Indexing Multimedia Databases

Christos Faloutsos

Contents

1 PROBLEM - APPLICATIONS

Multimedia System: A system that can store and retrieve objects/documents with text, voice, images, animation, slides show etc.

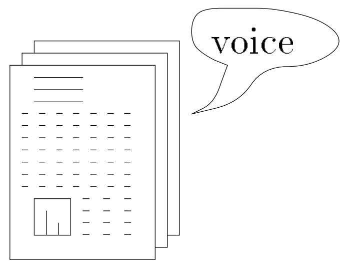

Problem definition:

- given a set of multimedia objects,
- find the ones containing a desirable pattern (or something similar to it)

GOALS: search by content

- Efficient and
- 'complete' (no false dismissals)

Applications:

- time series: financial, marketing, ECGs, voice/sound
- images: education, art, medicine
- higher-d signals: scientific db (eg., meteorology, astrophysics), medicine, entertainment (video)

Sample queries:

- *find companies whose stock prices move similarly*
- *find past days in which the solar magnetic wind showed patterns similar to today's pattern* [Vas93]
- *find X-rays similar to Smith's*
- *in LANDSAT images, find areas with cornfield colors*
- *find the average MRI brain scan of epileptics*

Similarity search, hypothesis testing, rule discovery, data mining [Agrawal et al., SIGMOD 93] [AIS93]

2 FRAMEWORK

- Whole matching vs. Sub-pattern matching
- Range queries vs Nearest Neighbor
- All-pairs queries ('spatial joins')

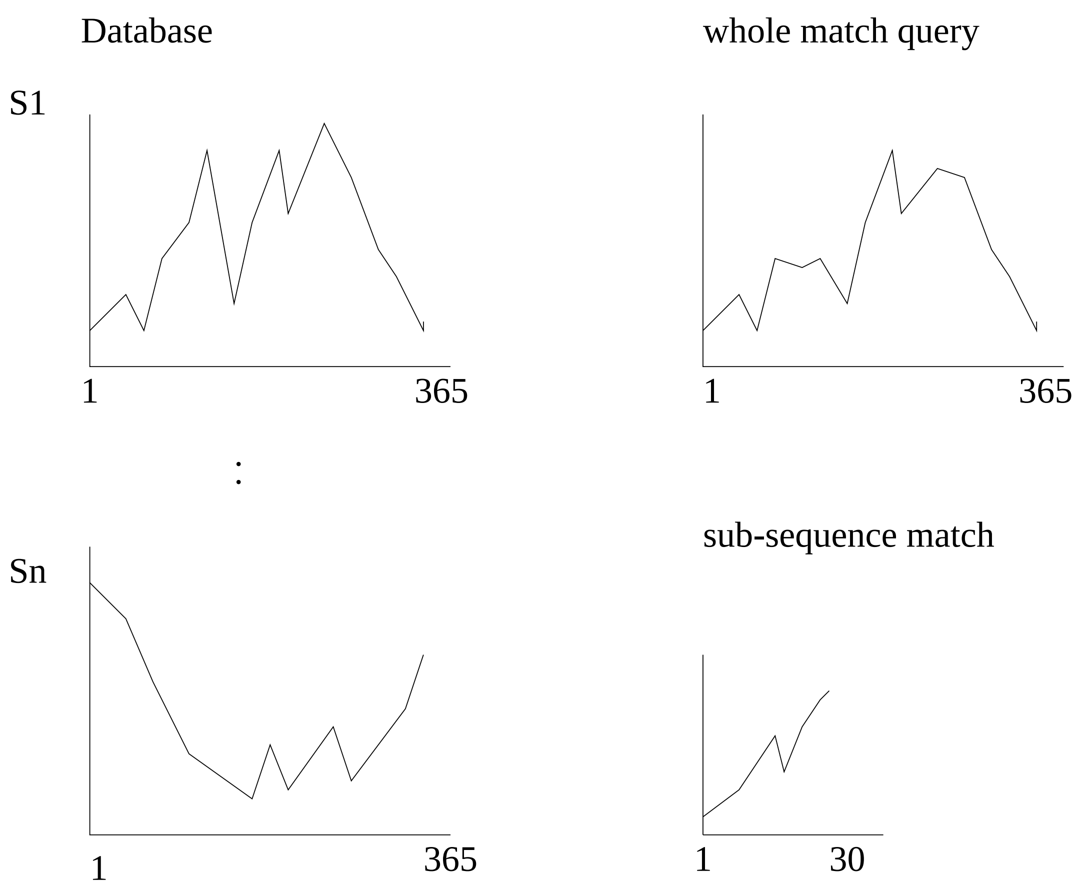

Distance function: by domain expert.
Eg., Euclidean

$$\mathcal{D}(S, Q) \equiv (\sum_{i=1}^{l} |S[i] - Q[i]|^2)^{1/2} \qquad (1)$$

(similarly for vector fields)

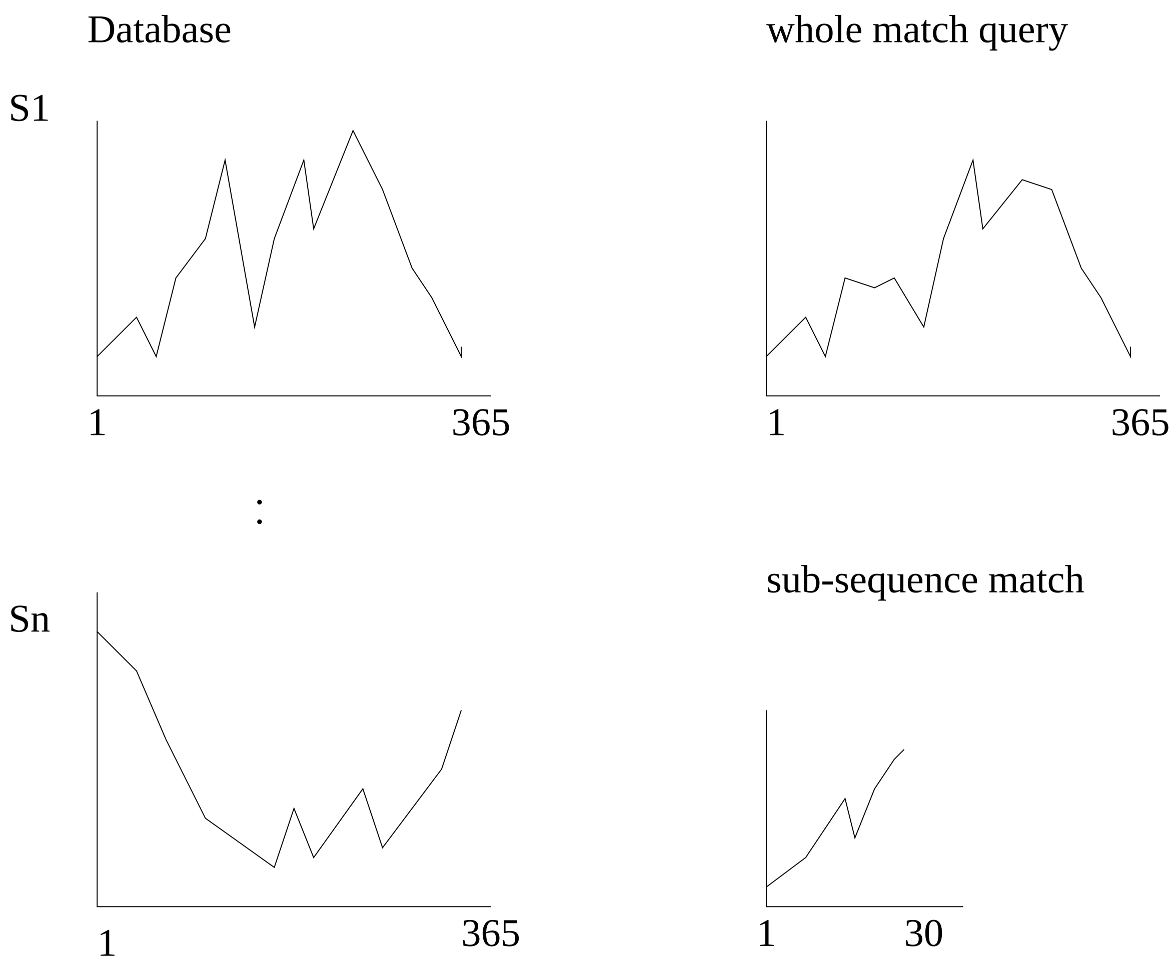
Database
whole match query
S1
1
365
1
365
:
sub-sequence match
Sn
1
365
1
30

Target method:

- fast
- no false dismissals (false alarms are OK)
- dynamic (insertions/deletions/appends)

3 SPATIAL ACCESS METHODS

Objects may be points, rectangles or arbitrary shapes.
Queries:

- *point* queries,
- *range* queries,
- *'all-pairs'/spatial join* queries [BKS93],
- nearest neighbor(s) queries [RKV95].

