AUTOMATIC CONSTRUCTION OF PERSONALIZED TV NEWS PROGRAMS

Bernard Merialdo, Kyung Tak Lee, Dario Luparello, Jeremie Roudaire
Multimedia Communications Department
Institut EURECOM
BP 193, 06904 Sophia-Antipolis, France
merialdo@eurecom.fr

ABSTRACT

In this paper, we study the automatic construction of personalized TV News programs, where we want to build a program with predefined duration and maximum content value for a specific user. We combine video indexing techniques to parse TV News recordings into stories, and information filtering techniques to select stories which are most adequate given the user profile. We formalize the selection process as an optimization problem, and we study how to take into account duration in the selection of stories. Experiments show that a simple heuristic can provide high quality selection with little computation. We also describe two prototypes, which implement two different mechanisms for the construction of user profiles:

- explicit specification, using a category-based model,
- implicit specification, using a keyword-based model.

1. INTRODUCTION

The development of Digital Television opens new perspectives in the distribution of audio-visual material to the general public. The immediate advantages of digital broadcast are the improvement in transmission quality and the increase in transmission capacity. But the major change for users will come from the future capacity of the video delivery chain to process this digital information to build new interaction paradigms, such as Interactive Television. In particular, it is expected that one important paradigm will be the construction of customized programs, programs that are specifically designed to fit the needs of each user. Within this perspective, TV News are a good candidate for such a customization process. First, because TV News are a very successful type of program (so successful that channels such as CNN are entirely devoted to this type). Second, because the current broadcast paradigm of TV News programs is very rigid: these programs have typically a almost fixed duration (generally from 30 to 45 minutes), a fixed schedule (the midday program, the introduction of the prime time, or the late night), similar contents (while different channels have different styles, the ordering of topics within TV News frequently follows general rules: highlights first, then important events, then invited person, terminate with sport and weather). In an ideal situation, every user would expect to have a TV News program at a time of his choice, with a duration that corresponds to the time that this user has available, and with a content that specifically matches this user's interests.

In this paper, we present two research prototypes which aim at investigating the technical issues involved in the construction of personalized news programs. Our focus is twofold:

- on one side, we are interested in the development of video indexing techniques, not necessarily specific to TV News programs, which allows to extract content information from the raw data, and the identification of useful metadata information to describe this content,
- on another side, we explore the customization process itself, which uses content information and knowledge of the user's interests to build a custom program.

We are currently not interested in other aspects such as efficient delivery mechanisms, optimization of the creation process, economical issues etc… The following figure illustrates the various components involved in the construction process. This paper concentrates on the Annotation and Selection steps.

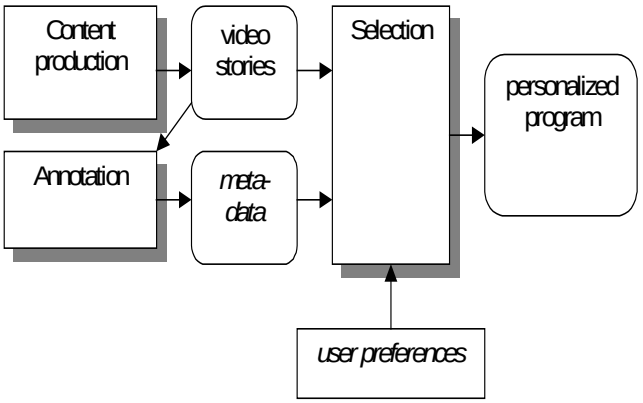

Our approach of Annotation is based on the automatic segmentation of TV News recordings into stories, plus two different content-description schemes:

- manual classification of stories into predefined categories,
- or automatic description by keywords through the automatic analysis of captions.

Personalization can be achieved through two paradigms:

- explicit: by which the user directly defines numerical values for his interest in categories,
- implicit: where a user profile is automatically updated based on the selection of stories made by the user.

Our work combines existing techniques both from video indexing and information filtering. One contribution of this paper is that we consider that the customized program has a predefined duration which has been specified by the user. The problem is therefore to select the most appropriate informations to fill this duration with a content of maximum value. We study closely the relation between the importance of a story and its duration, and we derive an evaluation criterion for inclusion. We also propose a heuristic approach for the selection of the best set of stories. This duration constraint is generally not taken into account in filtering systems, which rather aim at making a binary decision for relevant vs non-relevant information.

2. RELATED WORK

Video Indexing is a very active research field, which covers a large number of issues, among which the most fundamental are[2]:

- cut detection algorithms,
- shot classifications,
- keyframe selection,
- scene determination,
- navigation paradigms.

