

Visually Searching the Web for Content

John R. Smith and Shih-Fu Chang
Columbia University

New visual information in the form of images, graphics, animations, and videos is published on the Web at an incredible rate. However, cataloging it exceeds the capabilities of current text-based Web search engines. WebSeek provides a complete system that collects visual information from the Web by automated agents, then catalogs and indexes it for fast search and retrieval.

Many search engines index the plethora of documents on the World Wide Web. For example, recent systems by Lycos, Alta Vista, Infoseek, and Excite index Web documents by their textual content. These systems periodically scan the Web, analyze the documents, and create compact and searchable indexes. Users enter query terms and/or select subjects to more easily find the desired Web documents.

A large component of the content on the Web consists of visual information such as images, graphics, animations, and videos. The visual information on the Web proves highly volatile. Web site developers are constantly adding, removing, and replacing images and videos. Cataloging them requires a highly efficient automated system that regularly traverses the Web, detects visual information, and processes it and indexes it in such a way as to allow efficient and effective search and retrieval.

We developed WebSeek, a prototype image and video search engine located at <http://www.ctr.columbia.edu/webseek> to fulfill this need.¹ The system collects images and videos from the Web, catalogs them, and provides tools for searching and browsing. WebSeek uses text and visual information synergically to catalog and search. The complete system possesses several powerful functionalities. It

- automates collection of visual information from the Web,
- automates image and video subject classification,

- performs searches using innovative visual content-based techniques,
- conducts image and video subject searches and navigations, and text-based searches,
- compresses images and videos for displaying query results,
- performs search results list manipulations such as intersection, subtraction, and concatenation, and
- modifies queries using content-based relevance feedback.

Recently, researchers have reported on several World Wide Web image search engines.¹⁻⁴ The Webseer system from the University of Chicago³ provides a system for searching for images and animations. It detects faces within the images and lets users search by the number of faces. Webseer does not classify the images into subject categories, nor does it handle content-based searching. The Interpix⁴ image search engine does content-based image searching using color histograms. The Interpix system has been integrated with Yahoo's search engine for retrieving images related to some of Yahoo's Web categories.

Image and video collection process

In WebSeek, autonomous Web agents or "spiders" collect the images and videos. The spiders traverse the Web by following hyperlinks between documents. They detect images and videos, retrieve and process them, and add the new information to the catalog. The overall collection process, illustrated in Figure 1, is carried out using several distinct modules, including

- the Traversal Spider, which assembles lists of candidate Web documents that may include images, videos, or hyperlinks to them;
- the Hyperlink Parser, which extracts the Web addresses (called uniform resource locators or URLs) of the images and videos; and
- the Content Spider, which retrieves, analyzes, and iconifies the images and videos.

Image and video detection

In the first phase, the Traversal Spider traverses the Web looking for images and videos, as illus-

trated in Figure 2. Starting from seed URLs, the Traversal Spider follows a breadth-first search path. It retrieves Web documents via hypertext transfer protocol (http) and passes the Hypertext Markup Language (HTML) code to the Hyperlink Parser. In turn, the Hyperlink Parser detects new URLs—encoded as HTML hyperlinks—and adds them back to the queue of Web documents for the Traversal Spider to retrieve.

The Hyperlink Parser detects the hyperlinks in the Web documents and converts the relative URLs into absolute addresses. By examining the types of hyperlinks and filename extensions of the URLs, the Hyperlink Parser assigns each URL to one of several categories: image, video, or HTML. As illustrated in Table 1, the Multipurpose Internet Mail Extensions (MIME) content-type labels provide the mapping between filename extensions and Web object type.

