learning. In *Proceedings of the European conference on technology enhanced learning* (ECTEL) 2008 (pp. 234-244). Berlin: Springer.

Louwerse, M., & Van Peer, W. (2006). Waar het over gaat in cijfers: LSA als kwantitatieve benadering in tekst- en literatuurwetenschap. [Where it's about in numbers: LSA as a quantitative approach in text and literature studies]. *Tijdschrift voor Nederlandse Taal-en Letterkunde*, 122, 22-36.

Maedche, A. & Staab, S. (2002). Measuring similarity between ontologies. In Knowledge engineering and knowledge management: Ontologies and the semantic web. *Proceedings of the 13th European conference on knowledge acquisition and management* (EKAW), Sigüenza, Spain, October 1–4, 2002 (pp. 15-21) (Lecture Notes in Computer Science, 2473/2002). Berlin:Springer.

Manning, C. D., & Schutze, H. (2003). Foundations of statistical natural language processing (6 ed.) MIT Press.

Mason, J., & Sutton, S. (2005). Dublin Core Metadata Initiative Education Working Group. Draft Proposal.

McDonald, R., Boud, D., Francis, J., & Gonzci, A. (1995). New perspectives on assessment. UNESCO Studies in Technical and Vocational Education. Paris:UNESCO.

Melville, P., Mooney, R., & Nagarajan, R. (2002). Content-boosted collaborative filtering for improved recommendations. In *Proceedings of the National Conference on Artificial Intelligence* (pp. 187-192). John Wiley & Sons.

Merrifield, J., McIntyre, D., & Osaigbov, R. (2000). Mapping APEL: Accreditation of prior experiential learning in English higher education. London: Learning from Experience Trust.

Microsoft. (2005). Encarta 2005 Standard (Dutch version) [Computer software].

Minton, A. (2008, June). E-Apel Presentation. In Bernstein, K. (chair), Workshop accreditation of prior learning/APEL. London: West London Lifelong Learning Network.

Mishne, G. (2007). Applied text analytics for blogs (Doctoral Dissertation, University of Amsterdam) (p. 297). Amsterdam: Dutch Research School for Inofrmation and Knowledge Systems (SIKS).

Nakov, P., Popova, A., & Mateev, P. (2001). Weight functions impact on LSA performance. In *Proceedings of the Recent Advances in Natural language* (pp. 187--193). Tzigov Chark, Bulgaria: Bulgarian Academy of Science.

Ng, A., Hatala, M., & Gasevic, D. (2007). Ontology-based approach to formalization of competencies . In M.-A. Sicilia (Ed.), Competencies in organizational E-Learning. Concepts and tools (pp. 185-206). Hershey (PA): Idea Group Publishing.

NMC (2007). 2007 Horizon report. Austin (TX):New Media Consortium.

Nordeng, T., Lavik, S., & Meloy, J. (2004). E-Portfolios for meaningful learning and automated positioning. In McIlraith, S. A., Plexousakis, D., Harmelen, F. (Eds.), *Proceedings of the third international semantic web conference*, November 7-11, Hiroshima, Japan (pp. 7-11). (Lecture Notes in Computer Science; 3298). Berlin:Springer.

OAI (2005). The Open Archives Initiative Protocol for Metadata Harvesting v 2.0. Last retrieved 1. of Feburary 2009 at http://www. openarchives. org/, Open Archives Initiative.

Pellegrino, J., Chudowsky, N., & Glaser, R. (2001). Knowing what students know: The science and design of educational assessment. Washington, DC:National Academy Press.

Posea, V., & Harzallah, M. (2004). Building an ontology of competencies. In H/ Kühn (Ed.) *Proceedings of workshop on ontology and enterprise modelling: Ingredients for interoperability*. In conjunction with 5th international conference on practical aspects of knowledge management (pp. 24-34). Vienna:University of Vienna.

Ridgway, J., McCusker, S., & Pead, D. (2004). Literature review of e-assessment. Futurelab Series (p. 51). Bristol, UK:Futurelab.

Rohde, D. (2005). Doug Rohde's SVD C Library, Version 1.34 [Computer software].

Romero, C., & Ventura, S. (2007). Educational data mining: A survey from 1995 to 2005. *Expert Systems with Applications*, 33 (1), 135-146.

Sadler, D. (1989). Formative assessment and the design of instructional systems. *Instructional science*, 18 (2), 119-144.

Salton, G., Wong, A., & Yang, C. (1975). A vector space model for automatic indexing. *Communications of the ACM*, 18 (11), 613-620.

Sánchez-Alonso, S., & Sicilia, M.-A. (2005). Normative specifications of learning objects and learning processes: Towards higher levels of automation in standardized e-Learning. *International Journal of Instructional Technology and Distance Learning*, 2 (3), 3-12.

Shirky, C. (2003). The semantic web, syllogism and worldview, 2003, Retrieved 1st of July, 2006 from http://www.shirky.com/writings/semantic_syllogism.html

Skinner, K. (2005). Continuing professional development for the social services workforce in Scotland. Dundee:Scottish Institute for Excellence in Social Work Education.

Sluijsmans, D., Prins, F., & Martens, R. (2006). The design of competency-based performance assessment in e-learning. *Learning environments research*, 9 (1), 45-66.

Smith, K., & Tillema, H. (2003). Clarifying different types of portfolio use. *Assessment & Evaluation in Higher Education*, 28 (6), 625-648.

Sloep, P. (2008). Netwerken voor lerende professionals; hoe leren in netwerken kan bijdragen aan een leven lang leren. Unpublished inaugural address. Heerlen, The Netherlands: Open University of the Netherlands.

Soboro, I., & Nicholas, C. (1999). Combining content and collaboration in text filtering. In *Proceedings of the IJCAI workshop on machine learning in information* filtering (pp. 86-91). Stockholm, Sweden.

Steinhart, D. J. (2001). Summary Street : an intelligent tutoring system for improving student writing through the use of Latent Semantic Analysis (Doctoral dissertation, University of Colorado, 2001) Boulder, USA.

Stemler, S. (2004). A comparison of consensus, consistency, and measurement approaches to estimating interrater reliability. *Practical Assessment, Research & Evaluation*, 9 (4).

St.Pierre, M., & LaPlant, W. (1998). Issues in crosswalking content metadata standards. National Information Standards Organization (NISO). Bethesda (MD): NISO Press.

Su, X. (2002). A text categorization perspective for ontology mapping. Technical report, Trondheim, Norway: Norwegian University of Science and Technology.

Tillema, H. (2001). Portfolios as developmental assessment tools. *International journal of training and development,* 5 (2), 126 – 135.

Toyne, P. (1979). Educational credit transfer: feasibility study. London: Departement of Education and Science.

UNESCO. (2002). Forum on the impact of Open Courseware for higher education in developing countries. Final report. Paris:UNESCO.

Van Assche, F., Duval, E., Massart, D., Olmedilla, D., Simon, B., Sobernig, S., et al. (2006). Spinning interoperable applications for teaching & learning using the simple query interface. *Journal of Educational Technology & Society*, 9 (2), 51-67.

Van Bruggen, J., Sloep, P., Rosmalen, P. V., Brouns, F., Vogten, H., Koper, R., et al. (2004). Latent semantic analysis as a tool for learner positioning in learning networks for lifelong learning. *British Journal of Educational Technology*, 35 (6), 729-738.

Van Bruggen, J., Rusman, E., Giesbers, B., & Koper, R. (2006). Content based positioning in learning networks. *Proceedings of the 6th IEEE International conference on advanced learning technologies*, ICALT 2006, 5-7 July 2006 (pp. 366-368). Kerkrade, The Netherlands: IEEE Computer Society.

Van Bruggen, J., Kalz, M., & Joosten-Ten Brinke, D. (2009). Placement services for Learning Networks. In Koper, R. (Ed). Learning network services for professional development. Berlin:Springer.

Van der Vegt, W. (2009). Phplsa. Latent Semantic Analysis Extension for PHP.

Van Rosmalen, P., Sloep, P., Brouns, F., Kester, L., Kone, M., Koper, R., et al. (2006). Knowledge matchmaking in Learning Networks: Alleviating the tutor load by mutually connecting learning network users. *British Journal of Educational Technology*, 37 (6), 881-895.

Van Rosmalen, P. (2008). Supporting the tutor in the design and support of adaptive e-learning (Doctoral dissertation, Open University of the Netherlands, 2008), Heerlen, The Netherlands: Open University of the Netherlands.

Vuorikari, R., Ochoa, X. (2009). Exploratory Analysis of the Main Characteristics of Tags and Tagging of Educational Resources in a Multi-lingual Context. Special Issue of Social Information Retrieval for Technology Enhanced Learning. *Journal of Digital Information*,10 (2). Retrieved 2009-06-24, from http://journals.tdl.org/jodi/article/view/447/284

Vygotskii, L. S., & Cole, M. (1978). Mind in society: the development of higher psychological processes. Cambridge, Mass.: Harvard University Press.

Wiemer-Hastings, P. (1999). How latent is latent semantic analysis? In D. Thomas (Ed.) Proceedings *of the sixteenth international joint conference on artificial intelligence*, IJCAI 99, Stockholm, Sweden, July 31 - August 6, 1999 (pp. 932-937). 2 Volumes, 1450 pages, San Francisco, USA: Morgan Kaufmann.

Wiemer-Hastings, P., & Graesser, A. (2000). Select-a-Kibitzer: A computer tool that gives meaningful feedback on student compositions. *Interactive Learning Environments*, 8 (2), 149-169.