TV News have already been extensively studied[12][6][1], because they provide interesting structural information to segment a recording into stories. Most of the specific work on this material relies on finding the anchor shots. Quite often, TV news parsing can include audio, video and text (captions) processing at the same time, to build segmented stories which can be collected into a multimedia database[7][8].

Information filtering is a specific task within the scope of Information Retrieval. Many approaches have been proposed to achieve this task[9], from Natural language parsing to probabilistic networks, one of the most popular being the standard Vector Space model[10]. Many Internet sites allow for personalized delivery of (textual) news (CNN, PointCast), but the mechanisms that are currently used in regular services generally consist in a subscription to specific channels of information which provide news that are organized by topics.

This work is also related to the problem of summarization, because summarization involves a similar constraint of fitting the most important data within a predefined space or duration. Work is active in this area also, both on the video[5][11] and the textual side[4].

3. TV NEWS VIDEO PARSING

A number of projects have already addressed the problem of TV news indexing. TV News programs are of particular interest for video indexing experiments because they exhibit a natural structure: such a program is composed of news stories, introduced by one or more anchorpersons. The periodic appearance of the anchorperson greatly facilitates the automatic discovery of story boundaries. In previous work[6], we have presented our approach for the analysis of TV News programs, based on the following sequence of steps:

- cuts are detected in the video by computing the difference of color histograms of consecutive frames,
- keyframes are selected for each shot,

- similar keyframes are found by automatic clustering of the set of keyframes,
- shots containing persons are detected, based on the quantity and location of movement in the video),
- specific shots such as commercials are removed,
- anchor person shots are identified by combining shot similarity, person detection and a "high variance" factor (which reflects the fact that the anchor person occurs at regular spots throughout the video recording).

The result of this processing is a segmentation of the TV News program into a set of consecutive stories. A XML file is created to hold this information. An extract of such a file would look as in the following example (timing information is based on video frame numbers):

```
<STORY KEYFRAME=7645 DURATION=75>
<SHOT START=7106 END=7569 PERSON ID=0/>
<SHOT START=7569 END=7722 PERSON ID=6/>
<SHOT START=7722 END=7853/>
<SHOT START=7853 END=8028 PERSON ID=6/>
<SHOT START=8028 END=8340/>
<SHOT START=8340 END=8411/>
<SHOT START=8411 END=8455 PERSON ID=6/>
<SHOT START=8455 END=8473/>
<SHOT START=8473 END=8489 PERSON ID=6/>
<SHOT START=8489 END=8503 PERSON ID=0/>
<SHOT START=8503 END=8524/>
<SHOT START=8524 END=8944 PERSON ID=6/>
<SHOT START=8944 END=8985 PERSON ID=0/>
</STORY>
```

An illustration of the result of this processing is shown in the figure below, where a TV News recording has been automatically segmented into stories, after the anchor person has been found.

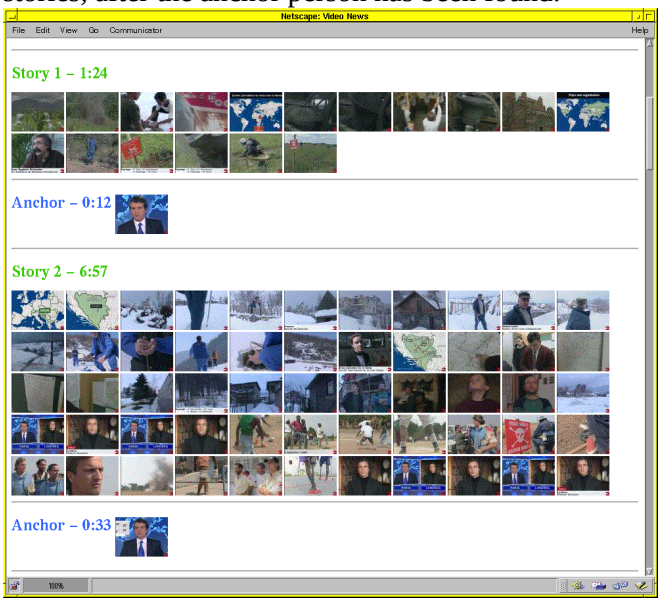

Such an analysis can be used to construct an interface which allows the user to have a visual overview of the contents of one or more recordings, as is shown in the figure below, where the most important stories of 6 consecutive issues of CNN World News as presented. Using hyperlinks, the user can immediately visualize the content of a particular story that he feels interested in.

