Model-Based Organ Segmentation: Recent Methods

Jiun-Hung Chen

General Exam Paper

2009

Problem Statement

- Learn how to segment new, unseen CT images from a set of training CT images with ground truth organs marked.

- Goal: Minimize the training errors while generalizing to the new CT images

Problem Organ: The Liver

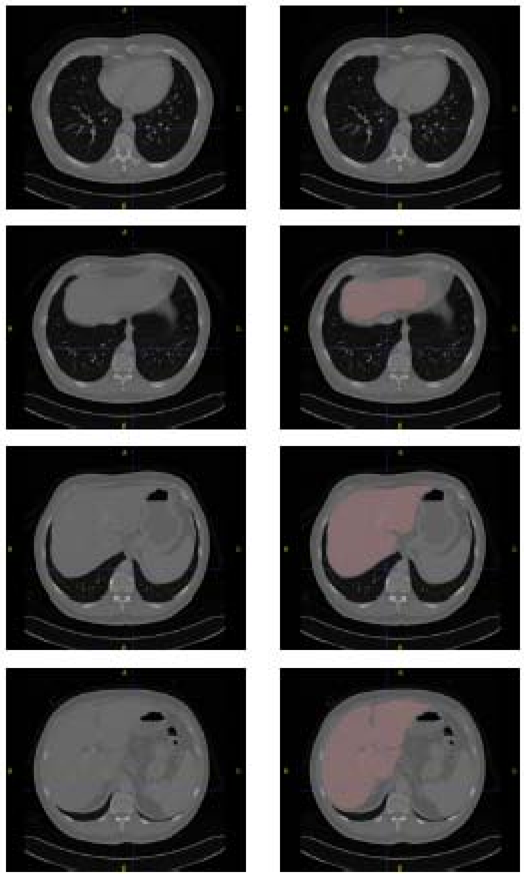

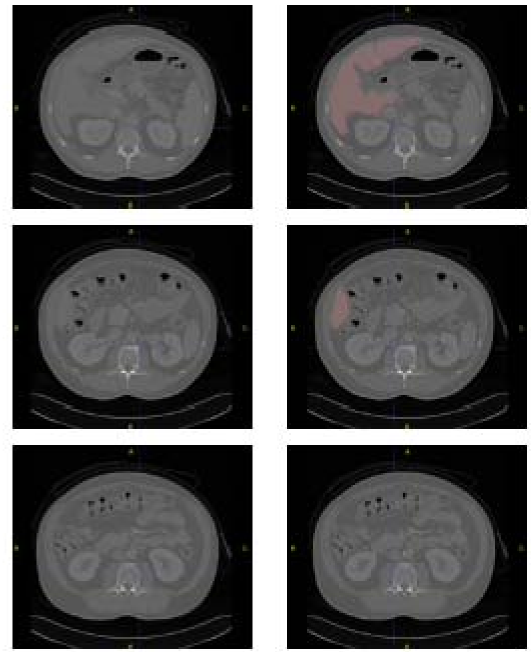

Why Difficult? (Shape Variations)

