Journal of Universal Computer Science, vol. 16, no. 16 (2010), 2214-2231
submitted: 31/3/10, accepted: 31/7/10, appeared: 28/8/10 © *J.UCS*

Content Recommendation in APOSDLE using the Associative Network

Hermann Stern
(Graz University of Technology, Graz, Austria
hermann.stern@tugraz.at)

Rene Kaiser
(JOANNEUM RESEARCH, Graz, Austria
rene.kaiser@joanneum.at)

Philip Hofmair
(JOANNEUM RESEARCH, Graz, Austria
philip.hofmair@joanneum.at)

Peter Kraker
(Know-Center GmbH, Graz, Austria
pkraker@know-center.at)

Stefanie N. Lindstaedt
(Know-Center GmbH, Graz, Austria
slind@know-center.at)

Peter Scheir
(Styria Media Group AG, Graz, Austria
peter.scheir@styria.com)

Abstract: One of the success factors of Work Integrated Learning (WIL) is to provide the appropriate content to the users, both suitable for the topics they are currently working on, and their experience level in these topics. Our main contributions in this paper are (i) overcoming the problem of sparse content annotation by using a network based recommendation approach called Associative Network, which exploits the user context as input; (ii) using snippets for not only highlighting relevant parts of documents, but also serving as a basic concept enabling the WIL system to handle text-based and audiovisual content the same way; and (iii) using the Web Tool for Ontology Evaluation (WTE) toolkit for finding the best default semantic similarity measure of the Associative Network for new domains. The approach presented is employed in the software platform APOSDLE, which is designed to enable knowledge workers to learn at work.

Keywords: Work Integrated Learning, Recommender Systems, Multimedia Information Systems, Associative Networks
Categories: H.5.1, M.7, M.8, L.2.1

1 Introduction

Work Integrated Learning (WIL) [Lindstaedt and Farmer 04] [Smith 03] enables knowledge workers to learn directly at their place of work, by doing their usual tasks instead of exercising in dedicated artificial learning situations, with working and learning seamlessly integrated [Lindstaedt et al. 09].

WIL raises the learner's awareness of content, people, learning opportunities and structure; all relevant to the work activities she is involved. One of the key factors of WIL environments is to provide appropriate content to the learner, suitable for (i) her current level of knowledge, and (ii) the current task at hand – both together forming the context. The approach presented in this paper derives the user context by exploiting the topics a knowledge worker is currently working with, and her experience level in these topics. This user context is then used to tailor the recommended content according to the learner's current needs.

In WIL, it is not sufficient to recommend whole documents; a knowledge worker does not benefit from browsing dozens of pages of a document just to find the one or two valuable pieces of information she needs. Instead, the learner has to be made aware of specific content sections relevant to her information needs. To assist the learner in finding relevant information within a document, we introduce *snippets* as semantically annotated parts of those documents. Besides the problem of too much content available and learners unable to select the right learning material for their needs, in WIL it is also often the case that there is not enough content available for a specific topic. We overcome this problem of sparse annotations by using a network based recommendation approach called *Associative Network*. This network exploits a domain model - which is a semantic description of the domain in terms of concepts, relations, and objects - to annotate snippets and relate them to various topics. These topics correlate to other topics in the model by means of their similarity. Using this *Associative Networ*k, the knowledge worker is able to select the appropriate content for her current working needs. In addition, the correlation between similar topics can be used to find more suitable content not belonging to the selected topic, but to very similar ones. In contrast to other learning environments, in this approach content is not limited anyhow to pure textual resources; other sources such as videos or tutorials may also serve as valuable learning materials.

Our main contributions in this paper are as follows: we overcome the problem of sparse content annotation by using a network based recommendation approach called Associative Network; we use snippets for not only highlighting relevant parts of documents, but also they are serving as a basic concept enabling the WIL system to handle text-based and audiovisual content the same way; and finally we use the Web Tool for Ontology Evaluation (WTE) toolkit to find the best default semantic similarity measure of the Associative Network for new domains.

In this paper, we employ a scenario showing how a fictitious employee called Eva prepares her first creativity workshop; her goal is to generate ideas for rear view mirrors. Eva uses a WIL system to get recommendations for mastering her current task. Instead of presenting the whole scenario, we will list the specific aspects directly affecting this scenario in every section. To keep the flow of the scenario easy traceable throughout the whole paper, we use an italic font for the scenario text.

2 Associative Network

The Associative Network - the approach itself can be found in [Scheir 08a] - is responsible for recommending content via snippets. Associative means that information is in some form associated to other information. This association is modelled as a graph, referred to as Associative Network. The idea behind this network is to return a list of most relevant snippets to the user; the input topics (one or more topics) like “Workshop” (i) can be chosen by the user, (ii) can also be detected automatically by some context detection functionality, or (iii) are inferred by a working task the user has chosen.

The principle of associating information is not new; [Salton 63] introduced a document retrieval system connecting books with the same authors. [Crestani 97] proposed Associative Retrieval as a form of Information Retrieval, where some given relevant information is used to find more related (“associated”) information.

[Crestani 97] defines an Associative Network as a generic network of information items represented by nodes. Those nodes are linked using edges, which can be named and weighted. According to [Heylighen et al. 02] an Associative Network is a weighted, directed graph, whose nodes represent documents, and whose weights represent the associations between nodes. The network itself is stored in a matrix, allowing to relate each node to all the other nodes. The edges themselves are not named. The Associative Network presented in this paper is defined according to [Crestani 97]. Every concept in the domain model and every snippet in the knowledge base are represented by nodes, optionally connected by edges to other nodes. These edges are weighted, representing the grade of similarity between the two connected nodes. Three directly dependent layers [see Figure 1] provide the connection between the input (one or more topics) and the list of snippets provided as result.

