Rock Fresco, Tassili N-Ajjer Plateau,
Cave of Tan, Zoumaitor, Algeria

ALGERIA/HOLTON COLLECTION/SUPERSTOCK

Shih-Fu Chang, John R. Smith, Mandis Beigi, and Ana Benitez

Visual Information Retrieval from Large Distributed Online Repositories

Try WebSEEk, an Internet search system that integrates textual and visual features for categorizing images and videos, and MetaSEEk, a meta-search system that links to various visual search systems.

Digital images and video are becoming an integral part of human communication. The ease of creating and capturing digital imagery has enabled its proliferation, making our interaction with online information sources largely visual. We increasingly use visual content to express ideas, report news, and educate and entertain each other. But how can we search for visual information?

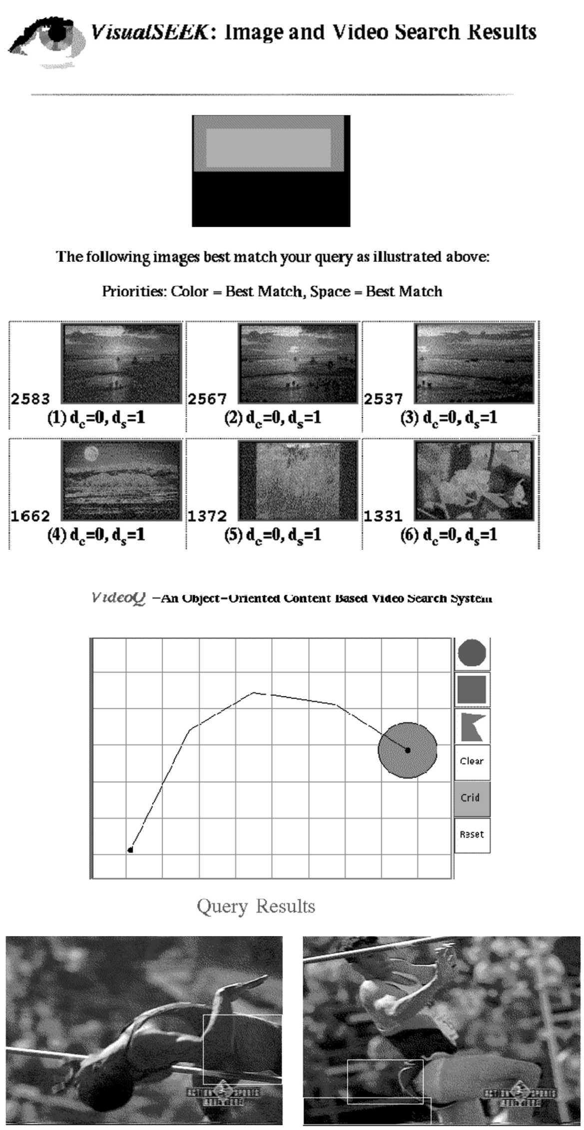

Figure 1. Feature-based visual query. Query 1 uses the local color regions and spatial structures to find images of a sunset or flower. Query 2 adds the motion feature to find relevant video clips, such as high jumps. (Videos courtesy Action Sports Adventure, Inc. and Hot Shots/Cool Cuts, Inc.)

Can solutions be developed that are as effective as existing text and nonvisual information search engines? With the increasing numbers of distributed repositories and users, how can we design scalable visual information retrieval systems?

Digital imagery is a rich and subjective source of information. For example, different people extract different meanings from the same picture. Their response also varies over time and in different viewing contexts. A picture also has meaning at multiple levels—description, analysis, and interpretation—as described in [10]. Visual information is also represented in multiple forms—still images, video sequences, computer graphics, animations, stereoscopic images—and expected in such future applications as multiview and 3D video. Furthermore, visual information systems demand large resources for transmission, storage, and processing. These factors make the indexing, retrieval, and management of visual information a great challenge.

Based on our experience developing visual information systems in Web-based environments, we have analyzed the present and future of such systems, focusing on search and retrieval from large, distributed, online visual information repositories. As a case study, we describe a recently developed Internet-based system called WebSEEk. We also describe the prototype of an Internet meta-visual information retrieval system called MetaSEEk, which is analogous to text-based meta-search engines on the Web. And we explore the challenges in developing scalable visual information retrieval systems for future online environments.

