



[Hopfgartner, F.](#), [Hannah, D.](#), [Gildea, N.](#) and [Jose, J. M.](#) (2008) Capturing Multiple Interests in News Video Retrieval by Incorporating the Ostensive Model. In: PersDB 2008 Second International Workshop on Personalized Access, Profile Management, and Context Awareness: Databases, Auckland, New Zealand, 23 Aug 2008, pp. 51-58.

This is the author's final accepted version.

There may be differences between this version and the published version. You are advised to consult the publisher's version if you wish to cite from it.

<http://eprints.gla.ac.uk/101464/>

Deposited on: 17 April 2018

Enlighten – Research publications by members of the University of Glasgow
<http://eprints.gla.ac.uk>

Capturing Multiple Interests in News Video Retrieval by Incorporating the Ostensive Model

Frank Hopfgartner, David Hannah, Nicholas Gildea, Joemon M. Jose
University of Glasgow
Glasgow, United Kingdom
{hopfgarf, hannahd, ngildea, jj}@dcs.gla.ac.uk

ABSTRACT

We propose an adaptive news video retrieval approach which is based on the Ostensive Model of developing information needs. We therefore introduce a news video retrieval system called NewsBoy which captures the users' implicit interactions with its graphical interface, extracts terms from visited video documents and stores them in user profiles. The terms are weighted based on the type of implicit feedback, multiple interests are identified by clustering the content of the profile. In this paper, we describe the architecture of the system and introduce our approach of adding the ostensive factor to capture the users' evolving interest. Preliminary results show the acceptance of the system and highlights drawbacks.

1. INTRODUCTION

Consuming information has a central impact on the development of our society, leading to the transformation from the industrial to the information age. Newspapers, television news broadcasts, the WWW and other sources provide the society with a vast amount of information, an increasing percentage of which is in digital format. However, facing this excessive supply of information sources might overwhelm information consumers. Hence, there is a need to provide personalised access to them.

Arezki et al. [1] provide an example to explain the need of a personalisation service: When a computer scientist enters the search query "java" into a search engine, he is most likely interested in finding information about the programming language. Other people, however, might expect results referring to the island of Java in Indonesia or a type of coffee beans bearing this name. Sebe and Tian [18] discuss that for providing personalised information based on multimedia content, sophisticated research in various areas is needed, including the acquisition of user preferences and how to filter information by exploiting the user's profile.

A classical approach to capture the user's preferences is profiling. User profiles can be used to create a simplified

model of the user which represents his interests on general topics. Commercial search engines incorporate such profiles, the most prominent being Google offering iGoogle and Yahoo! offering MyYahoo!. Query expansion is used to gather the user's interest and search results are re-ranked to match their interests.

These services rely on users' explicitly specifying preferences, a common approach in the text retrieval domain. By giving explicit feedback, users are forced to update their need, which can be problematic when their information need is vague [21]. Furthermore, users tend to provide not enough feedback on which to base an adaptive retrieval algorithm [8]. Deviating from the method of explicitly asking the user to rate the relevance of retrieval results, the use of implicit feedback techniques helps by learning user interests unobtrusively. The main advantage is that users are relieved from providing feedback. A disadvantage is that information gathered using implicit techniques are less accurate than information based on explicit feedback [14].

A challenging problem in user profiling is the users' evolving focus of interest. What a user finds interesting on day *A* might be completely uninteresting on day *B*, or even on the same day. The following example illustrates the problem: Joe Bloggs is rarely interested in sports. Thus, during Euro 2008, the European Football Championship, he is fascinated by the euphoria exuded by the tournament and follows all reports related to the event. After the cup final, however, his interest slowly abates again. How to capture and represent this dynamic user interest is an unsolved problem. Moreover, a user can be interested in multiple topics, which might evolve over time. Instead of being interested in only one topic at one time, users can search for various independent topics such as politics or sports, followed by entertainment or business.

In this paper, we introduce NewsBoy, a personalised multimedia application which is designed to capture the user's evolving interest in multiple aspects of news stories. NewsBoy automatically processes the daily BBC One news bulletin and recommends news stories by unobtrusively profiling the user based on his interactions with the system. The news aspects are identified by clustering the content of the profile. We introduce four different functions that incorporate the evolving interest of the user and evaluate the effect of these functions on the profiles.

The paper is structured as follows: Section 2 provides an overview of related work. Section 3 introduces the architecture of NewsBoy. In Section 4, we introduce our approach of capturing the user's interactions by extracting relevant

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the VLDB copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Very Large Data Base Endowment. To copy otherwise, or to republish, to post on servers or to redistribute to lists, requires a fee and/or special permission from the publisher, ACM.

VLDB '08, August 24-30, 2008, Auckland, New Zealand
Copyright 2008 VLDB Endowment, ACM 000-0-00000-000-0/00/00.

terms from results a user interacted with, combining them with a relevance weighting and storing them in a user profile. In order to capture the user's evolving interest, we adopt the ostensive model to manipulate the weighting of the terms in accordance to the iteration when they were added to the profile. Further, we introduce our methodology of clustering these terms to represent the user's multiple interests in different aspects and present a preliminary evaluation in Section 5. Finally, we discuss the system in Section 6.

