Reconstruction

EE/CSE 576
Linda Shapiro

3D model

- “Digital copy” of real object
- Allows us to
  - Inspect details of object
  - Measure properties
  - Reproduce in different material
- Many applications
  - Cultural heritage preservation
  - Computer games and movies
  - City modelling
  - E-commerce





Applications: cultural heritage

SCULPTEUR European project

Applications: art

Block Works Precipitate III 2004
Mild steel blocks 80 x 46 x 66 cm

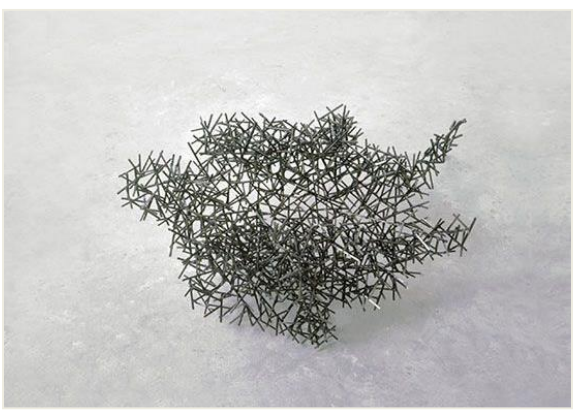

Domain Series Domain VIII Crouching
1999 *Mild steel bar* 81 x 59 x 63 cm

Applications: structure engineering

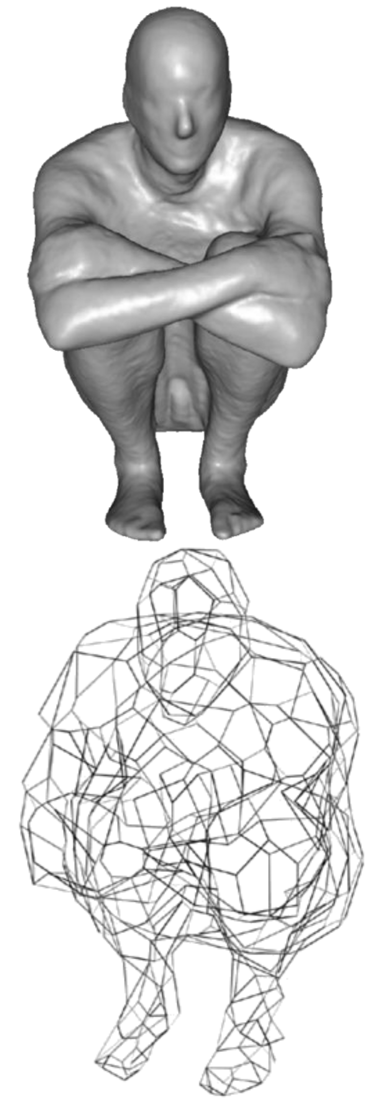

BODY / SPACE / FRAME, Antony Gormley, Lelystad, Holland

Applications: 3D indexation

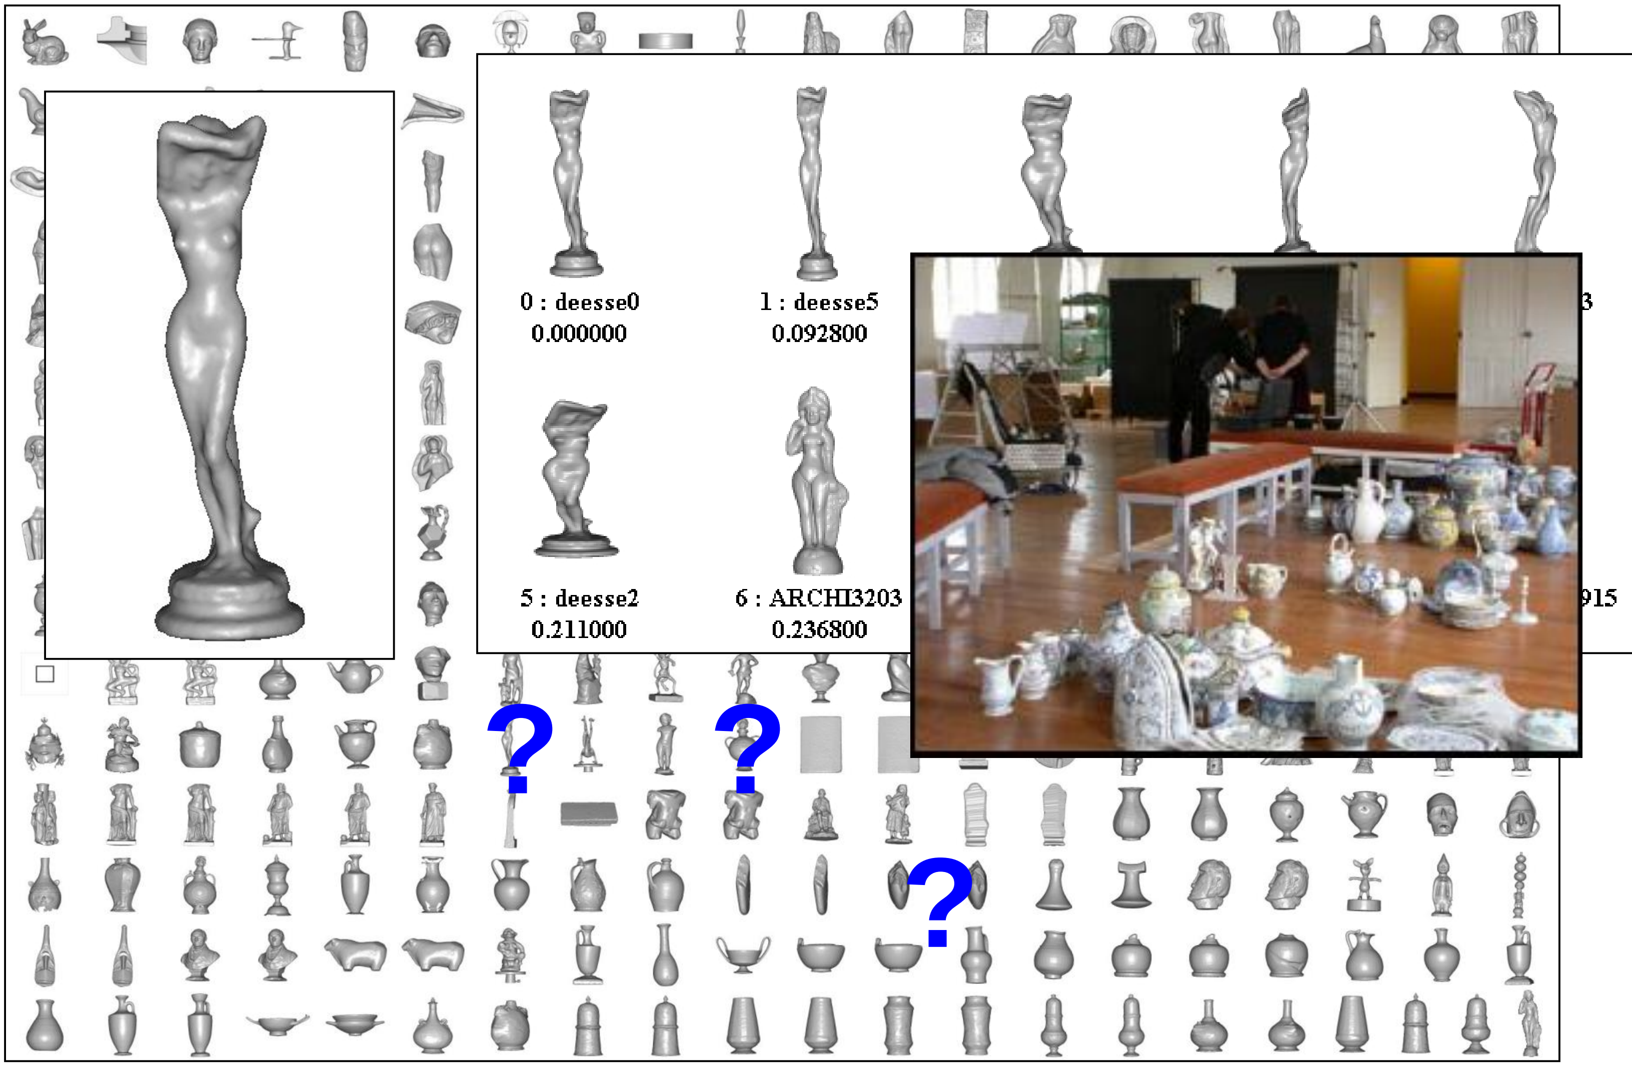

Applications: archaeology

- "forma urbis romae" project

Fragments of the City: Stanford's Digital Forma Urbis Romae Project
David Koller, Jennifer Trimble, Tina Najbjerg, Natasha Gelfand, Marc Levoy
Proc. Third Williams Symposium on Classical Architecture, Journal of Roman Archaeology supplement, 2006.

1186 fragments

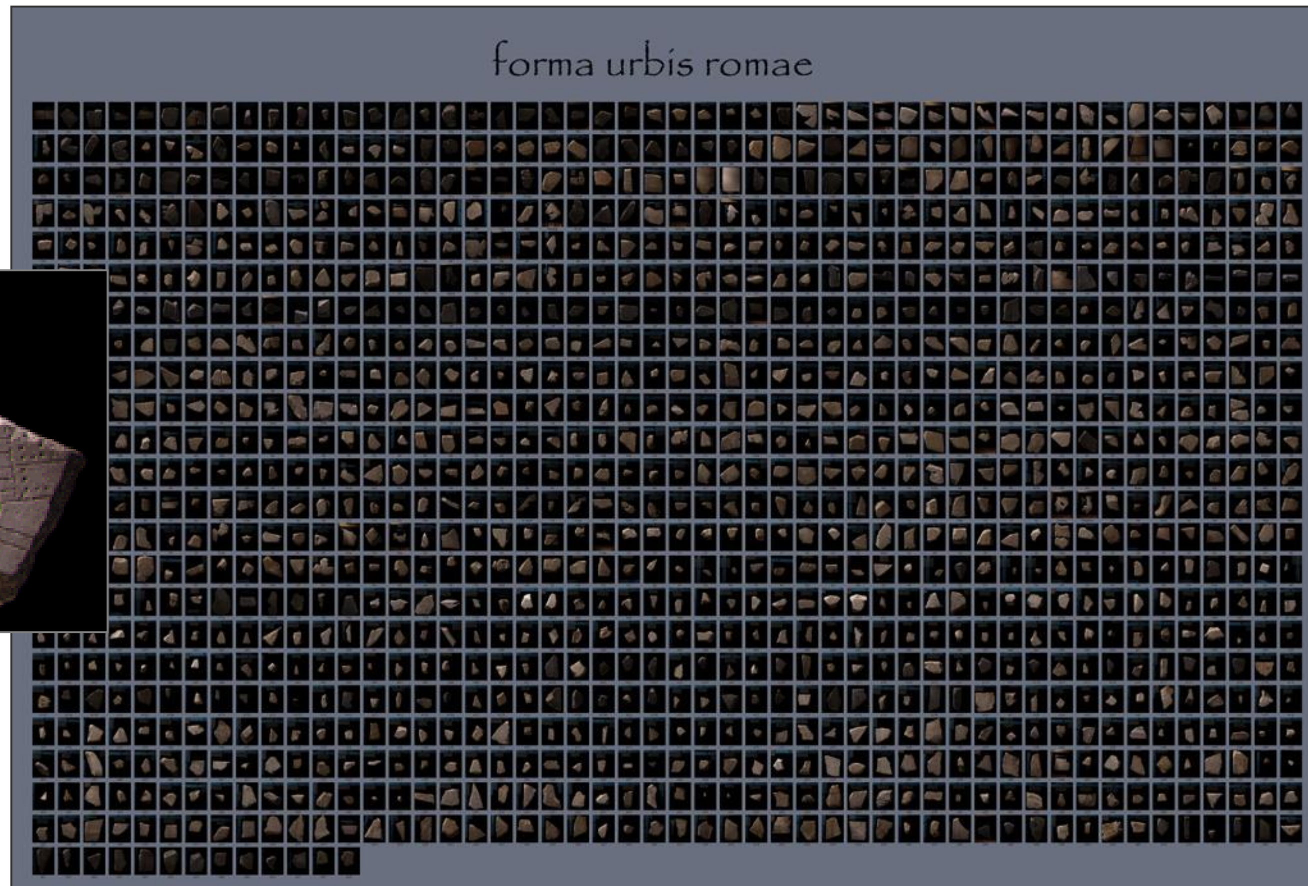

Applications: large scale modelling

[Furukawa10]

[Pollefeys08]

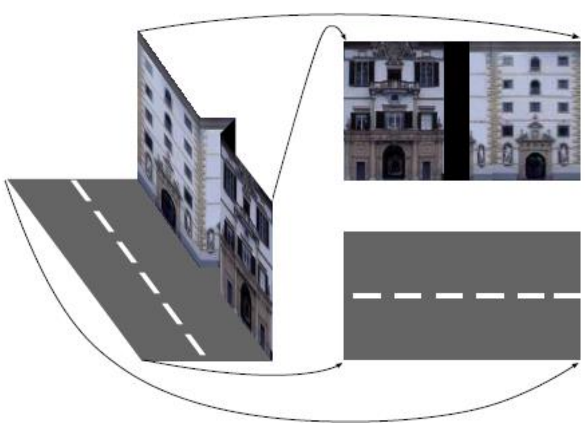

[Cornelis08]

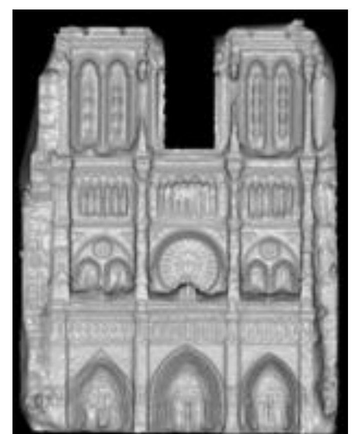

[Goesele07]

Applications: Medicine

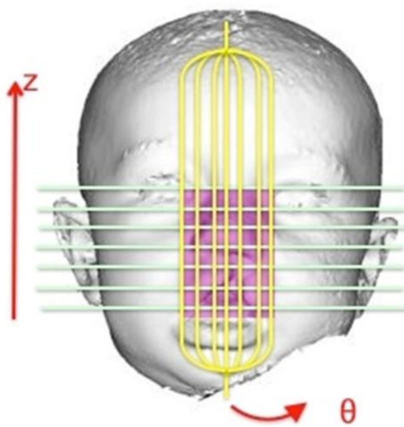

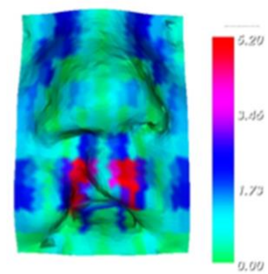

(a) Radius difference

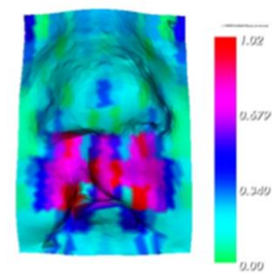

(b) Angle difference

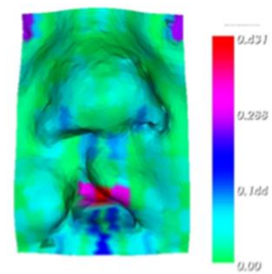

(c) Curvature difference

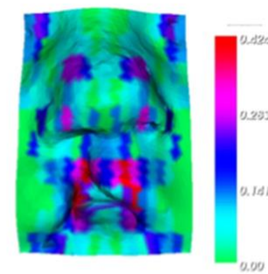

(d) Edge difference

expert's order	1	2	3	4	5	6	7	8	9	10
images	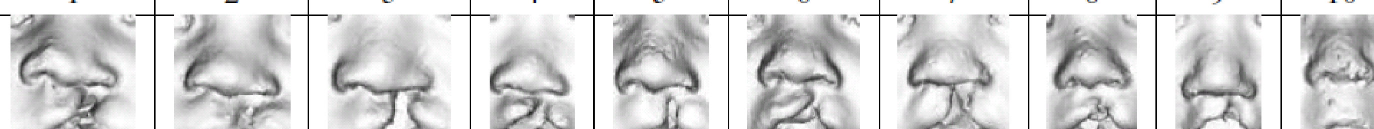									
learning	1	3	2	4	5	6	8	9	7	10
a-lmk	1	2	3	5	6	4	8	7	9	10
mirror	1	2	4	8	5	6	9	3	7	10
m-lmk	1	2	3	4	5	6	9	7	10	8
plane	1	2	3	5	4	6	7	9	10	8

Scanning technologies

- Laser scanner, coordinate measuring machine
  - Very accurate
  - Very Expensive
  - Complicated to use

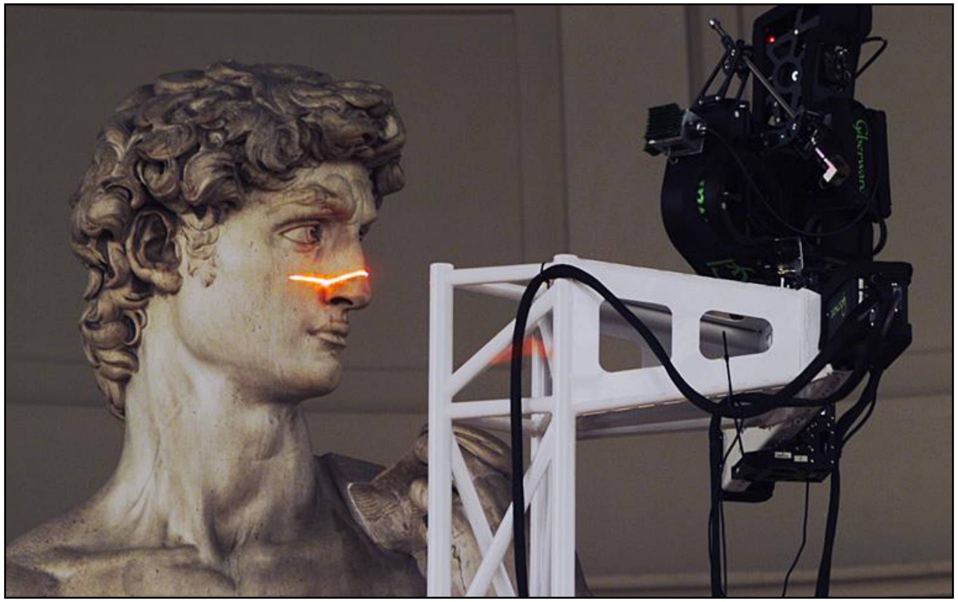

Minolta

Contura CMM

Medical Scanning System

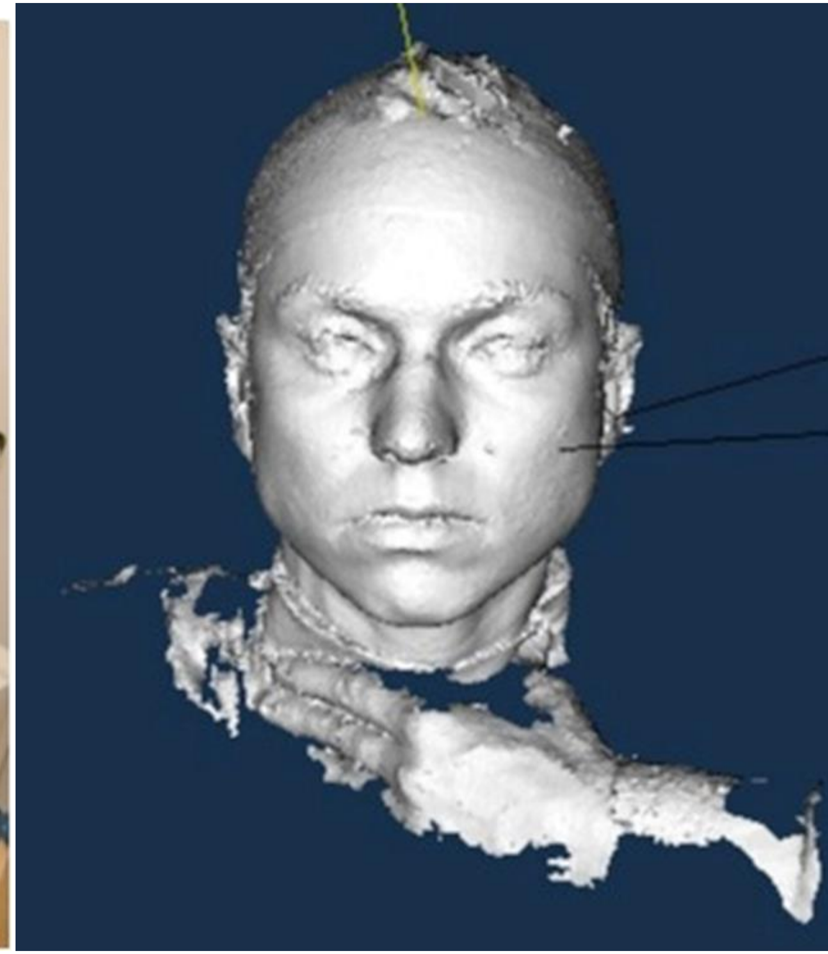

The “Us” Data Set (subset)

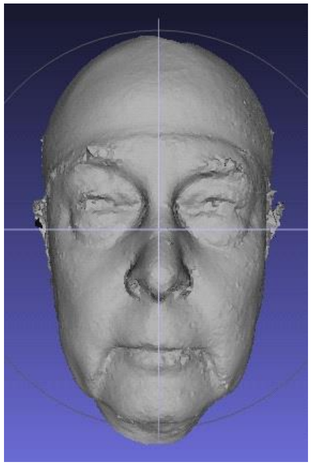
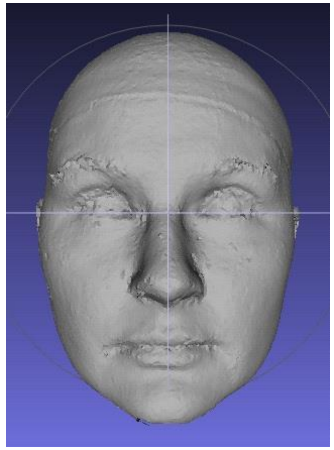
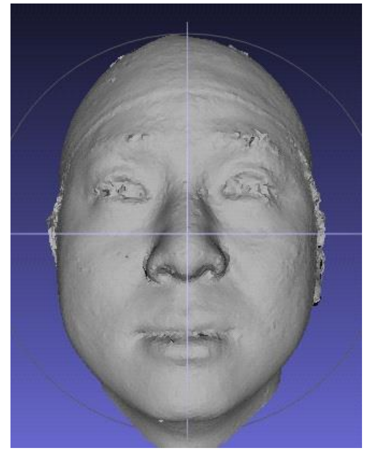
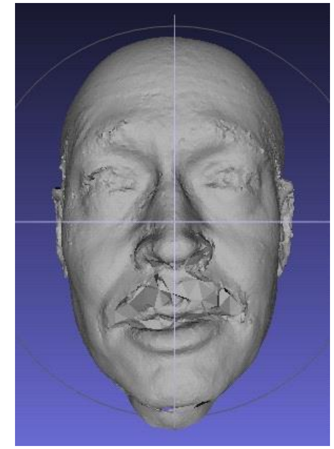
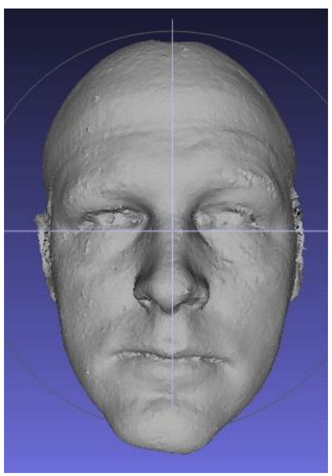
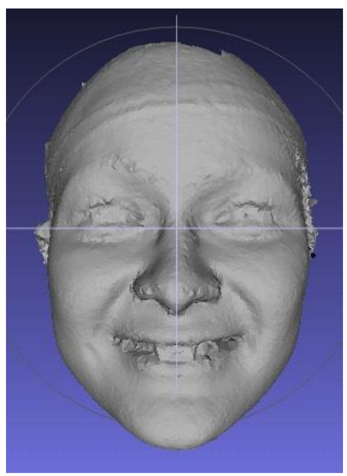
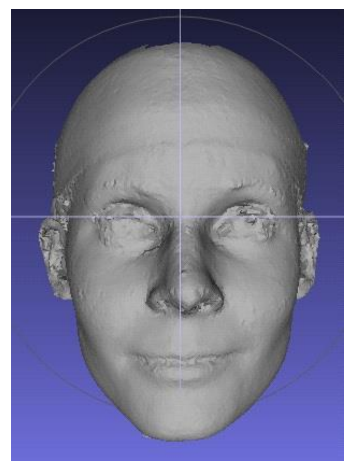
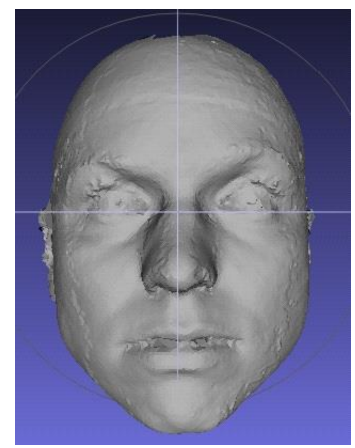