Wild, F., Stahl, C., Stermsek, G., & Neumann, G. (2005). Parameters driving effectiveness of automated essay scoring with LSA. In *Proceedings of the 9th CAA Conference*, June 5-6, 2005, Loughborough: Loughborough University.

Wild, F., Kalz, M., Van Bruggen, J., & Koper, R. (Eds.) (2007). Proceedings of the first European workshop on latent semantic analysis in technology enhanced learning (Vol. 12). Heerlen, The Netherlands: Open University of the Netherlands.

Witten, I., & Frank, E. (2000). Data mining. Practical machine learning tools and techniques (Morgan Kaufmann Series in Data Management Systems): (2 ed., p. 560). Morgan Kaufmann.

Woelk, D. (2002). E-Learning, Semantic web services and competency ontologies. In P. Barker & S. Rebelsky (Eds.), *Proceedings of World Conference on Educational Multimedia, Hypermedia and Telecommunications* 2002 (pp. 2077-2078). Chesapeake (VA): AACE.

Wolfe, M. B., Schreiner, M. E., Rehder, B., Laham, D., Foltz, P. W., Kintsch, W., et al. (1998). Learning from text: Matching readers and texts by Latent Semantic Analysis. *Discourse Processes*, 25, 309-336.

Yu, C., Cuadrado, J., Ceglowski, M., & Payne, J. S. (2002). Patterns in unstructured data. Discovery, aggregation, and visualization. A presentation to the Andrew W. Mellon Foundation. Retrieved 1st of Februrary, 2009 from http://www.knowledgesearch.org/lsi/cover_page.htm

Zampa, V., & Lemaire, B. (2002). Latent Semantic Analysis for user modeling. *Journal of Intelligent Information Systems*, 18 (1), 15-30.

Zeimpekis, D., & Gallopoulos, E. (2006). TMG: A MATLAB toolbox for generating term-document matrices from text collections. In J. Kogan, C. Nicholas, & M. Teboulle (Eds.), Grouping multidimensional data: Recent advances in Clustering (pp. 187-210). Berlin: Springer.

Zha, H., & Simon, H. (1999). On updating problems in latent semantic indexing. *SIAM Journal on Scientific Computing*, 21 (2), 782-791.

Summary

In this thesis a new bottom-up method to approximate the prior knowledge of learners via the similarity of content in their portfolios and content in a target study programme is evaluated. This process is referenced as the "placement problem" in this thesis. For this specific problem a long term research agenda is defined that is driven by the available data in a learning network or APL procedure. We differentiate in chapter 2.1 between three different situations shown in the following figure.

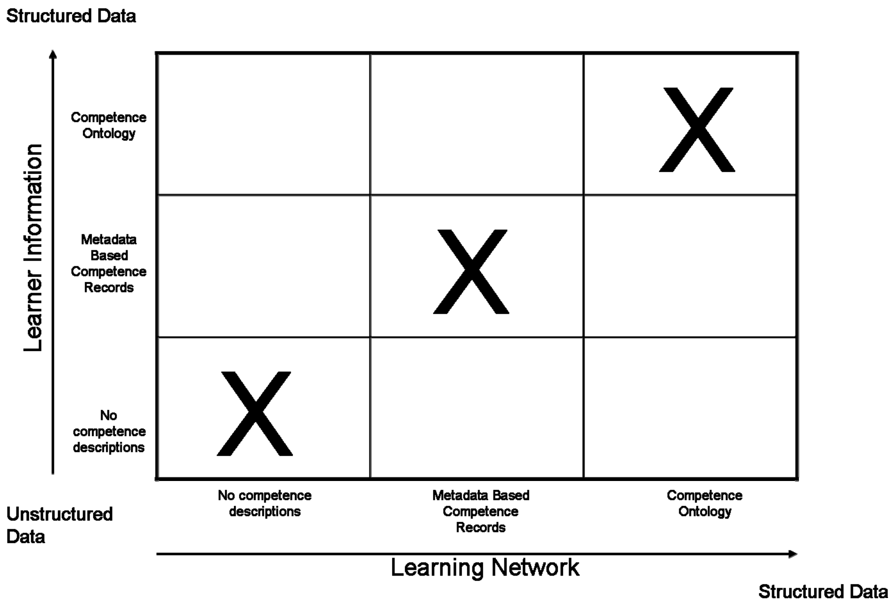

Figure 1: The Positioning Situations Matrix

While a matching of the current place in a study programme can be easily defined when very detailed ontological descriptions are available about the

learner and the learning content in a chosen study programme more reasoning is needed in the direction of unstructured data. In this thesis we pick the most complicated case where no competence descriptions of the learner and the learning network are available. For this situation we propose and evaluate a method to take content as a proxy for prior knowledge.
Based on a review of literature about eAssessment, learning technology standards, APL and electronic portfolios we have develop a model for a placement support web-service. This service should allow the implementation in different context and it should be possible to refine it to local thresholds expert opinions. For this purpose we expect that a "training phase" is needed before the service can reach its optimal performance. The service with the different layers is shown in the following figure.

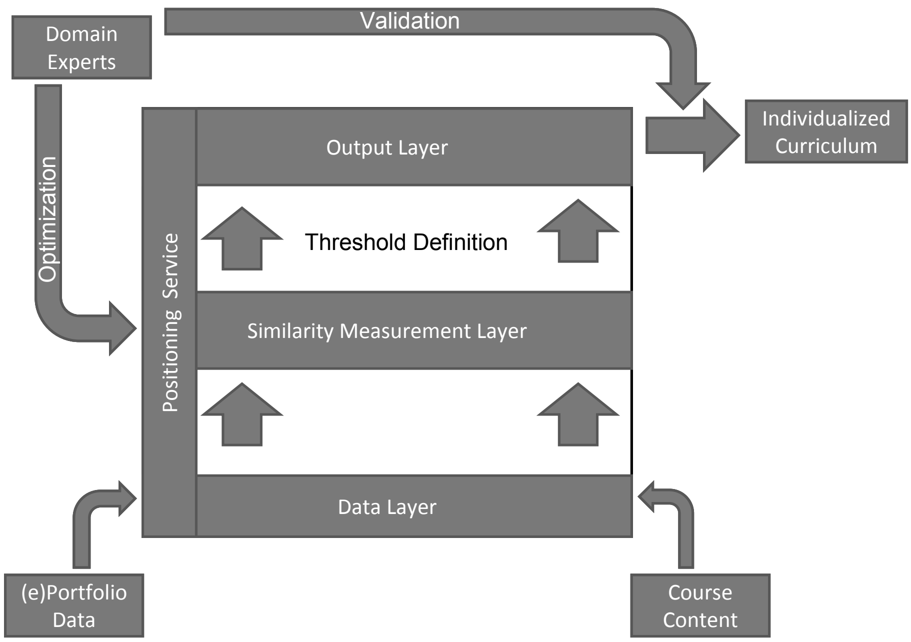

Figure 2: Placement Support Service

To evaluate this model a validation strategy is proposed in chapter 3.1 that is based on the aspects of sensitivity, reliability, validity and fit in APL procedures. In the first empirical study several issues are addressed that stem from the application of small scale corpora which are likely to be available in learning networks and APL procedures.

In Chapter 3.2 we present a study that addresses specific problems related to the application of the vector space model and dimensionality reduction to small scale corpora. In this study we have developed guidelines to reach sufficient discrimination and to identify the ideal number of dimensions to retain. Two corpora have been used in the study. While one corpus was constructed and prepared for our study we have used another corpus from course material from the OUNL. We could present in the study a method to identify the ideal number of singular values to retain based on the connection between singular values and the variance they account for in the Term-Document matrix (see figure 3).

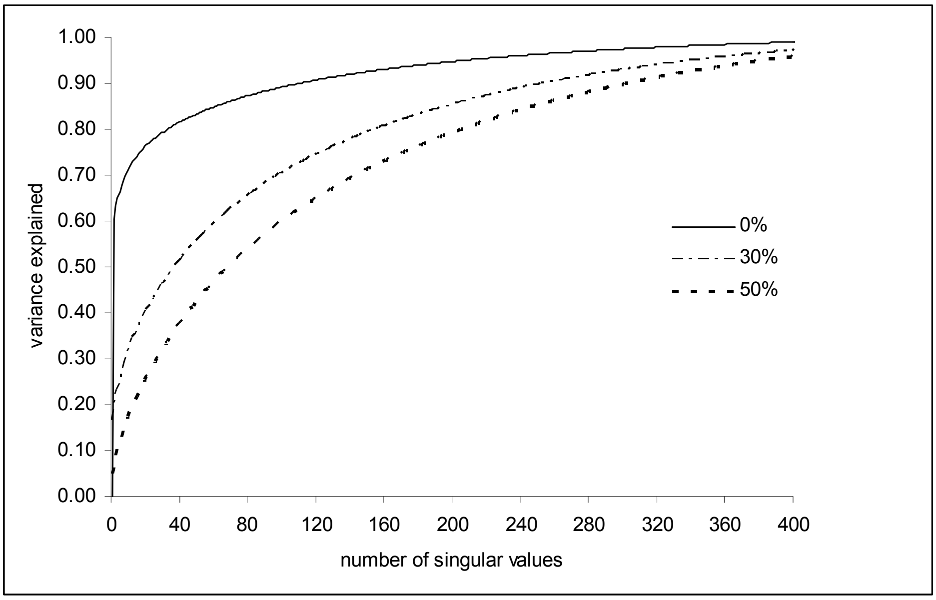

Figure 3. Number of singular values and variance explained under different stopping strategies for corpus 1

We have identified an interval between 70 % and 80 % of the variance of the data to retain as the area where we could filter enough noise while reaching at the same time sufficient discrimination with the use of a stopword list between 30 % and 50 %. The effect of using a 50 % stopword on the study material corpus is depicted in figure 4.