Applications

- Traditional data base systems.
- multimedia objects, after feature extraction [Jag91b]
- Cartography [Sam90a]
- Computer-Aided Design (CAD).
- Computer vision and robotics [BB82]
- Rule indexing in expert database systems [SSH86]

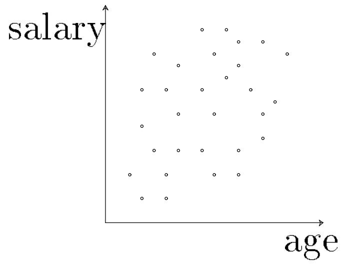

traditional data base

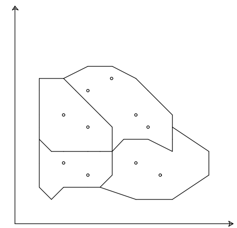

Cartographic data

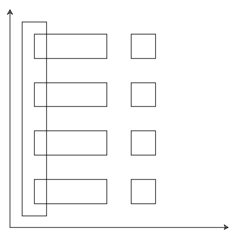

CAD / VLSI design

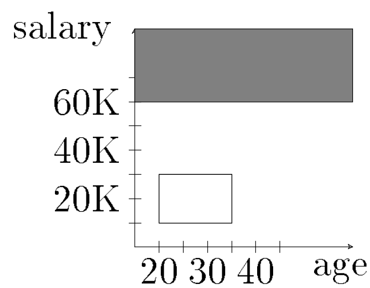

Rule indexing

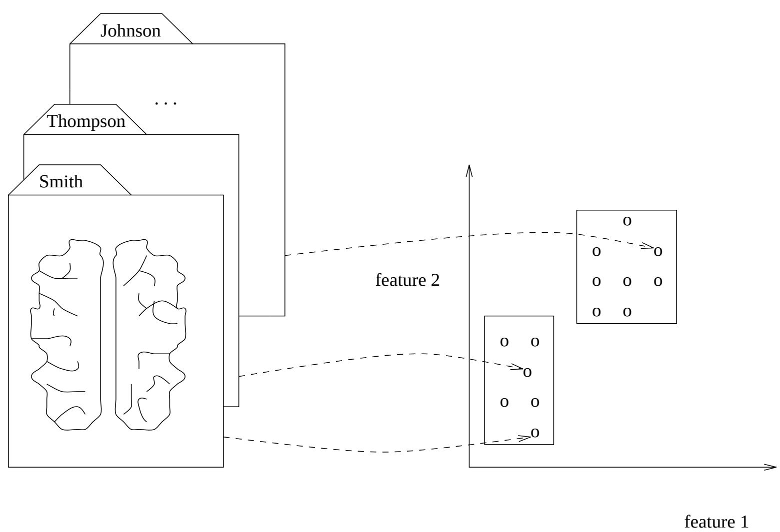

Overview

- space filling curves/ z-ordering / linear quadtrees
- grid files
- R-trees

3.1 Space filling curves

Proposed, among others, by. [Ore86] [SSN87] [OM88] [Ore89] [Ore90] [Sam90b]. Very similar linear quadtrees [Gar82].

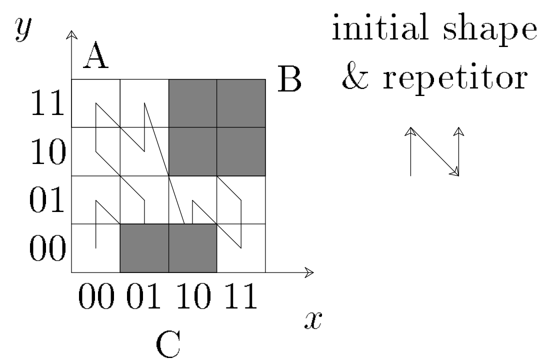

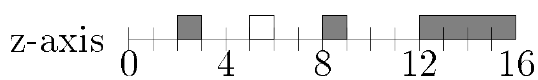

- z_A = Shuffle ("1,2,1,2", x_A, y_A) = Shuffle ("1,2,1,2", 00, 11) = 0101 = $(5)_10$
- z_B = 11 (common prefix of all its blocks)
- z_{C_1} = 0010 = 2
- z_{C_2} = 1000 = 8

NOTICE:

- Relations using z values $\Rightarrow$
  - excellent integration of geometric data bases with relational ones
  - fast processing of geometric queries, using **index** on the z values
- used by the U.S. Bureau of Census - TIGER project [Whi81]
- BUT: regions give too many pieces, unless we use approximations (='redundancy' [Ore90]).

Variations - improvements

- Best distance preserving mapping?

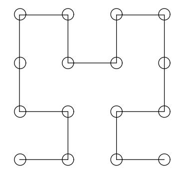

- Study of clustering properties of space filling curves:
  - Exhaustive enumeration, [FR89b]
  - formulas for partial match queries and 2x2 squares, [Jag90]
  - closed formula for z-ordering [RF91]

3.2 Grid File

Dynamic version of multi-attribute hashing [NHS84]

- 2 disk accesses for exact match queries
- symmetric with respect to the attributes
- adapting to non-uniform distributions

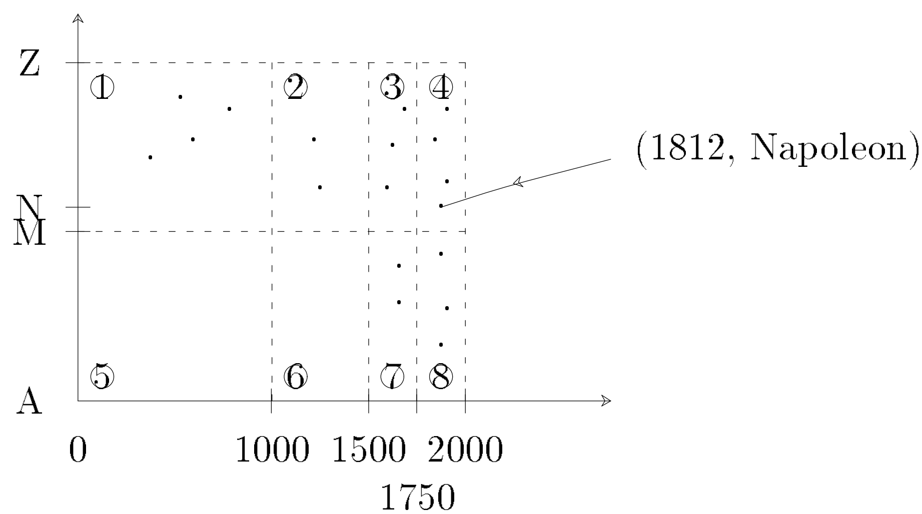

- Every cell → one disk page
- Cuts occur on predefined points (1/2, 1/4 etc of each axis)
- Cuts cut all the way (as opposed to k-d-B-trees)

Implementation

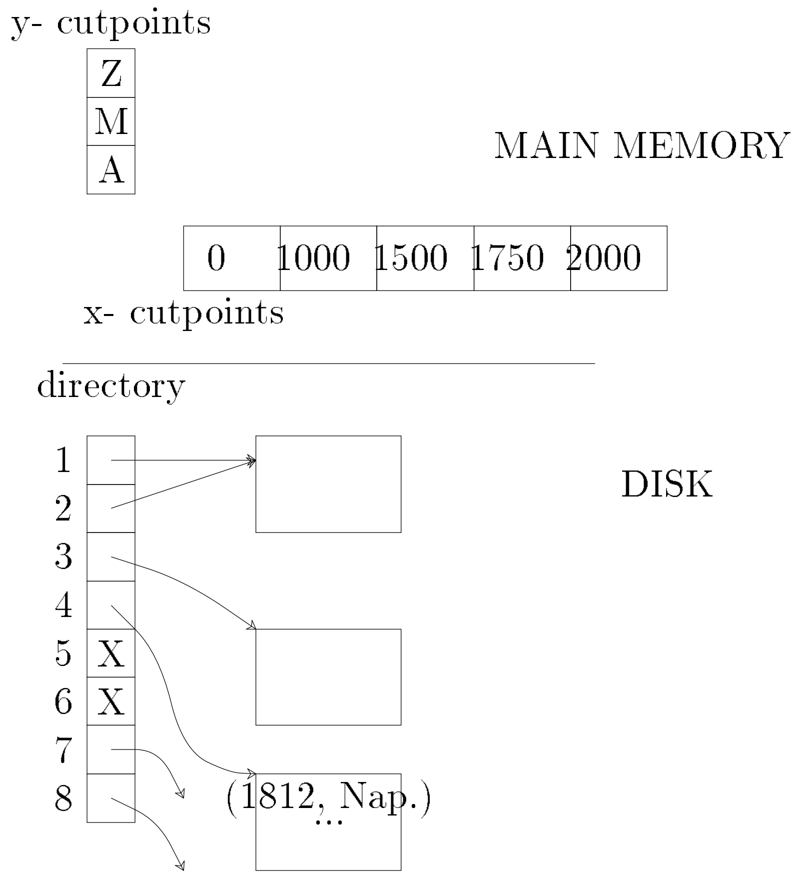

BUT: correlated attributes

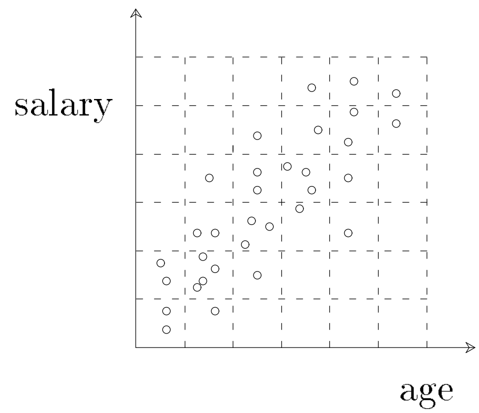

Solutions:

- Rotated grid file [HN83]
- tricell [FR89a]

Other Variations:

- Twin grid file [HSW88]

Notice that it handles *points*; rectangles can be handled by *transforming* them into points in 4-d [HN83]

3.3 R-trees

Idea: Group points in parents - allow parents to **overlap**. [Gut84]

- balanced
- ≥ 50% utilization

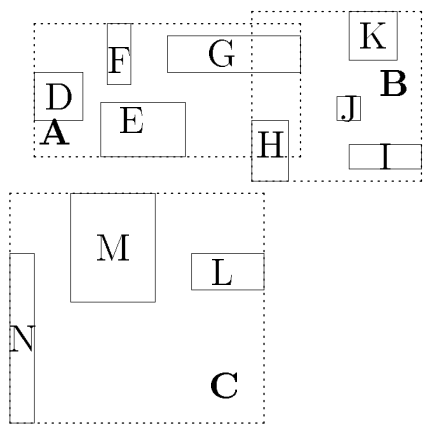

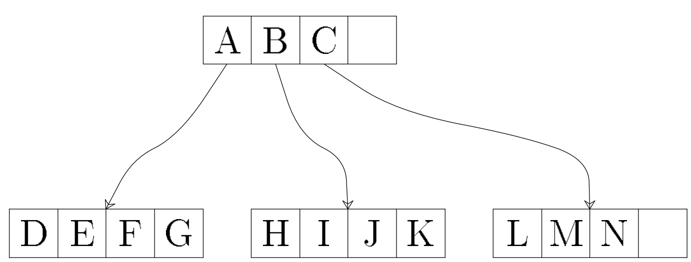

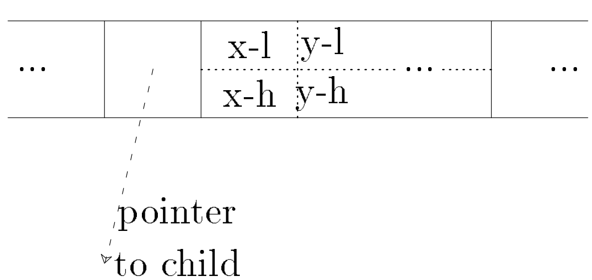

Variations / Improvements:

- Packed R-trees [RL85] [KF93] for static data.