3d shape from photographs

"Estimate a 3d shape that would generate the input photographs given the same material, viewpoints and illumination"

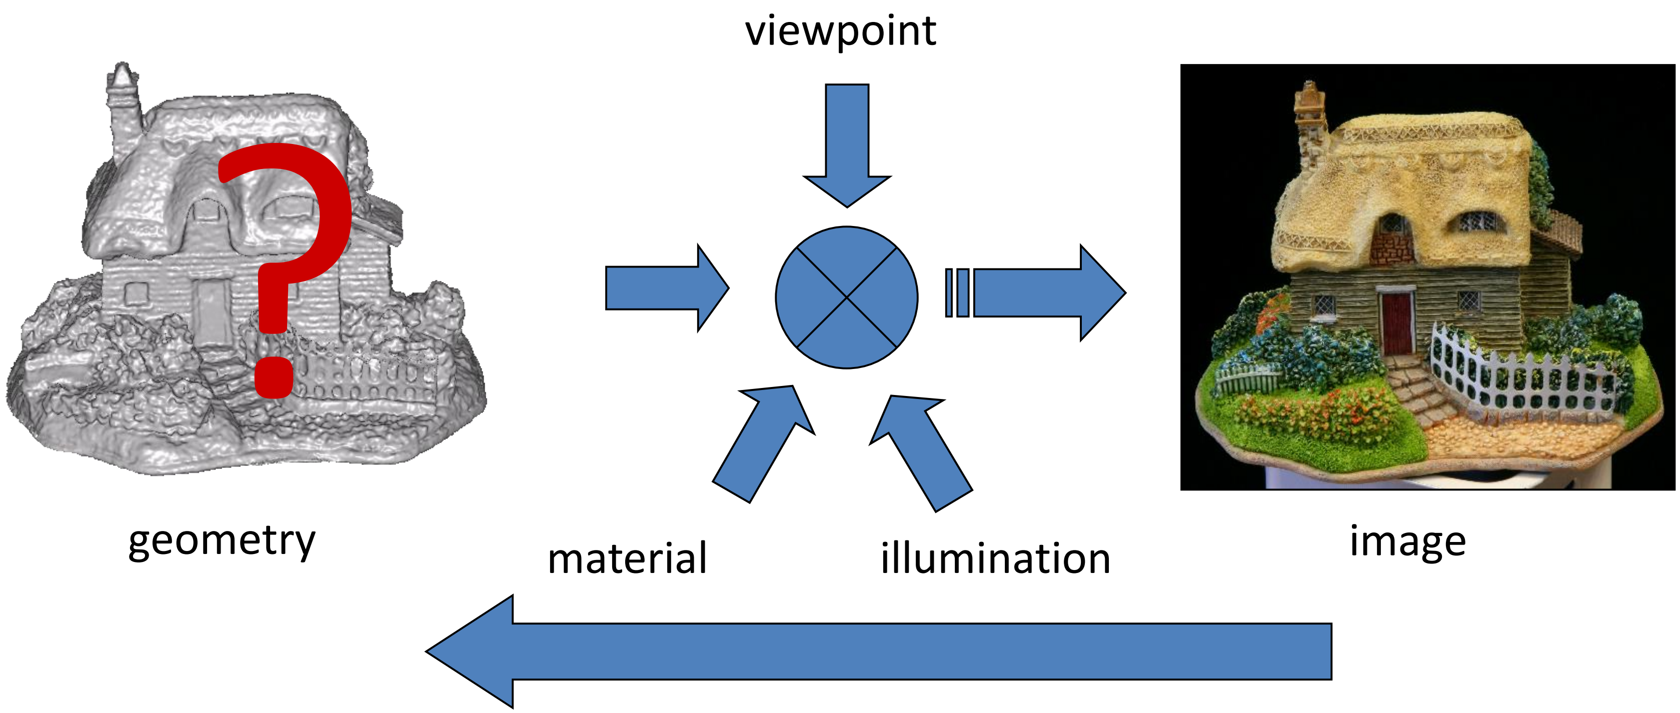

Photometric Stereo

- Estimate the surface normals of a given scene given multiple 2D images taken from the *same* viewpoint, but under *different lighting* conditions.
- Basic photometric stereo required a Lambertian reflectance model:

$$I = \rho\, \mathbf{n} \cdot \mathbf{v}$$

where I is pixel intensity, **n** is the normal, **v** is the lighting direction, and ρ is diffuse albedo constant, which is a reflection coefficient.

Basic Photometric Stereo

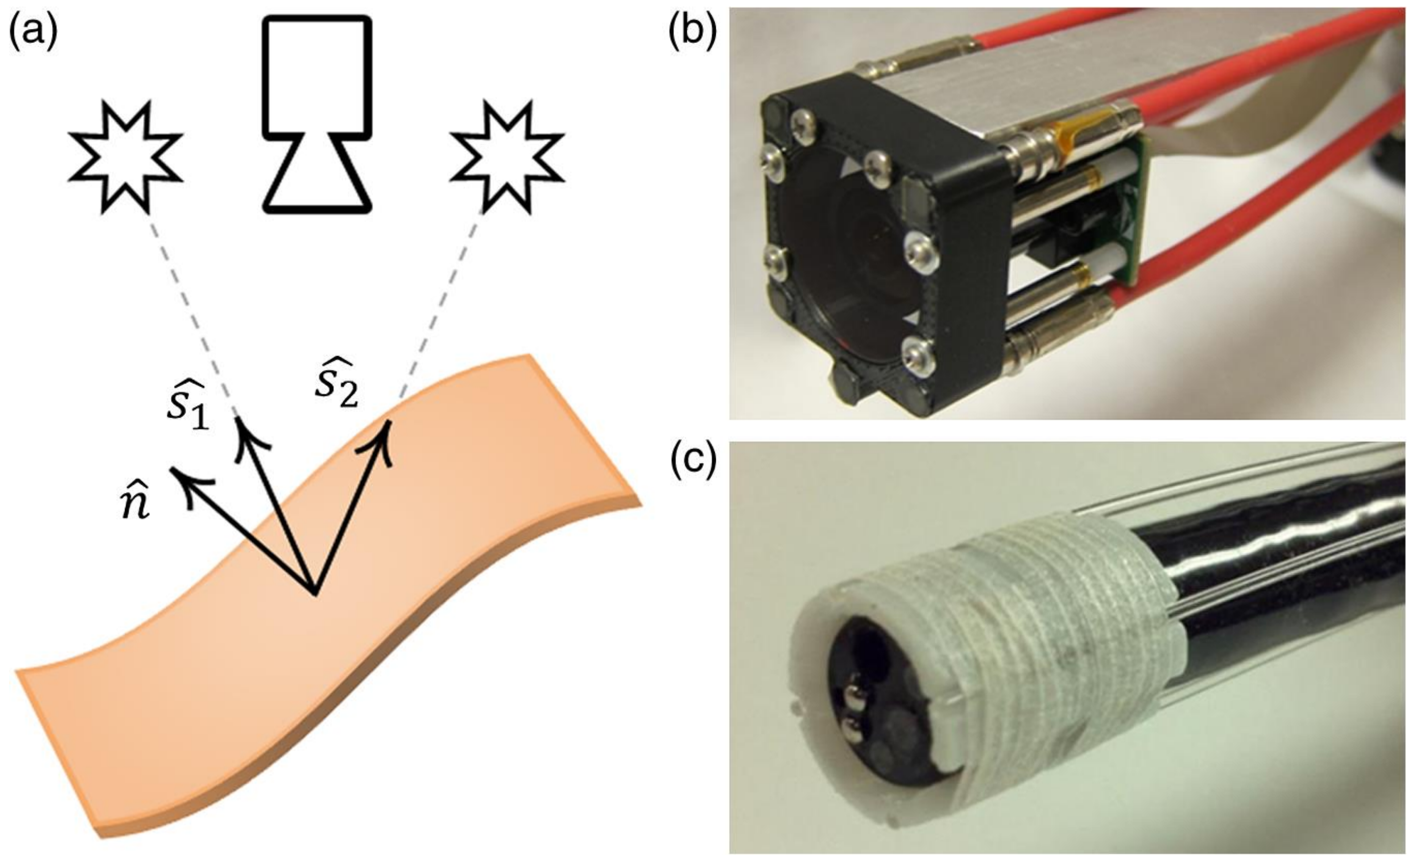

Basic Photometric Stereo

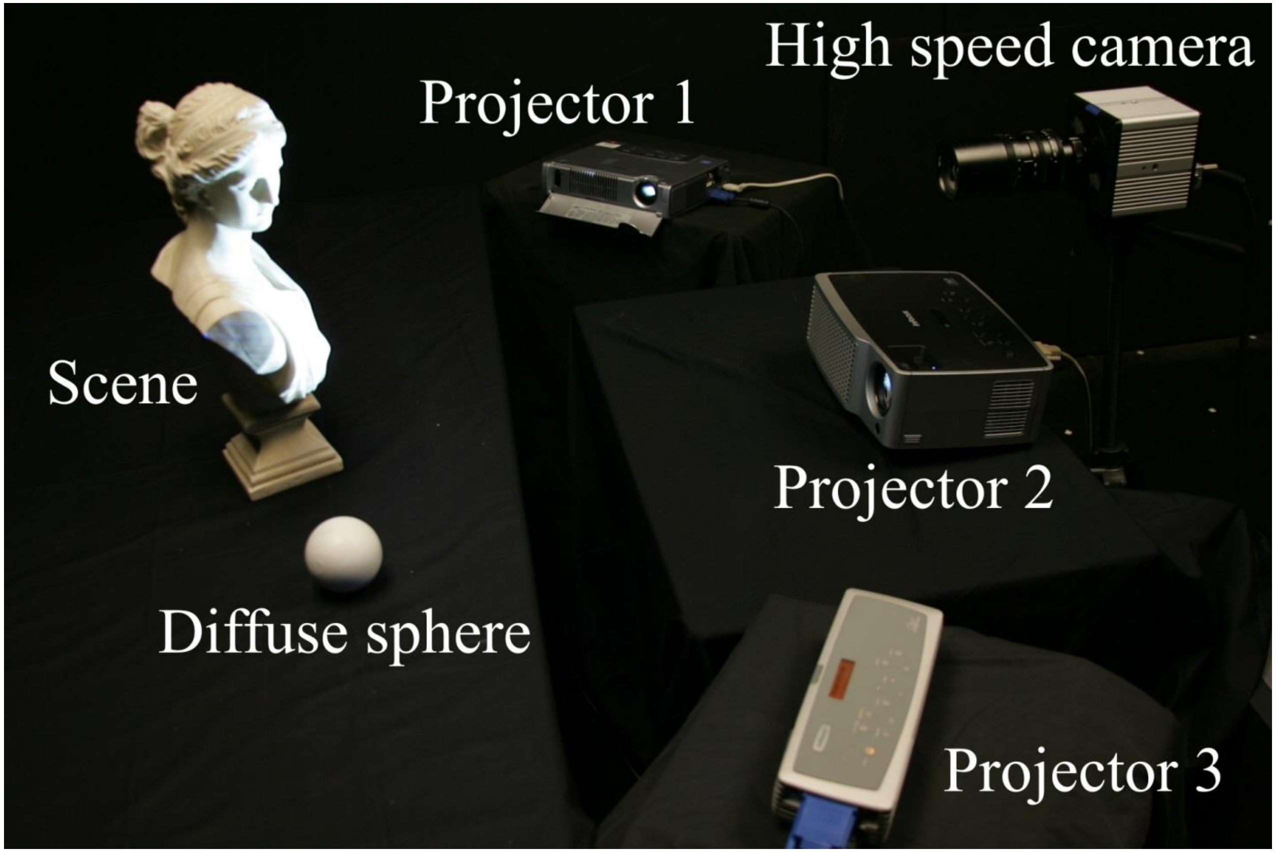

Basic Photometric Stereo

- K light sources
- Lead to K images $R_1(p,q)$, ...,$R_K(p,q)$ each from just one of the light sources being on
- For any (p,q), we get K intensities I_1,...I_K
- Leads to a set of linear equations of the form

$$I_k = \rho \mathbf{n} \bullet \mathbf{v}_k$$

- Solving leads to a surface normal map.

Photometric Stereo

Inputs

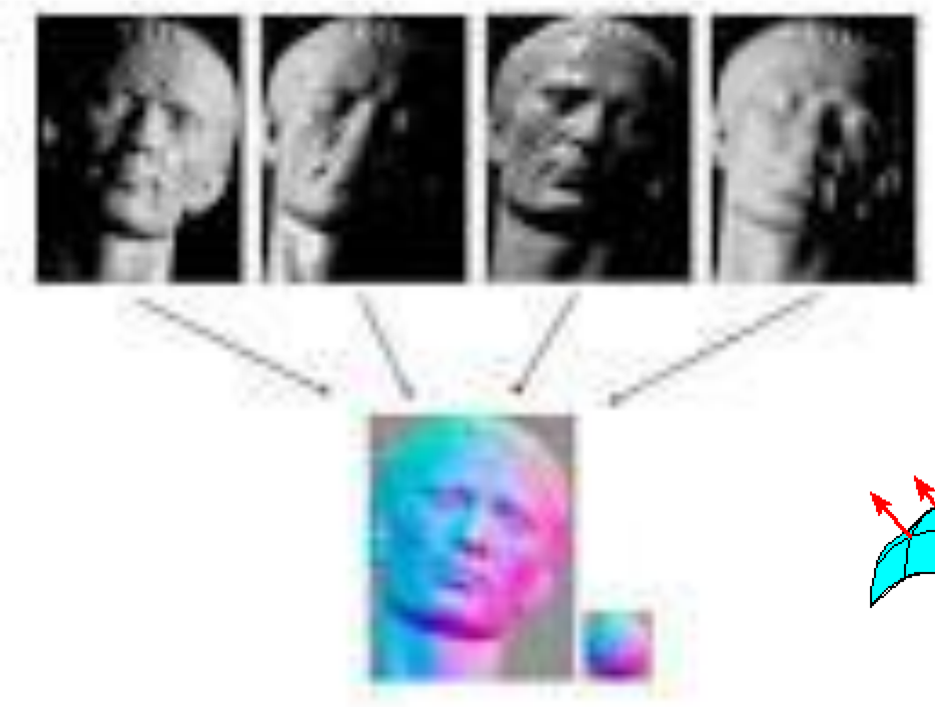

3D normals

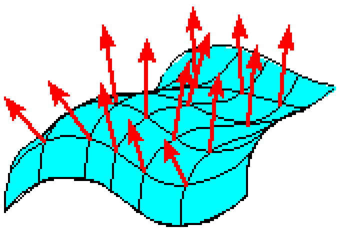

3d shape from photographs

Photograph based 3d reconstruction is:

- ✓ practical
- ✓ fast
- ✓ non-intrusive
- ✓ low cost
- ✓ Easily deployable outdoors
- ✗ "low" accuracy
- ✗ Results depend on material properties

Reconstruction

- Generic problem formulation: given several images of the same object or scene, compute a representation of its 3D shape

  

Reconstruction

- Generic problem formulation: given several images of the same object or scene, compute a representation of its 3D shape
- "Images of the same object or scene"
  - Arbitrary number of images (from two to thousands)
  - Arbitrary camera positions (camera network or video sequence)
  - Calibration may be initially unknown
- "Representation of 3D shape"
  - Depth maps
  - Meshes
  - Point clouds
  - Patch clouds
  - Volumetric models
  - Layered models

Multiple-baseline stereo

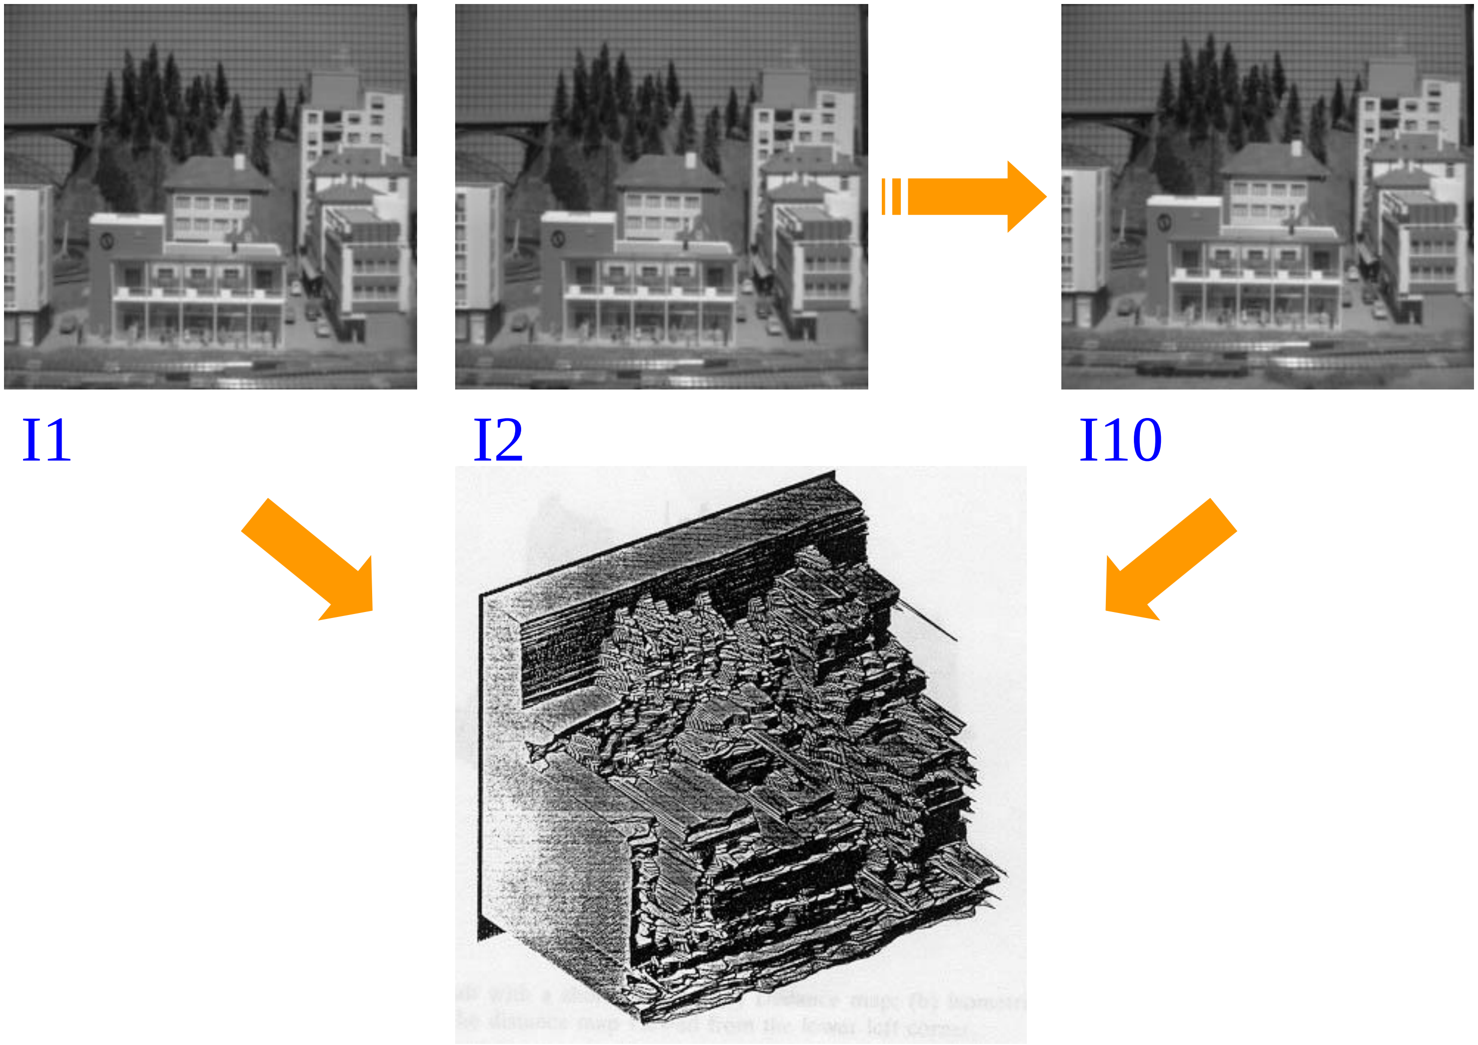

M. Okutomi and T. Kanade, "A Multiple-Baseline Stereo System," IEEE Trans. on Pattern Analysis and Machine Intelligence, 15(4):353-363 (1993).

