

Shape Descriptors II

Thomas Funkhouser
CS597D, Fall 2003
Princeton University



Taxonomy of Shape Descriptors



Structural representations

- Skeletons
- Part-based methods
- Feature-based methods

Statistical representations

- Voxels, moments, wavelets, ...
- Attributes, histograms, ...
- Point descriptors

Statistical Shape Descriptors



Alignment-dependent

- Voxels
- Wavelets
- Moments
- Extended Gaussian Image
- Spherical Extent Function
- Spherical Attribute Image

Alignment-independent

- Shape histograms
- Harmonic descriptor
- Shape distributions

Statistical Shape Descriptors



Alignment-dependent

- Voxels
- Wavelets
- Moments
- Extended Gaussian Image
- Spherical Extent Function
- Spherical Attribute Image

Alignment-independent

- Shape histograms
- Harmonic descriptor
- Shape distributions

Alignment

Translation (*Center of Mass*)

$$c = \frac{1}{n} \sum_{i=1}^n p_i$$

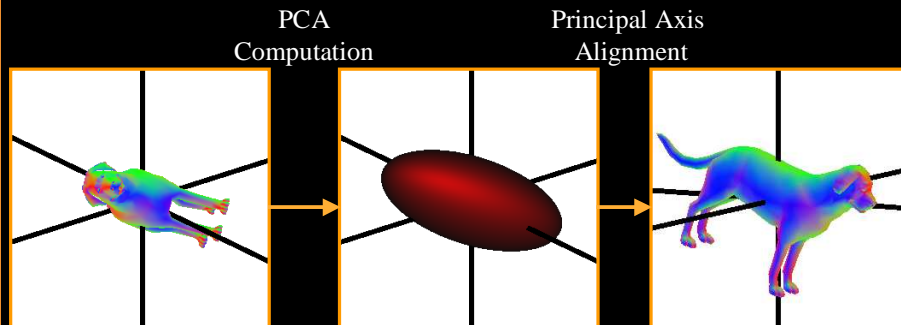
Scale (*Radial Deviation*)

$$s = \sqrt{\frac{1}{n} \sum_{i=1}^n \|p_i\|^2}$$

Alignment

Rotation (*PCA*)

- *Principal axes are eigenvectors associated with largest eigenvalues of 2nd order moments covariance matrix*




Alignment



Rotation (PCA)

- *Principal axes are eigenvectors associated with largest eigenvalues of 2nd order moments covariance matrix*




Not very robust!

Alignment



Mirror

- *PCA does not give directions for principal axes*




Need heuristics to determine positive axes!

Alignment-Independent Descriptors



Observation: it is difficult to normalize for differences in rotation and mirroring




Three mugs aligned automatically with PCA

Motivation: build a shape descriptor that is invariant to rotations and mirrors and as discriminating as possible

Shape Histograms

Image courtesy of Ankerst et al, 1999

Shape descriptor stores histogram of how much surface resides at different radii from center of mass




Shape Histograms (shells)
(Ankerst, 1999)

Shape Histograms

Image courtesy of
Misha Kazhdan

Shape descriptor stores histogram of how much surface resides at different radii from center of mass



Shape Histogram Experiment

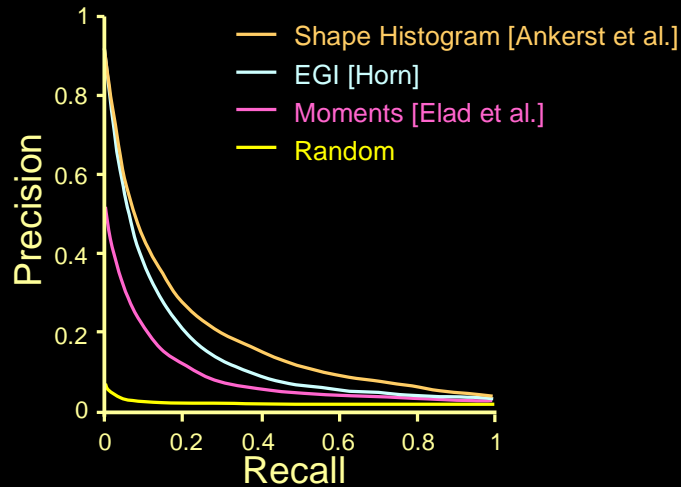
Test database is Viewpoint household collection
1,890 models, 85 classes



Shape Histogram Retrieval Results



Precision-recall curves (mean for all queries)



Shape Histograms



Properties


- ü Insensitive to noise
- ü Insensitive to topology
- ü Robust to degeneracies
- ü Quick to compute
- ü Efficient to match
- ü Concise to store
- ü Invariant to rotations
- Discriminating?



