Relevance Feedback

Tutorial Image Retrieval

Thomas Deselaers, Henning Müller

Outline

- What is relevance feedback
- Interfaces
- Instance-based relevance feedback
- Model-based relevance feedback
- Probabilistic relevance feedback

What is Relevance Feedback

- After an initial query, the user is presented with a set of results
- Then, the user gives feedback to the system
  - Marks images as relevant and/or
  - Marks images as irrelevant
- Relevance feedback can be given explicitly or implicitly
  - Explicit: marking of several images
  - Implicit: clicking on one image/stopping to search on

Positive and Negative Feedback

- Studies on strategies for relevance feedback
  - Positive feedback often a reordering of top results or one new query with a single image
    - Images have already much in common
  - Negative feedback is often the key to good results
    - Really new images are retrieved
    - Much more information is supplied
  - Problem with two much negative feedback!
    - Images with small number of features are returned
  - As much feedback as possible usually delivers best results

Problem with Negative Feedback

- Problem with too much negative feedback also in text retrieval
- Solution: Separately weighting positive and negative parts of feedback
  - Often positive=0.65, negative=0.35
- Compare Rocchio's Method [later]

Interfaces

- Simple interface

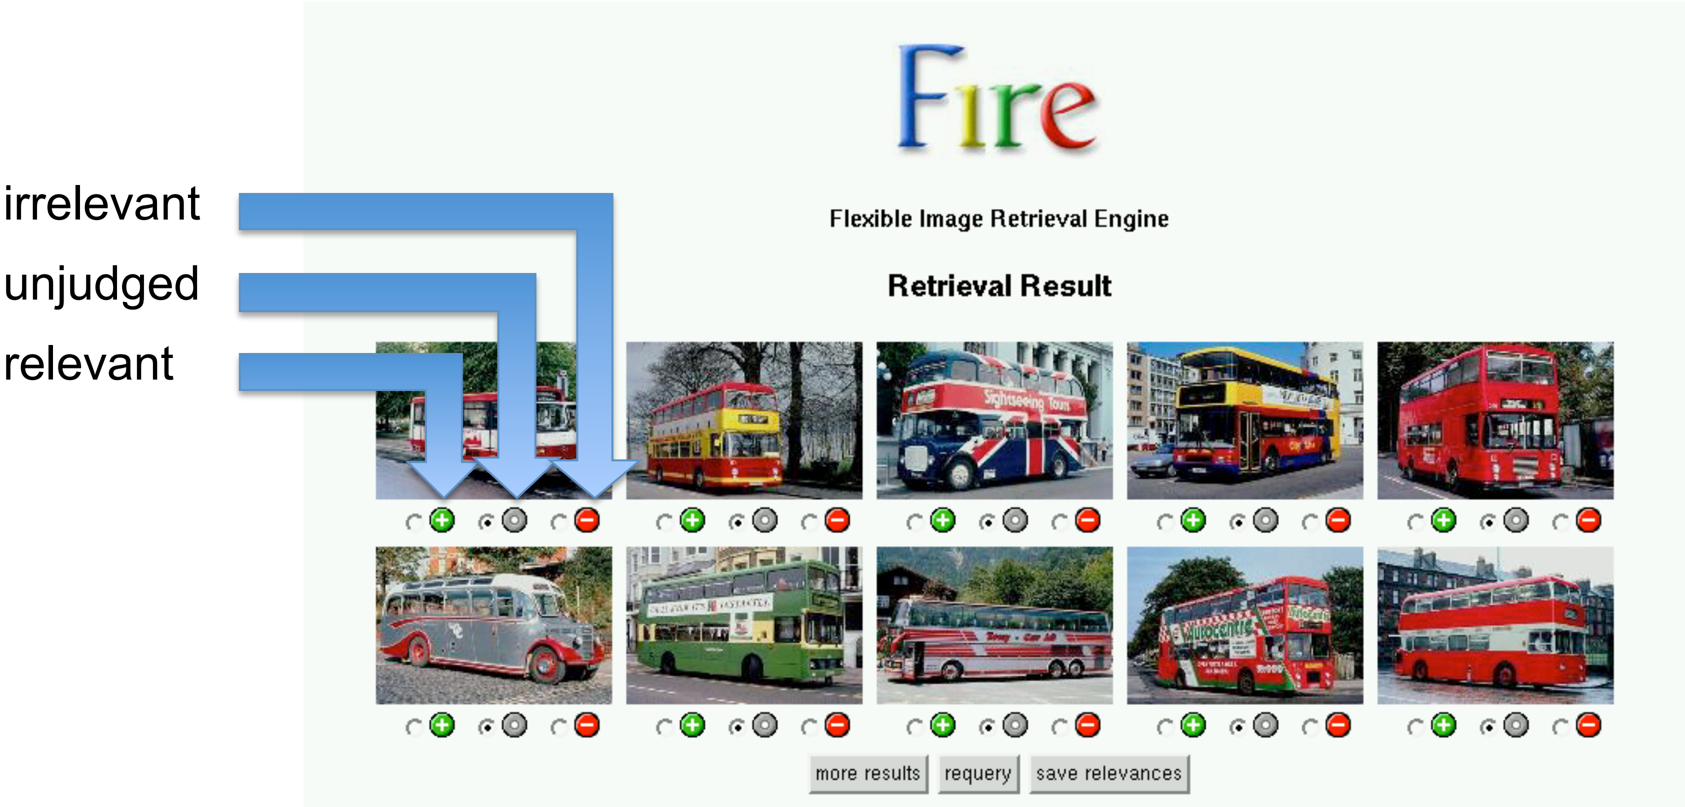

Interfaces: Tree-based Query Interface

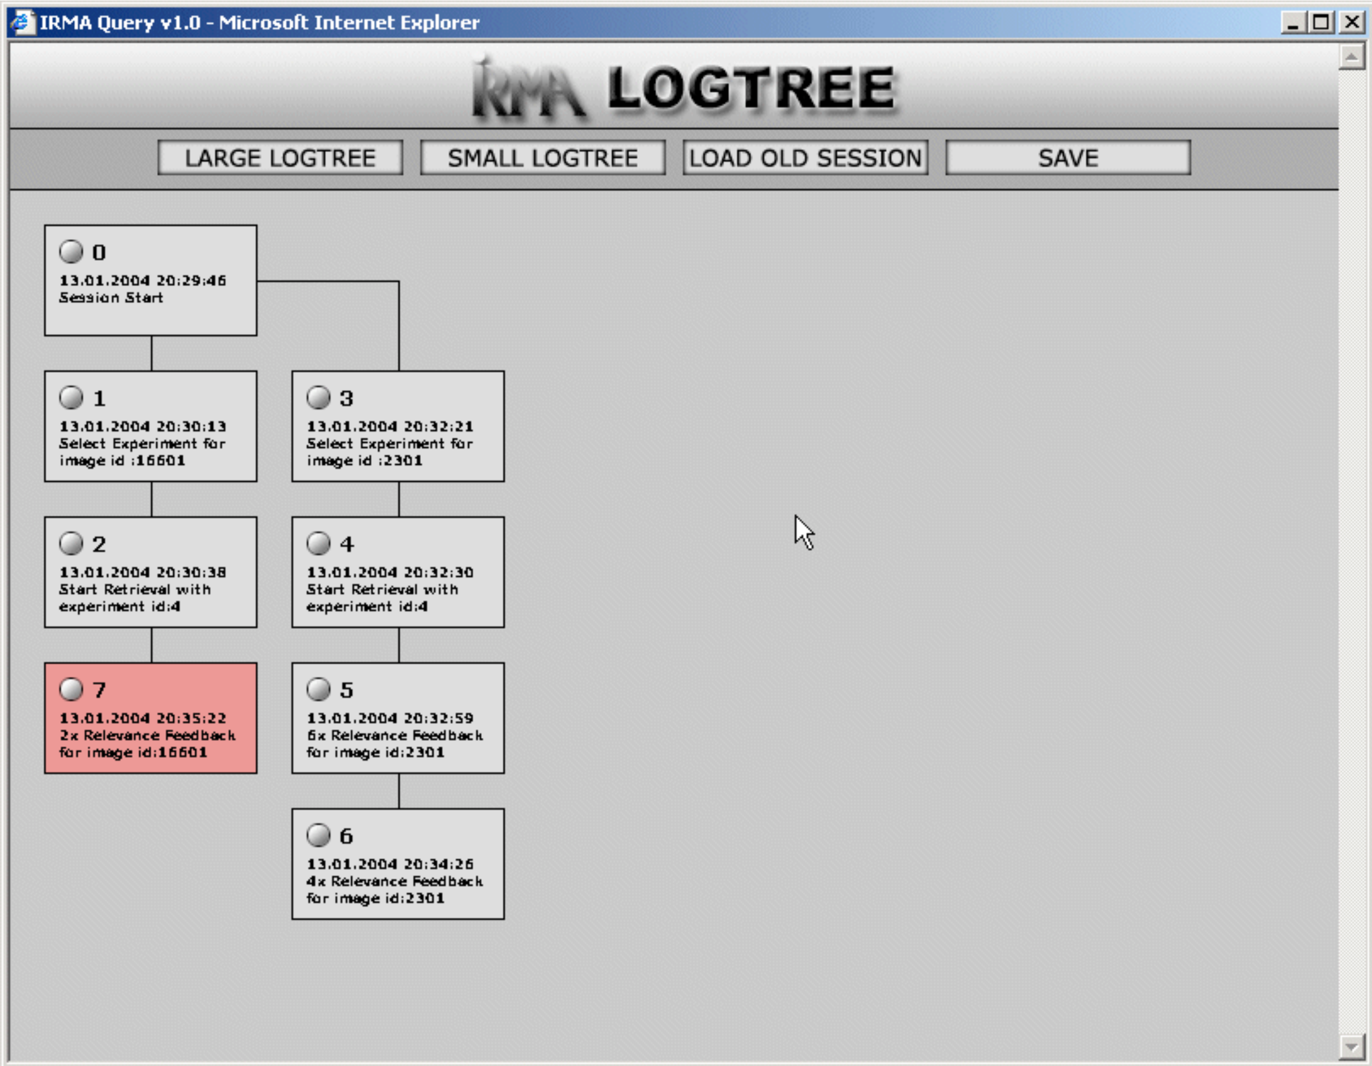

This is the IRMA (Image Retrieval in Medical Applications System from RWTH University Hospital