2. BACKGROUND

Our work builds on a number of research areas, including news video retrieval, personalised news delivery and techniques to capture evolving user needs. In the following, we introduce the state-of-the-art of these areas.

2.1 News Video Retrieval

Nowadays, almost every television channel has its own news bulletin, indicating that television is a widely accepted mass media to provide consumers with the latest news. Consequently, processing television news has been an important research area and much recent work, such as that represented by the TRECVID [20] research effort, aims to tackle the difficult problems of content based video retrieval. While some systems have a particular emphasis on the system side, other research efforts are looking towards improving state-of-the-art video retrieval techniques from the user's point of view, such as the Open Video Project¹.

A number of conclusions can be drawn from these efforts: first of all, video retrieval is not as sophisticated as its textual counterpart. The reason for this is the so-called "semantic gap" [10], the difference between the low-level representation of video and audio data, and the high-level semantics which the user would ideally like to associate with retrieved data. Furthermore, segmenting and indexing video is a challenge. Considering a news broadcast as a unit of retrieval will generate a result list containing whole video documents. A user must watch or browse through the whole video to finally find the information he wants, a demanding approach. Hence, it is necessary to split videos into smaller, semantically related, *segments* which should ease the access of the video data. In text retrieval, techniques have been developed to identify relevant sections of the text, e.g. [17] and to segment documents based on these sections. Hence, users can easily browse through short results to satisfy their information need. Boreczky et al. [3] argue that television news consists of a collection of *story units* which represent the different events being relevant for the day of the broadcast. An example story unit from the broadcasting news domain is a report on yesterday's football match, followed by another story unit about the weather forecast.

Indexing these segments, i.e. based on textual annotations or visual representations of the segments provides an easy access to the data collection. A challenging approach however is to identify these stories a user is really interested in. The problem will be introduced in the following section.

2.2 Personalised News Delivery

Web 2.0 facilities enable everyone to easily create their own content and to publish it online. Users can upload videos on platforms such as YouTube, share pictures on

Flickr or publish anything in a weblog. Two direct consequences of this development can be identified: first of all, it leads to a growing quantity of content presented in a multimedia format. Secondly, information sources are completely unstructured and finding interesting content can be an overwhelming task. Hence, there is a need to understand the user's interest and to customise information accordingly.

A common approach to capture and to represent these interests is user profiling. Using user profiles to create personalised online newspapers has been studied for a long time.

Chen and Sycara [6] join internet users during their information seeking task and explicitly ask them to judge the relevance of the pages they visit. Exploiting the created user profile of interest, they generate a personalised newspaper containing daily news. However, providing explicit relevance feedback is a demanding task and users tend not to provide much feedback [8].

Bharat et al. [2] create a personalised online newspaper by unobtrusively observing the user's web-browsing behaviour. Although their system is a promising approach to release the user from providing feedback, their main research focus is on developing user interface aspects, ignoring the sophisticated retrieval issues.

Smeaton et al. [19] introduced Físchlár-News, a news video recommendation system that captured the daily evening news from the national broadcaster's main TV channel. The web-based interface of their system provides a facility to retrieve news stories and recommends stories to the user based on his interest. According to Lee et al. [13], the recommendation of Físchlár-News is based on personal and collaborative explicit relevance feedback. The use of implicit relevance feedback as input has not been incorporated.

Profiling and capturing the users is an important step towards adapting systems to the user's evolving information need. In the following section, we introduce the problem of capturing this evolving need.

2.3 Evolving User Needs

In a retrieval context, profiles can be used to contextualise the user's search queries within their interests and to re-rank retrieval results. This approach is based on the assumption that the user's information interest is static, which is however, not appropriate in a retrieval context.

Campbell [4] argues that the users' information need can change within different retrieval sessions and sometimes even within the same session. He states that the user's search direction is directly influenced by the documents retrieved. The following example explains this observation: Imagine a user who is interested in red cars and uses an image retrieval system to find pictures showing such cars. His first search query returns him several images including pictures of red Ferraris. Looking at these pictures, he wants to find more Ferraris and adapts the search query accordingly. The new result list now consists of pictures showing red and green Ferraris. Fascinated by the rare colour for this type of car, he again re-formulates the search query to find more green Ferraris. Within one session, the user's information need evolved from red cars to green Ferraris. Based on this observation, Campbell and van Rijsbergen [5] introduce the ostensive model which incorporates this change of interest by considering *when* a user provided relevance feedback. In the ostensive model, providing feedback on a document is seen as ostensive evidence that this document is relevant for

¹<http://www.open-video.org>

the user's current interest. The combination of this feedback over several search iterations provides ostensive evidence about the user's changing interest.

There are different types of interaction feedback, usually divided into two categories: explicit and implicit feedback. Explicit feedback is given when a user actively informs a system what it has to do on purpose, such as selecting something or marking it as relevant. Implicit feedback is given unconsciously. An example is printing out a web page, which may indicate an interest in that web page. The basic assumption is that during a search, users' actions are used to maximise the retrieval of relevant information. Implicit indicators have been used and analysed in other domains, such as the WWW [7] and text retrieval [12], but rarely in the multimedia domain. However, traditional issues of implicit feedback can be addressed in video retrieval since digital video libraries facilitate more interactions and are hence amenable to implicit feedback.