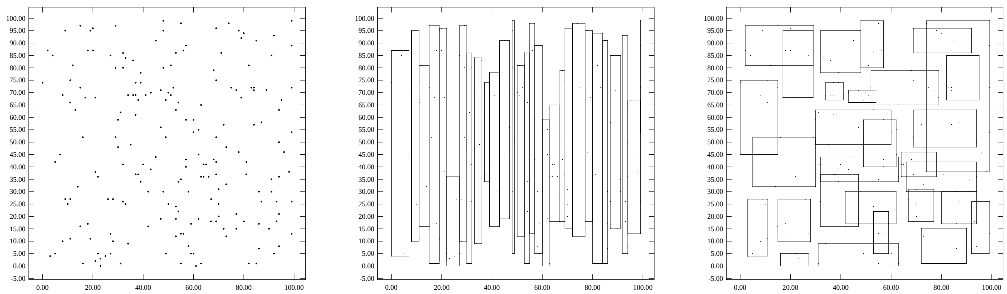

- cell trees [Gun86] introduce diagonal cuts, too.
- R+ trees: [FSR87] [SRF87] No overlap; balanced. Outperform R-trees when there are few large rectangles and several small ones.
- R^*-trees [BKSS90] Main idea: defer splitting, using forced-reinsert on 30% of the most remote rectangles.
- Hilbert R-trees [KF94] defer split, by pushing keys to the 'Hilbert' neighbor.
- Analysis for R-trees: Range queries [PSTW93]; using fractal dimensions [FK94]

3.4 Conclusions

- Z-ordering (Linear quadtrees) and R-trees seem the most promising methods.
- R-trees are more robust for high-d spaces.

4 TIME SERIES

[Agrawal et al., FODO 93] [AFS93]
Distance: Euclidean
Obvious solution: sequential scan.
Q: Something faster?
A: 'Quick-and-dirty' filter:

- extract n features (numbers, eg., avg., etc.)
- map a sequence into a point in n-d feature space
- organize points with a Spatial Access Method (SAM) [Jag91b]
- discard false alarms

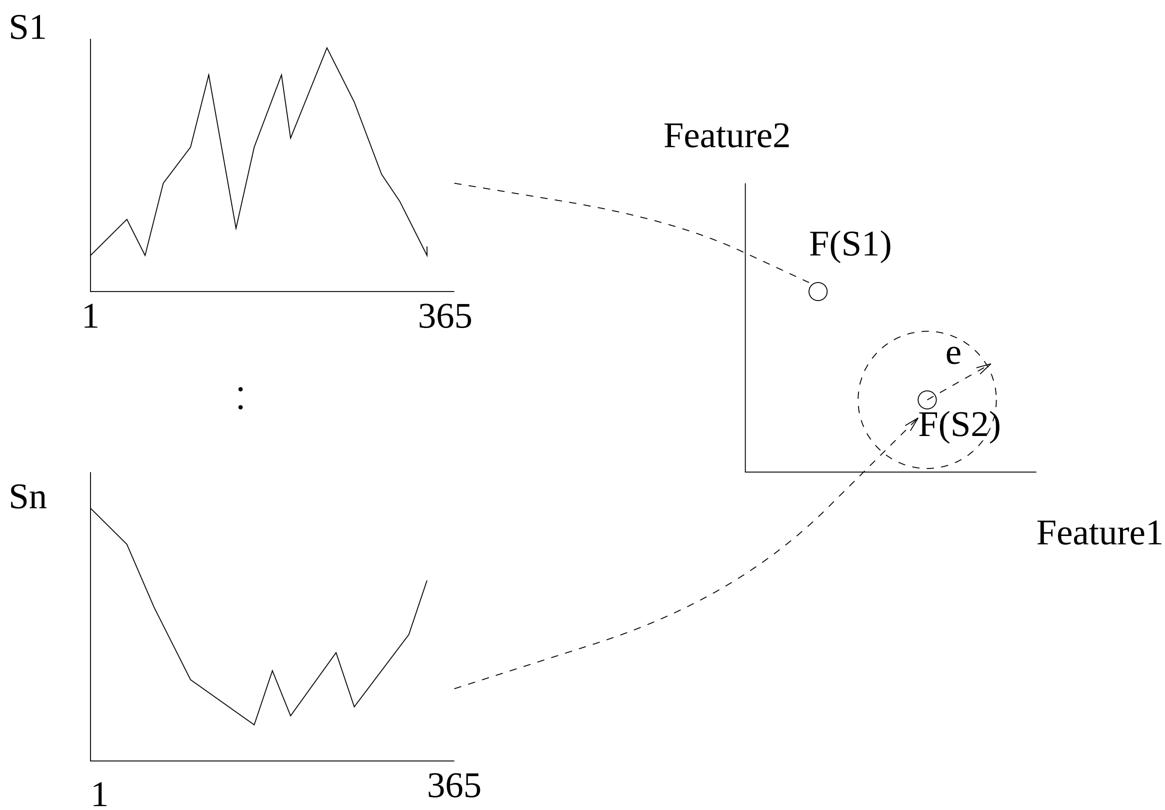

Intuitively,

$$D_{feature}(F(S_1), F(S_2)) \approx D_{actual}(S_1, S_2)$$

Ideally, "=". However equality might be difficult (eg., 'dimensionality curse').

Lemma: To guarantee no false dismissals, *lower-bound* the actual distance:

$$D_{feature}(F(S_1), F(S_2)) \leq D_{actual}(S_1, S_2)$$

I.e., '**it is OK to make things look closer**'.

Solution for the whole-matching problem:

- perform Discrete Fourier Transform (DFT).
- keep first few coefficients

It works well, because:

- DFT maintains distances (*Parseval's theorem*)
- concentrates the 'energy', for 'colored noises'
- Keeping the first few coefficients lower-bounds the distance

DFT formulas and definitions [Ham77] [OS75]:

$$X_f = 1/\sqrt{n} \sum_{i=0}^{n-1} x_i \exp\left(-j 2\pi f i/n\right) \quad f = 0, 1, \ldots, n-1$$

where $j = \sqrt{-1}$.

Energy of a signal $\vec{x}$:

$$E(\vec{x}) \equiv \| \vec{x} \|^2 \equiv \sum_{i=0}^{n-1} |x_i|^2 \tag{2}$$

Theorem (Parseval).

$$\sum_{i=0}^{n-1} |x_i|^2 = \sum_{f=0}^{n-1} |X_f|^2 \tag{3}$$

and also:

$$\| \vec{x} - \vec{y} \|^2 \equiv \| \vec{X} - \vec{Y} \|^2 \tag{4}$$

The first few (eg., 2) coefficients result in a lower bound:

$$(X_1 - Y_1)^2 + (X_2 - Y_2)^2 \leq (X_1 - Y_1)^2 + (X_2 - Y_2)^2 + (X_3 - Y_3)^2 \ldots$$

Colored noises

- brown noise ($1/f^2$ energy spectrum) $\equiv$ random walk (stock price movements, currency exchange rates) [Mandelbrot] [Man77]
- pink noise ($1/f$ energy spectrum) - works of art [Sch91]
- black noise ($1/f^b$ $b > 2$) water-level of rivers [Sch91]

Examples of colored noises:

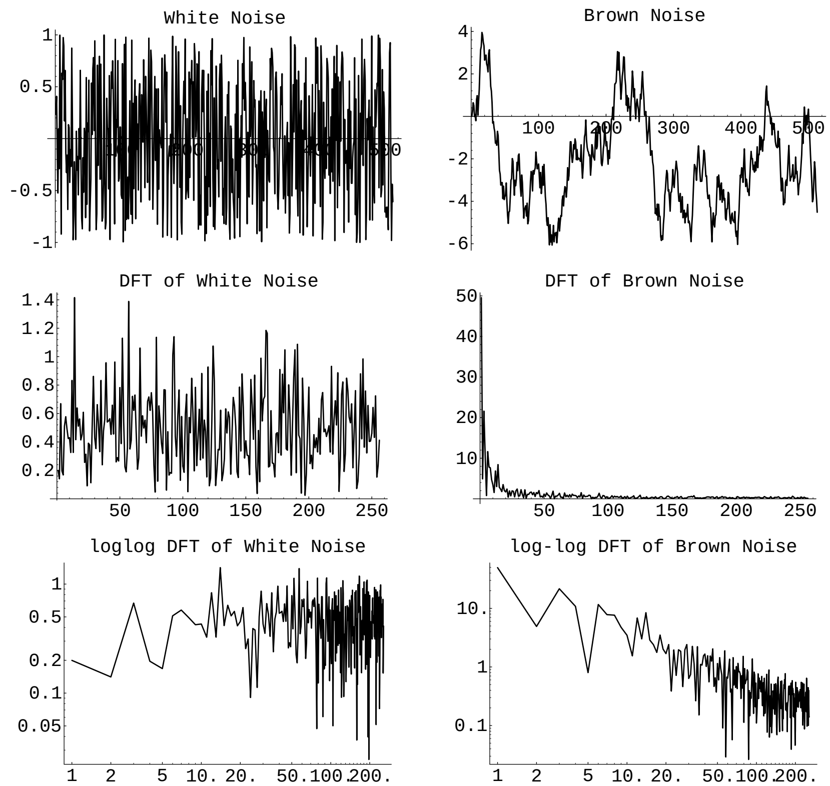

Other 'noises': skewed spectrum, too

Performance:

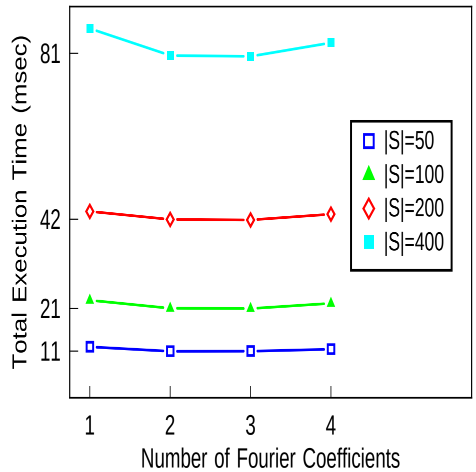

Figure 1: Time per query vs. # Fourier coefficients, for range queries

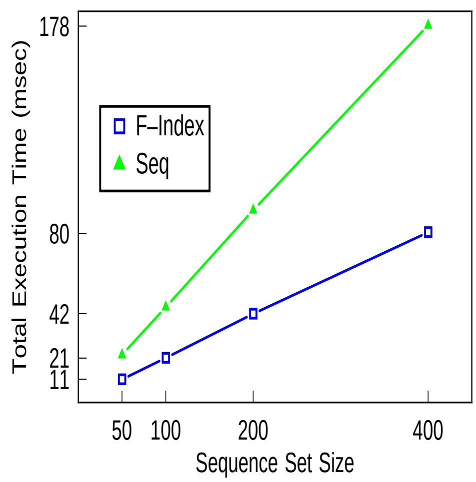

Figure 2: Time per query varying # sequences, for range queries

5 COLOR IMAGES

Much work on

- machine vision [BB82, DH73a], [TSSM89] [WSTM90] [CL91, CW92, LH90, LH92], [HK92], [IX90, Jag91a, KKS+91, CH91, MG89, GNM92, LW88], and [BGS92, SB91, Iok89];
- much work on fast searching;
- little communication between DB and MV communities [ACM91, JN92, NBE+93]

Except recently [HHLC92] [PO93] [?] [FBF+94]
Goal: Queries on color, shape, texture, eg.,

- find photos with color distribution similar to a sunset over the ocean
- find shapes similar to tropical fish

Queries (a) 'by example' (b) 'by sketch'
Support for combinations of color, shape, texture specifications; browsing; 'positional color'.

