

Content-Based Image Retrieval

Selim Aksoy

Department of Computer Engineering

Bilkent University

saksoy@cs.bilkent.edu.tr

Image retrieval

- Searching a large database for images that match a query:
 - What kind of databases?
 - What kind of queries?
 - What constitutes a match?
 - How do we make such searches efficient?

Applications

- Art Collections
 - Fine Arts Museum of San Francisco
- Medical Image Databases
 - CT, MRI, Ultrasound, The Visible Human
- Scientific Databases
 - Earth Sciences
- General Image Collections for Licensing
 - Corbis, Getty Images
- The World Wide Web
 - Google, Microsoft, Flickr


Corel data set




118011
WATER HARBOR
SKY CLOUDS



TIGER CAT WATER GRASS



1090
SUN CLOUDS
WATER SKY



1015
SUN TREE
PLAIN SKY



143078
MOUNTAINS TREES
aspens VALLEY



102042
MUSEUM memorial
FLAGS GRASS




119094
GARDEN BUILDING
FLOWERS TREES



131007
GARDEN FLOWERS
HOUSE TREES

60,000 images with annotated keywords

Fine Arts Museum of San Francisco



80,000 images

Query formulation

- Text description (keywords)

- Query by example



- Query by sketch



- Symbolic description (man and woman on a beach)
- Relevance feedback

Google query on "rose"



Two-Tone Rose
500 x 375 - 95k - jpg
gardening.about.com
[[More from z.about.com](#)]



Welcome to the Corpus Christi Rose
...
351 x 334 - 51k - gif
www.geocities.com




Frosty Rose
600 x 451 - 70k - jpg
blogs.warwick.ac.uk



Dominic Cavendish reviews The Rose
...
400 x 313 - 40k - jpg
www.telegraph.co.uk



We've been in the business of rose .
400 x 539 - 78k - jpg
www.witherspoonrose.com




Rose
512 x 512 - 19k - jpg
www.parabola.me.uk



ROSE Online
800 x 640 - 393k - jpg
www.gameogre.com



Tracey Rose The Thinker, 1996
500 x 451 - 26k - jpg
www.artthrob.co.za



Tracey Rose Span II, 1997
317 x 500 - 18k - jpg
www.artthrob.co.za



The rose apple is occasionally ...
350 x 307 - 57k - jpg
www.tradewindsfruit.com



Lewis Rose also known as Lewis N. ... Rose
347 x 348 - 18k - jpg
lewisrose.com



Rose Parade Seating Chart Seating ...
449 x 354 - 18k - gif
www.barrystickets.com



Rose Quartz point shape in Sterling ...
456 x 478 - 15k - jpg
www.wehug.com



The Exorcism of Emily Rose
300 x 375 - 29k
www.smh.com.au



rose rose 6499 hits
574 x 576 - 33k - jpeg
www.firstmonday.org

Corel query on "rose"



Corbis query on "rose"



42-17895937 (RM)
Season of Perfume by
Yang Ping

Price Image
Lightbox Cart



42-17895408 (RM)
Autumn Beauties by Wang
Chingho

Price Image
Lightbox Cart



42-17895882 (RM)
Little Black Cat by Wu
Yeizhao

Price Image
Lightbox Cart



42-17860276 (RM)
Home by Hugh Shurley

Price Image
Lightbox Cart



42-17895421 (RM)
Butterflies Among
Flowers by Wang Chingho

Price Image
Lightbox Cart



42-17895445 (RM)
Red Roses by Lu Bisa

Price Image
Lightbox Cart



42-17895962 (RM)
Beautiful Flowers by Cao
Jingen

Price Image
Lightbox Cart



42-16247767 (RM)
Peppers Stuffed with
Rosebuds

Price Image
Lightbox Cart



42-18240866 (RM)
David Bowie Smoking
Cigarette

Price Image
Lightbox Cart



42-17165934 (RM)
Woman Taking a Bath in
Rose Petals

Price Image
Lightbox Cart



42-16246447 (RM)
Oriental-Style Litchi
Salad

Price Image
Lightbox Cart



42-16801960 (RM)
Bouquet of Red Roses

Price Image
Lightbox Cart
See Image Set(s)



42-16801959 (RM)
Bouquet of Red Roses

Price Image
Lightbox Cart



42-16801939 (RM)
Bouquet of Red Roses

Price Image
Lightbox Cart



42-17529104 (RM)
Rose Red #9 Series by
Elisa Lazo de Valdez

Price Image
Lightbox Cart



42-17529137 (RM)
Rose Red #46 Series by
Elisa Lazo de Valdez

Price Image
Lightbox Cart



42-17529136 (RM)
Rose Red #57 Series by
Elisa Lazo de Valdez

Price Image



42-15766292 (RM)
Roses in Kohinoor Suite
Bathroom at Amarvilas...

Price Image



42-16249149 (RM)
Stuffed Red Onions

Price Image
Lightbox Cart



42-15944324 (RM)
Bridal Bouquet

Price Image
Lightbox Cart



42-16248521 (RM)
Red Onion

Price Image
Lightbox Cart



42-16248491 (RM)
Sliced Red Onion

Price Image
Lightbox Cart



42-16246499 (RM)
Glass of Rose Wine

Price Image
Lightbox Cart



42-15766394 (RM)
Rose Petal Bath at
Vanyavilas Resort

Price Image

Difficulties with keywords

- Images may not have keywords.
 - (An image is worth ... how many key-words?)
 - Query is not easily satisfied by keywords.
 - “A casually dressed couple gazing into each others eyes lovingly with dramatic clouds in the background.”
 - “Pretty girl doing something active, sporty in a summery setting, beach - not wearing lycra, exercise clothes - more relaxed in tee-shirt. Feature is about deodorant so girl should look active - not sweaty but happy, healthy, carefree - nothing too posed or set up - nice and natural looking.”
- Content-based image retrieval (CBIR)