This section introduced the research domains of our work and argued about the research problem of capturing the evolving user need in order to personalise news videos in accordance to the user's interest in multiple aspects of the news. In the next section, we introduce NewsBoy, a news video retrieval system which incorporates the previously introduced research domains.

3. NEWSBOY ARCHITECTURE

NewsBoy is a web based news video retrieval system based on AJAX technology. AJAX takes away the burden of installing additional software on each client (assuming that JavaScript is activated and a Flash Player running on the client side).

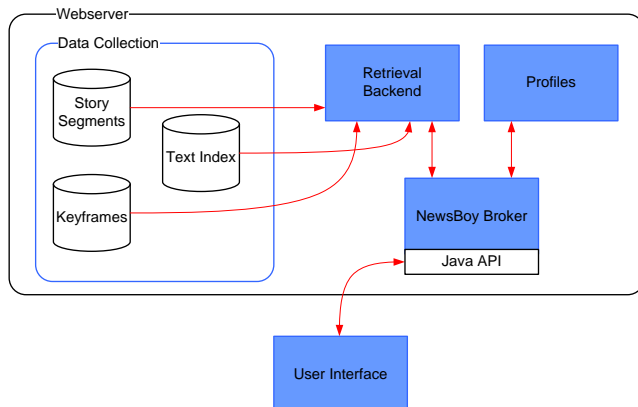


Figure 1: NewsBoy Architecture

Figure 1 illustrates the conceptual design of the system. As the graphic shows, NewsBoy can be divided into five main components, four running on a web server and one, the user interface, on the client side. The first component is the data collection which will be introduced in Section 3.1. The retrieval backend, the second component of NewsBoy, administers the data collection. We are using MG4J², an open source full-text search engine. The third component is the user interface, which runs on the client side. It will be introduced in Section 3.2.

²<http://mg4j.dsi.unimi.it/>

3.1 Data Collection

In the scope of this research, we focus on the regional version of the BBC One O'Clock news. The programme covers international, national (UK) and regional (Scotland) topics, which are usually presented by a single newsreader. The bulletin has a running time of 30 minutes and is broadcasted every day from Monday till Friday on BBC One, the nation's main broadcasting station. The BBC enriches its television broadcast with Ceefax, a closed caption (teletext) signal which provides televisual subtitles for the deaf. The data collection we used for this study consists of 115 editions of the daily news broadcast which have been recorded constantly over the past few months. Based on its textual, visual and audio features, we segmented the news videos into semantically related story segments, the unit of retrieval in our system. The index contains 2963 stories, which are aligned with 4.1 non-stopword-terms on average.

During the period of the recording, various main events have been dominant in the news. In relation to the evaluation date of this study (April 2008), these events can be classified into (1) latest, (2) recent and (3) past events. Here, we give some examples:

1. Latest events (current week): Discussions about air travel.
2. Recent events (2 month ago): Reports about the International Bank Crisis.
3. past events (>4 month ago): Christmas time

3.2 Interface

Figure 2 shows a screenshot of the NewsBoy interface, its features will be described in the following section. The interface can be divided into three main panels, search panel (A), result panel (B) and clustered search queries (C).

In the search panel (A) users can formulate and carry out their searches by entering a search query and clicking the button to start the search. BM25 [16] is used to rank the retrieved documents in accordance to their relevance to a given search query.

Once a user logs in, NewsBoy displays the latest news stories in the result panel (B). Moreover, this panel lists retrieval results. The panel displays a maximum of 15 results, further results can be displayed by clicking the annotated page number (1). The results can be sorted in accordance to their relevance to the query or chronologically by their broadcasting date (2). Results are presented by one keyframe and a shortened part of the text transcript. A user can get additional information about the result by clicking on either the text or the keyframe. This will expand the result and present additional information including the full text transcript, broadcasting date, time and channel and a list of extracted named entities³ such as persons, locations and relative times (3). In the example screenshot, the second search result has been expanded. The shots forming the news story are represented by animated keyframes of each shot. Users can browse through these animations by clicking on the keyframe. This action will center the selected keyframe and surround it by its neighboured keyframes. The keyframes are displayed in a fish-eye view (4), meaning that

³We use the General Architecture for Text Engineering (<http://gate.ac.uk>) for the extraction of named entities.

