UvA-DARE (Digital Academic Repository)

Size matters! How thumbnail number, size, and motion influence mobile video retrieval

Hürst, W.; Snoek, C.G.M.; Spoel, W.-J.; Tomin, M.

DOI

[10.1007/978-3-642-17829-0_22](https://doi.org/10.1007/978-3-642-17829-0_22)

Publication date

2011

Document Version

Author accepted manuscript

Published in

Advances in Multimedia Modeling

[Link to publication](#)

Citation for published version (APA):

Hürst, W., Snoek, C. G. M., Spoel, W.-J., & Tomin, M. (2011). Size matters! How thumbnail number, size, and motion influence mobile video retrieval. In K-T. Lee, W-H. Tsai, H-YM. Liao, T. Chen, J-W. Hsieh, & C-C. Tseng (Eds.), *Advances in Multimedia Modeling: 17th International Multimedia Modeling Conference, MMM 2011, Taipei, Taiwan, January 5-7, 2011 : proceedings* (Vol. 2, pp. 230-240). (Lecture Notes in Computer Science; Vol. 6524). Springer. https://doi.org/10.1007/978-3-642-17829-0_22

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

UvA-DARE is a service provided by the library of the University of Amsterdam (<https://dare.uva.nl>)

Size Matters! How Thumbnail Number, Size, and Motion Influence Mobile Video Retrieval

Wolfgang Hürst¹, Cees G. M. Snoek², Willem-Jan Spoel¹, and Mate Tomin¹

¹ Utrecht University, PO Box 80.089, 3508 TB Utrecht, The Netherlands

² University of Amsterdam, Science Park 904, 1098 XH Amsterdam, The Netherlands
huerst@cs.uu.nl, cgmsnoek@uva.nl, {cwjispoe | mtomin}@students.cs.uu.nl

Abstract. Various interfaces for video browsing and retrieval have been proposed that provide improved usability, better retrieval performance, and richer user experience compared to simple result lists that are just sorted by relevance. These browsing interfaces take advantage of the rather large screen estate on desktop and laptop PCs to visualize advanced configurations of thumbnails summarizing the video content. Naturally, the usefulness of such screen-intensive visual browsers can be called into question when applied on small mobile handheld devices, such as smart phones. In this paper, we address the usefulness of thumbnail images for mobile video retrieval interfaces. In particular, we investigate how thumbnail number, size, and motion influence the performance of humans in common recognition tasks. Contrary to widespread belief that screens of handheld devices are unsuited for visualizing multiple (small) thumbnails simultaneously, our study shows that users are quite able to handle and assess multiple small thumbnails at the same time, especially when they show moving images. Our results give suggestions for appropriate video retrieval interface designs on handheld devices.

Keywords: Mobile video, video retrieval interfaces, visual assessment tasks.

1 Introduction

Multimedia services like Internet browsing [3], music management [8], and photo organization [2] have become commonplace and frequently used applications on handheld devices – despite their limited screen sizes. Even for capturing, accessing, and displaying video, many effective mobile interfaces exist [12, 5]. These interfaces take the complete video as the unit for user interaction and offer means to navigate through the timeline with the aid of a finger [14], pen [11,15], or scroll wheel [13]. It has been predicted by many that users will soon demand facilities providing them direct access to the video content of interest without the need for intensive timeline navigation [20, 26]. With the help of social tagging and multimedia content analysis techniques, like speech recognition [27] and visual concept detection [24], textual labels can be added to video segments allowing for interactive retrieval. Although the video retrieval community has proposed several interactive browsers able to support the user in this task on the relatively large screen estate of desktop and laptop

machines [27, 24, 1] or as part of a collaborative (mobile) network [23], only few interfaces exist for single-user interaction on handheld devices [18, 9]. In this paper, we are interested in the question how advanced video retrieval browsers should be adapted to constraints imposed by the screen of the handheld device. In order to place this question in perspective, it is important to realize that the basic building block of advanced video retrieval interfaces are *thumbnails* extracted from a short piece of video assumed to be representative for the multimedia content. We distinguish between *static thumbnails* (i.e. a reduced-size version of a single static image) and *dynamic thumbnails* (i.e. a set of consecutive moving reduced-size images).