Content-based image retrieval

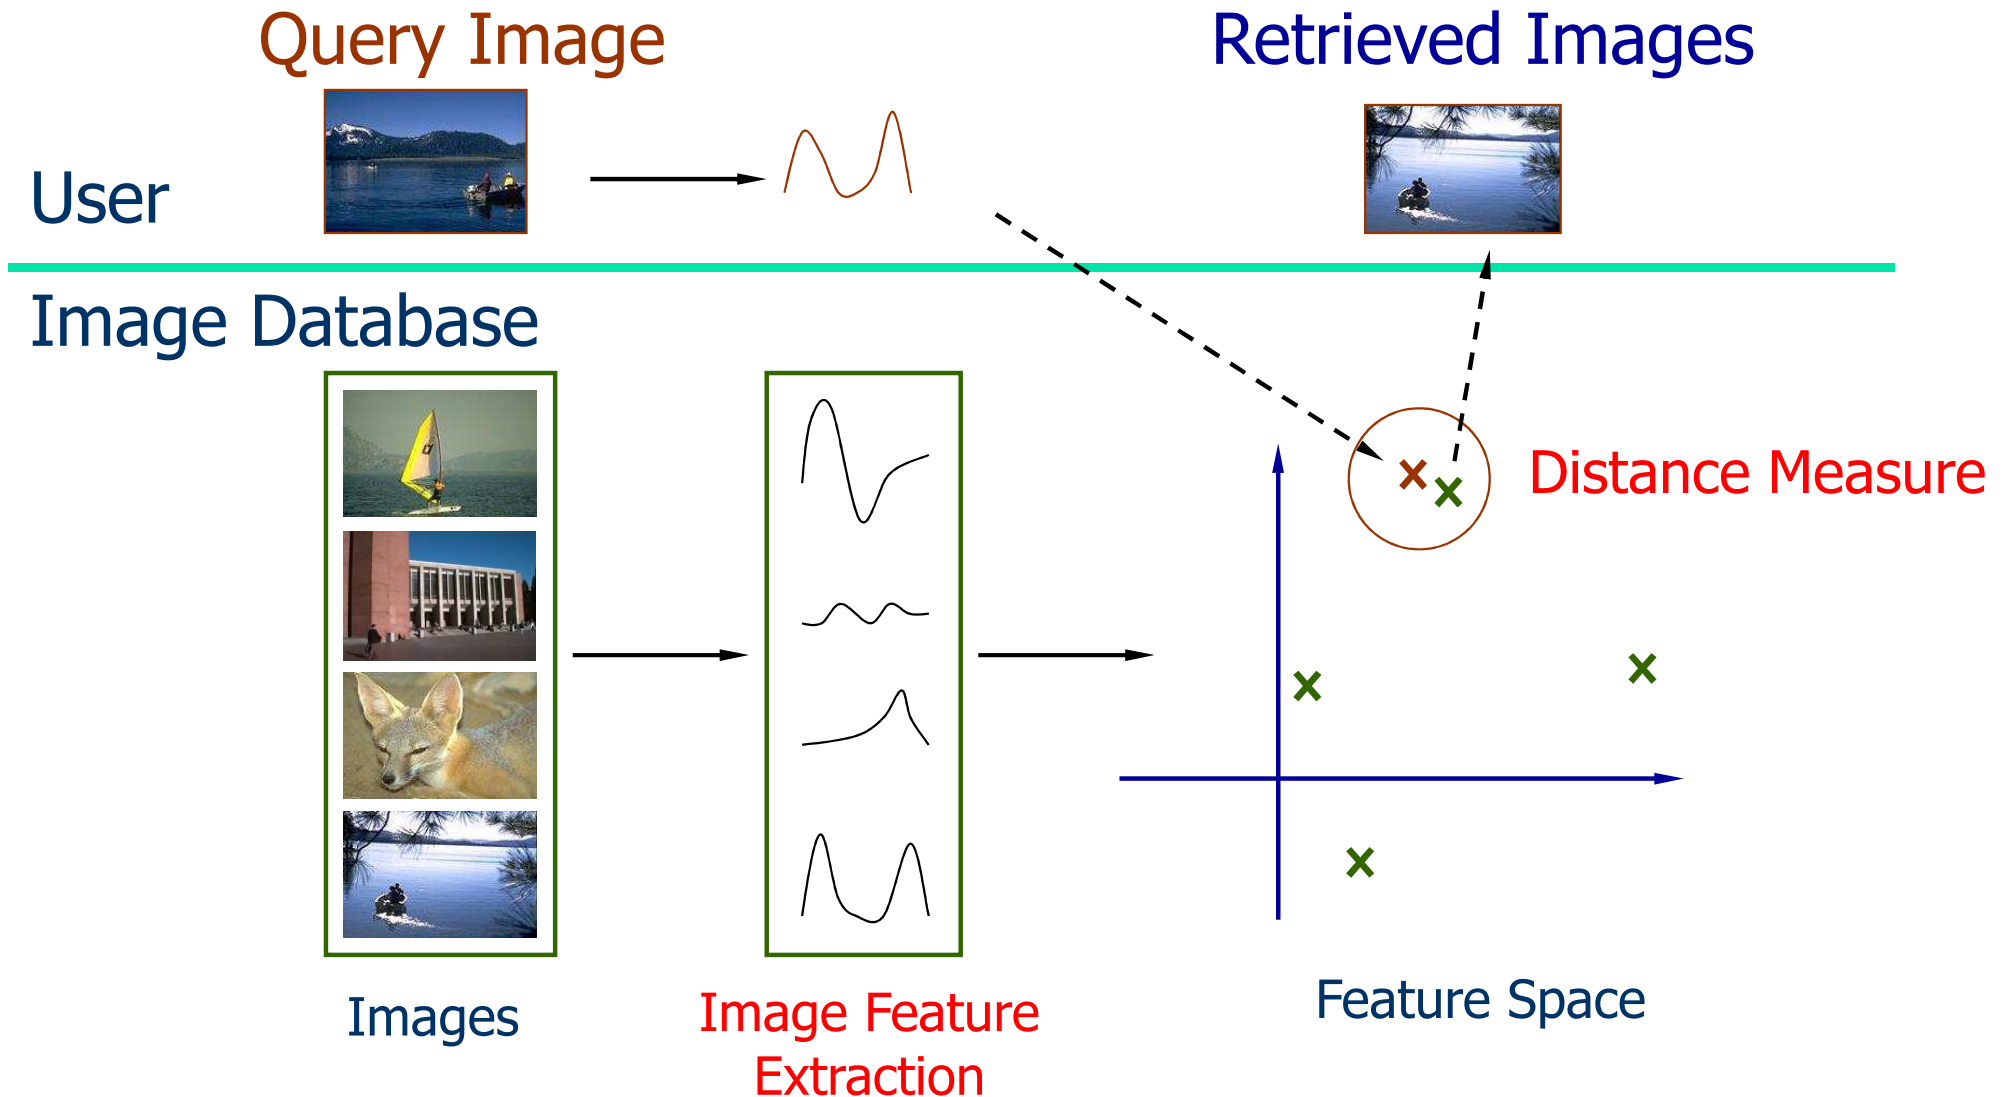


Image representations and features

- Image representations:
 - Iconic
 - Global
 - Region-based
 - Object-based
- Image features:
 - Color
 - Texture
 - Shape
 - Objects and their relationships
(this is the most powerful, but you have to be able to recognize the objects!)

Image similarity

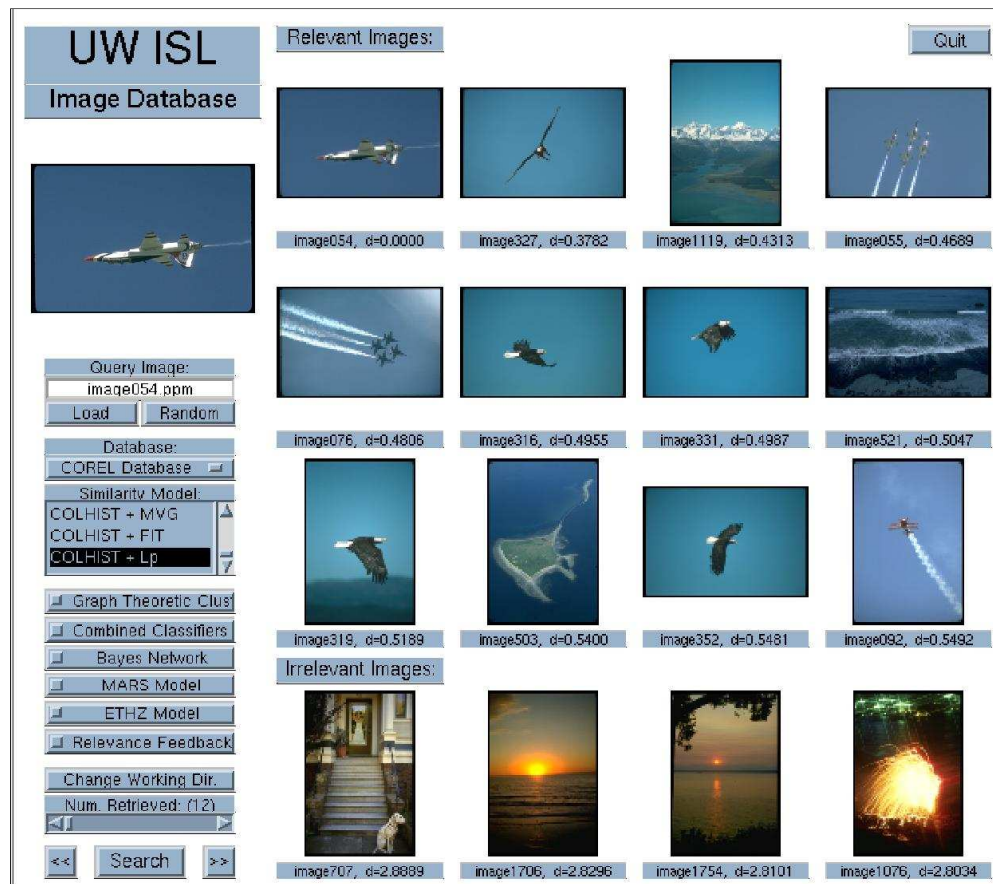
- Distance measures:
 - Euclidean distance
 - Other L_p metrics
 - Histogram intersection
 - Cosine distance
 - Earth mover's distance
- Probabilistic similarity measures:
 - $P(\text{relevance} \mid \text{two images})$
 - $P(\text{relevance} \mid \text{two images}) / P(\text{irrelevance} \mid \text{two images})$