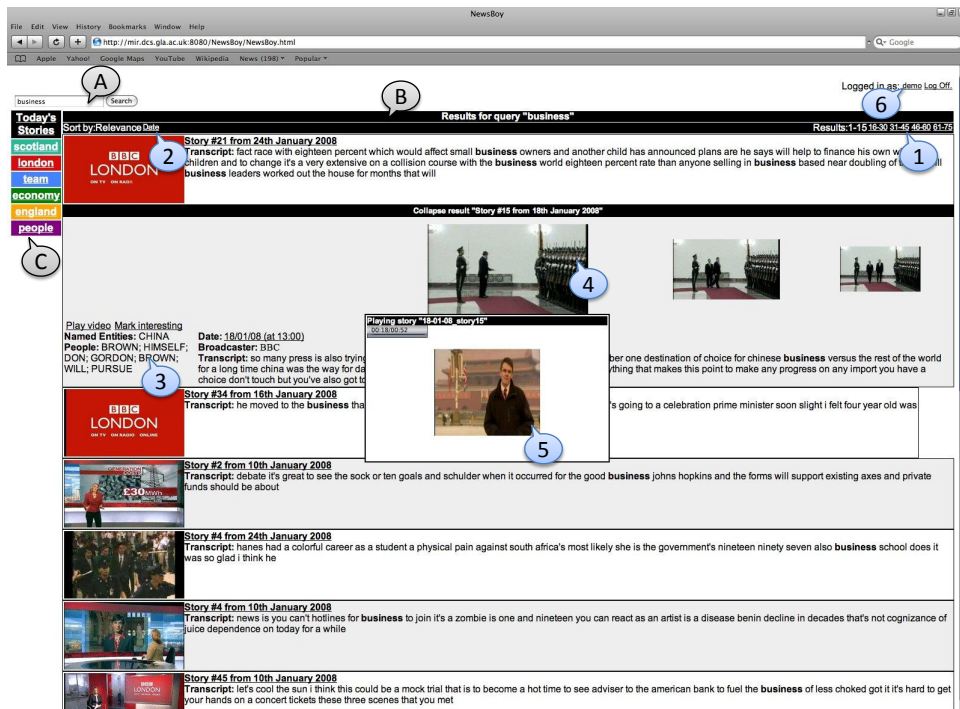


Figure 2: NewsBoy Interface

the size of the keyframe grows larger the closer it is to the focused keyframe. In the expanded display, a user can also select to play a video or to mark it as interesting. Clicking on “play video” starts playing the story video in a new panel (5).

NewsBoy recommends daily news videos based on the user’s multi-aspect preferences. These preferences are captured by unobtrusively observing the user’s interactions with the NewsBoy interface. By clustering the content of the profile NewsBoy identifies different topics of interest and recommends these topics to the user. The personalisation approach will be introduced in Section 4. The interface presents these topics as labelled clusters on the left hand side of the interface (C). Each cluster represents a group of terms, hence, when a user clicks on the term, a new search is triggered, using the selected terms as a new query. Results are displayed in the result panel.

On the top of the interface, the users can edit their profile by clicking on their username (6). This action will pop up a new frame where the top weighted terms of each cluster are listed, and the user can edit terms or the aligned weighting. Furthermore, the user can manually add new weighted terms.

In this section, we introduced the basic components of a video retrieval system, the frontend and the backend. These components enable the users to explore the indexed data collection. In the next section, we introduce our methodology of enhancing the users’ search sessions by adapting the output of the system to their personal interests.

4. PERSONALISATION

The aim of NewsBoy is to deliver daily news videos based on the user’s interest in multiple aspects of daily news. This

procedure raises some research questions. The main question is how the user’s interest can be captured and represented. Furthermore, we are interested how multiple interests can be identified. A common approach is to interpret the user’s interactions with the system’s interface and to represent this interest in a profile. The process of gathering these interactions will be introduced in Section 4.1. Our approach of incorporating the user’s evolving interest will be shown in Section 4.2. In Section 4.3, our approach of identifying multiple interests will be explained.

4.1 Relevance Feedback

O’Sullivan et al. [15] evaluated the use of explicit and implicit relevance feedback to recommend video stories. Their results indicate that user profiles created by exploiting implicit feedback are as valuable as profiles which are created by incorporating explicit feedback only. Hopfgartner and Jose [9] identified various implicit indicators of relevance in video retrieval when comparing the interfaces of state-of-the-art video retrieval tools. The most common features they identified were: clicking on a keyframe to start playing a video, browsing through a result list, using the sliding bar to go through a video, highlighting additional metadata and playing a video for a certain amount of time. However, analysing which of these implicit measures are useful to infer relevance has rarely been done.

NewsBoy tries to capture the users’ interests by exploiting the implicit relevance feedback captured from users interacting with the interface introduced in Section 3.2. The interface provides various possibilities to provide implicit relevance feedback. Users interacting with it can:

- Expand the retrieved results by clicking on it.
- Play the video of a retrieved story by clicking on “play

video”.

- Play the video for a certain amount of time.
- Browse through the keyframes.
- Highlight additional information by moving the mouse over the keyframes.

Any of these interface features can be seen as a possible indicator of relevance. Which one of these implicit measures are good indicators of relevance is not clear though. While Claypool et al. [7] identified time spend on a web site as being a valid implicit indicator of relevance in the text domain, Kelly and Belkin [11] criticise the time factor as indicator in the video domain. They assume that information-seeking behaviour is not influenced by contextual factors such as topic, task and collection. Their study cast doubt on the straightforward interpretation of dwell time as an indicator of interest or relevance. Hence, we decided to ignore the playing duration as a positive indicator and focus on the remaining indicators only.