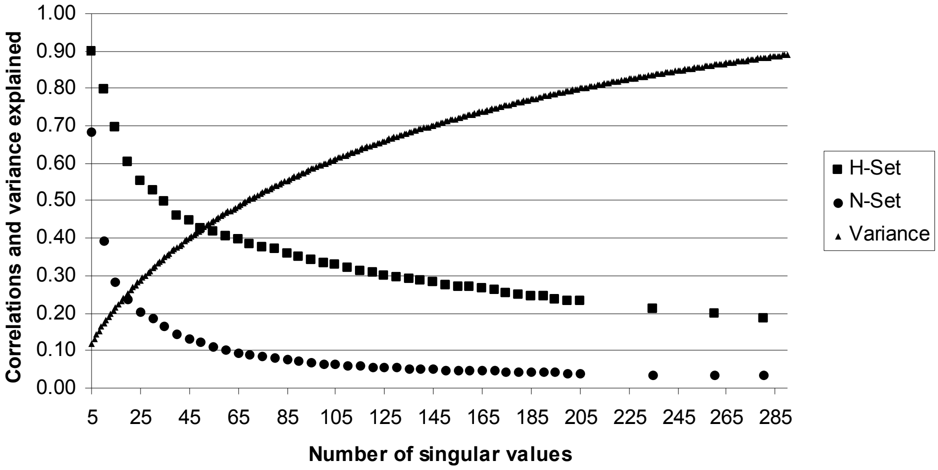

Figure 4 Query set 2 correlations and explained variance under 50% stopping for corpus 2

We demonstrated that we were able to detect correlations within homogenous content sets whilst discriminating between different sets.
In the second empirical study we could confirm results from the first study. The use of a 50% stopping strategy in combination with our dimensionality estimation approach has delivered enough discrimination to study the semantic similarity of the learner documents to the target learning units. Furthermore we have analyzed the impact of different weighting functions to the performance of our model for placement support.

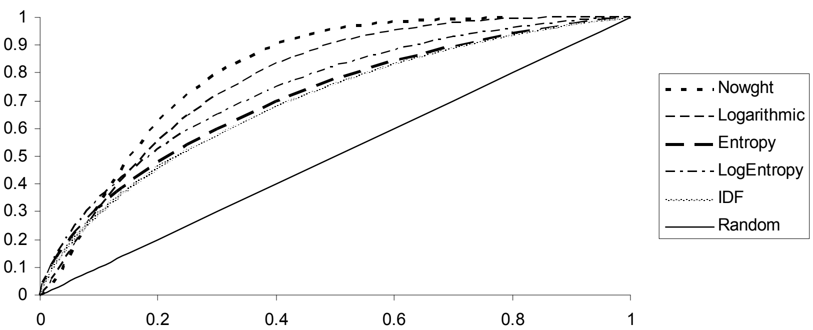

Figure 5 Receiver operating characteristics curve (ROC curve) for LSA with different weighting parameters

We could show that weighting did not improve the performance of our model although the application of no weighting did not show a significant difference in comparison to the application of weithing.

The performance assessment of our classification model in study 2 has shown that we could successfully classify 85% of the cases right.

Table 1: Classification results Human ratings vs. LSA ratings (n=504)

		Human rating	
		Relevant	Irrelevant
LSA rating	Relevant	28	40
	Irrelevant	12	424

In addition the confusion matrix in table 1 shows that we could reach our self-set target to reach a number of false and true positive cases under 10%. Based on the calculation of the Area und the ROC-curve (AUC) we could summarize that we have reached a good prediction model for classifying relevant and irrelevant document in APL procedures.

These results have lead to the development of two prototypes within the project. Figure 6 shows the architecture of the Placement Web-Service Prototype.

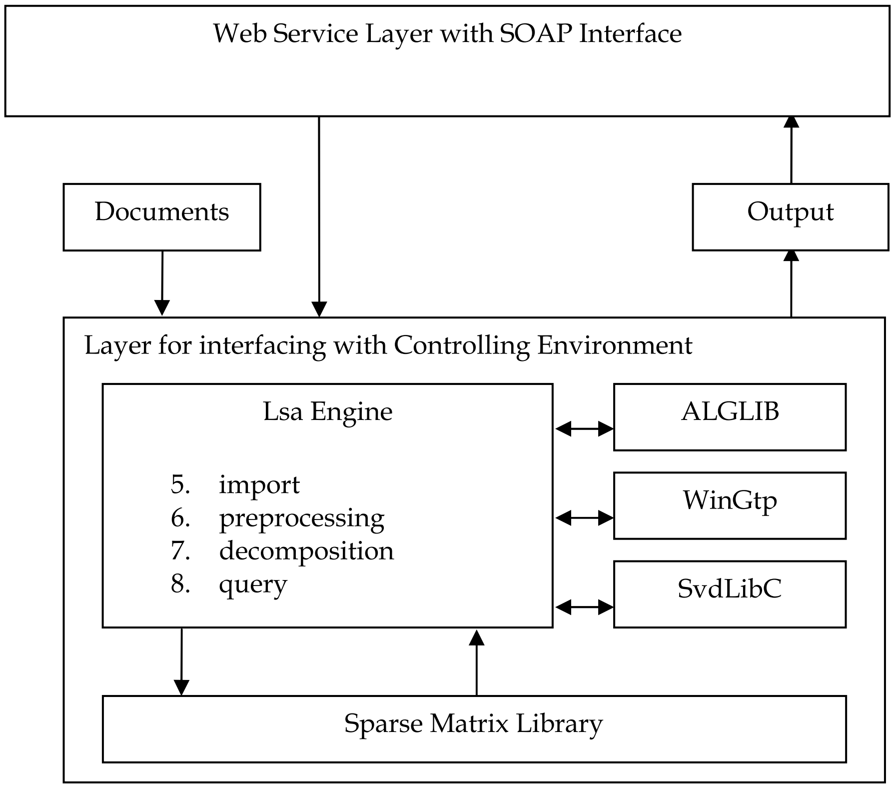

Figure 6: Architecture of the PWSP

This prototype has been evaluated from a technical perspective which has lead to several improvement ideas for the future. Furthermore we have developed a second prototype that should enable us in future research to address the problems of corpus construction and the application of other techniques from data and text mining to the placement problem in APL procedures and learning networks.

Samenvatting

In dit proefschrift wordt een nieuwe methode geëvalueerd om de eerder verworven kennis van lerenden te benaderen via de gelijkheid van portfolio-inhouden en de inhoud van een doelbewust studieprogramma. Aan dit proces wordt in dit proefschrift gerefereerd als 'plaatsingsprobleem'. Voor dit specifiek probleem is een lange-termijn onderzoeksagenda gedefinieerd, die wordt gebaseerd op de beschikbare data in een leernetwerk of op EVC-procedures (Erkenning van Verworven Competenties). We maken in hoofdstuk 2.1 verschil tussen drie verschillende situaties. Zie hiertoe onderstaande figuur.

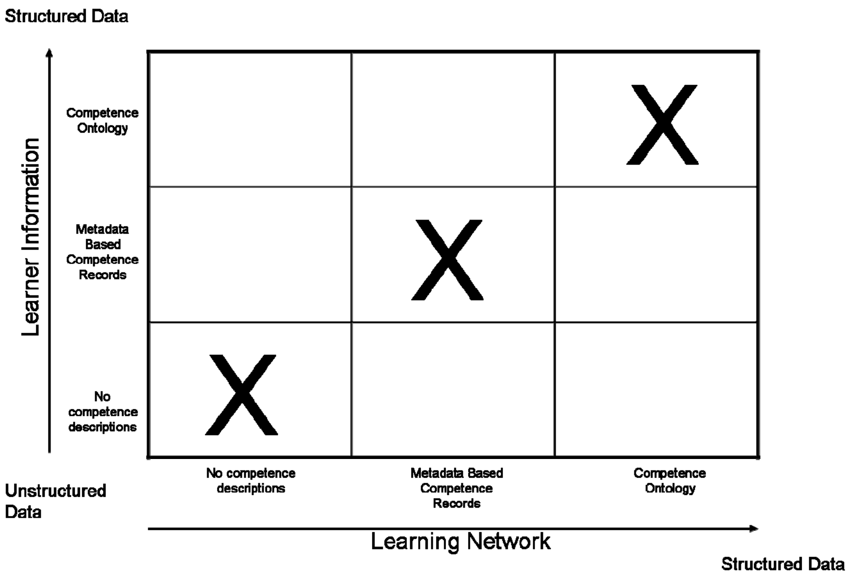

figuur 1: De matrix van plaatsingssituaties.

Terwijl de overeenkomst van de huidige plaats in een studieprogramma gemakkelijk bepaald kan worden als gedetailleerde ontologische beschrijvingen beschikbaar zijn over de lerenden en de leerinhoud van een gekozen studieprogramma, is er meer beredenering nodig als het ongestructureerde data betreft. In dit proefschrift kiezen we het meest gecompliceerde voorbeeld, waarbij geen beschrijving aanwezig is van de leerinhoud en de lerenden. Hiertoe evalueren wij een methode, die de inhoud van de portfolio's als basis neemt voor eerder verworven kennis.

Wij hebben een model ontwikkeld voor een plaatsingsondersteunende webservice, gebaseerd op literatuur over eAssessment, leertechnologie-standaarden, EVC en elektronische portfolio's. Deze service moet de implementatie in verschillende contexten en op verschillende niveaus mogelijk maken. We verwachten dat voor elke implementatie een 'trainingsfase' nodig zal zijn, voordat de service zijn optimale prestatie bereikt. Zie in onderstaande figuur de service bestaande uit verschillende lagen.

