

VISUAL OBJECT RECOGNITION

STATE-OF-THE-ART
TECHNIQUES AND
PERFORMANCE EVALUATION

LECTURE 6: TREE SEARCH AND HASHING

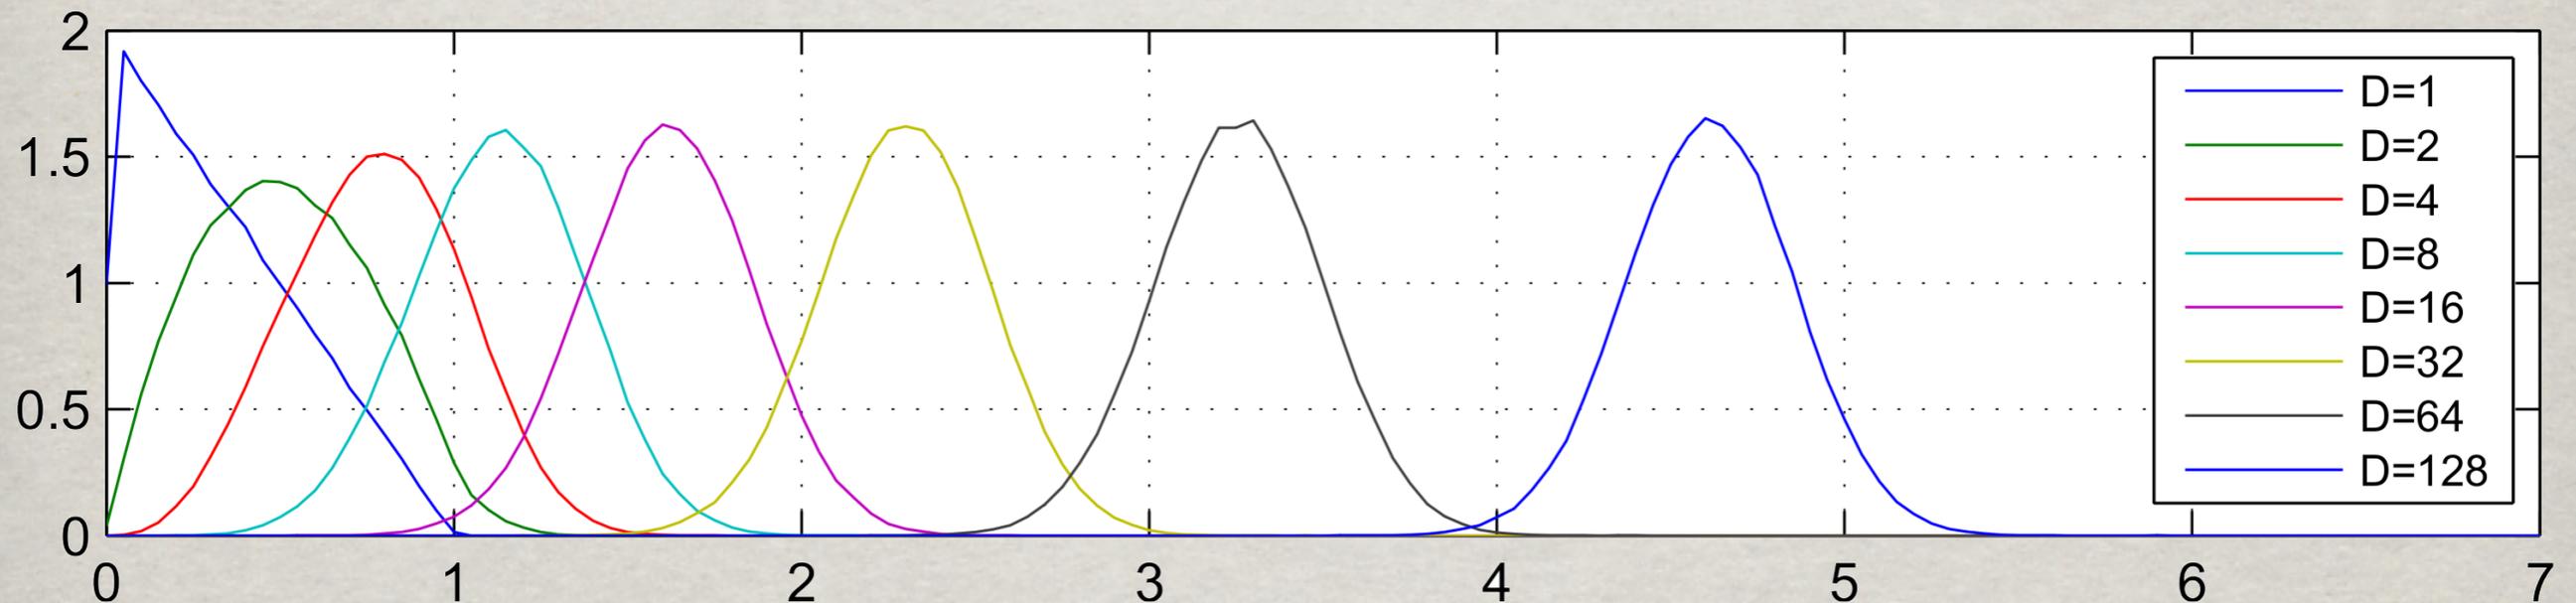
- ✻ High dimensional spaces
- ✻ kD-Trees and Best Bin First (BBF)
- ✻ Ball Trees
- ✻ K-means tree
- ✻ Geometric Hashing

MOTIVATION

- ✱ Finding *the best match* to a query descriptor q in a database with N prototypes $p_1 \dots p_N$ costs $O(N)$.
- ✱ For a database with thousands or millions of descriptors this is expensive.
- ✱ A tree can find *several good matches* (near neighbours) in $O(\log N)$ time.
- ✱ A hash table can find *a good match* in $O(1)$ time.

HIGH DIMENSIONAL SPACES

- ☼ Distances in high dimensional spaces are higher on average!



Expected distance for two points in D-dim unit cube

- ☼ Small distances are unlikely in \mathbb{R}^D for high D.