State-of-the-art video retrieval browsers display several of these static or dynamic thumbnails simultaneously in response to a user query, for example as a matrix-like storyboard or ordered in a grid [27]. One may be tempted to consider these browsers unsuited for mobiles, as the limited screen size of handheld devices would not allow displaying several thumbnails at the same time without loss of recognition by the user. However, a recent user study by Torralba *et al.* [25] revealed human participants were able to outperform computer vision algorithms in an image recognition task on the desktop, even when they were only able to see strongly reduced versions of the original images. In fact, the authors showed that at a size of 32x32 pixels, humans are still able to recognize 80% of the visual content accurately. In [16], we evaluated similar recognition tasks in a video retrieval context on a mobile device, confirming a surprisingly high recognition performance at relatively small sizes, especially when using dynamic thumbnails. These results suggest that despite the limited screen estate of handheld devices, we should still be able to design much more complex thumbnail-based interfaces than commonly assumed. However, our studies have been limited to single thumbnails shown in isolation. Thus, our findings can not necessarily be generalized to video retrieval application where multiple thumbnails are shown simultaneously. In this paper, we are presenting a user study that also investigates the number of thumbnails in addition to thumbnail size and type (i.e. static versus moving). In particular, we verify whether users are able to assess video retrieval results on a mobile phone when *multiple* static and dynamic thumbnails are displayed simultaneously (at different sizes) and whether this will influence their perception and verification performance.

In the remainder of this paper we survey related work on (mobile) video retrieval interfaces (section 2), present the experimental methodology used and detail our findings (section 3), and conclude with resulting design suggestions (section 4).

2 Related Work

2.1 Interfaces for Traditional Video Retrieval

Early video retrieval interfaces for desktop PCs and laptops simply presented a video to the user as a sequential sequence of static thumbnails, using metaphors like filmstrips [6]. Alternatively, researchers have merged several static thumbnails into Manga-like collages [4, 7] and storyboards [1]. Static thumbnails are well suited to summarize relatively short video shots. When shots are lengthy and contain a lot of

object or camera motion, a single still image might not be able to communicate the factual visual content appropriately. To summarize lengthy shots containing moving content, dynamic thumbnails such as skims [6] have been proposed. These summaries aim to capture the full information content as compact and efficient as possible. Despite the apparent advantage of dynamic thumbnails over static thumbnails [6, 19, 21], we are unaware of user studies other than [16] quantifying this advantage.

Recall that on desktop PCs and laptops, the thumbnail has always been the traditional building block for summarizing video content in video retrieval interfaces [27, 24, 1, 10, 19, 21]. The result of a video query is typically represented as a sorted sequence of relevant thumbnails ordered in a grid, or a simple list of results. Since content-based video retrieval is an unsolved problem, it has been shown by many that in addition to visualizing direct video retrieval results in the interface, displaying indirect video retrieval results such as the complete time line [24] of the retrieved video shot, visually similar video shots [10], semantically related shots [21], or pseudo-relevance feedback results [19] are all beneficial to increase video search performance [22]. These findings have resulted in effective but complex video retrieval interfaces, which maximize usage of screen estate to display as many thumbnails as possible. A good example is the CrossBrowser [24], shown in Figure 1, which displays shot-based video retrieval results on a vertical axis and utilizes the horizontal axis to display the timeline of the complete video for a selected shot. Naturally, it can be called into question if such advanced video retrieval interfaces exploiting the screen estate of desktop PCs to the max can be transferred to the relatively small display of handheld devices.