The first layer (*Concept to Concept layer, C2C)* contains only domain concept nodes (also referred to as topics) which are connected with edges by means of their *semantic similarity* outlined in Section 3; edge weights state the grade of similarity.

The second layer (*Concept to Document layer, C2D*) states the mappings between the topics and the snippets, referred to as *semantic annotation* explained in Section 4. Each snippet is annotated with exactly one topic. This is the only mandatory layer in the Associative Network since it is the core layer serving snippets for one or more topics; it cannot be turned off optionally.

The third layer (*Document to Document layer, D2D*) contains only snippet nodes and connects them with edges based on their *textual similarity*. If activated, this layer is searching for snippets with similar textual content and adding them to the list of snippets resulting from the second layer.

Textual similarity is not within the focus of this paper, therefore the *Concept to Concept layer* including the underlying semantic similarity measures and *Concept to Document layer* using semantically annotated snippets are presented in detail.

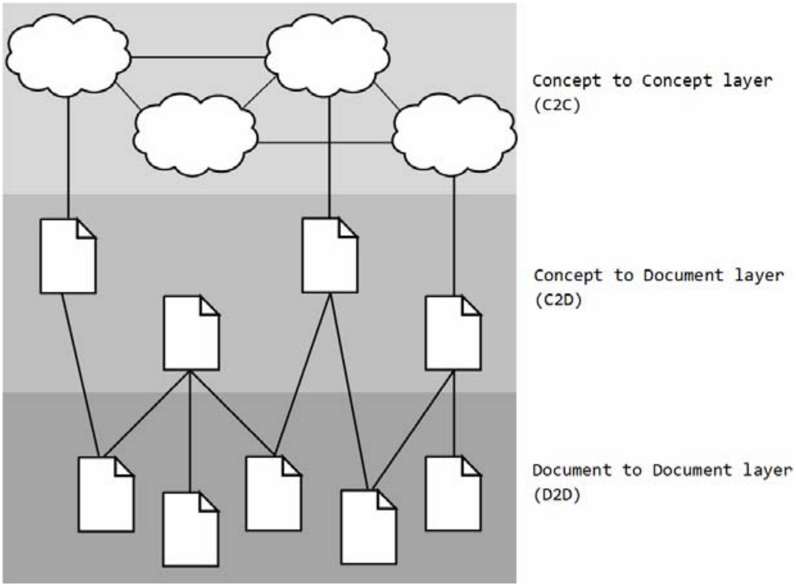

Figure 1: Anatomy of the Associative Network

Compared to a traditional retrieval engine, the Associative Network introduces several advantages when dealing with sparse annotations. Manual annotation of snippets is very time consuming and has to be done by domain experts; in most cases only few concepts from the models are used for annotation, and only few resources are really annotated. The Associative Network is able to (i) retrieve documents from concepts not directly selected by the user, but similar to them (*Concept to Concept layer*), and (ii) also documents that are not tagged with the queried concept, but similar to the found documents (*Document to Document layer*). Using snippets as information basis makes the Associative Network completely independent with regard to the documents. There is no difference in handling textual or multimedia based content, since snippets are only pointers to certain parts of information.

3 Concept Similarity

The *Concept to Concept layer* of the Associative Network is directly built from the application domain ontology. The edge weights stating the similarity between pairs of concepts of the domain ontology are derived using a semantic similarity function, which measures how much those two concepts are alike. Several semantic similarity functions have been proposed in literature, which correspond to different criteria of 'being alike'. The evaluation of these functions (hereafter, measures) has shown that no single measure always performs better than the others. Moreover, individual measures perform differently depending on the criteria of 'being alike' required.

This paper exploits the following semantic similarity measures: (i) Conceptual Similarity [Wu and Palmer 94], (ii) Shortest Path, (iii) Scaled Shortest Path [Leacock and Chodorow 98], (iv) Resnik [Resnik 95], (v) Lin [Lin 98], (vi) Jiang Conrath [Jiang and Conrath 97], and (vii) the Trento Vector [Scheir 08b]. Measures (i)-(iii) belong to the so-called *path-based measures*, as they measure similarity using the distance (or path length) between concepts in the is-a hierarchy. Measures (iv)-(vi) belong to the so-called *information content-based measures*, as they measure similarity using the closeness of the information content of different concepts. Measure (vii) calculates similarity by looking at how many relations concepts have in common.

The concept similarity based *Concept to Concept layer* is used to extend the list of input topics to also consider other topics, which are very similar but have not been selected by the user. For processing the information in the network representation of this layer, a technique called Spreading Activation is used. Starting from a set of initially activated nodes in the net, the activation spreads over the network [Sharifian and Samani 97]. During search, energy flows from a set of initially activated information items over the edges to their neighbours [Lindstaedt et al. 08]. Using several steps of Spreading Activation in the highly interconnected network would result in a list containing all of the topics; therefore, only one iteration of Spreading Activation is carried out, and only the topics directly connected to the starting topic are added.

Following our scenario, Eva selects the topic "Workshop" as starting point for her learning experience. Since there are not enough snippets available for the topic "Workshop", she also gets results from very similar topics (like "Workshop Agenda", "Workshop planning" or even "Creativity Techniques"). All those highly related documents are relevant for Eva's current working activity, even though she has not requested them explicitly.

4 Semantic Annotation of Snippets

The second layer (*Concept to Document layer*) of the Associative Network states the mappings between the topics and the snippets. This is the only mandatory layer in the Associative Network, since it is the core layer serving snippets for one or more topics; it cannot be turned off optionally.