HIGH DIMENSIONAL SPACES

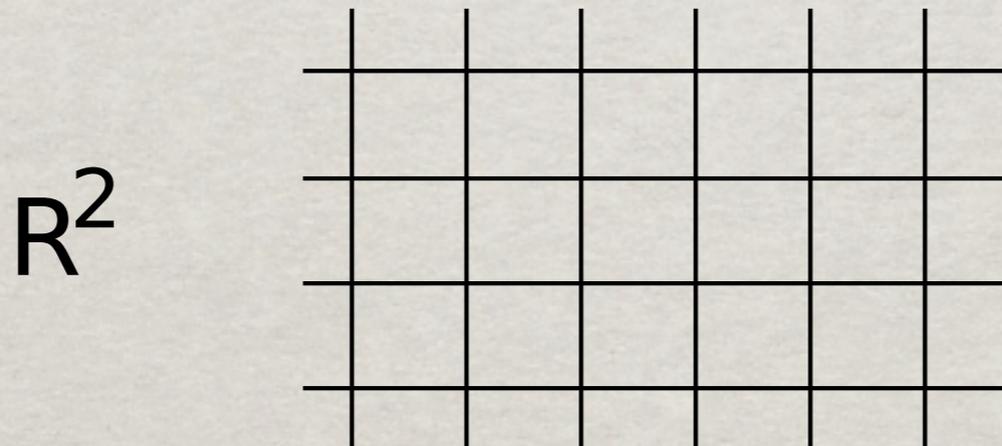
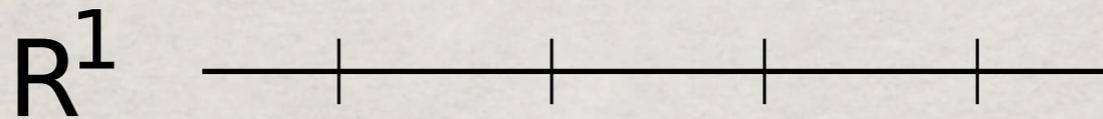
- ✱ Volume shrinks relative to area.
Example: unit “ball”

dimension	volume	area	volume/area
1	$2r$	2	r
2	πr^2	$2\pi r$	$r/2$
3	$4\pi r^3/3$	$4\pi r^2$	$r/3$
4	$\pi^2 r^4/2$	$2\pi^2 r^3$	$r/4$

- ✱ This means that a decision region in \mathbb{R}^D has increasingly more edges as D increases.

HIGH DIMENSIONAL SPACES

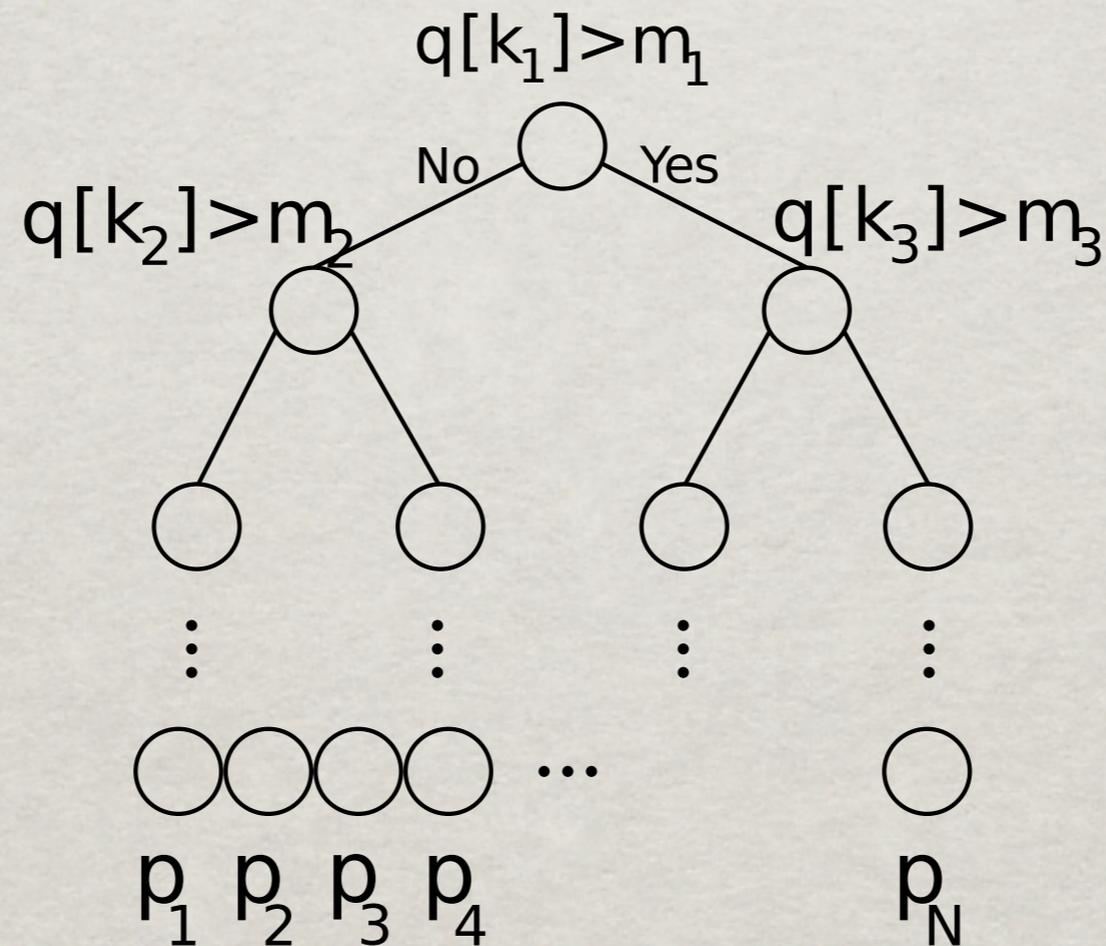
- ✱ Box decision regions have 2 edges in \mathbb{R}^1 , 4 in \mathbb{R}^2 , 6 in \mathbb{R}^3 , 8 in \mathbb{R}^4 , ... $2D$ in \mathbb{R}^D



...