Harmonic Shape Descriptor


Key idea:

- Decompose each sphere into irreducible set of rotation independent components
- Store "how much" of the model resides in each component



Step 1: Normalization

Normalize for translation and scale




3D Model

Step 2: Voxelization



Rasterize polygon surfaces into 3D voxel grid




3D Voxel Grid

Step 3: Spherical Decomposition



Intersect with concentric spheres



Spherical Functions

Step 4: Frequency Decomposition



Represent each spherical function as a sum of harmonic frequencies (orders)




Spherical Functions

Step 4: Frequency Decomposition



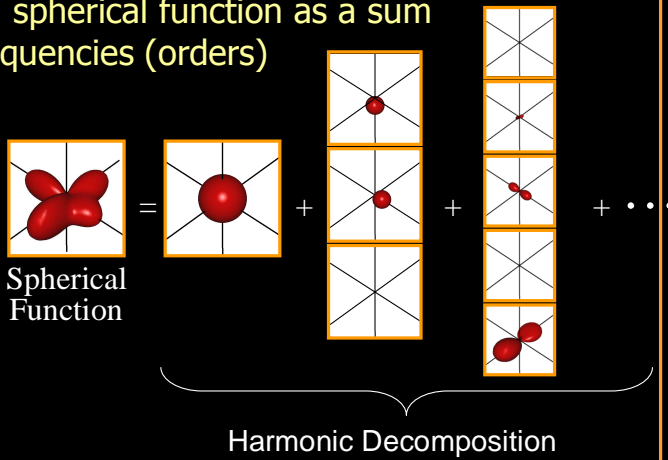
Represent each spherical function as a sum of harmonic frequencies (orders)



Spherical Functions

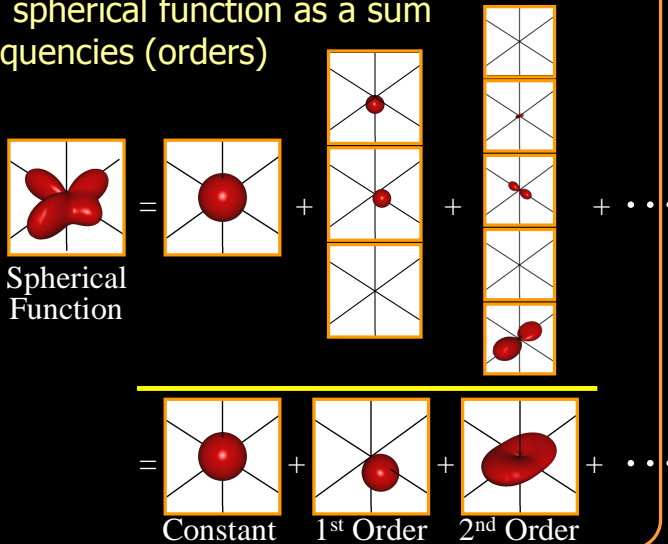
Step 4: Frequency Decomposition

Represent each spherical function as a sum of harmonic frequencies (orders)



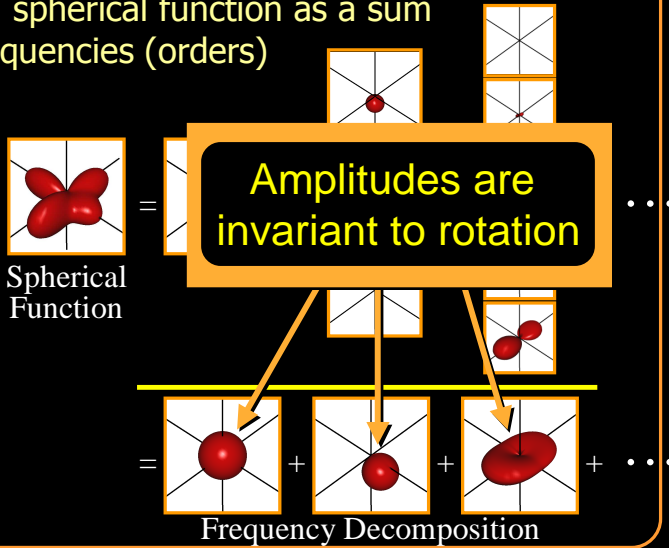
Step 4: Frequency Decomposition

Represent each spherical function as a sum of harmonic frequencies (orders)