Further research has to be done to find the strongest indicators in order to identify optimal an weighting for each interface feature. This is, however, not the focus of this work. Based on the analysis of implicit relevance feedback weight by Hopfgartner and Jose [9], we therefore define a static value for each possible feature:

$$W = \begin{cases} 0.1, & \text{when a user uses the highlighting feature} \\ 0.2, & \text{when a user starts playing a video} \\ 0.3, & \text{when a user browses through the keyframes} \\ 0.5, & \text{when a user expands a result} \end{cases}$$

For capturing the users’ interest, NewsBoy extracts the (non-stopword) query terms aligned with the story item a user interacted with, combines them with the feedback weighting and stores the weighted terms in the profile. The following example explains the process: A user retrieved a list of stories and decides to expand the first result. Capturing this action, NewsBoy extracts all terms aligned with this result, combines them with the weighting 0.5 in a vector and submits this vector to the profile. A more in-depth description of this profiling is given in the following section.

4.2 Profile

User profiling is the process of learning the user’s interest over a longer period of time. In this section, we introduce our approach of capturing the users’ interest and introduce the representation of this interest in the profile. The profile is the fourth component of the NewsBoy system illustrated in Figure 1. Furthermore, we introduce our approaches of representing the user’s evolving focus of interest.

4.2.1 Profile Learning and Representation

Several approaches have been studied to capture a user’s interest in a profile, the most prominent being the weighted keyword vector approach. In this approach, interests are represented as a vector of weighted terms where each dimension of the vector space represents a term aligned with a weighting. The weighting of the terms will be updated when the system submits a new set of weighted terms to the profile starting a new iteration j . Hence, we represent the interaction I of a user i at iteration j as a vector of weights

$$\vec{I}_{ij} = \{W_{ij1} \dots W_{ijv}\}$$

where v indexes the word in the whole vocabulary $|V|$.

We create a weighting W_{ij} by capturing the implicit relevance feedback provided by a user i in the iteration j with the interface introduced in Section 3.2. W has been introduced in detail in Section 4.1. Representative terms from relevant story segments will be extracted and assigned with an indicative weight to each term, which represents its weight in the term space. In our model, we extract non-stopwords v from the stories a user interacted with in the iteration i and assign these terms with the relevance weighting W_{ijv} .

Furthermore, we represent the profile \vec{P}_i of user i as a vector containing the profile weight PW of each term v of the vocabulary:

$$\vec{P}_i = \{PW_{i1} \dots PW_{iv}\}$$

4.2.2 Ostensive Factor

The simplest approach to create a weighting for each term in the profile is to combine the weighting of the terms over all iterations. This approach is based on the assumption that the user’s information interest is static, which is, however, not appropriate in a retrieval context. The users’ information need can change within different retrieval sessions.

Campbell and van Rijsbergen [5] propose in their ostensive model that the time factor has to be taken into account, i.e. by modifying the weighting of terms based on the iteration they were added to the user profile. They argue that more recent feedback is a stronger indicator of the user’s interest than older feedback. In our profile, the profile weight for each user i is the combination of the weighted terms v over different iterations j : $PW_{iv} = \sum_j a_j W_{ijv}$. We include the ostensive factor, denoted a_j , to introduce different weighting schemes based on the ostensive model. We have experimented with four different functions to calculate the weighting, depending on the nature of aging, the functions will be introduced in the following paragraphs.

4.2.2.1 Constant Weighting.

$$a_j = \frac{1}{j_{max}} \quad (1)$$

The constant weighting function does not influence the ostensive weighting. As Equation 1 illustrates, all terms will be combined equally, ignoring the iteration when a term was added or updated. The constant weighting can be seen as a baseline methodology which does not include any ostensive factor.

4.2.2.2 Exponential Weighting.

$$a_j = \frac{C^j}{\sum_{k=1}^{j_{max}} C^k} \quad (2)$$

The exponential weighting as defined in Equation 2 gives a higher ostensive weighting to terms which has been added or updated in older iterations. It is the most extreme function as the ostensive weighting of earlier iterations decreases distinctly.

4.2.2.3 Linear Weighting.

$$a_j = \frac{Cj}{\sum_{k=1}^{j_{max}} Ck} \quad (3)$$

Equation 3 defines the linear weighting function. The ostensive weighting of earlier iterations decreases linearly. This function linearly reduces the ostensive weighting of earlier iterations.

4.2.2.4 Inverse Exponential Weighting.

$$a_j = \frac{1 - C^{-j+1}}{\sum_{k=1}^{j_{max}} 1 - C^{-k+1}} \quad (4)$$

The inverse exponential weighting defined by Equation 4 is the most contained function. Compared to the other introduced functions, the ostensive weighting of early iterations decreases more slowly.