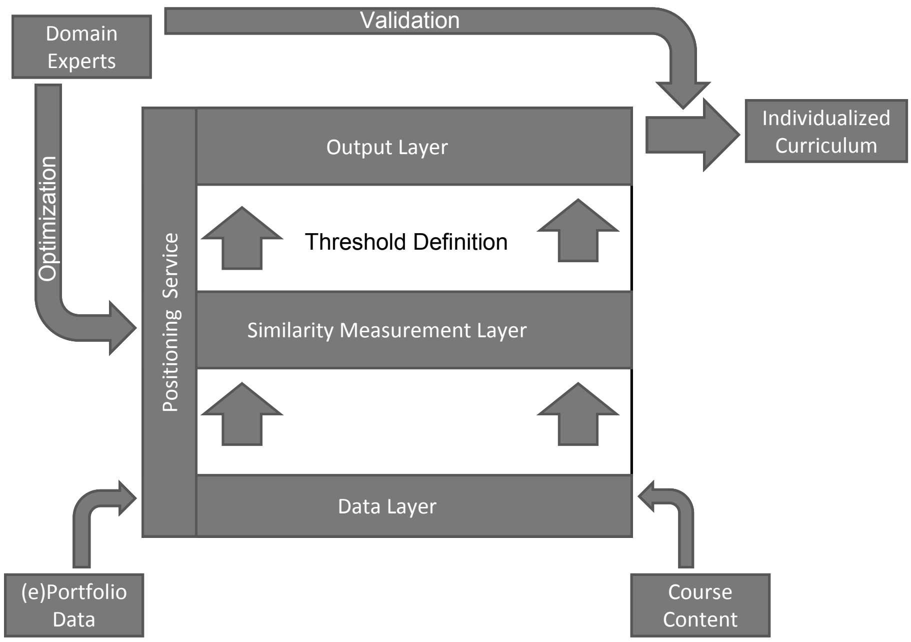

Figuur 2: Service voor Plaatsingsondersteuning

Om dit model te evalueren wordt in hoofdstuk 3.1 een validatiestrategie voorgesteld, die gebaseerd is op de volgende aspecten: gevoeligheid, betrouwbaarheid, validiteit en inpasbaarheid in de EVC-procedures. In de eerste empirische studie worden verschillende problemen behandeld, die voortkomen uit het toepassen van kleine corpora die naar alle waarschijnlijkheid verwacht kunnen worden in leernetwerken en EVC-procedures.

In hoofdstuk 3.2 beschrijven we een studie waarin specifieke problemen behandeld worden, die gerelateerd zijn aan het toepassen op kleine corpora van het vectorruimte-model en een dimensionale reductie. In deze studie hebben wij richtlijnen ontwikkeld om voldoende onderscheidingsvermogen te halen en het ideale aantal te behouden dimensies te identificeren. Twee corpora zijn in deze studie gebruikt.
Terwijl het ene corpus voor onze studie werd ontworpen en aangepast, hebben we een tweede corpus samengesteld uit cursusmateriaal van de OUNL. In deze studie werd een methode beschreven, die het ideale aantal te behouden singuliere waarden identificeert, gebaseerd op de relatie tussen singuliere waarden en de variantie in de term-documenten matrix die hiermee verklaard kan worden (zie figuur 3.)

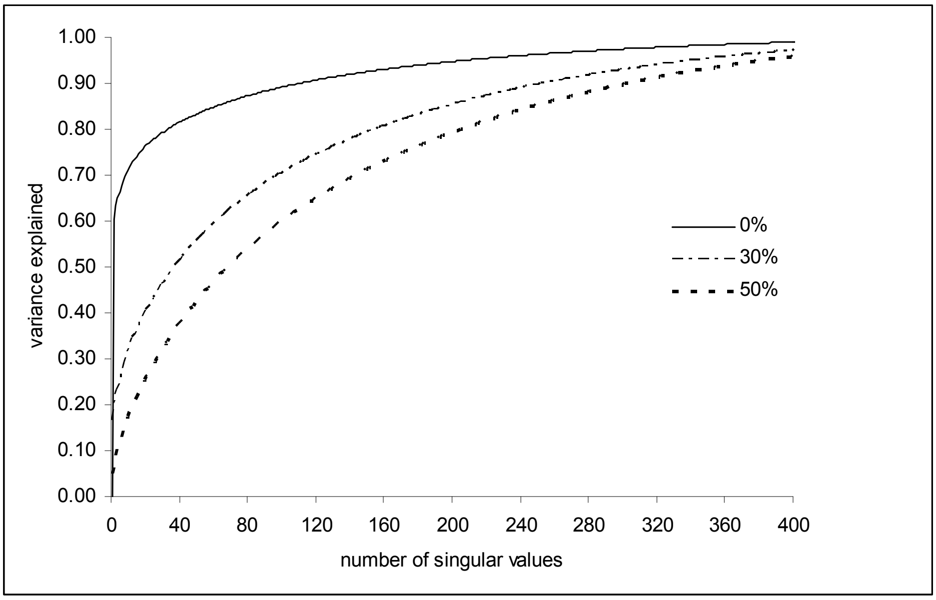

figuur 3. Aantal singuliere waarden en variantie verklaard voor het gebruik van verschillende stopping strategieën voor corpus 1.

We hebben een interval tussen 70% en 80% van de variantie in de data geïdentificeerd als het bereik dat behouden moest worden om er genoeg ruis uit te filteren terwijl tegelijkertijd genoeg discriminatie bereikt werd met een stopwoordenlijst tussen 30% en 50%. Het effect van het gebruik van meer dan 50% stopwoorden op het studiemateriaal corpus is uitgebeeld in figuur 4.

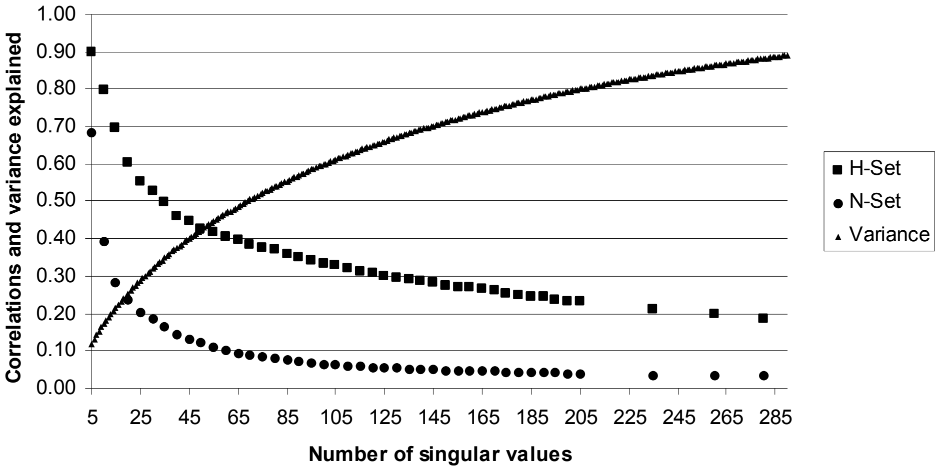

figuur 4: correlaties for query set 2 and verklaarde variantie met een 50 % stopping strategie.

We hebben aangetoond, dat we in staat waren verbanden vast te stellen tussen homogene inhouden, terwijl tegelijkertijd onderscheid werd gemaakt tussen heterogene inhouden.
In de tweede empirische studie konden we de resultaten van de eerste studie bevestigen. Het gebruik van een 50% stopwoordenlijst in combinatie met onze benadering die het aantal dimensies inschat, heeft voldoende onderscheid opgeleverd om de semantische gelijkheid van de portfolio's en leerinhouden te bestuderen. Verder hebben we de invloed van verschillende gewichtsfuncties op de prestatie van ons model voor plaatsingsondersteuning geanalyseerd.

We konden aantonen, dat het toepassen van gewichtsfuncties de prestatie van ons model niet verbeterde, maar ook niet significant verslechterde. De toetsing van de prestatie van ons classificatiemodel in studie 2 heeft aangetoond, dat we met succes 85% van de onderzoeksonderwerpen juist konden classificeren.

Tabel 1: classificatieresultaten domein experts vs. LSA experts (n=504)

		Human rating	
		Relevant	Irrelevant
LSA rating	Relevant	28	40
	Irrelevant	12	424

Verder toont de confusion matrix in tabel 1, dat we ons zelf-gestelde doel konden behalen om een aantal fout positief en fout negatief geclassificeerde gevallen onder 10% te bereiken.
Gebaseerd op de berekening van de Area under the ROC-curve (AUC) konden we aantonen dat we een goed voorspelbaar model hebben ontwikkeld voor de classificatie van relevante en irrelevante documenten in EVC-procedures. De resultaten hebben geleid tot de ontwikkeling van twee prototypes binnen het project. Figuur 5 laat de architectuur van het Placement Web-Service Prototype zien.