Content-Based Visual Query

Recent progress has been made in developing efficient and effective visual information retrieval systems. Some systems, such as Virage, QBIC, VisualSEEk, and VideoQ [1, 3, 6, 11], provide methods for retrieving digital images and videos by using examples and/or visual sketches. To query visual repositories, the visual features of the imagery, such as colors, textures, shapes, motions, and spatiotemporal compositions, are used in combination with text and other related information. Low-level visual features may be extracted with or without human intervention.

One characteristic of these systems is that the search is approximate, requiring a computed assessment of visual similarity. The items returned at the top of the list of query results have the greatest similarity with the query input. But the returned items rarely have an "exact" match to the attributes specified in the query. Figure 1 shows image and video search examples based on visual similarity queries.

Such systems can also use direct input from humans and other supporting data to better index visual information. For example, video icons in [4] are generated by a manual process for annotating objects in videos, like people and boats, and semantic events, like sunsets. Text indexes have also been generated from the captions and transcripts of broadcast video [8] for retrieving news video.

Visual summarization complements visual search. By decomposing the video, through, say, automated scene detection, a more spatially or temporally compact presentation of the video can be generated. For example, [12] has described news video summarization systems, with efficient browsing interfaces using video event detection and clustering. Others have developed techniques for automated video analysis of continuous video sequences to generate mosaic images for improved browsing and indexing.

Other researchers have begun seeking to automate the assignment of semantic labels to visual content. For example, through a process of learning from user interaction, the FourEyes system develops maps from visual features to semantic classes [9]. Furthermore, by manually developing specific models for visual classes using visual features, such as animals and nude people, techniques for automatically detecting these images are also being developed [7].

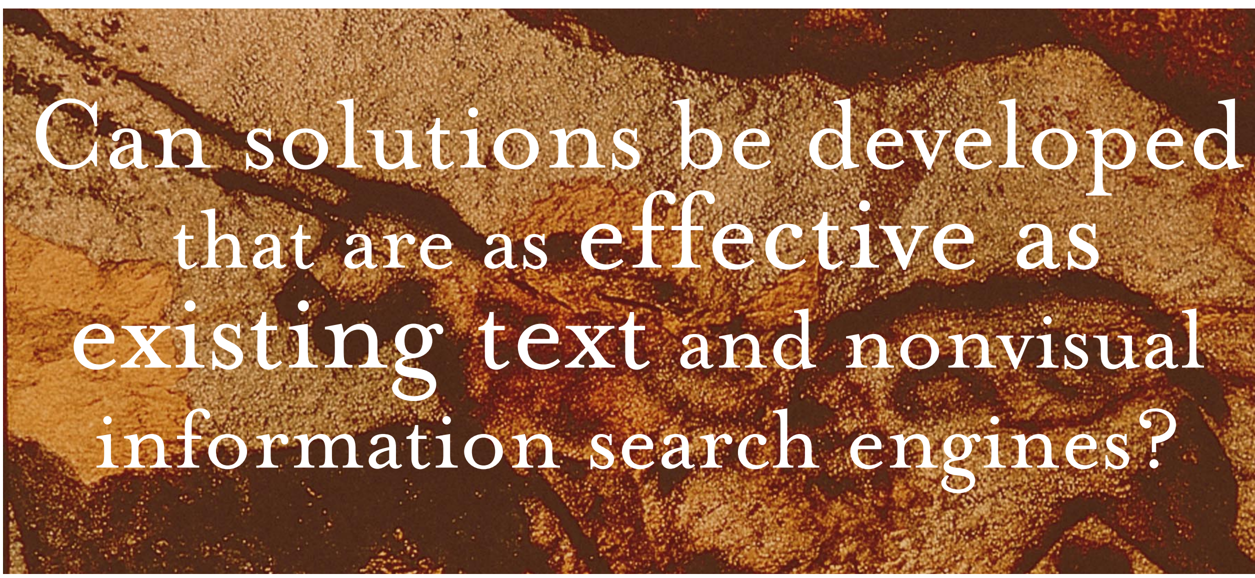

Classifying Retrieval Systems

Visual information retrieval systems have been used in many application domains, including libraries, museums, scientific data archives, photo stock houses, and Web search engines. We classify these systems using the following criteria:

- *Automation.* The visual features of the images and videos extracted and indexed by the system are used for interactive content-based visual querying. Visual information retrieval systems differ in degree of automation of feature extraction and index generation. Automatic methods are typically appropriate only for low-level feature extraction, involving, say, colors, textures, and motions. Generation of higher-level semantic indexes usually requires human input and/or system training. Some semiautomatic systems take initial input from humans, such as manual selection of image objects and features, and are then used to generate the feature indexes.
- *Multimedia features.* Multimedia content contains information in many modalities, including images, video, text, audio, and graphics. Visual information retrieval systems differ in their treatment of the multiple modalities. Typically, if multiple modalities are considered, they are indexed independently. Integration of multiple modalities has been investigated in a few systems [8, 11] but is not yet fully exploited.
- *Adaptability.* Most systems use a static set of previously extracted features. Selection of these features by the system designer involves trade-offs in indexing costs and search functionalities. However, due to the subjective nature of visual search, there is a need to be able to dynamically extract and index features to adapt to the changing needs of users and applications.
- *Abstraction.* Retrieval systems differ in the level of abstraction in which content is indexed. For example, images may be indexed at the feature level (color, texture, and shape), object level (moving foreground object), syntax level (video shot), and semantic level (image subject). Most automatic retrieval systems aim at low-level features, while high-level indexes are generated manually. Interaction among different levels is an exciting but unsolved issue.
- *Generality.* Retrieval systems differ as to specificity of the domain of visual information. For example, customized feature sets can be developed to incorporate specific domain knowledge, such as

for medical and remote-sensing applications. More general systems aim at indexing unconstrained visual information, such as that on the Internet.
- *Content collection.* Retrieval systems differ in the methods by which new visual information can be added. For example, in a dynamic visual information retrieval system, the content of the system's database may be collected by software robots, such as those that automatically traverse the Web. In other systems, such as those for online news archives and photo stock houses, visual information can be added manually.
- *Categorization.* Retrieval systems differ in their effort to categorize visual information into semantic ontologies. As visual information repositories have grown, we have found that interfaces allowing navigation through the semantic ontologies of visual information are very useful. However, effective image-categorization schemes have not been explored thus far.
- *Compressed domain processing.* Retrieval systems also differ in their approach to processing visual information. For example, it is possible to perform feature extraction directly on compressed images and videos. Compressed-domain processing avoids expensive expansion of the data.

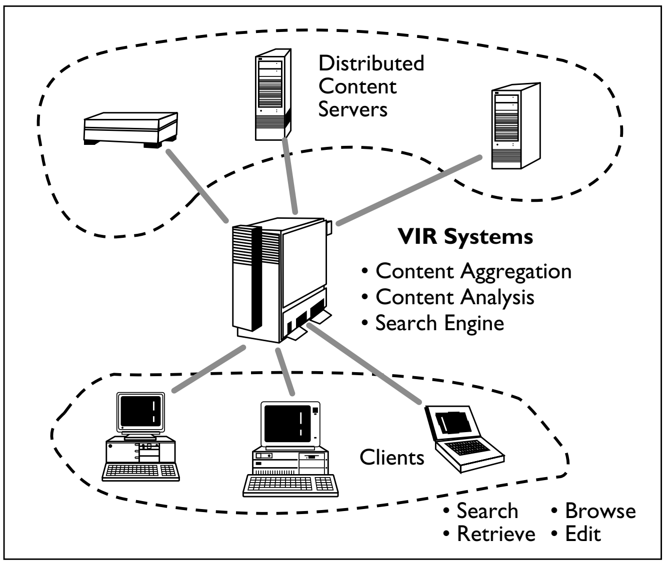

Figure 2. A high-level system architecture for Internet visual information retrieval

The criteria for federating content-based retrieval in multiple large distributed repositories may be different from those already described. We propose a framework for developing a visual information retrieval meta-search engine, motivated by the appearance of meta-search engines that federate Web-document searching [10].

WebSEEk for Internet Visual Information Retrieval

The Web's rich collection of visual information is integrated with a vast variety of nonvisual information. Although there are many popular search engines for nonvisual information, visual information search engines are only just appearing.