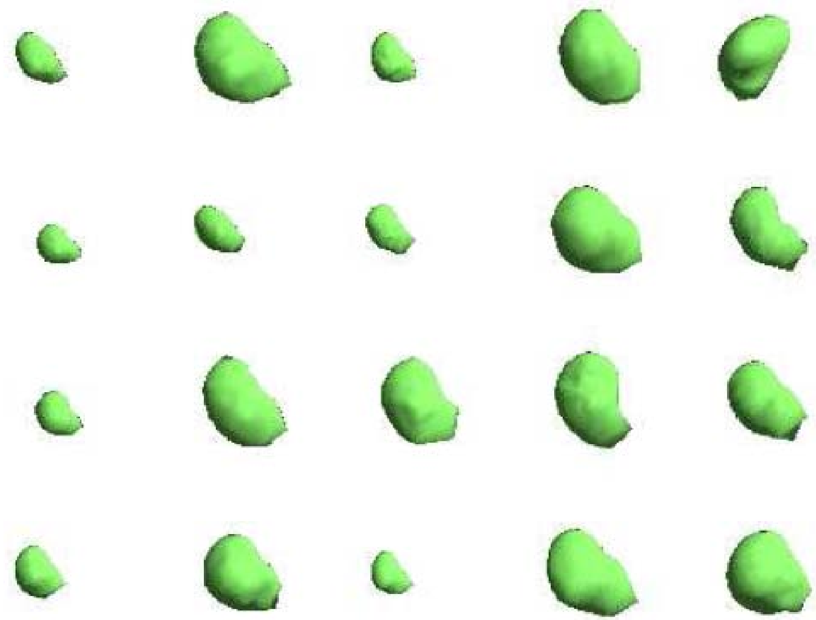

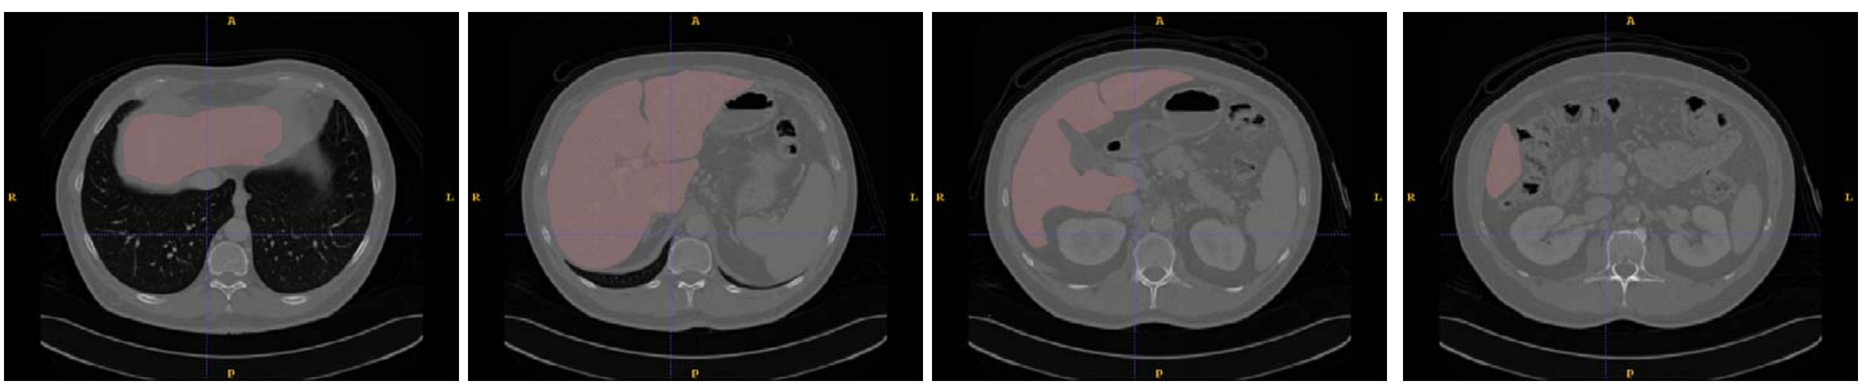

Why Difficult? (Similar Appearances)