Certain parts of documents are tagged with *concepts* and *material uses* yielding a list of annotated fragments; we refer to them as *snippets*. A *concept* denotes a topic, taken from the domain ontology; the selected part of the document is dealing with. *Material use* describes the resource type of the selected part, like "Definition", "Example", "Guideline" etc. Examples for snippets are a drawing in a PowerPoint document, a subsection of a requirements engineering document, a section of a video clip, a link to a person or a community of practice, etc. Every snippet is addressable via a unique identifier. Snippets of text-based documents denote a specific area (like an abstract, an image etc), while multimedia-based snippets denote a certain timeframe of the corresponding video or audio file.

Before the *Concept to Document layer* can be used, semantically annotated snippets have to be created. There are two possibilities to create snippets: (i) they may be entered manually by domain experts or (ii) they are automatically created using a classification engine based on the knowledge base that was created during manual annotation. The manual annotation continually delivers ground truth for automatic classification and annotation of new documents. However, to bootstrap the component, an initial dataset of snippets from representative documents has to be entered manually. This is necessary for the training of the automatic annotation functionality. According to our tests, good automatic classification results can be achieved for topics with at least one hundred assigned snippets. Depending on the number of topics used in the specific domain, this yields the need for thousands of manually annotated snippets. Since the two domains covered by this paper were maintained by only two domain experts each, and the resources test set consisted of a few hundred documents, not all of the topics were annotated with enough snippets. The text classification algorithm used can be applied to all text-based documents, but also to multimedia content by analyzing the textual information extracted from this content as metadata. The *Concept to Document layer* is using a Lucene[1] index storing all snippets and topics; this index is also available for full text search.

Various domain experts in Eva's company have manually created snippets and annotated them with topics. In our scenario, this layer provides the functionality to deliver all snippets associated to Eva's topic "Workshop", and all the (optional) topics added in the previous Concept to Concept layer.

5 Exploiting the Associative Network Approach in APOSDLE

The work presented within this paper originates from the EU project APOSDLE[2], which focuses on Work Integrated Learning. The detailed approach of the whole project is described in [Lindstaedt et al. 07]; more information regarding the project can be found at the project's website[3].

The aim of APOSDLE is to enhance the productivity of knowledge workers by supporting informal learning and teaching activities in the context of knowledge workers' everyday work processes. Knowledge workers are for example engineers, researchers, software developers, consultants, or designers. APOSDLE is designed to support knowledge workers in their three roles: learner, worker and expert. The project exploits synergies between learning and knowledge management by reusing content not originally intended for learning. For companies, this implies the chance to reuse existing content repositories, which potentially saves costs compared to the generation of new material. Moreover, integrating document-based and multimedia-based knowledge sources into the system requires relatively little effort.

An intelligent software component aims to identify and suggest knowledge that knowledge workers are lacking in order to complete their current task. This knowledge gap is estimated by taking into account the user's context, several semantic models containing work tasks, topics, and learning goals [see Section 5.1], as well as user competencies, and the Associative Network [see Section 2] that implements a heterogeneous recommendation mechanism.

---

[1] http://lucene.apache.org/
[2] Advanced Process Oriented and Self-Directed Learning Environment;
[3] http://www.aposdle.tugraz.at

5.1 Models used in APOSDLE

APOSDLE uses a model [Lindstaedt et al. 08] individually created for the respective application domain, which contains:

i. a *task model* (a formal description of work tasks to be supported, like "how to create a workshop");
ii. a *domain model* (a semantic description of the domain in terms of concepts, relations, and objects); and
iii. a *learning goal model* (providing a mapping between domain concepts, tasks and general learning goal types).

The creation of these models requires manual input and forms the basis for the automatic content recommendation mechanisms. Entities in the domain model are used to annotate content to make it available for retrieval. The effort for creating a new set of models from scratch is heavily dependent on domain specific boundaries, including the number of concepts, tasks or learning goals, and the company or department the models are created for. Depending on the number of domain experts and the documents serving as a basis for knowledge elicitation, the effort needed can vary between a few weeks and several months. The creation of the two domains CCI and ISN mentioned in this paper was carried out hand in hand with the development of the modelling tools used; therefore, the effort cannot be numbered exactly.

For a very small test domain containing less than one hundred concepts and only few domain experts involved, we have recorded the following list of tasks with the respective efforts:

Tasks on the company / department side:

(i)	Scope & Boundaries and Resources Collection:	2 hours
(ii)	Knowledge Elicitation from Digital Resources:	1 hour
(iii)	Knowledge Elicitation from Domain Experts:	10 hours
(iv)	Modelling of Topics and Tasks:	20 hours
(v)	Revision Support for Tasks and Topics:	2 hours
(vi)	Modelling of Learning Goals:	3 hours
(vii)	Formal Models Validation:	2 hours

Tasks on APOSDLE team side:

(viii)	Administration, Test Preparation, Verification etc.:	40 hours

This list denotes the minimum effort for a new model set, given that the domain experts already know how to use the modelling tools used. More on the semantic modelling process and the company-specific domain concepts can be found in [Christl et al. 08]. A typical application domain in the scope of this paper contains approximately one hundred different topics. The number of topics has to be chosen carefully, as the domain experts should be able to annotate the created snippets correctly; moreover, the users should be able to determine between the individual topics. In the domains CCI and ISN, the domain experts were able to annotate all snippets with about one hundred different topics. Using a model with hundreds of concepts to cover big domains with thousands of documents has little effect on the system performance, but significantly increases the effort needed for modelling the domain and correctly annotating the snippets.

The application domain model defines the working field of Eva; it contains the tasks she needs to execute during work and the topics she is dealing with. APOSDLE provides all the topics and tasks, and Eva can select some of them to get further assistance.