Interfaces

- 3D Browsing and Searching Interface in MARS

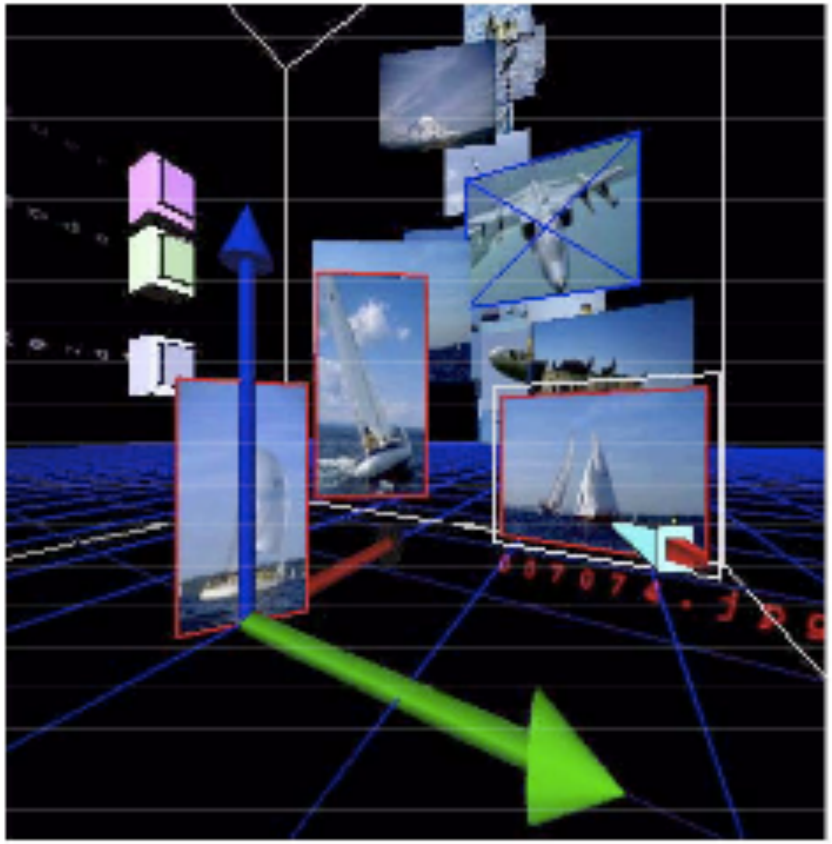

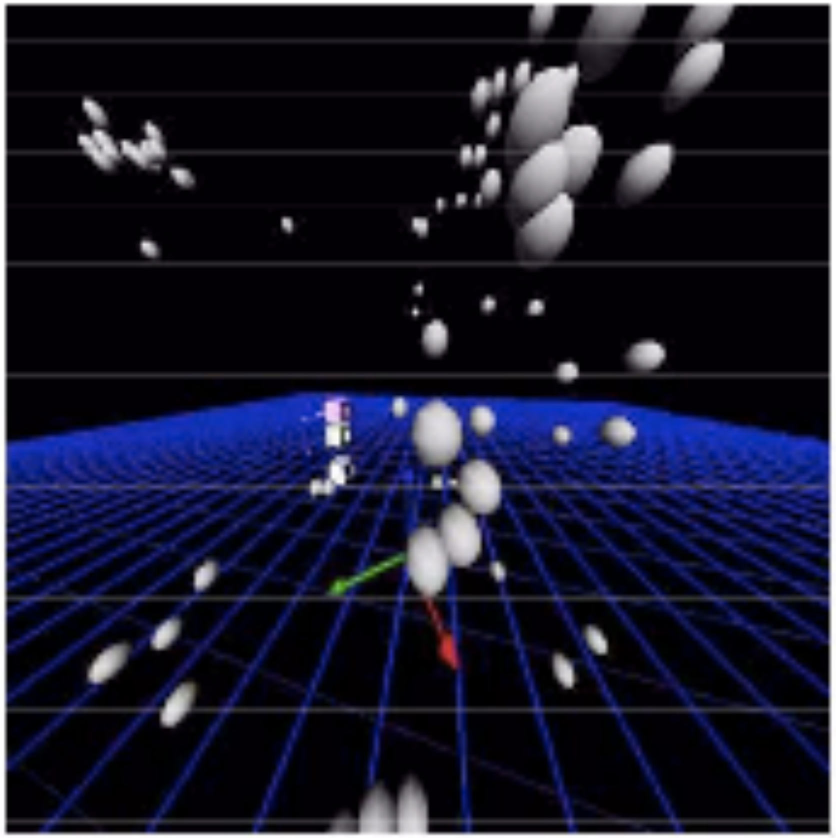

Interfaces

- Video Search and Retrieval Interface

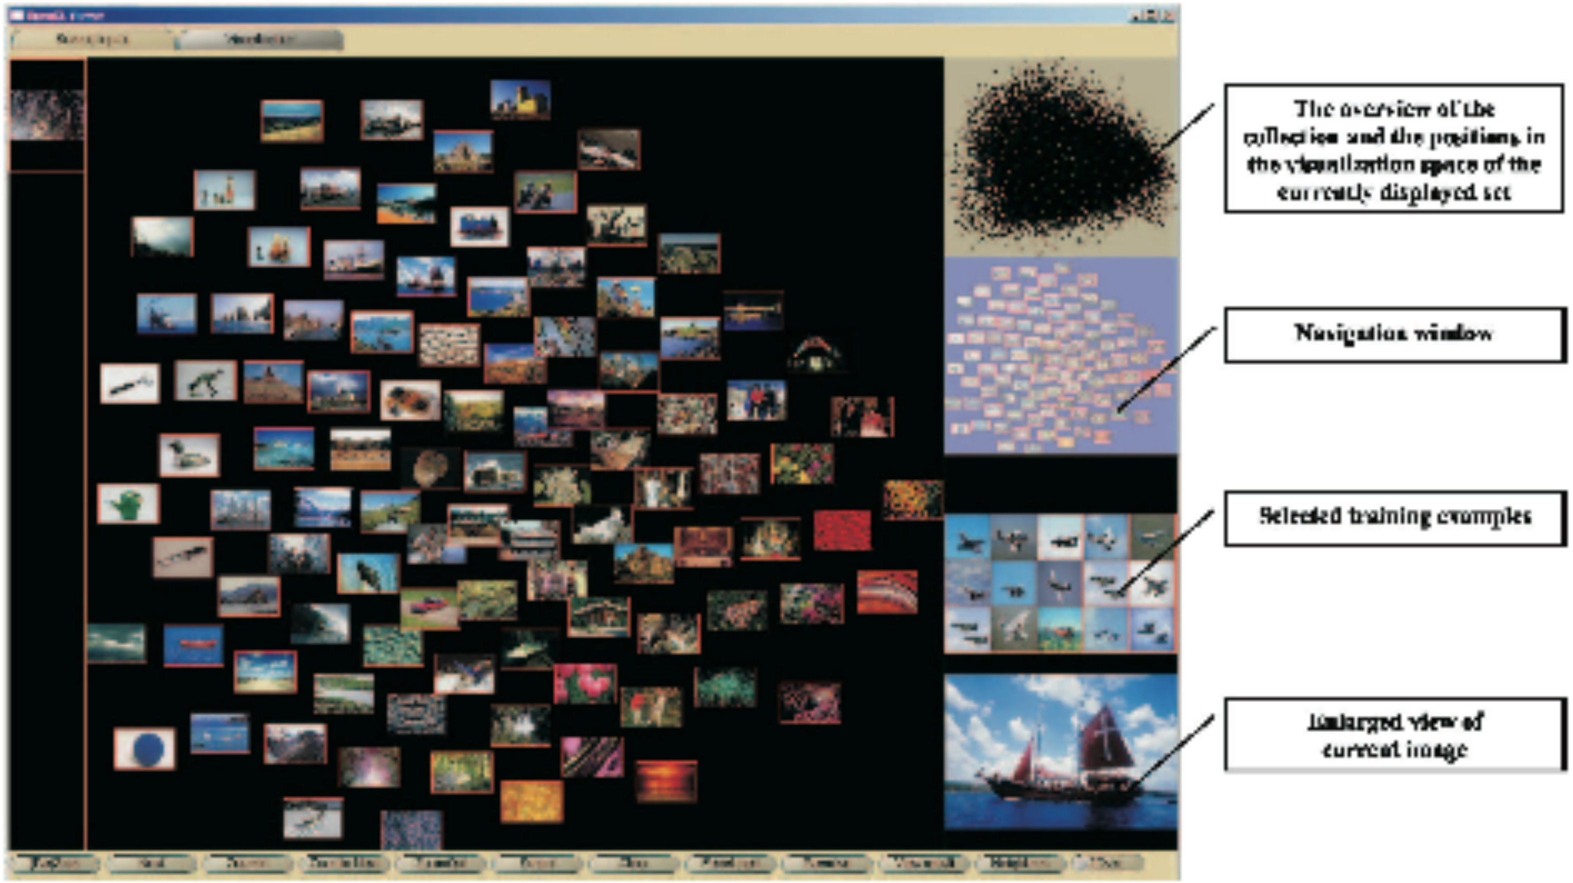

Nguyen et al. ACM Trans. MM 2008

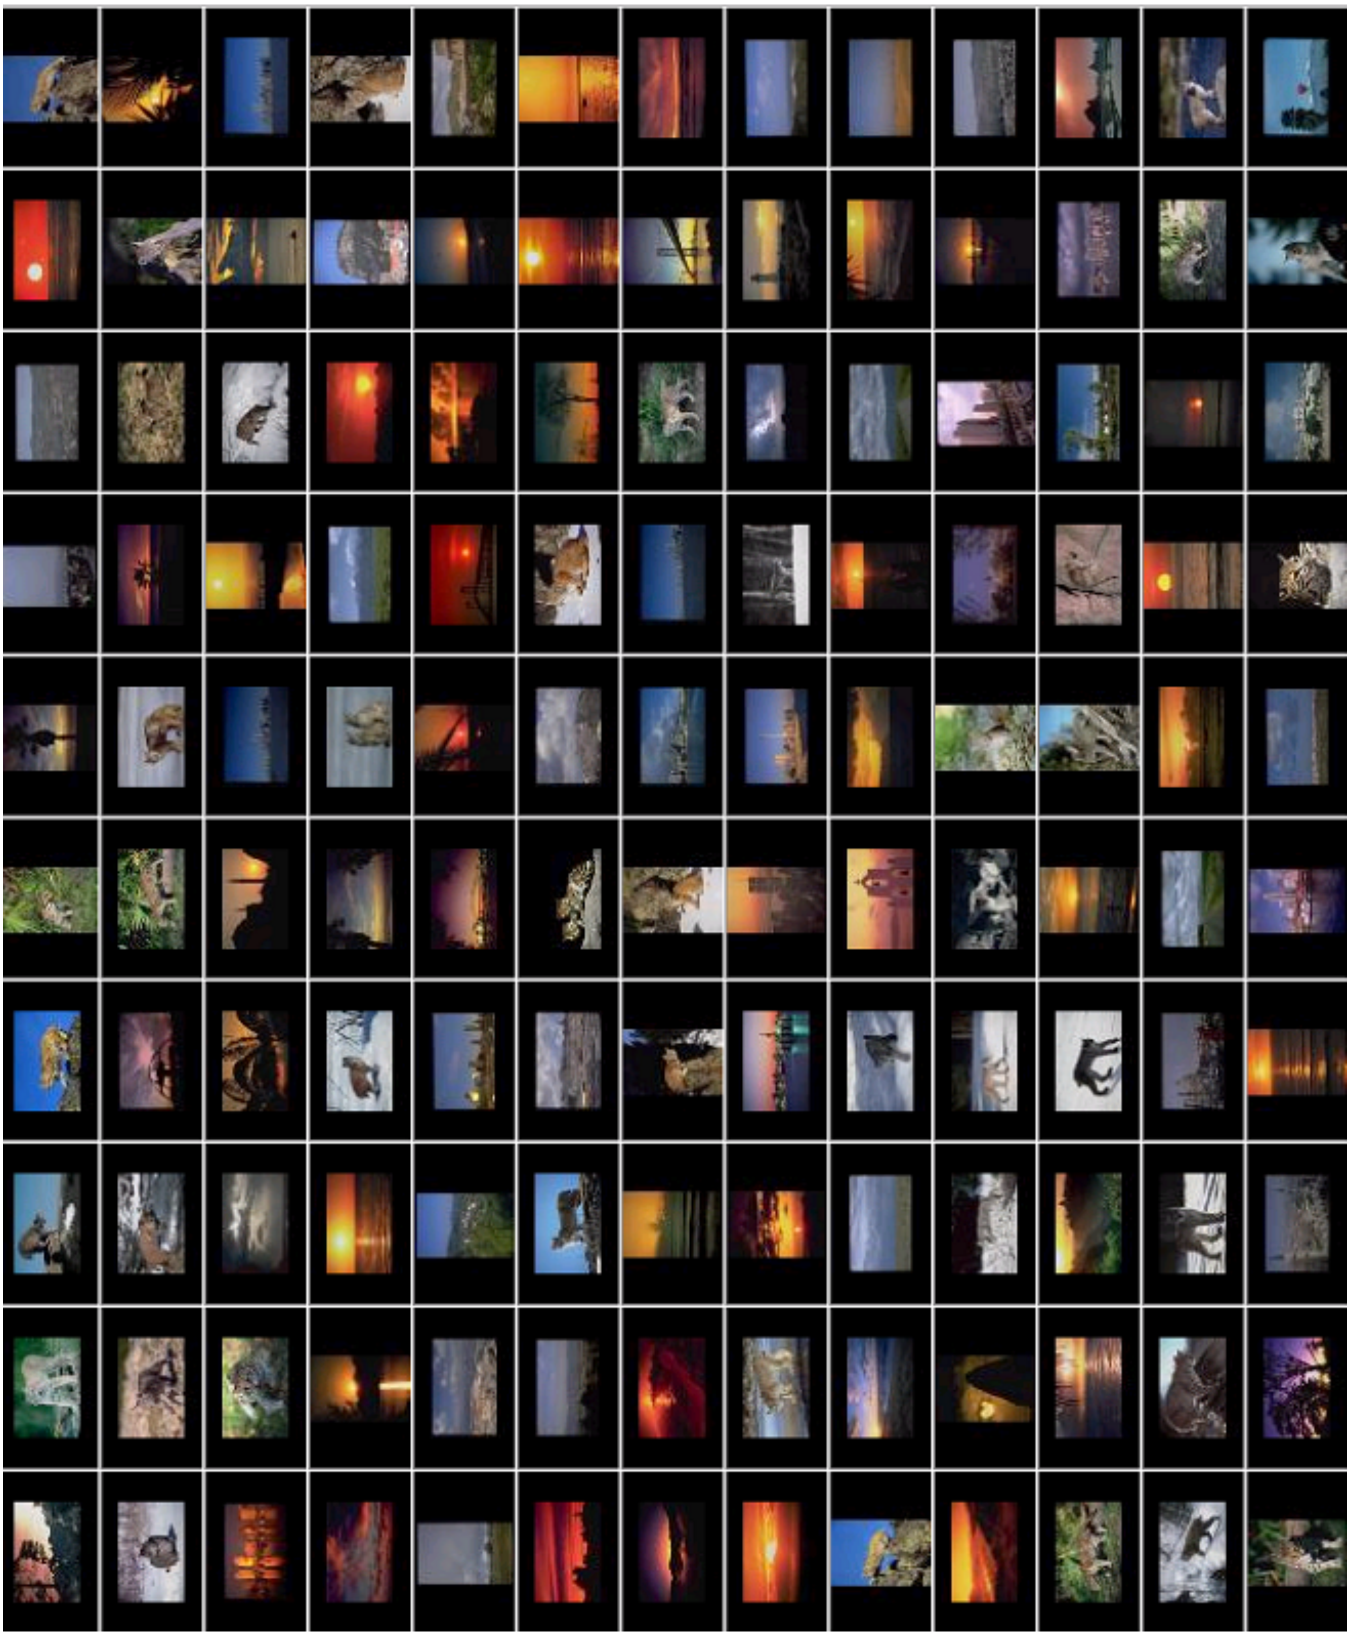

Collection guide

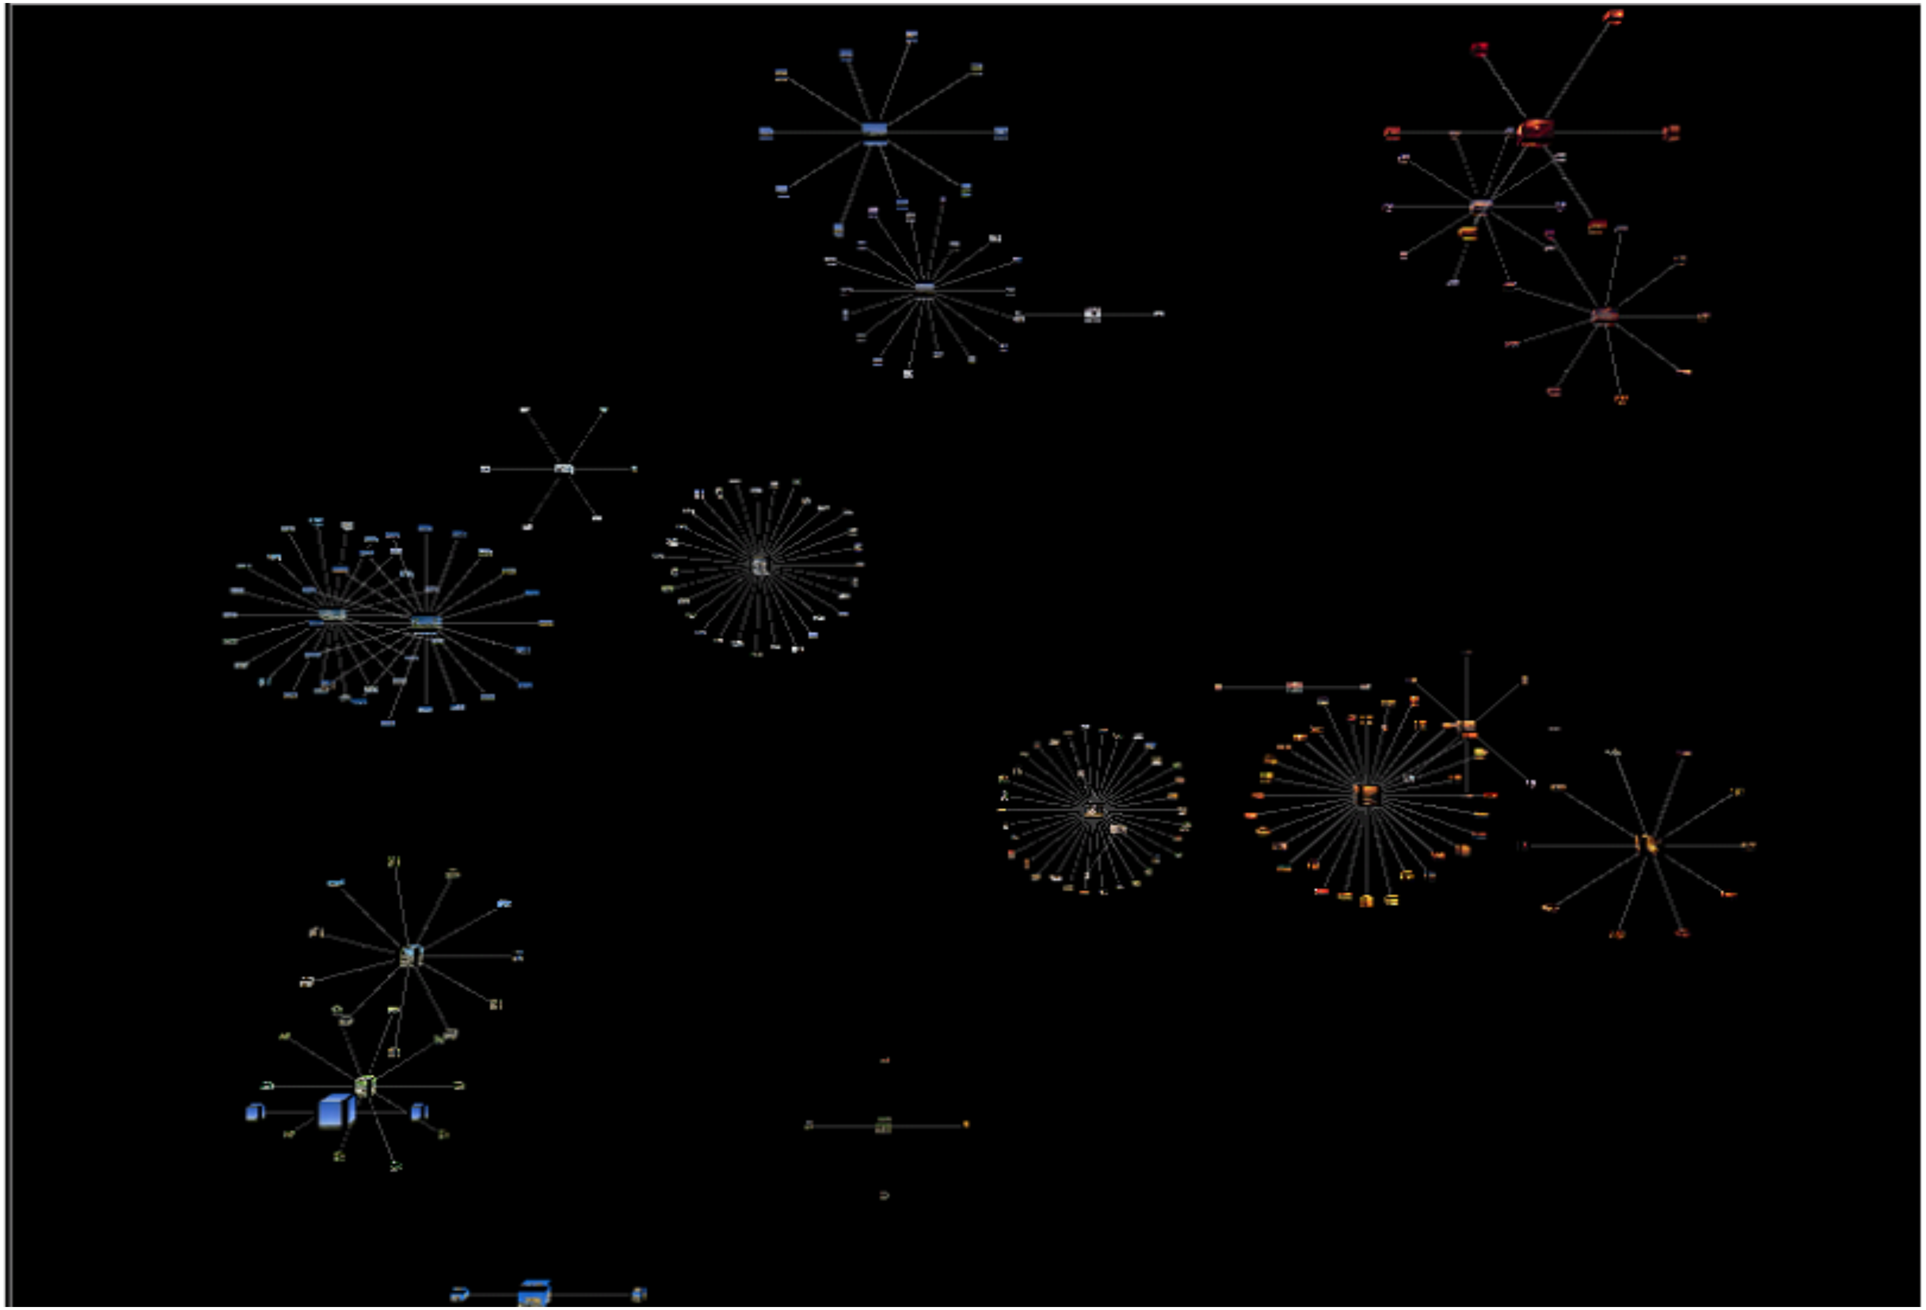

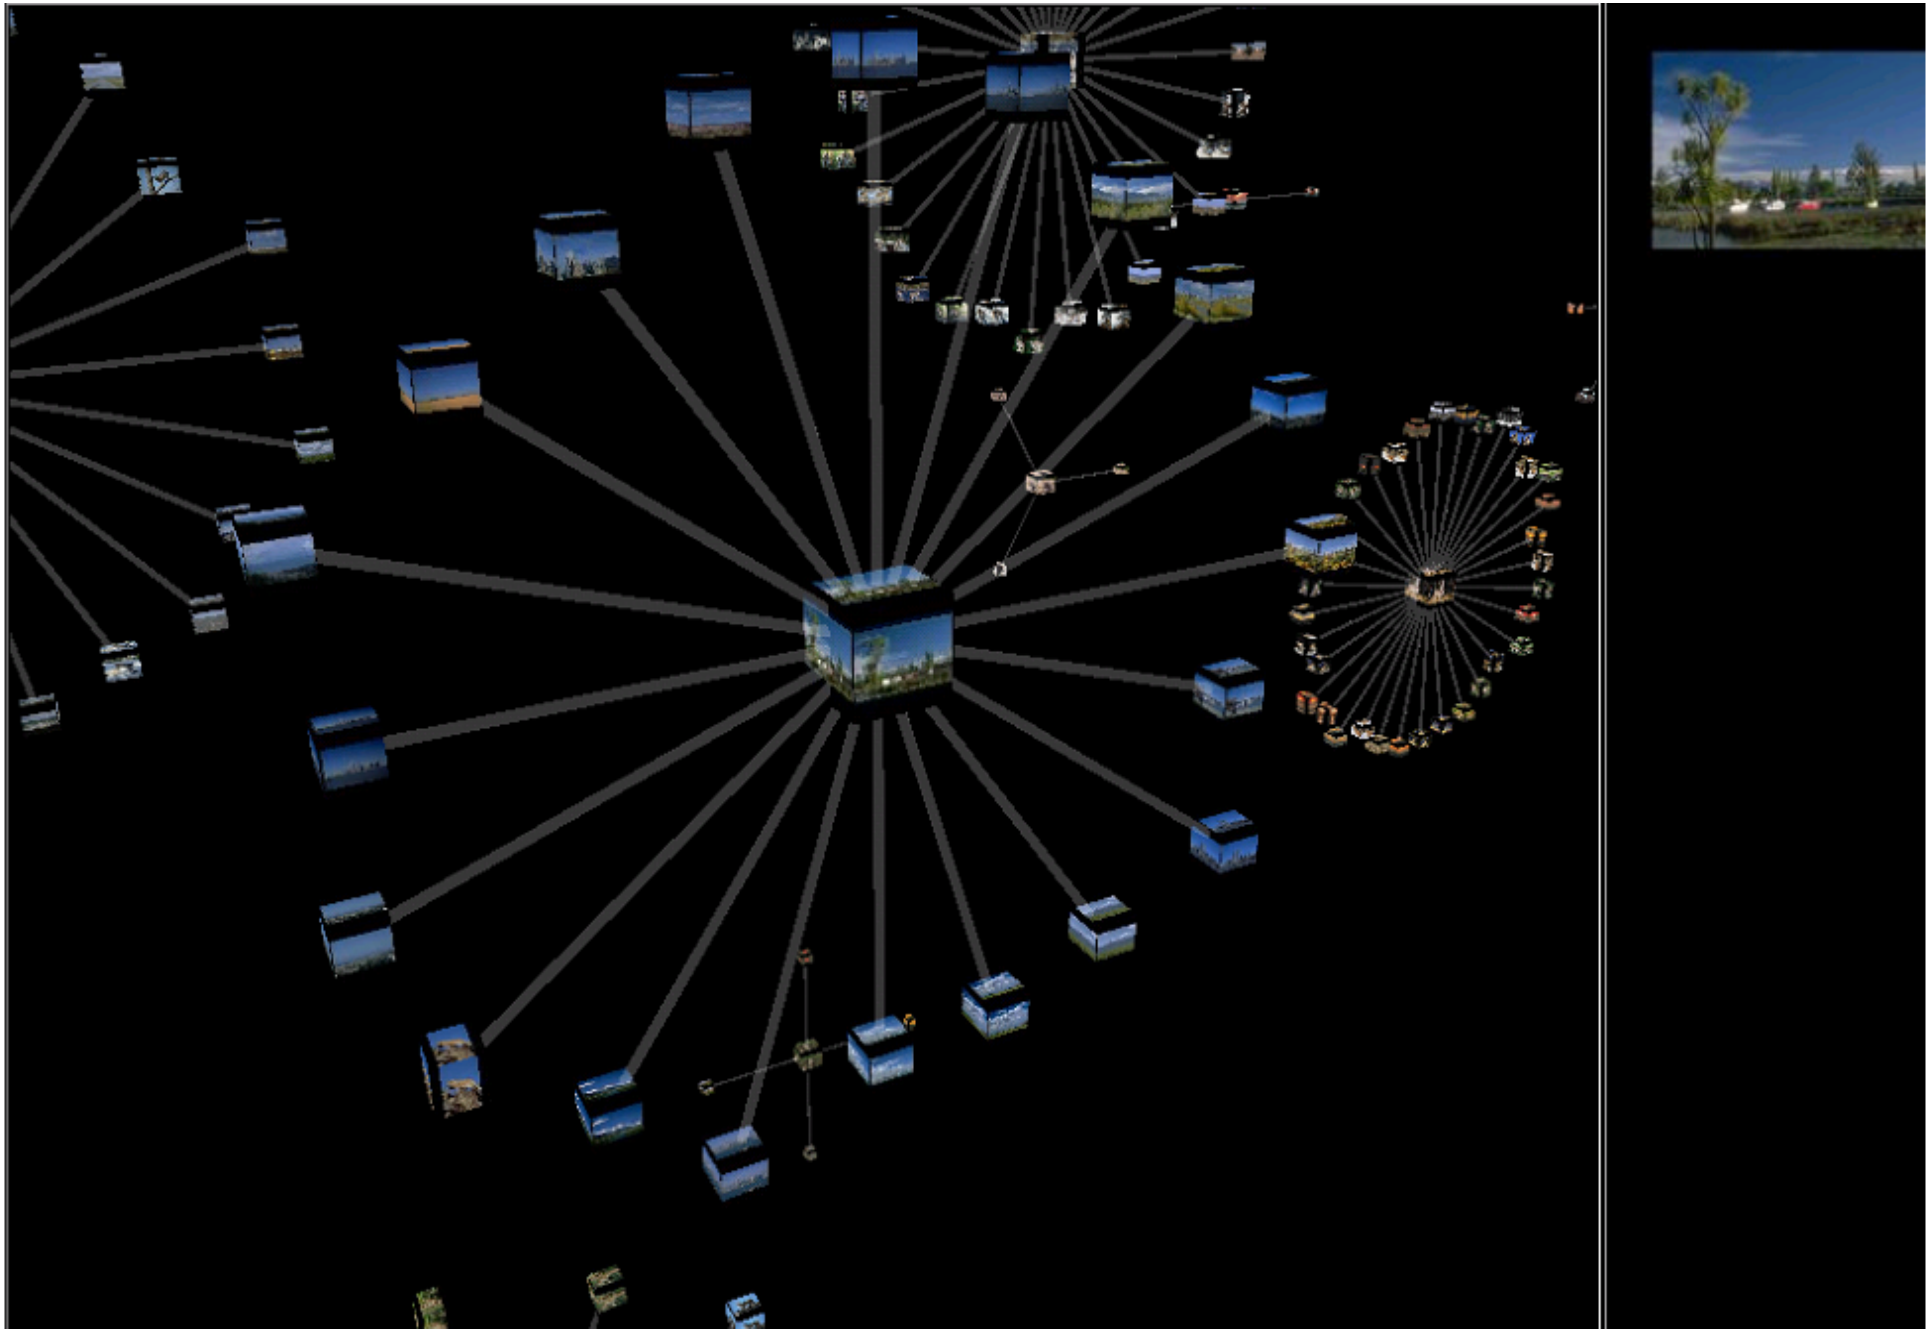

Combination Schemes

- Given a set of positive and a set of negative examples (possibly empty)
- How can we use this information?

- In the following: different schemes to fuse these information cues

Joint work with R. Paredes, UPV

Notation

- x image in the database
- q query image
- Q^+ set of positively marked images
- Q^- set of negatively marked images
- $p_x(r|q)$ probability that image x is relevant given query q
- $p_x(\bar{r}|q)$ probability that image x is irrelevant given query q

RF as Combination of Classifiers

- Consider relevance feedback images as training data for a kernel density classifier