Visual information on the Web is highly distributed, minimally indexed, and schemaless. To explore visual information on the Web, we developed WebSEEk, a semiautomatic image search and cataloging engine [11] whose objective is to provide a visual search gateway for collecting, analyzing, indexing, and searching for the Web's visual information. Figure 2 shows a high-level system diagram for WebSEEk, which serves as an aggregation point for the Web's distributed visual information and acts as a server for Web querying, retrieving, and manipulating of indexed visual information. WebSEEk stores the meta-data, visual summaries, and pointers to the visual information sources.

According to our classification criteria, WebSEEk is semiautomatic, uses static visual features and textual key terms, indexes unconstrained visual content, uses a customized semantic ontology, and collects visual information using autonomous software agents.

WebSEEk's image and video collection process is achieved by Web spiders. Visual information is detected by mapping the file name extensions to the object types according to the Multipurpose Internet Mail Extensions (MIME) labels (such as .gif, .jpg, .qt, .mpg, and .avi). Based on the images and video collected during a three-month period in 1996, 85% of visual information collected consists of color images, 14% contains gray-scale or black & white images, and 1.5% contains video. The current system has indexed approximately 650,000 images and 10,000 video sequences.

WebSEEk takes the following approaches to handling the challenges of Internet visual information retrieval:

- Multimedia meta-data. Most online visual information does not exist in isolation; it is usually accompanied by other related information. For example, WebSEEk uses Web URL addresses and html tags associated with the images and videos

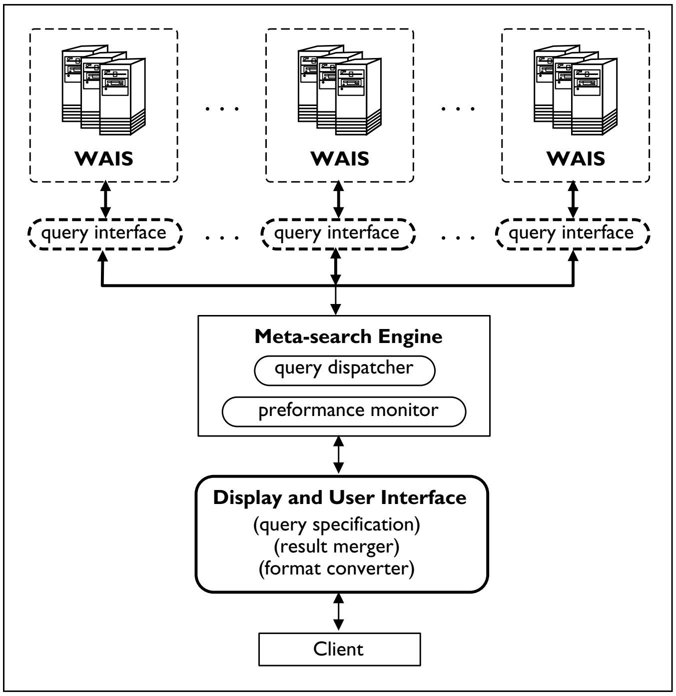

Figure 3. Basic components of a metasearch engine

to extract key terms for direct indexing and classification of visual content.

- Feature extraction and indexing. The large number of images and videos indexed by WebSEEk limits the amount of processing that can be performed in responding to a user's query. To achieve a content-based query response time of less than two seconds, WebSEEk uses a novel indexing scheme for fairly simple color features—binary color sets and color histograms.
- Image manipulation for efficient viewing. Visual query is an iterative, interactive process. Full-resolution images are not needed until the final stage, when the user issues an explicit request. Reduced representations of images or video can be efficiently extracted on the fly or in advance. For video, automatic video shot segmentation and key-frame selection are used.
- Semantic ontology and subject classification. The systems that require users to browse through pages of thumbnail views of images and videos are not suitable for large visual information retrieval systems. We are finding that users prefer to navigate within a clearly defined hierarchical semantic space. For example, by analyzing WebSEEk usage patterns, we found that semantic navigation is the most popular access method in this system. Once the user is within a smaller class of visual information, content-based methods are important for organizing, browsing, and viewing the content in that space.