5.2 APOSDLE Suggests

The user interface for presenting a list of recommended resources to the user is called APOSDLE Suggests. See Figure 2 for a recommendation screen showing relevant resources for the topic “Standard Deviation”.

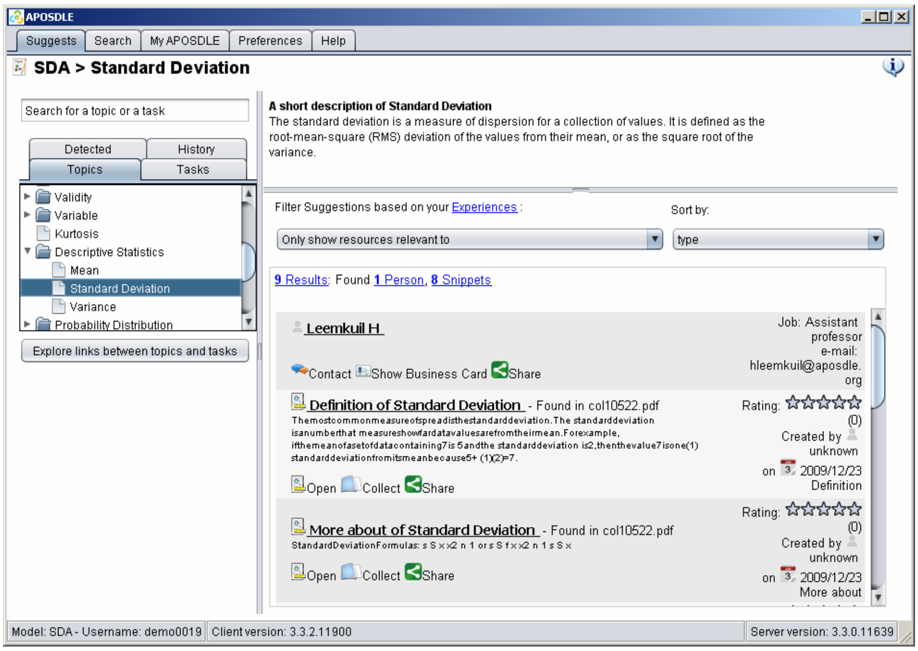

Figure 2: APOSDLE Suggests window

On the left side there is the navigational area where the knowledge worker can manually select *Topics* (“Standard Deviation”) or *Tasks* (like “Performing a Parametric Test” or “Interpreting a Test Value”) she is currently interested in; APOSDLE then provides her with a list of relevant resources – see Figure 2 for a listing of human experts and snippets, ordered by relevance.

APOSDLE automatically detects the tasks or topics the user is involved in, listing them up in *Detected*. This approach is keyword based; meaning that for instance using the text fragment “creating a workshop” in a Word document triggers the detection of topic “Workshop”. Further details about the APOSDLE task detection approach can be found in [Lokaiczyk et al. 08] and [Rath 10]. A small popup in the right lower corner of the desktop indicates that APOSDLE has recognized a task and prepared assistance, keeping user disturbance at a minimal level. A history of the last topics used for retrieval is available in *History***.** On the right side, the recommendations for the selected task or topic are presented, ordered by their relevance. All listed resources are shown including relevant metadata like contact details for human experts, or ratings and creation date for snippets. Clicking on a textual or multimedia resource opens the APOSDLE component for visualizing resources – the APOSDLE Reader.

Continuing our scenario of Eva searching for assistance preparing her workshop, APOSDLE automatically detects the topic Eva is working on (which is "Workshop" in this scenario) and informs Eva with a small notice popup. Eva can now decide whether she needs assistance or not; by clicking on the popup, APOSDLE presents all recommended resources. Eva now browses the list of recommended snippets, which are parts of Word files, Excel sheets, videos, audio resources etc.

5.3 APOSDLE Reader

APOSDLE treats all documents equivalently and uses a single component to view documents - the APOSDLE Reader.

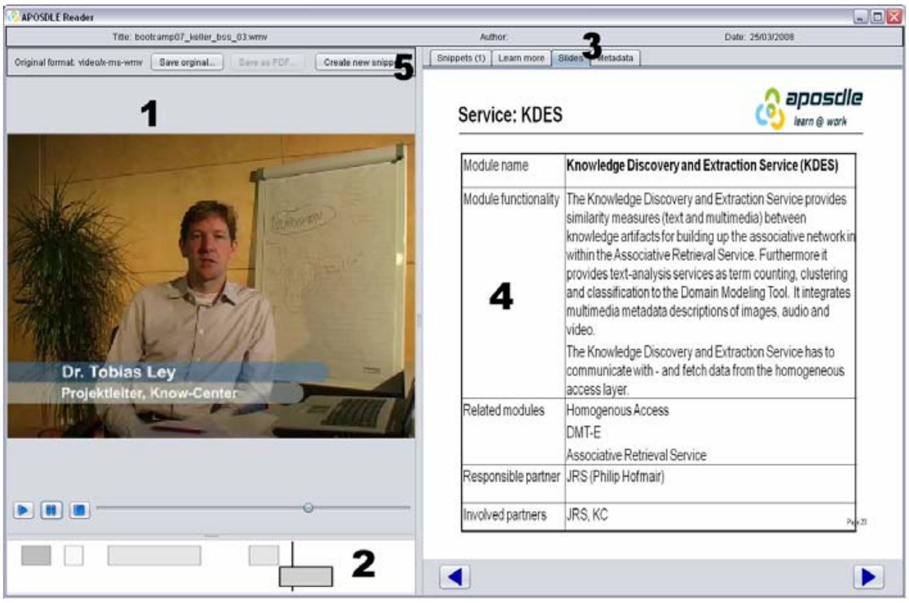

Figure 3: Screenshot of the APOSDLE Reader in video lecture mode.

