# Recognizing Deformable Shapes

Salvador Ruiz Correa
(CSE/EE576 Computer Vision I)

# Goal

- We are interested in developing algorithms for recognizing and classifying deformable object shapes from range data.

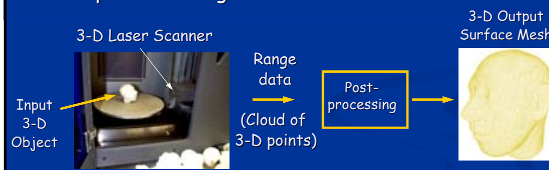


- This is a difficult problem that is relevant in several application fields.

# Applications

- Computer Vision:
  - Scene analysis
  - Industrial Inspection
  - Robotics
- Medical Diagnosis:
  - Classification and
  - Detection of craniofacial deformations.

# Basic Idea

- Generalize existing numeric surface representations for matching 3-D objects to the problem of identifying shape classes.

# Main Contribution

- An algorithmic framework based on symbolic shape descriptors that are robust to deformations as opposed to numeric descriptors that are often tied to specific shapes.

# What Kind Of Deformations?

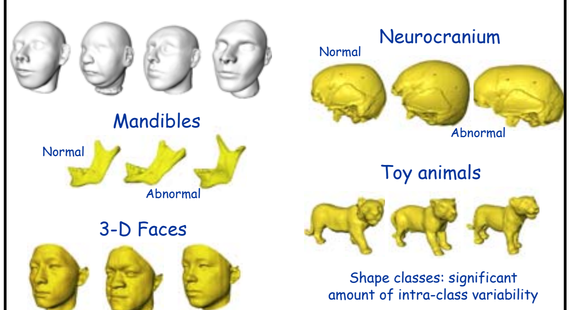

## Deformed Infants' Skulls

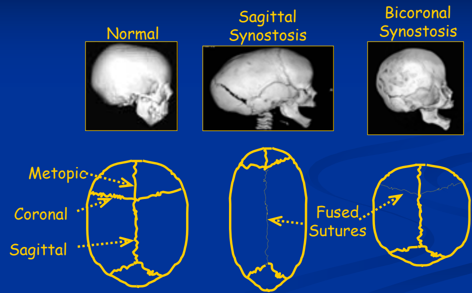


Occurs when sutures of the cranium fuse prematurely (synostosis).

## More Craniofacial Deformations

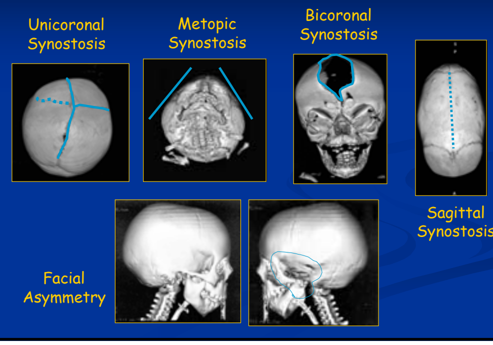


## Alignment-verification

- Find correspondences using numeric signature information.
- Estimate candidate transformations.

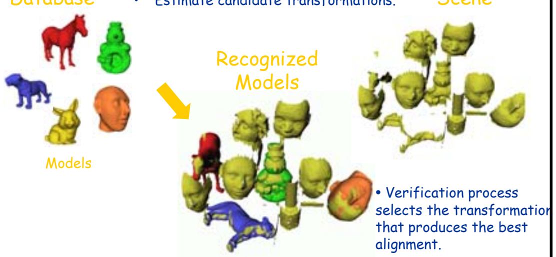


- Verification process selects the transformation that produces the best alignment.

## Related Literature (1)

- This approach has been used very successfully in industrial machine vision. Relevant investigations that use numeric signature representations for matching include:

- Splash representation – Stein and Medioni (IEEE PAMI, 1992)
- Spin image representation – Johnson and Hebert (IEEE PAMI, 1999).

## Related Literature (2)

- Spherical signatures – Ruiz-Correa et al. ( IEEE CVPR 2001).
- Shape distributions – Osada et al. (SMI, 2001,2002).
- Reflective symmetry descriptors – Kazhdan et al. (Algorithmica 2003).

## Alignment-Verification Limitations

The approach does not extend well to the problem of identifying classes of similar shapes. In general:

- Numeric shape representations are not robust to deformations.
- There are not exact correspondences between model and scene.
- Objects in a shape class do not align.