Fig. 1. Displaying video retrieval results using an advanced visualization for the desktop (left) and a handheld device (right). The CrossBrowser [24] on the left displays (static or dynamic) thumbnail results on the vertical axis and exploits the horizontal axis to display a thumbnail timeline of the complete video for selected clip. YouTube for the iPhone orders static thumbnail results in a linear list together with user-provided meta data. Note the waste of screen estate by the CrossBrowser and the limited number of static thumbnails on the mobile.

2.2 Interfaces for Mobile Video Retrieval

Most present-day video retrieval interfaces on handheld devices are based on an ordered list of thumbnails; see Figure 1 for a representative example. It appears that the size of the thumbnail in current commercial mobile video retrieval interfaces is merely decided by the amount of meta-data that services want to display next to it. This observation also holds for one of the earliest mobile video retrieval interfaces by Lee et al. [18]. In the reference, the authors study the design of a PDA user interface to video browsing, emphasizing the role of user interaction. Although the authors claim that spatial presentation of multiple thumbnails on a mobile display is unsuitable, no quantitative experiments are provided to support the claim.

Most existing work on mobile video retrieval has focused on mechanisms for browsing through a single video. The MobileZoomSlider [11], for example, provides the user with several sliders, overlaid on the video on demand, which mimic different granularity levels of the video. The same approach was integrated in Apple's iPhone in iPhoneOS version 3.0. The idea of using sliders has been developed further in [5], where the authors suggest combining a preview slider with markers, representing scene and speech information, on the timeline. In [15] the authors evaluate alternative approaches for browsing along the timeline such as flicking and elastic interfaces. [13] uses a virtual scroll wheel overlaid to the video. Despite their effectiveness for browsing a single video, none of these interfaces are optimized for browsing through thumbnails originating from several videos, as is typical in retrieval. Exceptions are [14] and [23]. In the iBingo system [23], several users are collaborating using different (mobile) devices to retrieve video fragments relevant to an information need. Special attention is placed on the collaborative retrieval backend which aims to eliminate redundancy and repetition among co-searchers to increase overall retrieval efficiency. Since iBingo emphasizes collaboration by using multiple devices, here also, no special attention is reserved for the role of the thumbnail in the video retrieval interface. In [14], the authors present an interface where thumbnails are used to access points of interest in a larger video. However, their work focuses on interface design for one handed operation and ignores related video retrieval or browsing questions.

2.3 Thumbnails for Mobile Video Retrieval Interfaces

Although static thumbnails are often used for mobile video browsing in both commercial systems (cf. Fig. 1) as well as research prototypes (cf. previous subsection), the related interfaces are generally much less complex than their desktop counterparts, and dynamic thumbnails are generally not applied in a mobile context at all. In addition, the size of used thumbnails is often rather large. Motivated by Torralba *et al.*'s findings [25] in a desktop context, we presented several experiments in [16] investigating the relevance of thumbnail size and type for human recognition performance in a thumbnail based video retrieval scenario. In different test runs, both static and dynamic thumbnails were extracted from videos and presented to subjects with increasing as well as random sizes. Participants of the experiments had to solve typical video retrieval tasks solely based on a single thumbnail. In terms of performance, we identified 90 pixels to be a reasonable threshold upon which most

tasks were solved successfully based on a static thumbnail. However, with dynamic thumbnails, a similar performance was achieved starting at sizes of 70 pixels. Most surprisingly, human performance was relatively high even at much lower sizes, with a successful recognition in almost 90% of the cases at the smallest thumbnail size of 30 pixels. These results indicate that we can indeed build much more complex thumbnail-based interfaces for mobile video retrieval despite the small screen estate of handheld devices. However, all these tests have been done based on a single thumbnail that was presented to the subjects on a black background. In this paper, we verify if and to what degree these findings generalize to more complex interfaces where multiple thumbnails are presented simultaneously. Motivated by the findings that have been observed with single thumbnails displayed in isolation, we present a user study with 24 participants where we investigate the influence of thumbnail size and motion when multiple thumbnails are displayed simultaneously.