Although this interface allows interactive browsing, it does not correspond to our requirement for a customized program, which we are going to handle now

4. CATEGORY-BASED SYSTEM

We now consider the question of the construction of a customized program. In this case, we assume that some video stories are available, together with some description of their content. The user is interested in the "best" program built out of these stories, which fits in the time that he has available. The problem is then to select a subset of the set of stories which:

- maximizes the interest of the user for the content of these stories,
- satisfies the constraint that the total duration should be less than or equal to the expected duration for the program.

We formalize this problem in the following way:

- each story s has a value $v_u(s)$ which measures the interest of this story for user u,
- each story s has a duration $d(s)$.

Since stories will be concatenated in the customized program, the duration of a set of stories is simply the sum of the duration of stories. We assume that the value of the set of stories will be also the sum of the value of stories (this contains two assumptions: that value is an additive measure, and that different stories are independent).

We are therefore looking for the subset S of stories which maximizes $v(S) = \sum_{s \in S} v_u(s)$ while respecting the constraint $d(S) = \sum_{s \in S} d(s) \leq D$ where D is the expected duration desired by the user.

With this formulation, the selection process is precisely equivalent to the resolution of a "knapsack" problem, which is unfortunately known to be a NP-complete problem. In order to reduce the complexity of the selection process, we will later propose a heuristic to construct reasonable solutions without a complete enumeration of all possible sets.

In our first approach, we want to implement an explicit interaction of the user on his profile. The user should have a limited number of controls to set, so we decided to rely on a category-based approach, where each story belongs to one category. A yes-no selection of categories, as is the case in most commercial systems, is not sufficient to define an adequate value measure $v_u(s)$, so that finer numerical weighting is required. More precisely, our first prototype has the following characteristics:

- a category and importance are manually assigned to each news story,
- the user can customize his own profile,
- a customized program can be build automatically.

For this experiment, we collected a set of TV News recordings (over one week, from a French TV channel). We applied the previous segmentation process to this data. This segmentation was manually checked, and corrected when necessary. In particular, we paid attention to announcements which were made by the anchor person himself, without the support of any video sequence. Such announcements cannot be detected by the video analysis component. This global analysis lead to a total of 137 stories. For each of these stories, a manual annotation was performed, to add the following information:

- a category which describes the general theme of the story. A set of 7 categories was designed:
  1. International politics,
  2. National politics,
  3. International society,
  4. National society,
  5. Economy,
  6. Culture,
  7. Sport.
- an importance factor, which is a number (from 0 to 100) which indicates the degree of importance of the subject covered by the story within the category.
- a title, which is useful for browsing applications, but could be omitted for custom program generation.

The figure below shows part of the repository containing these stories.

4.1 User customization

In the category-based scheme, the interest of the user for a news story is computed from the importance of the story within its category, and from the interest of the user in the category. This is conveniently modeled within a probabilistic formulation: if we assume that the interest of the user for stories defines a probability distribution over the set of stories

$$i_u(s) = p(s/u)$$

We can apply a simple decomposition formula to express $p(s/u) = p(s/c_s)\,p(c_s/u)$, where :

- $p(c_s/u)$ represents the interest of the user for a given category,
- $p(s/c_s)$ represents the importance ratio of the story within its category. It is computed by normalizing (so that they sum to one) the importance factors that have been previously defined.

The user profile is represented as a vector of weights which indicate the interest $p(c_s/u)$ of the user for the specific categories that have been defined. In our current prototype, we assume that the user provides these weights himself. The following interface shows how this can be practically done.

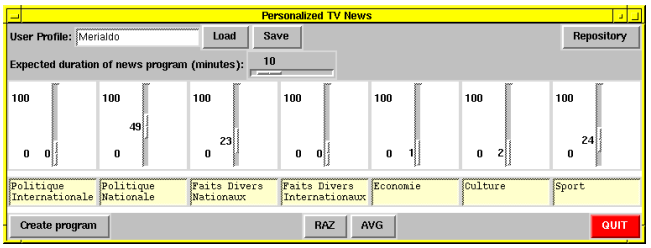

Each category is represented by a scale widget, which can be adjusted by the user. Since all weigths have to sum to 1, when the user manipulates a weight, this has an immediate effect on the other weights as well: if one weight is increased, all the others are decreased, and if one weight is decreased, all the others are increased. The amount of modification of the other weights is in proportion of their actual value, so, if the weight of category i_0 is changed from $w(i_0)$ to $w'(i_0)$ we use the update formula:

$$w'(i) = \left(1. - \frac{w'(i_0)}{\sum_{i \neq i_0} w(i)}\right) \cdot w(i) \qquad \forall i \neq i_0$$

In particular, it should be noted that if a weight is set to zero, which indicates that the user has no interest in this category, it remains set at zero by this formula.