4.3 Capturing Multiple Interests

All components introduced in the previous sections communicate through the NewsBoy Broker, the fifth component of the system illustrated in Figure 1. The task of the broker is to personalise the system by identifying the user’s multiple interests in different aspects. Our methodology of identifying these aspects is introduced in the following.

our approach is based on the assumption that news topics consist of a number of *unique* terms which appear in all stories about one topic. News stories about the topic *football* e.g. might consist of unique terms such as “goal”, “offside”, “match” or “referee”. We capture implicit feedback when a user interacts with these stories. The terms of these stories will be extracted and, combined with the implicit weighting, stored in the profile. Hence, as the particular terms are added with the same weighting, they are close neighbours in the profile’s vector space.

In this work, we sort the terms in the user’s profile according to their profile weighting and identify the terms which have the five biggest distances to the neighbouring terms. We use these identified weighted terms to cluster the remaining profile terms accordingly. Each cluster represents *one* aspect of the user’s interest.

The top weighted terms of each cluster are used as a label to visualise the aspect on the left hand side of the NewsBoy interface (marked (C) in Figure 2). In this work, we limited the number of terms to six. Clicking on this label hence triggers a retrieval with the top six weighted terms of this aspect being used as search query. The effect of the different weighting factors on the user’s profile will be illustrated in the following section.

4.4 Weighting Effect

In Section 4.2.2, we introduced four different profile weighting approaches that capture the evolving user need by incorporating the ostensive model. In Section 4.3 we introduced our approach of clustering the terms based on this weighting. In this section, we illustrate the effect of the different weighting factors on the user’s profile by simulating users interacting with the NewsBoy interface over several days. Simulations are an alternative methodology of evaluating different approaches to user-modelling. In this methodology, we assume that a user is interacting on the system. If such a user is available, he or she will carry out a set of actions to retrieve or look at relevant results. The aim of our simulation will be introduced in the following section.

4.4.1 Simulated User Interaction

C	Const.		Exp.		Lin.		Inv. Exp.	
	Term	W	Term	W	Term	W	Term	W
1	people	1	people	1	news	1	people	1
					flight	0.99		
					fuel	0.99		
					change	0.99		
					connecting	0.99		
					passengers	0.99		
houston	0.99							
2	christmas	0.74	christmas	0.77	morning	0.20	thousand	0.84
					people	0.19		
3	thousand	0.64	thousand	0.67	recent	0.14	christmas	0.76
					past	0.14		
4	house	0.48	twenty	0.49	sounds	0.11	twenty	0.65
					home	0.10		
					company	0.10		
					thousands	0.10		

Table 1: Top four clusters in the simulated user profile for the constant, exponential, linear and inverse exponential ostensive weighting functions.

In order to get an insight into the effect of the four different ostensive weighting functions on the clusters, we simulate a user interacting with stories of each day of our data collection. In a first step, we retrieve all stories on a particular date, starting with the oldest recording available. We then simulate a user interacting with these results by randomly selecting x stories, where $0 \leq x \leq (\# \text{ of stories})$. In the next step, the simulated user can (a) start playing a video, (b) expand a result and (c) use the highlighting feature. Each of these events has an equal probability of 33%. Furthermore, the simulated user browses up to ten times through the keyframes, each browsed with a probability of 10%. Each simulated action will start the profiling process which has been introduced in Section 4.1. The same simulation is repeated for all days of our data collection. A possible search session i.e. could be: A user expands a result, plays a video and browses through three keyframes.

We are aware that a user does not “randomly” select results and that the probability of using a feature is not always 50%. However, our aim is to evaluate the different ostensive weighting functions, which are not user-dependent. Hence, we decided to base our simulation on a simplified user model.

4.4.2 Profile Content

Exploiting the simulated user profile, we clustered the terms based on the different ostensive factors. Table 1 illustrates the four top clusters for the constant, exponential, linear and inverse exponential ostensive weighting functions.

Assuming that the profiles represent the interests of the simulated user, some observations can be stressed when analysing the top terms stored in the profile. While the profiles which are weighted using the *constant*, *exponential* and *inverse exponential* functions show similar clusters, the biggest difference can be spotted in the profile which incorporates the *linear* ostensive function. This profile seems to emphasise latest events, a news story related to air flights. Past events such as christmas do not appear in the top clusters of this profile, however, the term “christmas” has a high ranking in the remaining profiles, indicating that past events are still represented in the clusters.

The example profiles hence confirm our expectation that the introduced ostensive factors will set a different emphasis on added terms, based on the time when the terms were added. However, which of these profiles represents the user’s current interest cannot be answered by this study. A meaningful interpretation requires knowledge about the users preferences which can only be achieved by a user study. In the following section, we introduce a subsequent user study we

performed.

5. PRELIMINARY EVALUATION

In order to evaluate which clusters created by the different profiling methodologies can best represent the user’s multiple interests, we designed a user-centred evaluation. Nine participants of different nationality volunteered to include NewsBoy as an additional source in their daily news gathering process. For one month, the participants used NewsBoy to browse through the displayed news stories or to discover the data collection.

As we used our own data collection which grows every day and do not provide the users with pre-defined search topics, an evaluation based on precision and recall, as common in evaluation campaigns such as TRECVID is not possible. Therefore, we aimed to evaluate the models based on the user’s satisfaction.