Table 1. Textual terms and their mappings to the subject classes

Nondescriptive terms		Descriptive terms		
term	count	term	count	mapping
image	86380	planet	1175	astronomy/planets
gif	28580	music	922	entertainment/music
icon	14798	aircraft	458	transportation/aircraft
pic	14035	travel	344	travel
img	14011	porsche	139	transportation/ automobiles/porsches

Unlike ordinary image database applications, Internet visual information retrieval systems lack a well-defined schema or set of meta-data. Therefore, we use a set of the common attributes of the images and videos on the Web. WebSEEk extracts key terms from the URL addresses and html tags of images and videos on the Web and uses them to map the images and videos into one or more of the subject classes in WebSEEk's semantic ontology. Table 1 lists examples of descriptive and nondescriptive key terms and their mappings into classes of the semantic ontology.

WebSEEk's semantic ontology contains more than 2,000 classes and uses a multilevel hierarchy. It is constructed semiautomatically in that initially, human assistance is required in the design of the basic classes and their hierarchy. Then, periodically, additional candidate classes are suggested by the computer and verified with human assistance. Table 1 also includes example classes in the semantic ontology.

Because many textual terms are ambiguous, automatic subject classification using simple key terms is not perfect. However, its overall performance classifying visual information from the Web is quite good. We have found that WebSEEk's classification system is over 90% accurate in assigning images and videos to semantic classes. Performance is also verified by the popularity of WebSEEk's Web application; in its initial deployment, WebSEEk has processed more than

1,600,000 query and browse operations. The system currently serves thousands of queries daily.

As shown in Table 2, subject-based queries are the most popular search method for images and videos (accounting for 53.5% of queries). Content-based queries account for only 3.7% of all queries. However, this disparity may be partly influenced by the limitations of the content-based search functions in the current system.

Meta-Search Engines for Images

The proliferation of text search engines on the Web has motivated recent research in integrated search, or meta-search, engines [5]. Meta-search engines serve as common gateways linking users to multiple cooperative or competitive search engines. They accept query requests from users, sometimes along with user-specified query plans to select target search engines. Meta-search engines may also keep track of the past performance of each of the various search engines and use it in selecting target search engines for future queries.

A working meta-search engine includes three basic components (see Figure 3) [5]: The *dispatching* component selects target search engines for each query; the *query interface* component translates the user-specified query into scripts compatible with each target search engine; and the *display interface* component merges the query results from all the search engines. It may also perform format conversion or image manipulation to produce a list of displayable images for the client.

Our prototype meta-image search engine, MetaSEEk, has allowed us to research the issues involved in querying large, distributed online visual information sources [2]. MetaSEEk's target search engines include VisualSEEk/WebSEEk, QBIC, and Virage. The front-end interface allows browsing of random images from different engines. The user may select different visual features in the query and enter text (such as a URL or keywords) for the query. Figure 4 shows the user interface of this Web-based system, including color percentage, color layout, and texture, supported by the three target search engines. Users may search for images based on image examples or text [2]. Our experience with the initial MetaSEEk deployment has yielded several lessons:

Table 2. Usage patterns of an Internet visual information retrieval engine (WebSEEk)

	Browse			Query			
	total operations	subject navigations	visual browsing	subject-based queries	text-based queries	basic content-based queries (color distribution)	advanced content-based queries (VisualSEEk)
#	828,802	354,512	189,612	175,359	120,832	30,256	1,336
%	—	42.8%	22.9%	21.2%	14.6%	3.7%	0.2%

Figure 4. The MetaSEEk query interface, which allows users to brows images from and issue content-based queries to different target search engines. (Images in yellow are from QBIC; white from Virage; and red from VisualSEEk.)

- *Feature consistency.* Although the target search engines support querying based on color features, different feature sets and query methods are used. For example, color histograms are used by all target engines. However, the color spaces, distance metrics, and indexing methods are different in these systems. The mapping of each user-selected feature in MetaSEEk to those used in the target engines is based on the system's best effort.
- *Query results merging.* Image-match scores are not consistently returned by the target retrieval systems. Without match scores, the meta-search engine merges the result lists based on the ranks of the matches in each individual list. In a more advanced implementation, the meta-engine may

recompute simple visual features from the images returned from the various systems and rank them on the fly.