Global histograms

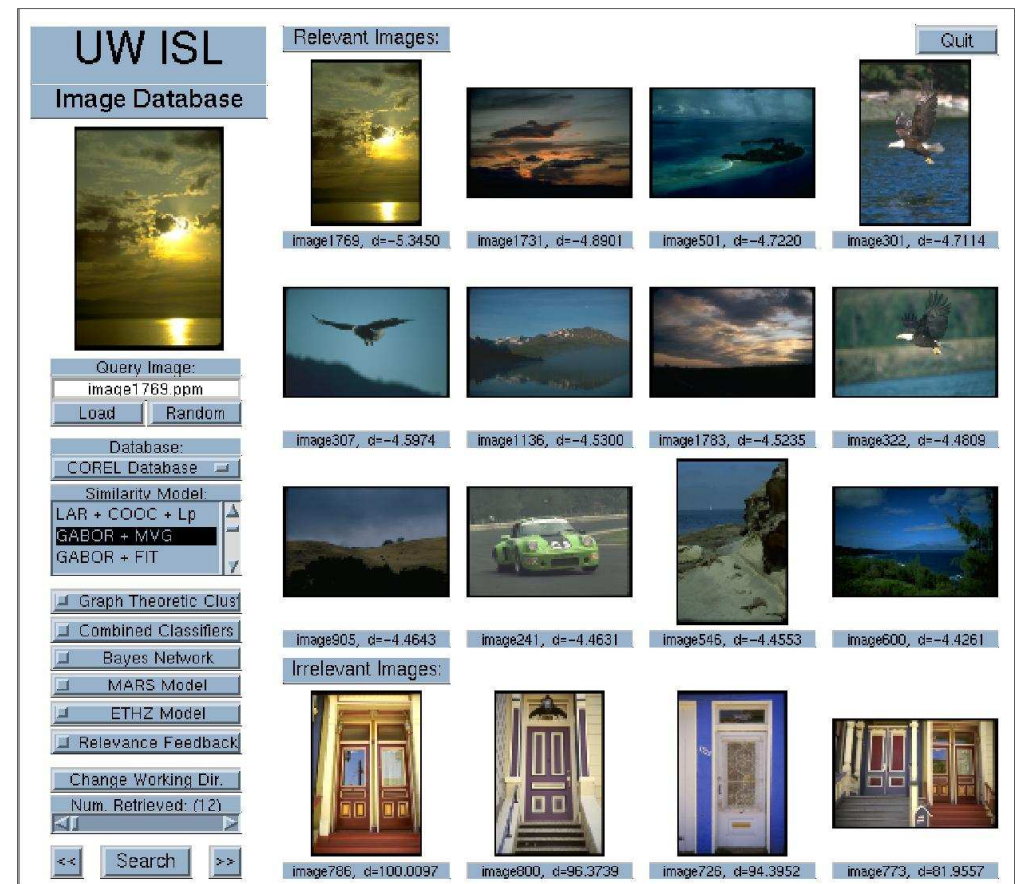
- Searching using global color histograms



Global histograms

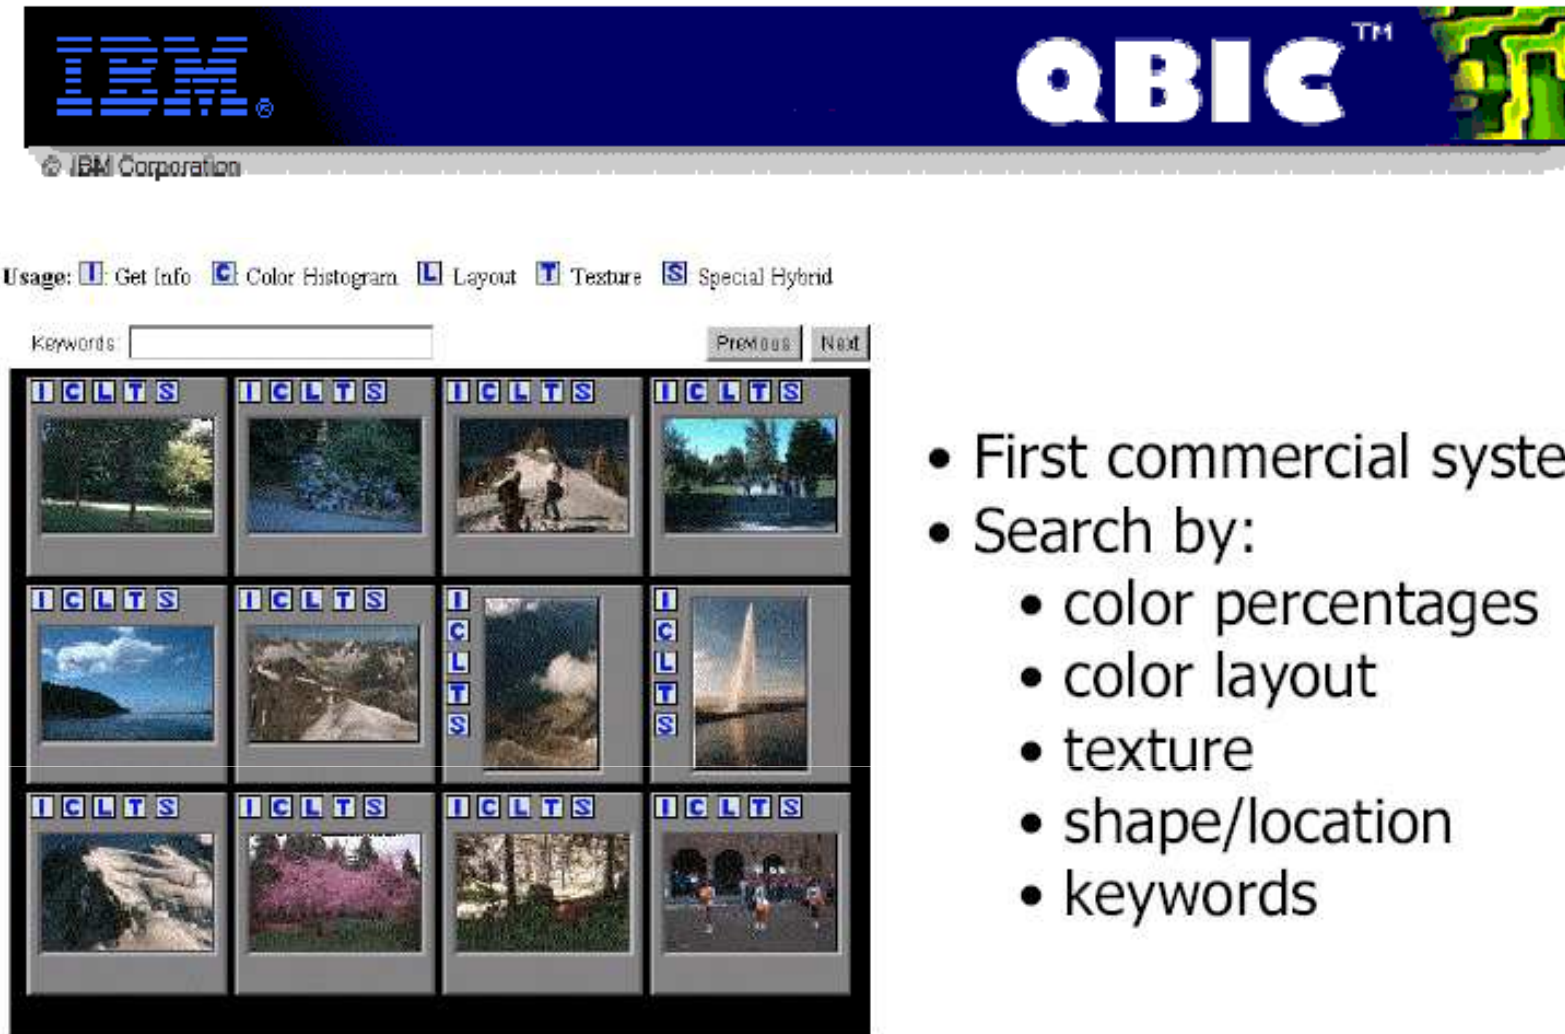


"Airplanes" using color histograms (4/12)



"Sunsets" using Gabor texture (3/12)

Query by image content (QBIC)



- First commercial system
- Search by:
 - color percentages
 - color layout
 - texture
 - shape/location
 - keywords