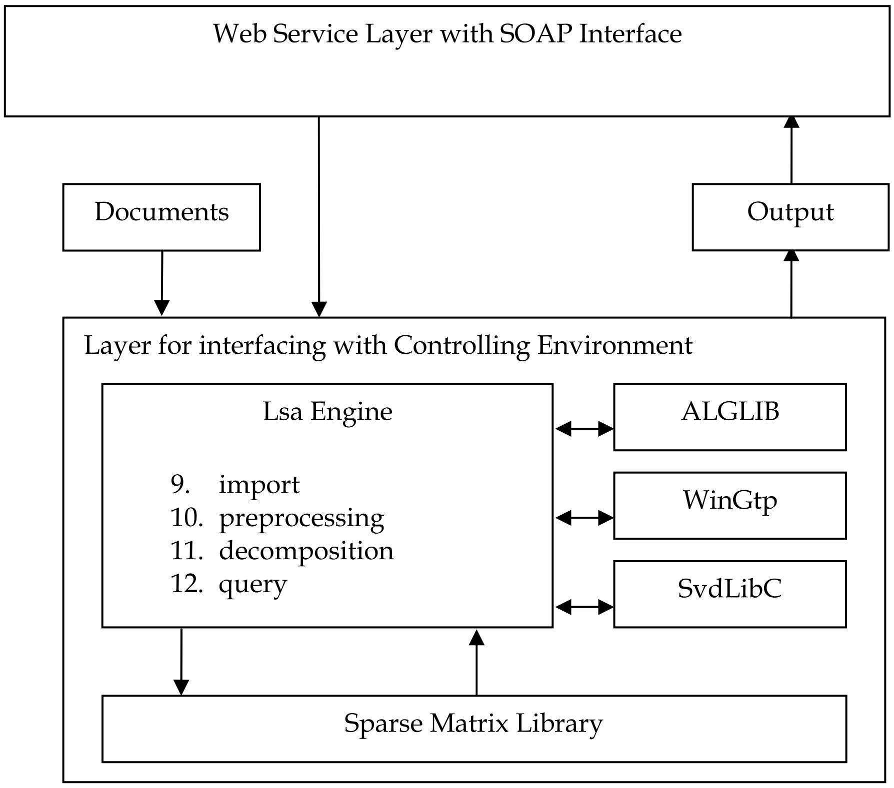

figuur 5: Architectuur van het PWSP

Dit prototype is vanuit een technisch perspectief geëvalueerd. Dit heeft geleid tot verschillende verbeterpunten voor de toekomst. Verder hebben we een tweede prototype ontwikkeld, dat het voor toekomstig onderzoek mogelijk moet maken om de problemen te behandelen met betrekking tot het samenstellen van corpora voor en de applicatie van andere technieken van data-en text-mining op het plaatsingsprobleem in EVC-procedures en leernetwerken.

Acknowledgements

This thesis is a marking point for me for the last 8 years in which I had the opportunity to focus on technology-enhanced learning form different perspectives, with different key aspects of activity and at least in two different countries. First of all I would like to thank my promoter Rob Koper who inspired and stimulated this work with his research on learning networks. Thank you Rob for the trust and unquestioned support throughout the whole project. This project would have not been possible without the continuous and reliable support of my daily supervisor Jan van Bruggen. Jan, I really enjoyed the numerous discussions and meetings we had and you offered me the right mix between supervision and inspiration. Thank you for your flexibility and openness in all respects.

The Center for Learning Sciences and Technologies (CELSTEC) was a very good place to conduct my research because it has offered me the right mix between structure and creative freedom, support and openness for creativity. Thanks to all my great colleagues there. A special thanks goes out to all my colleagues in the so called "Apecage". Although it was sometimes challenging to work in such a large office I really appreciated the open atmosphere and the mutual support there.

Some other colleagues need to be mentioned personally because they supported this project in the one or the other way: Hendrik Drachsler was always a source of constructive feedback and a great support during the technology development phases. Peter van Rosmalen had always an open door and was able to inspire and support the project by his unique way of questioning the questions. Mieke Haemers was always helpful in many organizational aspects and did a great job in proofreading some of my articles and this thesis. At the very beginning of the project I was very happy to get a very warm welcome and support by the members of the former "positioning project group" namely Bas Giesbers, Ellen Rusman and Howard Spoelstra. Wim van der Vegt supported the development of the technological artifacts.

Wim Waterink and Jannes Eshuis from the psychology faculty of the Open University of the Netherlands allowed me to collect data from their students and supported the empirical studies. In addition they were always open for questions. Thanks for your participation and support. Several colleagues from CELSTEC supported the study by providing me with material from their professional portfolios – thanks a lot for this.

Dimitris Zeimpekis from the University of Patras in Greece supported the project with his Text Matrix Generator (TMG) and he always reacted very fast on my questions and suggestions for improvement. Fridolin Wild (The Open University, UK) supported this project in several ways. During my second year we have organized the First European Workshop on Semantic Analysis in Technology-Enhanced Learning. This was a great opportunity to see the work of others in the field and I was very happy about the participation of some of the forerunners of the field like Thomas Landauer, Art Graesser, Dian Martin and Benoit Lemaire.

This thesis has been funded with support of the European Commission via the Integrated Project TENCompetence from the 6th Framework Programme, priority IST/Technology Enhanced Learning. Contract 027087 (www.tencompetence.org). I am very grateful for having been able to work in such a large European project which offers a lot of challenges and opportunities. Thanks to all colleagues in the project that cooperated with me and supported my project.

I am very grateful for the reliable support of my parents Monika & Anton Kalz who supported me in so many ways. Thanks so much! A very large contribution to this project comes from my family. My wife Verena has constantly supported me and offered me a source of inspiration and distraction that was needed to have enough energy for this endeavor. Thank you! My daughters Vivien & Lilly could also contribute in their own unique way to this project when they confronted the rationalistic way of thinking scientifically with children's logic. I love you girls.

This thesis was supported with approximately 2500 cups of finest Dutch coffee and the background noise was most of the time provided by

Monkeyradio, Motor.FM or ByteFM. Overall, during the time of writing my thesis 12177 tracks have been added to my music profile in Last.fm. Besides the music much inspiration has been provided by a large number of people who have decided to share their knowledge by means of the internet e.g. in form of making their publications available or by publishing their thoughts and reflection in the web via weblogs or other tools. Without naming all of them I am grateful for this altruistic behavior and I hope that this community will grow to realize a world-wide open space of knowledge sharing that helps us to address the urgent problems of the 21st century.

Curriculum Vitae

Marco Kalz, born 7th of December 1975 in Aachen, Germany has a state board examination and a teacher degree for special education and German language for secondary schools. After working as a consultant for content syndication he started a master programme for media education and instructional design which he finished in 2003. In 2002 he was involved in the development of a computer based training for geographical information systems and in the implementation of a university-wide learning environment on the basis of open source software.

In 2002 Marco started his academic career as a research assistant in the group for media education and knowledge management of the University of Duisburg where he participated in a national research project focusing on the use of mobile devices for teaching and learning in higher education. In 2004 he joined the new established educational technology research group at Fernuniversität in Hagen. In this time Marco was involved in the conception and implementation of a new bachelor programme for educational science and a new master programme eEducation. His research there was dedicated to the use of social software for self-directed competence development.

In February 2006 Marco joined the Open University of the Netherlands where he worked in the European funded Integrated Project TENCompetence (6th framework programme). Marco has more than 30 publications in the field of technology-enhanced learning and he edited several books, conference and workshop proceedings and special issues of internationally recognized journals. Marco serves a reviewer and programme committee member for international conferences and journals. In his current work Marco is focusing on new methods for formative assessment of learners and the use of social software for lifelong self-directed competence development. More information and a full list of publications and activities are available at http://www.marcokalz.de.

Publication List

Peer Reviewed Publications

Kalz, M., van Bruggen, J., Giesbers, B., Eshuis, J., Waterink, W. & Koper, R. (2008). A Model for New Linkages for Prior Learning Assessment. Campus Wide Information Systems. 25, 4. 233 -243.

Kalz, M., Drachsler, H., van Bruggen, J., Hummel, H., & Koper, R. (2008). Wayfinding Services for Open Educational Practices. International Journal of Emerging Technologies in Learning, 3(2).

Kalz, M, van Bruggen, J., Rusmann, El, Giesbers, B. & Koper, R. (2007). Positioning of Learners in Learning Networks with Content-Analysis, Metadata and Ontologies. Interactive Learning Environments,15, 191-200.

Kalz, Marco: Research and Development of a Positioning Service for Learning Networks for Lifelong Learning. In: Klamma, R., & Maillet, K. (2006). Proceedings of the Doctoral Consortium of the First European Conference on Technology Enhanced Learning. S. 26 - 31. Kreta, Greece.

Kalz, M., van Bruggen, J., Rusmann, E., Giesbers, B., & Koper, R. (2006). Positioning of Learners in Learning Networks with Content Analysis, Metadata and Ontologies. Proceedings of the International Workshop in Learning Networks for Lifelong Competence Development. March 30th-31st, Sofia, Bulgaria: TENCompetence. 77-81.

Conference Contributions

Kalz, M., Berlanga, A., Van Rosmalen, P., Stoyanov, S., Van Bruggen, J., & Koper, R. (2009). Semantic Networks as Means for Goal Directed Formative Feedback. In V. Hornung-Prähauser & M. Luckmann (Eds.), Proceedings of the Edumedia Conference "Creativity and Innovation Competencies in the Web" (pp. 88-95). May, 4-5, 2009, Salzburg, Austria.

Berlanga, A. J., Kalz, M., Stoyanov, S., Van Rosmalen, P., Smithies, A., & Braidman, I. (2009). Using Language Technologies to Diagnose Learner´s Conceptual Development. Proceedings of the IEEE Internacional Conference on Advanced Learning Technologies.

Kalz, M., Drachsler, H., Van Bruggen, J., Hummel, H.G.K., & Koper, R. (2007). Positioning and Navigation: Services for Open Educational Practices. In Auer, M. E. (Ed): International Conference on Computer Aided Learning 2007 [CD ROM]. Kassel: Kassel University Press.