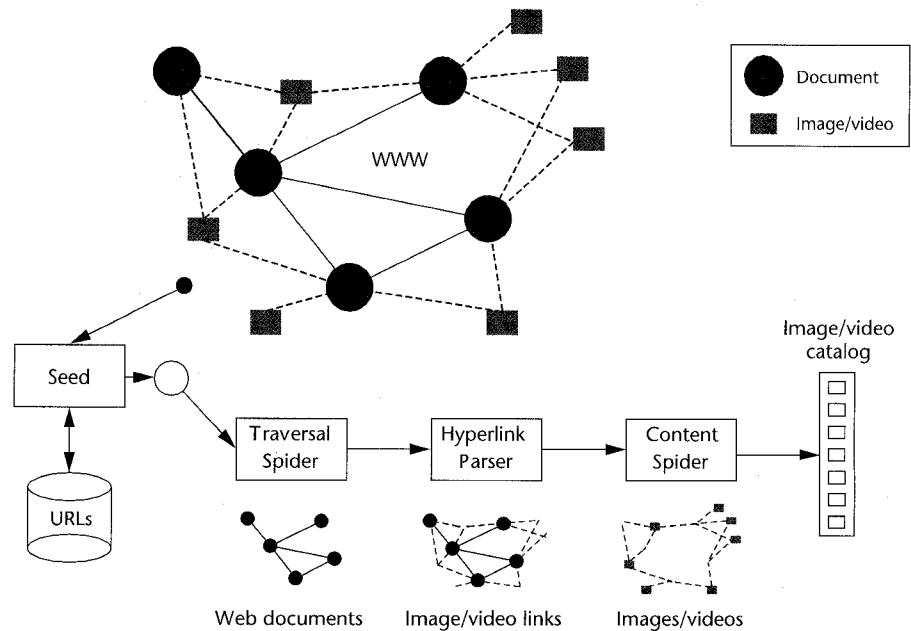


Figure 1. Image and video gathering process.

In the second phase, the list of image and video URLs from the Hyperlink Parser passes onto the Content Spider. The Content Spider retrieves

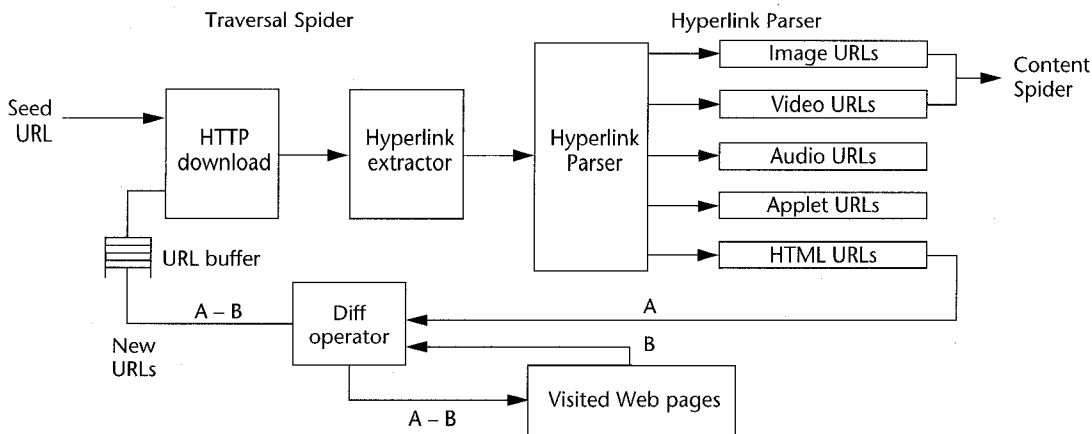


Figure 2. The Traversal Spider traverses the Web and assembles the list of URLs of images and videos.

Table 1. MIME mapping between extensions and object types.

Extension	Type	Description
gif	Image	CompuServe image format
jpg, jpeg, jpe, jfif, pjpeg, pjp	Image	JPEG image format
qt, mov, moov	Video	Quicktime video format
mpeg, .mpg, mpr, mpv, vbs, mpegv	Video	MPEG video format
avi	Video	Microsoft video format
htm, .html	HTML	Hypertext markup language

the images and videos, processes them, and adds them to the WebSeek catalog. The Content Spider performs three important functions. It

1. extracts visual features that allow for content-based searching, browsing, and grouping;
2. extracts other attributes such as width, height, number of frames, type of visual data (that is, photo, graphic, and video); and
3. generates an icon, or motion icon, which sufficiently compacts and represents the visual information to be used for browsing and displaying query results.

The process of extracting visual features generates a color histogram and a table of extracted color regions for each image and video. We discuss the process in more detail in the section "Content-based techniques." The other attributes of the images and videos populate the WebSeek database tables, which we define in the "Catalog database" section. Finally, the Content Spider generates coarse and highly compressed images and videos to provide pictorial data in the query output.

Image and video presentation

We obtained image icons by spatially reducing the original images. And by spatially and temporally reducing the original videos, we generated the video icons. Achieving temporal reduction requires three steps: First, we retained only one frame per every one second of video. Next, we identified the sequence's key frames using automated shot detection.⁵ Finally, we reanimated the video from the key frames and packaged it as an animated GIF file.