- An image is relevant with probability

$$p_x(r|q) \quad \propto \exp\left(-d(x, q_+)\right)$$

- Analog: irrelevant

$$p_x(\overline{r}|q_-) \quad \propto \exp\left(-d(x, q_-)\right)$$

- Combine these classifiers by averaging (bagging)

$$p_x(r|(Q^+, Q^-)) =$$

$$\frac{\alpha}{|Q^+|} \sum_{q_+ \in Q^+} p_x(r|q_+) + \frac{1-\alpha}{|Q^-|} \sum_{q_- \in Q^-} \left(1 - p_x(\overline{r}|q_-)\right)$$

Default method in FIRE

Relevance Score

- Giacinto et al. propose an instance-based relevance feedback mechanism
- Relevance Score:

$$RS(x, (Q^+, Q^-)) \quad = \left(1 + \frac{\min\limits_{q_+ \in Q^+} d(x, q_+)}{\min\limits_{q_- \in Q^-} d(x, q_-)}\right)^{-1}$$

- Advantages: supports inhomogeneous sets Q^+ and Q^-

Rocchio's Method

- Proposed for textual information retrieval by Rocchio in 1972
- Reformulate the query by going
  - Into the direction of the positive feedback Q^+
  - Away from the negative feedback Q^-
- Advantage: very fast, only one query has to be performed