5.1 Participants

The participants were mostly postgraduate students and research assistants with a background in computing science. The group consisted of eight male and one female with an average age of 27.7 years and advanced proficiency with English.

Prior to the experiment, each participant was asked to fill out a questionnaire so that we could measure their experience of dealing with news media and personalisation systems. The most cited topics of interest are sports, politics, science and entertainment. The group follow news by watching television once or twice a month and mainly use the internet as their daily source of information. The BBC and Google News websites were mostly cited as favourite source, followed by websites of national newspapers such as www.spiegel.de or www.elmundo.es. Furthermore, they use the internet to occasionally watch news videos online. The most common search strategy that the participants mentioned was browsing their favourite websites.

The questionnaire revealed a clear tendency towards news personalisation systems. However, the participants stated that they are sensitive about the type of feedback they have to provide in order to get personalised news. Privacy is considered to be an important issue, hence, the participants would not agree to provide details about the income or the private address. For the group, providing explicit relevance feedback is an unpopular approach, supporting our methodology of relying on implicit relevance feedback to adapt to the user’s interests.

Summarising, the participants mostly rely on the internet to follow daily news and are open-minded towards personalisation systems.

5.2 Objective Profile Evaluation

During the study, the users performed an average of 9.8 implicit actions each day which triggered the profiling process introduced in Section 4.3. Hence, the ostensive weighting of the profile terms introduced in Section 4.2.2 is computed over 190 iterations on average which strongly influences the different weighting approaches.

Table 2 shows the average number of terms when clustering the terms of the users’ profiles based on their ostensive weighting as introduced in Section 4.3. As can be seen, clustering the terms based on their ostensive weighting results in different clusters, supporting the assumption that the previ-

C	Const.		Exp.		Lin.		Inv. Exp.	
	Mean	Median	Mean	Median	Mean	Median	Mean	Median
1	1	1	1.125	1	5.625	1	1	1
2	1.125	1	2.375	1	9.625	7	1	1
3	2.375	1	7.25	1	16.25	19.5	1.75	2
4	184.375	2	34.125	4	8.375	3.5	2.25	2
5	166.375	7.5	434.5	9.5	24.125	22.5	3.75	2.5

Table 2: Number of terms of the top four clusters C for each ostensive weighting

ously introduced ostensive factor influences the users’ profile. However, two drawbacks can be seen in the table. First of all, many clusters consist of few words only, indicating that the introduced methodology of identifying multiple interests is not appropriate and needs to be further investigated. Moreover, some clusters consist of a large amount of terms which hardly represent any specific interest of the users.

While the table confirms the effect of the ostensive model on user profiling, a conclusion about the quality of these profiles cannot be drawn. Therefore, we further focused on the users’ subjective opinion about the content of the profile. The evaluation will be introduced in the following section.

5.3 Subjective Profile Evaluation

First of all, we were interested to identify which ostensive weighting function best represents the users’ interests. Hence, we used the constant weighting factor introduced in Section 4.2.2 to create the user profile. Thus, the weighting of the terms in the profile is not influenced by an ostensive factor. At the end of the experiment, we clustered the user profile based on the exponential, linear and inverse exponential factor, respectively, and asked the participants to judge, which of these clustered profiles represents their information need best.

In the first question, we asked our participants to judge which profile is the most efficient one in clustering the terms in accordance to their semantic meaning. This question was aimed to analyse whether our assumption that news stories consist of a number of unique terms which appear in all stories of the topic can be applied to identify semantically related terms. The participants did not highlight any particular profile, indicating that the different weighting schemes semantically cluster terms in a similar way.

In the next question, we asked them to judge which profile identified best their interest. Here, the participants preferred the profile created using the inverse exponential weighting function, followed by the constant weighting and linear weighting. The exponential weighting received the lowest ranking, indicating that the approach of giving a higher weight to most recent feedback does not cover the user’s long term interest.

In a follow up question, we were interested if the order of the clusters in the profiles represent the participants’ interest accordingly. Again, the users showed a tendency towards the inverse exponential weighting function, followed by the constant and linear weightings.

Concluding, the questionnaires revealed a slight preference towards the model which privileges most recent feedback.

6. DISCUSSION AND CONCLUSIONS

In this paper, we address two main research challenges in the field of information retrieval. The first problem we

introduce is how to capture and represent a user's evolving information need. As Campbell [4] argues, the users' information need can change within different retrieval sessions and sometimes even within the same session. The user's search direction is directly influenced by the documents retrieved. So far, capturing and representing this dynamic user interest is an unsolved problem. Another question is how the different aspects of a user's interest can be represented. A user can be interested in various aspects, which also might evolve over time.

In order to study these problems, we introduced NewsBoy, a news video retrieval system which delivers news videos based on the user's interest. NewsBoy captures the user's interactions by extracting relevant terms from results a user interacted with. These terms are combined with an explicit relevance weighting and stored in a user profile. A user's interest in multiple aspects is identified by clustering the profile based on this weighting. We introduced four different models that incorporate the ostensive model to capture the evolving interest of the user. For each model we show their effect on user profiling by conducting a simulated user study. In addition to this, we performed a user study to evaluate these models based on the user's satisfaction. The study indicates the user's preferences against the model which privileges most recent feedback.