5.1 Color - features and distance function

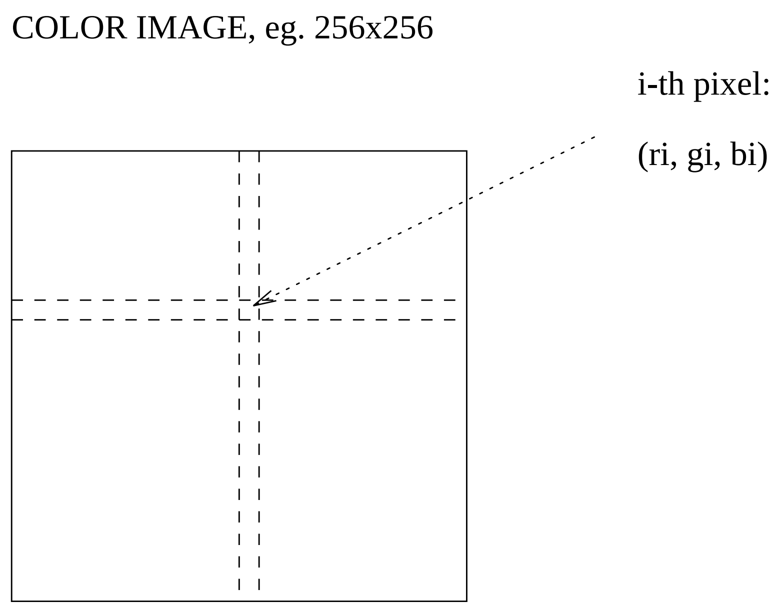

i-th pixel: (r_i, g_i, b_i) (Red, Green, Blue), $0 \leq r_i, g_i, b_i \leq 255$

eg. pink = (200, 60, 60) ($\pm$)

black = (0, 0, 0)

white = (255, 255, 255)

Feature vector: histogram with, say 64 colors (bright red, pink, orange, ..., light blue, dark blue, ...)

Eg., sunset photo:

$$\vec{x} = (80, 85, 75, \ldots, 90, 110, \ldots)$$

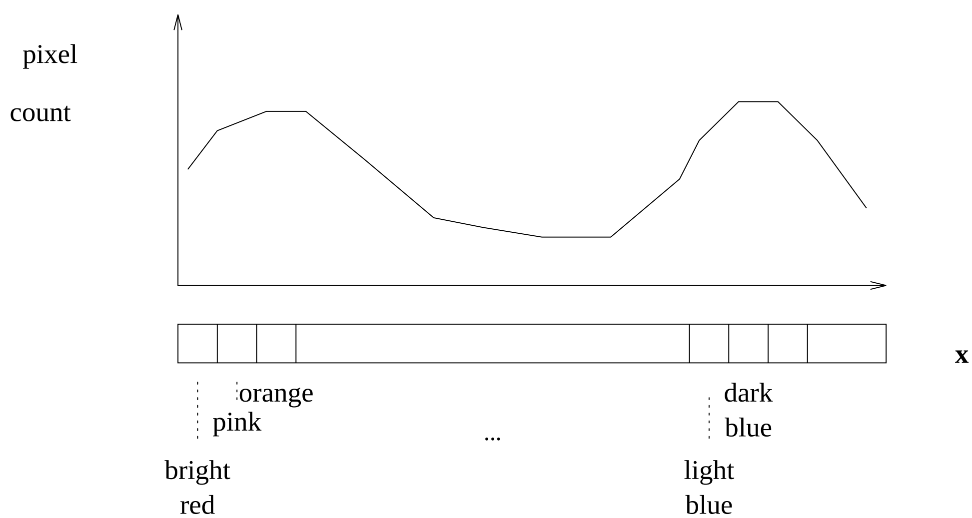

Distance of two histograms $\vec{x}$ and $\vec{q}$: **CROSS TALK**

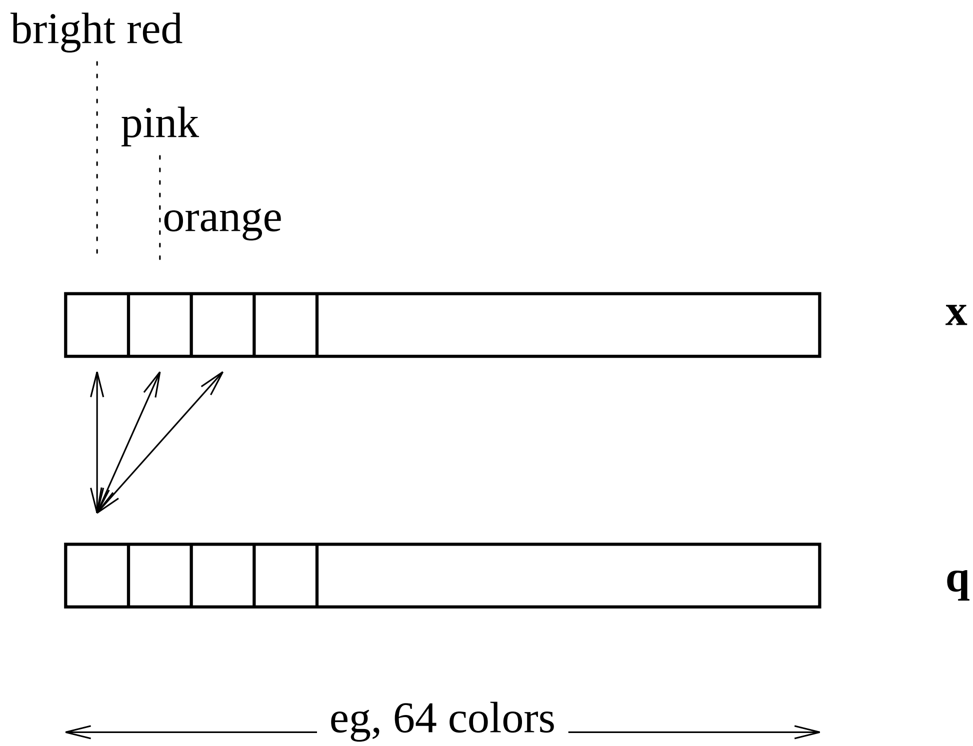

$$distance_{histogram}(\vec{x}, \vec{q}) = (\vec{x} - \vec{q}) \begin{bmatrix} a_{RR} & a_{RP} & \dots \\ a_{PR} & a_{PP} & \dots \\ \dots & \dots & \dots \end{bmatrix} (\vec{x} - \vec{q})^t$$

$$\dots = (\vec{x} - \vec{q}) \mathcal{A} (\vec{x} - \vec{q})^t$$

NONE of the S.A.Ms can handle crosstalk.

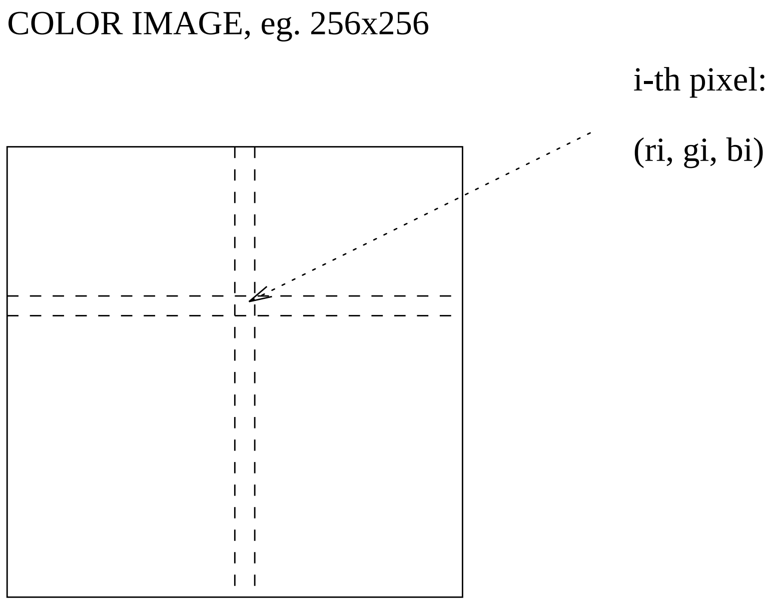

Solution:

- Use a simpler feature vector, eg., average or total R, G, B

$$\vec{x'} = (\sum r_i, \sum g_i, \sum b_i)$$

  with

$$distance_{RGB}(\vec{x'}, \vec{q'}) = \text{Euclidean distance}$$

- exploit the theorem [FBF+94]

$$distance_{RGB}(\vec{x'}, \vec{q'}) \leq distance_{histogram}(\vec{x}, \vec{q})$$

Thus, the $distance_{RGB}()$:

- is euclidean $\Rightarrow$ SAMs apply
- lowerbounds actual distance $\Rightarrow$ no false dismissals