Reconstruction from silhouettes

Can be computed robustly

Can be computed efficiently

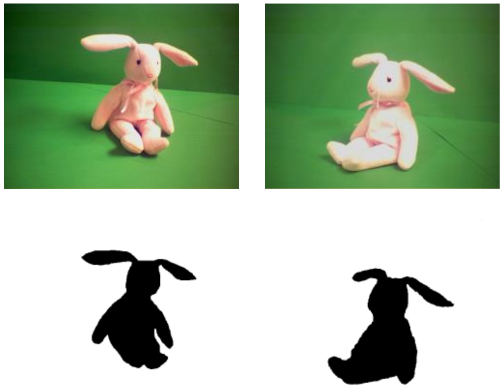

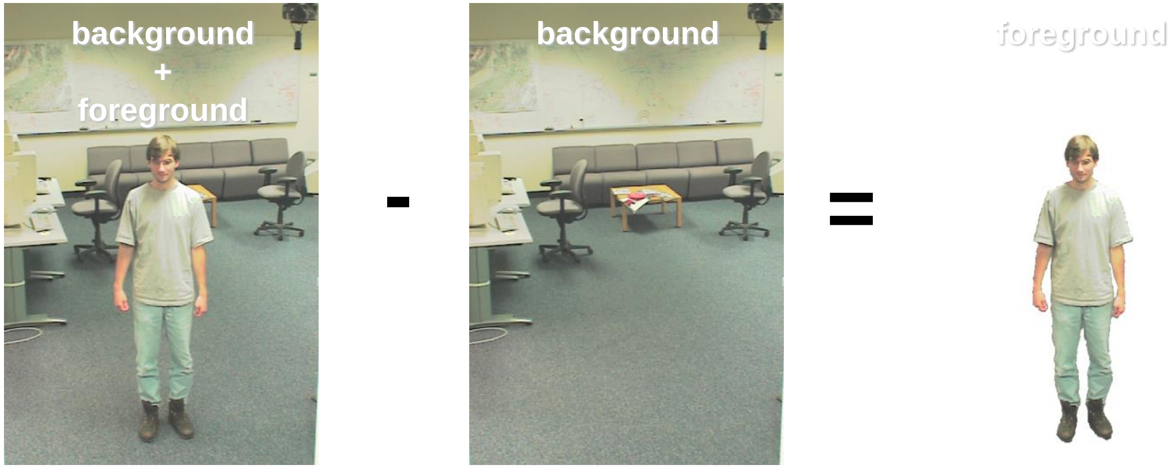

Reconstruction from Silhouettes

- The case of binary images: a voxel is photo-consistent if it lies inside the object's silhouette in **all** views

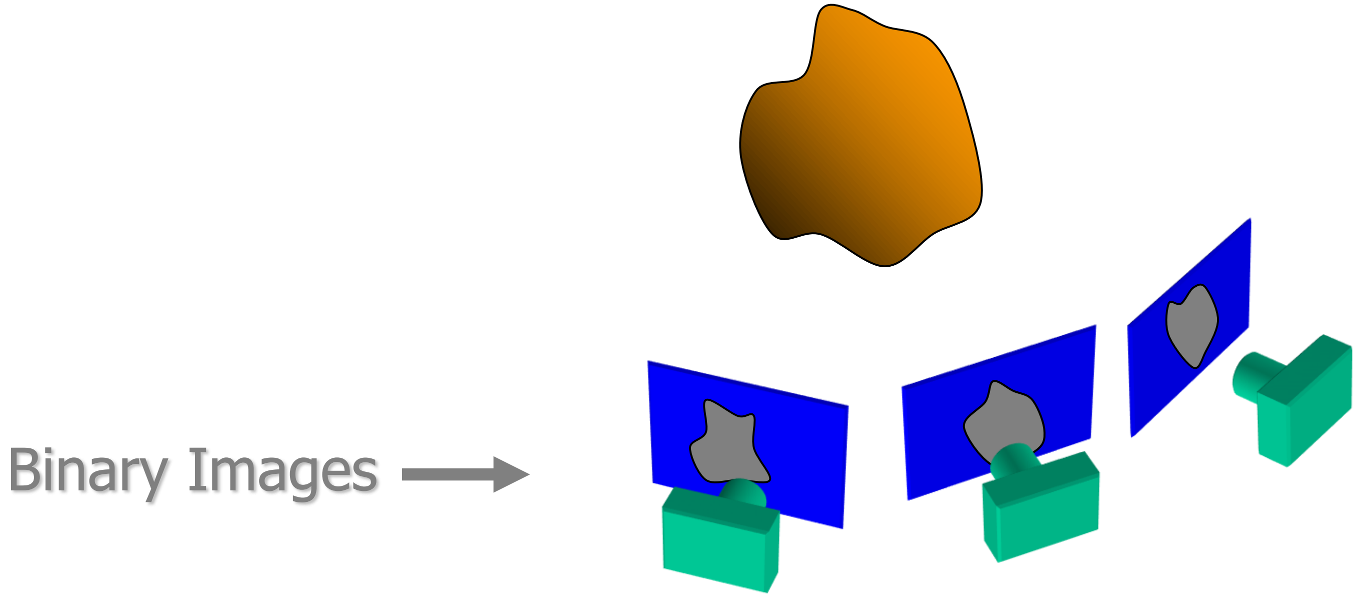

Reconstruction from Silhouettes

- The case of binary images: a voxel is photo-consistent if it lies inside the object's silhouette in all views

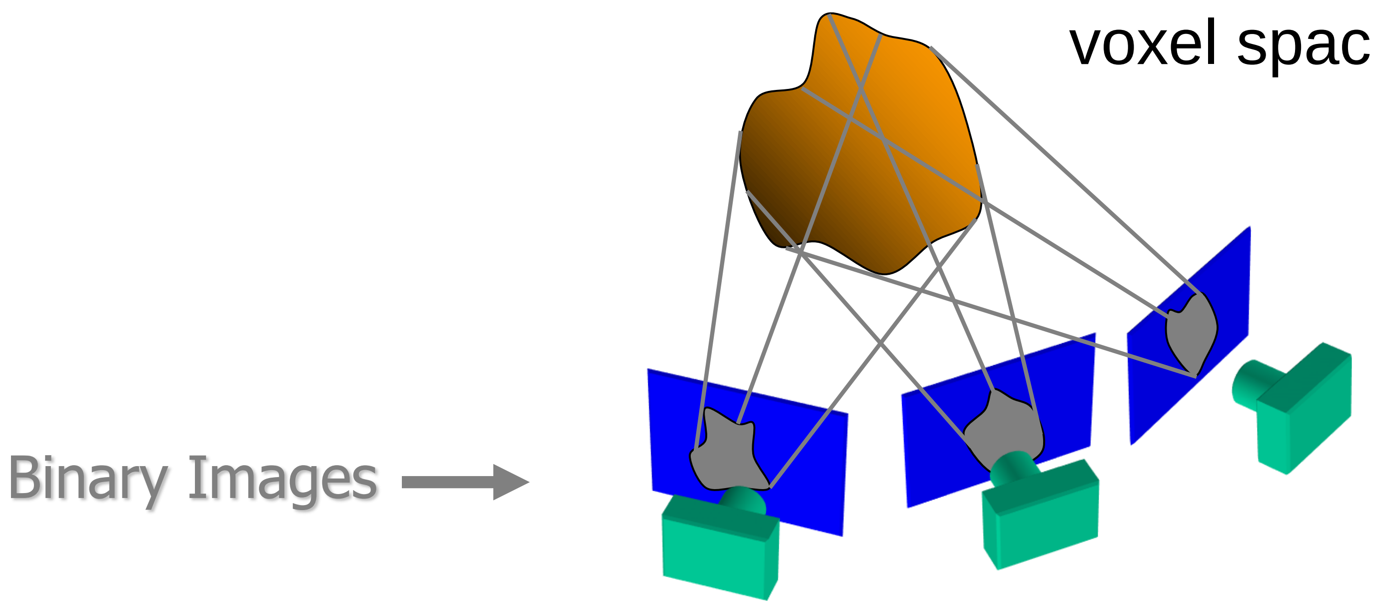

Finding the silhouette-consistent shape (*visual hull*):

- *Backproject* each silhouette
- Intersect backprojected volumes

Calibrated Image Acquisition

Calibrated Turntable

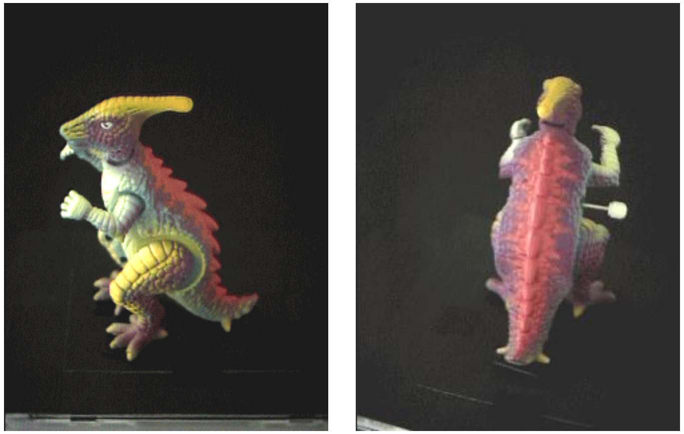

Selected Dinosaur Images

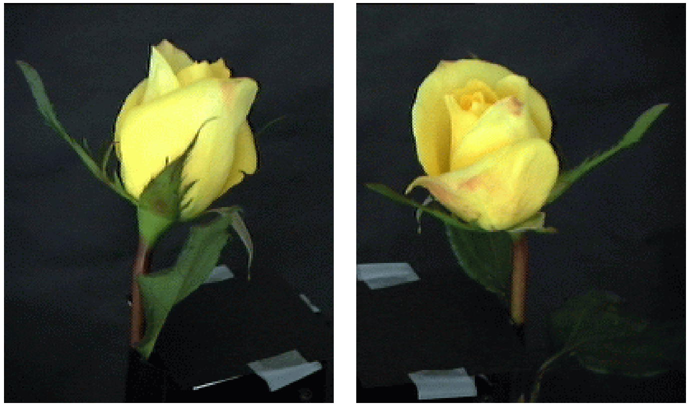

Selected Flower Images

Space Carving in General

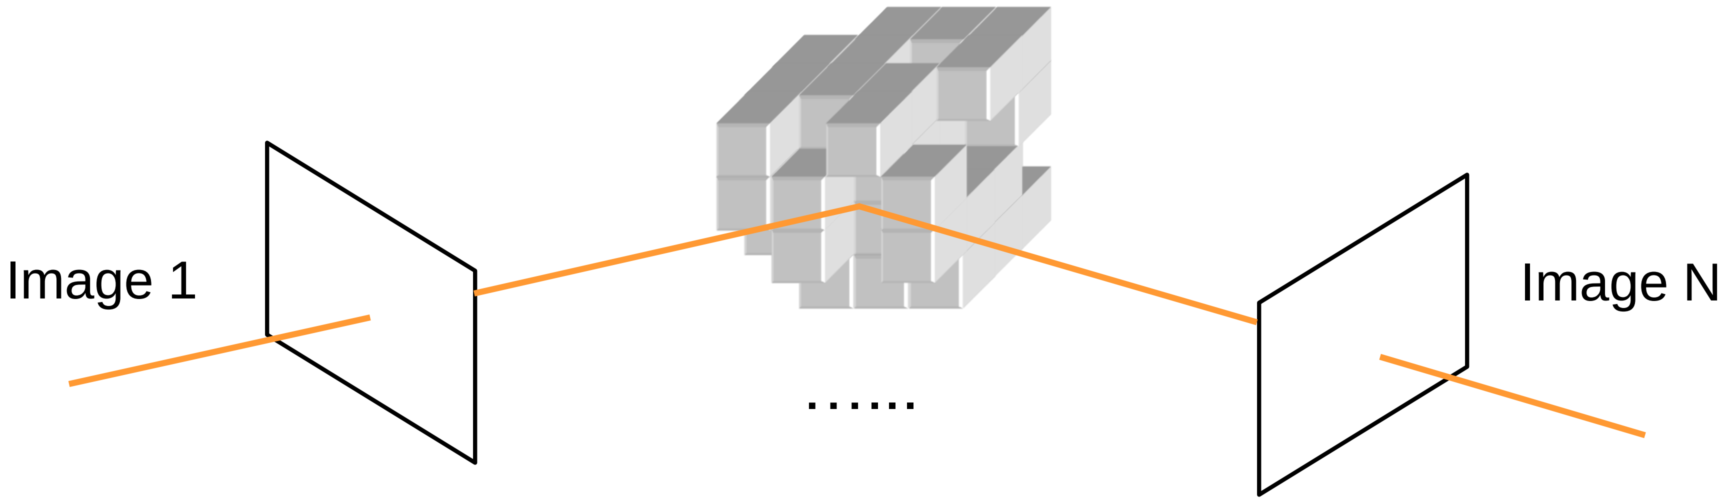

Space Carving Algorithm

- Initialize to a volume V containing the true scene
- Choose a voxel on the outside of the volume
- Project to visible input images
- Carve if not photo-consistent (inside object's silhouette)
- Repeat until convergence

K. N. Kutulakos and S. M. Seitz, **A Theory of Shape by Space Carving**, *ICCV* 1999

Our 4-camera light-striping stereo system

(now deceased)

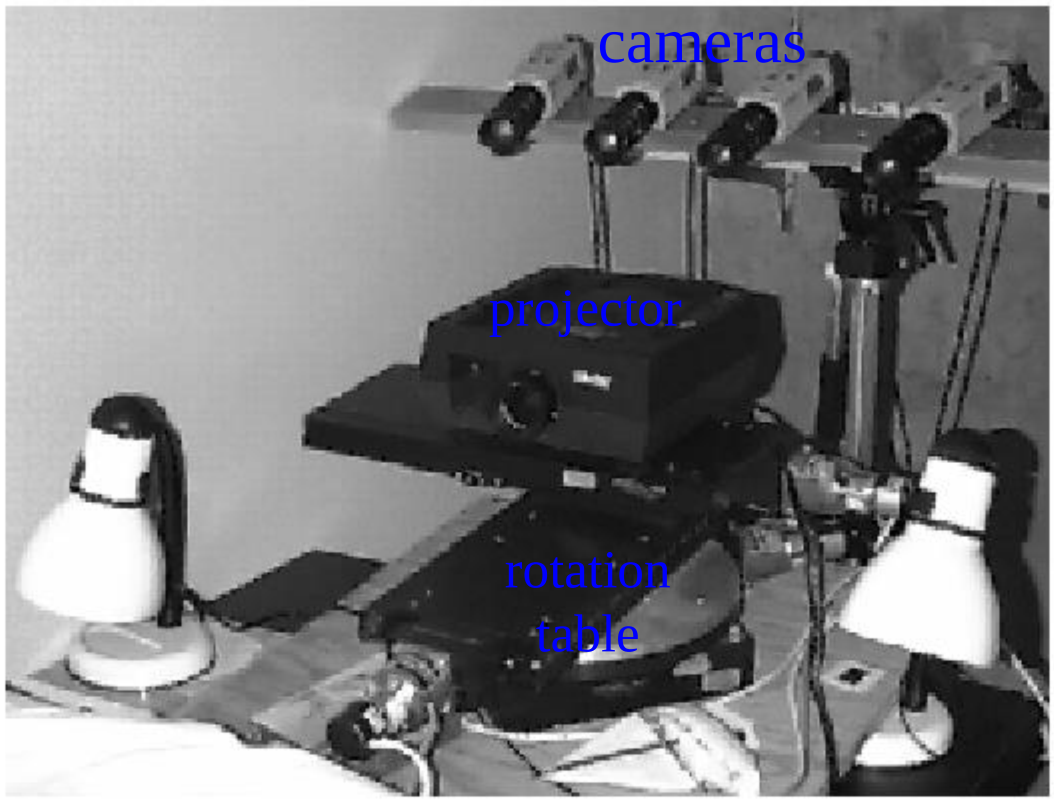

3D
object

Calibration Object

The idea is to snap images at different depths and get a lot of 2D-3D point correspondences.

Surface Modeling and Display from Range and Color Data

Kari	Pulli	UW
Michael	Cohen	MSR
Tom	Duchamp	UW
Hugues	Hoppe	MSR
John	McDonald	UW
Linda	Shapiro	UW
Werner	Stuetzle	UW

UW = University of Washington
Seattle, WA USA

MSR = Microsoft Research
Redmond, WA USA

Introduction

Goal

- develop robust algorithms for constructing 3D models from range & color data
- use those models to produce realistic renderings of the scanned objects

Surface Reconstuction

Step 1: Data acquisition

Obtain range data that covers the object. Filter, remove background.

Step 2: Registration

Register the range maps into a common coordinate system.

Step 3: Integration

Integrate the registered range data into a single surface representation.

Step 4: Optimization

Fit the surface more accurately to the data, simplify the representation.

Carve space in cubes

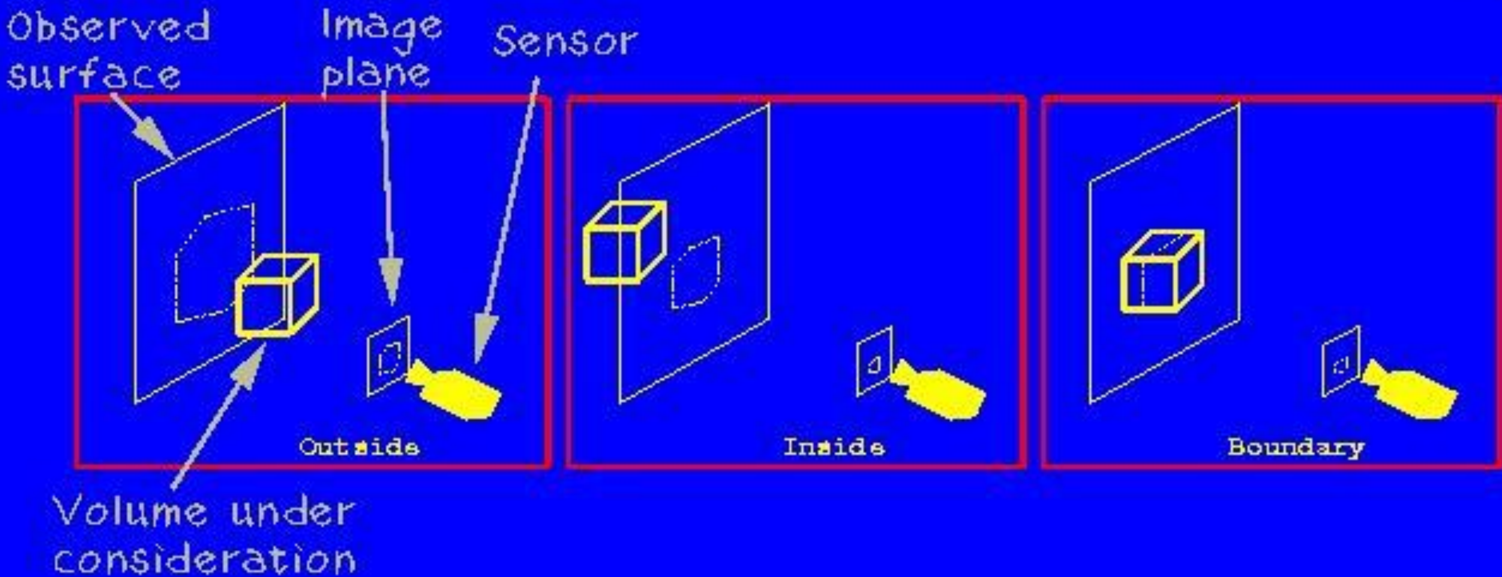

Label cubes

- Project cube to image plane (hexagon)
- Test against data in the hexagon

3D space is made up of many cubes.

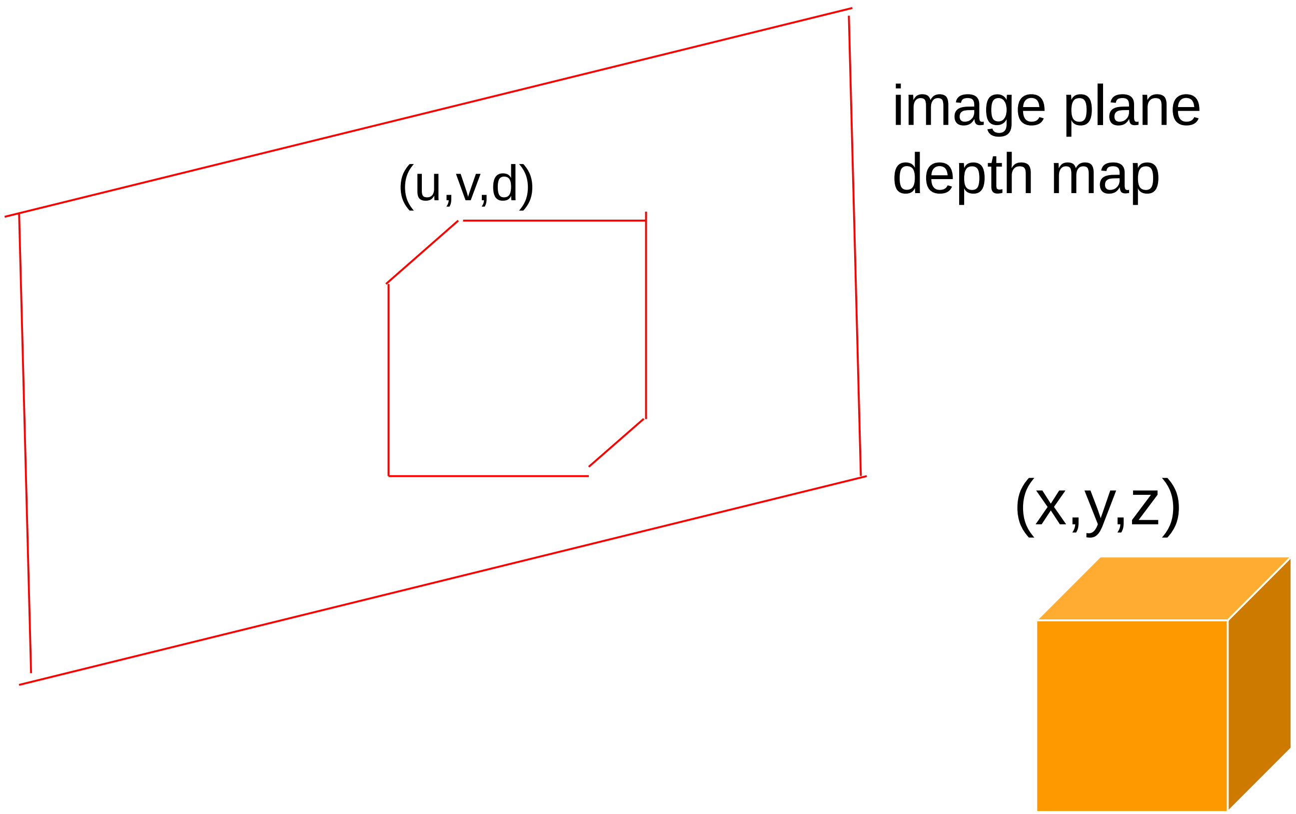

one of many cubes
in virtual 3D cube space

OUTSIDE

Several views

Processing order:

FOR EACH cube
    FOR EACH view

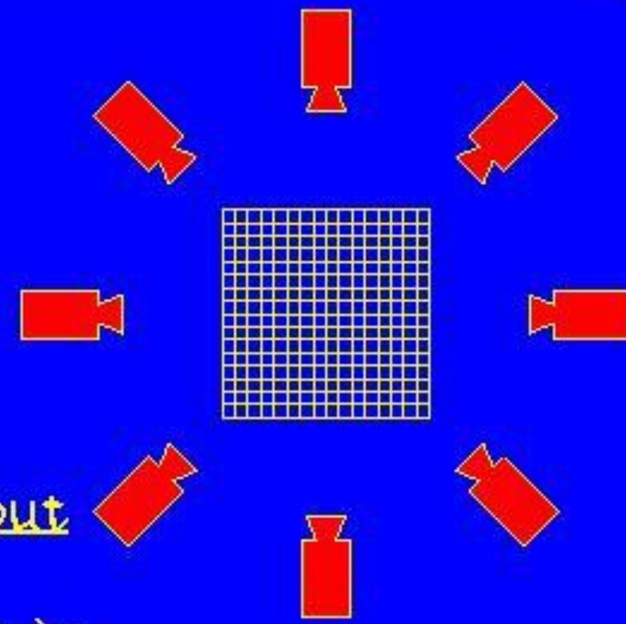

Rules:

any view thinks cube's out
 => it's out
every view thinks cube's in
 => it's in
else
 => it's at boundary

Hierarchical space carving

- Big cubes => fast, poor results
- Small cubes => slow, more accurate results
- Combination = octrees

RULES:
- cube's out => done
- cube's in => done
- else => recurse

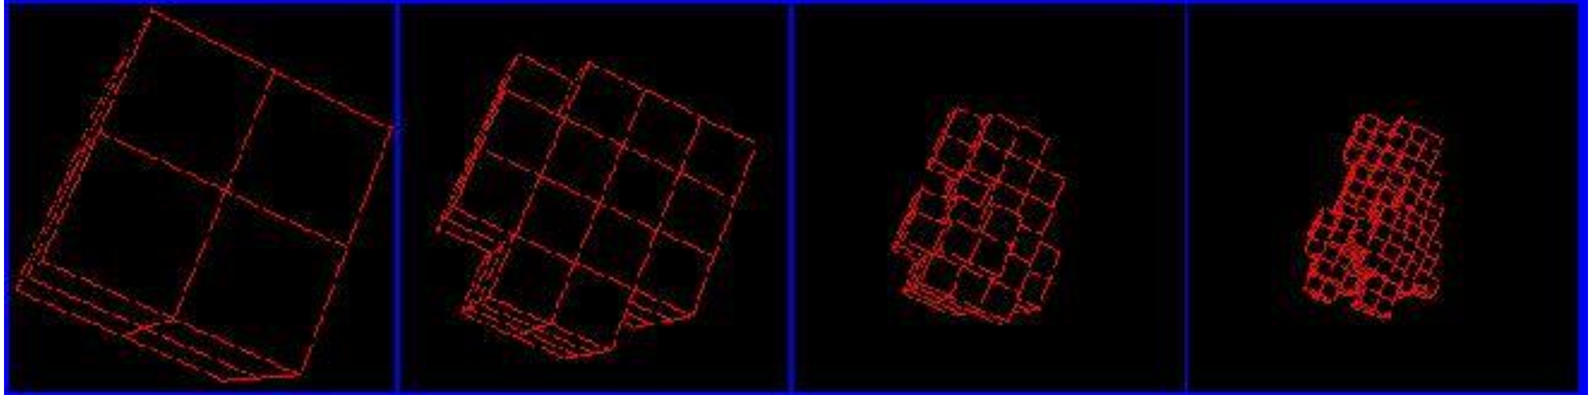

Hierarchical space carving

- Big cubes => fast, poor results
- Small cubes => slow, more accurate results
- Combination = octrees