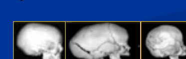

## Component-Based Methodology

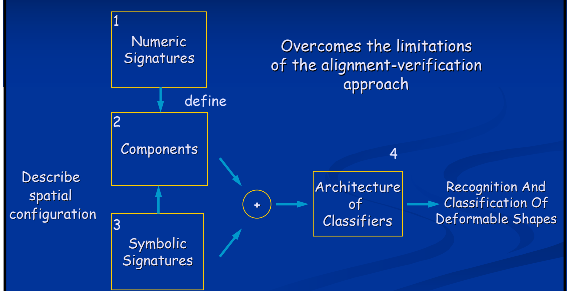


## Efficient Object Recognition (1)

- Developed spherical spin image representation (SSI) : computational complexity O(ms).
- Standard spin image (SI): computational complexity O(nms) (n~10³, m~10⁴, s~10³).
- Developed compressed SSI representation that requires O(mk) floats, k~40.
- Standard SIC (PCA) algorithm also requires O(md) floats but the proportionality constant is ~10 bigger.

## Efficient Object Recognition (3)

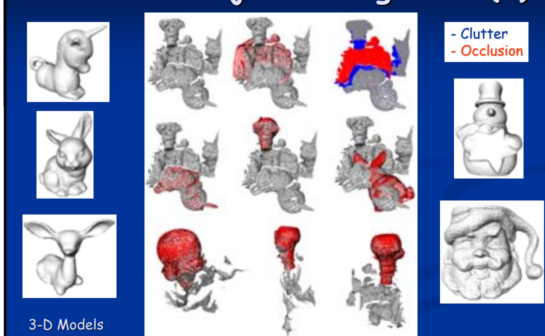


3-D Models

## Efficient Object Recognition (2)

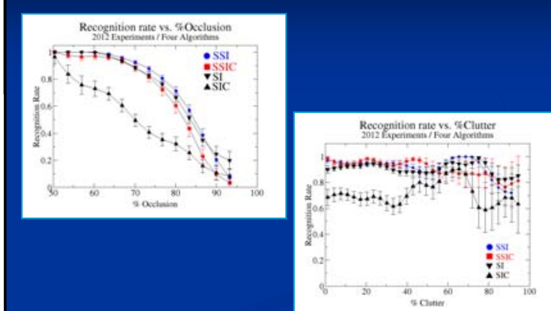


## Outline

- Mathematical Background.
- Formalize recognition and classification problems.
- Approach and implementation.
- Experimental validation: recognition and classification experiments.
- Discuss future work.
- Conclude.

## SVMs - Geometry (1)

Input Space

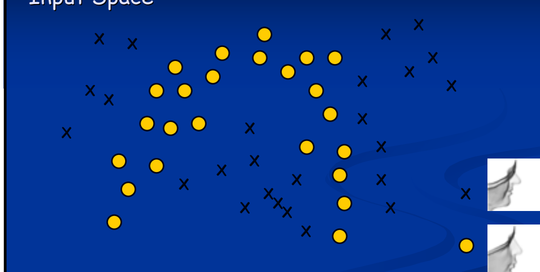

Separate classes using a hyperplane

## SVMs (2)

Feature Space

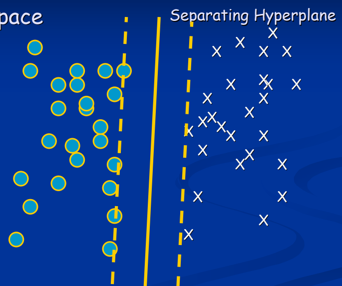


Plane is a function of the inner product of the points in feature space

## SVMs (3)

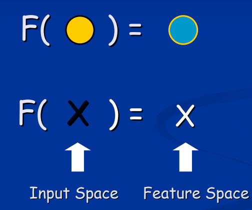


## SVMs – Kernel Trick (4)

F( Y ) = Z

F( Y′ ) = Z′

Input Space  Feature Space

<F( Y ), F( Y )> = <Z,Z′> = K(Y,Y′)

K, kernel function

## SVMs – Kernel Trick (5)

<F( Y ), F( Y )> = <Z,Z′> = K(Y,Y′)

K, kernel function example:

$K(Y,Y') = \exp(-g^2||Y-Y'||^2)$

## SVMs (2)

Feature Space

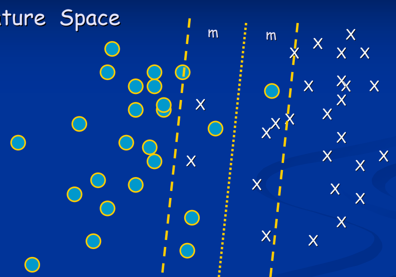


SVMs successful even in non-separable cases

## Recognition Problem (1)

- We are given a set of surface meshes $\{C_1, C_2, \ldots, C_n\}$ which are random samples of two *shape classes* ***C***

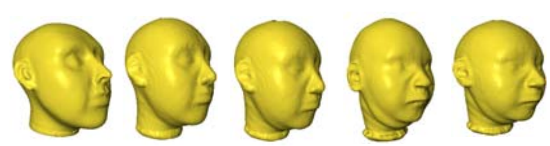

$C_1$ $C_2$ … $C_k$ … $C_n$

## Recognition Problem (2)

- The problem is to use the given meshes and labels to construct an algorithm that determines whether shape class members are present in a single view range scene.