The APOSDLE Reader primarily visualizes recommended resources - which includes presentation of textual PDF resources and replay of multimedia content as seen in Figure 3. Secondly this interface is used for the creation of the underlying snippets themselves. Providing only one tool for browsing the APOSDLE knowledge base and creating new snippets lowers the entry barrier for new users and makes usage of the APOSDLE system quite easy.

Textual documents, internally converted to the PDF format, are displayed using a PDF Reader component, whereas multimedia documents are displayed using the *JMediaPlayer* from JavaFx[4] which utilizes codecs already available on the platform and provides classic video player functionality.

If a certain part of video is recommended, the playback component will be positioned at the snippet's start; given a textual resource, APOSDLE automatically selects the specific text snippet in the document. As stated before, new snippets can be easily created by defining the boundaries (temporal start and end positions in of videos respectively selecting an area in a text document) via pressing the button *create snippets* [see Number 5 in Figure 3]. When dealing with video lecture content, additional metadata and PowerPoint slides may also be displayed in additional tabs if available.

Figure 3 shows the most relevant areas of the APOSDLE Reader in multimedia content mode: In the centre of the left column [see Number 1], there is the video player component showing a video lecture from a project tutorial. Below [see Number 2], there is a *ThemeRiver* view; a specific graphical overview illustrating the temporal position of all the snippets in the multimedia file. Each snippet represents one bar on the river. By clicking on one of the bars, the video player goes to the position of the selected snippet. On the right side there are multiple tabs [see Number 3] - *Snippets*, *Learn more*, *Slides* and *Metadata* in the current case. On the screenshot, the slides tab is active, therefore the slides of the lectures are shown in the right area [see Number 4]. Using the left and right arrow buttons on the bottom navigates through the different slides.

Eva chooses a recommended snippet and the related document is displayed in the APOSDLE Reader. To avoid the need for reading whole documents, relevant parts of the recommended resource are highlighted; they have been annotated with snippets. After reading some textual documents, Eva opens a recommended video of the last company workshop and gets valuable clues about how to moderate a workshop and plan the agenda.

5.4 APOSDLE Text-based and Multimedia Content Types

APOSDLE handles two types of text-based snippets, (i) text files (in various forms) and (ii) URLs. A text snippet describes a specific section of a textual resource, which can be for example a plain text file, a Microsoft Office or Open Office document, HTML pages, a PDF file etc. An URL snippet is used to point to a specific Web resource; the content of this Web resource is not being indexed because internal storage of huge amounts of information usually available on large Web sites is out of scope for a learning environment. URL snippets are used like bookmarks to attach relevant Web sites to topics, APOSDLE recommends those URLs when the corresponding topic is selected.

Further, APOSDLE supports multiple types of audiovisual content: audio files, videos, and specifically video lecture casts. As a consequence, user interface may vary: the viewing component (e.g. a video player for a video in a specific encoding), metadata granularity and certain features such as slide navigation for video lectures are adapted accordingly. The following categories are currently supported:

---

[4] http://javafx.com

Educational legacy videos: During the installation process, a company's existing suitable content can be imported and indexed. This material may range from an image film to any footage of a business meeting or a product presentation video. Any (temporal) subsection may be used to create a snippet of valuable knowledge, which can later help knowledge workers to complete tasks. Because the creation of multimedia content is typically expensive, efficient re-use is economically reasonable.

Screen casts: APOSDLE enables the recording of screen contents via an integrated screen cast tool called CamStudio[5]. This recorder also captures videos played on screen in combination with audio. The main purpose of this tool is to help experts to document knowledge with very little effort. Any explanation, instruction or consultation created that way can be imported into the system and be made available to other knowledge workers.

Webcam Tutorials: Similar to screen casts, streams from webcams can be captured, not limiting the content to the computer interface itself. This feature enables experts to create simple tutorials at their workplace possibly using a flipchart or demonstrating practical actions. By annotating snippets carefully, the expert may not have to answer the same question again and can save precious working time.

Video Lectures: Typically video lectures consist of the video content itself, some metadata and a set of slides. Any video lecture, such as content from the Videolectures.net[6] portal, could be imported including metadata. The functionality for including external material extensively is a future work task; the current import functionality is rather simple. Slides will be extracted as images for displaying them in a separate slides tab, so that the content of the slides is visible even when the projected area in the background of the speaker is not, which is typically the case in lecture casts.

6 Evaluation of semantic similarity measures used in the Associative Network

In this evaluation, we focus on the *Concept to Concept layer*, which is responsible for handling the semantic similarity between the concepts (topics) available in an application domain. There are several semantic similarity measures available in APOSDLE, and it is very important to choose the best one for a specific domain as early as possible. Using the evaluation tool presented in this section, we are able to identify the best measures for a specific domain while still creating the models. The tool is able to import concepts and their connections automatically, which makes continuous evaluation possible, even on a daily basis.

The *Concept to Document layer* basically describes the manual annotations done by domain experts. Evaluating this layer would require a larger number of domain experts, since they would have to rate the snippets created by their colleagues. Due to the very limited number of domain experts available in the two domains (two at most), this evaluation was not applicable here.

---

[5] http://camstudio.org/

[6] http://videolectures.net/

The *Document to Document layer* of the Associative Network is based on already established and evaluated text content similarity measures (tf-idf) and is also not applicable for multimedia content. For a detailed evaluation discussing the optimal layer configurations of the Associative Network refer to [Scheir et al. 07]. For evaluating the Associative Network as a whole, there are plans to include rating elements (such as thumbs up or down) to enable the user to give feedback about the benefit of certain documents or snippets for her current information need.