We also found that it was useful to have two special buttons:

- one sets all weights to zero. While in this case the sum of the weights can no longer be one (it is zero), we found that it was useful when the user wanted to constructively indicate the categories he was interested in.
- Another button set all weights equal, this defines an average situation from which the user can specialize his own interests.

This interface is intended as a prototype tool, not a production system, and in particular, it does not have the attractive graphical design that would make it suitable for a large population of users. However, it contains the basic functionalities that are required to create custom programs:

- creation, storage or retrieval of a user profile,
- indication of the desired duration for the customized program,
- weights for the various categories.

4.2 Content value

We now discuss how to compute the content value $v_u(s)$ of a story from the interest $i_u(s)$ and its duration.

We have already indicated how the interest of a user for a given news story was computed from the importance factor of the story and the user profile. We noticed that the importance factor that was manually assigned relates to the topic, which is covered, but that it was not an adequate indicator for the selection of a story in a personalized program, because it does not take into account the duration of the story. It is clear that, to optimize the content of the customized program, it might not be optimal to include a story with very high interest, but also long duration. Duration relates to the amount of details, which are given, about this particular event. Although duration should be related to importance and value, because journalists should decide not to spend too much time on less interesting events, it is also the case that the amount of details which is provided for an event may depend on other factors such as competing events in the same journal, or amount of material which is available (when for example the event has just occurred but no image is yet available).

As an example, the graph below indicates the repartition of duration of stories in our repository.

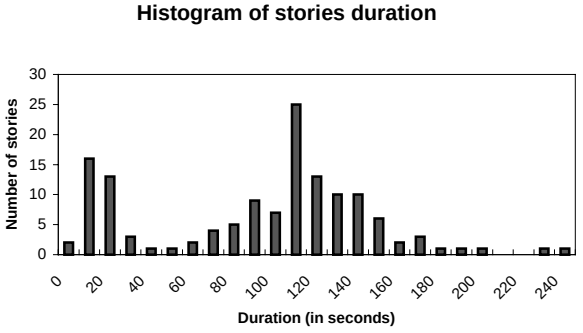

The previous arguments motivate that the value of a story should be computed both from the importance of the topic and the duration of the story:

$$v_u(s) = f(i_u(s), d(s))$$

We have seen that the importance is computed as

$$i_u(s) = p(s/c_s)\,p(c_s/u)$$

To estimate a suitable f, we make the following remarks:

- it is reasonable to expect that the value is proportional to the importance, for stories with the same duration, so that:

$$f(\lambda i, d) = \lambda f(i, d)$$

this implies that $f(i,d) = if(1,d)$, so that the value can be rewritten as $v_u(s) = i_u(s)\,f(d(s))$

- $f(d)$ should increase with d, because more details can be provided when duration is increased. However, for stories with the same interest, we should prefer two short stories to a long one, which means that f should satisfy:

$$f(d_1) + f(d_2) \geq f(d_1 + d_2)$$

For example, we can assume that the following holds:

$$f(2d) = \alpha \cdot 2 f(d), \text{ for some } \frac{1}{2} \leq \alpha \leq 1$$

this implies that:

$$f(2^n d) = \alpha^n 2^n f(d) = 2^{n(1+Log\alpha)} f(d),$$

and suggests the following form for f:

$$f(d) = d^{1+Log\alpha} = d^{\beta}, \text{ for some } 0 \leq \beta \leq 1$$

In our experiments, we assumed that

$$f(d) = \left(\frac{d}{d_0}\right)^{\beta}, \text{ with } d_0 = 1\,\text{min and } \beta = 0.9$$

4.3 Story selection

The selection of the best set of stories corresponds to the resolution of the "knapsack" problem. This problem is NP-complete, which means that the optimal solution can be found only by enumerating a large number of possible combinations. In order to simplify the selection process, we use a heuristic by which we sort stories by score, enumerate stories by decreasing score and select those which fit within the admissible duration. This heuristic selection makes the amount of computation linear with respect to the number of stories (n Log n if we consider the time needed to sort the scores). We considered two possible scores:

- the content value $v_u(s)$
- the content density $density_u(s) = \dfrac{v_u(s)}{duration(s)}$.