$$\hat{q} = q + \beta \left(\sum_{q_+ \in Q^+} q_+ \right) - \gamma \left(\sum_{q_- \in Q^-} q_- \right)$$

Default method in GIFT

Quotient of Sums

- Try to define a sound probabilistic model for relevance feedback
- Determine probability for an image to be relevant (according to Bayes’ decision rule)

$$p(r|x) = \frac{P(r)p(x|r)}{P(r)p(x|r) + P(\overline{r})p(x|\overline{r})}$$

- The different terms used here are defined in the following
- Inspired by RelevanceScore and FIRE’s method

Quotient of Sums

- Prior probabilities for images to be relevant or irrelevant:

$$P(r) = \frac{|Q^+|}{|Q^+ \cup Q^-|} \qquad P(\overline{r}) = \frac{|Q^-|}{|Q^+ \cup Q^-|}$$

- Emission probability for an image given relevant or irrelevant

$$p(x|r) \propto \frac{1}{|Q^+|} \sum_{q \in Q^+} 1/d(x,q)$$

$$p(x|\overline{r}) \propto \frac{1}{|Q^-|} \sum_{q \in Q^-} 1/d(x,q)$$

Quotient of Sums

- The likelihood for an image to be relevant then is

$$S_{(Q^+,Q^-)}(x) \quad = \quad \frac{\sum\limits_{q \in Q^+} 1/d(x,q)}{\sum\limits_{q \in \{Q^+ \cup Q^-\}} 1/d(x,q))}$$

- This is used to rank images in retrieval

Tuning the system

- Apart from combining the queries, they can be used to tune parameters of the system
  - Compare to the learning of feature combinations
- Here:
  - Refine image comparison measures

Learning Weighted Distances

- Relevance feedback allows to use machine learning techniques to tune parameters of a system
- One possibility is to tune the distance function
  - Add weight for each component of the vectors which are compared

$$d(x,q) \quad = \sum_{i=1}^{D} w_i |x_i - q_i|$$

Joint work with R. Paredes, UPV

Learning Weighted Distances

- Optimisation:
  - Minimise the distances from each relevant (positive) image to all other relevant images

$$\sum_{x \in Q^+} \sum_{q_+ \in Q^+ \setminus \{x\}} \sum_{q_- \in Q^- \setminus \{x\}} \frac{d(x, q_+)}{d(x, q_-)}$$

  - Analogously: Maximise the distances from each irrelevant (negative) image to all relevant images

Joint work with R. Paredes, UPV

Experimental Evaluation: WANG

WANG database: 1000 images, 10 classes, “easy task”