Kalz, M., Van Bruggen, J., Giesbers, B., Waterink, W., Eshuis, J. & Koper, E.J.R. (2008). A New Linkage for Prior Learning Assessment. In European Institute for E-Leatrning (Ed.) Proceedings of the conference ePortfolio2007 - Employability and Lifelong Learning in the Knowledge Society. Maastricht, The Netherlands.

Kalz, M., van Bruggen, J., Giesbers, B., & Koper, E.J.R. (2007). Prior Learning Assessment with Latent Semantic Analysis. In Wild, F., Van Bruggen, J., Kalz, M. & Koper, R. (Eds.). Proceedings of the First European Workshop on Latent Semantic Analysis in Technology Enhanced Learning (pp. 24-25). Heerlen, the Netherlands:Open University of the Netherlands.

Book chapters

Van Bruggen, J., Kalz, M., & Joosten-Ten Brinke, D. (2009). Placement Services for Learning Networks. In Koper, R. (Ed.). Learning Network Services for Professional Development. Springer 2009.

Van der Vergt, W., Kalz, M., Giesbers, B., Wild, F., & Van Bruggen, J. (2009). Tools and Techniques for Placement Experiments. In Koper, R. (Ed.). Learning Network Services for Professional Development. Springer 2009.

Kalz, M., van Bruggen, J., Giesbers, B., Rusman, E., Eshuis, J., & Waterink, W. (2009). A Validation Scenario for a Placement Service in Learning Networks. In Koper, R. (Ed.). Learning Network Services for Professional Development. Springer 2009.

Schaffert, S., & Kalz, M. (2009). Persönliche Lernumgebungen: Grundlagen, Möglichkeiten und Herausforderungen eines neuen Konzepts. In K. Wilbers & A. Hohenstein (Hrsg.), Handbuch E-Learning (Gruppe 5, Nr. 5.16, pp. 1-24). Köln, Germany: Deutscher Wirtschaftsdienst (Wolters Kluwer Deutschland), 27. Erg.-Lfg. Januar 2009.

Hornung-Prähauser, V., Luckmann, M., & Kalz, M. (2008). Eine Landkarte internetgestützten Lernens. In: Hornung-Prähauser, V., Luckmann, M., & Kalz, M. (2008). Selbstorganisiertes Lernen im Internet. Einblick in die Landschaft der webbasierten Bildungsinnovationen. Studienverlag:Innsbruck. ISBN 978-3-7065-4641-6. S. 13 - 25.

Kalz, Marco; Specht, Marcus; Klamma, Ralf; Chatti, Mohamed Amine; Koper, Rob: Kompetenzentwicklung in Lernnetzwerken für das Lebenslange Lernen. Schwarz, Christine, Kindt, Michael & Dittler, Ulrich (2007). Online Communities als soziale Systeme. Hannover: Waxmann Verlag. S.181 – 197.

Edited books and conference proceedings

Hornung-Prähauser, V., Luckmann, M., & Kalz, M. (2008). Selbstorganisiertes Lernen im Internet. Einblick in die Landschaft der webbasierten Bildungsinnovationen. Studienverlag:Innsbruck. ISBN 978-3-7065-4641-6.

Wild, F., Kalz, M. & Palmér, M. (2008). Proceedings of the First International Workshop on Mashup Personal Learning Environments. September 17, Maastricht, The Netherlands: CEUR Workshop Proceedings, ISSN 1613-0073. Available at http://ceur-ws.org/Vol-388.

Kalz, M., Koper, R., Hornung-Prähauser, V., & Luckmann, M. (Eds.) (2008). Proceedings of the 1st Workshop on Technology Support for Self-Organized Learners. June, 2-3, 2008, Salzburg, Austria: CEUR Workshop Proceedings, ISSN 1613-0073. Available at http://ceur-ws.org/Vol-349.

Wild, F., Van Bruggen, J., Kalz, M. & Koper, R. (Eds.) (2007). Proceedings of the First European Workshop on Latent Semantic Analysis in Technology Enhanced Learning. Heerlen, the Netherlands:Open University of the Netherlands.

SIKS Dissertatiereeks

====
1998
====

1998-1 Johan van den Akker (CWI)
DEGAS - An Active, Temporal Database of Autonomous Objects

1998-2 Floris Wiesman (UM)
Information Retrieval by Graphically Browsing Meta-Information

1998-3 Ans Steuten (TUD)
A Contribution to the Linguistic Analysis of Business Conversations within the Language/Action Perspective

1998-4 Dennis Breuker (UM)
Memory versus Search in Games

1998-5 E.W.Oskamp (RUL)
Computerondersteuning bij Straftoemeting

====
1999
====

1999-1 Mark Sloof (VU)
Physiology of Quality Change Modelling; Automated modelling of Quality Change of Agricultural Products

1999-2 Rob Potharst (EUR)
Classification using decision trees and neural nets

1999-3 Don Beal (UM)
The Nature of Minimax Search

1999-4 Jacques Penders (UM)
The practical Art of Moving Physical Objects

1999-5 Aldo de Moor (KUB)
Empowering Communities: A Method for the Legitimate User-Driven Specification of Network Information Systems

1999-6 Niek J.E. Wijngaards (VU)
Re-design of compositional systems

1999-7 David Spelt (UT)
Verification support for object database design

1999-8 Jacques H.J. Lenting (UM)
Informed Gambling: Conception and Analysis of a Multi-Agent Mechanism for Discrete Reallocation.

====
2000
====

2000-1 Frank Niessink (VU)
Perspectives on Improving Software Maintenance

2000-2 Koen Holtman (TUE)
Prototyping of CMS Storage Management

2000-3 Carolien M.T. Metselaar (UVA)
Sociaal-organisatorische gevolgen van kennistechnologie;
een procesbenadering en actorperspectief.

2000-4 Geert de Haan (VU)
ETAG, A Formal Model of Competence Knowledge for User Interface Design

2000-5 Ruud van der Pol (UM)
Knowledge-based Query Formulation in Information Retrieval.

2000-6 Rogier van Eijk (UU)
Programming Languages for Agent Communication

2000-7 Niels Peek (UU)
Decision-theoretic Planning of Clinical Patient Management

2000-8 Veerle Coup, (EUR)
Sensitivity Analyis of Decision-Theoretic Networks

2000-9 Florian Waas (CWI)
Principles of Probabilistic Query Optimization

2000-10 Niels Nes (CWI)
Image Database Management System Design Considerations, Algorithms and Architecture

2000-11 Jonas Karlsson (CWI)
Scalable Distributed Data Structures for Database Management

====
2001
====

2001-1 Silja Renooij (UU)
Qualitative Approaches to Quantifying Probabilistic Networks

2001-2 Koen Hindriks (UU)
Agent Programming Languages: Programming with Mental Models

2001-3 Maarten van Someren (UvA)
Learning as problem solving

2001-4 Evgueni Smirnov (UM)
Conjunctive and Disjunctive Version Spaces with Instance-Based Boundary Sets

2001-5 Jacco van Ossenbruggen (VU)
Processing Structured Hypermedia: A Matter of Style

2001-6 Martijn van Welie (VU)
Task-based User Interface Design

2001-7 Bastiaan Schonhage (VU)
Diva: Architectural Perspectives on Information Visualization

2001-8 Pascal van Eck (VU)
A Compositional Semantic Structure for Multi-Agent Systems Dynamics.

2001-9 Pieter Jan 't Hoen (RUL)
Towards Distributed Development of Large Object-Oriented Models, Views of Packages as Classes

2001-10 Maarten Sierhuis (UvA)
Modeling and Simulating Work Practice
BRAHMS: a multiagent modeling and simulation language for work practice analysis and design

2001-11 Tom M. van Engers (VUA)
Knowledge Management:
The Role of Mental Models in Business Systems Design

====
2002
====

2002-01 Nico Lassing (VU)
Architecture-Level Modifiability Analysis

2002-02 Roelof van Zwol (UT)
Modelling and searching web-based document collections

2002-03 Henk Ernst Blok (UT)
Database Optimization Aspects for Information Retrieval

2002-04 Juan Roberto Castelo Valdueza (UU)
The Discrete Acyclic Digraph Markov Model in Data Mining

2002-05 Radu Serban (VU)
The Private Cyberspace Modeling Electronic
Environments inhabited by Privacy-concerned Agents

2002-06 Laurens Mommers (UL)
Applied legal epistemology; Building a knowledge-based ontology of the legal domain

2002-07 Peter Boncz (CWI)
Monet: A Next-Generation DBMS Kernel For Query-Intensive Applications

2002-08 Jaap Gordijn (VU)
Value Based Requirements Engineering: Exploring Innovative E-Commerce Ideas

2002-09 Willem-Jan van den Heuvel(KUB)
Integrating Modern Business Applications with Objectified Legacy Systems

2002-10 Brian Sheppard (UM)
Towards Perfect Play of Scrabble

2002-11 Wouter C.A. Wijngaards (VU)
Agent Based Modelling of Dynamics: Biological and Organisational Applications

2002-12 Albrecht Schmidt (Uva)
Processing XML in Database Systems

2002-13 Hongjing Wu (TUE)
A Reference Architecture for Adaptive Hypermedia Applications

2002-14 Wieke de Vries (UU)
Agent Interaction: Abstract Approaches to Modelling, Programming and Verifying Multi-Agent Systems

2002-15 Rik Eshuis (UT)
Semantics and Verification of UML Activity Diagrams for Workflow Modelling

2002-16 Pieter van Langen (VU)
The Anatomy of Design: Foundations, Models and Applications

2002-17 Stefan Manegold (UVA)
Understanding, Modeling, and Improving Main-Memory Database Performance