To evaluate the performance of this heuristic, we compared it with optimal selection on a set of randomly generated profiles. For each profile, we construct the optimal set by enumeration and the heuristic set by the previous method. We compare the value found by the heuristic set with the optimal one (of course the heuristic value is always lower than the optimal), and we observe the average ratio of these values for various duration. The corresponding graph is shown below.

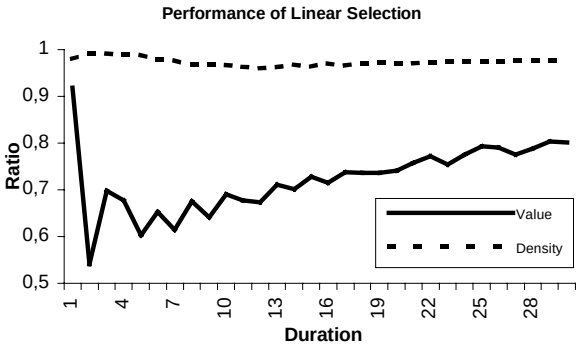

These experiments show that sorting by density provides a very efficient selection, which has an average value of more than 95% of the optimal value. The reason for this is probably that the duration of news stories has a regular repartition around an average value of two minutes, so that stories selected for the first minutes of the program are quite good anyway, and only for the last minutes is the selection not optimal (note that if all story durations were equal, the heuristic would guarantee to provide the optimal result).

4.4 Construction of personalized program

Once the selection of stories has been made, the construction of the personalized program simply consists in concatenating the recordings of selected stories into a single video file. In our prototype, we introduce an intermediate step by which the selected stories are displayed, and a preview mode in which the first 20 seconds of each story are presented (this is mostly used for demonstration purposes). The following figure is an example of this intermediate interface.

4.5 Pros and cons

A big disadvantage of this approach is that it contains a manual annotation scheme which is crucial for its implementation.
Manual annotation consists of two aspects:

- category assignment,
- importance evaluation.

The category assignment is almost a deterministic process, so it is not very dependent on specific evaluation of the annotator. The importance evaluation is much more sensitive to personal evaluation of the annotator.
It could be very interesting to automate this step. While we have not experimented in this way (but in another approach that will be described in the next section) here is a description of what could be achieved to automate the process:

- category assignment: this is essentially similar to text categorization, so many automatic methods can be applied.
- importance evaluation: since this factor represents an average value, it can be estimated by collecting evaluation from a limited set of users, for example users who access a news database from the Internet. The relative number of access to a specific story (when compared to the number of access to the category) could be used as an importance evaluation.

A big advantage of this approach is that the user profile is indicated explicitly. This makes it easy for the user to introduce any modification he might want. This also allows a single user to have several profiles, for example for business and leisure usage.

5. KEYWORD-BASED SCHEME

In our second prototype, we want to implement an implicit control of the user, by which the user does not directly have access to his profile, but simply provides information which is used to update this

profile. Furthermore, we want to investigate the construction of a fully automatic system, where no operator manipulation is required. To avoid the manual step of classification, we need some information of semantic nature about the news stories. We are using caption information that is captured with a special teletext decoder. Teletext is a technique which sends digital information during invisible period of the TV signal (the Vertical Blanking Interface). Some of this information contains subtitles (captions) which are mostly intended to be viewed by deaf or hard-of-hearing people (many modern TV sets contain a built-in Teletext decoder). We record TV News programs where caption information is available (only one French TV channel currently has this feature, because a group of journalists has to be specially devoted to the preparation of captions in parallel to the preparation of the news program itself). The teletext recording consists of text fragments, typically one or two lines of at most 32 characters, with the time at which this text was presented on the TV screen, and the time at which it was removed. It can therefore be easily synchronized with the audio-video recording. The captions are mostly a simplified transcription of what the anchor person is saying, since full text transcription would be too verbose to read. Also, not all of the program is closed-captioned, (when there is no time to prepare text for live interviews and stories which are recorded late), but the major part of the program generally is.

5.1 Content value

In the keyword scheme, the user profile is expressed as a measure $p_u(w)$ of the interest of the user for certain words. A story is represented as a vector of keyword counts $tf_s(w)$ which counts the number of occurrences of each keyword in the caption of the story. The interest of the user for a story is computed according to the standard Vector Space model, using the Tf.Idf formula, as:

$$i_u(s) = \frac{\sum_w v_s(w)\, p_u(w)}{\sqrt{\sum_w v_s(w)^2}\ \sqrt{\sum_w p_u(w)^2}}$$

where $v_s(w) = tf_s(w).idf(w)$.