- *Special functionalities.* Each of the retrieval systems has special functionalities and limitations. For example, Virage allows the user to weight the multiple visual features in a query—color, texture, composition, and structure. On the other hand, QBIC and VisualSEEk support custom searches using visual sketches. Since functionality differs among target search engines, it is difficult to develop a common interface for all the systems in a meta-search engine.
- *Performance evaluation.* The lack of effective evaluation metrics or benchmarks for retrieval systems is a critical issue. Difficulty in evaluating retrieval effectiveness is due to the subjectivity of visual similarity and the lack of a standard visual information corpus in the research community. Without standard performance metrics, it is difficult for the meta-engine to monitor the performance of different target engines and make recommendations for subsequent queries. For example, the labels *No Results* and *Visit* are used in [5] to report query effectiveness. But in visual information retrieval systems, the queries usually return a fixed-size set of nearest neighbors without regard for the actual threshold of similarity. In practice, determination of an appropriate threshold of similarity is difficult, especially when the queries involve multiple features and arbitrary weightings.

Toward Scalable Internet Visual Retrieval Systems

Although visual information retrieval systems are still only emerging, it is important to understand the challenges in developing scalable solutions. A scalable Internet retrieval solution needs to solve three main problems: heterogeneity, complexity, and bandwidth.

Heterogeneity. Unlike text documents, visual material does not share consistent formats, indexes, or meta-data. This issue was brought to the fore in our initial experiments with meta-search engines. Dozens of formats are used for representing images and videos on the Web. Many different techniques are used for implementing content-based searching in retrieval systems. There is also no standard for interoperability between systems. For example, even at the semantic level, ontologies are custom-developed for the target systems.

These problems must be solved to improve interoperability. Standardization of the representation of meta-data should provide a uniform method for labeling visual information at the semantic level. Such standardization will enhance interoperability among different retrieval systems and improve the effectiveness of individual systems. As with the development of standards for meta-data for text documents, such recent efforts as that of CNI/OCLC Metadata have been made to extend the meta-data schemes to images. A standard taxonomy for graphic materials has been proposed by the Library of Congress. And the audiovisual research community has started to investigate development of a standard for describing multimedia information in MPEG-7.

The diversity of visual features and indexing schemes could be resolved by developing a distributed database query protocol, such as that used in information retrieval—the `EXPLAIN` operation proposed in Z39.50. Such a facility would allow clients to dynamically configure to the types of content and visual features indexed by the retrieval systems. For example, we found that the typical visual features indexed in visual information retrieval systems can refer to global or local content, be pixel- or region-based, describe intraobject attributes and/or interobject relationships, and include spatial and/or temporal features. These systems should have methods of describing the indexed information and query facilities to a generic client.

The various visual features and query metrics can also include count-based feature sets, such as histograms, vector-based feature sets and/or use Euclidean-norms, quadratic-norms, and/or transform-domain metrics. By developing a standard for communicating the types of features and the associated metrics used by retrieval systems, functional mappings may be generated that allow transparent querying among multiple retrieval systems. Having the means of translating and interpreting query methods and results, a visual information retrieval meta-search engine could integrate the results of queries of multiple retrieval systems more effectively.

Complexity. The process of searching through a repository of visual information is complex. For example, users may not know what they are looking for or how to describe the content they seek. The cataloging of the images or video into a fixed subject hierarchy provides a partial solution. But cataloging is limited, in that it provides only a limited number of directions in which the content can be found.

Online visual information retrieval systems can help solve this problem by supporting interactive browsing and dynamic querying. Users would be able to preview the initial search results, provide feedback as to the relevance of the returned items,

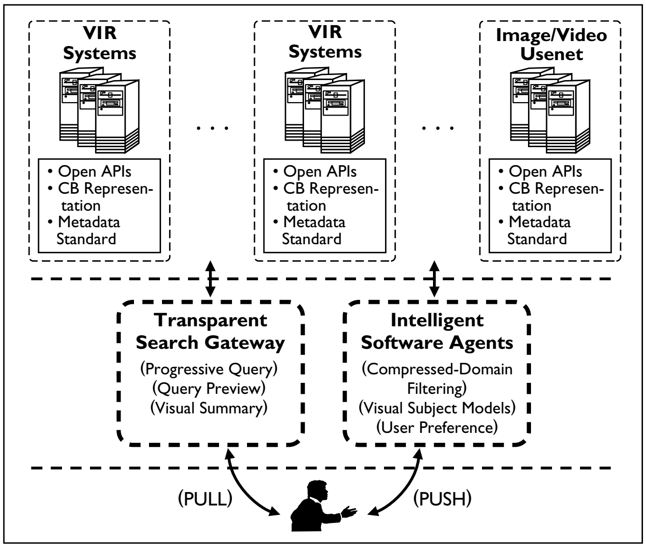

Figure 5. Scalable solutions for Internet visual information retrieval

and refine queries efficiently. Particularly in the Internet context, previews of query results reduce the amount of transmission required for full-resolution images and videos.