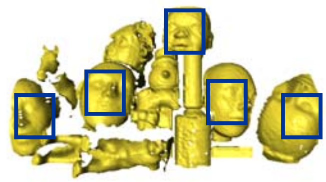

## Classification Problem (1)

- We are given a set of surface meshes $\{C_1, C_2, \ldots, C_n\}$ which are random samples of two *shape classes* $C^{+1}$ *and* $C^{-1}$,
- where each surface mesh is labeled either by +1 or -1.

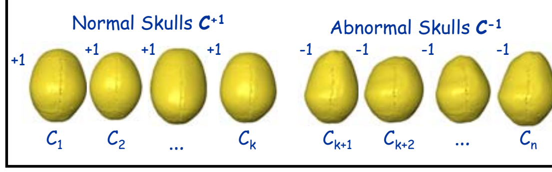


## Classification Problem (2)

- The problem is to use the given meshes and labels to construct an algorithm that predicts the label of a new surface mesh $C_{new}$.

$C_{new}$

Is this skull normal (+1) or abnormal (-1)?

## Classification Problem (3)

- We also consider the case of “missing” information:

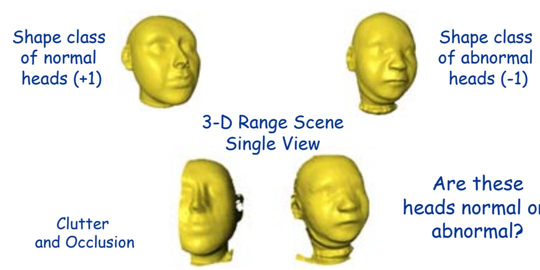


## Assumptions

- All shapes are represented as oriented surface meshes of fixed resolution.

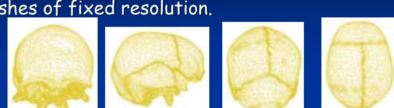

- The vertices of the meshes in the training set are in full correspondence.
- Finding full correspondences : hard problem yes … but it is approachable ( use morphable models technique: Blantz and Vetter, SIGGRAPH 99; C. R. Shelton, IJCV, 2000; Allen et al., SIGGRAPH 2003).

## Four Key Elements To Our Approach

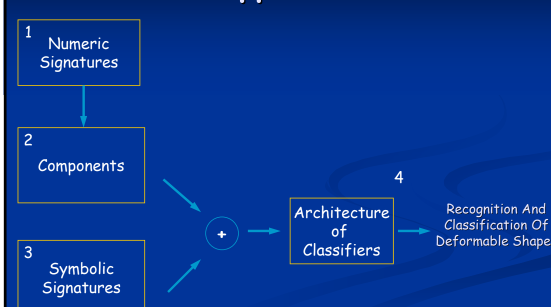

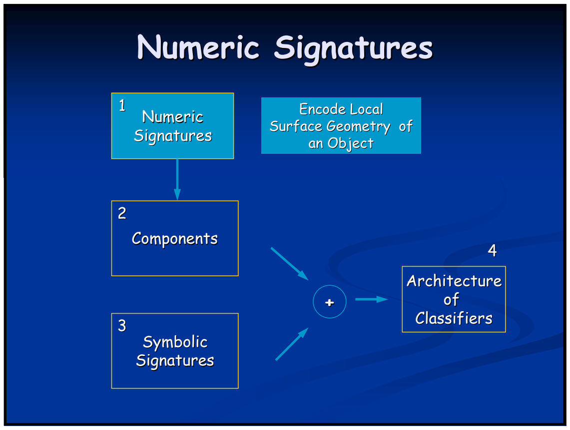
Numeric Signatures
1 Numeric Signatures
Encode Local Surface Geometry of an Object
2 Components
3 Symbolic Signatures
4 Architecture of Classifiers
+

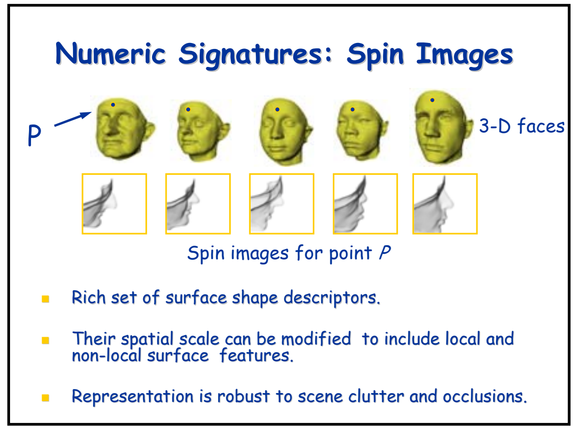
Numeric Signatures: Spin Images
P
3-D faces
Spin images for point P
Rich set of surface shape descriptors.
Their spatial scale can be modified to include local and non-local surface features.
Representation is robust to scene clutter and occlusions.