RULES:
- cube's out => done
- cube's in => done
- else => recurse

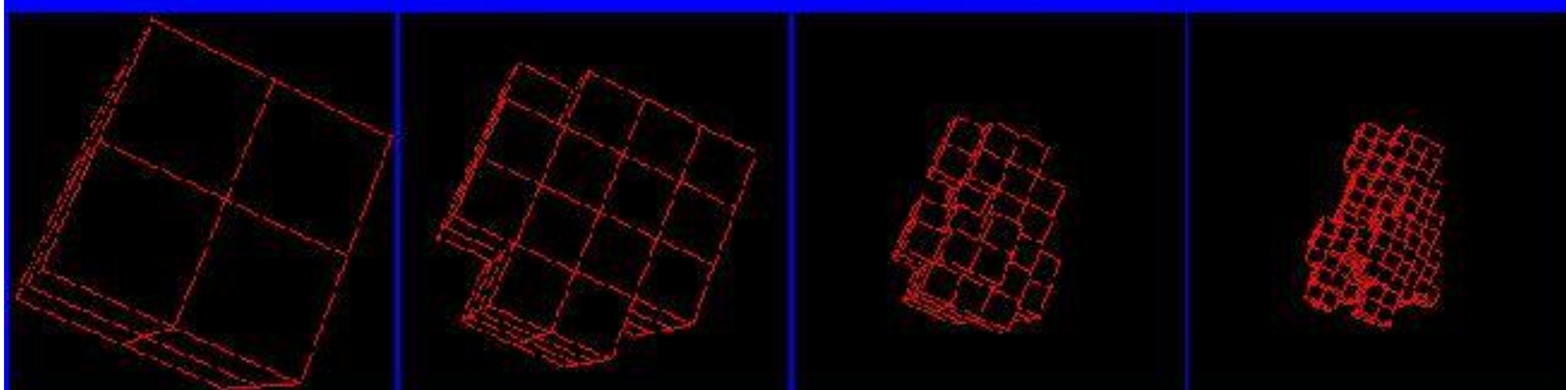

The rest of the chair

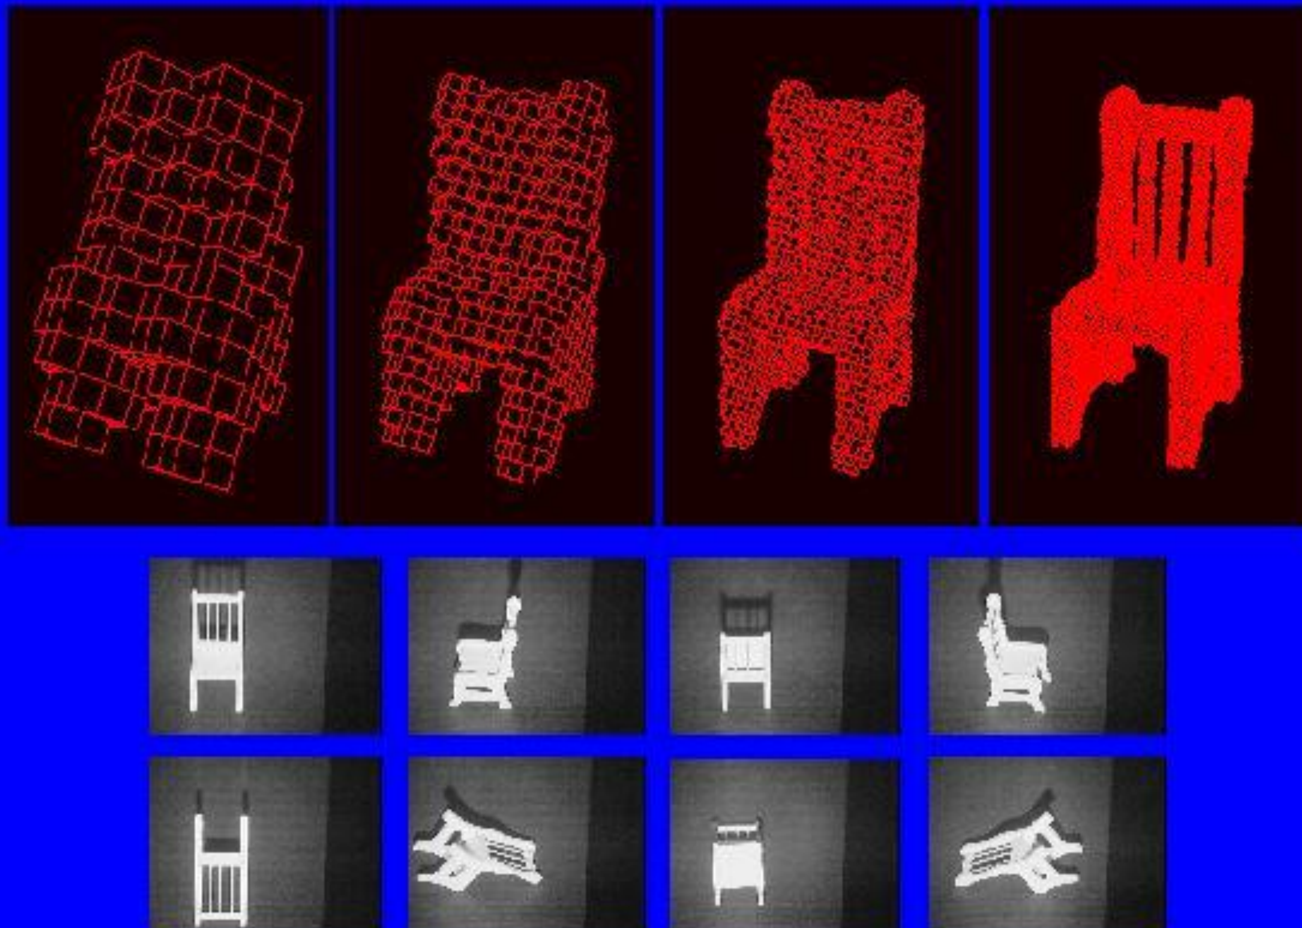

Same for a husky pup

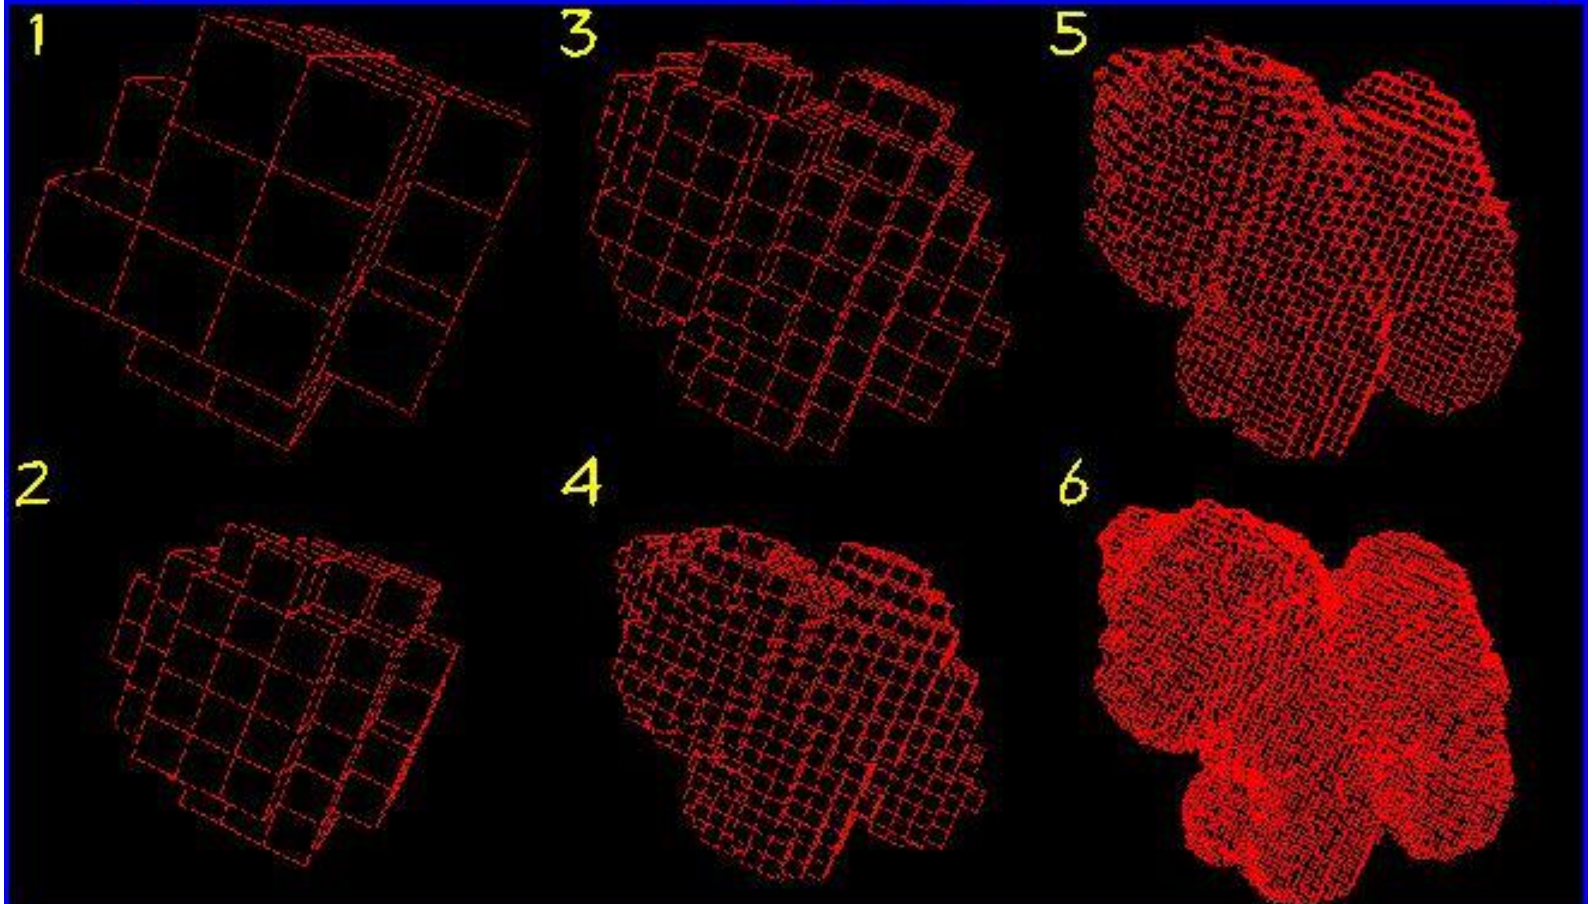

Optimizing the dog mesh

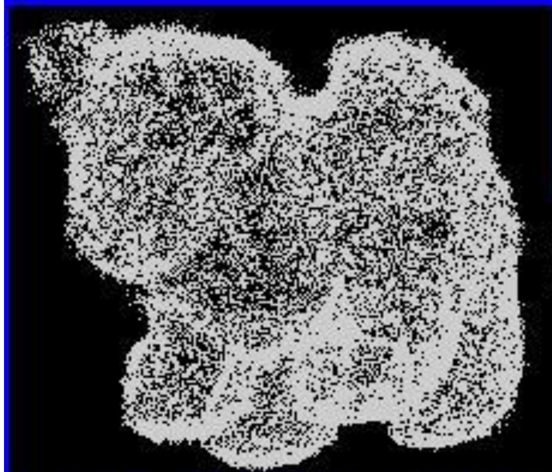

Registered points

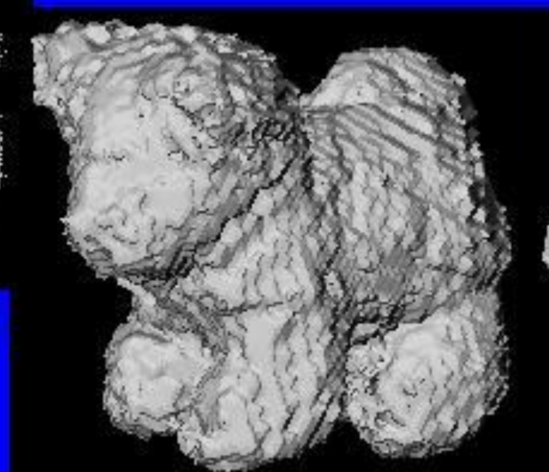

Initial mesh

Optimized mesh

View dependent texturing

Our viewer

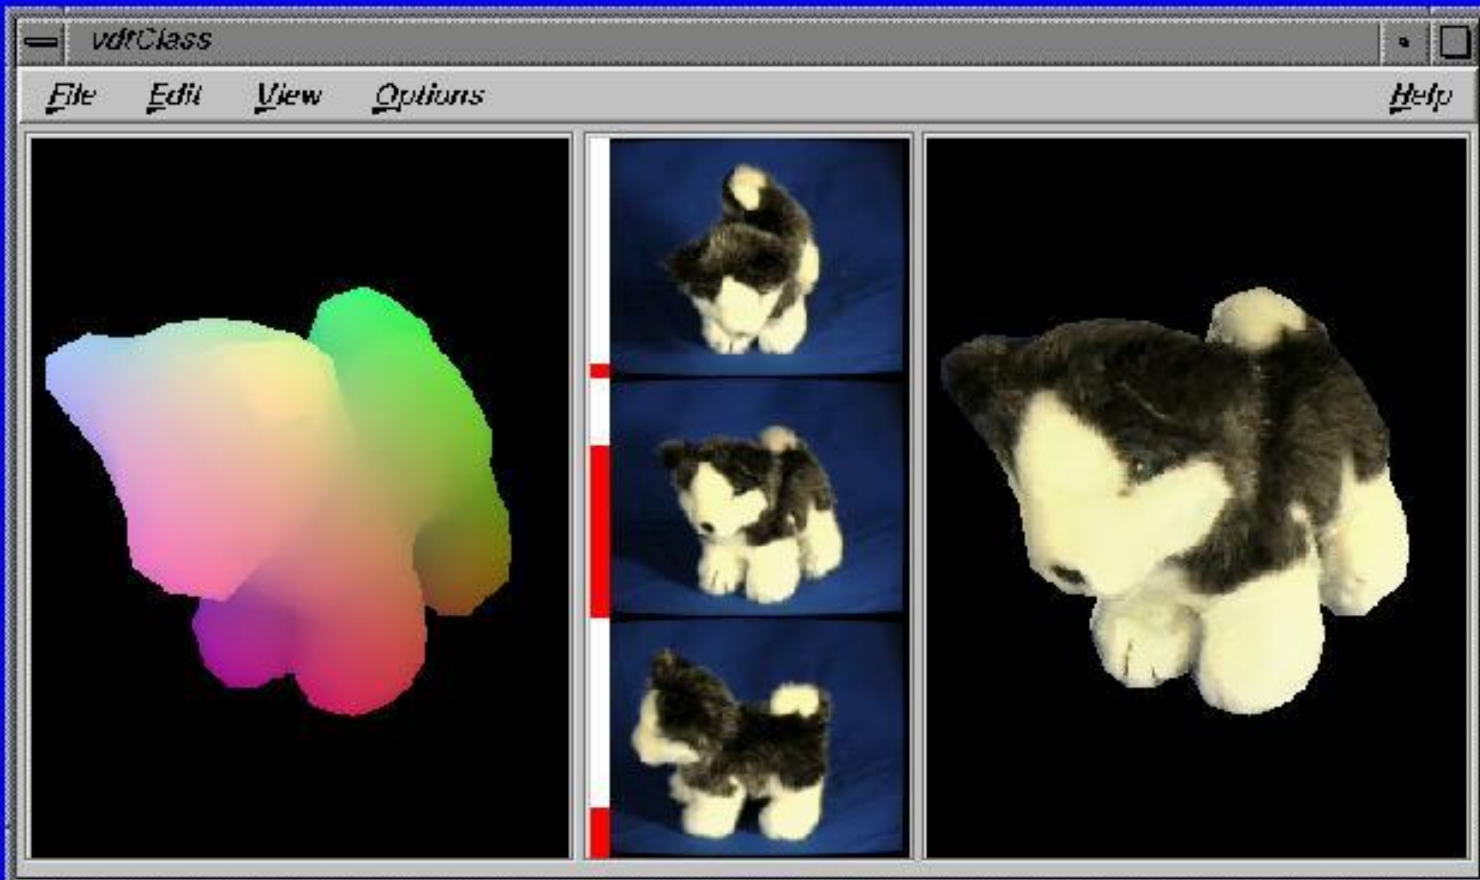

More: Space Carving Results: African Violet

Input Image (1 of 45)

Reconstruction

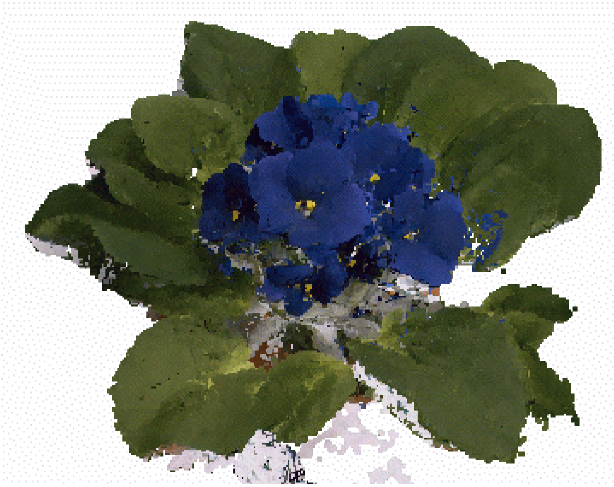

Reconstruction

Reconstruction

Source: S. Seitz

More: Space Carving Results: Hand

Input Image
(1 of 100)

Views of Reconstruction

Stereo from community photo collections

- Up to now, we've always assumed that camera calibration is known
- For photos taken from the Internet, we need *structure from motion* techniques to reconstruct both camera positions and 3D points. (SEE POSTED VIDEO)