The factor $idf(w)$ is the Inverse Document Frequency of the keyword and weights the importance of keywords depending of their frequency of occurrence in documents. It should be computed from a large number of documents. Since the number of stories per TV news program is limited, we compute this factor using a database of wirenews (see next section), using the formula:

$$idf_{wirenews}(w) = Log_2\left(\frac{N_{wirenews}}{N_{wirenews}(w)}\right)$$ where

$N_{wirenews}$ (resp. $N_{wirenews}(w)$) is the number of news articles (resp. containing w) in the wirenews database.

As in the category-based scheme, the content value of a story is computed by combining importance and duration: $v_u(s) = i_u(s)\left(\frac{d(s)}{d_0}\right)^{\beta}$

5.2 WireNews collection

The captions that are captured with the TV News program are limited in quantity. In order to compute a more reliable estimation of the Idf coefficient for keywords, we take advantage of the availability of similar textual data on the Internet. We use one WireNews site (actually the web site of the France 2 TV Channel) which contains specific pages with selected copies of wirenews information. These news are much more extensive that what is presented in the TV news program, yet most news from the TV program also appear in the wirenews, so that wirenews can be consider as a valid database for collecting keywords statistics. We automatically record these stories (pages are automatically updated every day), which provides us with text that is similar with the one in captions, but in much larger quantities.

The capture process has been going on for the last 6 months, for a total of 8000 news.

The following figure indicates the global architecture of this prototype, and shows the place of wirenews collection within the complete processing