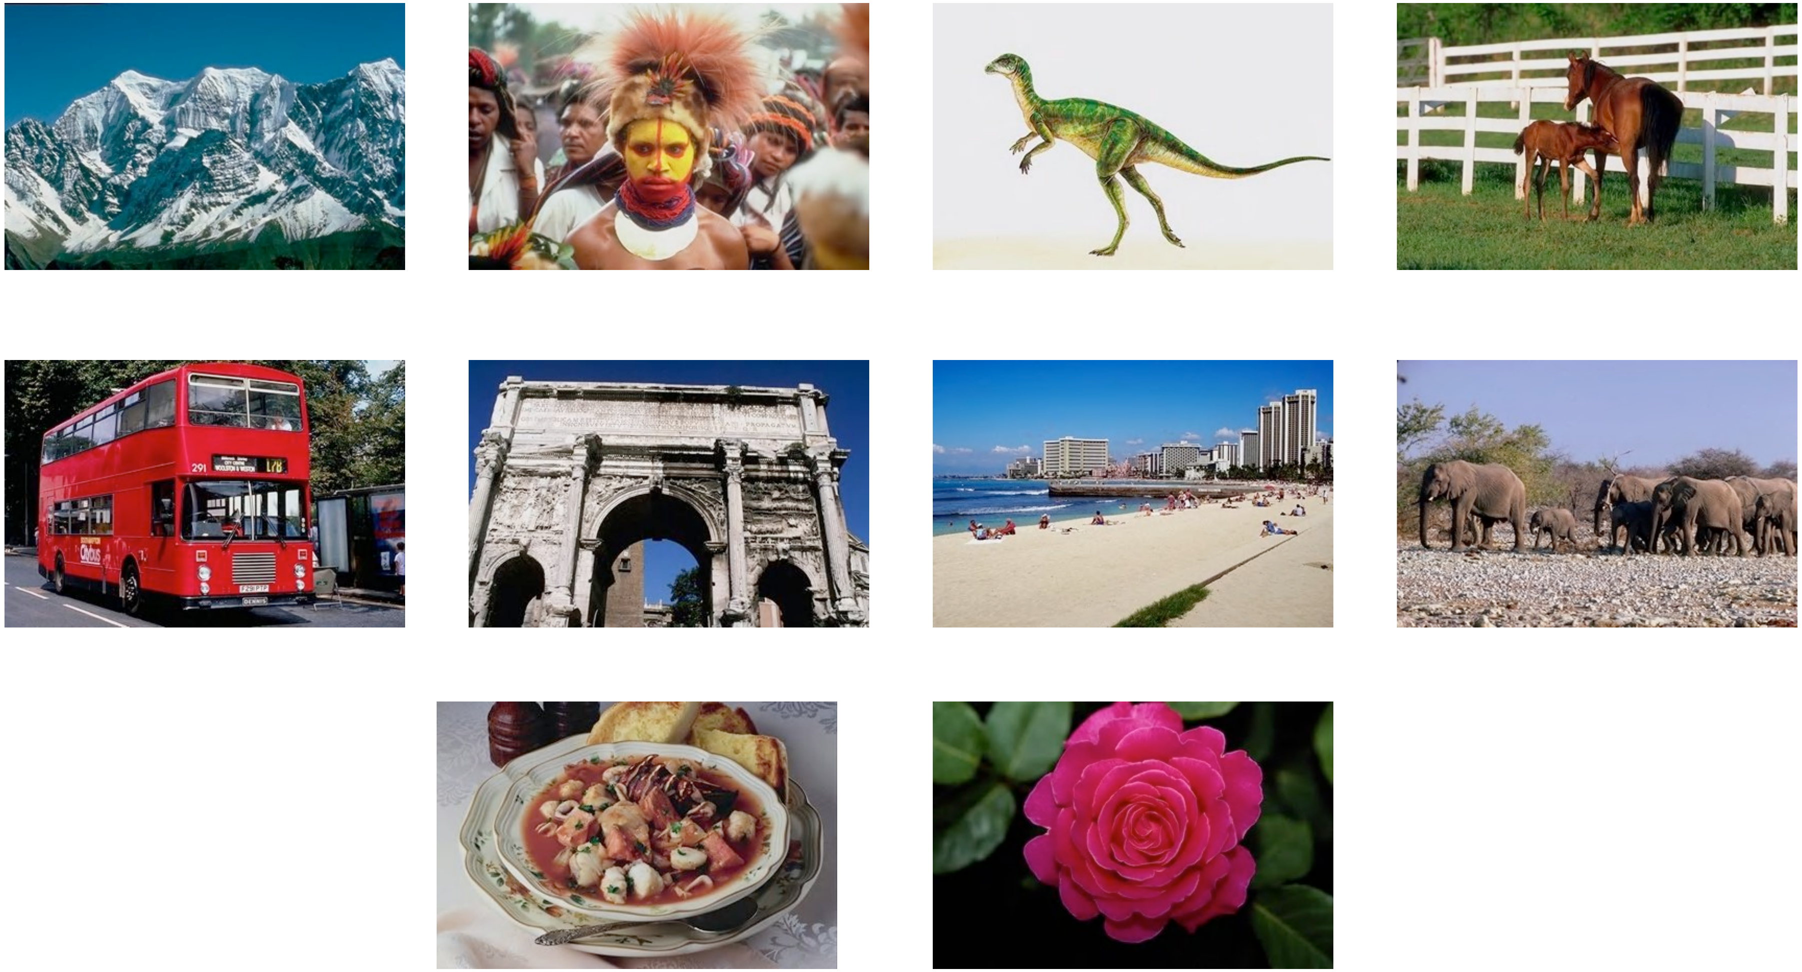

Evaluation on WANG Data

- P(20)

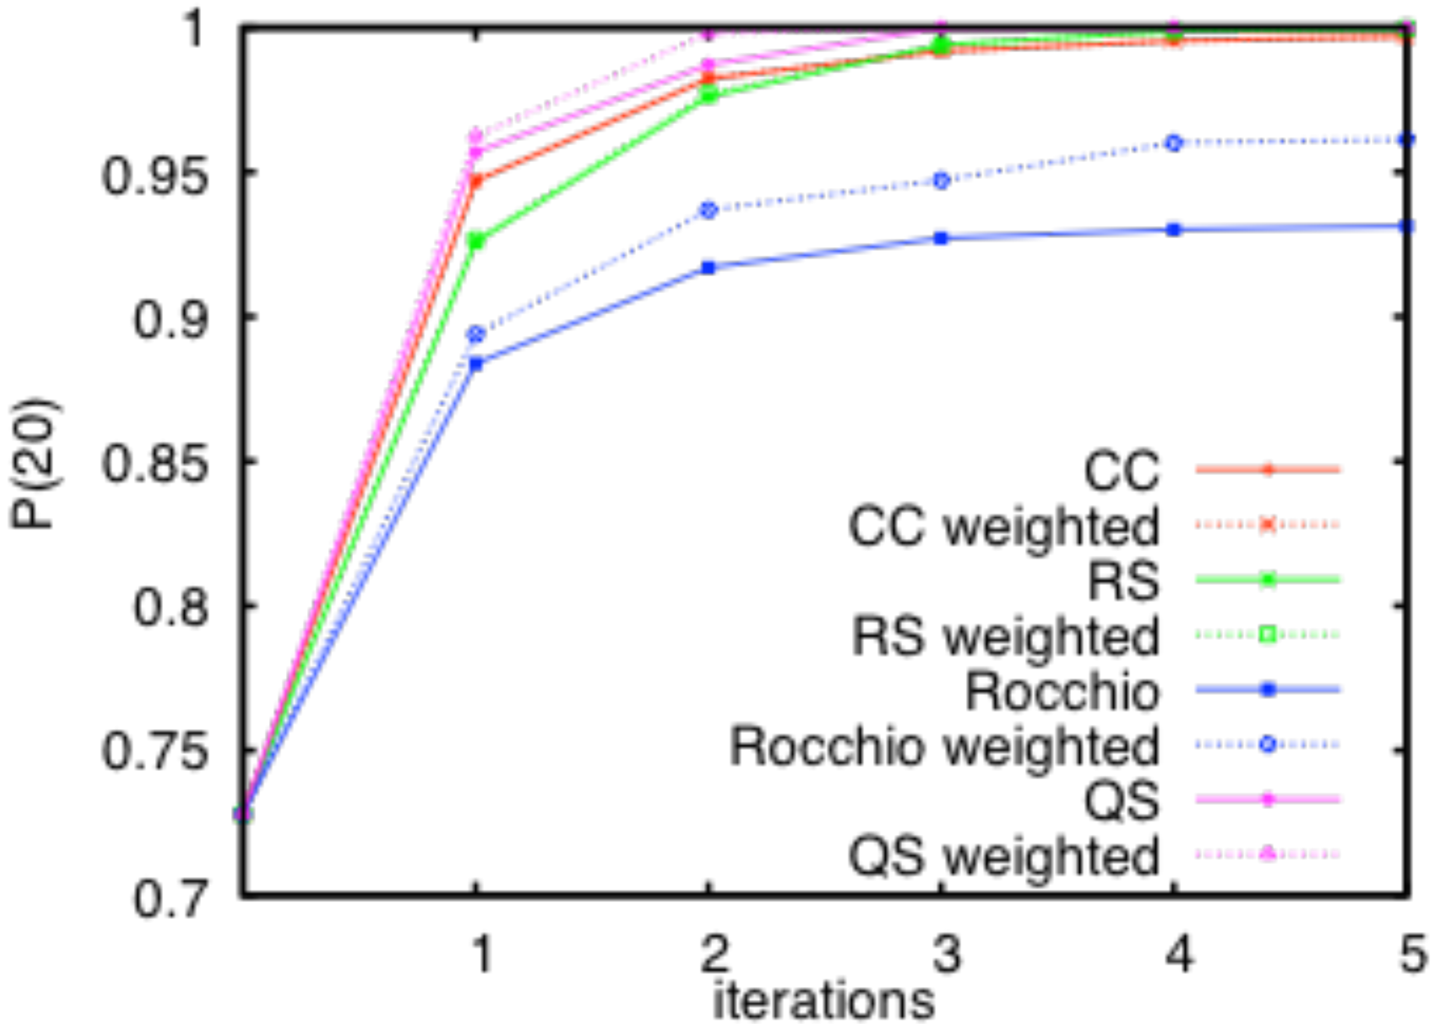

ImageCLEF Database

- 20,000 colour photographs
- Accompanied by semi-structured captions
    - English and Random
- Many images have similar visual content but varying
    - illumination
    - viewing angle
    - background
- Used in ImageCLEF in 2006 – 2008
- Publicliy available from www.imageclef.org