Try their demo: <http://www.qbic.almaden.ibm.com>

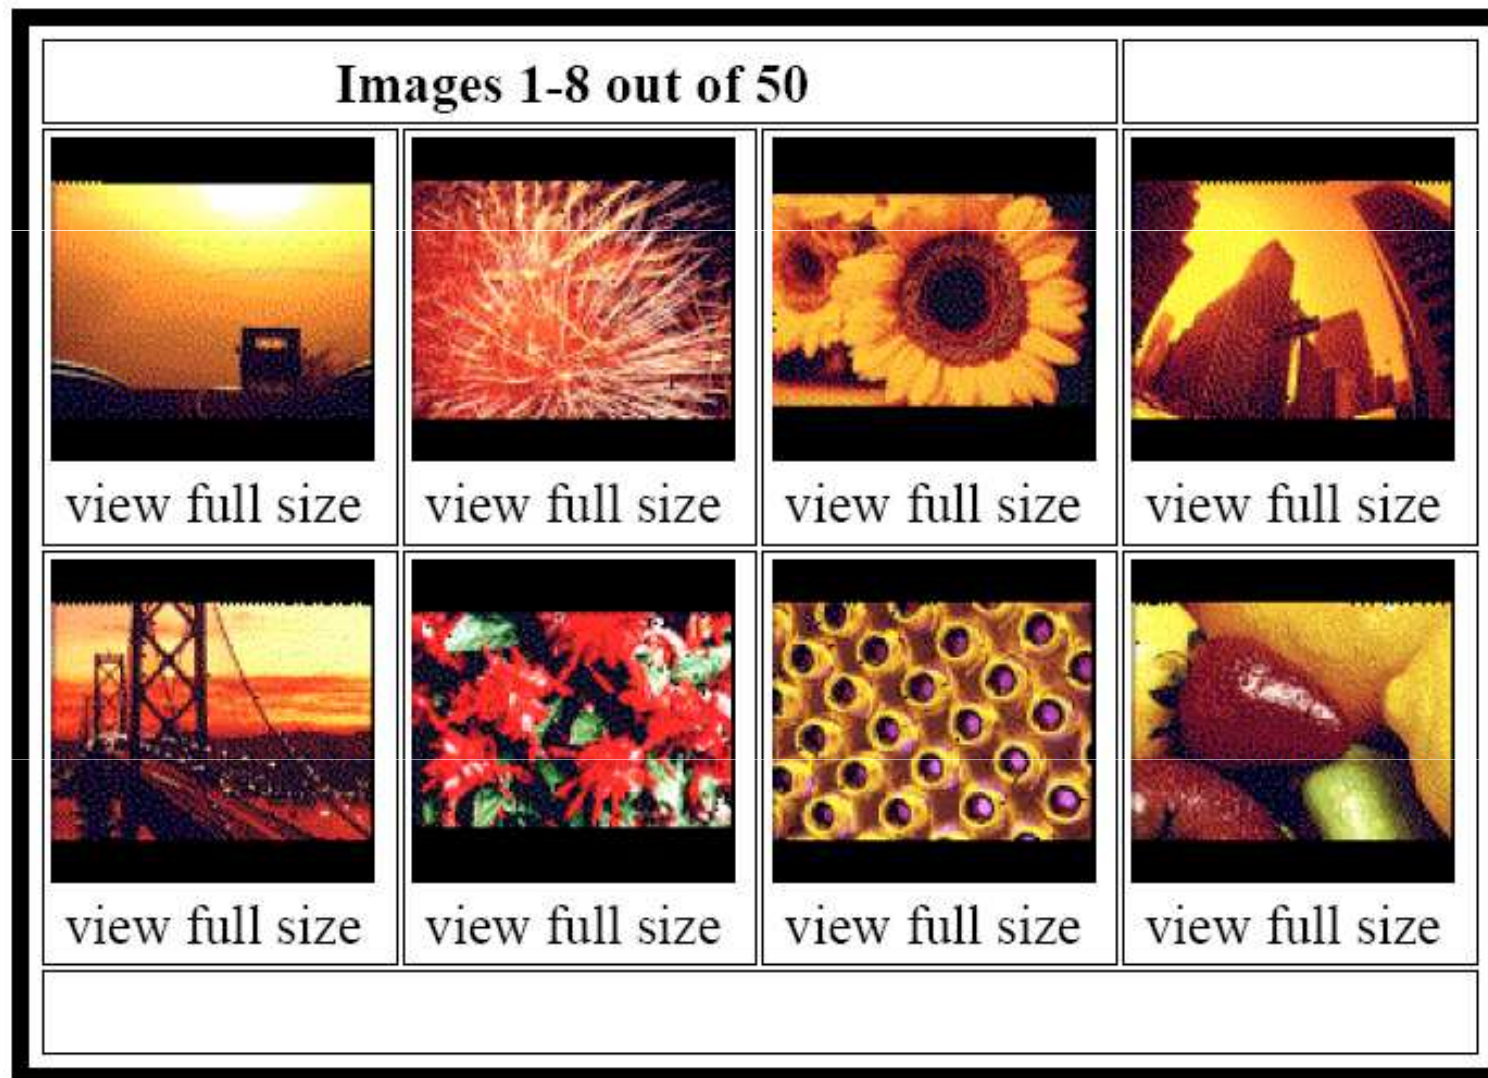
Color histograms in QBIC

- The QBIC color histogram distance is:
$$d_{\text{hist}}(I, Q) = (h(I) - h(Q))^T A (h(I) - h(Q)).$$
- $h(I)$ is a K-bin histogram of a database image.
- $h(Q)$ is a K-bin histogram of the query image.
- A is a K x K similarity matrix.

	R	G	B	Y	C	V
R	1	0	0	.5	0	.5
G	0	1	0	.5	.5	0
B	0	0	1		?	
Y				1		
C		?			1	
V						1


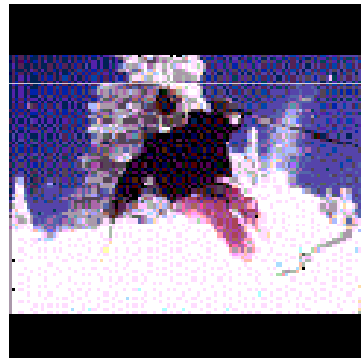
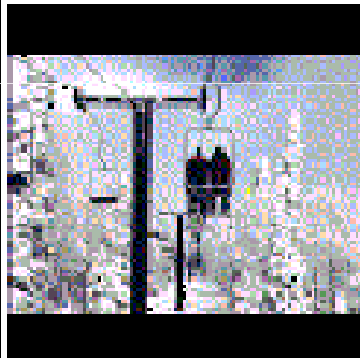

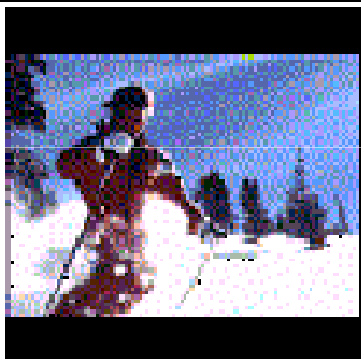
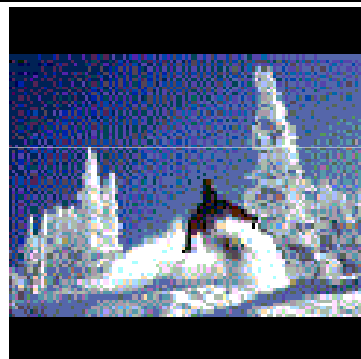
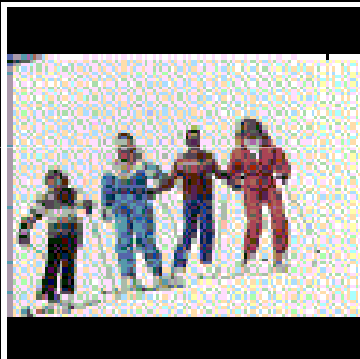
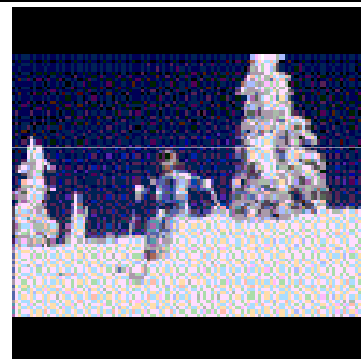
How similar is blue to cyan?

Color percentages in QBIC



%40 red, %30 yellow, %10 black

Color layout in QBIC

Images 1-8 out of 41			
			
view full size	view full size	view full size	view full size
			
view full size	view full size	view full size	view full size
Columns: Rows:			

Earth mover's distance

For each image, compute color signature:



Define distance between two color signatures to be the minimum amount of “work” needed to transform one signature into another.



Earth mover's distance

- Visualization using EMD and multidimensional scaling



Probabilistic similarity measures

- Two classes:
 - Relevance class A
 - Irrelevance class B
- Bayes classifier
 - Assign (ξ_i, ξ_j) to $\begin{cases} A & \text{if } P(A | (\xi_i, \xi_j)) > P(B | (\xi_i, \xi_j)) \\ B & \text{otherwise} \end{cases}$

- Discriminant function for classification

$$\Delta(\xi_i, \xi_j) = \frac{P(A | (\xi_i, \xi_j))}{P(B | (\xi_i, \xi_j))} = \frac{P((\xi_i, \xi_j) | A) P(A)}{P((\xi_i, \xi_j) | B) P(B)}$$

- Rank images according to posterior ratio values based on feature differences.

Probabilistic similarity measures

UW ISL Image Database

Query Image: image1562.ppm
Load Random

Database: COREL Database

Similarity Model:
LAR + COOC + MVG
LAR + COOC + FIT
LAR + COOC + Lp

Graph Theoretic Clus
Combined Classifiers
Bayes Network
MARS Model
ETHZ Model
Relevance Feedback

Change Working Dir.
Num. Retrieved: (12)

Search

Relevant Images:

Quit

image1562, d=-5.1543 image1578, d=-5.1400 image1522, d=-4.8981 image1571, d=-4.7950

image1511, d=-4.7196 image1566, d=-4.6189 image1528, d=-4.5649 image1568, d=-4.4428

image1552, d=-4.3891 image1525, d=-4.3362 image1577, d=-4.1172 image1597, d=-4.0424

Irrelevant Images:

image647, d=56.2184 image946, d=56.2083 image380, d=54.3261 image796, d=52.8174

"Residential interiors" (12/12)

UW ISL Image Database

Query Image: image946.ppm
Load Random

Database: COREL Database

Similarity Model:
LAR + COOC + MVG
LAR + COOC + FIT
LAR + COOC + Lp

Graph Theoretic Clus
Combined Classifiers
Bayes Network
MARS Model
ETHZ Model
Relevance Feedback

Change Working Dir.
Num. Retrieved: (12)

Search

Relevant Images:

Quit

image946, d=-5.1543 image978, d=-4.8187 image909, d=-4.7620 image949, d=-4.6530

image957, d=-4.6435 image1000, d=-4.4657 image983, d=-4.2742 image985, d=-4.0550

image950, d=-4.0355 image931, d=-4.0309 image990, d=-4.0242 image920, d=-4.0066

Irrelevant Images:

image752, d=68.1365 image736, d=61.8827 image725, d=60.9700 image716, d=59.6958


"Fields" (12/12)

Shape-based retrieval


Find more shapes like this



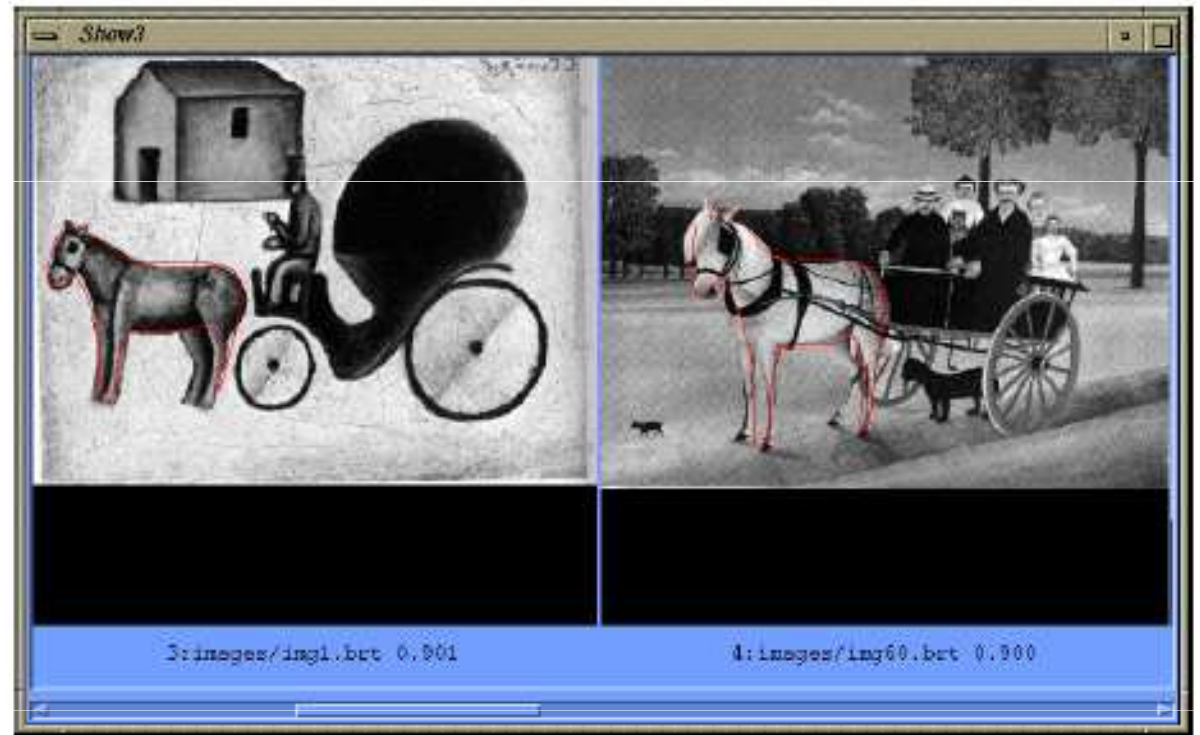
Shape-based retrieval



Elastic shape matching



Sketch-based
query



retrieved images

Iconic matching

Example applications:

- Copyright and trademark protection
- Duplicate removal
- Linking images used in evidence, for example child pornography

Problems in finding “exact” matches:

- Lossy compression, image scanning
- Color space conversion
- Photoshop-style transforms: blur, scale, rotate, warp, crop, cut, etc.

Iconic matching

- Wavelet-based image compression
- Quantization of coefficients

painted



scanned



target





From Jacobs, Finkelstein, & Salesin
Fast Multi-Resolution Image Querying, SIGGRAPH 1995

Iconic matching



Example taken from Jacobs, Finkelstein, & Salesin
Fast Multi-Resolution Image Querying, SIGGRAPH 1995

Region-based retrieval: Blobworld



- Images are segmented on color plus texture
- User selects a region of the query image
- System returns images with similar regions
- Works really well for tigers and zebras

Demo: <http://elib.cs.berkeley.edu/photos/blobworld>

Region-based retrieval: Blobworld

Query image: 108019

Query blobs

blob and feature importance:					
	blob (overall)	color	texture	location	shape
blob 2	very	very	somewhat	not	not
blob 1	somewhat	very	somewhat	not	not

Querying from 10000 images (full search).

1: 10804 (score = 0.9411) [New query](#)

2: 108029 (score = 0.9399) [New query](#)

3: 108024 (score = 0.9375) [New query](#)

4: 108056 (score = 0.9294) [New query](#)

5: 108044 (score = 0.9294) [New query](#)

6: 108051 (score = 0.9294) [New query](#)

7: 108004 (score = 0.9174) [New query](#)

8: 208042 (score = 0.9109) [New query](#)

Query image: 108019

Query blobs

blob and feature importance:					
	blob (overall)	color	texture	location	shape
blob 2	very/val	very	somewhat	not	not
background	very	very	not	not	not

Querying from 10000 images (full search).

1: 10804 (score = 0.9411) [New query](#)

2: 108029 (score = 0.9399) [New query](#)

3: 108024 (score = 0.9375) [New query](#)

4: 108056 (score = 0.9294) [New query](#)

5: 108044 (score = 0.9294) [New query](#)

6: 108051 (score = 0.9294) [New query](#)


7: 108004 (score = 0.9174) [New query](#)

8: 208042 (score = 0.9109) [New query](#)

Retrieval using spatial relationships




image



abstract regions

- Build graph using regions and their spatial relationships.
- Similarity is computed using graph matching.



Combining multiple features

Text query
on
"rose"




Combining multiple features

Visual query
on



Combining multiple features


Text query
on
"rose"
and visual query
on




Video Google: object matching



Video Google



Viewpoint invariant
descriptors



Visual vocabulary

Video Google

Document 1

Now is the time
for all good men
to come to the aid
of their country.