5.2 Shapes

- Features: area, perimeter, moments (≈ 20)
- Distance: (weighted) Euclidean
- Problem: too many features
- Solution: Karhunen-Loeve (K-L) transform ([Fuk90] [DH73b]) $\Rightarrow$ 2-3 coefficients are enough

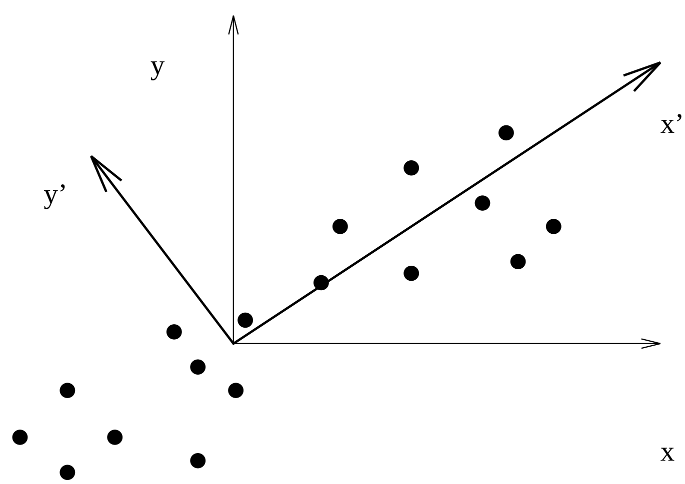

5.3 Performance

For color

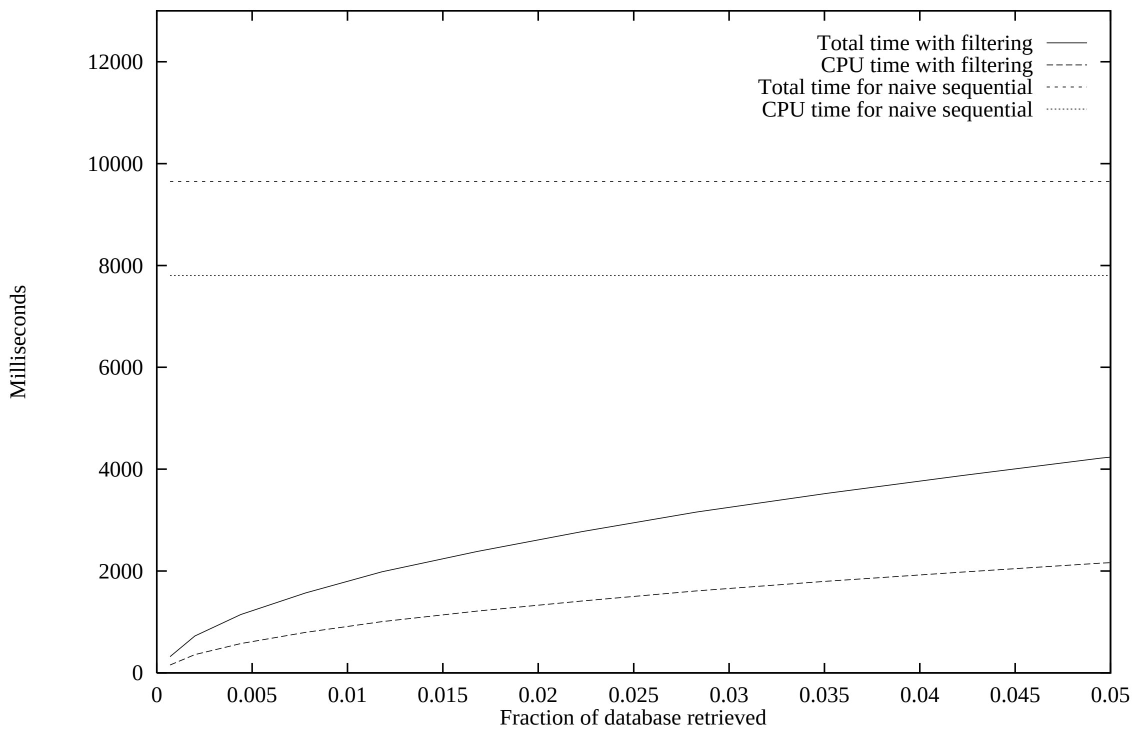

Figure 3: Time spent with sequential retrieval vs. filtered retrieval

Observations: resolving crosstalk

- allows indexing
- saves CPU time ($distance_{histogram}$ is $O(k^2)$).

Performance for shapes

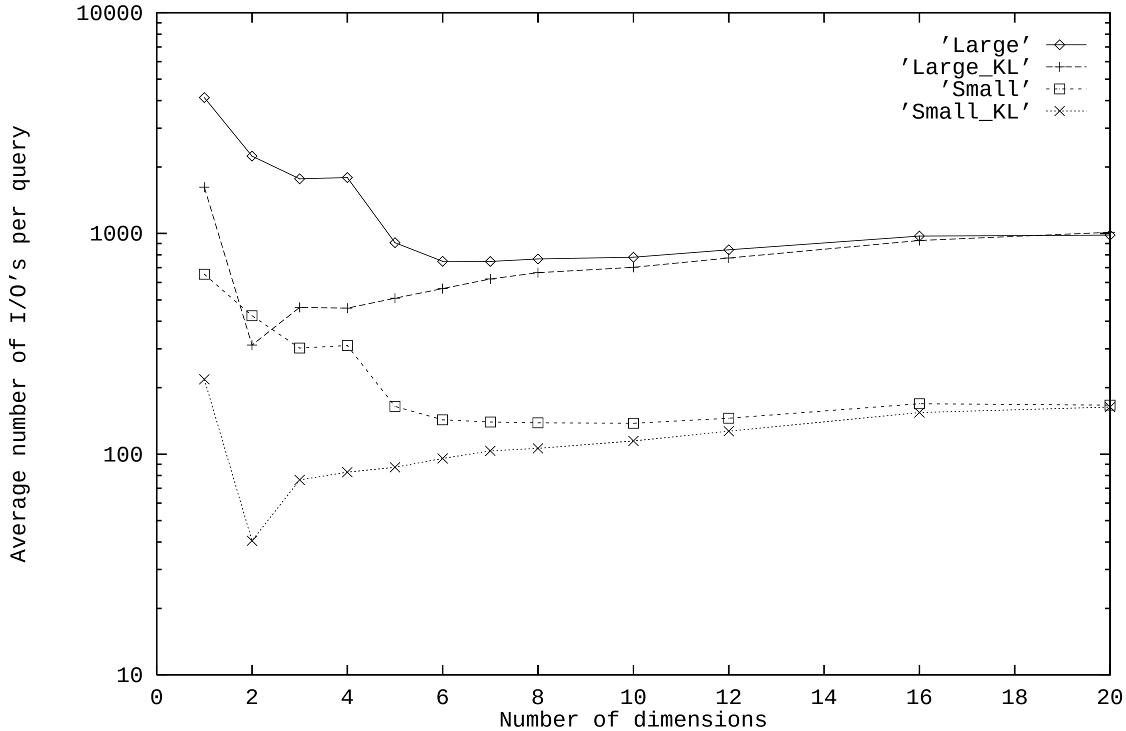

Figure 4: Average Disk I/O's per query, vs. dimensions kept

Observations

- first 2 K-L coefficients are best ($\approx 70\%$ of energy)
- similar performance for larger db

6 3-D MEDICAL IMAGES

Query by content in 2-d medical image databases [HHLC92] [PO93] [?]
Case study for 3-d medical images: QBISM [ACF+93]
Goal: DB support for 3-d medical images (and specifically, for the *Human Brain Mapping* project). Ie.,:

given 3-d brain scans (PET, MRI etc) + demographic characteristics

build a system

to allow exploratory research

- a PET study: $(x, y, z,$ intensity$)$ tuples
- intensity = glucose consumption rate = brain activity

QBISM: Typical of multimedia retrieval system (handles scalar fields = n-d signals)

Sample queries:

- fetch a patient's PET study
- show visual system
- rotate and/or slice
- show areas of high intensity
- list organs within 10mm from hippocampus
- *typical/average scan for 20-yr old left-handed females*
- find brain scans similar to "Smith's"

System architecture:

- DBMS, with 'long fields', extended SQL with *contains*, *intersects* etc., plus
- visualization package (eg., 'Data Explorer') as front end

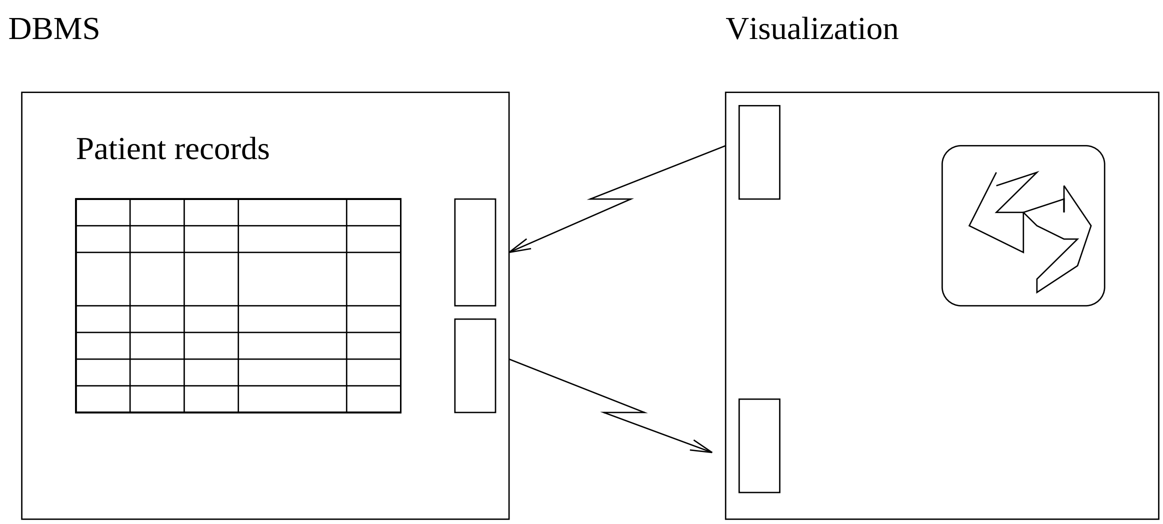

Good news for similarity searches:

- Euclidean distance is a good first step
- DFT leads to a skewed spectrum

Conclusions: Compared to a file-based system, a DBMS allows:

- convenient querying over multiple images
- better performance through early filtering