ImageCLEF Database: Image

```
<DOC>
<DOCNO>annotations/17/17405.eng</DOCNO>
<TITLE>Group photo with Machu Picchu and
Huayna Picchu in the background</TITLE>
<DESCRIPTION>tourists are sitting on a grey
gravel road in the foreground;
a ruin with grey walls and many green
terraces and a distinctive, rocky, steep
mountain behind it;
a wooden mountain range and white clouds in
the background; </DESCRIPTION>
<NOTES></NOTES>
<LOCATION>Machu Picchu, Peru</LOCATION>
<DATE>26 October 2004</DATE>
<IMAGE>images/17/17405.jpg</IMAGE>
<THUMBNAIL>thumbnails/17/17405.jpg</
THUMBNAIL>
</DOC>
```

ImageCLEF Database: Query

- 39 topics with full information
  - Based on realistic topics (log-file analysis and interviews)
- Available in English only
- Augmented by a cluster tag
  - defines how the rel. images should be clustered

```
<top>
<num> Number: 5 </num>
<title> animal swimming </title>
<cluster> animal </cluster>
<narr> Relevant images will show one or
more animals (fish, birds, reptiles, etc.)
swimming in a body of water. Images of
people swimming in water are not relevant.
Images of animals that are not swimming
are not not relevant. </narr>
<image> SampleImages/05/3739.jpg </image>
<image> SampleImages/05/4986.jpg </image>
<image> SampleImages/05/30823.jpg </image>
</top>
```

Sample topic images:



Evaluation on ImageCLEF 2007 task

- P(20)

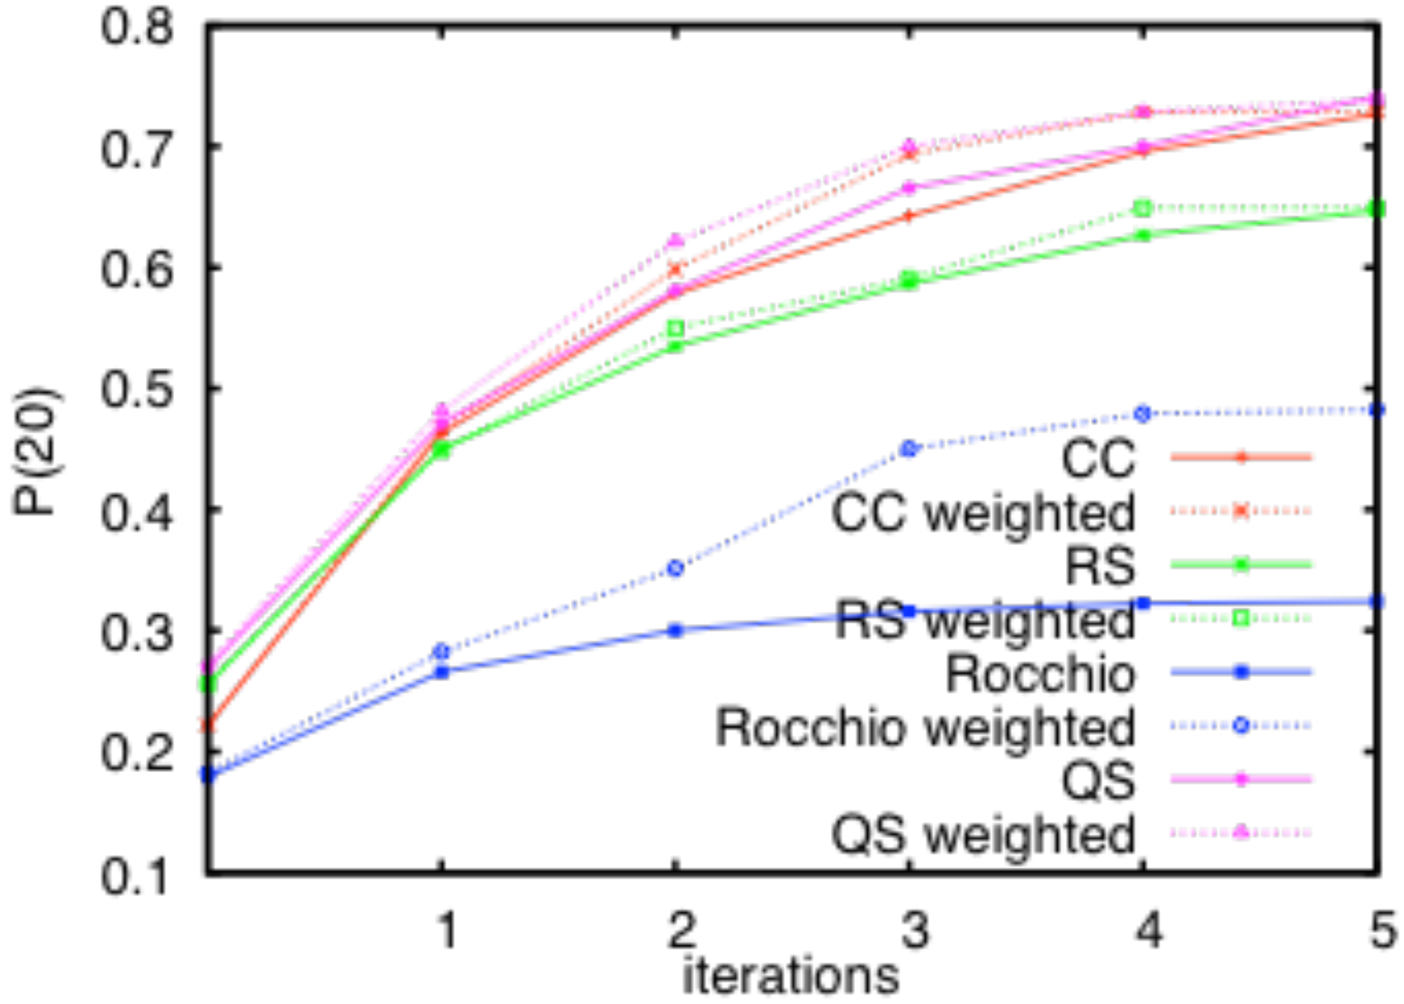

iCLEF: interactive multi-lingual image retrieval

- Simultaneous search in multiple languages
- User interaction
  - Have users participate
  - Competition
  - Researchers work on log-files
- Mainly text-based, visual extension would be interesting but difficult