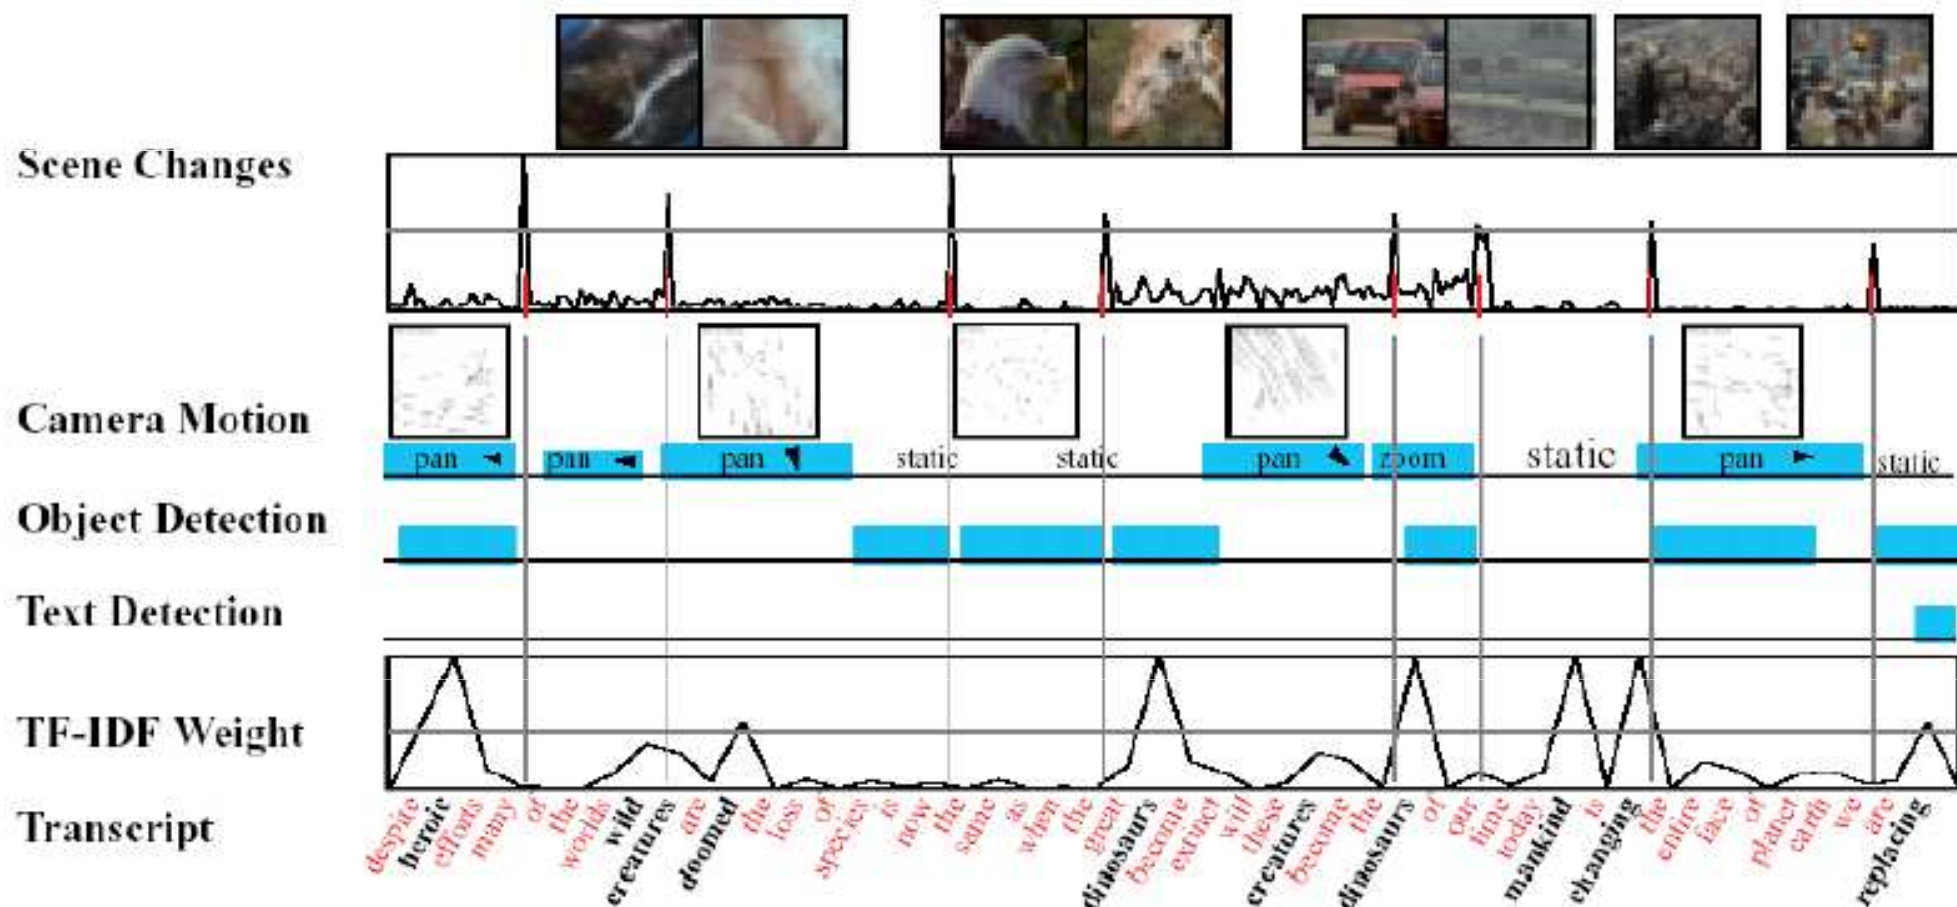
Document 2

Summer has come
and passed. The innocent
can never last.

Inverted index

Word	Document
aid	1
all	1
and	2
can	2
come	1, 2
country	1
for	1
good	1
...	...
the	1, 1, 2
...	...

Video skimming

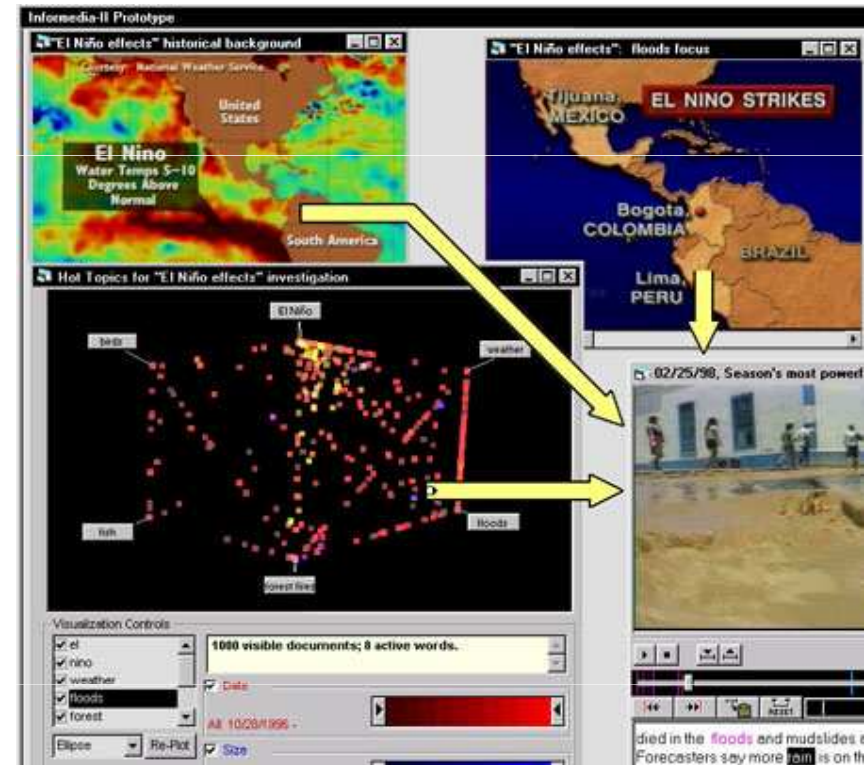


Event detection, indexing, retrieval

Assign semantic labels to significant events in video:

- Explosion, car crash, door slam (audio/video track)
- Marilyn Monroe enters scene
- Pele scores goal
- Jay Leno tells joke and then delivers punch-line
- Two people exchange a briefcase in park
- etc.

Informedia Digital Video Library



IDVL interface returned for "El Nino" query along with different multimedia abstractions from certain documents.

Informedia Digital Video Library





IDVL interface returned for "bin ladin" query.

The results can be tuned using many classifiers.

Relevance feedback

- In real interactive CBIR systems, the user should be allowed to interact with the system to “refine” the results of a query until he/she is satisfied.



Relevance feedback

- Example methods:
 - Query point movement
 - Query point is moved toward positive examples and moved away from negative examples.
 - Weighting features
 - The CBIR system should automatically adjust the weight that were given by the user for the relevance of previously retrieved documents.
 - Weighting similarity measures
 - Feature density estimation
 - Probabilistic relevance feedback

Relevance feedback

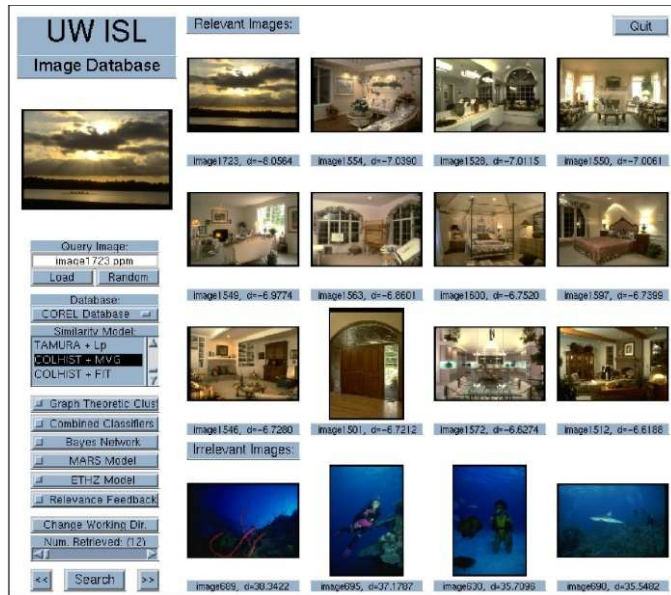
■ Positive feedback

$$\begin{aligned} p(A \mid \xi_{(0)}, \xi_{(1)}^+, \dots, \xi_{(n)}^+) &\propto p(\xi_{(n)}^+ \mid A) p(A \mid \xi_{(0)}, \xi_{(1)}^+, \dots, \xi_{(n-1)}^+) \\ p(B \mid \xi_{(0)}, \xi_{(1)}^+, \dots, \xi_{(n)}^+) &\propto p(\xi_{(n)}^+ \mid B) p(B \mid \xi_{(0)}, \xi_{(1)}^+, \dots, \xi_{(n-1)}^+) \end{aligned}$$