====
2003
====

2003-01 Heiner Stuckenschmidt (VU)
Ontology-Based Information Sharing in Weakly Structured Environments

2003-02 Jan Broersen (VU)
Modal Action Logics for Reasoning About Reactive Systems

2003-03 Martijn Schuemie (TUD)
Human-Computer Interaction and Presence in Virtual Reality Exposure Therapy

2003-04 Milan Petkovic (UT)
Content-Based Video Retrieval Supported by Database Technology

2003-05 Jos Lehmann (UVA)
Causation in Artificial Intelligence and Law - A modelling approach

2003-06 Boris van Schooten (UT)
Development and specification of virtual environments

2003-07 Machiel Jansen (UvA)
Formal Explorations of Knowledge Intensive Tasks

2003-08 Yongping Ran (UM)
Repair Based Scheduling

2003-09 Rens Kortmann (UM)
The resolution of visually guided behaviour

2003-10 Andreas Lincke (UvT)
Electronic Business Negotiation: Some experimental studies on the interaction between medium, innovation context and culture

2003-11 Simon Keizer (UT)
Reasoning under Uncertainty in Natural Language Dialogue using Bayesian Networks

2003-12 Roeland Ordelman (UT)
Dutch speech recognition in multimedia information retrieval

2003-13 Jeroen Donkers (UM)
Nosce Hostem - Searching with Opponent Models

2003-14 Stijn Hoppenbrouwers (KUN)
Freezing Language: Conceptualisation Processes across ICT-Supported Organisations

2003-15 Mathijs de Weerdt (TUD)
Plan Merging in Multi-Agent Systems

2003-16 Menzo Windhouwer (CWI)
Feature Grammar Systems - Incremental Maintenance of Indexes to Digital Media Warehouses

2003-17 David Jansen (UT)
Extensions of Statecharts with Probability, Time, and Stochastic Timing

2003-18 Levente Kocsis (UM)
Learning Search Decisions

====
2004
====

2004-01 Virginia Dignum (UU)
A Model for Organizational Interaction: Based on Agents, Founded in Logic

2004-02 Lai Xu (UvT)
Monitoring Multi-party Contracts for E-business

2004-03 Perry Groot (VU)
A Theoretical and Empirical Analysis of Approximation in Symbolic Problem Solving

2004-04 Chris van Aart (UVA)
Organizational Principles for Multi-Agent Architectures

2004-05 Viara Popova (EUR)
Knowledge discovery and monotonicity

2004-06 Bart-Jan Hommes (TUD)
The Evaluation of Business Process Modeling Techniques

2004-07 Elise Boltjes (UM)
Voorbeeldig onderwijs; voorbeeldgestuurd onderwijs, een opstap naar abstract denken, vooral voor meisjes

2004-08 Joop Verbeek(UM)
Politie en de Nieuwe Internationale Informatiemarkt, Grensregionale politiële gegevensuitwisseling en digitale expertise

2004-09 Martin Caminada (VU)
For the Sake of the Argument; explorations into argument-based reasoning

2004-10 Suzanne Kabel (UVA)
Knowledge-rich indexing of learning-objects

2004-11 Michel Klein (VU)
Change Management for Distributed Ontologies

2004-12 The Duy Bui (UT)
Creating emotions and facial expressions for embodied agents

2004-13 Wojciech Jamroga (UT)
Using Multiple Models of Reality: On Agents who Know how to Play

2004-14 Paul Harrenstein (UU)
Logic in Conflict. Logical Explorations in Strategic Equilibrium

2004-15 Arno Knobbe (UU)
Multi-Relational Data Mining

2004-16 Federico Divina (VU)
Hybrid Genetic Relational Search for Inductive Learning

2004-17 Mark Winands (UM)
Informed Search in Complex Games

2004-18 Vania Bessa Machado (UvA)
Supporting the Construction of Qualitative Knowledge Models

2004-19 Thijs Westerveld (UT)
Using generative probabilistic models for multimedia retrieval

2004-20 Madelon Evers (Nyenrode)
Learning from Design: facilitating multidisciplinary design teams

====
2005
====

2005-01 Floor Verdenius (UVA)
Methodological Aspects of Designing Induction-Based Applications

2005-02 Erik van der Werf (UM))
AI techniques for the game of Go

2005-03 Franc Grootjen (RUN)
A Pragmatic Approach to the Conceptualisation of Language

2005-04 Nirvana Meratnia (UT)
Towards Database Support for Moving Object data

2005-05 Gabriel Infante-Lopez (UVA)
Two-Level Probabilistic Grammars for Natural Language Parsing

2005-06 Pieter Spronck (UM)
Adaptive Game AI

2005-07 Flavius Frasincar (TUE)
Hypermedia Presentation Generation for Semantic Web Information Systems

2005-08 Richard Vdovjak (TUE)
A Model-driven Approach for Building Distributed Ontology-based Web Applications

2005-09 Jeen Broekstra (VU)
Storage, Querying and Inferencing for Semantic Web Languages

2005-10 Anders Bouwer (UVA)
Explaining Behaviour: Using Qualitative Simulation in Interactive Learning Environments

2005-11 Elth Ogston (VU)
Agent Based Matchmaking and Clustering - A Decentralized Approach to Search

2005-12 Csaba Boer (EUR)
Distributed Simulation in Industry

2005-13 Fred Hamburg (UL)
Een Computermodel voor het Ondersteunen van Euthanasiebeslissingen

2005-14 Borys Omelayenko (VU)
Web-Service configuration on the Semantic Web; Exploring how s emantics meets pragmatics

2005-15 Tibor Bosse (VU)
Analysis of the Dynamics of Cognitive Processes

2005-16 Joris Graaumans (UU)
Usability of XML Query Languages

2005-17 Boris Shishkov (TUD)
Software Specification Based on Re-usable Business Components

2005-18 Danielle Sent (UU)
Test-selection strategies for probabilistic networks

2005-19 Michel van Dartel (UM)
Situated Representation

2005-20 Cristina Coteanu (UL)
Cyber Consumer Law, State of the Art and Perspectives

2005-21 Wijnand Derks (UT)
Improving Concurrency and Recovery in Database Systems by Exploiting Application Semantics

====
2006
====

2006-01 Samuil Angelov (TUE)
Foundations of B2B Electronic Contracting

2006-02 Cristina Chisalita (VU)
Contextual issues in the design and use of information technology in organizations

2006-03 Noor Christoph (UVA)
The role of metacognitive skills in learning to solve problems

2006-04 Marta Sabou (VU)
Building Web Service Ontologies

2006-05 Cees Pierik (UU)
Validation Techniques for Object-Oriented Proof Outlines

2006-06 Ziv Baida (VU)
Software-aided Service Bundling - Intelligent Methods & Tools for Graphical Service Modeling

2006-07 Marko Smiljanic (UT)
XML schema matching -- balancing efficiency and effectiveness by means of clustering

2006-08 Eelco Herder (UT)
Forward, Back and Home Again - Analyzing User Behavior on the Web

2006-09 Mohamed Wahdan (UM)
Automatic Formulation of the Auditor's Opinion

2006-10 Ronny Siebes (VU)
Semantic Routing in Peer-to-Peer Systems

2006-11 Joeri van Ruth (UT)
Flattening Queries over Nested Data Types

2006-12 Bert Bongers (VU)
Interactivation - Towards an e-cology of people, our technological environment, and the arts

2006-13 Henk-Jan Lebbink (UU)
Dialogue and Decision Games for Information Exchanging Agents

2006-14 Johan Hoorn (VU)
Software Requirements: Update, Upgrade, Redesign - towards a Theory of Requirements Change

2006-15 Rainer Malik (UU)
CONAN: Text Mining in the Biomedical Domain

2006-16 Carsten Riggelsen (UU)
Approximation Methods for Efficient Learning of Bayesian Networks

2006-17 Stacey Nagata (UU)
User Assistance for Multitasking with Interruptions on a Mobile Device

2006-18 Valentin Zhizhkun (UVA)
Graph transformation for Natural Language Processing

2006-19 Birna van Riemsdijk (UU)
Cognitive Agent Programming: A Semantic Approach

2006-20 Marina Velikova (UvT)
Monotone models for prediction in data mining

2006-21 Bas van Gils (RUN)
Aptness on the Web

2006-22 Paul de Vrieze (RUN)
Fundaments of Adaptive Personalisation

2006-23 Ion Juvina (UU)
Development of Cognitive Model for Navigating on the Web

2006-24 Laura Hollink (VU)
Semantic Annotation for Retrieval of Visual Resources

2006-25 Madalina Drugan (UU)
Conditional log-likelihood MDL and Evolutionary MCMC

2006-26 Vojkan Mihajlovic (UT)
Score Region Algebra: A Flexible Framework for Structured Information Retrieval

2006-27 Stefano Bocconi (CWI)
Vox Populi: generating video documentaries from semantically annotated media repositories

2006-28 Borkur Sigurbjornsson (UVA)
Focused Information Access using XML Element Retrieval

====
2007
====

2007-01 Kees Leune (UvT)
Access Control and Service-Oriented Architectures

2007-02 Wouter Teepe (RUG)
Reconciling Information Exchange and Confidentiality: A Formal Approach

2007-03 Peter Mika (VU)
Social Networks and the Semantic Web

2007-04 Jurriaan van Diggelen (UU)
Achieving Semantic Interoperability in Multi-agent Systems: a dialogue-based approach

2007-05 Bart Schermer (UL)
Software Agents, Surveillance, and the Right to Privacy: a Legislative Framework for Agent-enabled Surveillance