7 SUB-PATTERN MATCHING

Problem: [Faloutsos et. al., SIGMOD 94] [FRM94]

- Given a collection of sequences (of variable duration)
- find the ones that contain a desirable pattern (within distance ϵ)

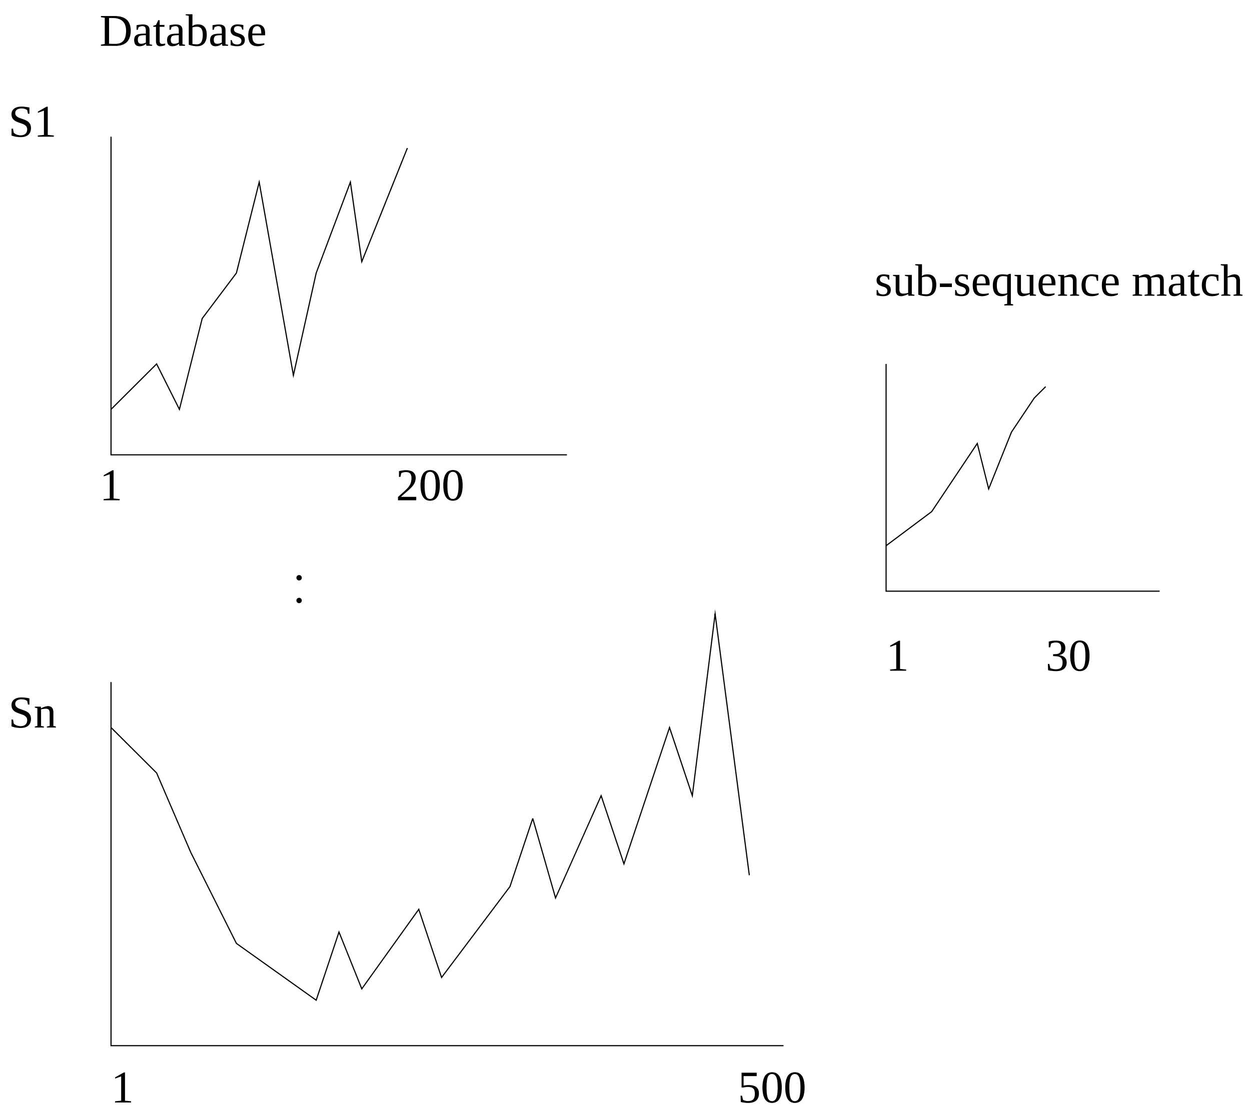

Assumption: queries have length $\geq w$
Proposed method: use sliding, overlapping windows to create trails in feature space.
Features: first few coefficients of the w-point DFT

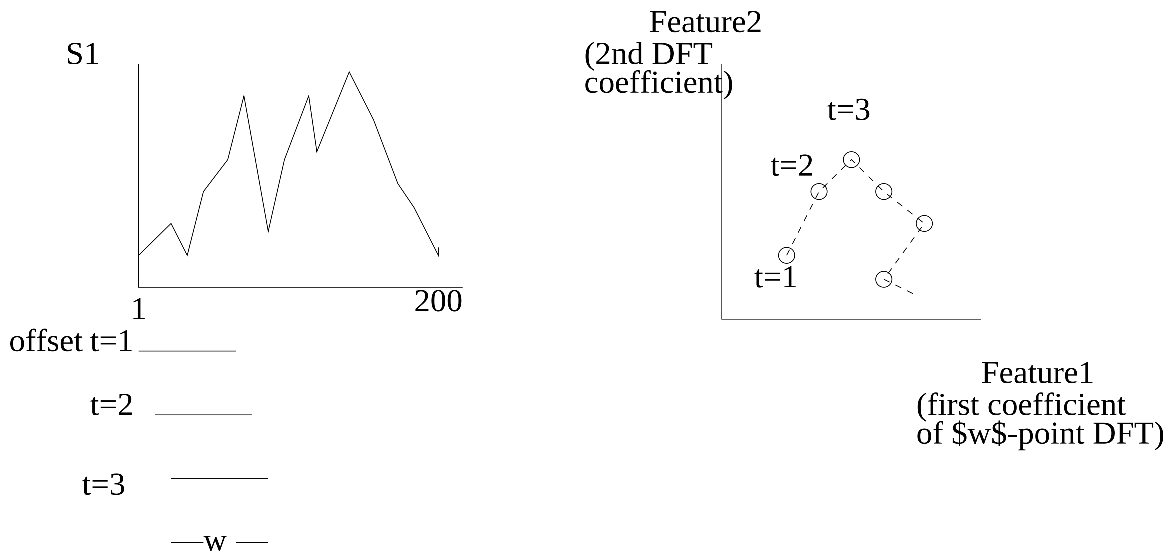

More details on the method:

- subdivide trails and bound them by MBRs
- store MBRs in a S.A.M.

Q: Why not store *all* the points of the trail?
A: Too much space overhead!

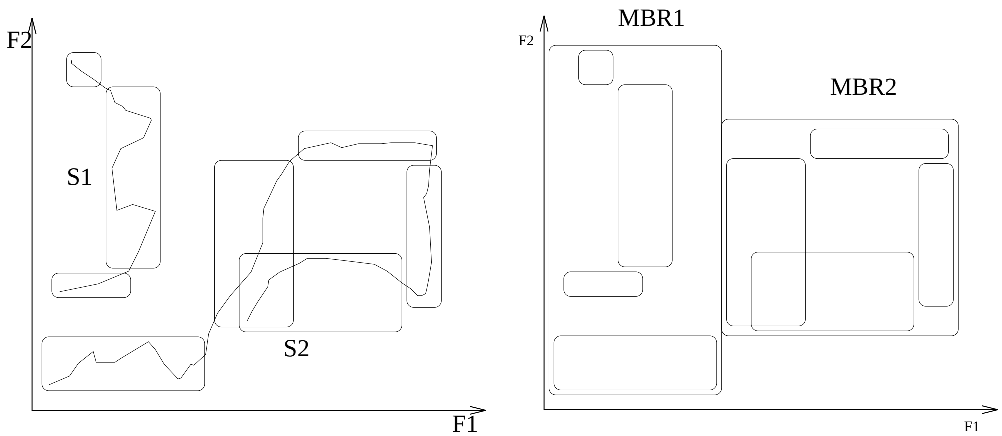

Leaf nodes: set of (seq-id, t-start, t-end, MBR)
Non-leaf nodes: set of (MBR, node-ptr)

More details - how to divide a trail into subtrails?

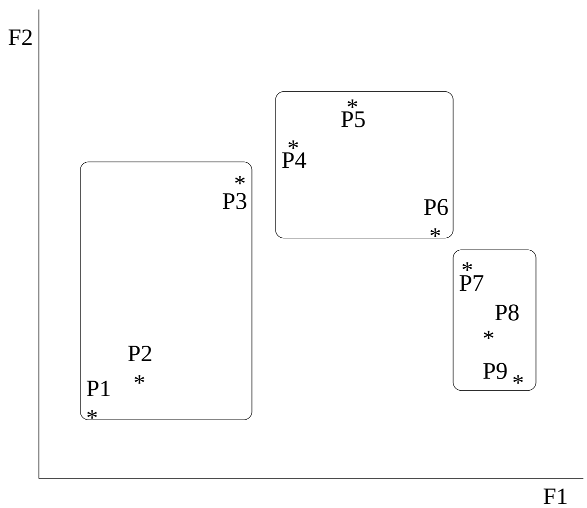

(a) 'fixed'

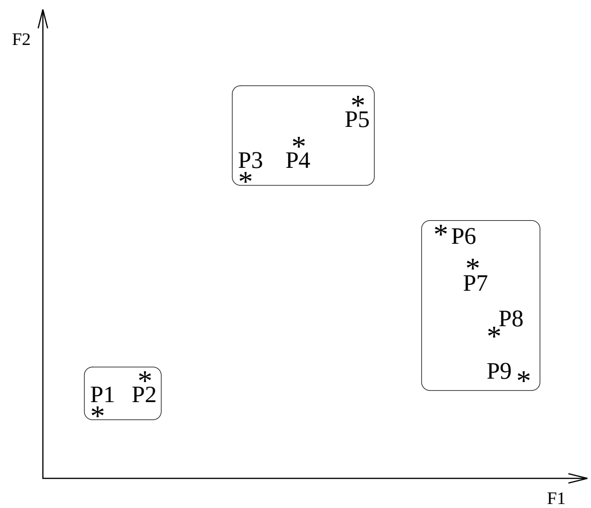

(b) 'adaptive'