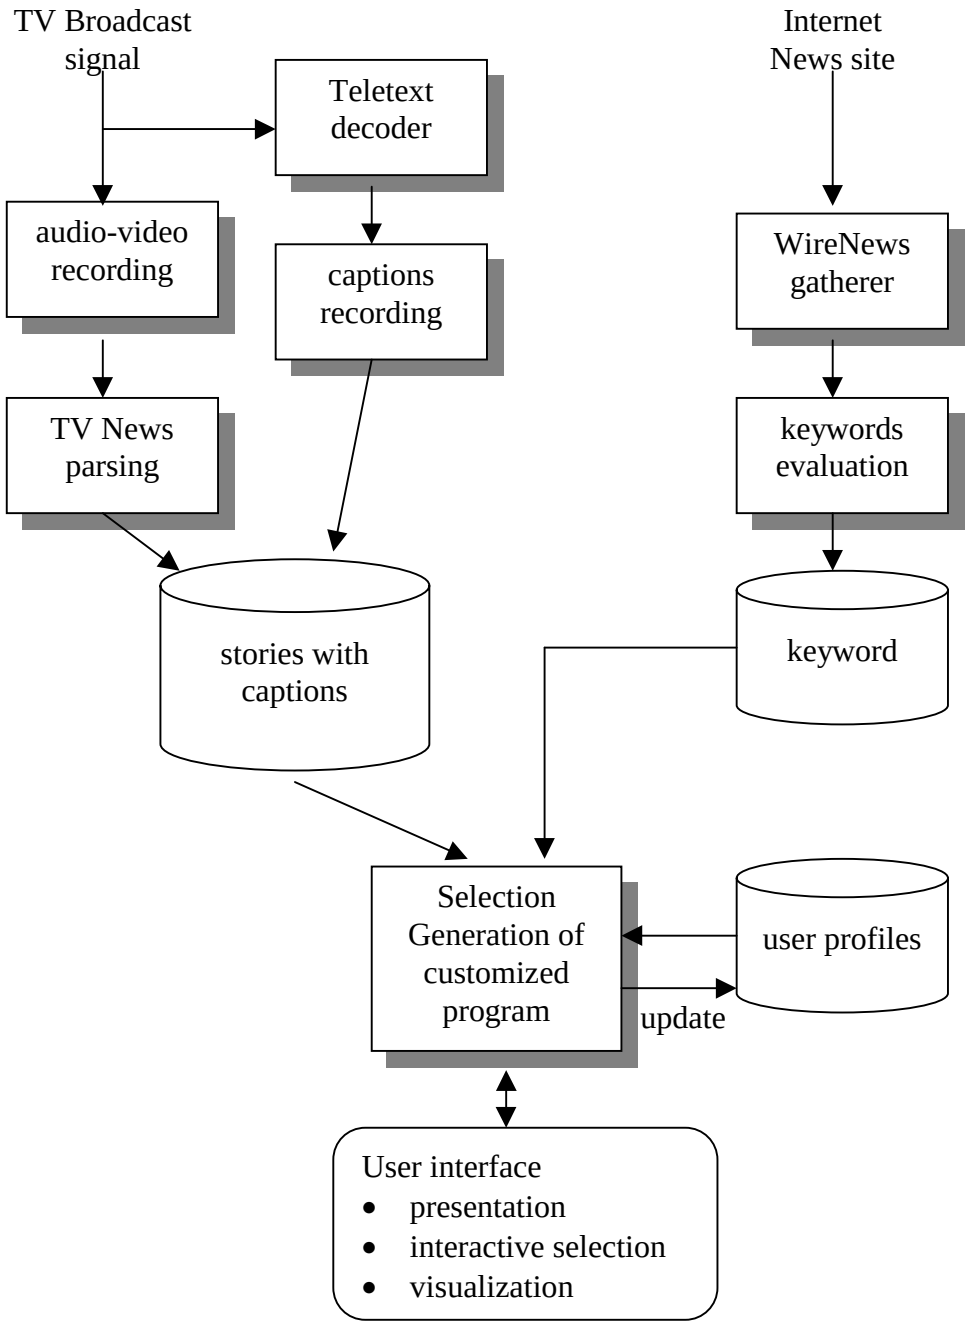

5.3 User interaction

The process of recording TV News and captions, segmenting, collecting wirenews, is performed everyday and continuously updates the story and wirenews databases. In the current version of our prototype, the user who asks for a personalized program is first asked the expected duration, plus the precise dates of the stories that he wants to be considered (for example, he might want to get a customized version of yesterday's program, or a selection from the previous week).

Implicit interaction of the user can be implemented in several ways. One solution could be to automatically generate the customized program, play it to the user, and get on-the-fly feedback on which stories the user appreciates. This requires a player

with adequate functionality, so we implemented a simpler paradigm where the user makes a selection from partial information on stories, these stories being arranged in decreasing preference order. This partial information is composed of a keyframe, and some textual information extracted from the captions. This textual information could be the most important fragment extracted of the captions, but we estimated that this would be too restrictive. Rather, we display the most important keywords in the caption. An example of the corresponding interface is shown in the figure below. The best stories within the expected duration are preselected. The user is allowed to select or deselect any story other story. The list of keywords is provided as a short indication of the content of the story.

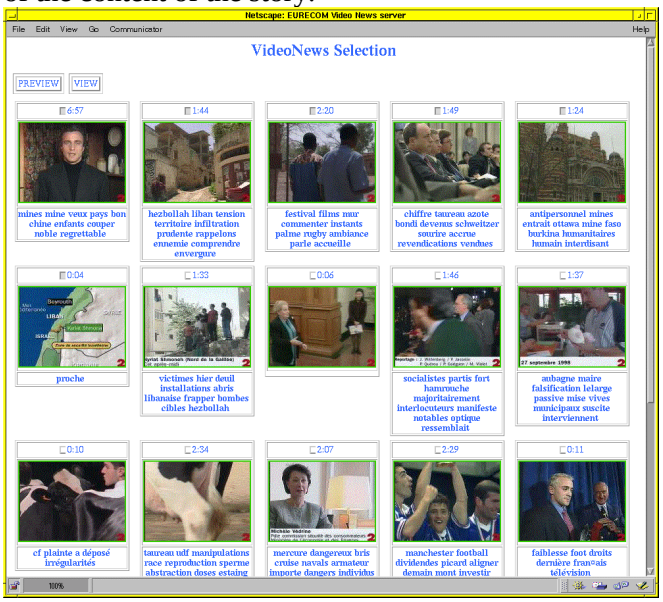

When the user is satisfied with the selection, he can press the VIEW button for a visualization of the customized program. This currently starts a process, which concatenates the Mpeg files corresponding to various stories into a single program file. This also updates frequency counts in the user profile with the counts for keywords in the selected stories. For demonstration purposes, a PREVIEW button will concatenates the first 20 seconds of each story only.

It is intended that the selection process will be only used in the initial phase, when the user has to build his profile, or at regular intervals when the user wants to provide new modifications to his profile (and has sufficient time to do it). In occasions where the user is satisfied with the preselection, or does not have time to refine his own program, the VIEW button can be activated right away and the default program is automatically constructed and presented.