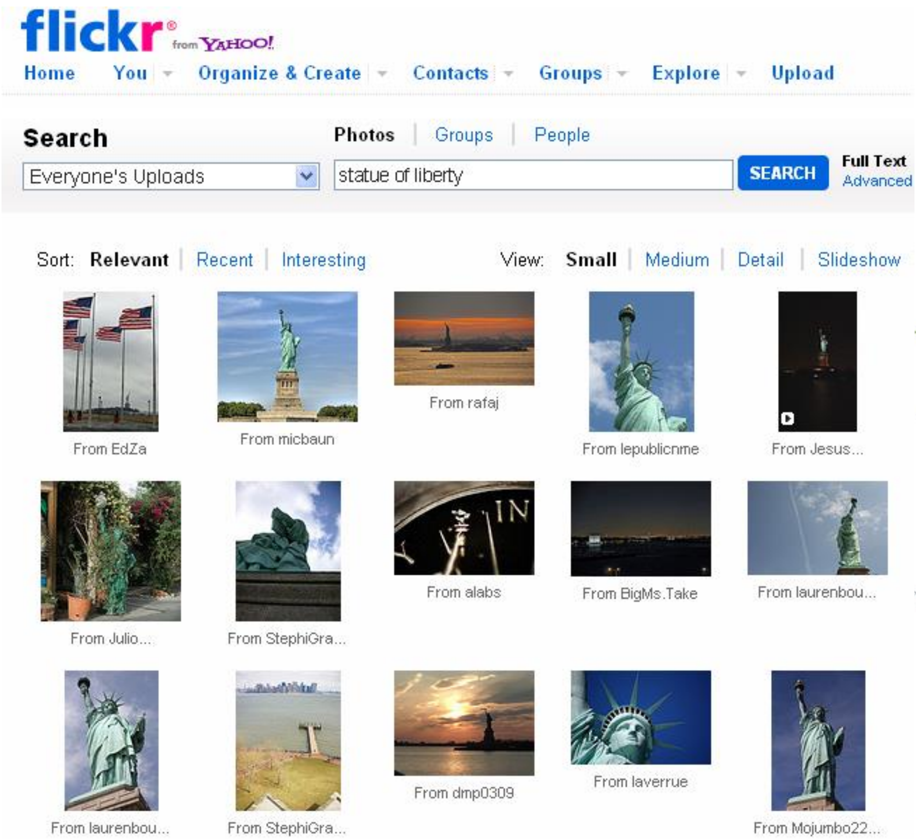

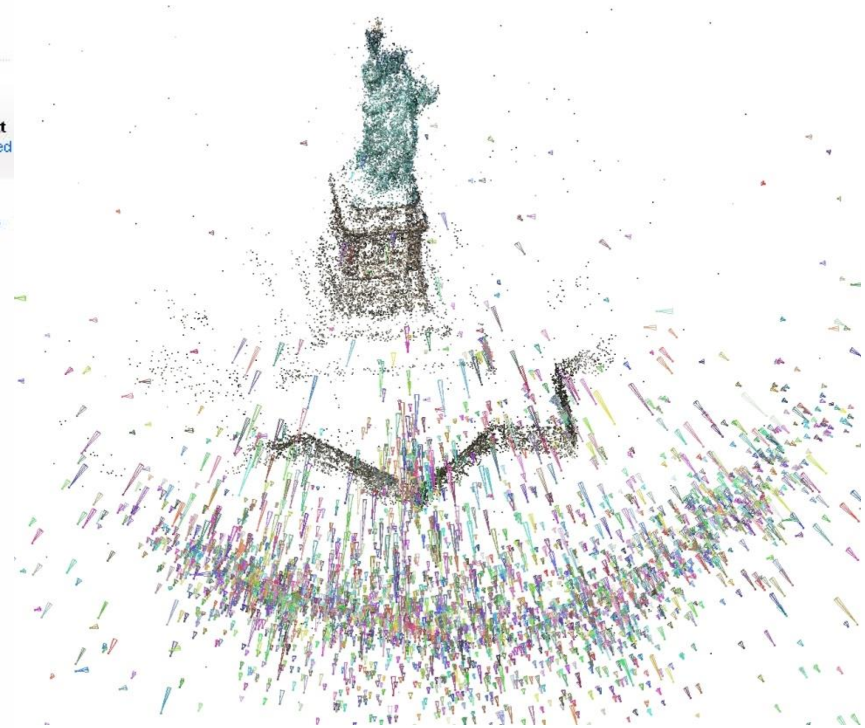

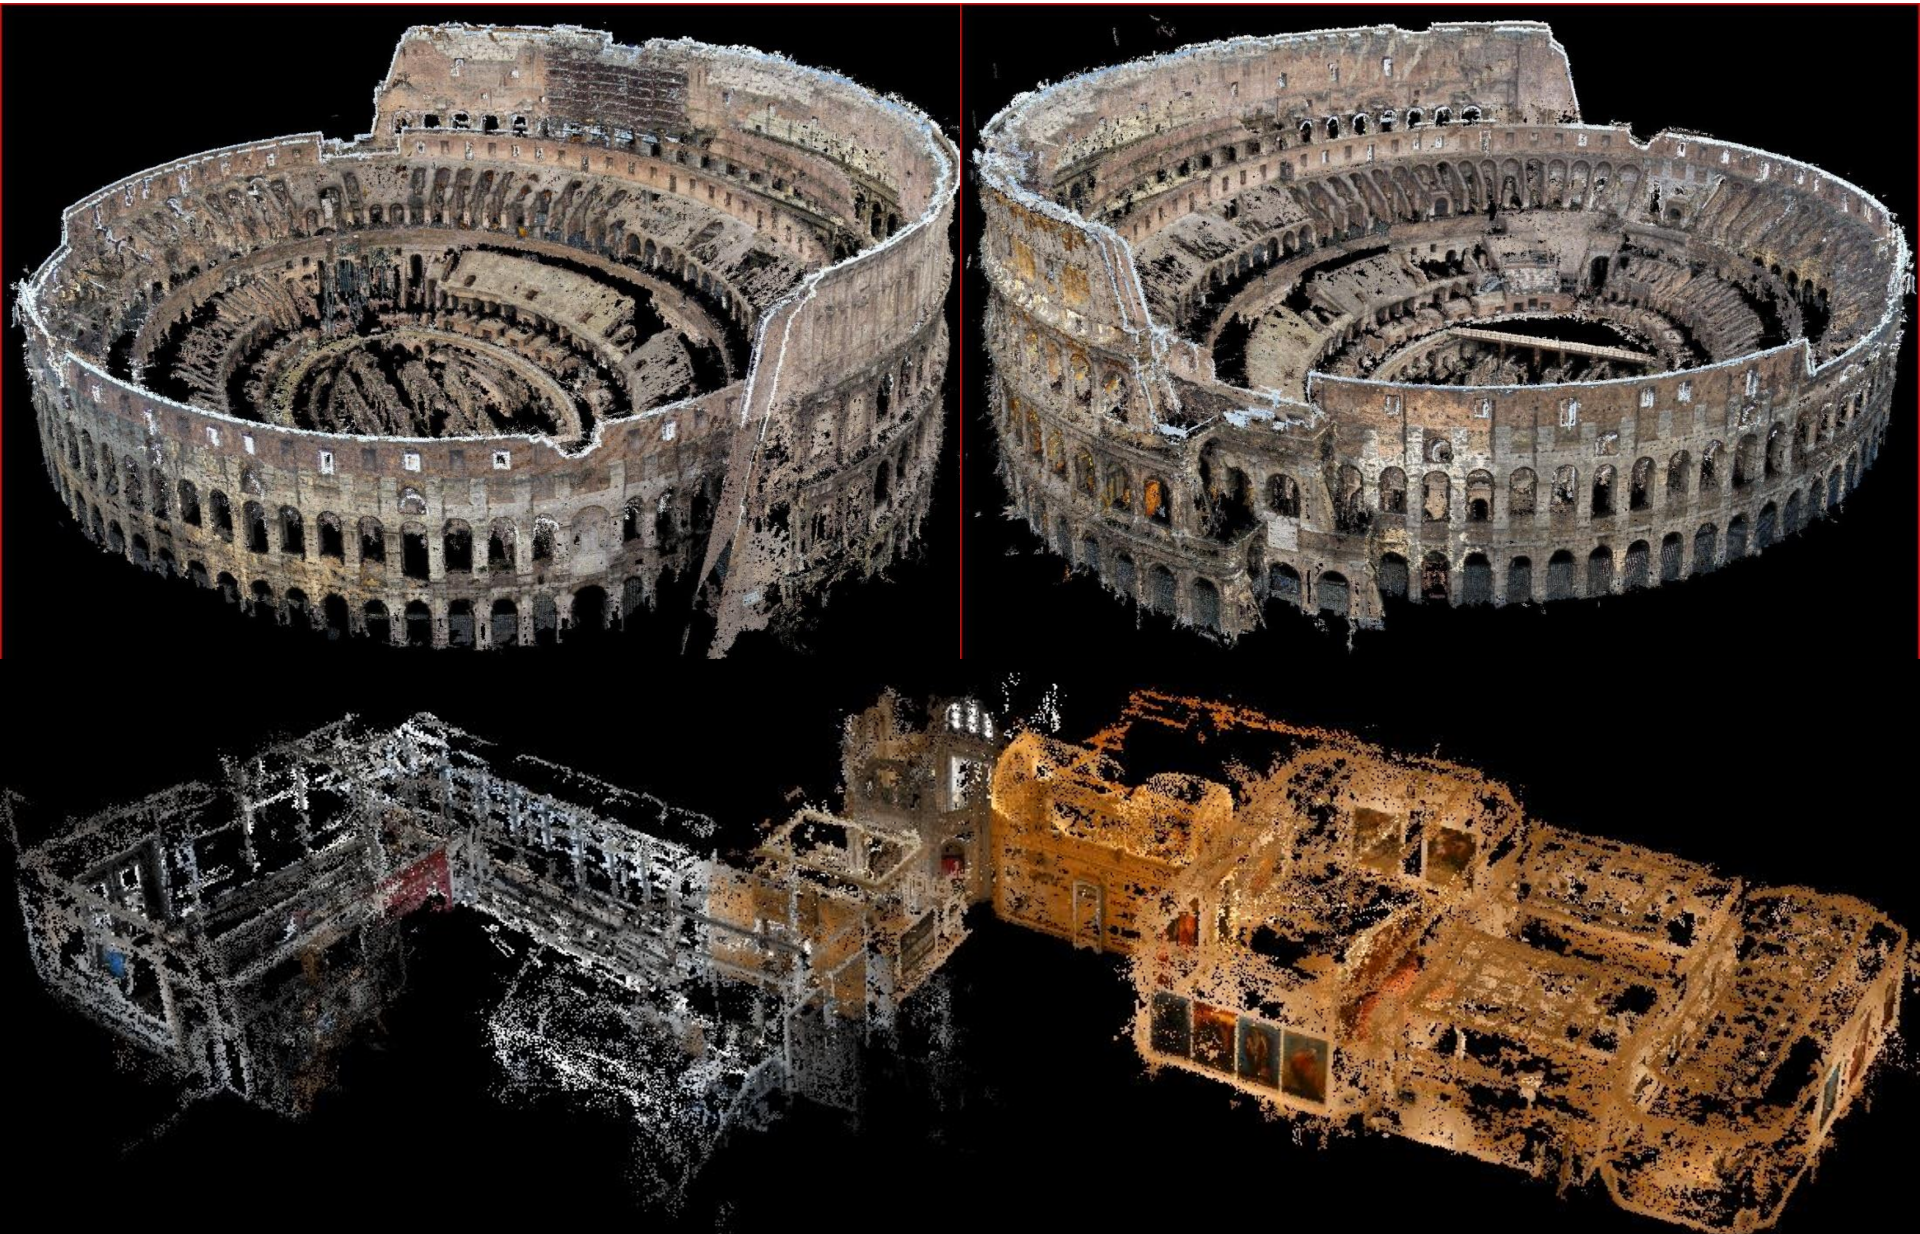

Head Reconstruction from Uncalibrated Internet Photos

Input: Internet photos in different poses and expressions

Output: 3D model of the head

work of
Shu Liang

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.

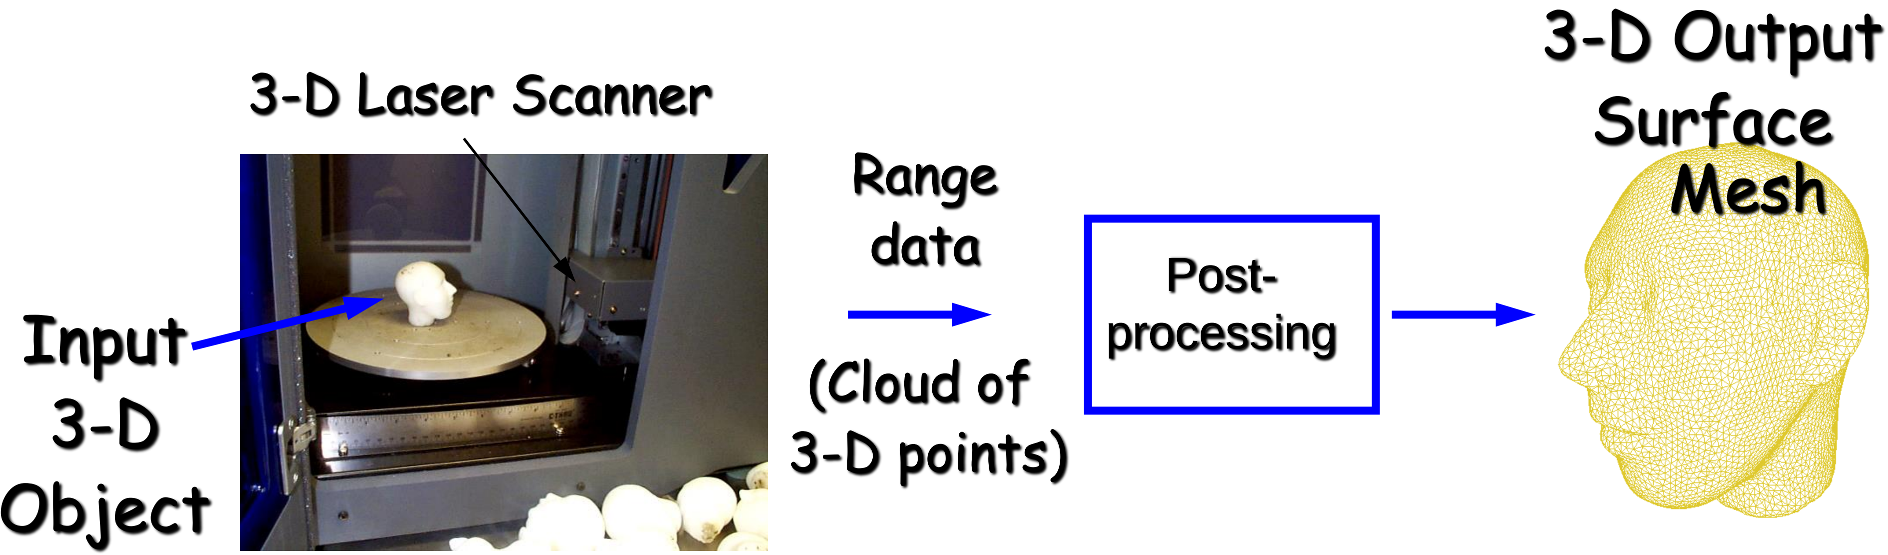

- This is a difficult problem that is relevant in several application fields.

What Kind Of Deformations?

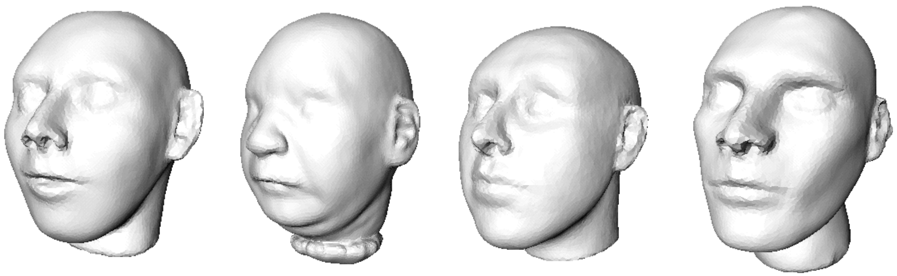

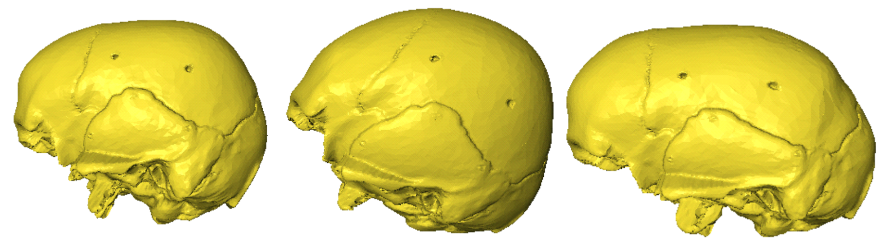

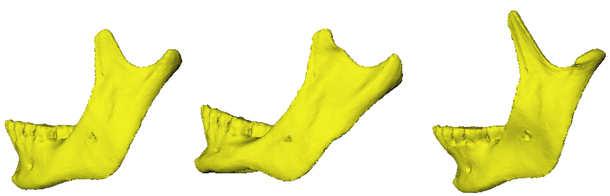

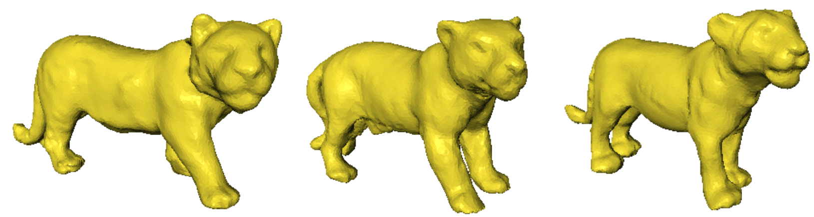

Component-Based Methodology

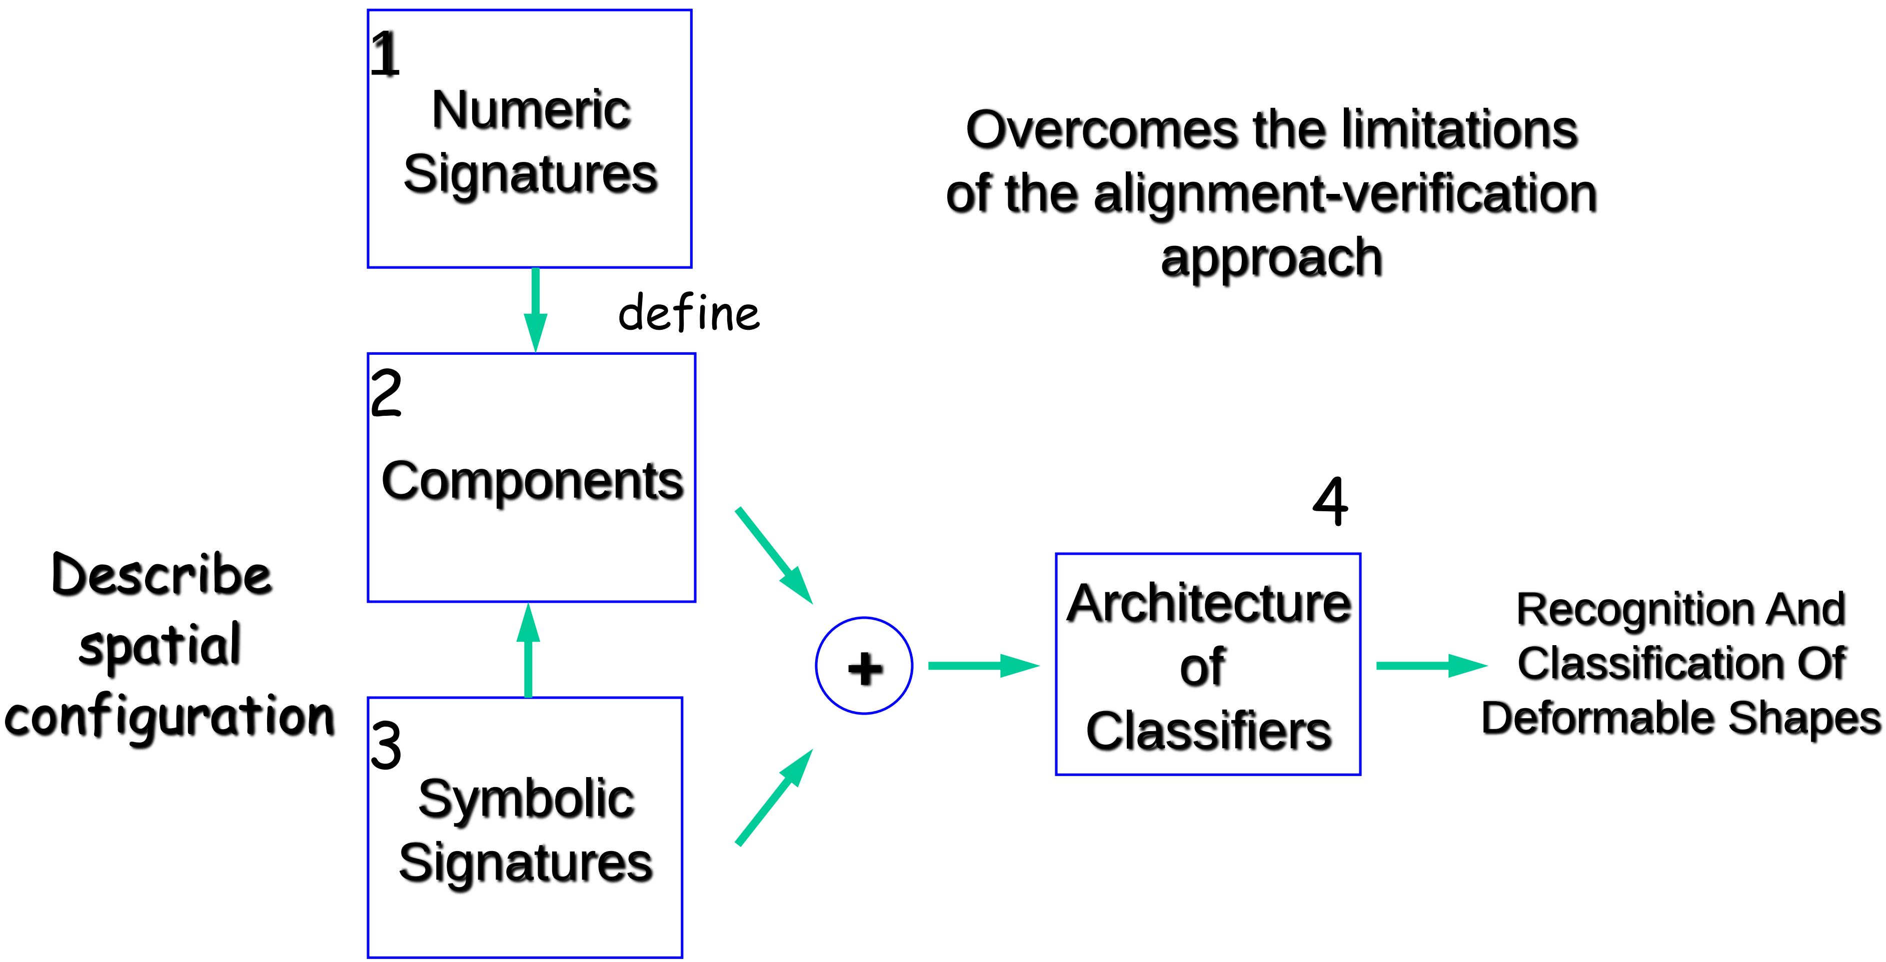

Assumptions

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

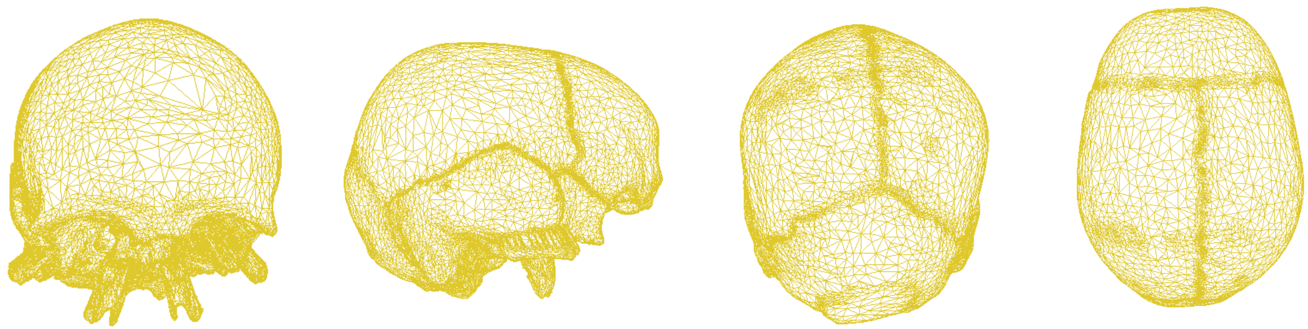

The vertices of the meshes in the training set are in full correspondence.