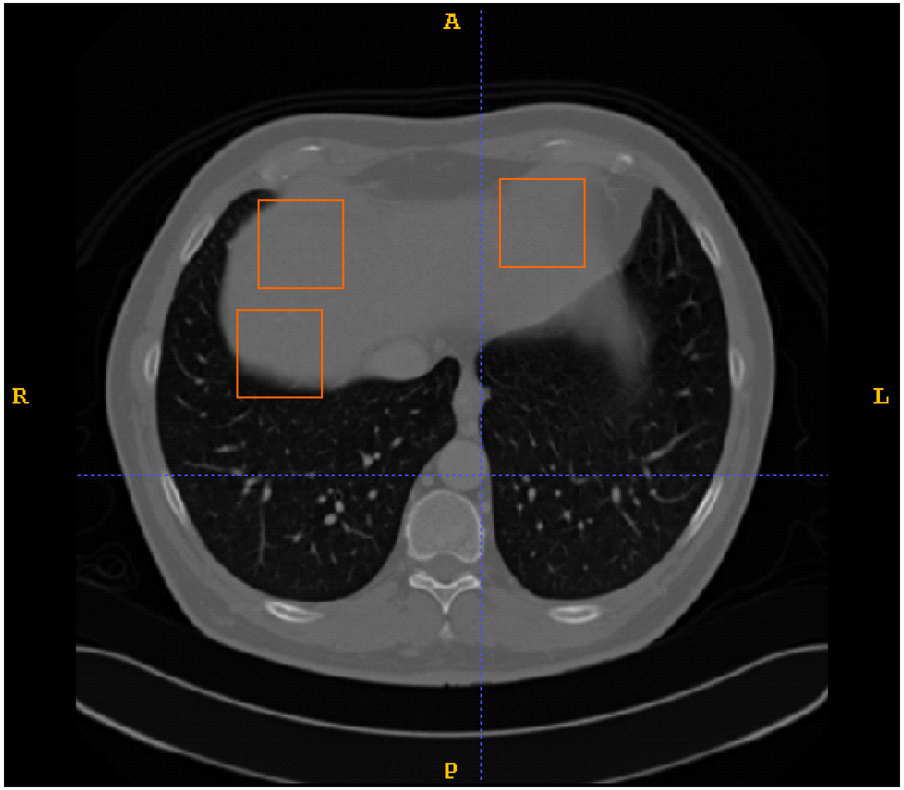

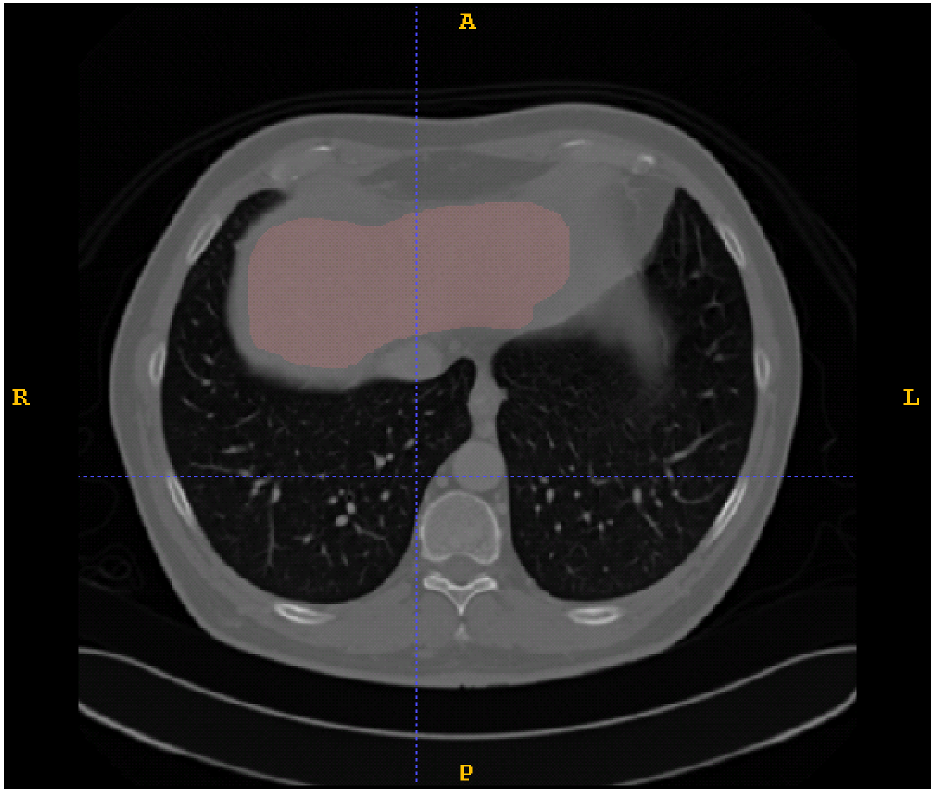

Why Difficult? (Appearance Changes)