3 User Study with Varying Numbers and Sizes of Thumbnails

3.1 Motivation and Setup

Small thumbnails are generally used in video retrieval interfaces because they allow us to present a large amount of information at a time – either from various documents (e.g. the shot-based retrieval results shown in the vertical axis of the CrossBrowser, cf. Fig. 1, left) or from one single video (e.g. the timeline-based representation of selected shots shown in the horizontal axis of the CrossBrowser). Hence, in addition to evaluated thumbnail size and type for a single shot (as done in [16]), it is also important to investigate these characteristics in context with several thumbnails, since we can expect that the concurrent representation of multiple thumbnails will influence users' perception and verification performance. The main purpose of our user study was therefore to evaluate sizes and the relevance of dynamic versus static thumbnails when multiple shots of one video are represented in a timeline – similar to the horizontal axis in the CrossBrowser.

Experiments in this study have been done with a *Motorola Droid/Milestone* phone running the Android operating system version 2.0. The phone features a touch screen with a relatively large resolution of 854x480 pixels. Although we expect most phones to increase in screen resolution in the future, this is at the upper end even for smart phones and clearly above the current state-of-the-art for the majority of devices. Therefore, we decided to implement and run all experiments in compatibility mode with older phones and Android OS versions, resulting in a screen resolution of 569x320 pixels that was used in all tests presented in this paper.

Based on this resolution, we created three setups: first, a timeline with nine thumbnails at size 60 pixels each; second, one with seven thumbnails at size 75 pixels each; and third, one with five thumbnails at size 110 pixels each (sizes indicate width of the thumbnails, and heights are adapted according to the video's aspect ratio). The reason for the selection of these three thumbnail sizes is explained below. We show examples of the three interfaces in Figure 2.

Videos and thumbnails were taken from the TRECVID benchmark [4], and realistic questions were selected from [17]. Some questions needed to be adapted in order to fit to a “yes/no” answer scheme (which was chosen to focus on the independent variables thumbnail size and type) but were similar in spirit to the ones used in the literature. Questions were chosen randomly, but under consideration of covering different retrieval tasks – in particular: object and subject verification (e.g. “Does the clip contain any police car?”) versus scene and event verification (e.g. “Does the clip contain any moving black car?”). Overall, we took twelve video clips plus thumbnails from [4] and associated questions from [17] to create twelve different examples – four for each of the three setups shown in Figure 2.

Fig. 2. Interfaces used in the user study: 3 multiple thumbnail timelines representing different shorts from one video with 9, 7, and 5 thumbnails at width of 60, 75, and 110 pixels, respectively. Shorts are sorted along the timeline.

3.2 Experiment

Experiments have been done in a quiet place with no distractions and subjects sitting comfortably on a chair. They involved 24 subjects (22 male, 2 female, ages 2 from 15-20, 15 from 21-30, 5 from 31-40, and 2 from 41-50). Human recognition performance can be extremely high when the device is hold unnaturally close to ones face (even at very small thumbnail sizes, cf. [16]). Therefore, participants were asked to “hold the device in a natural and comfortable way”, for example by resting their arms on a table (cf. Fig. 3). A neutral observer reminded them of this guideline when an awkward position was recognized during the evaluations.

Based on the sequence of thumbnails shown in the interfaces, subjects had to perform verification tasks that are typical in common video retrieval situations. For this, the participants had to answer the 12 questions (4 for each interface setup). The order of interfaces as well as the order of questions for each individual setup was randomized across the users in order to avoid any related influence on the results.