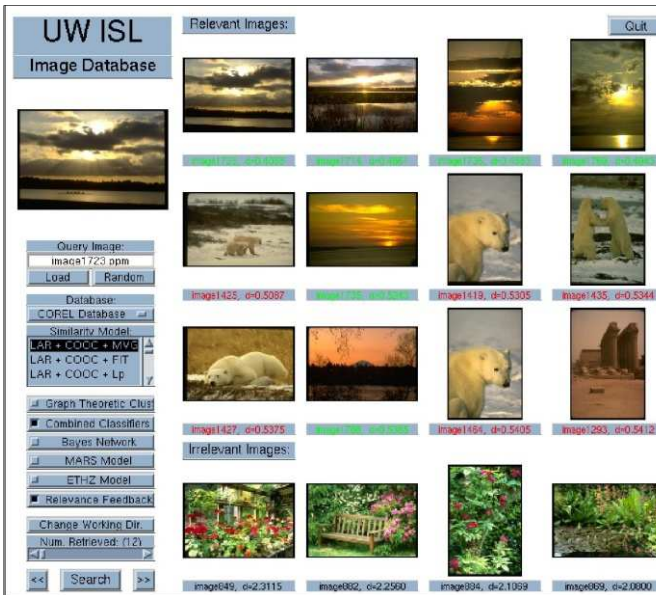
■ Negative feedback

$$\begin{aligned} p(A \mid \xi_{(0)}, \xi_{(1)}^+, \dots, \xi_{(n)}^+, \xi_{(1)}^-, \dots, \xi_{(m)}^-) &\propto \\ &p(\xi_{(m)}^- \mid B) p(A \mid \xi_{(0)}, \xi_{(1)}^+, \dots, \xi_{(n)}^+, \xi_{(1)}^-, \dots, \xi_{(m-1)}^-) \\ p(B \mid \xi_{(0)}, \xi_{(1)}^+, \dots, \xi_{(n)}^+, \xi_{(1)}^-, \dots, \xi_{(m)}^-) &\propto \\ &p(\xi_{(m)}^- \mid A) p(B \mid \xi_{(0)}, \xi_{(1)}^+, \dots, \xi_{(n)}^+, \xi_{(1)}^-, \dots, \xi_{(m-1)}^-) \end{aligned}$$

Relevance feedback



"Sunsets" using color histograms
(1/12)

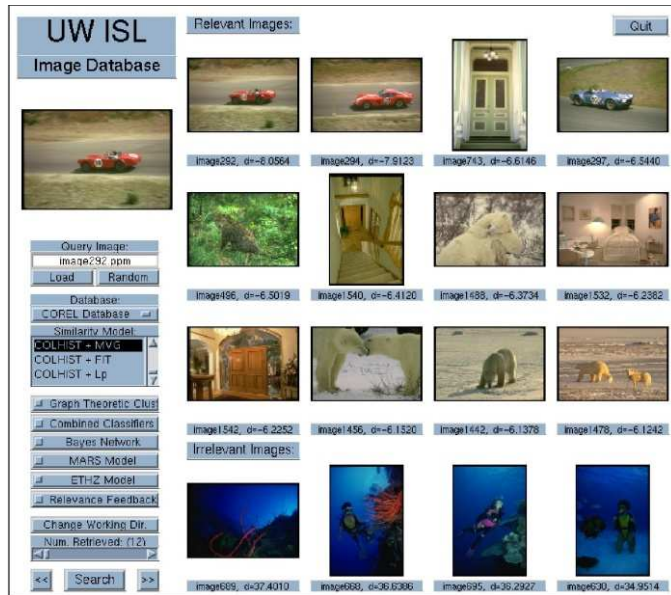


Using combined features (6/12)

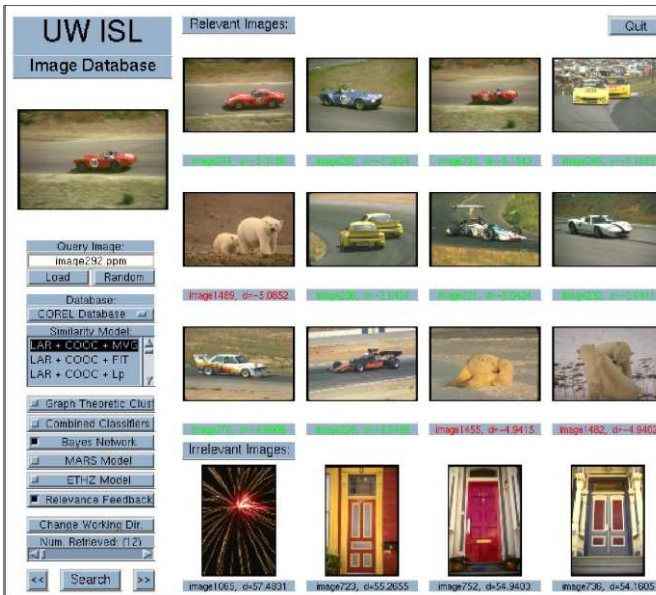


After 1st feedback (12/12)

Relevance feedback



"Auto racing" using color histograms (3/12)




Using combined features (9/12)



After 1st feedback (12/12)

Indexing for fast retrieval

- Use of **key images** and the **triangle inequality** for efficient retrieval.



Indexing for fast retrieval

- **Offline**


1. Choose a small set of key images.
2. Store distances from database images to keys.

- **Online** (given query Q)

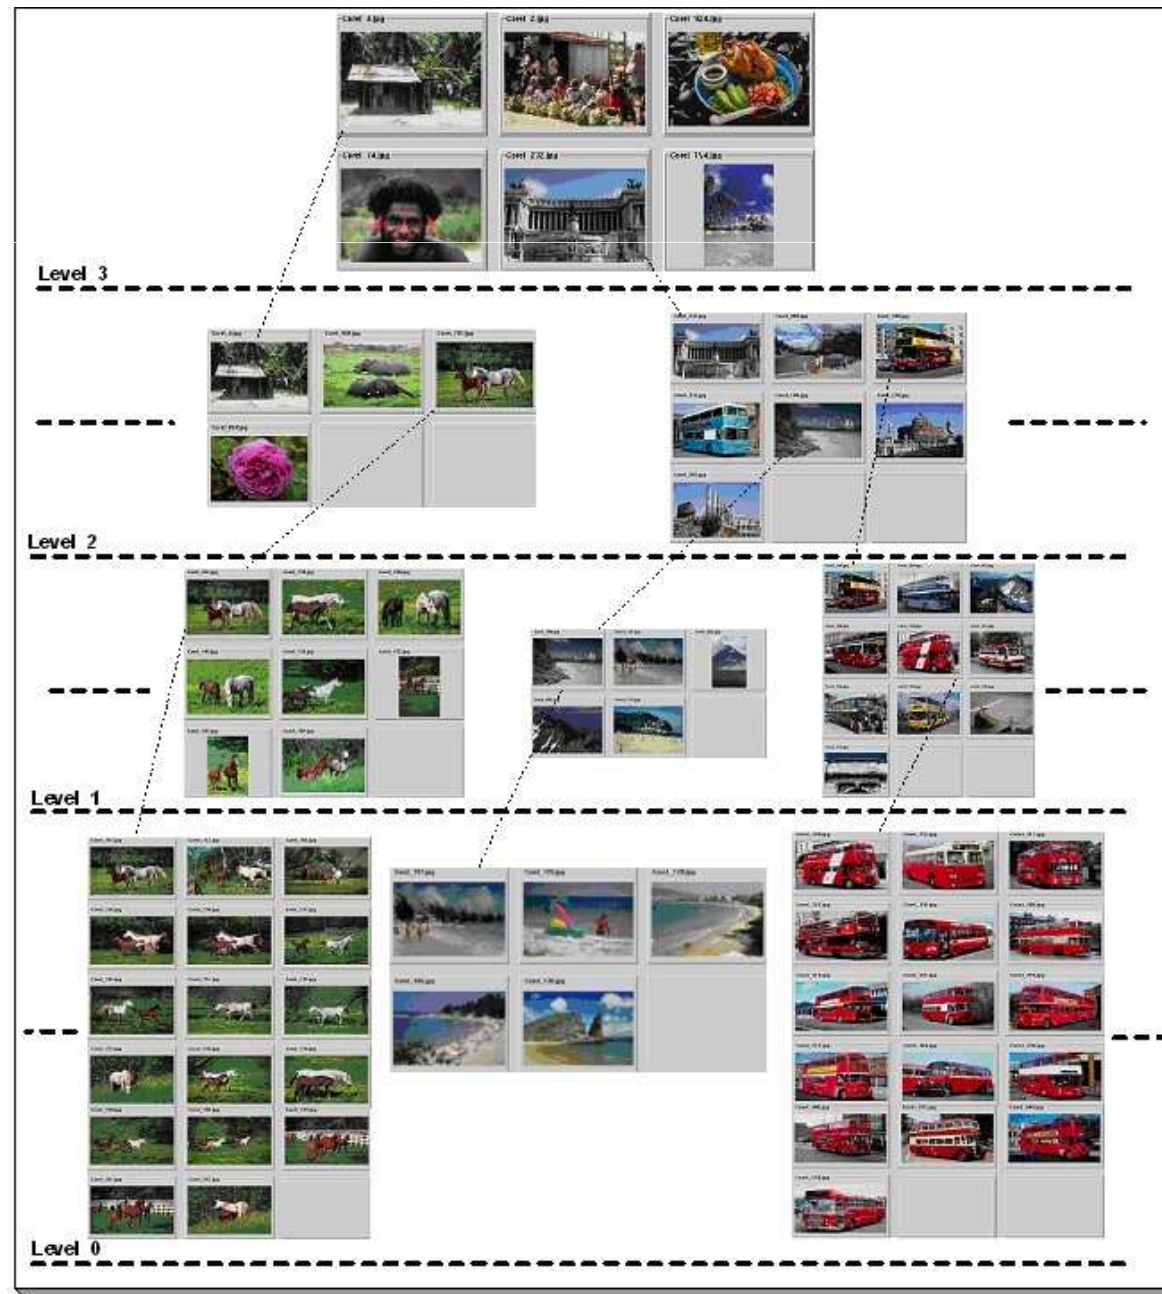
1. Compute the distance from Q to each key.
2. Obtain lower bounds on distances to database images.
3. Threshold or return all images in order of lower bounds.