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

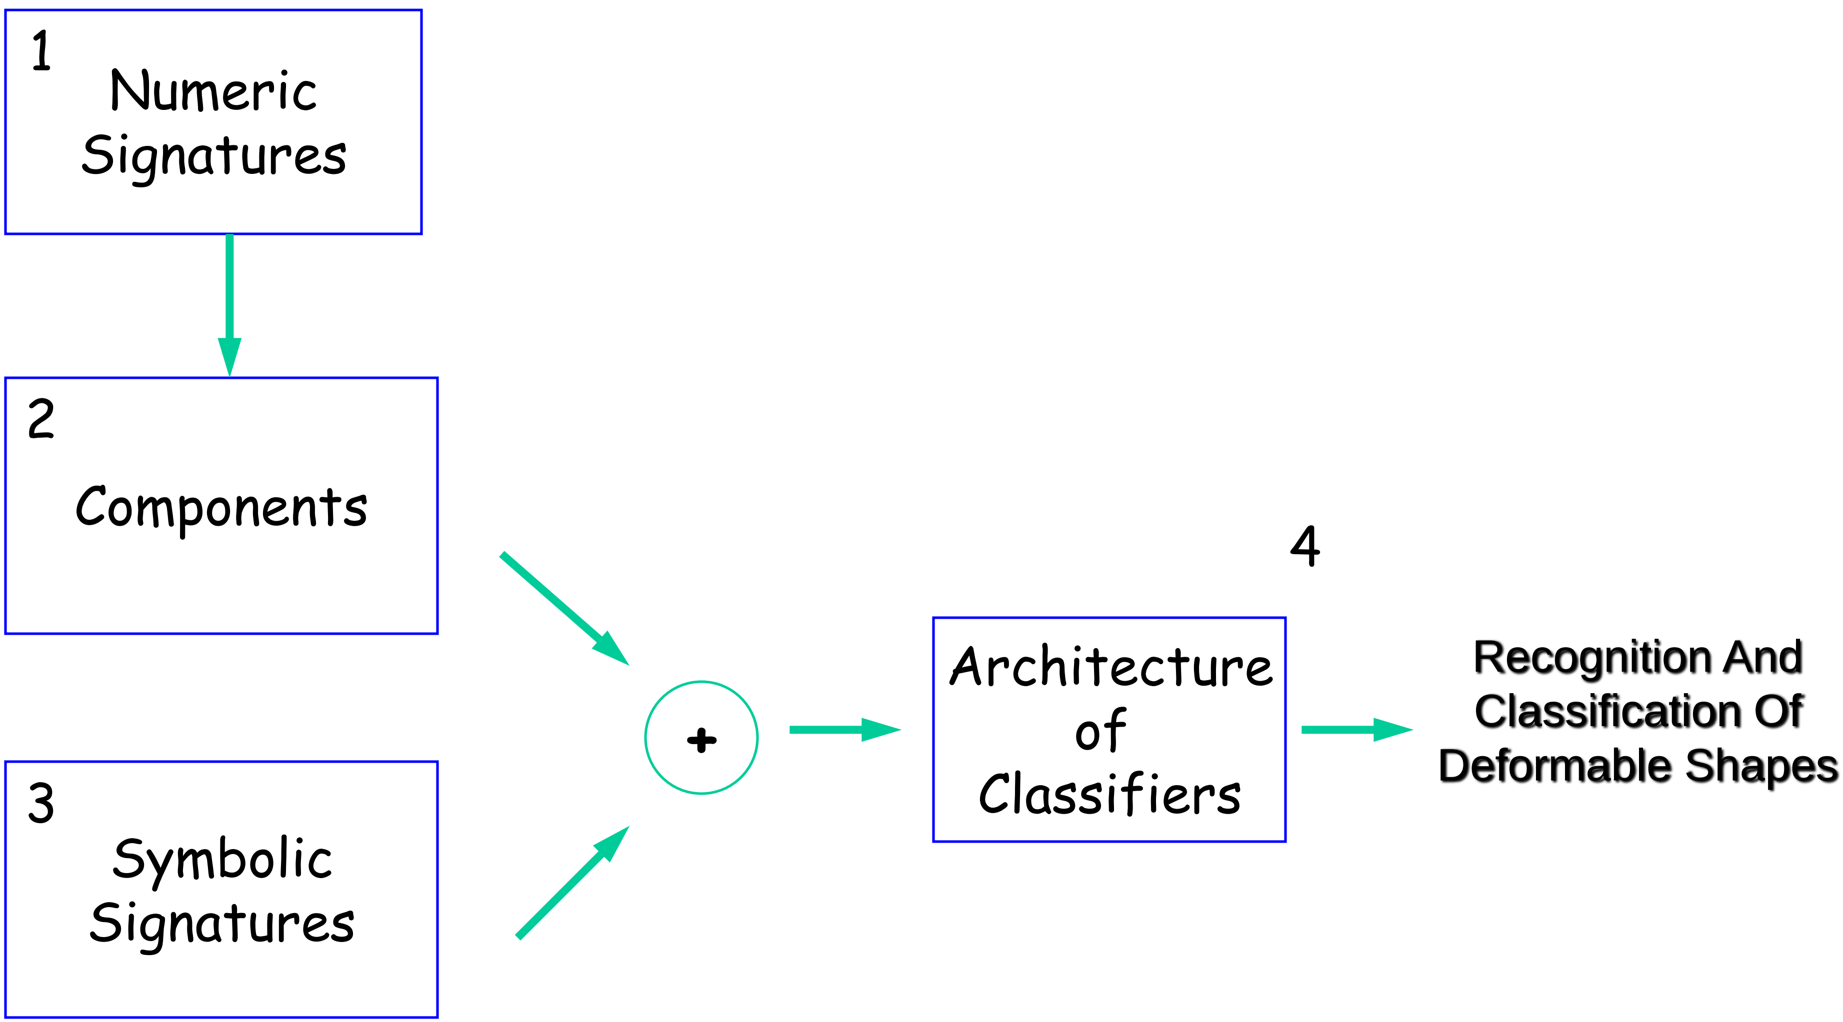

Numeric Signatures

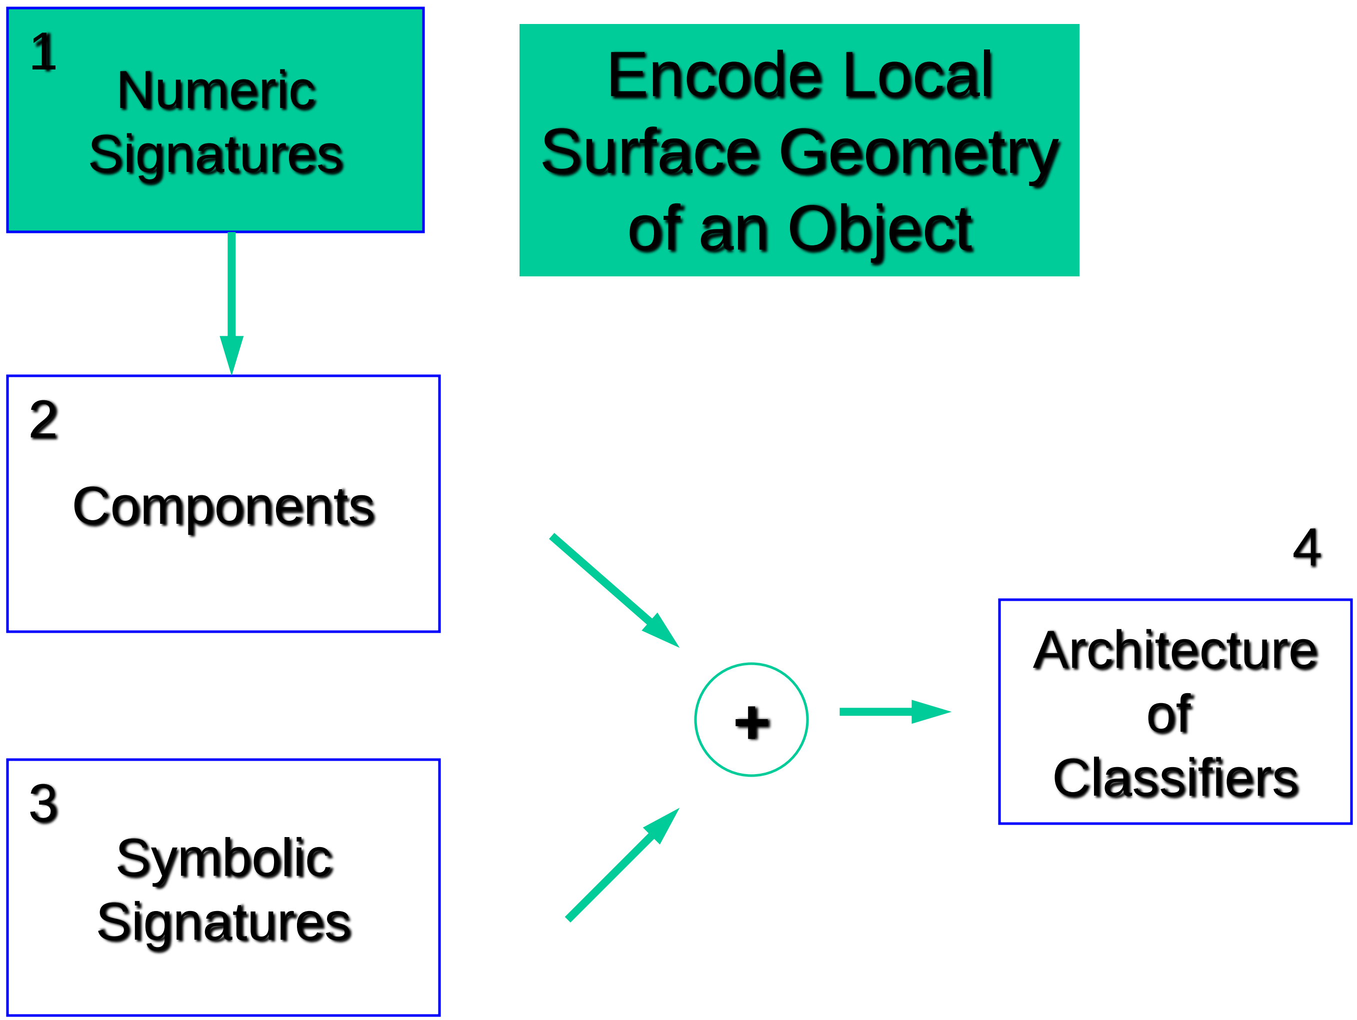

The Spin Image Signature

P is the selected vertex.

X is a contributing point
of the mesh.

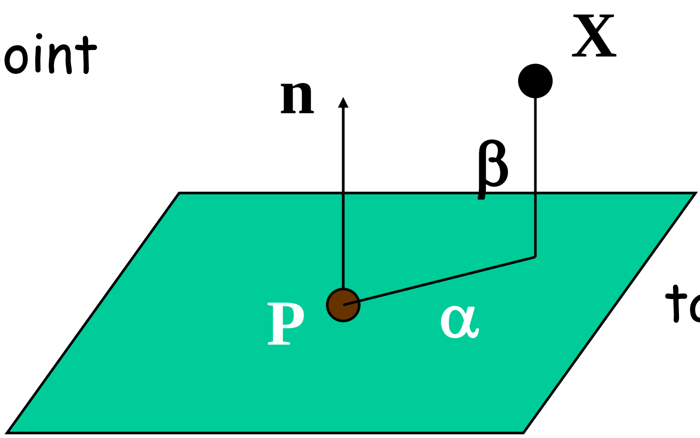

tangent plane at P

α is the perpendicular distance from X to P's surface normal.

β is the signed perpendicular distance from X to P's tangent plane.

Spin Image Construction

- A spin image is constructed
  - about a specified oriented point o of the object surface
  - with respect to a set of contributing points C, which is controlled by maximum distance and angle from o.

- It is stored as an array of accumulators $S(\alpha,\beta)$ computed via:

- For each point c in C(o)

  1. compute α and β for c.
  2. increment $S(\alpha,\beta)$

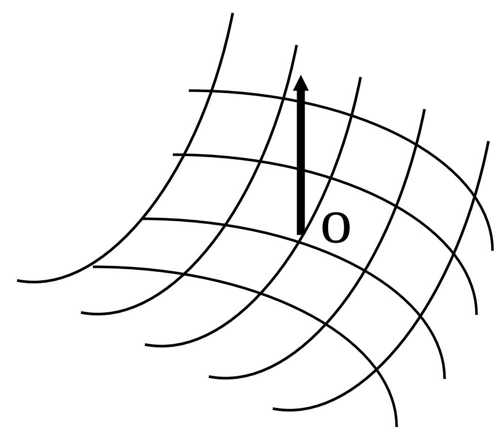

Numeric Signatures: Spin Images

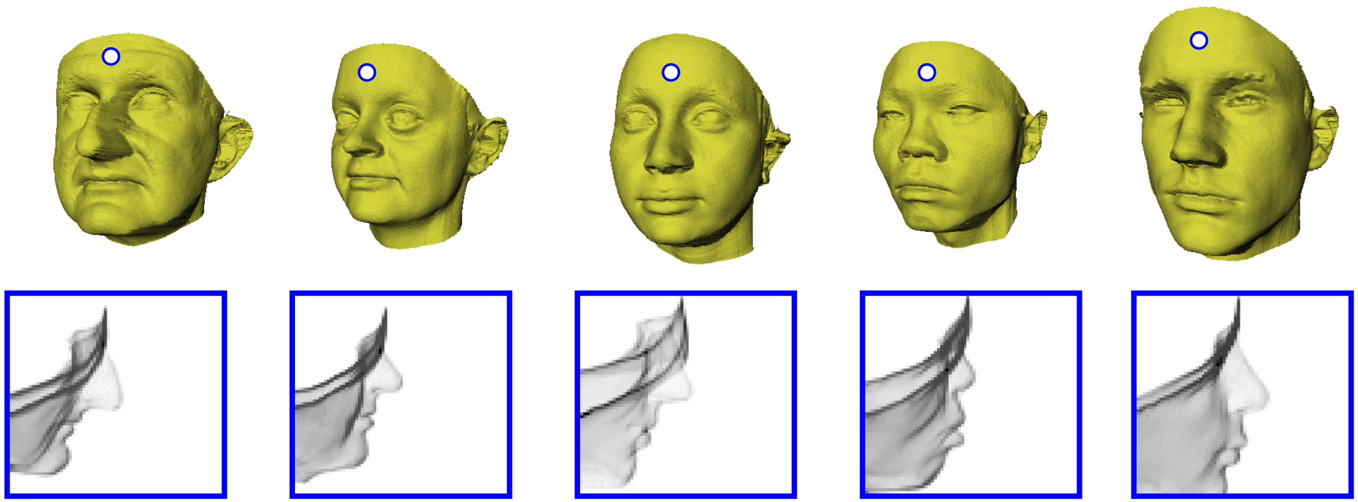

Rich set of surface shape descriptors.

Their spatial scale can be modified to include local and non-local surface features.

Representation is robust to scene clutter and occlusions.

Components

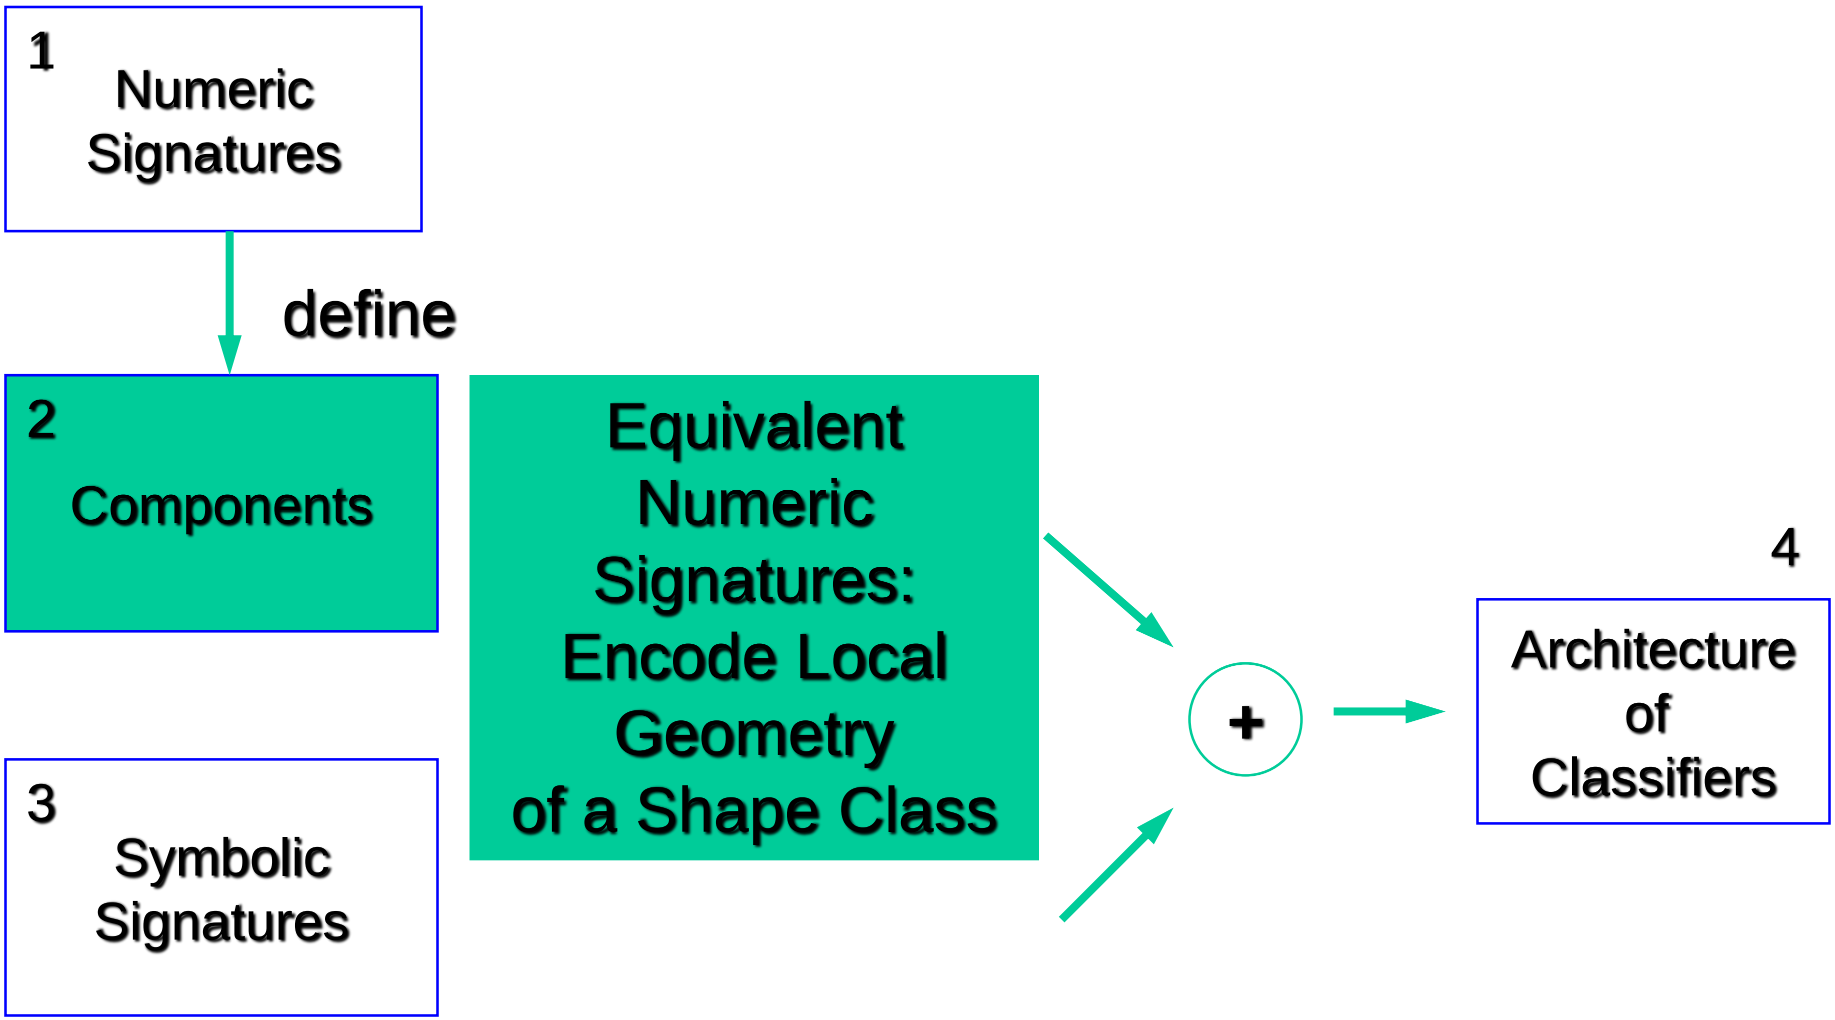

1
Numeric Signatures
define
2
Components
Equivalent Numeric Signatures: Encode Local Geometry of a Shape Class
3
Symbolic Signatures
+
4
Architecture of Classifiers

How To Extract Shape Class Components?

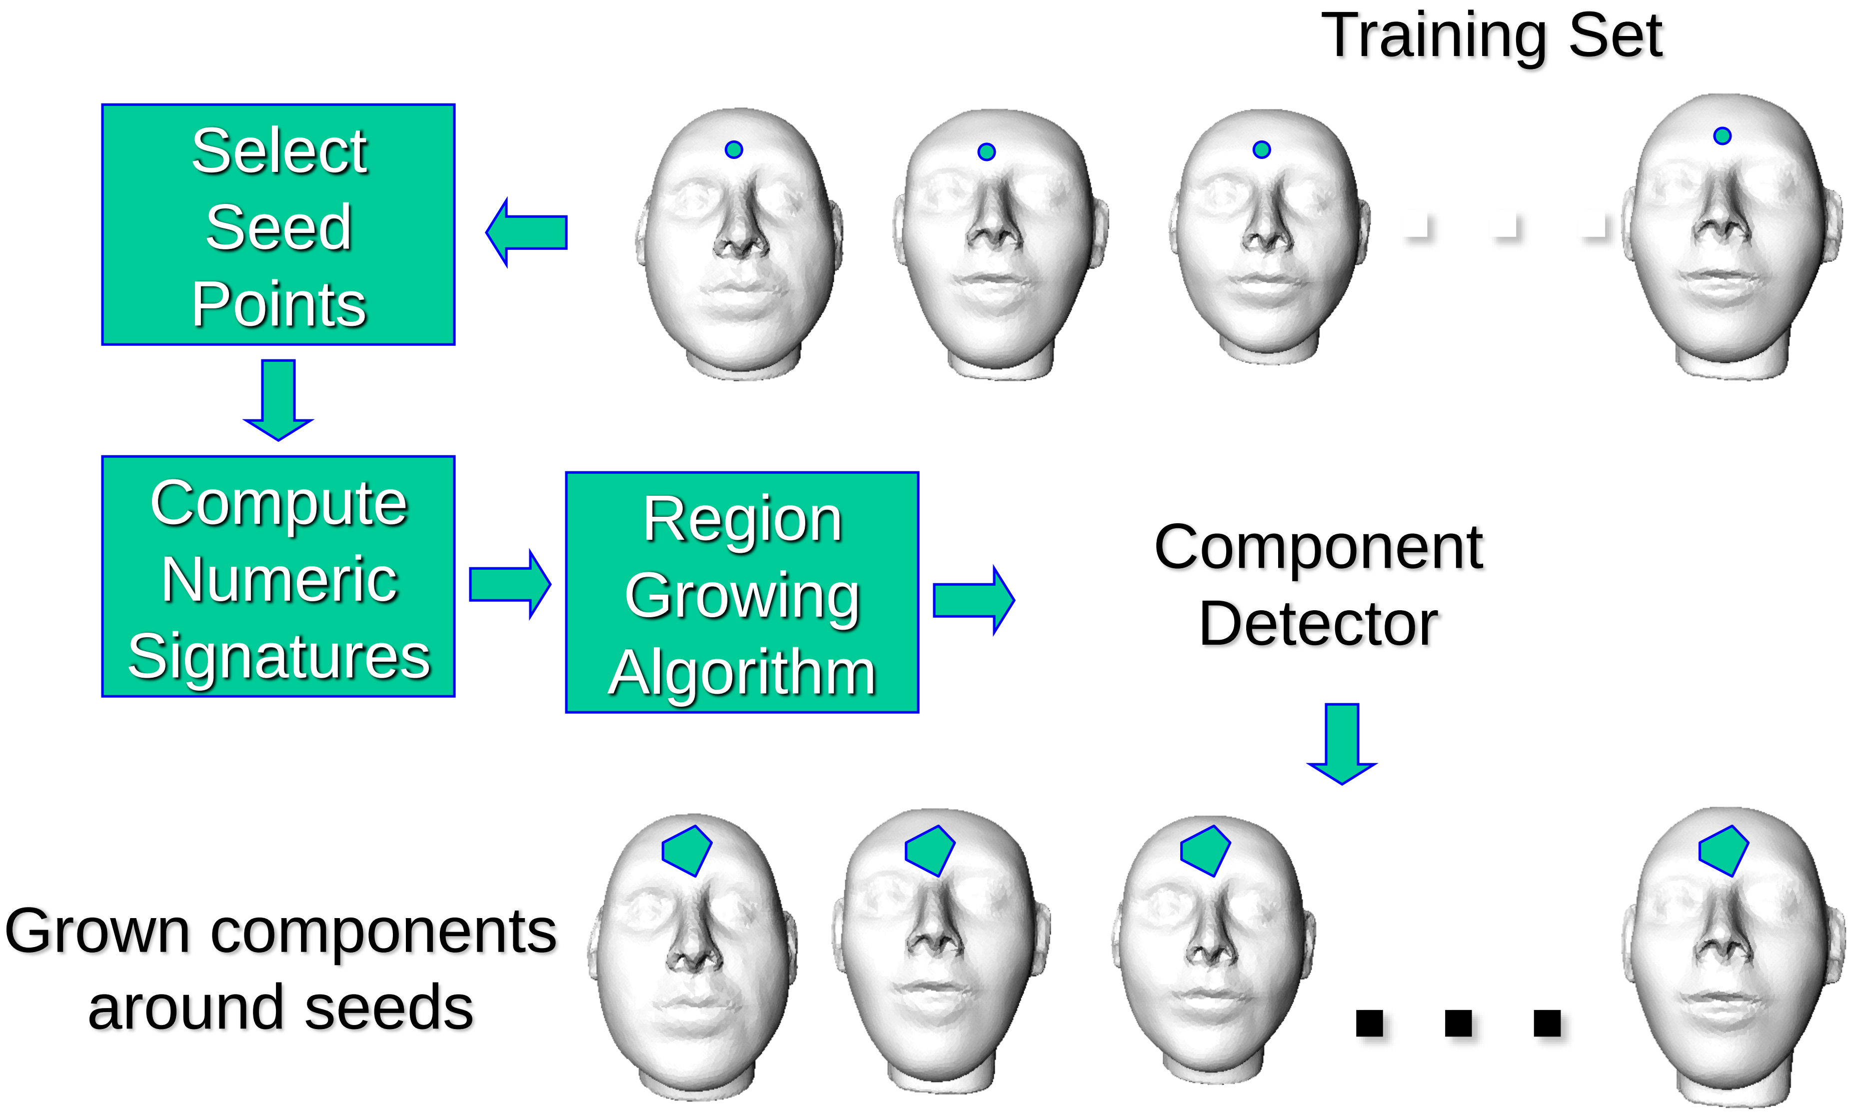

Component Extraction Example

Selected 8 seed points by hand

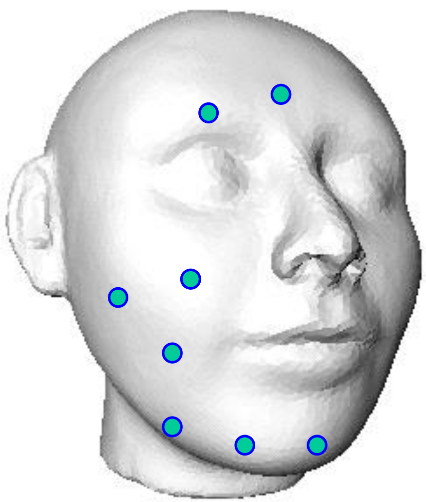

Grow one region at the time
(get one detector
per component)

Region
Growing

Labeled
Surface Mesh

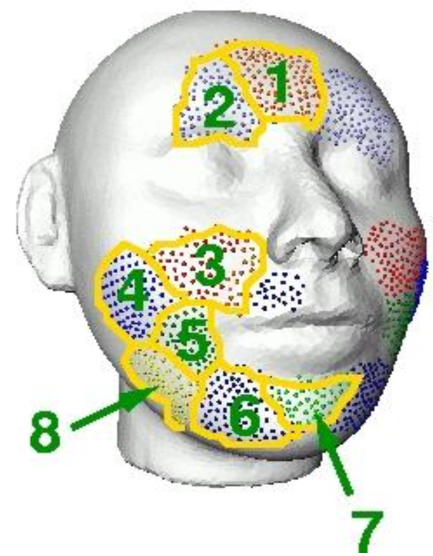

Detected
components on a
training sample

How To Combine Component Information?