iCLEF: user interface

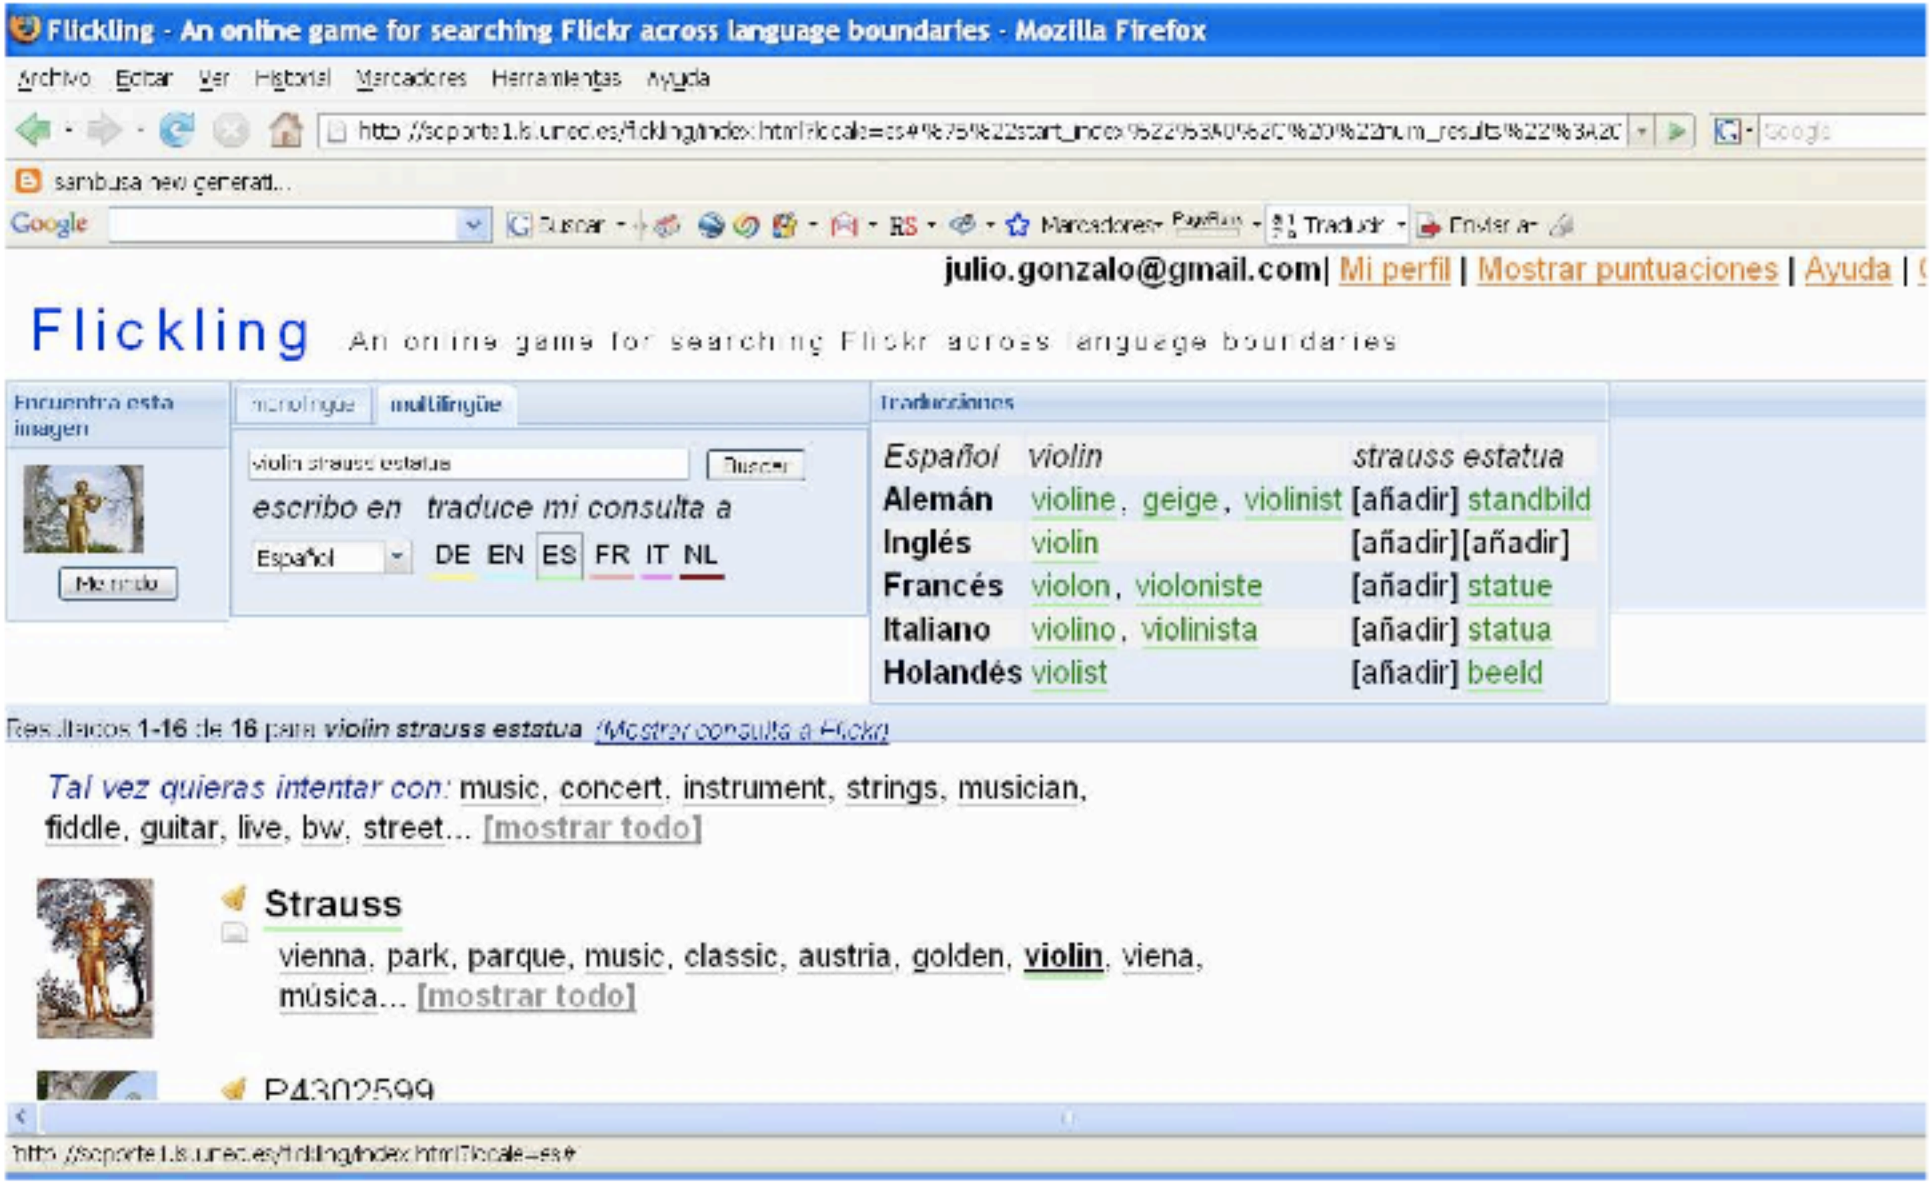

Literature

- E. di Sciascio and M. Mongiello. Query by sketch and relevance feed- back for content-based image retrieval over the web. Journal of Visual Languages and Computing, 10(6):564–584, Dec. 1999.
- G. Giacinto and F. Rolli. Instance-based relevance feedback for image retrieval. In Neural Information Processing Systems Conference, Vancouver, Canada, Dec. 2004.
- J. Gonzalo, P. Clough, and J. Karlgren. Overview of iCLEF 2008: Search log analysis for multilingual image retrieval. In Working Notes of the CLEF Workshop, Aarhus, Denmark, Sept. 2008.
- D. Harman. Relevance feedback revisited. In ACM Special Interest Group on Information Retrieval (SIGIR), pages 1–10, New York, NY, USA, 1992.
- M. J. Huiskes and M. S. Lew. Performance evaluation of relevance feedback methods. In Conference On Image And Video Retrieval, pages 239–248, New York, NY, USA, 2008.
- D. Kelly and J. Teevan. Implicit feedback for inferring user preference: a bibliography. SIGIR Forum, 37(2):18–28, 2003.

Literature

- H. Müller, W. Müller, S. Marchand-Maillet, and D. M. Squire. Strate- gies for positive and negative relevance feedback in image retrieval. In International Conference on Pattern Recognition, volume 1 of Computer Vision and Image Analysis, pages 1043–1046, Barcelona, Spain, Sept. 2000.
- H. Müller, W. Müller, D. M. Squire, S. Marchand-Maillet, and T. Pun. Learning features weights from user behavior in content-based image retrieval. In S. Simoff and O. Zaiane, editors, ACM SIGKDD Interna- tional Conference on Knowledge Discovery and Data Mining (Workshop on Multimedia Data Mining MDM/ KDD2000), Boston, MA, USA, Aug. 2000.
- M. Nakazato and T. S. Huang. 3d mars: immersive virtual reality for content-based image retrieval. In IEEE International Conference on Multimedia and Expo, pages 44– 47, Tokyo, Japan, aug 2001.
- G. P. Nguyen and M. Worring. Optimization of interactive visual- similarity-based search. ACM Transactions on Multimedia Computing, 4(1), Jan. 2008
- G. P. Nguyen, M. Worring, and A. W. Smeulders. Interactive search by direct manipulation of dissimilarity space. IEEE Transactions on Multimedia, 9(7):1404–1415, Nov. 2007.

Literature

- R. Paredes, T. Deselaers, and V. Enrique. A probabilistic model for user relevance feedback on image retrieval. In Workshop on Machine Learning and Multimodal Interaction, Utrecht, The Netherlands, Sept. 2008.
- J. Rocchio. Relevance feedback in information retrieval. In G. Salton, editor, The SMART Retrieval System: Experiments in Automatic Document Processing, pages 313–323. Prentice-Hall, Englewood Cliffs, NJ, USA, 1971.
- Y. Rui, T. S. Huang, and S. Mehrotra. Content-based image retrieval with relevance feedback in mars. In International Conference on Image Processing, volume 2, pages 815–818, Santa Barbara, CA, USA, Oct. 1997.
- L. Setia, J. Ick, and H. Burkhardt. SVM-based relevance feedback in image retrieval using invariant feature histograms. In IAPR Workshop on Machine Vision Applications (MVA), Tsukuba Science City, Japan, May 2005.
- C. G. Snoek, M. Worring, J.-M. Geusenbroek, D. C. Kolma, F. J. Seinstra, and A. W. Smeulders. The semantic pathfinder: Using an authoring metaphor for generic multimedia indexing. IEEE Transaction on Pattern Analysis and Machine Intelligence, 28(10):1678–1689, Oct. 2006.
- Z. Su, H. Zhang, S. Li, and S. Ma. Relevance feedback in content-based image retrieval: Bayesian framework, feature subspaces, and progressive learning. IEEE Transactions on Image Processing, 12(8):924–937, Aug. 2003.
- X. S. Zhou and T. S. Huang. Relevance feedback in image retrieval: A comprehensive review. Multimedia Systems, 8:536–544, 2003.