Indexing for fast retrieval

- Hierarchical cellular tree



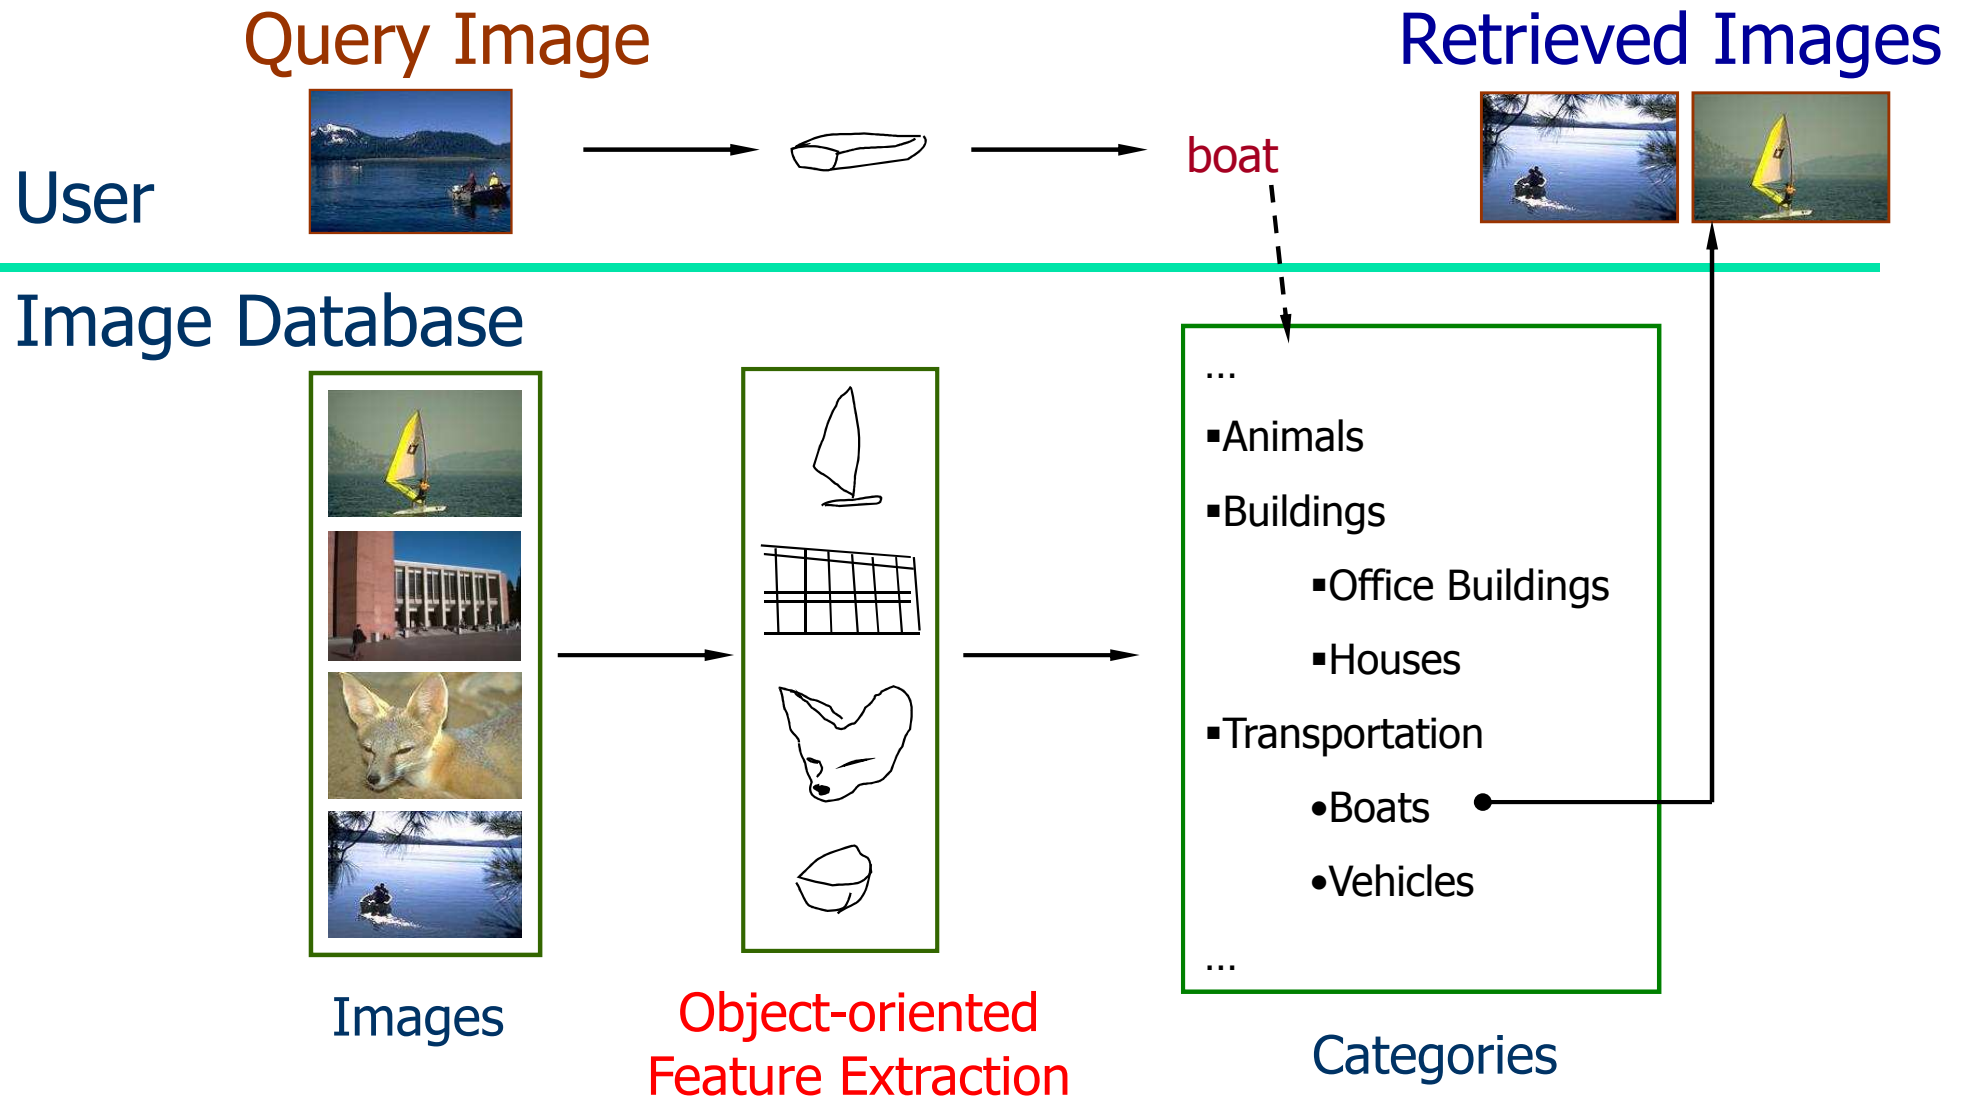
Indexing for fast retrieval



Performance evaluation

- Two traditional measures for retrieval performance in the information retrieval literature are precision and recall.
- Given a particular number of images retrieved,
 - **precision** is defined as the percentage of retrieved images that are actually relevant, and
 - **recall** is defined as the percentage of relevant images that are retrieved.

Current research objective



Demos

- Blobworld (<http://elib.cs.berkeley.edu/blobworld/>)
- Video Google (<http://www.robots.ox.ac.uk/~vgg/research/vgoogle/index.html>)
- FIDS (<http://www.cs.washington.edu/research/imagedatabase/demo/fids/>)
- Like Visual Shopping (<http://www.like.com/>)
- Google Image Search (<http://images.google.com/>)
- Yahoo Image Search (<http://images.search.yahoo.com/>)
- Flickr (<http://flickr.com/>)
- The ESP game (<http://www.espgame.org/>)

Demos

- Vitalas

- <http://vitalas.ercim.org/>

- Google Similar Images

- <http://similar-images.googlelabs.com/>

- Google Image Swirl

- <http://image-swirl.googlelabs.com/>

- Microsoft Bing

- <http://www.bing.com/>

First use keywords, then mouse over an image and click on show similar images