Searching, for minimum-length queries:

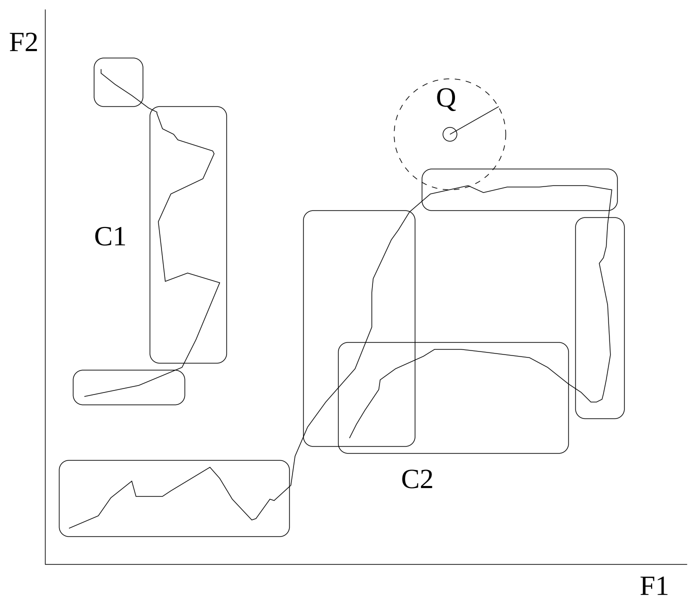

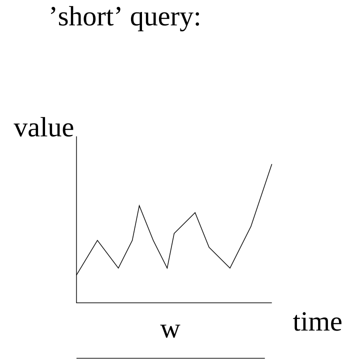

Searching, for longer queries: '*Multi-piece*' algorithm:

- Break it in p pieces of length w
- Search for each piece, with tolerance $\epsilon/\sqrt{p}$
- 'OR' the results and cleanup false alarms

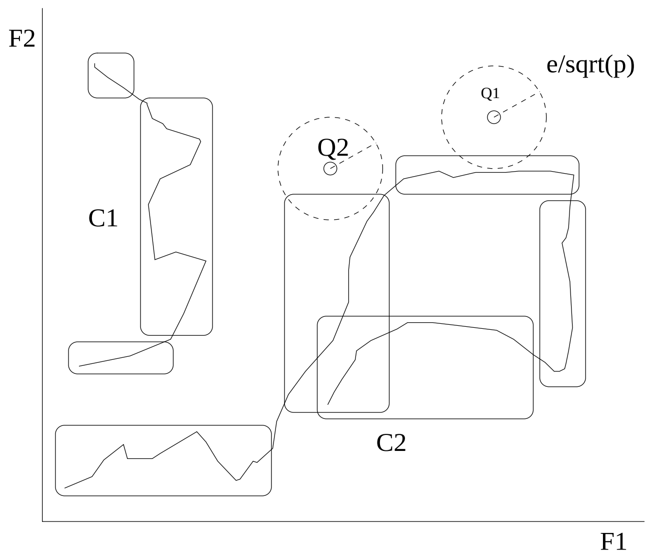

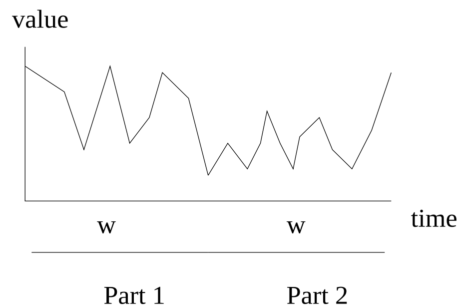

Experiments

- real data (stock prices) - 329,000 points; also, synthetic data.
- implementation in C, AIX, IBM RS/6000
- features: first 3 (complex) DFT coefficients
- 'adaptive' heuristic for sub-trails
- R*-tree [BKSS90] for S.A.M.

Search time experiments - ’short’ queries

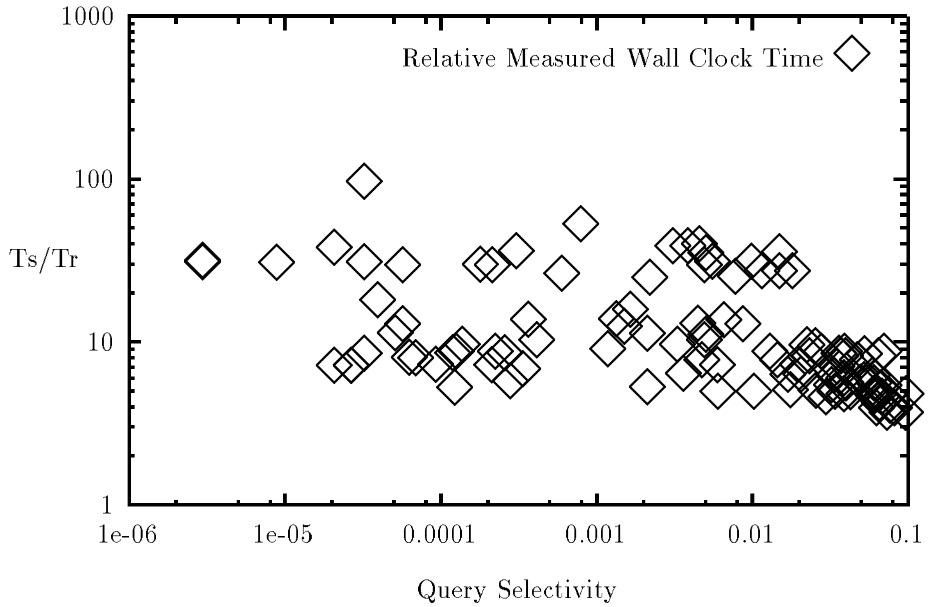

Relative wall clock time vs. selectivity in log-log scale ($Len(Q) = w$=512 points).
Conclusion: 3 to 100 times better than seq. scanning.

8 CONCLUSIONS

Problem: Fast and 'complete' indexing for multimedia databases
Solution: 'Quick and dirty' filter:

- map objects into points in feature space, **lower-bounding** the actual distance. ($\Rightarrow$ 'completeness').
- use any spatial access method ($\Rightarrow$ efficiency).

Main idea, for 'sub-pattern matching'

- map an object into a set of (hyper)-rectangles in feature space
- use a S.A.M.

Occasional mismatches between features and SAMs:

- 'dimensionality curse': Use *distance preserving, energy concentrating* transforms (eg., DFT, DCT, wavelet [RBC+92] etc.)
- 'cross-talk': *diagonalization*

Challenges:

- specific distance function / features for each application
- handling of non-Euclidean distance functions (eg., as in DNA strings)

References

[ACF+93] Manish Arya, William Cody, Christos Faloutsos, Joel Richardson, and Arthur Toga. Qbism: a prototype 3-d medical image database system. *IEEE Data Engineering Bulletin*, 16(1):38–42, March 1993.

[ACM91] ACM SIGIR. *Proceedings of International Conference on Multimedia Information Systems*, Singapore, 1991.

[AFS93] Rakesh Agrawal, Christos Faloutsos, and Arun Swami. Efficient similarity search in sequence databases. In *Fourth Int. Conf. on Foundations of Data Organization and Algorithms (FODO)*, pages 69–84, Evanston, Illinois, October 1993. also available through anonymous ftp, from olympos.cs.umd.edu: ftp/pub/TechReports/fodo.ps.

[AIS93] Rakesh Agrawal, Tomasz Imielinski, and Arun Swami. Mining association rules between sets of items in large databases. *Proc. ACM SIGMOD*, pages 207–216, May 1993.

[BB82] D. Ballard and C. Brown. *Computer Vision.* Prentice Hall, 1982.

[BGS92] Elizabeth Binaghi, Isabella Gagliardi, and Raimondo Schettini. Indexing and fuzzy logic-based retrieval of color images. In *Visual Database Systems, II, IFIP Transactions A-7*, pages 79–92. Elsevier Science Publishers, 1992.

[BKS93] Thomas Brinkhoff, Hans-Peter Kriegel, and Bernhard Seeger. Efficient processing of spatial joins using r-trees. *Proc. of ACM SIGMOD*, pages 237–246, May 1993.

[BKSS90] N. Beckmann, H.-P. Kriegel, R. Schneider, and B. Seeger. The r*-tree: an efficient and robust access method for points and rectangles. *ACM SIGMOD*, pages 322–331, May 1990.

[CH91] Zen Chen and Shinn-Ying Ho. Computer vision for robust 3d aircraft recognition with fast library search. *Pattern Recognition*, 24(5):375–390, 1991.

[CL91] C. C. Chang and S. Y. Lee. Retrieval of similar pictures on pictorial databases. *Pattern Recognition*, 24(7):675–680, 1991.

[CW92] Chin-Chen Chang and Tzong-Chen Wu. Retrieving the most similar symbolic pictures from pictorial databases. *Information Processing and Management*, 28(5):581–588, 1992.

[DH73a] R. Duda and P Hart. *Pattern Classification and Scene Analysis.* Wiley, New York, 1973.

[DH73b] R.O. Duda and P.E. Hart. *Pattern Classification and Scene Analysis.* Wiley, New York, 1973.

[FBF+94] C. Faloutsos, R. Barber, M. Flickner, J. Hafner, W. Niblack, D. Petkovic, and W. Equitz. Efficient and effective querying by image content. *Journal of Intell. Inf. Systems*, 3(3/4):231–262, July 1994.

[FK94] Christos Faloutsos and Ibrahim Kamel. Beyond Uniformity and Irndependence: Analysis of R-trees Using the Concept of Fractal Dimension. *Proc. ACM SIGACT-SIGMOD-SIGART PODS*, pages 4–13, May 1994. Also available as CS-TR-3198, UMIACS-TR-93-130.

[FR89a] C. Faloutsos and W. Rego. Tri-cell: a data structure for spatial objects. *Information Systems*, 14(2):131–139, 1989. early version available as UMIACS-TR-87-15, CS-TR-1829.