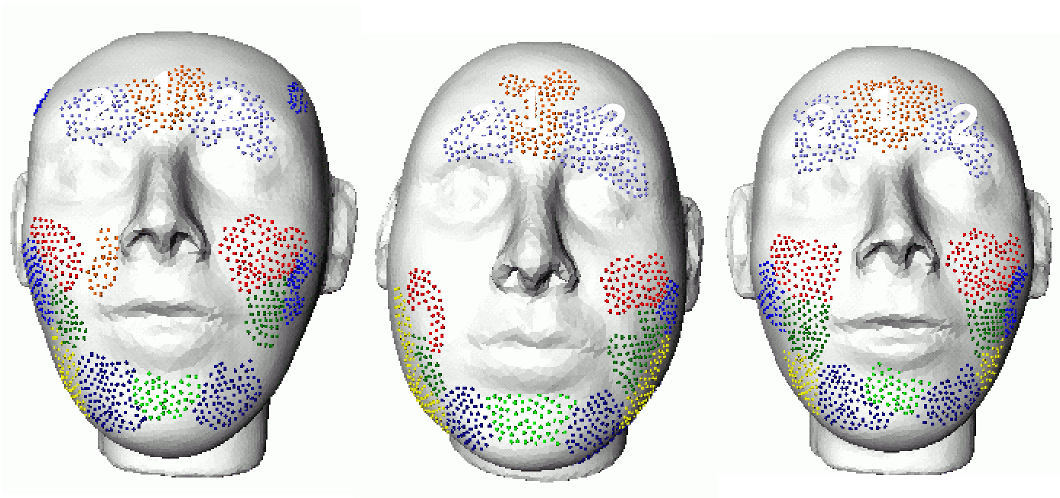

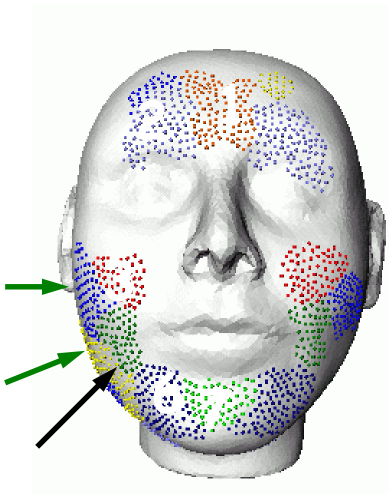

Symbolic Signatures

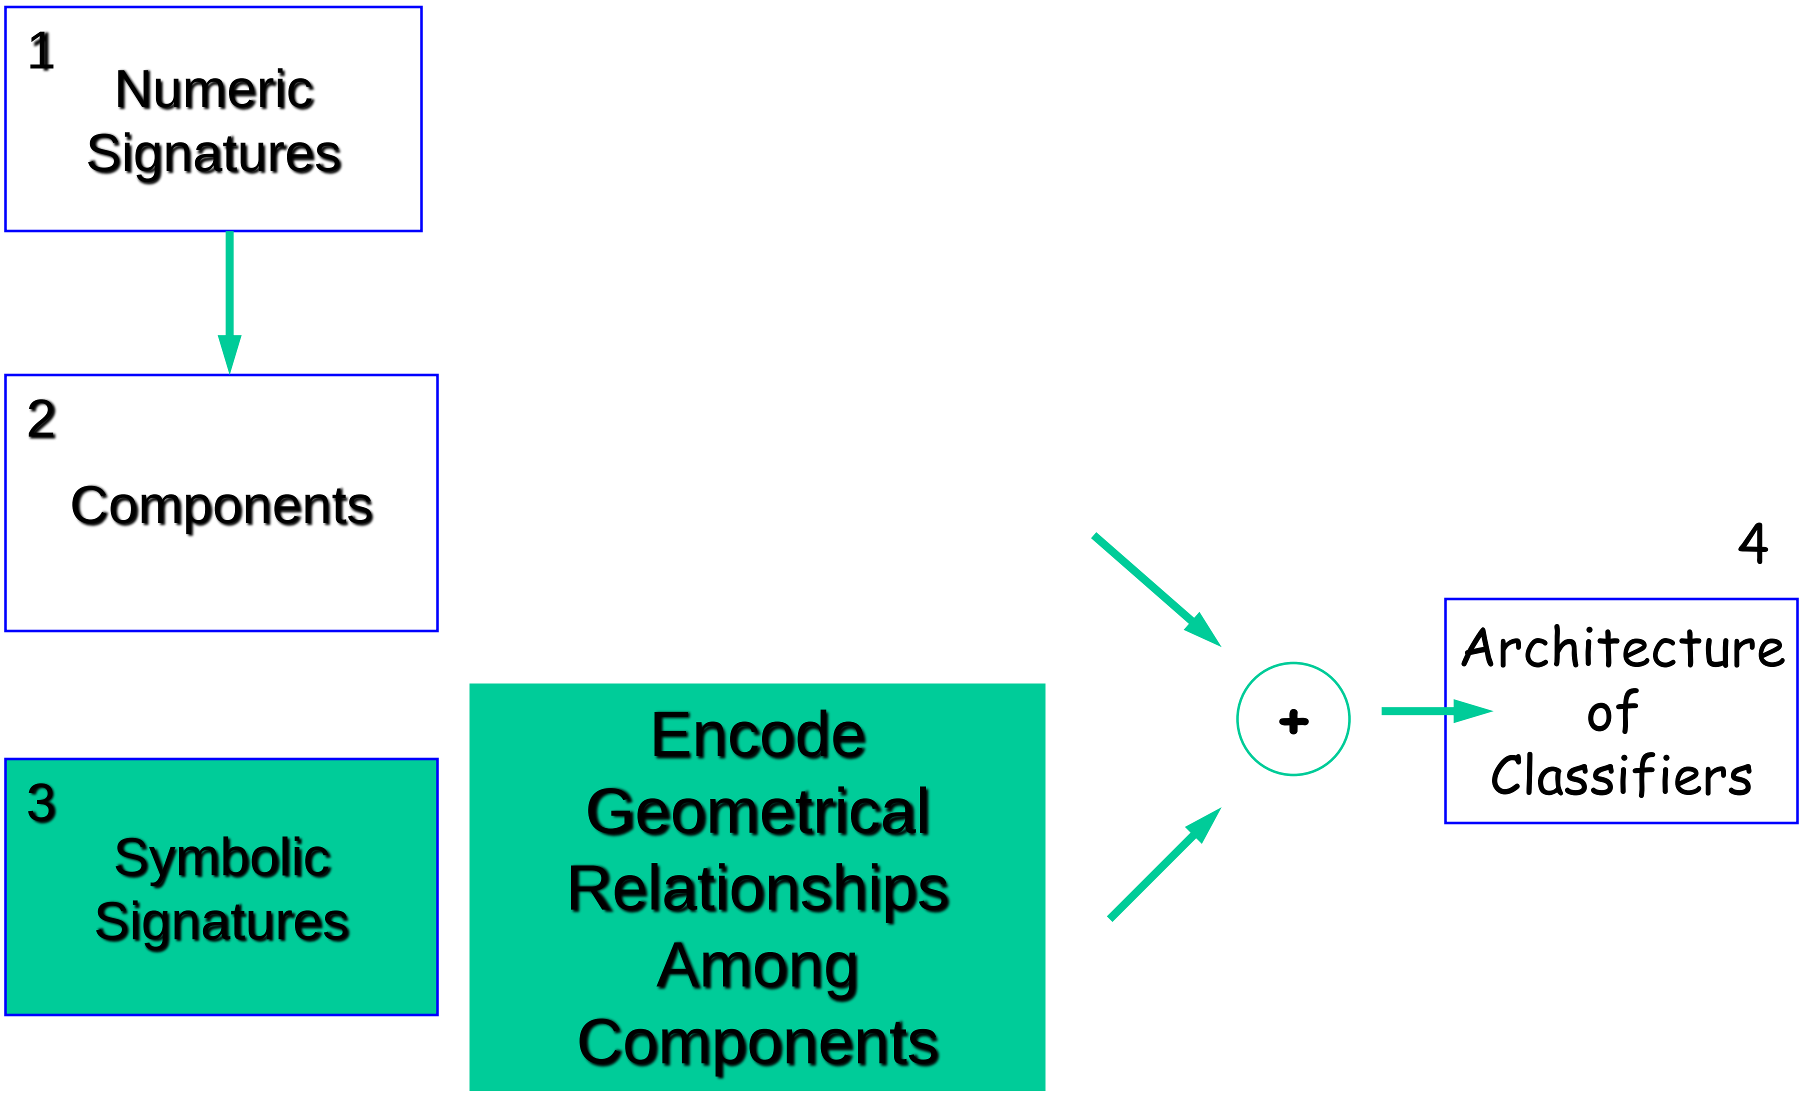

Symbolic Signature

Labeled
Surface Mesh

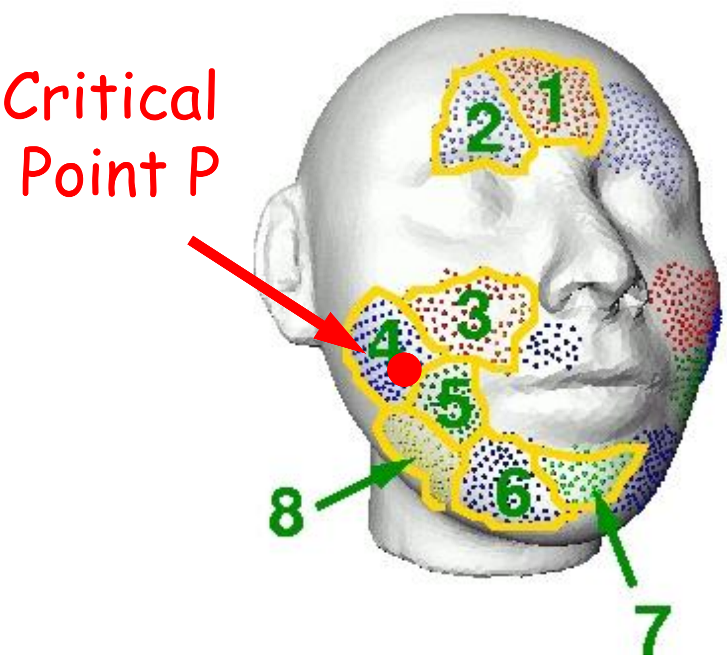

Encode
Geometric
Configuration

Symbolic
Signature at P

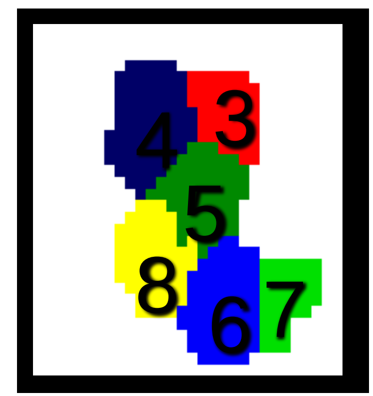

Matrix storing
component
labels

Symbolic Signatures Are Robust To Deformations

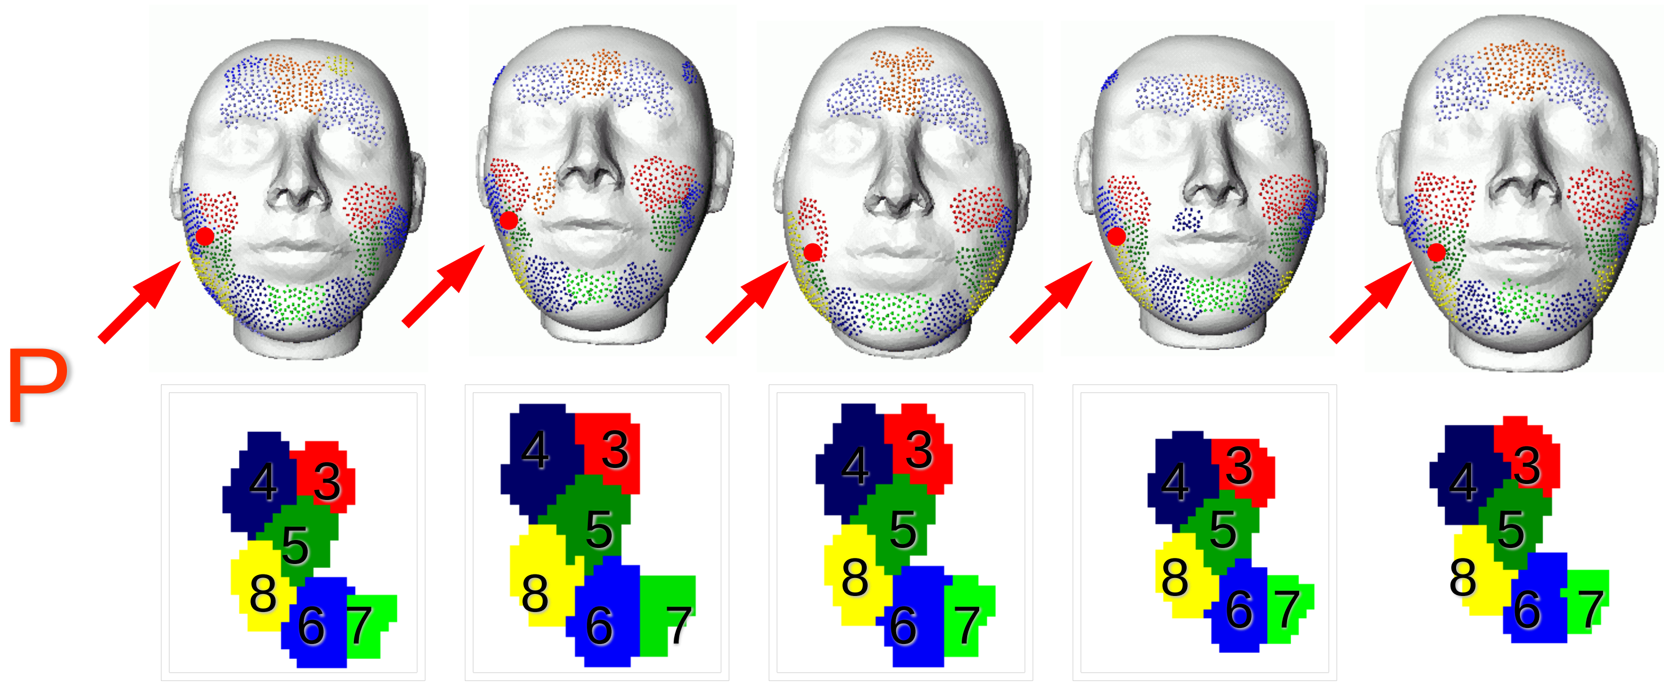

Relative position of components is stable across deformations: experimental evidence

Architecture of Classifiers

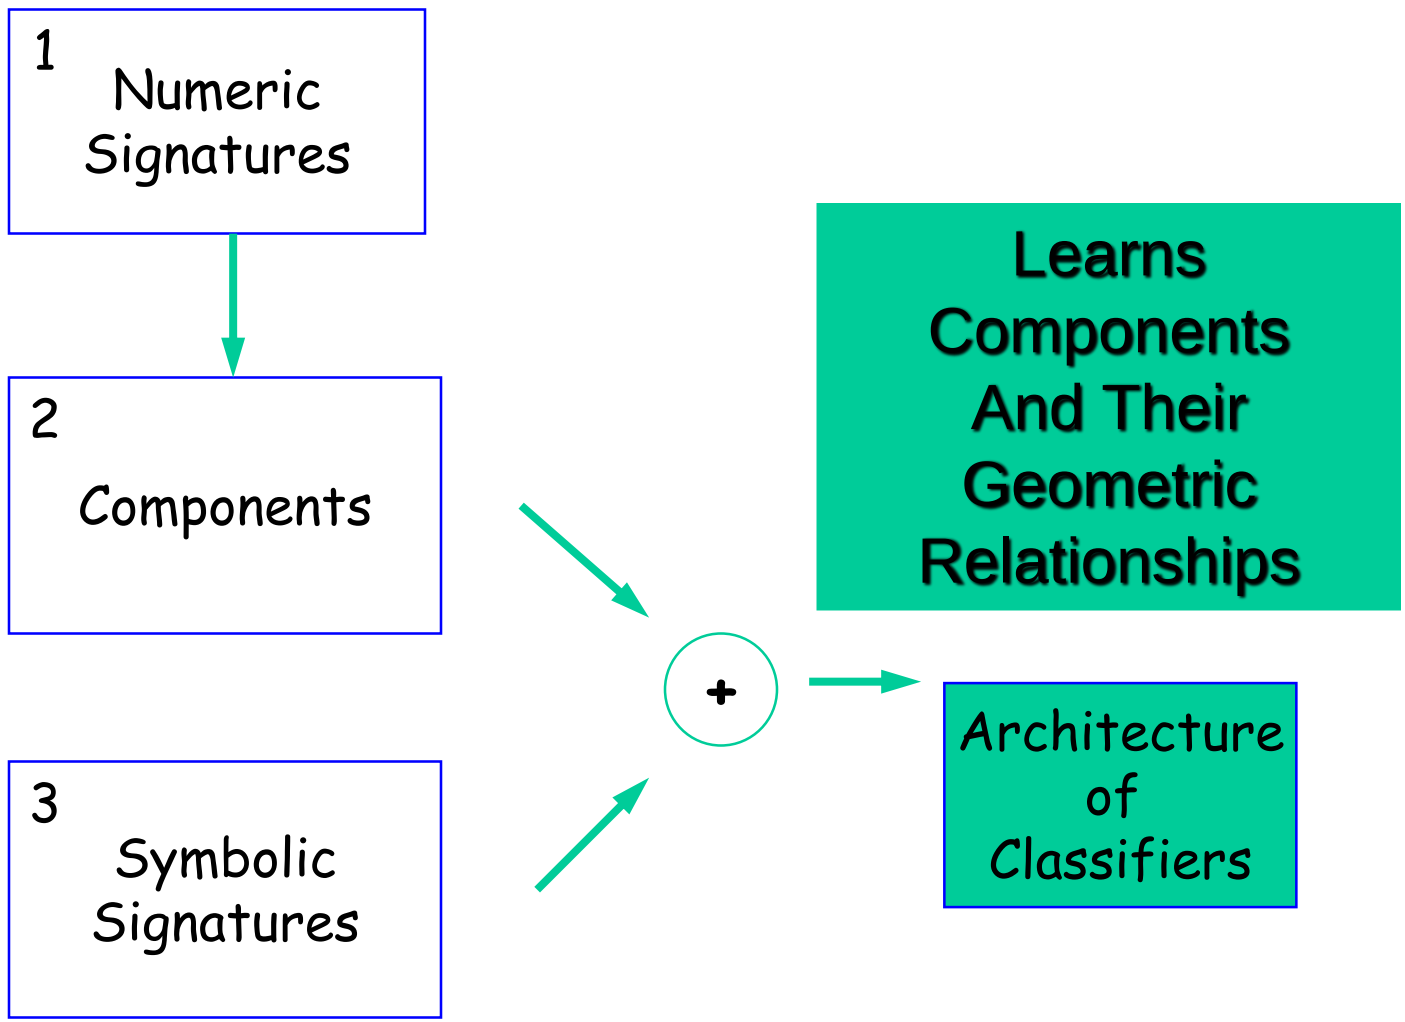

Proposed Architecture

Verify spatial configuration
of the components

Input → **Identify Components** → Labeled Mesh → **Identify Symbolic Signatures** → Class Label