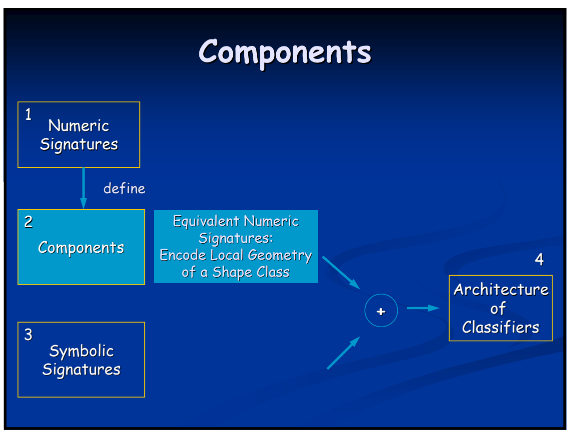
Components
1 Numeric Signatures
define
2 Components
Equivalent Numeric Signatures: Encode Local Geometry of a Shape Class
3 Symbolic Signatures
4 Architecture of Classifiers
+

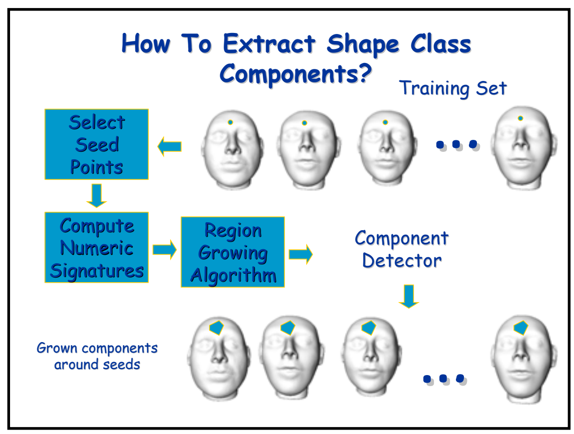
How To Extract Shape Class Components?
Training Set
Select Seed Points
Compute Numeric Signatures
Region Growing Algorithm
Component Detector
Grown components around seeds

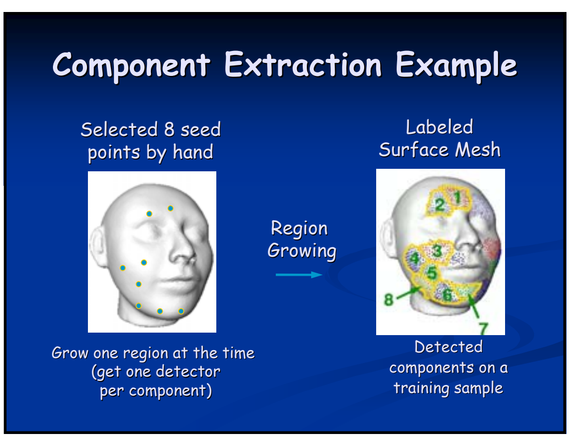
Component Extraction Example
Selected 8 seed points by hand
Region Growing
Labeled Surface Mesh
Grow one region at the time (get one detector per component)
Detected components on a training sample

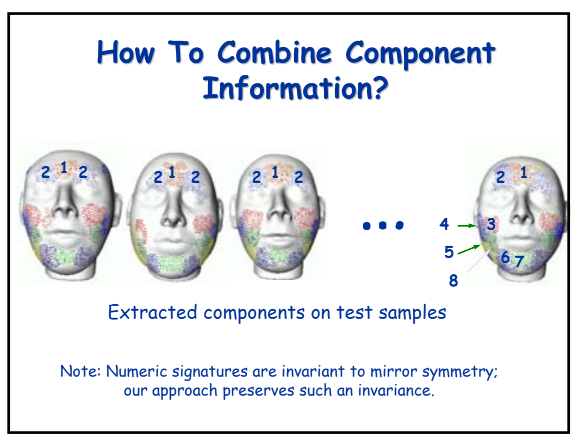
How To Combine Component Information?
Extracted components on test samples
Note: Numeric signatures are invariant to mirror symmetry; our approach preserves such an invariance.