2007-06 Gilad Mishne (UVA)
Applied Text Analytics for Blogs

2007-07 Natasa Jovanovic' (UT)
To Whom It May Concern - Addressee Identification in Face-to-Face Meetings

2007-08 Mark Hoogendoorn (VU)
Modeling of Change in Multi-Agent Organizations

2007-09 David Mobach (VU)
Agent-Based Mediated Service Negotiation

2007-10 Huib Aldewereld (UU)
Autonomy vs. Conformity: an Institutional Perspective on Norms and Protocols

2007-11 Natalia Stash (TUE)
Incorporating Cognitive/Learning Styles in a General-Purpose Adaptive Hypermedia System

2007-12 Marcel van Gerven (RUN)
Bayesian Networks for Clinical Decision Support: A Rational Approach to Dynamic Decision-Making under Uncertainty

2007-13 Rutger Rienks (UT)
Meetings in Smart Environments; Implications of Progressing Technology

2007-14 Niek Bergboer (UM)
Context-Based Image Analysis

2007-15 Joyca Lacroix (UM)
NIM: a Situated Computational Memory Model

2007-16 Davide Grossi (UU)
Designing Invisible Handcuffs. Formal investigations in Institutions and Organizations for Multi-agent Systems

2007-17 Theodore Charitos (UU)
Reasoning with Dynamic Networks in Practice

2007-18 Bart Orriens (UvT)
On the development an management of adaptive business collaborations

2007-19 David Levy (UM)
Intimate relationships with artificial partners

2007-20 Slinger Jansen (UU)
Customer Configuration Updating in a Software Supply Network

2007-21 Karianne Vermaas (UU)
Fast diffusion and broadening use: A research on residential adoption and usage of broadband internet in the Netherlands between 2001 and 2005

2007-22 Zlatko Zlatev (UT)
Goal-oriented design of value and process models from patterns

2007-23 Peter Barna (TUE)
Specification of Application Logic in Web Information Systems

2007-24 Georgina Ramírez Camps (CWI)
Structural Features in XML Retrieval

2007-25 Joost Schalken (VU)
Empirical Investigations in Software Process Improvement

====
2008
====

2008-01 Katalin Boer-Sorbán (EUR)
Agent-Based Simulation of Financial Markets: A modular, continuous-time approach

2008-02 Alexei Sharpanskykh (VU)
On Computer-Aided Methods for Modeling and Analysis of Organizations

2008-03 Vera Hollink (UVA)
Optimizing hierarchical menus: a usage-based approach

2008-04 Ander de Keijzer (UT)
Management of Uncertain Data - towards unattended integration

2008-05 Bela Mutschler (UT)
Modeling and simulating causal dependencies on process-aware information systems from a cost perspective

2008-06 Arjen Hommersom (RUN)
On the Application of Formal Methods to Clinical Guidelines, an Artificial Intelligence Perspective

2008-07 Peter van Rosmalen (OU)
Supporting the tutor in the design and support of adaptive e-learning

2008-08 Janneke Bolt (UU)
Bayesian Networks: Aspects of Approximate Inference

2008-09 Christof van Nimwegen (UU)
The paradox of the guided user: assistance can be counter-effective

2008-10 Wouter Bosma (UT)
Discourse oriented summarization

2008-11 Vera Kartseva (VU)
Designing Controls for Network Organizations: A Value-Based Approach

2008-12 Jozsef Farkas (RUN)
A Semiotically Oriented Cognitive Model of Knowledge Representation

2008-13 Caterina Carraciolo (UVA)
Topic Driven Access to Scientific Handbooks

2008-14 Arthur van Bunningen (UT)
Context-Aware Querying; Better Answers with Less Effort

2008-15 Martijn van Otterlo (UT)
The Logic of Adaptive Behavior: Knowledge Representation and Algorithms for the Markov Decision Process Framework in First-Order Domains

2008-16 Henriëtte van Vugt (VU)
Embodied agents from a user's perspective

2008-17 Martin Op 't Land (TUD)
Applying Architecture and Ontology to the Splitting and Allying of Enterprises

2008-18 Guido de Croon (UM)
Adaptive Active Vision

2008-19 Henning Rode (UT)
From Document to Entity Retrieval: Improving Precision and Performance of Focused Text Search

2008-20 Rex Arendsen (UVA)
Geen bericht, goed bericht. Een onderzoek naar de effecten van de introductie van elektronisch berichtenverkeer met de overheid op de administratieve lasten van bedrijven.

2008-21 Krisztian Balog (UVA)
People Search in the Enterprise

2008-22 Henk Koning (UU)
Communication of IT-Architecture

2008-23 Stefan Visscher (UU)
Bayesian network models for the management of ventilator-associated pneumonia

2008-24 Zharko Aleksovski (VU)
Using background knowledge in ontology matching

2008-25 Geert Jonker (UU)
Efficient and Equitable Exchange in Air Traffic Management Plan Repair using Spender-signed Currency

2008-26 Marijn Huijbregts (UT)
Segmentation, Diarization and Speech Transcription: Surprise Data Unraveled

2008-27 Hubert Vogten (OU)
Design and Implementation Strategies for IMS Learning Design

2008-28 Ildiko Flesch (RUN)
On the Use of Independence Relations in Bayesian Networks

2008-29 Dennis Reidsma (UT)
Annotations and Subjective Machines - Of Annotators, Embodied Agents, Users, and Other Humans

2008-30 Wouter van Atteveldt (VU)
Semantic Network Analysis: Techniques for Extracting, Representing and Querying Media Content

2008-31 Loes Braun (UM)
Pro-Active Medical Information Retrieval

2008-32 Trung H. Bui (UT)
Toward Affective Dialogue Management using Partially Observable Markov Decision Processes

2008-33 Frank Terpstra (UVA)
Scientific Workflow Design; theoretical and practical issues

2008-34 Jeroen de Knijf (UU)
Studies in Frequent Tree Mining

2008-35 Ben Torben Nielsen (UvT)
Dendritic morphologies: function shapes structure

====
2009
====

2009-01 Rasa Jurgelenaite (RUN)
Symmetric Causal Independence Models

2009-02 Willem Robert van Hage (VU)
Evaluating Ontology-Alignment Techniques

2009-03 Hans Stol (UvT)
A Framework for Evidence-based Policy Making Using IT

2009-04 Josephine Nabukenya (RUN)
Improving the Quality of Organisational Policy Making using Collaboration Engineering

2009-05 Sietse Overbeek (RUN)
Bridging Supply and Demand for Knowledge Intensive Tasks - Based on Knowledge, Cognition, and Quality

2009-06 Muhammad Subianto (UU)
Understanding Classification

2009-07 Ronald Poppe (UT)
Discriminative Vision-Based Recovery and Recognition of Human Motion

2009-08 Volker Nannen (VU)
Evolutionary Agent-Based Policy Analysis in Dynamic Environments

2009-09 Benjamin Kanagwa (RUN)
Design, Discovery and Construction of Service-oriented Systems

2009-10 Jan Wielemaker (UVA)
Logic programming for knowledge-intensive interactive applications

2009-11 Alexander Boer (UVA)
Legal Theory, Sources of Law & the Semantic Web

2009-12 Peter Massuthe (TUE, Humboldt-Universitaet zu Berlin)
Operating Guidelines for Services

2009-13 Steven de Jong (UM)
Fairness in Multi-Agent Systems

2009-14 Maksym Korotkiy (VU)
From ontology-enabled services to service-enabled ontologies (making ontologies work in e-science with ONTO-SOA)

2009-15 Rinke Hoekstra (UVA)
Ontology Representation - Design Patterns and Ontologies that Make Sense

2009-16 Fritz Reul (UvT)
New Architectures in Computer Chess

2009-17 Laurens van der Maaten (UvT)
Feature Extraction from Visual Data

2009-18 Fabian Groffen (CWI)
Armada, An Evolving Database System

2009-19 Valentin Robu (CWI)
Modeling Preferences, Strategic Reasoning and Collaboration in Agent-Mediated Electronic Markets

2009-20 Bob van der Vecht (UU)
Adjustable Autonomy: Controling Influences on Decision Making

2009-21 Stijn Vanderlooy (UM)
Ranking and Reliable Classification

2009-22 Pavel Serdyukov (UT)
Search For Expertise: Going beyond direct evidence

2009-23 Peter Hofgesang (VU)
Modelling Web Usage in a Changing Environment

2009-24 Annerieke Heuvelink (VUA)
Cognitive Models for Training Simulations

2009-25 Alex van Ballegooij (CWI)
"RAM: Array Database Management through Relational Mapping"

2009-26 Fernando Koch (UU)
An Agent-Based Model for the Development of Intelligent Mobile Services

2009-27 Christian Glahn (OU)
Contextual Support of social Engagement and Reflection on the Web

2009-28 Sander Evers (UT)
Sensor Data Management with Probabilistic Models

2009-29 Stanislav Pokraev (UT)
Model-Driven Semantic Integration of Service-Oriented Applications

2009-30 Marcin Zukowski (CWI)
Balancing vectorized query execution with bandwidth-optimized storage

2009-31 Sofiya Katrenko (UVA)
A Closer Look at Learning Relations from Text

2009-32 Rik Farenhorst (VU) and Remco de Boer (VU)
Architectural Knowledge Management: Supporting Architects and Auditors

2009-33 Khiet Truong (UT)
How Does Real Affect Affect Affect Recognition In Speech?

2009-34 Inge van de Weerd (UU)
Advancing in Software Product Management: An Incremental Method Engineering Approach

2009-35 Wouter Koelewijn (UL)
Privacy en Politiegegevens; Over geautomatiseerde normatieve informatie-uitwisseling