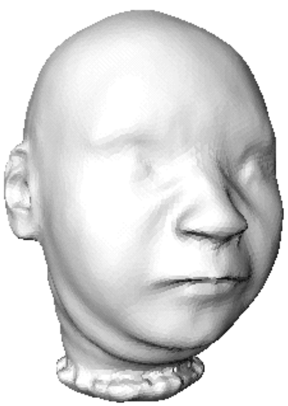

Surface
Mesh

-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

Rotary Table

Recognition

Laser

Classification

Shape Classes

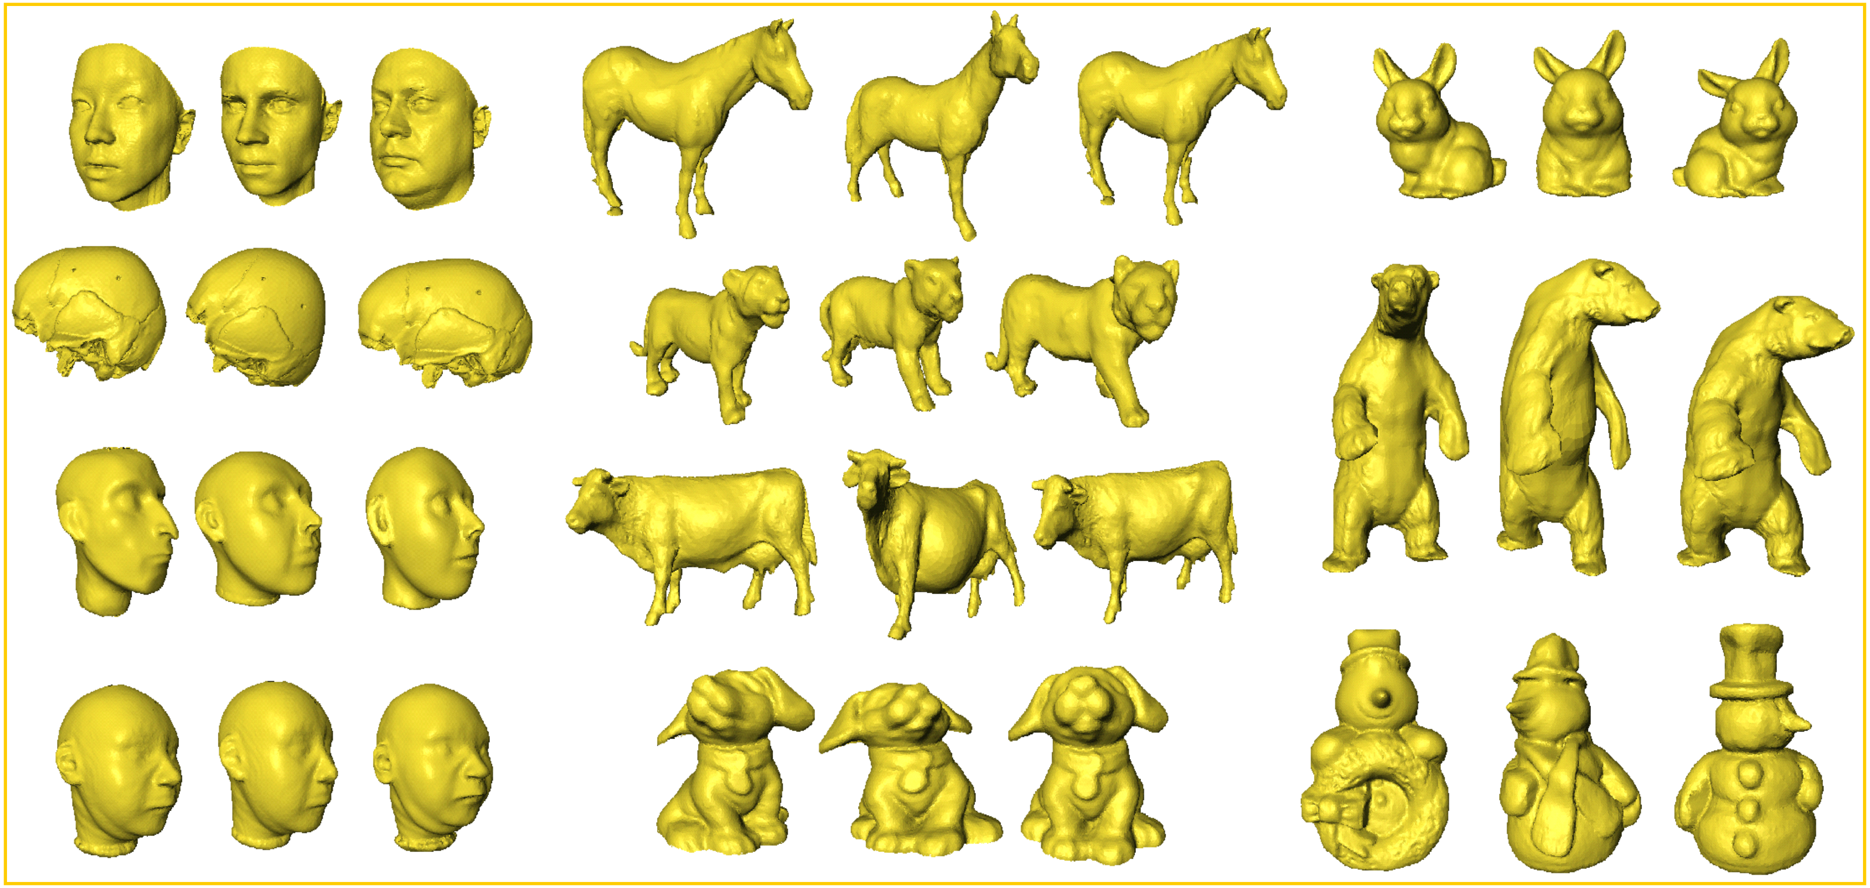

Enlarging Training Sets Using Virtual Samples

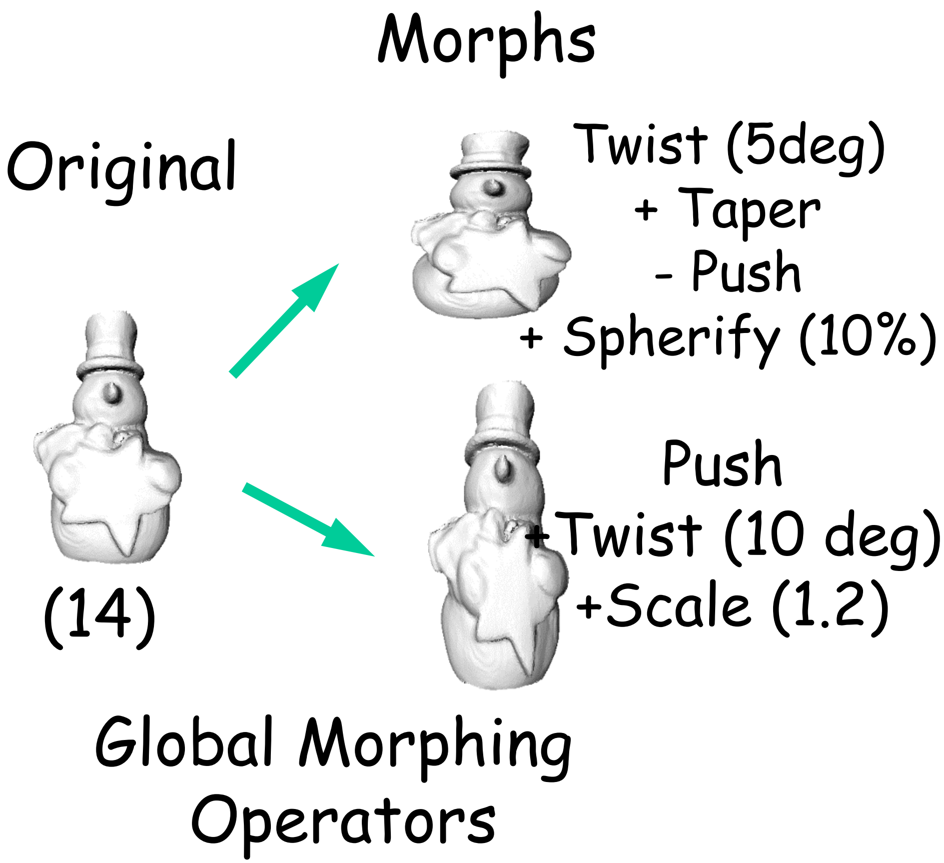

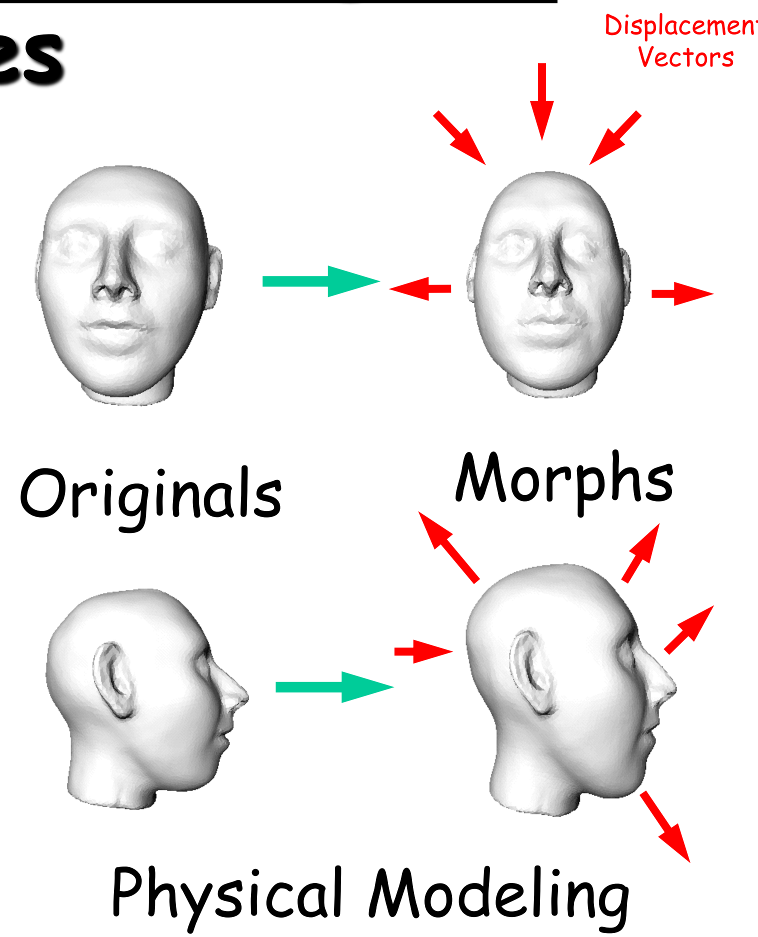

Task 1: Recognizing Single Objects (1)

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%.

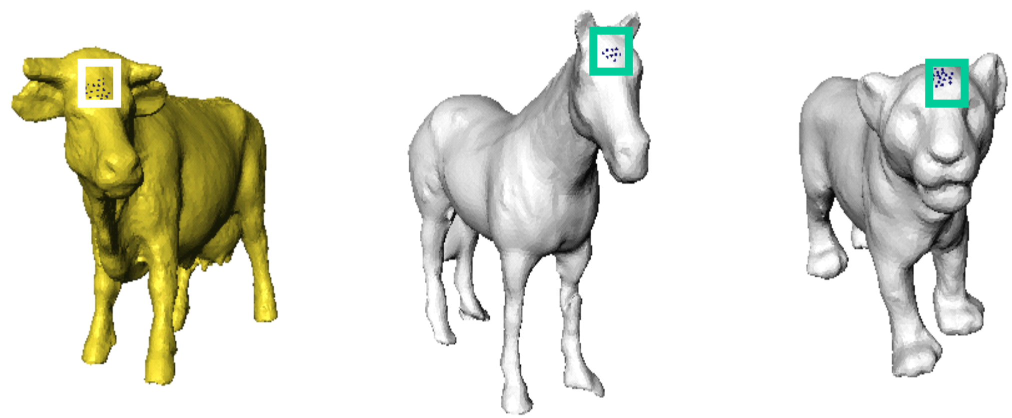

Recognition rates (true positives)

(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)

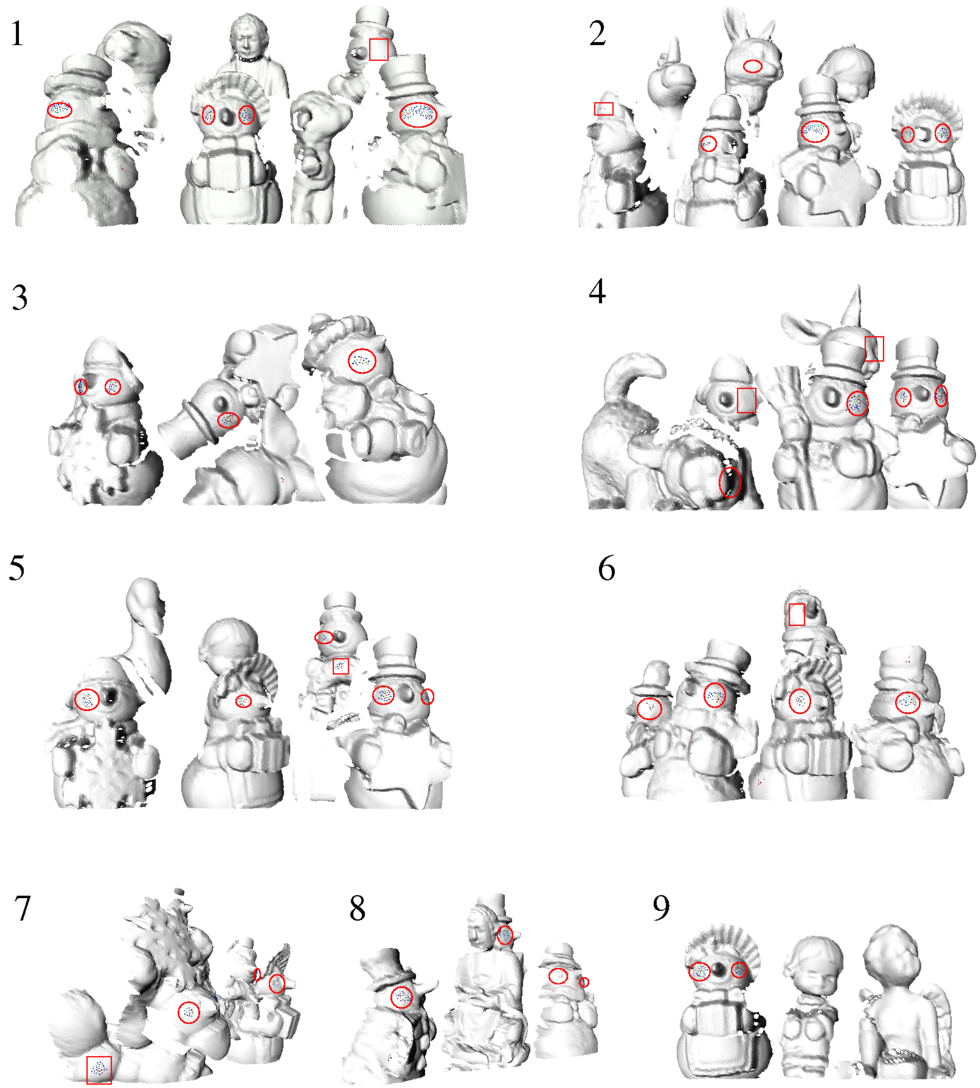

Task 4: Recognizing Human Heads (3)

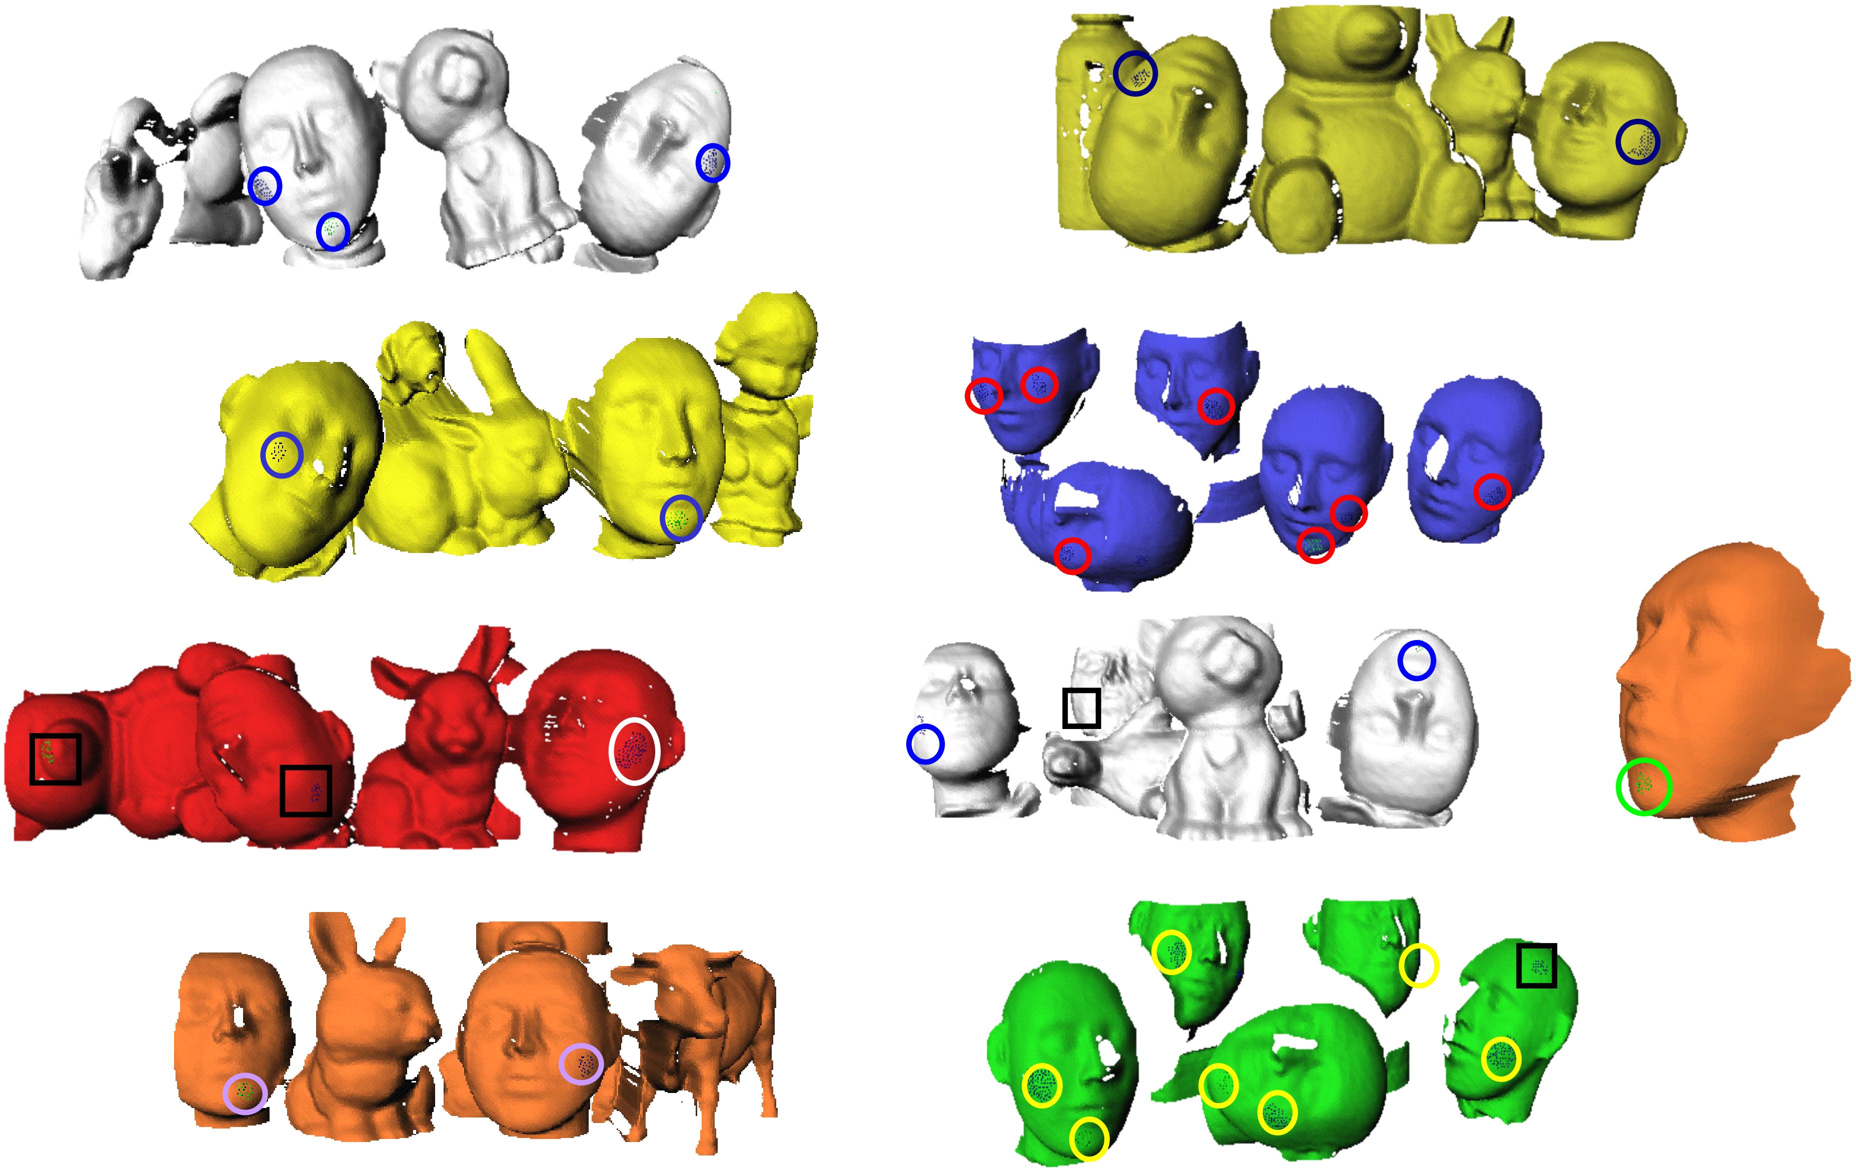

Task 5: Classifying Normal vs. Abnormal Human Heads (1)

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

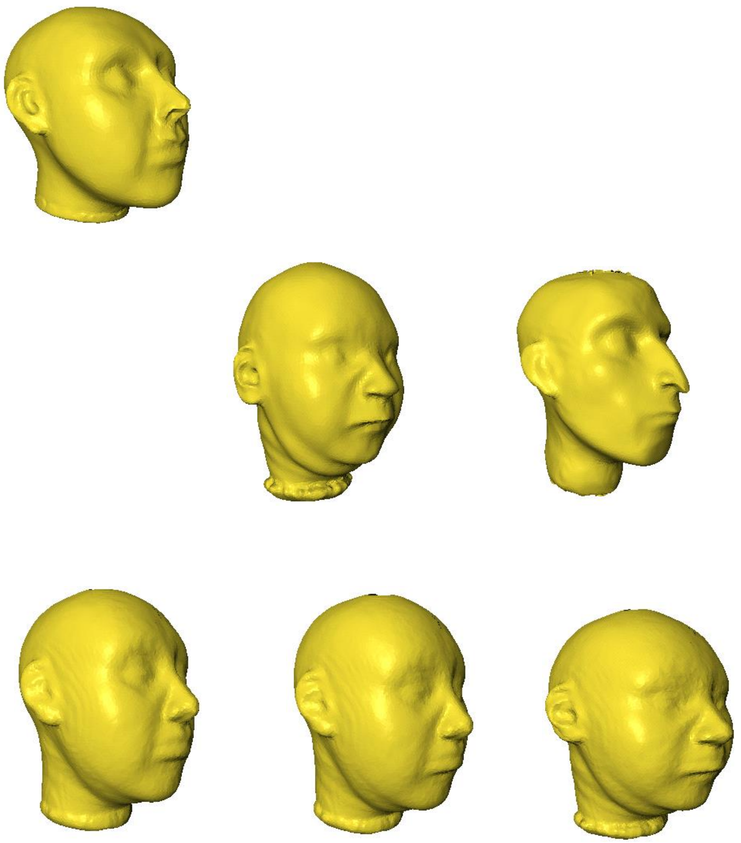

Task 6: Classifying Normal vs. Abnormal Human Heads In Complex Scenes(1)

Shape Classes	Classification Accuracy %
Normal vs. Abnormal 1	88

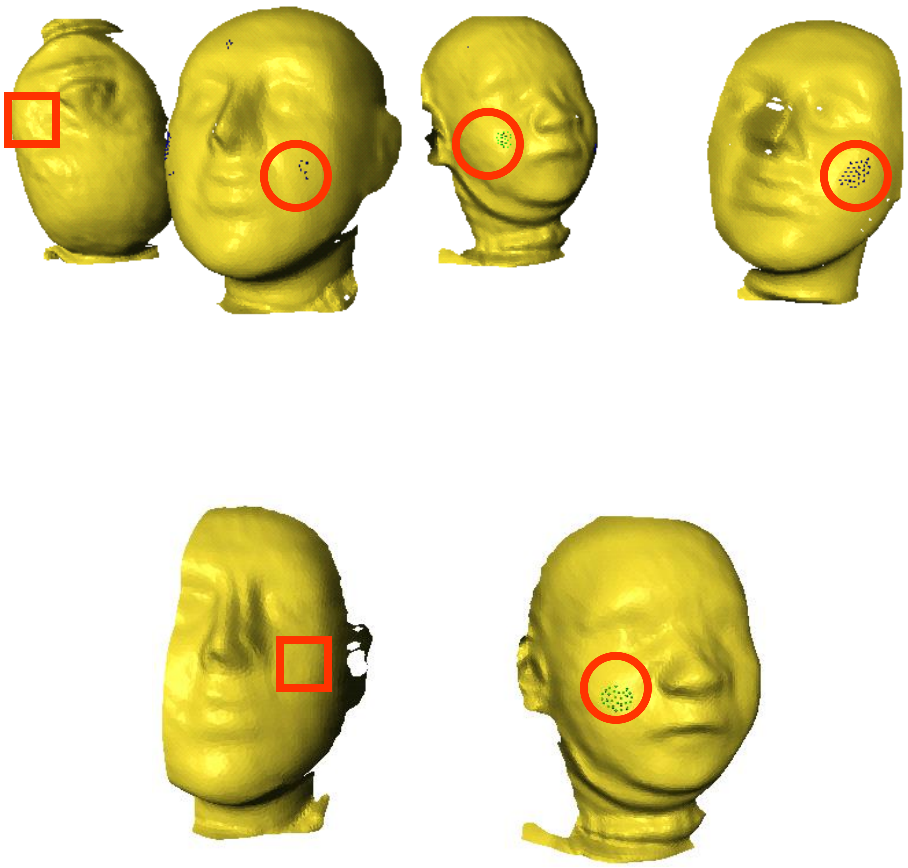

Task 7: Classifying Normal vs. Abnormal Neurocranium (2)

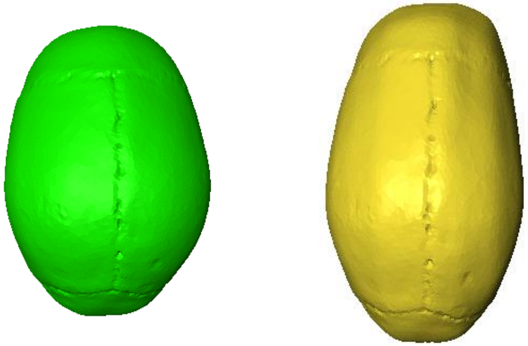

Shape Classes	Classification Accuracy %
Normal vs. Abnormal	89

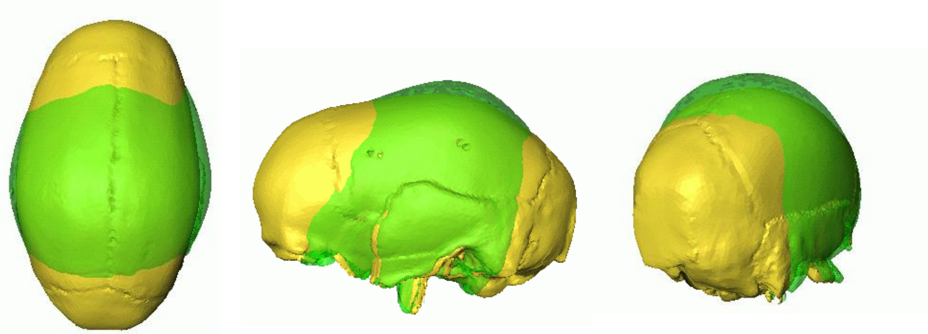

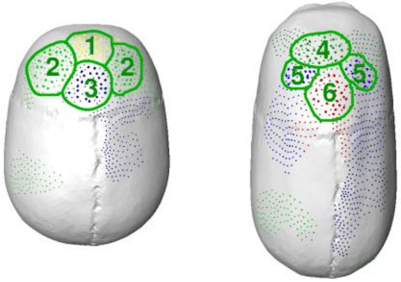

Main Contributions (1)

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.