As introduced in Section 3, the first layer of the Associative Network can exploit several different semantic similarity measures for weighting the edges between topics in the *Concept to Concept layer*. These measures include: (i) Conceptual Similarity [Wu and Palmer 94], (ii) Shortest Path, (iii) Scaled Shortest Path [Leacock and Chodorow 98], (iv) Resnik [Resnik 95], (v) Lin [Lin 98], (vi) Jiang Conrath [Jiang and Conrath 97], and (vii) the Trento Vector [Scheir 08b].

In this evaluation the following research questions are covered:

(i) Which semantic similarity measure is the best one for a specific application domain?

(ii) Is there a specific similarity measure which is performing well on every domain tested and could therefore be proposed as default measure? and

(iii) How much differ the covered measures actually?

To examine the behaviour in different application domains, two individual domains are evaluated: ISN (dealing with innovations) and CCI (German Chambers of Commerce). The basic idea of the evaluation is to get an ideal semantic similarity sample (“golden sample”) created by human experts and to compare it to the resulting layers using different measures. The data required for this evaluation is gathered using a Web application called Web Tool (for Ontology) Evaluation (WTE), see Figure 4 for a typical screen during a rating session.

Two domain experts of every domain were rating Concept Pairs in two or more individual sessions, yielding a test set of at least four individual sessions. To reduce memory effects, the experts were asked to wait at least a few days before they start the next session. During the test, a certain number of Concept Pairs were presented and the domain expert had to grade the degree of interconnection between those two concepts. Using these ratings, an “ideal” model of the *Concept to Concept layer* was created.

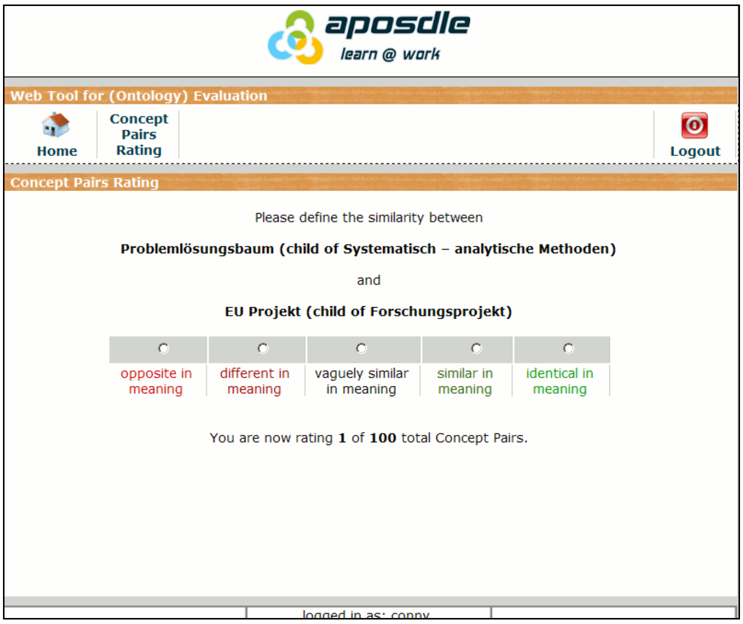

Figure 4: Screenshot of WTE

To ensure a correct understanding of the concepts shown, the parent concepts (if available) for both concepts were shown. According to [Charles 00] we have chosen a five point similarity scale, containing the following values:

Scale Description	**Value**
identical in meaning	1
similar in meaning	0.75
vaguely similar in meaning	0.55
different in meaning	0.27
opposite in meaning	0

Table 1: Similarity Scale used in WTE

For the domain to be evaluated, APOSDLE provides the derived measures for each edge connecting two concepts. A domain containing 120 concepts is therefore yielding a maximum of 7140 individual Concept Pairs to be rated. Recursive edges (an edge from one concept back to the same one) are not evaluated; also edges between two concepts are symmetric: *concept1* to *concept2* is the same as *concept2* to *concept1* and therefore has the same similarity values assigned. Since it is impossible to evaluate all the available Concept Pairs, there is the need for some pre-selection.

The number of Concept Pairs to be rated in one session can be defined individually for each application domain: (i) *fixed concepts* denote predefined Concept Pairs which are queried in every session – allowing us to compare ratings from different experts and over time; (ii) *variable concepts* describe Concept Pairs which are randomly selected from the pool of all possible Concept Pairs for each session. Fixed concepts were preselected manually by inspecting the similarity matrix of the domain including all the measures with the aim of being able to discriminate between measures. For example, Concept Pairs scoring very different values for the different measures are helping us to see which measures are more similar to the human evaluation. Having more individual rating results for significant Concept Pairs tends to be more useful than having a larger number of totally randomly selected Concept Pairs. The number of total Concept Pairs to be rated was set to 100, taking approximately 15 to 20 minutes per session. In the ISN domain we preselected 32 Concept Pairs, in the CCI domain 28.

Only completed sessions were stored by WTE, ensuring that sessions always contain the same number of ratings. All tasks (import of measures, import of concepts, import of Concept Pairs etc.) except the pre-selection of Concept Pairs can be carried out automatically, which makes it very easy to test other domains or work with different measure settings. The results of the ratings then were evaluated against the computed measures. All measures were normalized (values between 0 and 1) and the average deviations of each measure compared to the expert ratings were derived.