Step 4: Frequency Decomposition

Represent each spherical function as a sum of harmonic frequencies (orders)



Step 5: Amplitude Computation

Store "how much" (L_2 -norm) of the shape resides in each harmonic frequency of each sphere



Harmonic Shape Descriptor

Matching Harmonic Descriptors

Define similarity as L_2 -distance between descriptors


- Enables nearest neighbor indexing and fast search
- Provides lower bound for L_2 -distance between models

$$\text{Sim} \left[\begin{array}{c} \text{Car Model 1} \\ \text{Car Model 2} \end{array} \right] = \left\| \begin{array}{c} \text{Harmonic Descriptor 1} \\ \text{Harmonic Descriptor 2} \end{array} \right\| \leq \left\{ \begin{array}{l} \text{Car Model 1} - \text{Car Model 3} \\ \text{Car Model 1} - \text{Car Model 4} \\ \text{Car Model 1} - \text{Car Model 5} \\ \vdots \end{array} \right.$$

Harmonic Shape Descriptor

Properties

- **Ø Concise to store?**
- Quick to compute?
- Insensitive to noise?
- Insensitive to topology?
- Robust to degeneracies?
- Invariant to transforms?
- Efficient to match?
- Discriminating?



2048 bytes per model
(16 frequencies x 32 radii x 4 bytes)

Harmonic Shape Descriptor



Properties

- ü Concise to store
- Ø Quick to compute?
 - Insensitive to noise?
 - Insensitive to topology?
 - Robust to degeneracies?
 - Invariant to transforms?
 - Efficient to match?
 - Discriminating?

1.6 seconds (on average)




Polygons




Voxels



Spherical Decomposition



Frequency Decomposition



Harmonic Shape Descriptor

Harmonic Shape Descriptor



Properties

- ü Concise to store
- Ø Quick to compute?
 - Insensitive to noise?
 - Insensitive to topology?
 - Robust to degeneracies?
 - Invariant to transforms?
 - Efficient to match?
 - Discriminating?

1.6 seconds (on average)




Polygons




Voxels



Spherical Decomposition



Frequency Decomposition




Harmonic Shape Descriptor

Harmonic Shape Descriptor



Properties

- ü Concise to store
- ü Quick to compute
- Ø Insensitive to noise
- Ø Insensitive to topology
- Ø Robust to degeneracies
- Invariant to transforms?
- Efficient to match?
- Discriminating?



Rasterize polygon surfaces
(no solid reconstruction)

Harmonic Shape Descriptor



Properties

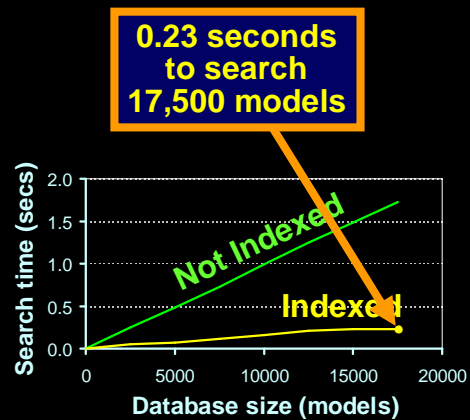
- ü Concise to store
 - ü Quick to compute
 - ü Insensitive to noise
 - ü Insensitive to topology
 - ü Robust to degeneracies
 - Ø Invariant to transforms
 - Efficient to match?
 - Discriminating?
- ü Rotation
 - ü Mirror
 - ü Translation (w/ normalization)
 - ü Scale (w/ normalization)

Harmonic Shape Descriptor



Properties

- ü Concise to store
- ü Quick to compute
- ü Insensitive to noise
- ü Insensitive to topology
- ü Robust to degeneracies
- ü Invariant to transforms
- Ø Efficient to match?
- Discriminating?



Harmonic Shape Descriptor



Properties

- ü Concise to store
- ü Quick to compute
- ü Insensitive to noise
- ü Insensitive to topology
- ü Robust to degeneracies
- ü Invariant to transforms
- ü Efficient to match?
- Ø Discriminating?