KD-TREES

✱ A binary tree for search in \mathbb{R}^k



KD-TREES

- ✱ Binary search only works in 1D, in higher dimensions the kD-tree gives a *near neighbour*.
- ✱ Tree construction algorithm:
 1. Select dimension k_n with largest variance
 2. Split dataset in two along selected dimension at median value, m_n .
 3. Repeat for each of the subsets.

KD-TREES

- ✱ Search for one neighbour is just one pass down the tree, and thus computation time is proportional to tree depth, d
- ✱ Tree depth $d = \lceil \log_2 N \rceil$
- ✱ To find more neighbours, the original algorithm suggested a depth-first search with branch pruning.
- ✱ If $e_{\text{curr}} < q[k_n] - m_n$ then skip branch.

BEST-BIN-FIRST

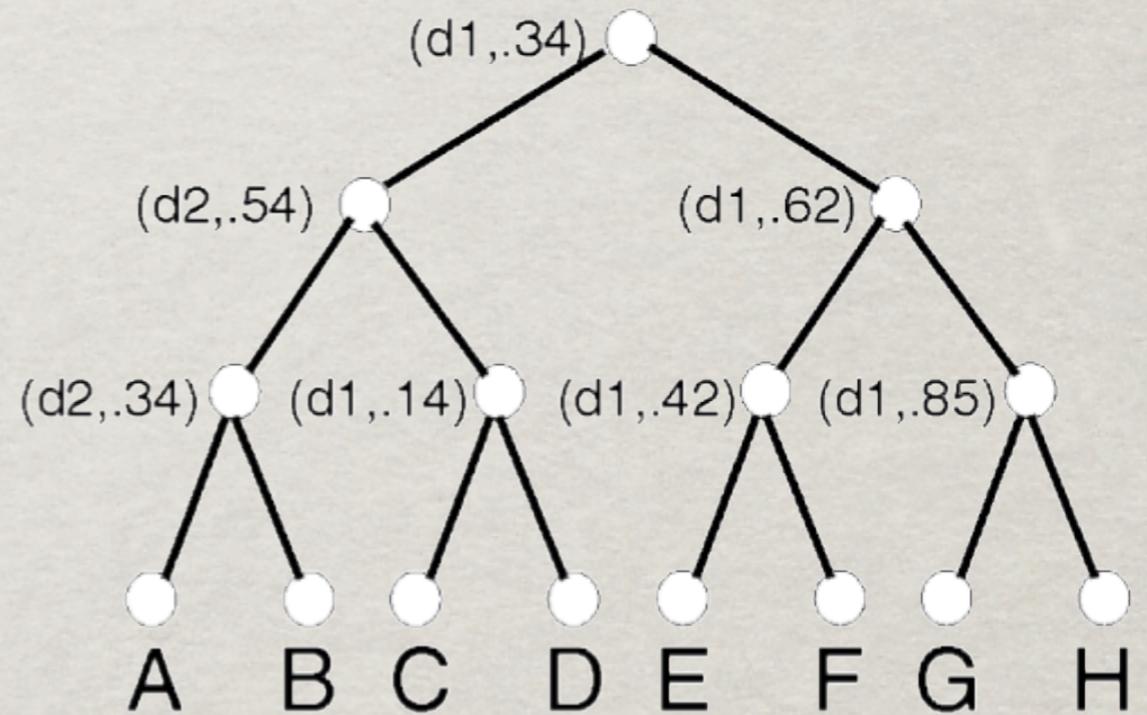
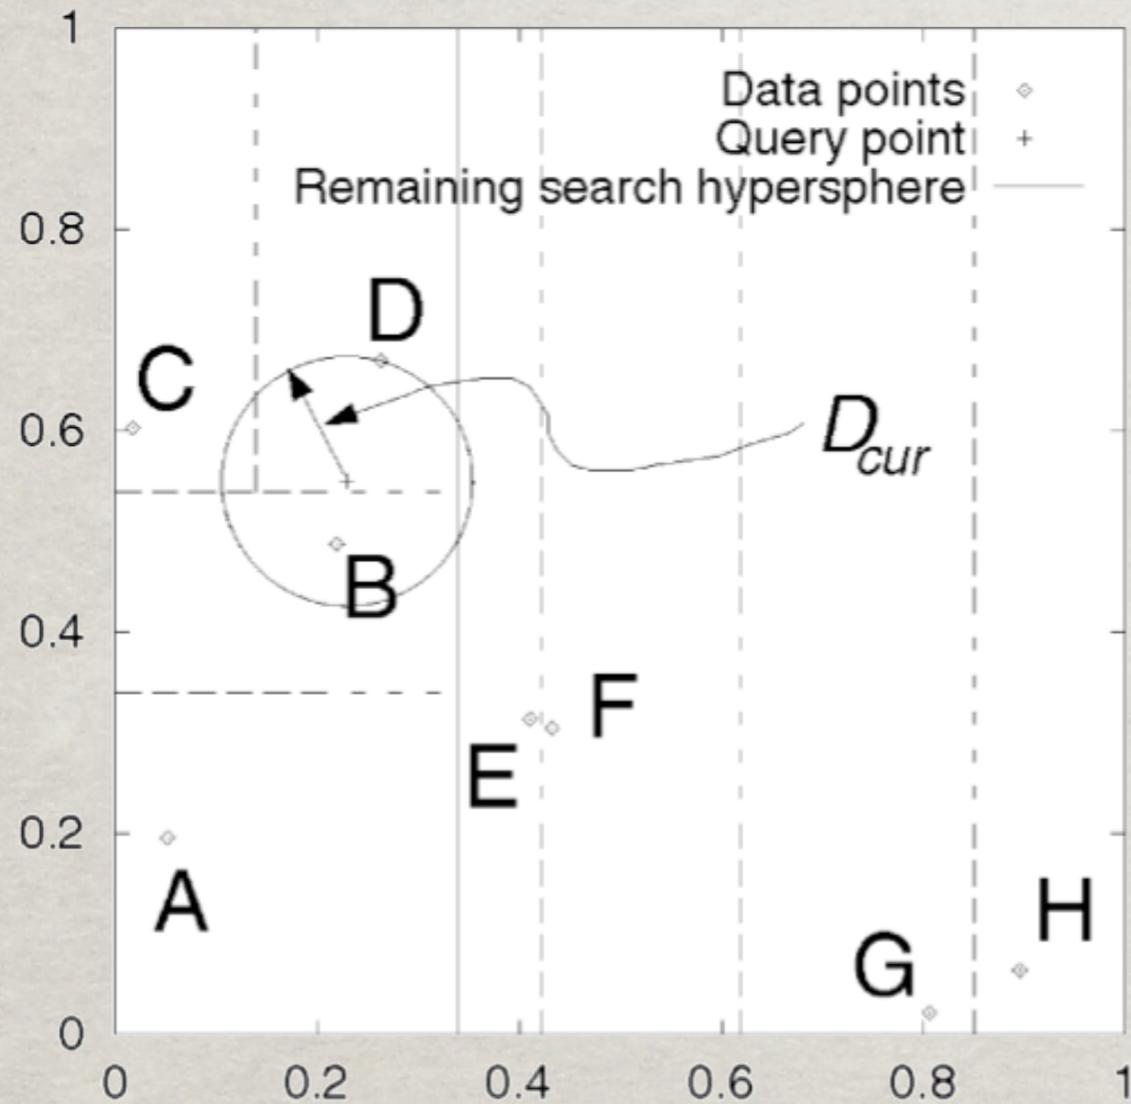
✱ Depth-first search works poorly in high dimensional spaces, and thus Beis&Lowe suggest a best-first search instead.