The evaluation yielded the following results: In the ISN domain, two experts created 600 individual ratings in six sessions, effectively rating 435 different Concept Pairs. The 32 fixed Concept Pairs have been rated in every session (so they had 6 ratings); 5 Concept Pairs have been rated twice (meaning that they had been randomly selected twice). The 6 individual ratings of the fixed Concept Pairs were fairly consistent, they only differed by one rating at the most; the average value was taken to clear out some possible erroneous entries.

The best three measures compared to the expert ratings in ISN are:

1. *Conceptual Similarity:*
   Average Deviation: *0.1899*

2. *Jiang Conrath:*
   Average Deviation: *0.2032*

3. *Shortest Path:*
   Average Deviation: *0.2123*

In the *CCI* domain two experts created 400 individual ratings in four sessions, effectively rating 306 different Concept Pairs. The 28 fixed Concept Pairs have been rated in every session (so they had 4 ratings), 10 Concept Pairs have been rated twice (meaning that they had been randomly selected twice). The best three measures compared to the expert ratings in CCI are:

1. *Scaled Shortest Path:*
   Average Deviation: *0.1247*

2. *Shortest Path:*
   Average Deviation: *0.1579*

3. *Conceptual Similarity:*
   Average Deviation: *0.1929*

The results in ISN and CCI domains yield the following answers to the evaluation questions:

(i) Which semantic similarity measure is the best one for a specific application domain?

In the ISN domain, the Conceptual Similarity measure is the best one compared to the expert ratings. Jiang Conrath and Shortest Path rank second and third. In the CCI domain, Scaled Shortest Path delivers the best results; Shortest Path and Conceptual Similarity do follow with a quite significant distance.

(ii) Is there a specific similarity measure, which is performing well on every domain, tested and could therefore be proposed as default measure?

There is no similarity measure, which performs perfectly on both domains, so there is no clear winner. However, considering both domains, the path-based measures seem to be the most appropriate for the kind of ontologies, which constitute the domain models. In fact, both the Conceptual Similarity and the Scaled Shortest Path belong to this family of measures. Focusing on a single measure the Shortest Path measure seems to be the best one suitable for a default setting.

(iii) How much differ the covered measures actually?

The difference between the different measures can be seen clearly in the top 3 ranking of the results section. While the differences are not that big in the case of ISN, Scaled Shortest Path is the clear winner in the case of CCI. The worst measures in both domains deliver very high average deviations of up to 0.35, meaning that choosing the wrong measure may significantly degrade the results.

7 Conclusion and Future Work

In this paper, we presented our approach of a context aware recommendation engine based on the Associative Network, which is able to overcome several problems when dealing with manual annotation of resources by human experts, like sparse annotations and annotation interfaces not intuitively enough for easy use. Snippets are relevant parts of documents annotated with concepts from a domain ontology, saving the knowledge worker from browsing the whole document or watching the whole video by highlighting the right learning material for her needs.

While typical personal learning environments focus on text-based content, APOSDLE is a self-directed Work Integrated Learning system supporting both text-based and audiovisual content. As most of the functionality is fully automatic, economic re-use of existing content is supported; easy recording and ingestion of new material enables experts to share their knowledge rapidly.

The system currently does not support the automatic import of video lectures metadata, a feature which certainly users would benefit from. The increasingly attractive integration of online videos (lectures) in learning environments is also discussed in [Auinger et al. 09]. Semantic interlinking between snippets and links to external, related (Web) resources may also be subject of future work. While the current import component demonstrates how external data can be integrated, a component with more sophisticated features providing extensive metadata would obviously be interesting. To improve the system's response time when opening a large video file, streaming could be utilized instead of downloading the whole file locally.

The Associative Network approach will be used in other Web based projects, with the following improvements planned: the network will be able to store not only snippets, but also other types of resources; the order of topics submitted as input will be taken into account and additional weighting algorithms and similarity measures will be developed.

The APOSDLE project has come to its end in the beginning of 2010. The application is still in prototype status but will be enhanced according to the needs of the application partners where APOSDLE is still running and being used. A demonstrator version is available at http://www.aposdle.tugraz.at for two of the application domains, ISN and SDA. As a final step in the APOSDLE project, the source code will be released under Open Source License in the mid of 2010.

For in-depth information regarding APOSDLE and the used technology, refer to project deliverables and various publications on the APOSDLE website. APOSDLE is a rather closed system at the moment, recommendations from other learning environments or other tools for instance cannot be included. By releasing APOSDLE under Open Source License, we foresee the development of an APOSDLE API to provide interface functionality to and from other systems.

The evaluation of the *Concept to Concept layer* of the Associative Network shows that path based measures seem to be the most appropriate measures for the APOSDLE models and the very simple shortest path measure performs well on both tested domains. Nevertheless, the best measure has to be evaluated on a case-by-case manner for each domain. In fact, the evaluation has emphasized the existence of big differences between the top scorer and the other measures, making the selection of the correct similarity function a crucial task for a good performance of the Associative Network. Since WTE can derive those measure evaluations on the fly via the Web interface, it is also possible to have a look at the results while the evaluation is still ongoing, for example after each individual session or user. Within APOSDLE, the use of a path based measures was accurate enough. In addition to just using one single measure type, the combination of good performing measures is also an interesting aspect for future work.

Acknowledgements

APOSDLE is partially funded under the FP6 of the European Commission within the IST Workprogramme (project number 027023). The Know-Center is funded within the Austrian COMET Program - Competence Centers for Excellent Technologies - under the auspices of the Austrian Ministry of Transport, Innovation and Technology, the Austrian Ministry of Economics and Labor and by the State of Styria. COMET is managed by the Austrian Research Promotion Agency FFG.