Harmonic Matching Results




Test database is Viewpoint household collection
1,890 models, 85 classes



153 dining chairs



25 livingroom chairs




16 beds



12 dining tables




8 chests



28 bottles



39 vases

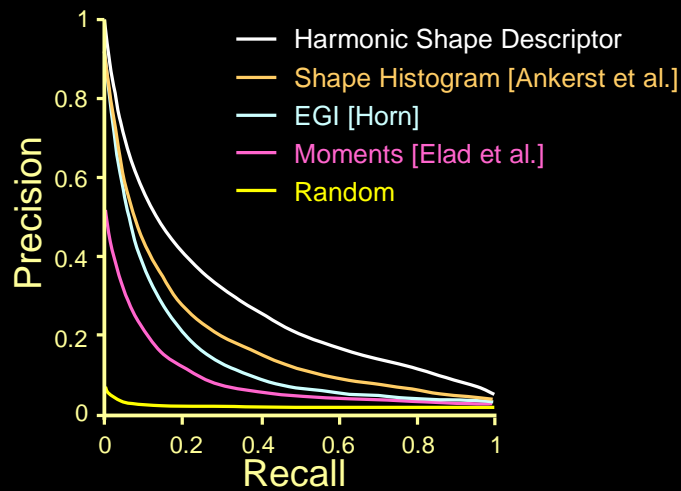


36 end tables

Harmonic Retrieval Results



Precision-recall curves (mean for all queries)



Statistical Shape Descriptors



Alignment-dependent

- Voxels
- Wavelets
- Moments
- Extended Gaussian Image
- Spherical Extent Function
- Spherical Attribute Image

Alignment-independent

- Shape histograms
- Harmonic descriptor
- \emptyset Shape distributions

Shape Distributions



Motivation: general approach to finding a common parameterization for matching



Audio



3D Surface



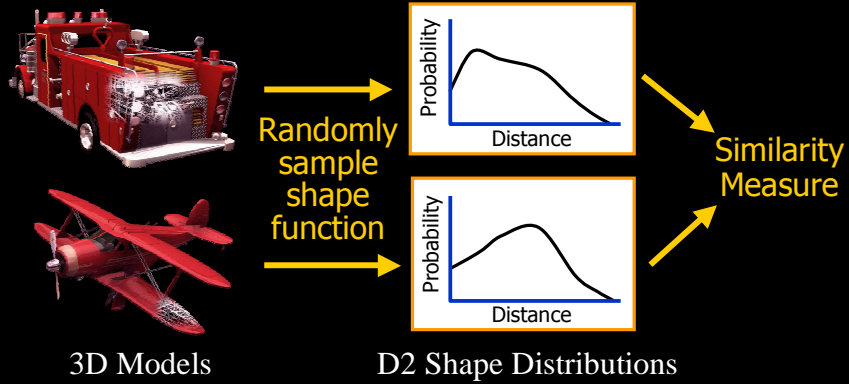
2D Contour



3D Volume


Shape Distributions

Key idea: map 3D surfaces to common parameterization by randomly sampling shape function



Which Shape Function?

Implementation: simple shape functions based on angles, distances, areas, and volumes



[Ankerst 99]

D2 Shape Distribution



Properties


- Concise to store?
- Quick to compute?
- Invariant to transforms?
- Efficient to match?
- Insensitive to noise?
- Insensitive to topology?
- Robust to degeneracies?
- Discriminating?

D2 Shape Distribution



Properties

- Ø Concise to store?
- Ø Quick to compute?
- Invariant to transforms?
- Efficient to match?
- Insensitive to noise?
- Insensitive to topology?
- Robust to degeneracies?
- Discriminating?



512 bytes (64 values)


0.5 seconds (10^6 samples)

D2 Shape Distribution



Properties

- ü Concise to store
- ü Quick to compute
- Ø **Invariant to transforms?**
 - ü Translation
 - ü Rotation
 - ü Mirror
 - ü Scale (w/ normalization)
- Efficient to match?
- Insensitive to noise?
- Insensitive to topology?
- Robust to degeneracies?
- Discriminating?




Normalized Means

D2 Shape Distribution



Properties

- ü Concise to store
- ü Quick to compute
- ü Invariant to transforms
- Ø **Efficient to match?**
 - Insensitive to noise?
 - Insensitive to topology?
 - Robust to degeneracies?
 - Discriminating?




D2 Shape Distribution



Properties

- ü Concise to store
- ü Quick to compute
- ü Invariant to transforms
- ü Efficient to match
- Ø Insensitive to noise?
- Ø Insensitive to topology?
- Ø Robust to degeneracies?
- Discriminating?



D2 Shape Distribution



Properties

- ü Concise to store
- ü Quick to compute
- ü Invariant to transforms
- ü Efficient to match
- ü Insensitive to noise
- ü Insensitive to topology
- ü Robust to degeneracies
- Ø Discriminating?