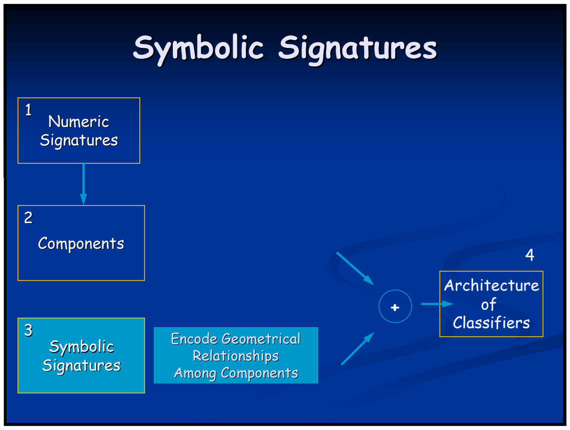
Symbolic Signatures
1 Numeric Signatures
2 Components
3 Symbolic Signatures
Encode Geometrical Relationships Among Components
+
4
Architecture of Classifiers

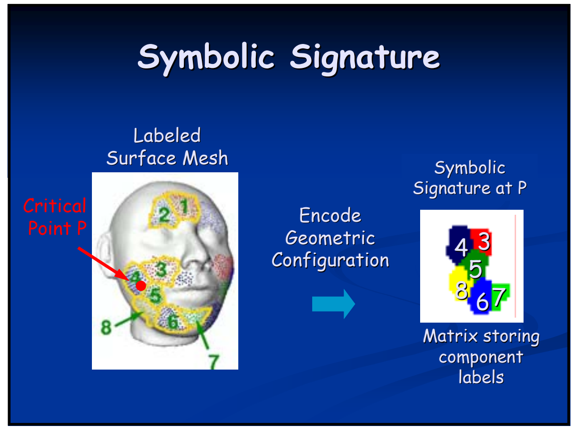
Symbolic Signature
Labeled Surface Mesh
Critical Point P
Encode Geometric Configuration
Symbolic Signature at P
Matrix storing component labels
1 2 3 4 5 6 7 8

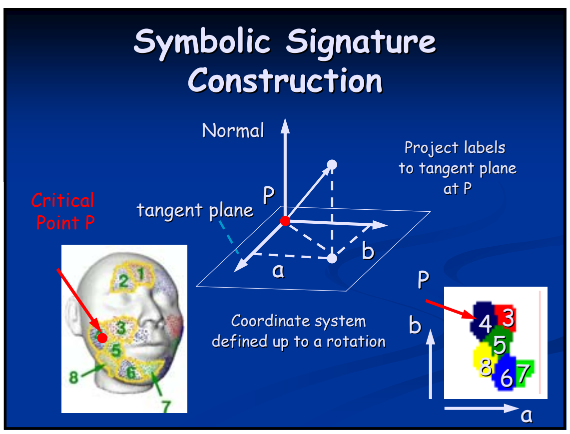
Symbolic Signature Construction
Normal
Project labels to tangent plane at P
Critical Point P
tangent plane
P
a
b
Coordinate system defined up to a rotation
P
b
a
1 2 3 4 5 6 7 8

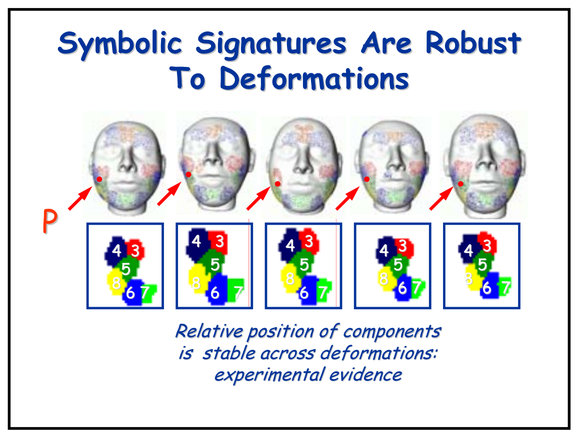
Symbolic Signatures Are Robust To Deformations
P
Relative position of components is stable across deformations: experimental evidence

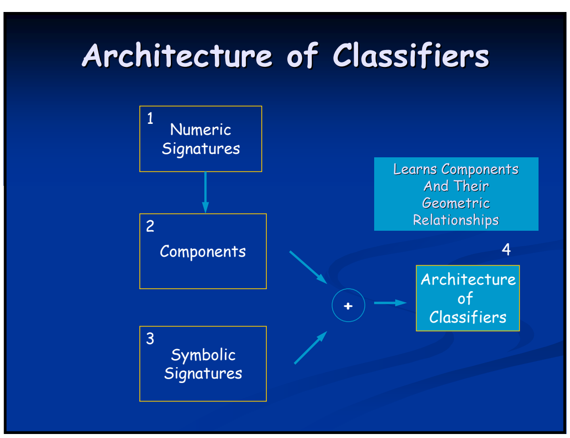
Architecture of Classifiers
1 Numeric Signatures
2 Components
3 Symbolic Signatures
Learns Components And Their Geometric Relationships
+
4
Architecture of Classifiers