After some informal testing with an initial implementation, it was obvious to us that a version where all dynamic thumbnails are playing at the same time would

create too much distraction and a cognitive overload. Hence, we decided that initially only the center thumbnail (which corresponds to the major thumbnail shown in the center of the CrossBrowser's horizontal axis; cf. Fig. 1, left) is playing (in an endless loop) whereas all other thumbnails are static ones. Users could however start playing any related dynamic thumbnail by just clicking on its static representation. This also stopped the currently playing thumbnail and replaced it with its static version. Hence, only one dynamic thumbnail was shown at a time.

Fig. 3. Participants during the user study assessing the relevance of video retrieval results on the mobile phone with varying sizes and numbers of static and dynamic thumbnails.

3.3 Results

The sizes of the thumbnails used in all three setups were motivated by the results of our previous study with single thumbnails shown in isolation [16]. The smallest thumbnail size of 60 pixels was chosen because it proved to be large enough to achieve a good performance although both static and dynamic thumbnails also had a large amount of errors at this size. 70 pixels have been identified to be a good threshold for dynamic thumbnails, i.e. they are smaller than the size suggested for static thumbnails, but large enough for dynamic ones. Finally, 110 pixels is a size where both static and dynamic thumbnail types should lead to a good human verification performance. The number of thumbnails in each timeline has been assigned to take full advantage of the whole width of the mobile's screen at a resolution of 569x320.

Figure 4 illustrates the number and correctness of answers with respect to the number of dynamic thumbnails played for each of the three timelines used in the experiments. The diagrams illustrate that displaying 7 or 9 clips at smaller sizes results in a relative high number of wrong answers – especially compared to displaying 5 clips at a size of 110 pixels. Summed over all questions, we observed 30 and 44 mistakes for thumbnail sizes of 60 and 75, respectively, but only 7 for a size of 110 pixels. Interestingly, this number of mistakes seems inversely related to how often dynamic thumbnails have been played. If we divide the total number of played dynamic thumbnails by the number of clips, we observe an average number of plays per thumbnail of 36 and 37 for the two smallest thumbnail sizes versus 57 plays per thumbnail for a size of 110 pixels. However, in terms of absolute numbers, users click

far more often on static thumbnails at the smallest size (232 clicks at size 60) than for larger ones (164 and 190 for sizes 75 and 110, respectively).

Fig. 4. Results for each interface setup used in the experiments. Human performance in relation to how many dynamic thumbnails have been played for timelines at different sizes and thumbnail numbers summed over all answers and participants. Note an extremely low number of mistakes for the largest thumbnails in the shortest timeline.

These observations confirm the relevance of dynamic thumbnails compared to static ones that we already identified in our experiments with single thumbnails [16]. In addition, there was a lower average play rate per thumbnail for smaller sizes. This indicates that users were able to gain some dynamic information from the static representation of the timeline despite their small sizes. For example, if three thumbnails in a row show an airplane in the sky, people can conclude that the related video contains a flying airplane, even without actually playing any dynamic thumbnail. However, relying on static thumbnails came at the price of lower recognition performance: There were a relatively high number of mistakes despite the context represented by the larger number of static thumbnails. This can be explained with the verification problems with static thumbnails at smaller sizes that have been identified in our first study.

4 Conclusion and Design Suggestions

In this paper, we investigated the human recognition performance for typical video retrieval tasks based on a time-ordered sequence of multiple thumbnails. In particular, our experiments evaluated the relevance of thumbnail size, number, and type for human retrieval performance on a handheld device. Considering thumbnail size, we compared our results to previously identified thresholds in experiments where single thumbnails are shown in isolation. The observations of our user study confirmed the advantage of dynamic over static thumbnails, especially at very small sizes. However, the most important result of the study presented in this paper is that the previously identified thresholds for optimal thumbnail sizes do not transfer when they are not shown in isolation but represented in combination with several other thumbnails (even from the same video). Potential reasons for this are manifold (cognitive overload, distraction and additional clutter, etc.) and worth further investigation.