Subject classification process

Because text provides clues about the semantic content of the visual information, it is important to the cataloging process. In particular, every image and video on the Web has a unique Web address, or URL, and possibly other HTML tags. We process the URLs and HTML tags in the following ways to index the images and videos:

- Extract terms t_k
- Extract directory names d_i
- Map key terms t_k^* to subjects s_m using a key-term dictionary M_{km}

- Map directory names d_i to subjects s_m

Text processing

Images and videos are published on the Web by two distinct methods, inline and reference. The HTML syntax differs in the two cases, as follows:

1. To inline, or embed, an image or video into a Web document, includes the following code in the document: ``, where the URL gives the relative or absolute address of the image or video. The optional `alt` tag specifies the text that substitutes for the inlined image or video when the image or video is not displayed.
2. Alternatively, images and videos may be referenced from parent Web documents using the following code: `[hyperlink text]`, where the optional `[hyperlink text]` provides the highlighted text that describes the image or video pointed to by the hyperlink.

Term extraction. The terms, t_k , are extracted from the image and video URLs, `alt` tags, and hyperlink text by chopping the text at nonalpha characters. For example, the URL of an image or video has the form `URL = http://host.site.domain[:port]/[user]/[directory]/[file].[extension]`. Here [...] denotes an optional argument. For example, typical URLs are

$URL_1 = \text{http://www.mynet.net:80/animals/domestic-beasts/dog37.jpg}$

$URL_2 = \text{http://camille.gsfc.nasa.gov/rsd/movies2/Shuttle.gif}$

$URL_3 = \text{http://www.arch.columbia.edu/DDL/projects/amiens/slide6b.gif}$

Terms are extracted from the directory and file strings using functions F_{key} and F_{chop} where

$$F_{key}(URL) = F_{chop}(\text{directory/file}),$$

and $F_{chop}(\text{string})$ gives the set of substrings delimited by nonalpha characters. For example,

$$F_{key}(URL_1) = F_{chop}(\text{"animals/domestic-beasts1/dog37"}) = \text{"animals," "domestic," "beasts," "dog."}$$

The process of extracting terms produces an

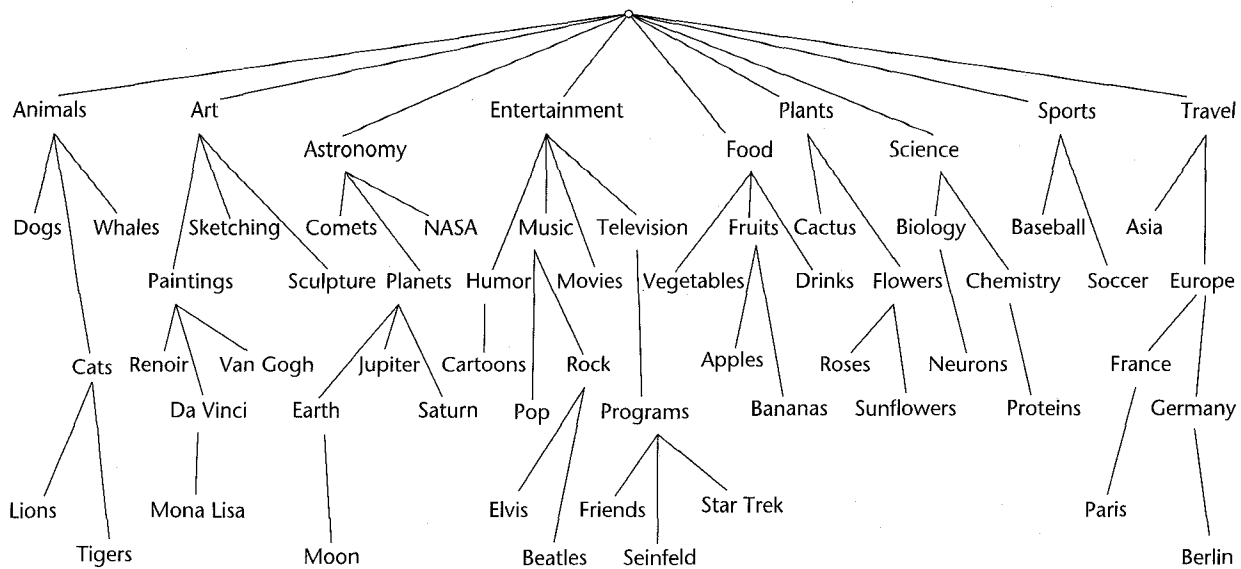


Figure 3. Portion of the image and video subject taxonomy.

overall set of terms $\{t_k\}$ for the image and video collection. The system indexes the images and videos directly using $\{t_k\}$. In addition, certain terms, such as key terms, t_k^* , will map the images and videos to subject classes, as we explain shortly.