✱ Algorithm:

1. At each node, store the distance $e_n = q[k_n] - m_n$ in a *priority queue*. Always insert lowest value first.

2. Go down alternate branch of the first node in the queue if $e_n < e_{curr}$

BEST-BIN-FIRST



BALL TREES

- ✻ Omohundro 89, *Metric Tree* Uhlmann 91
- ✻ *A range search* method:
find all neighbours within a distance ϵ
from query vector
- ✻ Each node in tree has a centre \mathbf{p} , and a radius r

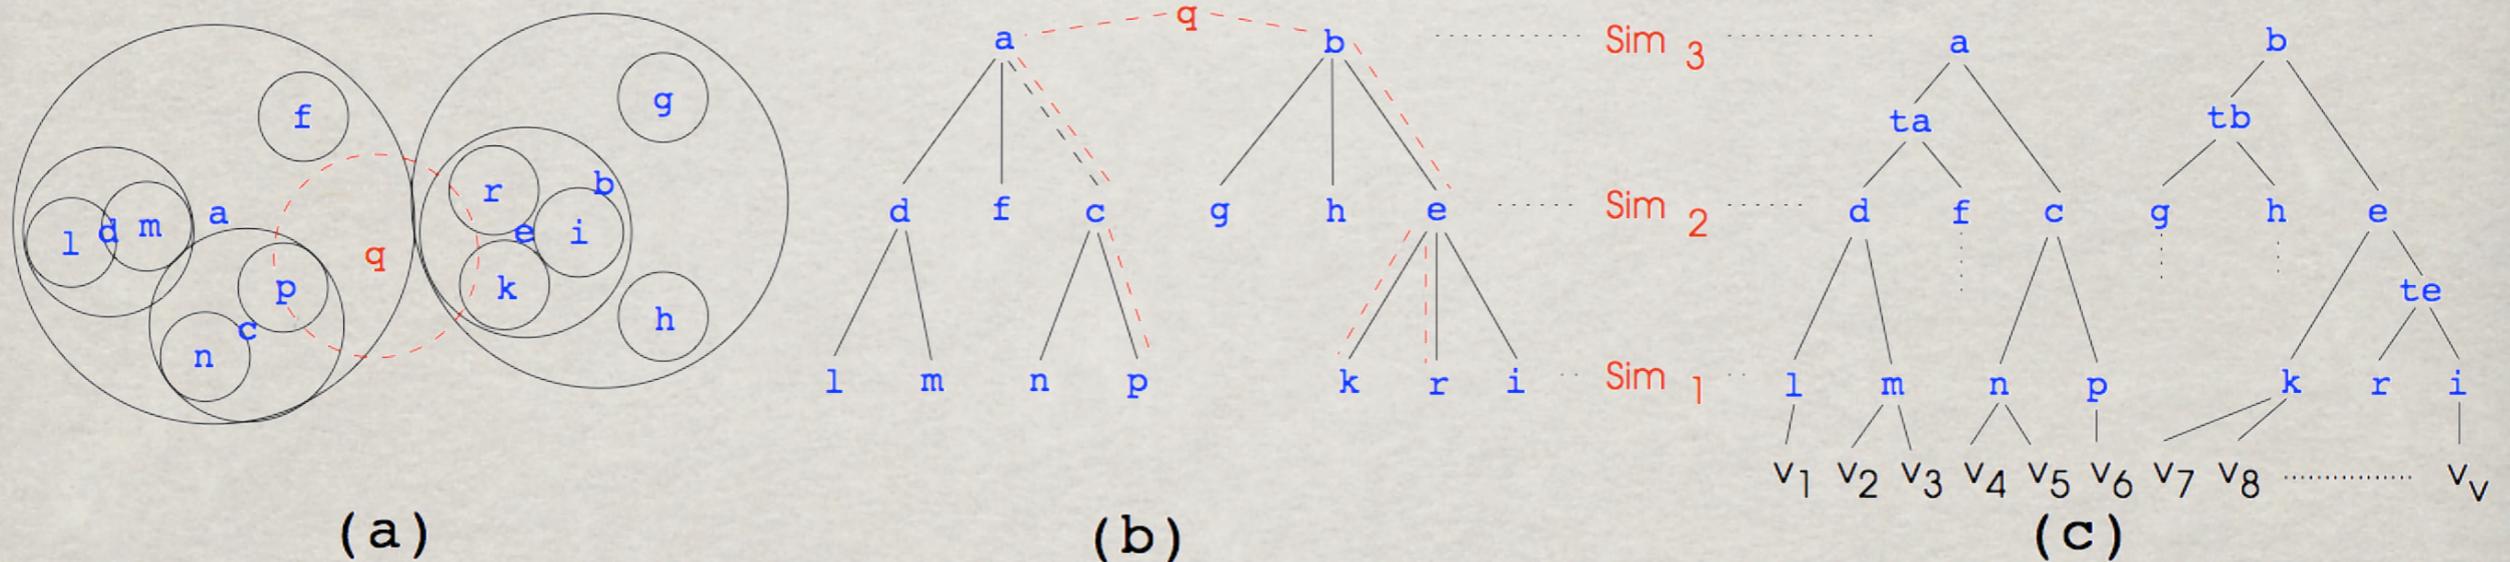
 \mathbf{p} is average of all leaves
 r is maximum distance from \mathbf{p} to a leaf