5.4 User profile update

The idea is now that the user indicates its feedback by selecting the specific stories that he wants to be part of his news program. The intention is that this is a learning process: initially, the user profile is set to an average value, and will exhibit poor performance in predicting the user's interests. At that point, the user has to select stories by browsing and looking at keyframes and keywords. This is an extensive work for him, but as time goes on, the performance of the user profile improves, and there is gradually a larger correspondence between the user's choice and the first N stories proposed. The user can therefore choose to rely on the proposed selection, and ask for its program to be generated right away. Since this process is done at every step, he might at any later time, decide to reuse manual selection and update his profile.

When a story is selected by the user, statistics about the words included in the corresponding captions are added in the user profile. The user profile is simply a probability distribution over words $p_u(w)$. When a new document is added, the relative frequency distribution of words in this document is used to update the user profile:

$$p'_u(w) = \lambda p_u(w) + (1-\lambda) tf_s(w)$$

λ is a coefficient which is used to smooth the evolution of the user profile. Because of this coefficient, the importance of words from older stories decreases as new stories are added. This helps in tracking the evolution of the user's interests as they evolve over time (this method is known as the exponential moving average).

6. CONCLUSION

In this paper, we have shown how video indexing and information filtering techniques can be combined to construct automatically personalized TV News programs. We have studied how duration can be taken into account in the definition of the story content value and the selection of the best stories. We have presented two paradigms (explicit and implicit) for the construction and update of the user profile.

Our ultimate goal is the construction of a fully automated system. Our second prototype is currently operational and used daily to parse news recordings. In this phase, it is not yet exploited by real users, but is used mostly to test the robustness of the parsing procedures on new data. We expect that limited experimentation will provide useful indications on the applicability of such processing on a larger scale. Among the open questions are issues about the user interface, and the best mechanism to capture user feedback. Another important issue is the dependence of news stories: we have assumed independence, which is not true when several journals are considered. When an event is important, and spans over a certain period of time, it can be described in several successive stories.

The automatic construction of personalized programs is one of the great challenges of future television systems. We hope that the experiments presented in this paper will illustrate how such mechanisms can be put in practice.

7. REFERENCES

[1] Alfred Aho, Shih-Fu Chang, Kathleen McKeown, Dragomir Radev, John Smith, and Kazi Zaman. Columbia digital news system: An environment for briefing and search over multimedia information. In *Proceedings of IEEE International Conference on the Advances in Digital Libraries*, Washington, DC, 1997.

[2] Philippe Aigrain, HongJiang Zhang, and Dragutin Petkovic. Content-based representation and retrieval of visual media: A state-of-the-art review. *Multimedia Tools and Applications*, 3(3):179-202, November 1996.

[3] Alexander G. Hauptmann and Danny Lee. Topic labeling of broadcast news stories in the informedia digital video library. In *DL'98: Proceedings of the 3rd ACM International Conference on Digital Libraries*, pages 287-288, 1998.

[4] Karen Sparck Jones. What might be in a summary? In Knorz, Krause, and Womser-Hacker, editors, *Information Retrieval 93: Von der Modellierung zur Anwendung*, pages 9-26, Konstanz, DE, 1993. Universitätsverlag Konstanz.

[5] Rainer Lienhart, Silvia Pfeiffer, and Wolfgang Effelsberg. Video abstracting. *Communications of the ACM*, 40(12):54-62, December 1997.

[6] Bernard Merialdo. Automatic indexing of tv news. *Workshop on Image Analysis for Multimedia Integrated Services*, June 1997.

[7] A. Merlino, D. Morey, and Mark Maybury. Broadcast news navigation using story segmentation. In *Proceedings of The Fifth ACM International Multimedia Conference (MULTIMEDIA '97)*, pages 381-392, New York/Reading, November 1998. ACM Press/Addison-Wesley.

[8] Y. Nakamura and Takeo Kanade. Semantic analysis for video contents extraction spotting by association in news video. In *Proceedings of The Fifth ACM International Multimedia Conference (MULTIMEDIA '97)*, pages 393-402, New York/Reading, November 1998. ACM Press/Addison-Wesley.

[9] Douglas W. Oard. The state of the art in text filtering. *User Modeling and User-Adapted Interaction*, 7(3):141-178, 1997.

[10] Salton, Allan, Buckley, and Singhal. Automatic analysis, theme generation, and summarization of machine-readable texts. *SCIENCE: Science*, 264, 1994.

[11] Michael Smith and Takeo Kanade. Video skimming for quick browsing based on audio and image characterization. In IEEE, editor, *CVPR*, February 1997.

[12] HongJiang Zhang, S. Yeo Tan, Stephen Smoliar, and Gong Yihong. Automatic parsing and indexing of news video. *Multimedia Systems*, 2:256-266, 1995.