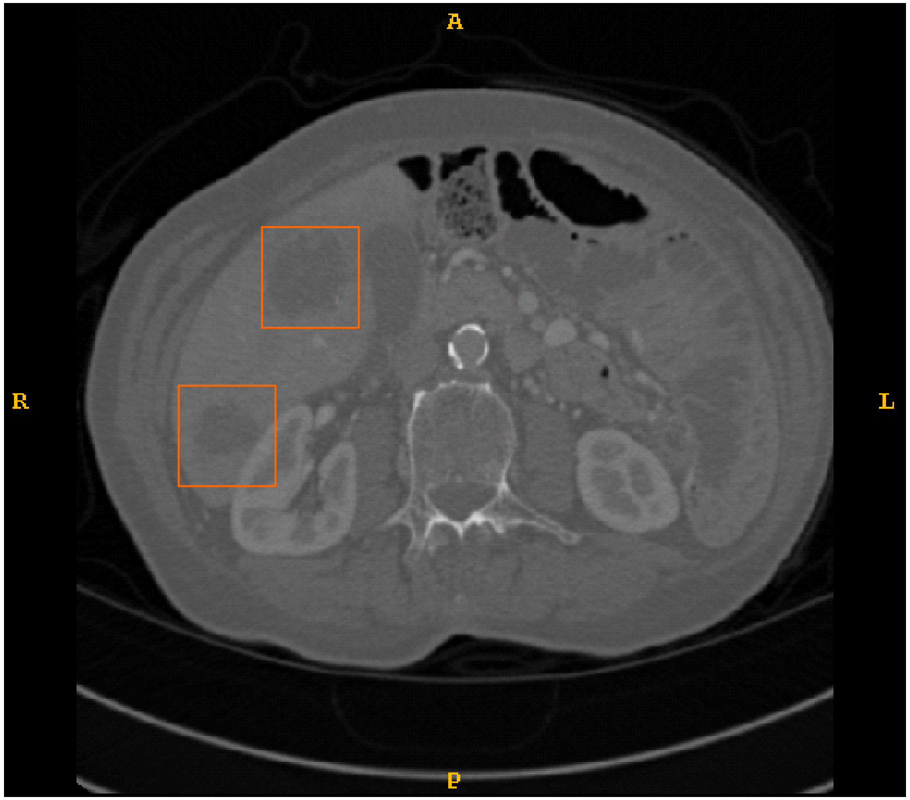

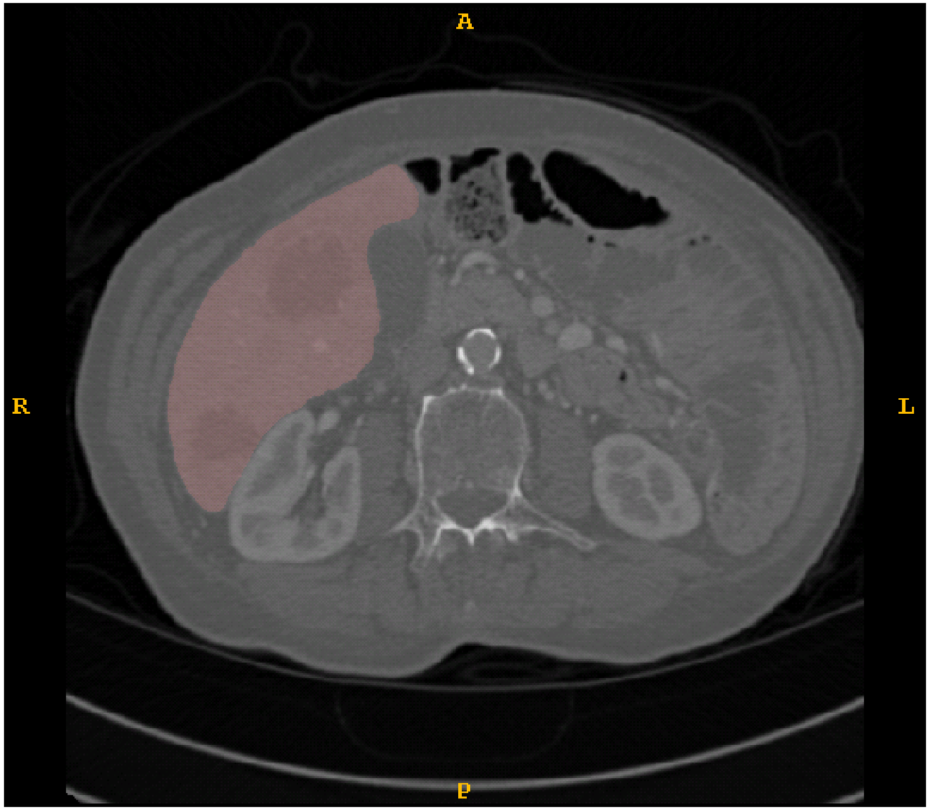

Why Difficult? (Position Changes)