BALL TREES

- ✻ An optimal ball tree is constructed bottom up. Very expensive. E.g. using *agglomerative clustering*:
 1. Set each sample to be one cluster
 2. Merge the two most similar clusters
 3. Repeat step 2 until no clusters are left.
- ✻ Agglomerative clustering generates a *dendrogram*, or similarity tree. This can be pruned using various heuristics to form ball tree.

BALL TREES

☀ Example of a search:



Leibe&Mikolajczyk&Schiele, BMVC'06

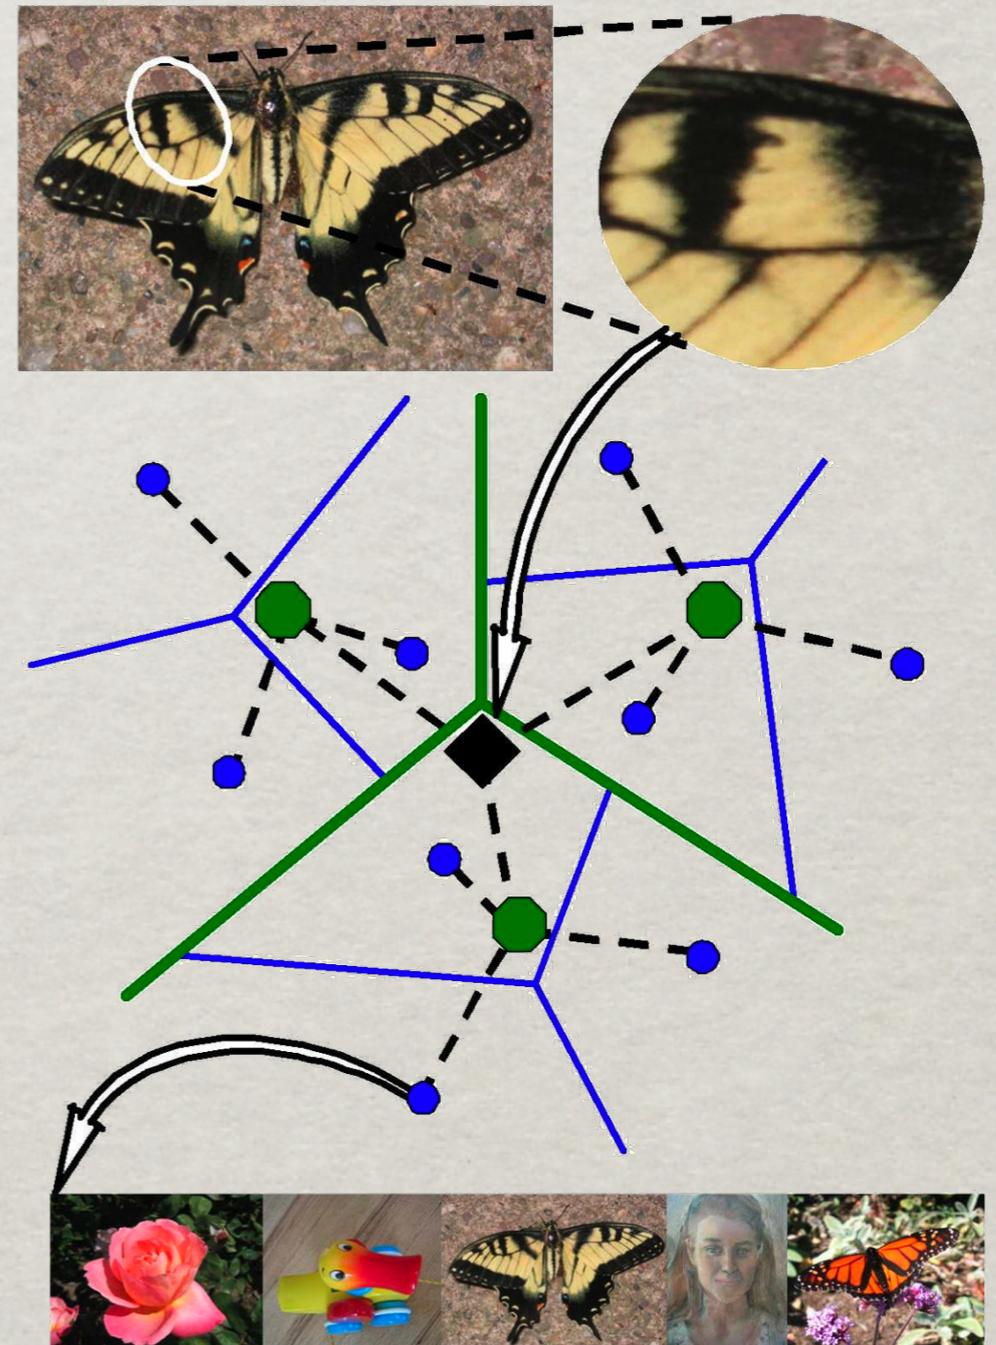
☀ At each node, the distances to circle centres are computed, and compared to the radius.

BALL TREES

- ✱ Advantage: Good if range search is needed.
I.e. find all neighbours with $d < d_{\max}$
- ✱ Disadvantages:
 - ✱ Tree construction algorithm does not scale to very large datasets
 - ✱ A ball in \mathbb{R}^D is not such a useful region shape if sample density varies in the feature space.