Directory name extraction. A directory name, d_i , is a phrase extracted from the URLs that groups images and videos by location on the Web. The directory name contains the directory portion of the URL, namely $F_{dir}(URL) = \text{directory}$. For example, $F_{dir}(URL_1) = \text{"animals/domestic-beasts."}$ The directory names $\{d_i\}$ are also used by the system to map images and videos to subject classes.

Image and video subject taxonomy

A subject class or subject, s_m , is an ontological concept that represents the semantic content of an image or video, for example, subject = "The Beatles." A subject taxonomy is an arrangement of the set of subject classes $\{s_m\}$ into a hierarchy, that is "The Beatles" \in "Rock music" \in "Music." Figure 3 shows a portion of the subject taxonomy that we're developing for image and video topics. When we detect potentially descriptive terms, such as $t_k = \text{"dog"}$ or $t_k = \text{"canine,"}$ we add an appropriate subject class, such as $s_m = \text{"animals/dogs,"}$ to the subject taxonomy.

Key-term dictionary

The key terms, t_k^* , are terms manually identified as semantically related to one or more subject classes, s_m . The key-term dictionary contains the set of key terms $\{t_k^*\}$ and the related subject classes $\{s_m\}$. As such, the key-term dictionary provides a

set of mappings $\{M_{km}\}$ from key terms to subject classes, where $M_{km} : t_k^* \rightarrow s_m$.

We build the key-term dictionary in a semiautomated process. In the first stage, the term histogram for the image and video archive is computed such that each term t_k is assigned a count number f_k , which indicates how many times the term appeared. Then, we rank the terms by highest count f_k and present them in this order of priority for manual assessment.

The goal of manual assessment is to assign qualified terms $t_k^* \in \{t_k\}$ to the key-term dictionary. The qualified terms should be descriptive and not homonymic. Ambiguous words make poor key terms. For example, "rock" is a not a good key term due to its ambiguity. "Rock" can refer to a large mass of stone, rock music, a diamond, or several other things. Once we add a term and its mappings to the key-term dictionary, it's used to assign the images and videos in the collection to subject classes.

From the initial experiments of cataloging more than 500,000 images and videos, we extracted 11,550 unique terms. Some examples are listed in Tables 2 and 3. Notice in Table 2 that some of the most frequent terms are not sufficiently descriptive of the visual information, such as "image" and "picture." The terms in Table 3 (next page) unambiguously define the subject of the images and videos, like "aircraft," "gorilla," and "porsche."

Table 2. Sample terms and their counts $\{t_k : f_k\}$ taken from the assessment of more than 500,000 images and videos.

Term: t_k	Count: f_k
Image	86,380
Gif	28,580
Icon	14,798
Pic	14,035
Img	14,011
Graphic	10,320
Picture	10,026
Small	9,442
Art	8,577

Table 3. Sample key terms counts $\{t_k, f_k\}$ with subjects s_m and mappings $M_{km} : t_k \rightarrow s_m$ taken from the assessment of more than 500,000 images and videos.

Key term: t_k	Count: f_k	Mapping M_{km} to subject s_m
Planet	1,175	Astronomy/Planets
Music	922	Entertainment/Music
Aircraft	458	Transportation/Aircraft
Travel	344	Travel
Gorilla	273	Animals/Gorillas
Starwars	204	Entertainment/Movies/Films/Starwars
Soccer	195	Sports/Soccer
Dinosaur	180	Animals/Dinosaurs
Porsche	139	Transportation/Automobiles/Porsches

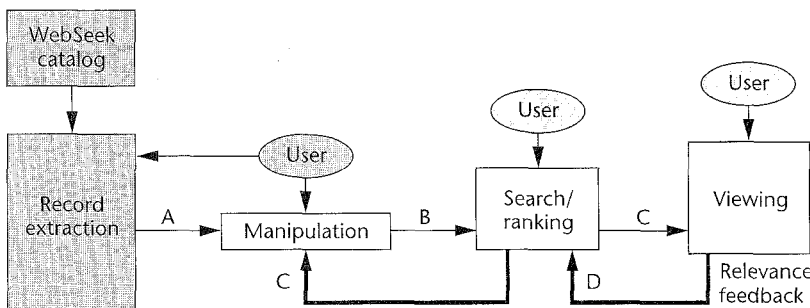


Figure 4. Search, retrieval, and search results list manipulation processes.

In a similar process, the directory names d_i map to subject classes s_m . For example, the directory $d_i = \text{"gallery/space/planets/saturn"}$ is mapped to subject $s_m = \text{"astronomy/planets/saturn."}$ Similar to the process for key-term extraction, the system computes a histogram of directory names $\{d_i : f_i\}$ and presents it for manual inspection. The directories that sufficiently group images and videos related to particular topics are then mapped to the appropriate subject classes.