Furthermore, summarization of visual information for improved browsing and querying may be performed at several content levels. For example, the database statistics for each attribute or subject, say, the number of images with red regions or images in the subject class "nature," can be cues for users while also being used to optimize queries. Summaries can also be visual. Examples include mosaic images, video key-frame summaries, and video program transitional graphs. Querying may be improved by using only the dominant attributes of the visual information summaries.

The challenging application environment of Internet visual information retrieval makes effective visual information searching even more critical and difficult. For example, consider in an image metasearch system the difficulty of giving relevance feedback to individual retrieval sources. How can the query results from distributed sources be manipulated to conduct follow-up queries, so images from one retrieval system can be used to query other retrieval systems? And how can the visual metasearch engine initiate a content-based query using features not globally available in those systems?

Solving this problem requires that more "intelligence" be added directly to the image and video representations. For example, images should not be simply passive arrays of pixel values. Instead, they should be "active," including methods that generate, manipulate, and describe their content. The images and videos may also be rendered in different ways and at different resolution levels, without requiring modification of the raw data. The coding algorithms should be dynamically self-configuring. And new features should be available for dynamic extraction and use in visual matching at query time. Recent development along these lines involves MPEG-4 for object-oriented multimedia representation and FlashPix for efficient image transport, display, and rendering.

As the amount of visual information on the Internet increases, the methods for accessing and disseminating it will change. For example, we have discussed only a "pull"-based paradigm for visual information retrieval systems in which users are active in connecting to visual information servers and in specifying the search and/or browse operations. These users are also directly involved in manipulating search results and retrieving the images and/or videos of interest. However, the pull paradigm for visual information is limited.

In a more advanced system, the visual information servers and/or agents may learn to detect users' visual information needs. For example, the system may observe the types of visual information a particular user typically retrieves. In this case, a "push" system can automatically provide visual information using a subscription-based model. New visual information would simply arrive at the user's desktop without requiring the user's involvement in the process.

Bandwidth. Slow Internet response time is a critical constraint for today's visual information retrieval systems. Image and video compression reduce the amount of bandwidth and storage required by these systems. But network and storage constraints are still fundamental limitations in their performance. For example, downloading a 10-second MPEG-1 video clip takes 10 seconds over a T1 dedicated line, two minutes over a 128Kbps ISDN line, and up to nine minutes over a 28.8Kbps modem. Download time clearly influences user interaction with Internet visual information retrieval.

The progressive storage and retrieval of images and videos provides a method for viewing the query results from a rough-to-fine scale. Although not widely used, several scalable compression schemes have been included in the high-profile MPEG-2 video standard. Open image standards, like FlashPix, also include multitile, multiresolution features. In future video standards, such as MPEG-4, content-based scalability is being considered as a way of supporting independent access to objects of interests in video scenes.

For such compression standards as JPEG, MPEG-1, and MPEG-2, functionalities can be developed in the compressed domain. Assuming that visual information is stored and transmitted in compressed form, great savings in computation and bandwidth can be achieved by avoiding decompression of the images and videos. Using compressed-domain approaches, operations, such as key-frame selection, feature extraction, and visual summarization, are performed without fully expanding the data.

In summary, visual information retrieval from large, distributed, online repositories requires eight critical components: a meta-data representation standard; open content-based query protocols; effective methods for summarizing visual information; further exploitation of query previewing techniques; exploration of both push and pull methods; better methods of progressive storage and retrieval of visual information; improved functionalities of the image and video representations; and methods of operating directly on compressed visual information. Figure 5 shows an architecture including these components.

Conclusions

The Internet contains a growing wealth of visual information. But the technologies for searching for it are inadequate. In the short term, the Internet's limited bandwidth constrains some of the functionality of visual information retrieval systems. However, we already see deployment of such sophisticated retrieval applications as virtual museums, online news reporting, stock-photography systems, and on-demand streaming video retrieval.