K-MEANS TREE

- David Nistér and Henrik Stewénus, *Scalable Recognition with a Vocabulary Tree*, CVPR06
- Hierarchical modification of the visual words idea from LE5



K-MEANS TREE

✻ Building the tree:

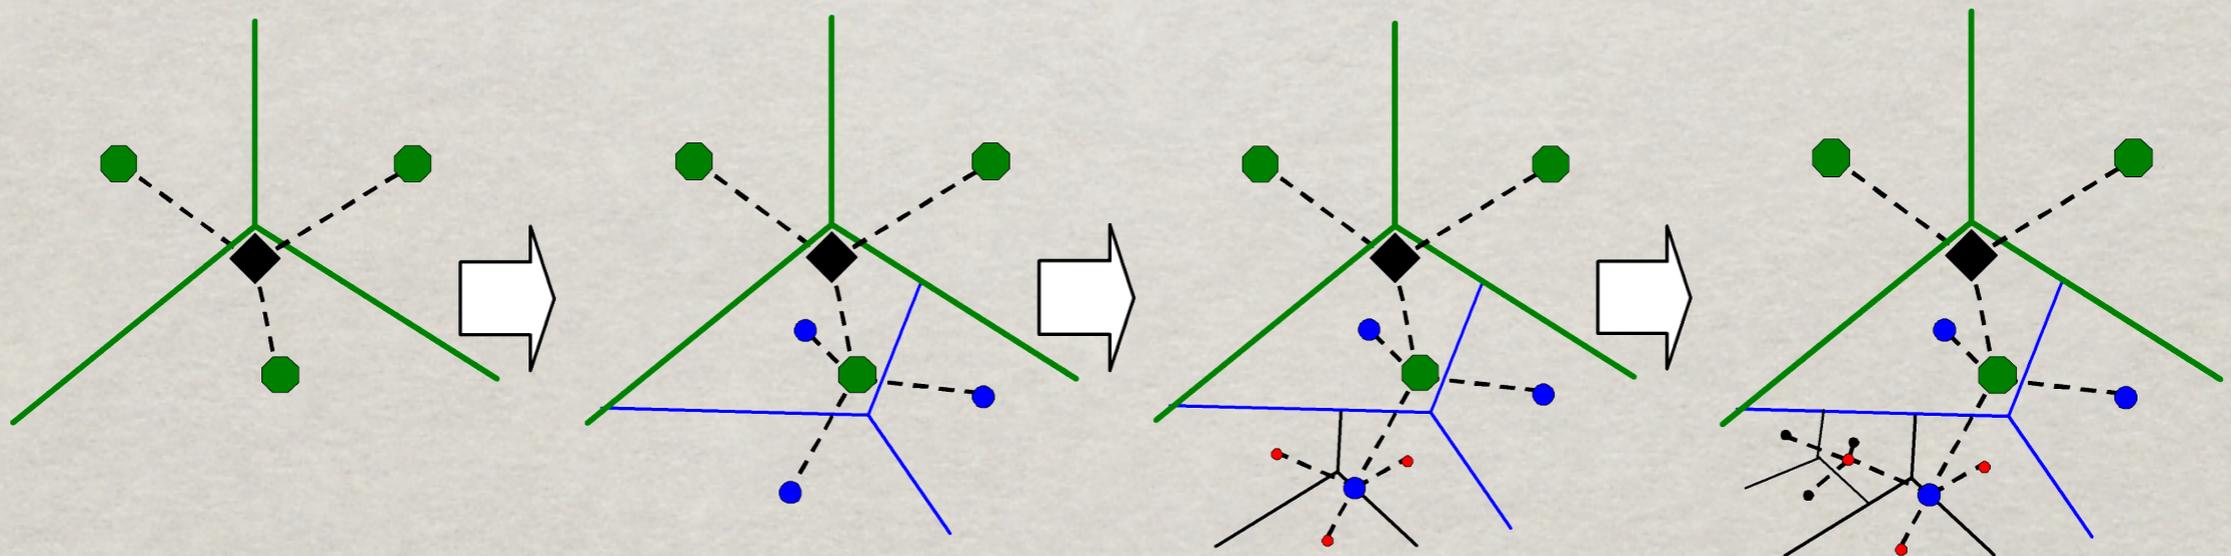
1. Run K-means with e.g. $K=10$ on whole dataset.
2. Partition dataset into K subsets using Voronoi regions
3. Apply algorithm recursively on subsets.

✻ The tree gets branching factor K .

K-MEANS TREE

☼ Using the tree:

1. Compare query vector to prototypes at current level.
2. Go down best branch

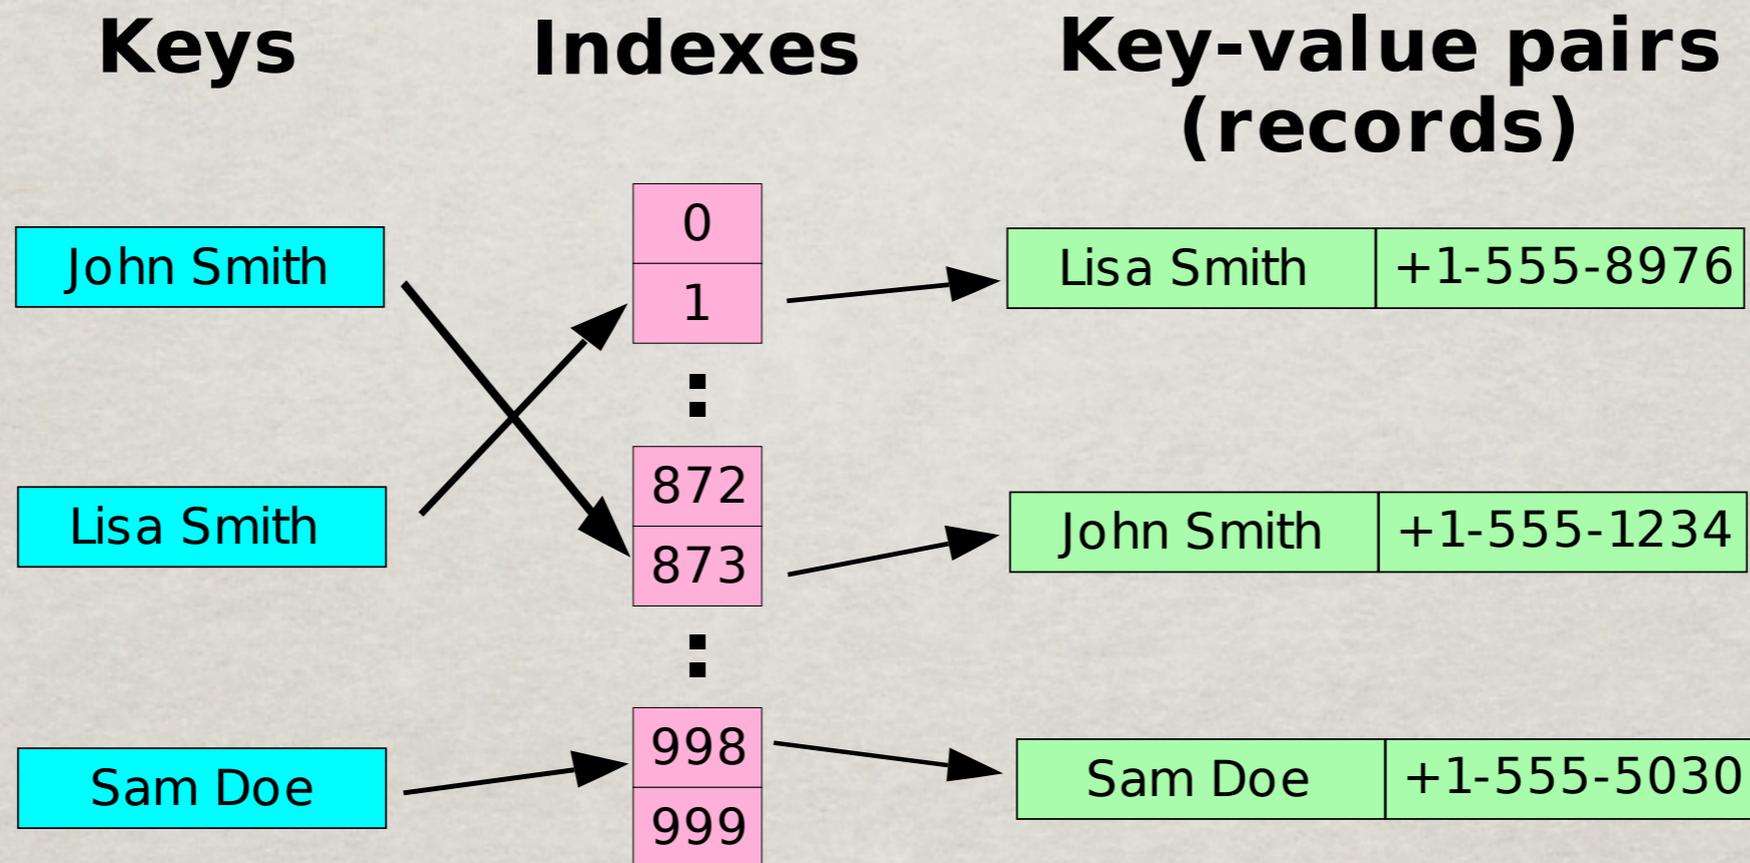


K-MEANS TREE

- ✱ Used to compute a TF-IDF bag-of-words vector quickly.
- ✱ Much faster than non-hierarchical visual words algorithm.
- ✱ As in the kD-tree, the terminal leaf node is a near neighbour.
- ✱ Not so successful as a near neighbour search, as prototypes are far from region edge in \mathbb{R}^D .

HASH TABLES

- ☼ An efficient way to perform lookup.



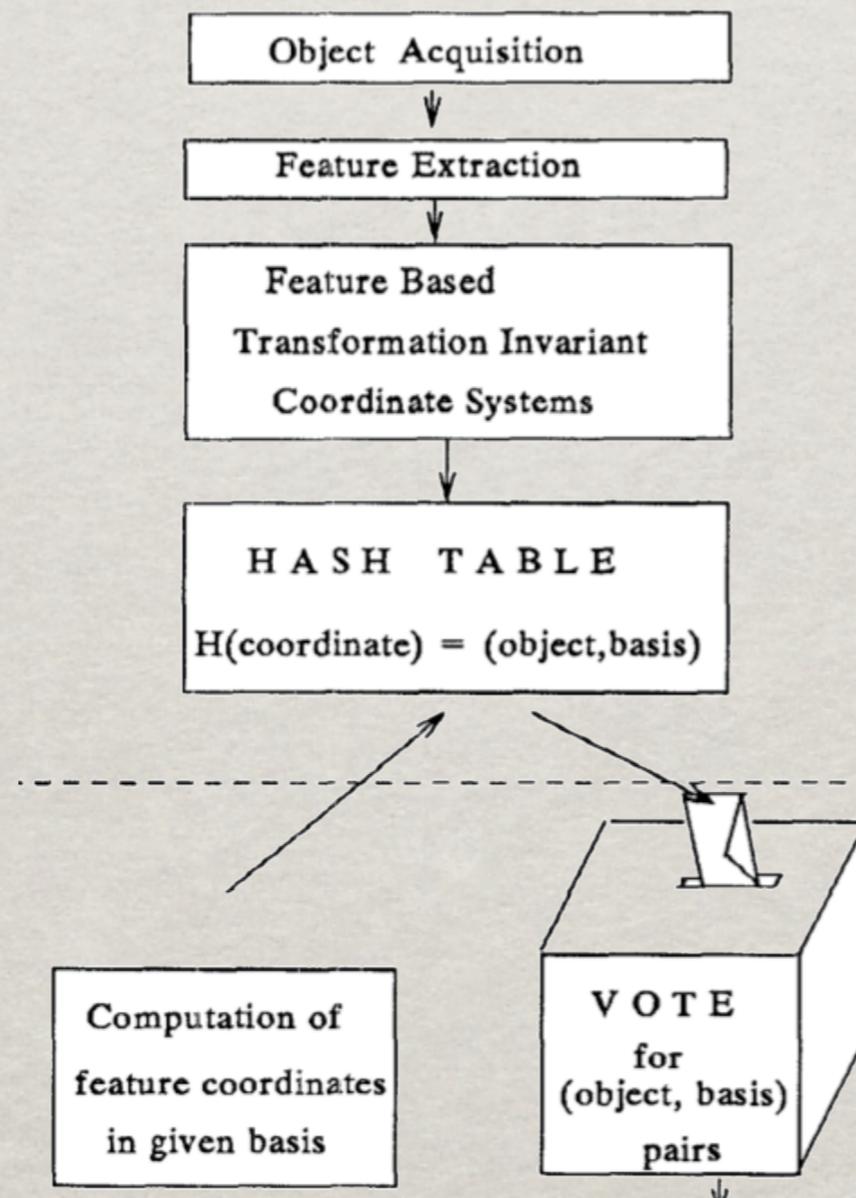
- ☼ Each key is converted to an index using a *hashing function*: $\text{index} = H(\text{key})$