In the section "Subject classification evaluation," we demonstrate that these methods of key-term and directory-name extraction coupled with subject mapping provide excellent performance in classifying the images and videos by subject. By incorporating some results of natural language processing in addition to using visual features, we hope to further improve and automate the subject classification process.

Catalog database

As described above, each retrieved image and video is processed and the following information tables become populated:

- Images — Imid URL name format width height frames date
- Types — Imid type
- Subjects — Imid subject
- Text — Imid term
- FV — Imid color histogram
- Regions — Imid regid color-set X Y width height

where the special (nonalphanumeric) data types are

- Type \in {Color photo, Color graphic, Video, Black and white image, Gray image}
- Subject \in {Subject classes from taxonomy $\{s_m\}$, partially depicted in Figure 3}
- Color histogram $\in R^{166}$ (166-bin histogram)
- Color set $\in B^{166}$ (166-element binary set)

We explain the automated assignment of Type to the images and videos using visual features in the section "Automated type assessment." Queries on the database tables Images, Types, Subjects, and Text are performed using standard relational algebra. For example, the query, give me all records with Type = "video," Subject = "news," and Term = "basketball" can be carried in SQL as follows:

```

    Select Imid
    From Types, Subjects, Text
    Where Type = "video" and Subject = "news"
    and Term = "basketball"
    
```

However, content-based queries, which involve table FV and Regions, require special processing, which we will discuss later.

Search and retrieval processes

To retrieve images and videos, users may enter search terms t_k or select an image or video subject s_m . Figure 4 shows the search process and model for user interaction. The query for images and videos produces search results A. For example, Figure 5 illustrates the search results for a query for images and videos related to $s_m = \text{"nature."}$ That is, A = Query (Subject = "nature"). Users may manipulate, search, or view A.

Search results list manipulation

The results of a prior search may be fed back to the manipulation module as list C, as illustrated in Figure 5. Users manipulate C by adding or removing records. This happens by issuing a new

query that creates a second, intermediate list A. Users then generate the new list B by selecting one of the following manipulations on C using A—for example, define $C = \text{Query}(\text{Subject} = \text{"nature"})$ and $A = \text{Query}(\text{Term} = \text{"sunset"})$, then

- Union: $B = A \cup C$, that is, $B = \text{Query}(\text{Term} = \text{"sunset"} \text{ or } \text{Subject} = \text{"nature"})$
- Intersection: $B = A \cap C$, that is, $B = \text{Query}(\text{Term} = \text{"sunset"} \text{ and } \text{Subject} = \text{"nature"})$
- Subtraction: $B = C - A$, that is, $B = \text{Query}(\text{Subject} = \text{"nature"} \text{ and } \text{Term} = \text{"sunset"})$
- Replacement: $B = A$, that is, $B = \text{Query}(\text{Term} = \text{"sunset"})$
- Keep: $B = C$, that is, $B = \text{Query}(\text{Subject} = \text{"nature"})$

Content-based visual query

Users may browse and search list B using both content-based and text-based methods. In content-based searching, the output is a new list C, where $C \subseteq B$ becomes an ordered subset of B such that C is ordered by highest similarity to the image or video selected by the user (as described in the section "Content-based techniques"). The search and browse operations may be conducted on input list B or the entire catalog.

For example, $C = B \approx D$, where \approx means "visually similar," ranks list B in order of highest similarity to the selected item in D. For instance, the content-based visual query

$C = \text{Query}(\text{Subject} = \text{"nature"}) \approx D$ ("mountain scene image")

ranks the "nature" images and videos in order of highest visual similarity to the selected "mountain scene image." In Figure 6, the query $C = B \approx D$ ("red race car") retrieves the images and videos from the subject class $\text{Subject} = \text{"transportation/automobiles/ferraris"}$ that are visually similar to the selected image of a "red race car."

Alternatively, users can select one of the items and use it to search the entire catalog for similar items. For example, $C = A \approx D$ ranks list A, where in this case A is the full catalog, in order of highest visual similarity to the selected item in D.

Content-based techniques

Our system provides two methods for content-based searching: by color histograms and by spa-




Figure 5. Search results for Subject = "nature."