D2 Shape Distribution Results



Question

- How discriminating are D2 shape distributions?

Test database

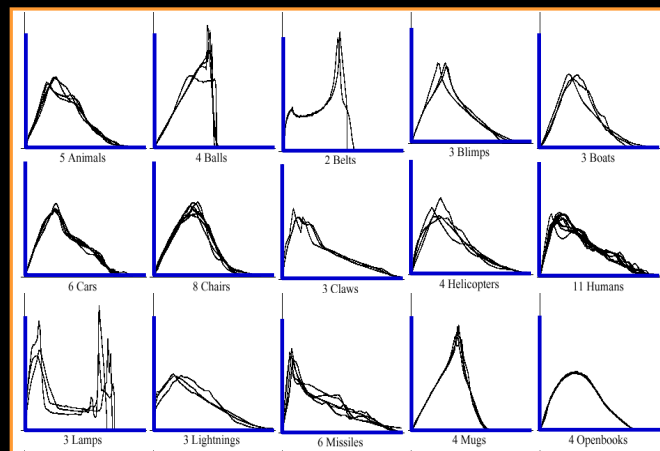
- 133 polygonal models
- 25 classes



D2 Shape Distribution Results




D2 distributions are different across classes



D2 shape distributions for 15 classes of objects

D2 Shape Distribution Results



D2 distributions for 5 tanks (gray) and 6 cars (black)

D2 Shape Distribution Results




Similarity Matrix

- Darkness represents similarity

Blocks

- Tanks, cars
- Airplanes
- Humans
- Helicopters



D2 Retrieval Experiment




Test database is Viewpoint household collection
1,890 models, 85 classes



153 dining chairs



25 livingroom chairs



16 beds



12 dining tables




8 chests



28 bottles



39 vases

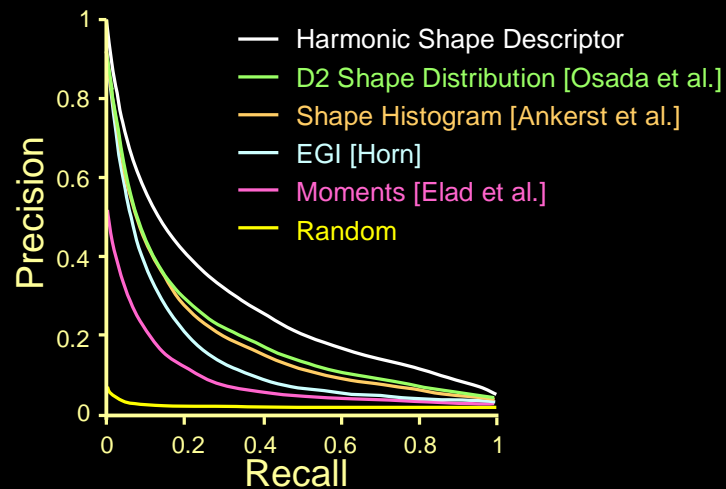


36 end tables

D2 Retrieval Results



Precision-recall curves (mean for all queries)

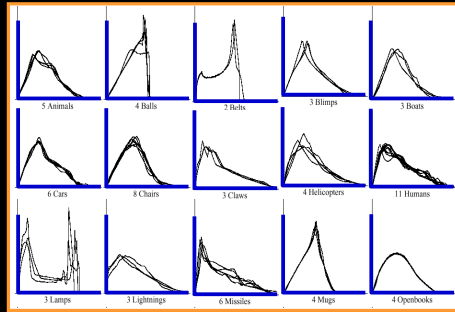


Shape Distributions



Next steps:

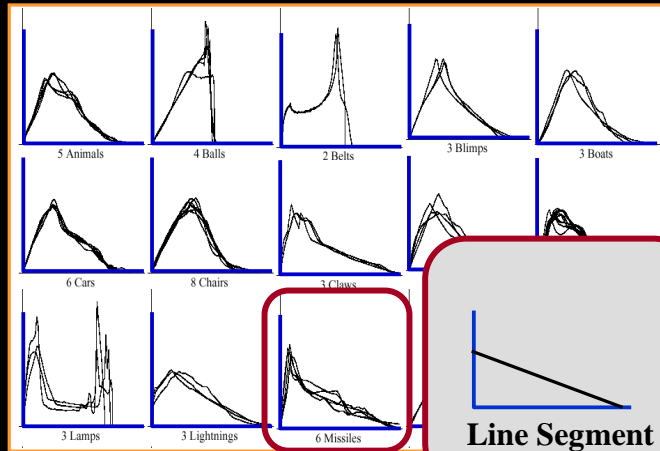
- Better shape functions
- Better comparison methods
- Analysis apps