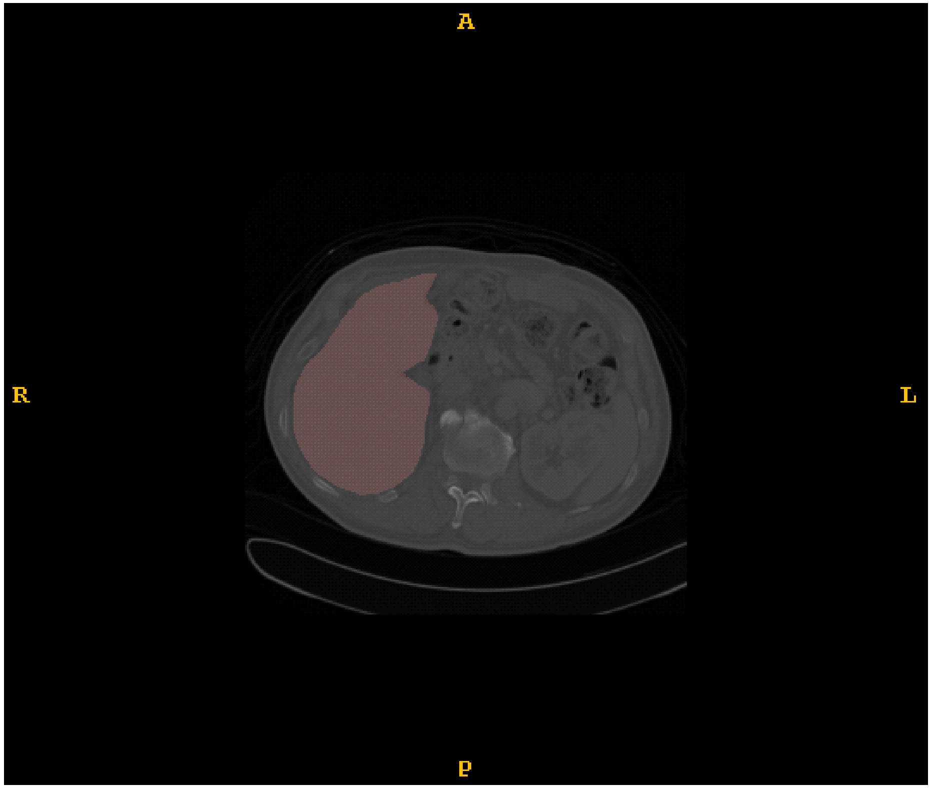

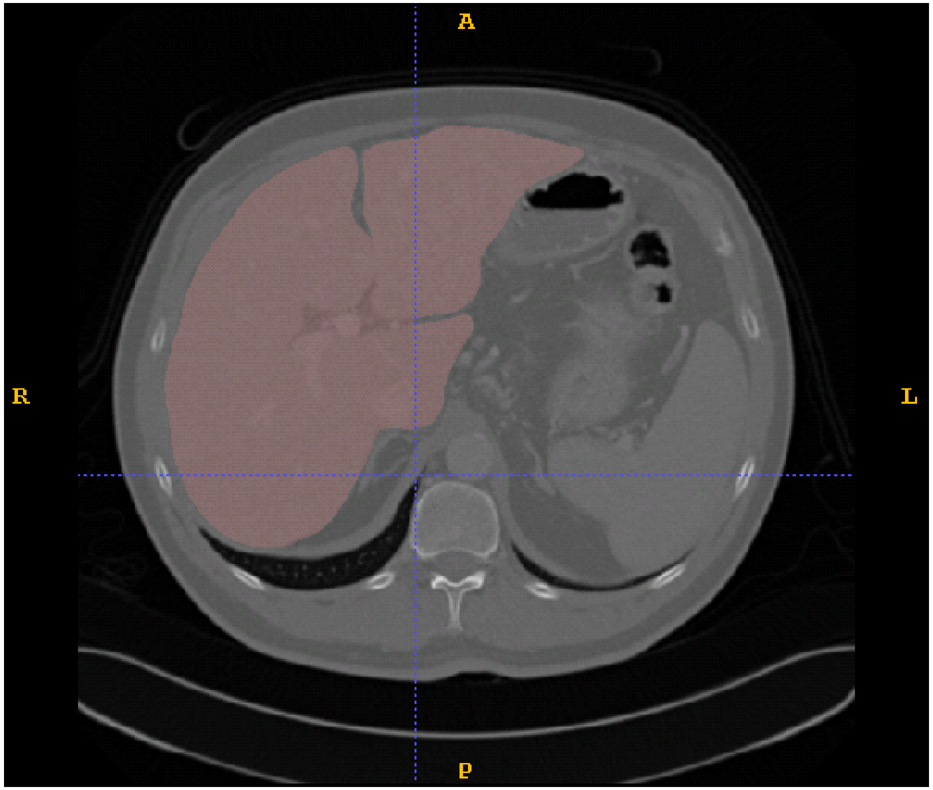

Active Shape Model Based Framework

Training phase

Testing phase

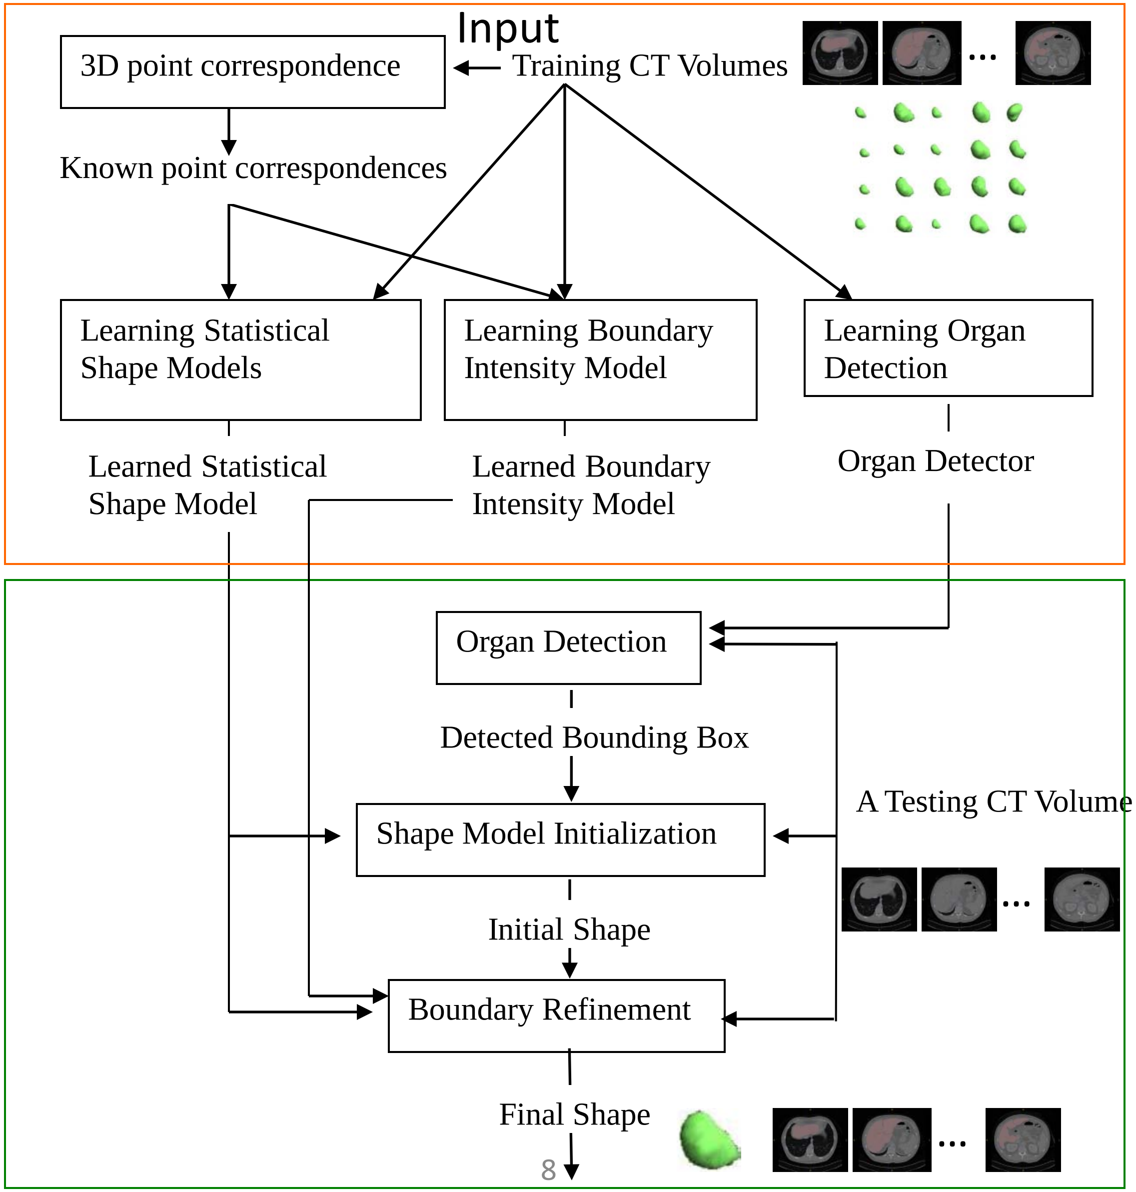

Active Shape Models: Training

- Shapes are modeled in a training phase using a set of CT volumes whose ground truth segmentations are given.

- There are 4 steps to the training phase.

1. Find 3D point correspondences on training meshes.
2. Learn a statistical 3D shape model of the shapes.
3. Learn a boundary intensity model for each vertex.
4. Learn an organ detector that finds bounding boxes.