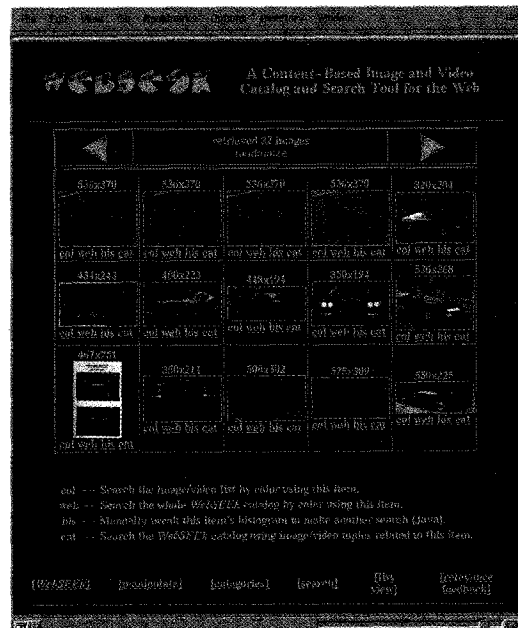


Figure 6. Content-based visual query results for images and videos.

tial locations and arrangements of color regions. We adopted these techniques to use domain-independent approaches.⁶ The content-based methods developed here for indexing, searching, and navigation can be applied—in principle—to other types of features, such as texture and motion, and to other application domains.

Global feature query

The color histograms describe the distribution of colors in each image or video. We define the

Figure 7. Spatial and feature (SaFe) query specifies color regions on the query grid. The system finds the images that contain similar regions in similar spatial arrangements to the query.

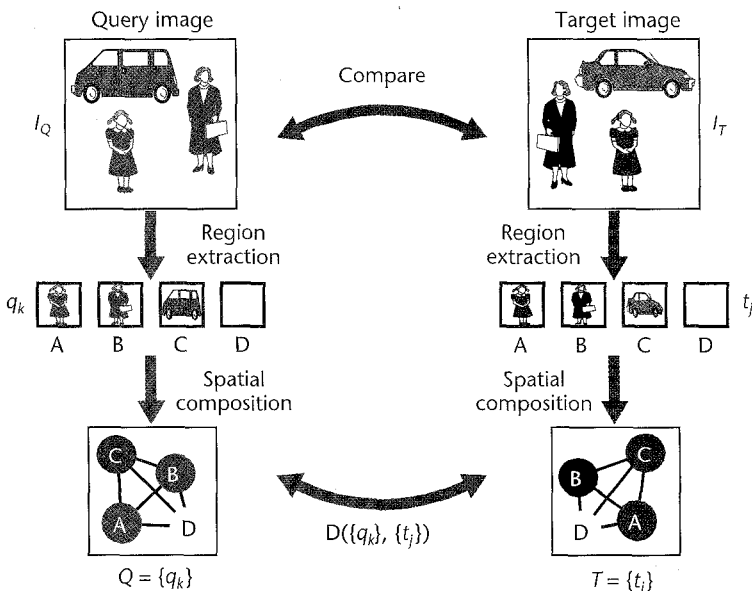
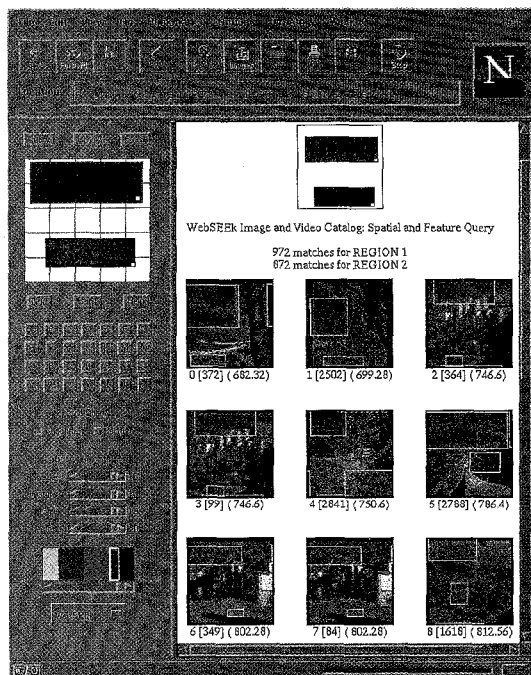


Figure 8. Integrated spatial and feature query process. The query image I_Q defined by users contains a set of color regions $A = \{q_k\}$. The target images I_T in the database decompose into sets of color regions $T = \{t_j\}$. The similarity between the query image and target images is measured using the joint spatial and feature distance function $D(\{q_k\}, \{t_j\})$.

color histograms as discrete, 166-bin distributions in a quantized hue saturation value (HSV) color space.⁶ The system computes a color histogram for each image and video scene, which is used to assess its similarity to other images and video

scenes. The color histograms also automatically assign the images and videos to type classes using Fisher discriminant analysis, as described in the section "Automated type assessment."