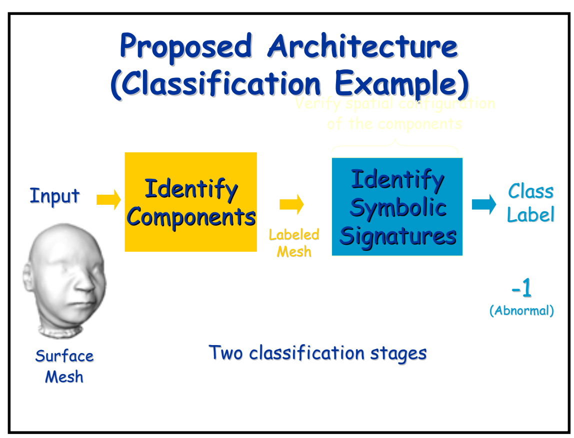
Proposed Architecture (Classification Example)
Input
Surface Mesh
Identify Components
Labeled Mesh
Identify Symbolic Signatures
Class Label
-1
(Abnormal)
Two classification stages

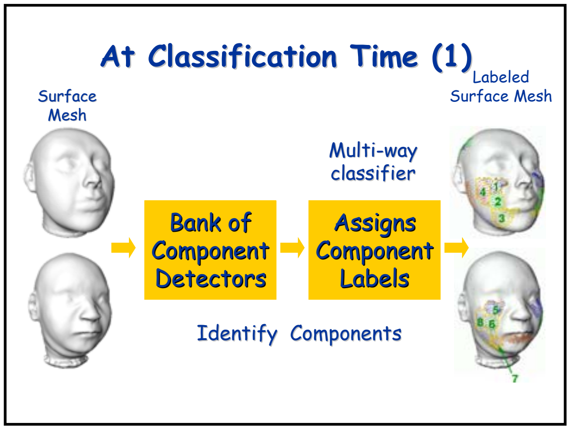


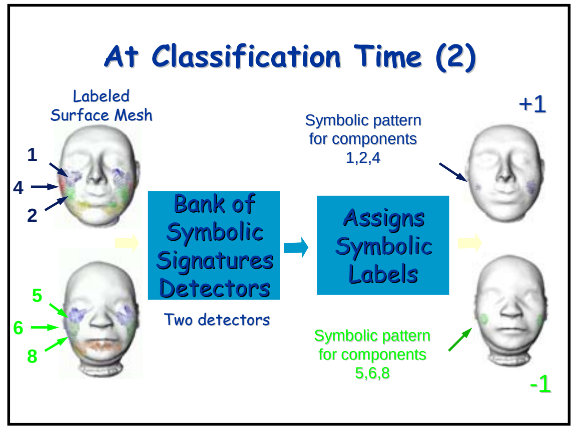


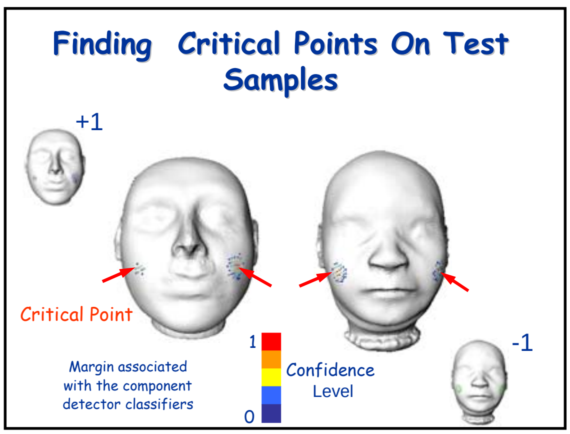


## Architecture Implementation

- ALL our classifiers are (off-the-shelf) ν-Support Vector Machines (ν-SVMs) (Schölkopf et al., 2000 and 2001).
- Component (and symbolic signature) detectors are one-class classifiers.
- Component label assignment: performed with a multi-way classifier that uses pairwise classification scheme.
- Gaussian kernel.

## Experimental Validation

Recognition Tasks: 4 (T1 - T4)
Classification Tasks: 3 (T5 – T7)
No. Experiments: 5470

Setup


Recognition


Classification

## Shape Classes

## Enlarging Training Sets Using Virtual Samples

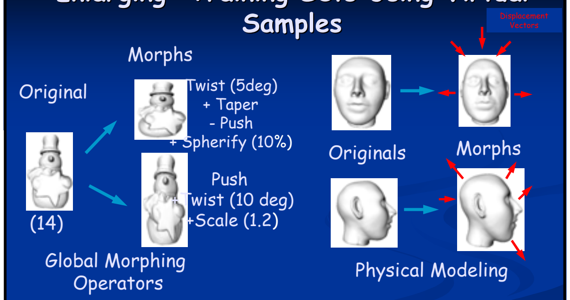




## Other Approaches

- Tried standard alignment-verification.
- Alignment-verification with PCA.
- However, no systematic comparison was performed due to poor performance.
- Existing methods for classifying shapes do not use range data.