D2 Shape Distribution Results



Recognizing gross shapes with D2 distributions

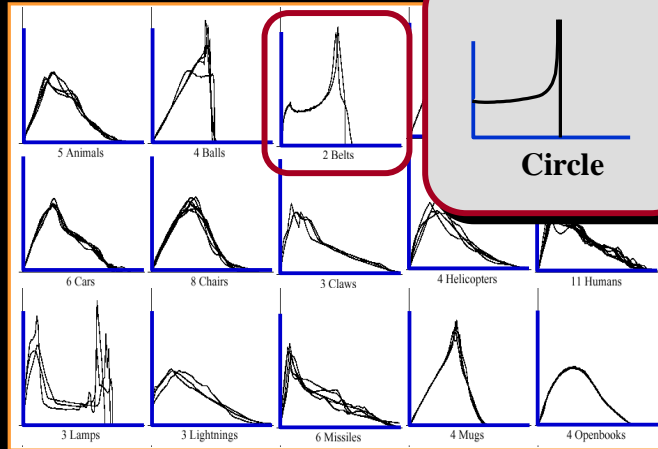


D2 shape distributions for 15 c...

D2 Shape Distribution Results



Recognizing gross shapes with D2 distributions

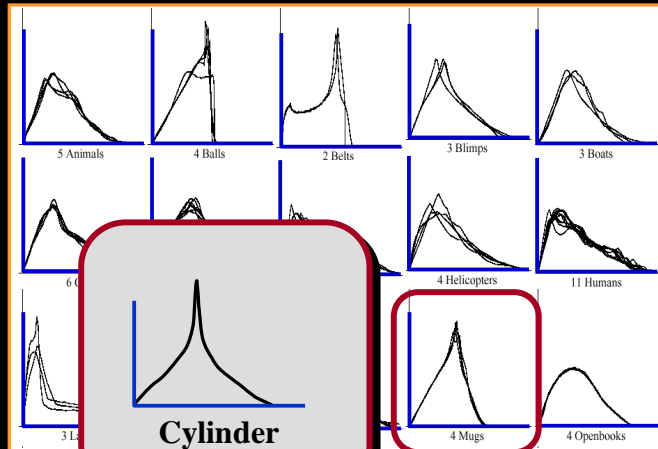


D2 shape distributions for 15 classes of objects

D2 Shape Distribution Results



Recognizing gross shapes with D2 distributions

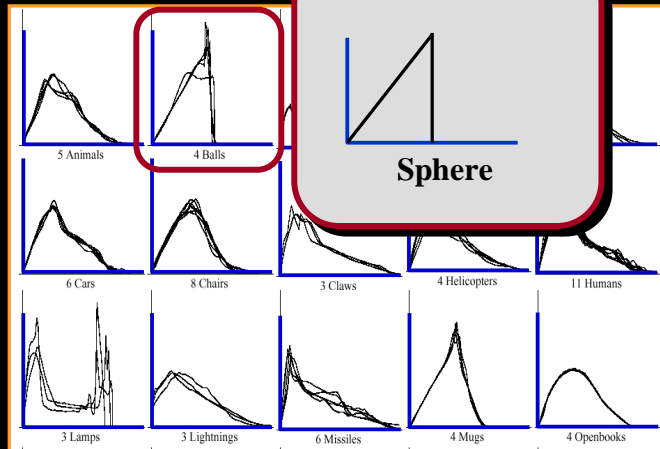


D2 shape distributions for 15 classes of objects

D2 Shape Distribution Results



Recognizing gross shapes with D2 distributions




D2 shape distributions for 15 classes of objects

D2 Shape Distribution Results



Recognizing gross shapes with D2 distributions



D2 shape distributions for 15 classes of objects

Taxonomy of Shape Descriptors



Structural representations

- Skeletons
- Part-based methods
- Feature-based methods

Statistical representations

- Voxels, moments, wavelets, ...
- Attributes, histograms, ...
- Point descriptors

Taxonomy of Shape Descriptors



Structural representations

- Skeletons
- Part-based methods
- Feature-based methods

Statistical representations

- Voxels, moments, wavelets, ...
- Attributes, histograms, ...
- **Point descriptors** ← **Next Time!**