References

[Auinger et al. 09] Auinger, A., Ebner, M., Nedbal, D., Holzinger, A.: "Mixing content and endless collaboration | mashups: Towards future personal learning environments"; UAHCI '09: Proceedings of the 5th International Conference on Universal Access in Human-Computer Interaction, Part III, Springer-Verlag, Berlin, Heidelberg (2009), 14-23.

[Bedert et al. 08] Biedert, R., Schwarz, S., Roth-Berghofer, T.R.: "Designing a context-sensitive dashboard for an adaptive knowledge worker assistant", Workshop on Modeling and Reasoning in Context, HCP '08, Delft, The Netherlands (2008).

[Charles 00] Charles, W.G.: "Contextual correlates of meaning", Applied Psycholinguistics, v21:4, Cambridge University Press (2000), 505-524.

[Christl et al. 08] Christl, C., Ghidini, C., Guss, J., Pammer, V., Rospocher, M., Lindstaedt, S., Scheir, P., Serafini, L.: "Deploying semantic web technologies for work integrated learning in industry. a comparison: Sme vs. large sized company", Proceedings of the 7th International Semantic Web Conference (ISWC 2008), Semantic Web In Use Track, Volume 318/-1., Karlsruhe, Germany, Springer Berlin / Heidelberg (28-30 October 2008) 709-722.

[Crestani 97] Crestani, F.: "Application of Spreading Activation Techniques in Information Retrieval", Artificial Intelligence Review 11, 6 (1997), 453-482.

[Heylighen et al. 02] Heylighen, F., Bollen, J.: "Hebbian Algorithms for a Digital Library Recommendation System", 31st International Conference on Parallel Processing Workshops (ICPP 2002 Workshops), 20-23 August 2002, Vancouver, BC, Canada (2002), 439-446.

[Jiang and Conrath 97] Jiang, J.J. and Conrath, D.W.: "Semantic similarity based on corpus statistics and lexical taxonomy", Proceedings of International Conference on Research in Computational Linguistics (1997), 19–33.

[Leacock and Chodorow 98] Leacock, C. and Chodorow, M.: "Combining local context and wordnet similarity for word sense identifica", MIT Press (1998).

[Lin 98] Lin, D.: "An information-theoretic definition of similarity", Proceedings of the 15th International Conf. on Machine Learning, Morgan Kaufmann, San Francisco, CA (1998), 296-304.

[Lindstaedt and Farmer 04] Lindstaedt, S.N. & Farmer, J.: „Kooperatives Lernen in Organisationen", CSCLKompendium - Lehr- und Handbuch für das computerunterstützte kooperative Lernen, Oldenbourg Wissenschaftsverlag, München, Germany.

[Lindstaedt et al. 07] Lindstaedt, S. N., Ley, T., Mayer, H.: "APOSDLE - New Ways to Work, Learn and Collaborate", Proceedings of the 4th Conference on Professional Knowledge Management WM2007, 28. - 30. März 2007, GITO-Verlag, Berlin, Potsdam, Germany (2007), 381-382.

[Lindstaedt et al. 08] Lindstaedt, S. N., Scheir, P., Lokaiczyk, R., Kump, B., Beham, G., Pammer, V.: "Knowledge Services for Work-integrated Learning", Proceedings of the European Conference on Technology Enhanced Learning (ECTEL) 2008, Maastricht, The Netherlands (2008), 234-244.

[Lindstaedt et al. 09] Lindstaedt, S. N., Beham, G., Kump, B., Ley, T.: "Getting to Know Your User – Unobtrusive User Model Maintenance within Work-Integrated Learning Environments", Proceedings of the 4th European Conference on Technology Enhanced Learning, ECTEL 2009, Nice, France, September/October (2009), 73-87.

[Lokaiczyk et al. 08] Lokaiczyk, R. and Goertz, M.: "ContextAggregator: A heuristic-based approach for automated feature construction and selection", Proceedings of ECAI 2008 (2008), 827-828.

[Rath 10] Rath, A.: "User Interaction Context - Studying and Enhancing Automatic User Task Detection on the Computer Desktop via an Ontology-based User Interaction Context Model", Dissertation, Technische Universitaet Graz, Graz, Austria (2010).

[Resnik 95] Resnik, P.: "Using information content to evaluate semantic similarity in a taxonomy", IJCAI (1995), 448–453.

[Salton 63] Salton, G.: "Associative Document Retrieval Techniques Using Bibliographic Information", Journal of the ACM 10, Nr. 4 (1963), 440-457.

[Scheir et al. 07] Scheir, P., Granitzer, M., Lindstaedt, S.: "Evaluation of an information retrieval system for the semantic desktop using standard measures from information retrieval", LWA 2007, Lernen - Wissensentdeckung - Adaptivitaet, Halle/Saale. (2007).

[Scheir 08a] Scheir, P.: „Assoziative Suche fuer das Semantic Web", Dissertation, Technische Universitaet Graz, Graz, Austria (2008).

[Scheir 08b] Scheir, P., Lindstaedt, S.N., Ghidini, C.: "A network model approach to retrieval in the semantic web", Int. J. Semantic Web Inf. Syst. 4(4) (2008), 56-84.

[Sharifian and Samani 97] Sharifian, F. & Samani, R.: "Hierarchical spreading of activation", in Farzad Sharifian, ed., Proc. of the Conference on Language, Cognition, and Interpretation, IAU Press (1997), 1-10.

[Smith 03] Smith, P. J.: "Workplace Learning and Flexible Delivery", Review of Educational Research, 73, 1 (2003), 53-88.

[Wu and Palmer 94] Wu, Z. and Palmer, M.: "Verb semantics and lexical selection", 32nd Annual Meeting of the Association for Computational Linguistics (1994), 133-138.