In conclusion, our results have highlighted that the ostensive model can be incorporated to represent the users' interests in video retrieval. While we present in this paper a preliminary user evaluation, we plan to further analyse the users' feedback, i.e. by exploiting the log files, which should help to investigate in the introduced research questions.

7. ACKNOWLEDGMENTS

This research was supported by the European Commission under the contracts FP6-027026-K-SPACE and FP6-027122-SALERO. It is the view of the authors but not necessarily the view of the community.

8. REFERENCES

- [1] R. Arezki, P. Poncet, G. Dray, and D. W. Pearson. *Adaptive Hypermedia and Adaptive Web-Based Systems*, chapter Web Information Retrieval Based on User Profiles, pages 275–278. Springer Berlin / Heidelberg, 2004.
- [2] K. Bharat, T. Kamba, and M. Albers. Personalized, interactive news on the web. *Multimedia Systems*, 6(5):349–358, 1998.
- [3] J. S. Boreczky and L. A. Rowe. Comparison of Video Shot Boundary Detection Techniques. In *Storage and Retrieval for Image and Video Databases (SPIE)*, pages 170–179, 1996.
- [4] I. Campbell. Supporting information needs by ostensive definition in an adaptive information space. In *MIRO'95 – Workshops in Computing*. Springer Verlag, 1995.
- [5] I. Campbell and C. J. van Rijsbergen. The ostensive model of developing information needs. In *Proc. of CoLIS-96, 2nd Int. Conf. on Conceptions of Library Science*, pages 251–268, 1996.
- [6] L. Chen and K. Sycara. WebMate: A personal agent for browsing and searching. In K. P. Sycara and M. Wooldridge, editors, *Proceedings of the 2nd International Conference on Autonomous Agents (Agents'98)*, pages 132–139, New York, 9–13, 1998. ACM Press.
- [7] M. Claypool, P. Le, M. Wased, and D. Brown. Implicit interest indicators. In *Intelligent User Interfaces*, pages 33–40, 2001.
- [8] M. Hancock-Beaulieu and S. Walker. An evaluation of automatic query expansion in an online library catalogue. *J. Doc.*, 48(4):406–421, 1992.
- [9] F. Hopfgartner and J. Jose. Evaluating the Implicit Feedback Models for Adaptive Video Retrieval. In *ACM MIR '07*, pages 323–332, 09 2007.
- [10] A. Jaimes, M. Christel, S. Gilles, S. Ramesh, and W.-Y. Ma. Multimedia Information Retrieval: What is it, and why isn't anyone using it? In *MIR '05*, pages 3–8, New York, NY, USA, 2005. ACM Press.
- [11] D. Kelly and N. J. Belkin. Display time as implicit feedback: understanding task effects. In *SIGIR '04: Proceedings of the 27th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 377–384, New York, NY, USA, 2004. ACM Press.
- [12] D. Kelly and J. Teevan. Implicit Feedback for Inferring User Preference: A Bibliography. *SIGIR Forum*, 32(2), 2003.
- [13] H. Lee, A. F. Smeaton, N. E. O'Connor, and B. Smyth. User evaluation of Físchlár-News: An automatic broadcast news delivery system. *ACM Trans. Inf. Syst.*, 24(2):145–189, 2006.
- [14] D. M. Nichols. Implicit rating and filtering. In *Proceedings of 5th DELOS Workshop on Filtering and Collaborative Filtering*, pages 31–36. ERCIM, 1998.
- [15] D. O'Sullivan, B. Smyth, and D. C. Wilson. Explicit vs Implicit Profiling - A Case-Study in Electronic Programme Guides. In *IJCAI'03 – Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence, Acapulco, Mexico*, pages 1351–1353, 08 2003.
- [16] S. E. Robertson, S. Walker, S. Jones, M. Hancock-Beaulieu, and M. Gatford. Okapi at TREC-3. In *Proceedings of the Third Text Retrieval Conference (TREC 1994), Gaithersburg, USA, 1994*.
- [17] G. Salton, J. Allan, and C. Buckley. Approaches to passage retrieval in full text information systems. *ACM SIGIR conference on research and development in Information Retrieval*, pages 49–58, 1993.
- [18] N. Sebe and Q. Tian. Personalized Multimedia Retrieval: The New Trend? In *MIR '07*, pages 299–306, New York, NY, USA, 2007. ACM.
- [19] A. F. Smeaton. The Físchlár Digital Library: Networked Access to a Video Archive of TV News. In *TERENA Networking Conference 2002, Limerick, Ireland, 3-6 June 2002*, 2002.
- [20] A. F. Smeaton, P. Over, and W. Kraaij. Evaluation campaigns and trecvid. In *MIR '06: Proceedings of the 8th ACM International Workshop on Multimedia Information Retrieval*, pages 321–330, New York, NY, USA, 2006. ACM Press.
- [21] A. Spink, H. Greisdorf, and J. Bateman. From highly relevant to not relevant: examining different regions of relevance. *Inf. Process. Manage.*, 34(5):599–621, 1998.