P. Golland, NIPS 2001, J. Matrin et al. IEEE PAMI 1998.

## Task 1: Recognizing Single Objects (1)

- No. Shape classes: 9.
- Training set size: 400 meshes.
- Testing set size: 200 meshes.
- No. Experiments: 1960.
- No. Component detectors:3.
- No. Symbolic signature detectors: 1.
- Numeric signature size: 40x40.
- Symbolic signature size: 20x20.
- No clutter and occlusion.

## Task 1: Recognizing Single Objects (2)

- Snowman: 93%.
- Rabbit: 92%.
- Dog: 89%.
- Cat: 85.5%.
- Cow: 92%.
- Bear: 94%.
- Horse: 92.7%.
- Human head: 97.7%.
- Human face: 76%.

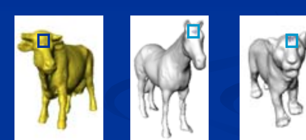

Recognition rates (true positives)

(No clutter, no occlusion, complete models)

## Tasks 2-3: Recognition In Complex Scenes (1)

- No. Shape classes: 3.
- Training set size: 400 meshes.
- Testing set size: 200 meshes.
- No. Experiments: 1200.
- No. Component detectors:3.
- No. Symbolic signature detectors: 1.
- Numeric signature size: 40x40.
- Symbolic signature size: 20x20.
- T2 – low clutter and occlusion.

## Task 2-3: Recognition in Complex Scenes (2)

| Shape Class | True Positives | False Positives | True Positives | False Positives |
|---|---|---|---|---|
| Snowmen | 91% | 31% | 87.5% | 28% |
| Rabbit | 90.2% | 27.6% | 84.3% | 24% |
| Dog | 89.6% | 34.6% | 88.12% | 22.1% |
| | Task 2 | | Task 3 | |

## Task 2-3: Recognition in Complex Scenes (3)

## Task 4: Recognizing Human Heads (1)

- No. Shape classes: 1.
- Training set size: 400 meshes.
- Testing set size: 250 meshes.
- No. Experiments: 710.
- No. Component detectors:8.
- No. Symbolic signature detectors: 2.
- Numeric signature size: 70x70.
- Symbolic signature size: 12x12.

## Task 4: Recognizing Human Heads (2)

Recognition Rate

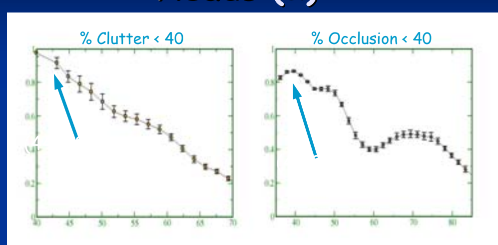


% Occlusion % Clutter

FP rate: ~1%,

## Task 4: Recognizing Human Heads (3)

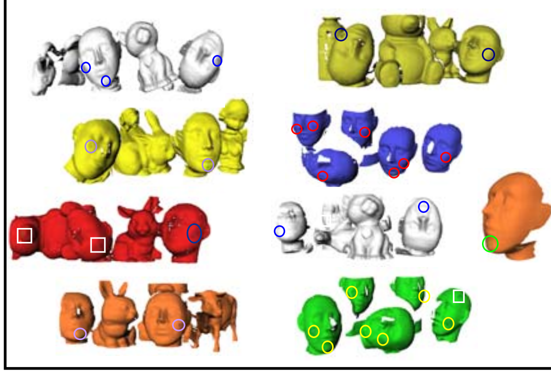

## Task 5: Classifying Normal vs. Abnormal Human Heads (1)

- No. Shape classes: 6.
- Training set size: 400 meshes.
- Testing set size: 200 meshes.
- No. Experiments: 1200.
- No. Component detectors:3.
- No. Symbolic signature detectors: 1.
- Numeric signature size: 50x50.
- Symbolic signature size: 12x12.

## Task 5: Classifying Normal vs. Abnormal Human Heads (1)

Five Cases

| Shape Classes | Classification Accuracy % |
|---|---|
| Normal vs. Abnormal 1 | 98 |
| Normal vs. Abnormal 2 | 100 |
| Abnormal 1 vs. 3 | 98 |
| Abnormal 1 vs. 4 | 97 |
| Abnormal 1 vs. 5 | 92 |

Full models

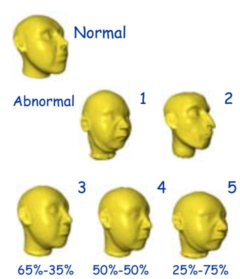


(convex combinations of Normal and Abnormal 1)

## Task 6: Classifying Normal vs. Abnormal Human Heads In Complex Scenes(1)

- No. Shape classes: 2.
- Training set size: 400 meshes.
- Testing set size: 200 meshes.
- No. Experiments: 1200.
- No. Component detectors:3.
- No. Symbolic signature detectors: 1.
- Numeric signature size: 100x100.
- Symbolic signature size: 12x12.