Despite this decrease in performance compared to [16], our new experiments showed that users are able to achieve good verification results – even at small to medium thumbnail sizes. The participants also performed well in recognizing and selecting the most promising ones for playback as dynamic thumbnails. Contrary to the widespread believe that screens of handheld devices are unsuited for simultaneously visualizing several (small) thumbnails, our results therefore suggest that users are quite able to handle and assess multiple thumbnails. This is especially true when using moving images. These results suggest promising avenues for future research related to the design of advanced video retrieval interfaces on mobile devices.

Let's first reconsider the mobile YouTube interface shown on the right side of Figure 1. Here, independent thumbnails are presented in a sorted list of results. In such an interface, we suggest that it is good to play them, because our results from [16] and the ones presented in this paper showed an increased human performance when using dynamic thumbnails. For more complex interfaces, as they are common for the desktop, the user study we presented here suggests that it is unwise to display the thumbnails on a mobile as small as suggested by our previous studies. However, a timeline with up to five thumbnails is no problem at all, even for the small display of

a handheld device (also remember that we purposely used a rather low resolution for our tests). Reconsidering the CrossBrowser for example, our results suggest that we could successfully use a mobile version where the vertical axis as well as horizontal axis are both visualized at the same time (horizontally though and with less thumbnails). Figure 5 illustrates design drafts for an implementation of a mobile video retrieval interface featuring a similar functionality as the CrossBrowser interface. The results of our studies suggest a high human performance in common retrieval tasks with such a setup despite the small screen size. Our current and future work includes implementations of such designs on a smart phone and evaluations comparing human assessment performance with the traditional CrossBrowser interface on a PC.

In conclusion, as the title “Keep moving!” of our previous work [16] suggests, moving images have a high relevance in mobile recognition tasks. Thumbnail sizes on the other hand are almost negligible – if and only if thumbnails are shown in isolation. The study presented in this paper confirmed the relevance of moving images over stills. However, it also showed that size is important if multiple thumbnails are used simultaneously to create more complex interface designs. Hence, when designing mobile video retrieval interfaces, we not only recommend to “Keep moving!” but also to keep in mind that “Size matters!”

Fig. 5. Design drafts for a mobile video retrieval interface considering the results of our user studies (center and right). By rearranging the vertical and horizontal axis in an appropriate way (and considering the size recommendations from our user study), we are able to visualize about the same information as with the CrossBrowser interface (left) and offer almost the same functionality to the user – something that is generally considered impossible due to the mobile’s small screen size. Note that in both cases there is even room for the representation of additional meta information (filename, etc.) at the bottom of the screen.

References

1. J. Adcock, M. Cooper, A. Girgensohn, and L. Wilcox. Interactive Video Search Using Multilevel Indexing. In Proc. CIVR (2005)
2. M. Ames, D. Eckles, M. Naaman, M. Spasojevic, and N. Van House. Requirements for Mobile Photoware. *Personal and Ubiquitous Computing*, 14(2):95-109 (2010)
3. S. Björk, L.E., Holmquist, J. Redström, I. Bretan, R. Danielsson, J. Karlgren, and K. Franzén. WEST: a Web Browser for Small Terminals. In Proc. ACM UIST (1999)
4. J. Boreczky, A. Girgensohn, G. Golovchinsky, and S. Uchihashi. An Interactive Comic Book Presentation for Exploring Video. In Proc. ACM CHI, pp. 185-192. The Hague, Netherlands (2000)