Color histograms dissimilarity. A color histogram dissimilarity metric determines the color dissimilarity between a query image and a target image.^{7,8} To achieve high efficiency in the color histogram query process, we developed the technique of query optimized distance (QOD) computation.⁹ This technique efficiently retrieves target histograms and computes histogram dissimilarities. In the QOD technique, the retrieval of candidate target histograms is prioritized in the order of the most significant colors in the query histogram during retrieval. Smith gives a more complete description elsewhere.⁹

We found that setting the "color significance" threshold adaptively to select only the 80- to 90-percent most significant colors in the query histogram decreases query time dramatically without degrading retrieval effectiveness. By using the QOD computation method, we can greatly reduce the query processing time, as demonstrated in the section "Efficiency."

Integrated spatial and color query

While color histograms prove useful for indexing images by global color, a significant aspect of discriminating among images depends on the sizes, spatial locations, and relationships of color regions within the images. Therefore, we developed a system for querying by spatial locations and arrangements of color regions.⁶ With this integrated spatial and feature image-query engine, users can graphically construct a query by placing color regions on a query grid, as illustrated in Figure 7.

We computed the similarity between the query image and each target image by assessing the similarities of the spatial locations, colors and sizes of the regions in the query, and target images (as depicted in Figure 8). The integrated spatial and features system provides greater power and flexibility in querying for images than the color histogram-based methods.⁶ We found that retrieval effectiveness improves substantially using this method.⁹

Automated type assessment

By training on the color histograms of image and video samples, we developed a process of automated type assessment using Fisher discrim-

inant analysis (FDA). FDA constructs a series of uncorrelated linear weightings of the color histograms, which gives maximum separation between training classes.¹⁰ The images are automatically assigned to the nearest Type \in {color photo, color graphic, video, black and white image, gray image} class.¹

In the section "Type classification evaluation" we show that this approach offers excellent performance in automatically classifying the images and videos into these broad type classes. We hope to increase the number of type classes and improve the type-classification system by incorporating other visual features such as spatial and color region information into the classification process.

Evaluation

In the initial trials, the system cataloged 513,323 images and videos from 46,551 directories at 16,773 Web sites. The process took several months and was performed simultaneously with development of the user application. Overall, the system has cataloged more than 129 Gbytes of visual information. Locally storing the information, which includes coarse versions of the data and color histogram feature vectors, requires approximately 2 Gbytes.

Subject classification evaluation

The cataloging process assigned 68.23 percent of the images and videos into subject classes using key-term and directory-name mappings. To evaluate the subject classification system, we assessed the classification rates for several subject classes. Table 4 summarizes the results. The classification performance was excellent—the system provided a classification precision of approximately 92 percent. For this assessment we chose nine classes, depicted in Table 4, at random from the subject taxonomy of 2,128 classes. We established the ground truth by manually verifying each image's and video's subject in the nine test sets.

We observed that errors in classification resulted for several reasons:

1. the key terms defined by the WebSeek system are sometimes used out of context by the publishers of images or videos,
2. the system relies on some ambiguous key

Table 4. Rates of correct subject classification (precision) for a random set of classes.

Subject	Number of Sites	Count	Rate
Art/Illustrations	29	1,410	0.978
Entertainment/Humor/Caroons/Daffy Duck	14	23	1.000
Animals/Possums	2	9	0.444
Science/Chemistry/Proteins	7	40	1.000
Nature/Weather/Snow/Frosty	9	13	1.000
Food	403	2,460	0.833
Art/Paintings/Pissarro	3	54	1.000
Entertainment/Music/MTV	15	87	0.989
Horror	366	2,454	0.968

terms, such as t_k^* = "madonna," and

3. the semantics of the images and videos change when viewed outside of the context of their Web sites.

For example, in this last category, the precision of subject class s_m = "animals/possums" remains low, as depicted in Table 4, because five out of the nine items are not images or videos of possums. These items were classified incorrectly because the key term "possum" appeared in the directory name. While some of the images in that directory depict possums, others depict only the forests to which possums are indigenous. When viewed outside of the context of the "possum" Web site, the images of forests are not expected to be in the class "animals/possums."

Type classification evaluation

Table 5 summarizes the precision of the automated type classification system. For this evaluation, the training and test samples contained 200 images from each type class. We found the automated type assessment for these five simple classes was excellent. The system gave a classification precision of 95 percent. In future work we will extend WebSeek to include a larger number of classes including new type classes, such as fractal images, cartoons, faces, art paintings, and subject classes from the subject taxonomy.