Developing and integrating the potential variety of search methods for visual information is a great challenge. Every day we find new retrieval systems that provide new types of search methods we call generically "Query by X." X can refer to an image example; a visual sketch; a specification of visual features, such as color, texture, shape, and motion; a keyword; or a semantic class. The emergence of visual information retrieval search engines on the Web is accelerating the pace of R&D for content-based visual querying. However, deploying this technology in the unconstrained and distributed Web environment is especially challenging. Several critical barriers remain, including heterogeneity of the visual information, lack of interoperable retrieval systems, the distributed and transient nature of the Web's visual information, and the limited bandwidth for transmitting visual information. Development of scalable solutions for visual information systems requires contributions from many disciplines, as well as better assessments of user needs and application requirements. C

References

1. Bach, J., Fuller, C., Gupta, A., Hampapur, A., Horowitz, B., Humphrey, R., Jain, R., and Shu, C. Virage image search engine: An open framework for image management. In *Proceedings of the Symposium on Storage & Retrieval for Image and Video Databases IV, IS&T/SPIE* (San Jose, Calif., Feb. 1996).
2. Beigi, M., Benitez, A., and Chang, S.-F. MetaSEEk: A content-based meta-search engine for images. To appear in *Proceedings of the Symposium on Storage and Retrieval for Image and Video Databases VI, IS&T/SPIE* (San Jose, Calif., Jan. 1998). (See demo at http://www.ctr.columbia.edu/metaseek.)
3. Chang, S.-F., Chen, W., Meng, H., Sundaram, H., and Zhong, D. VideoQ: An automatic content-based video search system using visual cues. In *Proceedings of ACM Multimedia 1997* (Seattle, Wa., Nov. 1997). (See demo at http://www.ctr.columbia.edu/videoq.)
4. Davis, M. Media Streams: An iconic language for video annotation. *Telektronikk 89*, 4 (1993), 59–71.
5. Drelinger, D., and Howe, A. Experiences with selecting search engines using meta-search. To appear in *ACM Trans. Inf. Syst.*
6. Flickner, M., Sawhney, H., Niblack, W., Ashley, J., Huang, Q., Dom, B., Gorkani, M., Hafner, J., Lee, D., Petkovic, D., Steele, D., and Yanker, P. Query by image and video content: The QBIC system. *IEEE Comput. 28*, 9 (Sept. 1995), 23–32.
7. Forsyth, D., Malik, J., Fleck, M., Greenspan, H., Leung, T., Belongie, S., Carson, C., and Bregler, C. Finding pictures of objects in large collections of images, Tech. Rep. CSD-96-905, Dept. of EECS, Computer Science Div., Univ. of California, Berkeley, 1996.
8. Hauptmann, A., and Smith, M. Text, speech, and vision for video segmentation: The Informedia project. In *Proceedings of AAAI Fall Symposium, Computational Models for Integrating Language and Vision* (Boston, Nov. 10–12, 1995).
9. Minka, T., and Picard, R. An image database browser that learns from user interaction. Tech. Rep. 365, MIT Media Laboratory and Modeling Group, 1996.
10. Shatford, S. Analyzing the subject of a picture: A theoretical approach. *Cat. Classif. Q.* 6, 3 (Spring 1986).
11. Smith, J., and Chang, S.-F. Visually searching the Web for content. *IEEE Multimedia 4*, 3 (Summer 1997), 12–20. (See also Tech. Rep. 459-96-25 and demo., Columbia Univ., http://www.ctr.columbia.edu/webseek.)
12. Zhang, H., Low, C., and Smoliar, S. Video parsing and browsing using compressed data. *J. Multimedia Tools Appl. 1*, 1 (Mar. 1995), 89–111.

Shih-Fu Chang (sfchang@ee.columbia.edu) is an associate professor in the Department of Electrical Engineering and the New Media Technology Center at Columbia University in New York.
John R. Smith (jrsmith@watson.ibm.com) is a member of the research staff at the IBM Thomas J. Watson Research Center in Hawthorne, N.Y.
Mandis Beigi (mandis@ee.columbia.edu) is a graduate student in the Department of Electrical Engineering in Columbia University in New York.
Ana Benitez (ana@ee.columbia.edu) is a graduate student in the Department of Electrical Engineering in Columbia University in New York.