5. C. Brachmann and R. Malaka. Keyframe-less Integration of Semantic Information in a Video Player Interface. In Proc. EITC. Leuven, Belgium (2009)
6. M.G. Christel, A. Warmack, A.G. Hauptmann, and S. Crosby. Adjustable Filmstrips and Skims as Abstractions for a Digital Video Library. In Proc. IEEE Advances in Digital Libraries Conference 1999, Baltimore, MD. pp. 98-104 (1999)
7. M.G. Christel, A.G. Hauptmann, H.D. Wactlar, and T.D. Ng. Collages as Dynamic Summaries for News Video. In Proc. ACM Multimedia (2002)
8. R. Dachsel and M. Frisch. Mambo: A Facet-Based Zoomable Music Browser. In Proc. ACM Mobile and Ubiquitous Multimedia (2007)
9. C. Gurrin, L. Brenna, D. Zagorodnov, H. Lee, A.F. Smeaton, and D. Johansen. Supporting Mobile Access to Digital Video Archives Without User Queries. In Proc. MobileHCI (2006)
10. D. Heesch and S. Rüger. Image Browsing: A semantic analysis of NNk networks. In Proc. CIVR (2005)
11. W. Hürst, G. Götz, and M. Welte. Interactive Video Browsing on Mobile Devices. In Proc. ACM Multimedia (2007)
12. W. Hürst. Video Browsing on Handheld Devices—Interface Designs for the Next Generation of Mobile Video Players. *IEEE MultiMedia*, 15(3):76-83 (2008)
13. W. Hürst, G. Götz. Interface Designs for Pen-Based Mobile Video Browsing. In Proc. Designing Interactive Systems, DIS (2008)
14. W. Hürst, P. Merkle. One-Handed Mobile Video Browsing. In Proc. uxTV (2008)
15. W. Hürst, K. Meier. Interfaces for Timeline-based Mobile Video Browsing. *Proc. ACM Multimedia* (2008)
16. W. Hürst, C.G.M. Snoek, W.-J. Spoel, and M. Tomin. Keep Moving! Revisiting Thumbnails for Mobile Video Retrieval. In Proc. ACM Multimedia (2010)
17. B. Huurnink, C.G.M. Snoek, M. de Rijke, and A.W.M. Smeulders. Today's and Tomorrow's Retrieval Practice in the Audiovisual Archive. In Proc. ACM CIVR (2010)
18. H. Lee, A.F. Smeaton, N. Murphy, N. O'Connor, and S. Marlow. Fischlar on a PDA: Handheld User Interface Design to a Video Indexing, Browsing, and Playback System. In Proc. UAHCI (2001)
19. H. Luan, S. Neo, H. Goh, Y. Zhang, S. Lin, and T. Chua. Segregated Feedback with Performance-based Adaptive Sampling for Interactive News Video Retrieval. In Proc. ACM Multimedia (2007)
20. K. O'Hara, A. Slayden, A.S. Michell, and A. Vorbau. Consuming Video on Mobile Devices. In Proc. CHI (2007)
21. O. de Rooij, C.G.M. Snoek, and M. Worring. Balancing Thread Based Navigation for Targeted Video Search. In Proc. CIVR (2008)
22. A.F. Smeaton, P. Over, and W. Kraaij. Evaluation Campaigns and TRECVID. In Proc. ACM Multimedia Information Retrieval, pp. 321-330, Santa Barbara, USA (2006)
23. A.F. Smeaton, C. Foley, D. Byrne, and G.J.F. Jones. iBingo Mobile Collaborative Search. In Proc. CIVR (2008)
24. C.G.M. Snoek, M. Worring, D.C. Koelma, and A.W.M. Smeulders. A Learned Lexicon-Driven Paradigm for Interactive Video Retrieval. *IEEE Transactions on Multimedia*, 9(2):280-292 (2007)
25. A. Torralba, R. Fergus, and W.T. Freeman. 80 Million Tiny Images: A Large Data Set for Nonparametric Object and Scene Recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 30(11):1958-1970 (2008)
26. S.J. Vaughan-Nichols. The Mobile Web Comes of Age. *IEEE Computer*, 41(11):15-17 (2008)
27. H.D. Wactlar, M.G. Christel, Y. Gong, and A.G. Hauptmann. Lessons Learned from Building a Terabyte Digital Video Library. *IEEE Computer*, 32(2):66-73 (1999)