HASH TABLES

- ✱ Lookup is $O(1)$ instead of e.g. $O(N)$ in a list, $O(\log N)$ in a sorted list/tree etc.
- ✱ Collisions can happen. i.e. different keys get the same index. Solved e.g. using *chaining* (linked lists), or *linear probing* (insertion at next free slot).
- ✱ Linear probing typically wants a $<80\%$ filled table.
- ✱ Hashing has poor cache locality.

GEOMETRIC HASHING

✻ Introduced in Lamdan & Wolfson ICCV'88

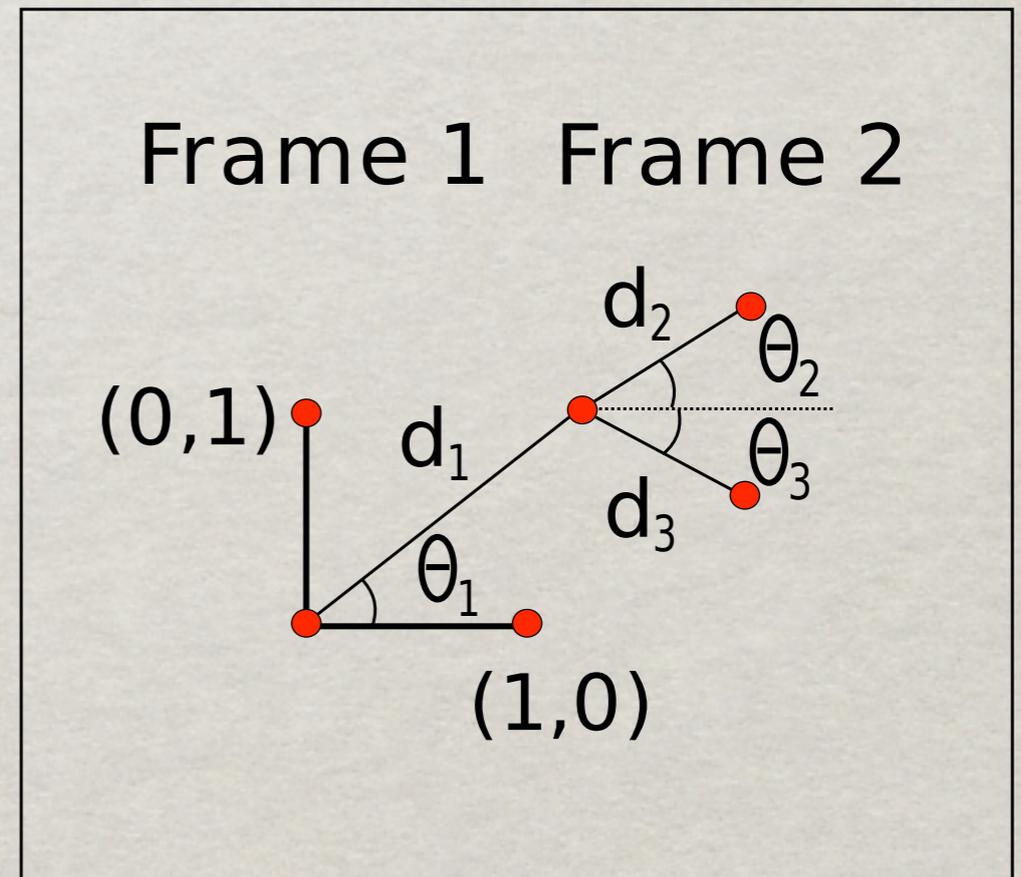


GEOMETRIC HASHING

✱ Modern example: Used for matching frames *without descriptors* by Chum & Matas, *Geometric Hashing with Local Affine Frames*, CVPR'06

✱ Use pairs of affine frames. Express frame 2 in frame 1. 25 bins for angle 16 for d_1 , 6 for d_2 & d_3

✱ $9 \cdot 10^6$ unique values for key to hash.



GEOMETRIC HASHING

- ✻ Design of $H(\text{key})$ is not discussed further.
- ✻ Hash tables suffer from the same basic problem as trees: Neighbouring bins might contain the closest match. More neighbours in high dimensional spaces.
- ✻ To deal with the neighbour problem, Chum&Matas construct 6 different tables (for 6 different frame constructions) and run them in parallel.

PROJECTS

- ✻ Course is 8hp:
 - 5hp for lectures+articles+exam
 - 3hp for project.
- ✻ Project part is 2 weeks of:
 - programming&research
 - writing a small report.

PROJECT SUGGESTIONS

1. Comparison of kD-tree and K-means tree in terms of speed and accuracy.
2. Implementation of a bag-of-words recognition system (using existing feature detector code). Test how system parameters affect result.
3. Implementation and test of a new descriptor for a given detector.

PROJECT SUGGESTIONS

4. Implement and test a voting scheme for global geometric deformation (e.g. similarity, or affine transform) of feature locations.
5. Learn a matching metric, and compare it to least squares matching.
6. Compare Chi^2 , EMD and least-squares on a problem of choice.
7. Your own suggestion.

EXAM

Everyone should bring calendar next week, so we can decide on a date for the written exam.

DISCUSSION

✻ Questions/comments on paper and lecture.