[FR89b] C. Faloutsos and S. Roseman. Fractals for secondary key retrieval. *Eighth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems (PODS)*, pages 247–252, March 1989. also available as UMIACS-TR-89-47 and CS-TR-2242.

[FRM94] Christos Faloutsos, M. Ranganathan, and Yannis Manolopoulos. Fast subsequence matching in time-series databases. *Proc. ACM SIGMOD*, pages 419–429, May 1994. 'Best Paper' award; also available as CS-TR-3190, UMIACS-TR-93-131, ISR TR-93-86.

[FSR87] C. Faloutsos, T. Sellis, and N. Roussopoulos. Analysis of object oriented spatial access methods. *Proc. ACM SIGMOD*, pages 426–439 426–439, May 1987. also available as SRC-TR-87-30, UMIACS-TR-86-27, CS-TR-1781.

[Fuk90] Keinosuke Fukunaga. *Introduction to Statistical Pattern Recognition.* Academic Press, 1990. 2nd Edition.

[Gar82] I. Gargantini. An effective way to represent quadtrees. *Comm. of ACM (CACM)*, 25(12):905–910, December 1982.

[GNM92] William I. Grosky, Peter Neo, and Rajiv Mehrotra. A pictorial index mechanism for model-based matching. *Data and Knowledge Engineering*, 8:309–327, 1992.

[Gun86] O. Gunther. The cell tree: an index for geometric data. Memorandum No. UCB/ERL M86/89, Univ. of California, Berkeley, December 1986.

[Gut84] A. Guttman. R-trees: a dynamic index structure for spatial searching. *Proc. ACM SIGMOD*, pages 47–57, June 1984.

[Ham77] Richard Wesley Hamming. *Digital Filters.* Prentice-Hall Signal Processing Series, Englewood Cliffs, N.J., 1977.

[HHLC92] Hou, Hsu, Liu, and Chiu. A content-based indexing technique using relative geometry features. *SPIE 92*, 1662:59–68, 1992.

[HK92] Kyoji Hirata and Toshikazu Kato. Query by visual example. In *Advances in Database Techonology EDBT '92, Third International Conference on Extending Database Technology*, Vienna, Austria, March 1992. Springer-Verlag.

[HN83] K. Hinrichs and J. Nievergelt. The grid file: a data structure to support proximity queries on spatial objects. *Proc. of the WG'83 (Intern. Workshop on Graph Theoretic Concepts in Computer Science)*, pages 100–113, 1983.

[HSW88] A. Hutflesz, H.-W. Six, and P. Widmayer. Twin grid files: Space optimizing access schemes. *Proc. of ACM SIGMOD*, pages 183–190, June 1988.

[Iok89] Mikihiro Ioka. A method of defining the similarity of images on the basis of color information. Technical report RT-0030, IBM Tokyo Research Lab, 1989.

[IX90] M. A. Ireton and C. S. Xydeas. Classification of shape for content retrieval of images in a multimedia database. In *Sixth International Conference on Digital Processing of Signals in Communications*, pages 111 – 116, Loughborough, UK, 2-6 Sept., 1990. IEE.

[Jag90] H.V. Jagadish. Linear clustering of objects with multiple attributes. *ACM SIGMOD Conf.*, pages 332–342, May 1990.

[Jag91a] H. V. Jagadish. A retrieval technique for similar shapes. In *International Conference on Management of Data, SIGMOD 91*, pages 208–217, Denver, CO, May 1991. ACM.

[Jag91b] H.V. Jagadish. A retrieval technique for similar shapes. *Proc. ACM SIGMOD Conf.*, pages 208–217, May 1991.

[JN92] Ramesh Jain and Wayne Niblack. NSF Workshop on Visual Information Management, February 1992.

[KF93] Ibrahim Kamel and Christos Faloutsos. On packing r-trees. *Second Int. Conf. on Information and Knowledge Management (CIKM)*, November 1993.

[KF94] Ibrahim Kamel and Christos Faloutsos. Hilbert R-tree: An Improved R-tree Using Fractals. In *Proceedings of VLDB Conference,*, pages 500–509, Santiago, Chile, September 1994.

[KKS+91] T. Kato, T. Kurita, H. Shimogaki, T. Mizutori, and K. Fujimura. A cognitive approach to visual interaction. In *International Conference of Multimedia Information Systems, MIS'91*, pages 109–120. ACM and National University of Singapore, January 1991.

[LH90] Suh-Yin Lee and Fang-Jung Hsu. 2d c-string: A new spatial knowledge representation for image database systems. *Pattern Recognition*, 23(10):1077–1087, 1990.

[LH92] Suh-Yin Lee and Fang-Jung Hsu. Spatial reasoning and similarity retrieval of images using 2d c-string knowledge representation. *Pattern Recognition*, 25(3):305–318, 1992.

[LW88] Yehezkel Lamdan and Haim J. Wolfson. Geometric hashing: A general and efficient model-based recognition scheme. In *2nd International Conference on Computer Vision (ICCV)*, pages 238–249, Tampa, Florida, 1988. IEEE.

[Man77] B. Mandelbrot. *Fractal Geometry of Nature*. W.H. Freeman, New York, 1977.

[MG89] Rajiv Mehrotra and William I. Grosky. Shape matching utilizing indexed hypotheses generation and testing. *IEEE Transactions on Robotics and Automation*, 5(1):70–77, 1989.

[NBE+93] W. Niblack, R. Barber, W. Equitz, M. Flickner, E. Glasman, D. Petkovic, P. Yanker, C. Faloutsos, and G. Taubin. The QBIC project: Querying images by content using color, texture, and shape. *SPIE 1993 International Symposium on Electronic Imaging: Science & Technology, Conference 1908, Storage and Retrieval for Image and Video Databases*, February 1993.

[NHS84] J. Nievergelt, H. Hinterberger, and K.C. Sevcik. The grid file: an adaptable, symmetric multikey file structure. *ACM TODS*, 9(1):38–71, March 1984.

[OM88] J.A. Orenstein and F.A. Manola. Probe spatial data modeling and query processing in an image database application. *IEEE Trans. on Software Engineering*, 14(5):611–629, May 1988.

[Ore86] J. Orenstein. Spatial query processing in an object-oriented database system. *Proc. ACM SIGMOD*, pages 326–336, May 1986.

[Ore89] J.A. Orenstein. Redundancy in spatial databases. *Proc. of ACM SIGMOD Conf.*, May 1989.

[Ore90] J.A. Orenstein. A comparison of spatial query processing techniques for native and parameter spaces. *Proc. of ACM SIGMOD Conf.*, pages 343–352, 1990.

[OS75] Alan Victor Oppenheim and Ronald W. Schafer. *Digital Signal Processing*. Prentice-Hall, Englewood Cliffs, N.J., 1975.

[PO93] Euripides G.M. Petrakis and Stelios C. Orphanoudakis. Methodology for the Representation, Indexing and Retrieval of Images by Content. *Image and Vision Computing*, 11(8):504–521, October 1993.

[PSTW93] B. Pagel, H. Six, H. Toben, and P. Widmayer. Towards an analysis of range query performance. *Proc. of ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems (PODS)*, pages 214–221, May 1993.

[RBC+92] Mary Beth Ruskai, Gregory Beylkin, Ronald Coifman, Ingrid Daubechies, Stephane Mallat, Yves Meyer, and Louise Raphael. *Wavelets and Their Applications.* Jones and Bartlett Publishers, Boston, MA, 1992.

[RF91] Yi Rong and Christos Faloutsos. Analysis of the clustering property of peano curves. Techn. Report CS-TR-2792, UMIACS-TR-91-151, Univ. of Maryland, December 1991.

[RKV95] Nick Roussopoulos, Steve Kelley, and F. Vincent. Nearest Neighbor Queries. *Proc. of ACM-SIGMOD*, pages 71–79, May 1995.

[RL85] N. Roussopoulos and D. Leifker. Direct spatial search on pictorial databases using packed R-trees. *Proc. ACM SIGMOD*, May 1985.

[Sam90a] H. Samet. *Applications of Spatial Data Structures Computer Graphics, Image Processing and GIS.* Addison-Wesley, 1990.

[Sam90b] H. Samet. *The Design and Analysis of Spatial Data Structures.* Addison-Wesley, 1990.

[SB91] Michael J. Swain and Dana H. Ballard. Color indexing. *International Journal of Computer Vision*, 7(1):11–32, 1991.

[Sch91] Manfred Schroeder. *Fractals, Chaos, Power Laws: Minutes From an Infinite Paradise.* W.H. Freeman and Company, New York, 1991.

[SRF87] T. Sellis, N. Roussopoulos, and C. Faloutsos. The r+ tree: a dynamic index for multi-dimensional objects. In *Proc. 13th International Conference on VLDB*, pages 507–518, England,, September 1987. also available as SRC-TR-87-32, UMIACS-TR-87-3, CS-TR-1795.

[SSH86] M. Stonebraker, T. Sellis, and E. Hanson. Rule indexing implementations in database systems. In *Proceedings of the First International Conference on Expert Database Systems*, Charleston, SC, April 1986.

[SSN87] C.A. Shaffer, H. Samet, and R.C. Nelson. Quilt: a geographic information system based on quadtrees. Technical Report CS-TR-1885.1, Univ. of Maryland, Dept. of Computer Science, July 1987. to appear in the International Journal of Geographic Information Systems.

[TSSM89] Satoshi Tanaka, Mitsuhide Shima, Jun'ichi Shibayama, and Akira Maeda. Retrieval method for an image database based on topological structure. In *Applications of Digital Image Processing*, volume 1153, pages 318–327. SPIE, 1989.

[Vas93] Dimitris Vassiliadis. The input-state space approach to the prediction of auroral geomagnetic activity from solar wind variables. *Int. Workshop on Applications of Artificial Intelligence in Solar Terrestrial Physics*, September 1993.

[Whi81] M. White. *N-Trees: Large Ordered Indexes for Multi-Dimensional Space.* Application Mathematics Research Staff, Statistical Research Division, U.S. Bureau of the Census, December 1981.

[WSTM90] Koji Wakimoto, Mitsuhide Shima, Satoshi Tanaka, and Akira Maeda. An intelligent user interface to an image database using a figure interpretation method. In *9th Int. Conference on Pattern Recognition*, volume 2, pages 516–991, 1990.