Efficiency

Another important factor in the image and video search system is the speed at which users perform operations and queries. In the initial system, the overall efficiency of various database

Table 5. Rates of correct automated type classification.

Type	Rate
Color photo	0.914
Color graphic	0.923
Gray image	0.967
Black and white image	1.000

manipulation operations was excellent, even on the large catalog (greater than 500,000 images and videos). In particular, the content-based visual query methods' good performance resulted from the strategy for indexing the 166-bin color histograms described in the section "Color histograms dissimilarity." For example, the system identified the $N = 60$ most similar visual scenes in the catalog of 513,323 images and videos to a selected query scene in only 1.83 seconds.

Future work

In future work, we will incorporate other visual dimensions such as texture, shape, and motion to enhance the system's content-based components. We are also incorporating automated techniques for detecting faces.¹¹ We plan to investigate new techniques for exploiting text and visual features independently and jointly to improve the process of cataloging the images and videos. **MM**

References

1. J.R. Smith and S.-F. Chang, "An Image and Video Search Engine for the World Wide Web," *Proc. SPIE 2670 Storage and Retrieval for Still Image and Video Databases IV*, I. Sethi and R. Jain, eds., SPIE, Bellingham, Wash., 1996, pp. 84-95.
2. S. Sclaroff, "World Wide Web Image Search Engines," Tech. Report TR 95-016, Computer Science Dept., Boston Univ., Boston, 1995.
3. C. Frankel, M. Swain, and V. Athitsos, *Webseer: An Image Search Engine for the World Wide Web*, Tech. Report TR-96-14, Computer Science Dept., Univ. of Chicago, July 1996, <http://webseer.cs.uchicago.edu/>.
4. Interpix Image Surfer, <http://www.interpix.com/>.
5. J. Meng and S.-F. Chang, "CVEPS—A Compressed Video Editing and Parsing System," *Proc. ACM Int'l Conf. Multimedia*, ACM Press, New York, Nov. 1996, pp. 43-53, <http://www.ctr.columbia.edu/cveps>.
6. J. R. Smith and S.-F. Chang, "VisualSeek: A Fully Automated Content-Based Image Query System," *Proc. ACM Int'l Conf. Multimedia*, ACM Press, New York, Nov. 1996, pp. 87-93, <http://www.ctr.columbia.edu/VisualSEEK>.
7. M. Flickner et al., "Query by Image and Video Content: The QBIC system," *Computer*, Vol. 18, No. 9, Sept. 1995, pp. 23-32.
8. J. R. Bach et al., "Virage Image Search Engine: An Open Framework for Image Management," *Proc. SPIE 2670 Storage and Retrieval for Still Image and Video Databases IV*, I. Sethi and R. Jain, eds., SPIE, Bellingham, Wash., 1996, pp. 76-87.
9. J. R. Smith, *Integrated Spatial and Feature Image Systems: Retrieval, Analysis, and Compression*, PhD thesis, Columbia Univ., Graduate School of Arts and Sciences, New York, February 1997.
10. K. Fukunaga, *Introduction to Statistical Pattern Recognition*, Harcourt Brace-Jovanovich, New York, 1990.
11. H. Wang and S.-F. Chang, *A Highly Efficient System for Automatic Face Region Detection in MPEG Video Sequences*, Tech. Report CU/CTR/TR 466-96-32, Center for Telecommunications Research, Columbia Univ., Oct. 1996.



John R. Smith is a research staff member at the IBM T.J. Watson Research Center, Hawthorne, New York. He received his M. Phil and PhD. degrees in electrical engineering from Columbia University in 1994 and 1997, respectively. He has developed several prototype image retrieval systems, including the WebSeek image and video search engine, VisualSeek, and the SaFe integrated spatial and feature image system. In 1997, he received the Eliahu I. Jury award from Columbia University for outstanding achievement as a graduate student in the areas of systems communication or signal processing.



Shih-Fu Chang is an assistant professor at the Department of Electrical Engineering and Center for Telecommunications Research at Columbia University. His current research interests include visual information systems, content-based visual query, and digital video processing and communications. His group has developed several online prototypes in visual information systems, including a network-based compressed video editing and search system, WebClip, and WebSeek. At Columbia, he is also co-leading the development of a video-on-demand/multimedia testbed. Chang received an NSF Career award 1995-98 and an IBM UPP (Faculty Development) Award 1995-97.

Contact Smith at IBM Thomas J. Watson Research Center, 30 Saw Mill River Road, Hawthorne, N.Y. 10532, e-mail jrsmith@watson.ibm.com.