## Task 6: Classifying Normal vs. Abnormal Human Heads In Complex Scenes(1)

| Shape Classes | Classification Accuracy % |
|---|---|
| Normal vs. Abnormal 1 | 88 |

Clutter < 15%
and occlusion < 50%

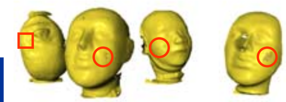

Range scenes – single view

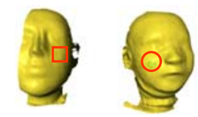

## Task 7: Classifying Normal vs. Abnormal Neurocranium (1)

- No. Shape classes: 2.
- Training set size: 400 meshes.
- Testing set size: 200 meshes.
- No. Experiments: 2200.
- No. Component detectors:3.
- No. Symbolic signature detectors: 1.
- Numeric signature size: 50x50.
- Symbolic signature size: 15x15.

## Task 7: Classifying Normal vs. Abnormal Neurocranium (2)

100 Experiments

| Shape Classes | Classification Accuracy % |
|---|---|
| Normal vs. Abnormal | 89 |

No clutter and occlusion

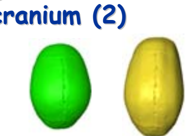

Normal
Abnormal
(sagittal synostosis )

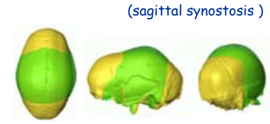

Superimposed models

## Main Contributions (1)

- A novel symbolic signature representation of deformable shapes that is robust to intra-class variability and missing information, as opposed to a numeric representation which is often tied to a specific shape.
- A novel kernel function for quantifying symbolic signature similarities.

## Main Contributions (2)

- A region growing algorithm for learning shape class components.
- A novel architecture of classifiers for abstracting the geometry of a shape class.
- A validation of our methodology in a set of large scale recognition and classification experiments aimed at applications in scene analysis and medical diagnosis.

## Main Contributions (3)

- Our approach:
- Is general can be applied to a variety of shape classes.
- Is robust to clutter and occlusion
- It Works in practice
- Is a step forward in 3-D object recognition research.

## Future Work (1)

- Encouraging results but need to make a more extensive quantification in order to characterize the algorithm wrt:
- sensor noise and mesh resolution,
- numeric and symbolic signature parameters,
- intra-class variability.

## Future Work (2)

- Need to find the break points.
- Investigate semi-automatic selection of seed points and critical points. At least, provide guidelines.
- Combine our approach with the alignment-verification technique.
- Simultaneous training of all classification stages.

## Thanks!

Now your questions …

## Spare slides

## So What Is A Component?

- Classification function, an outlier detector (one-class SVM) that defines two half-spaces in feature space:

All mesh points of the shape class whose numeric signatures lie on this "side" belong to the component.

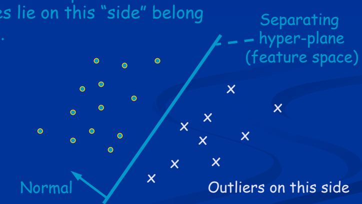

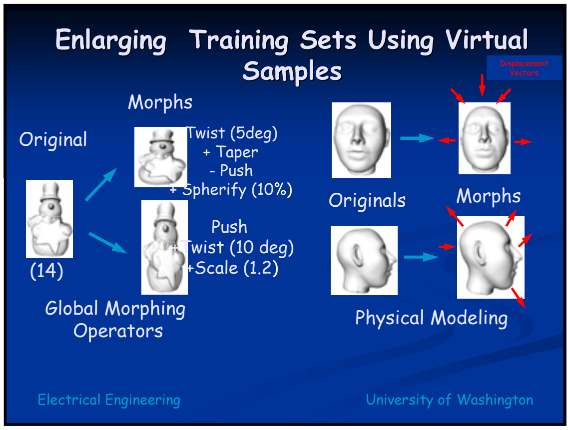


## Complexity (Worst Case)

- Numeric (symbolic) signature construction: $O(ns)$
- Bank of detectors: $O(nsc)$
- Label Assigner: $O(nsc^2)$

Where:

$n$ - number of scene points: ~$10^4$

$s$ - signature size: ~$10^2$-$10^3$

$c$ - number of detectors: ~10