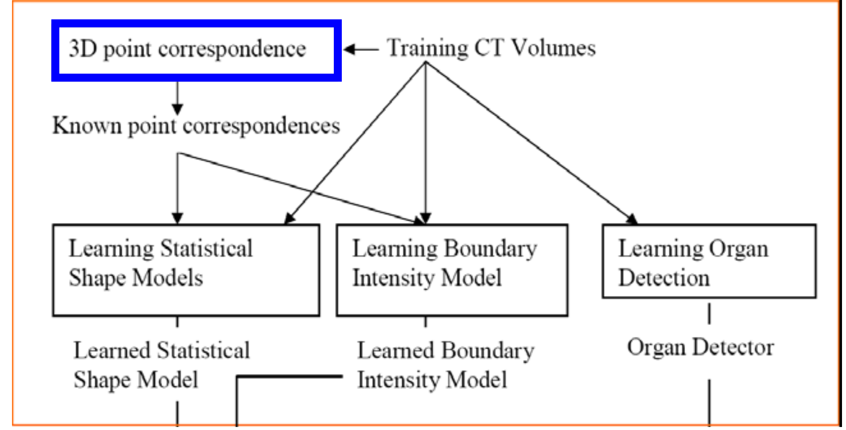

Input

Output

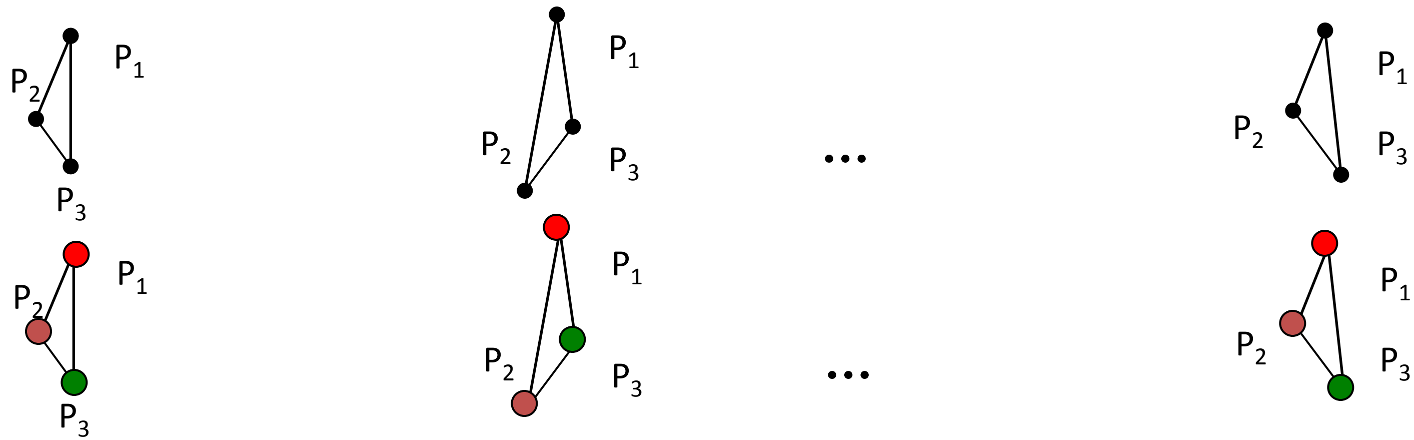

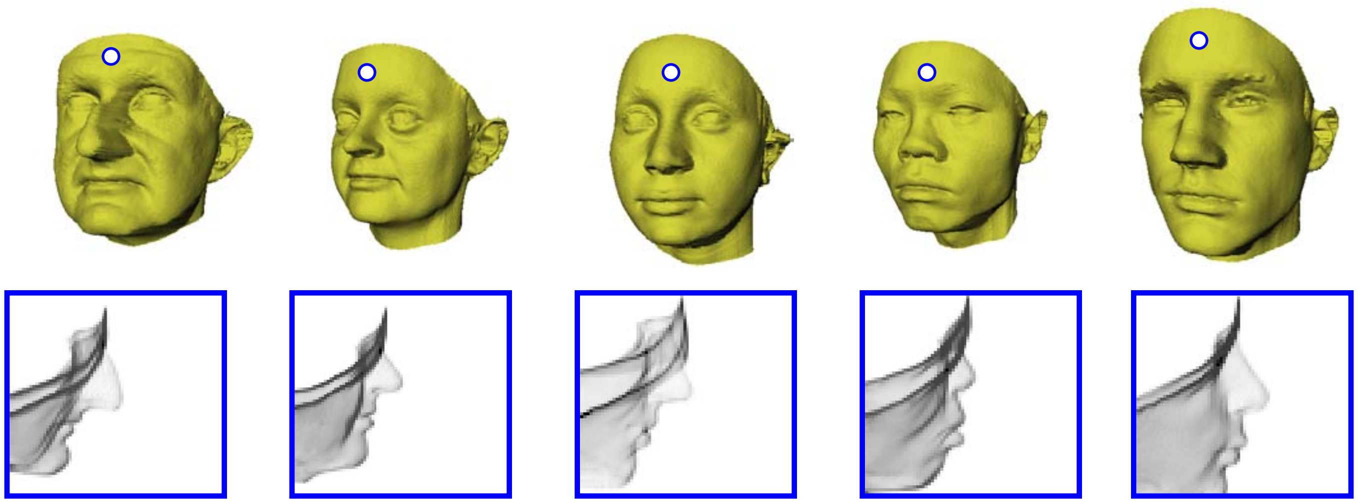

Corresponding points on head meshes plus their numeric (spin image) signatures from the work of Salvador Ruiz Correa.

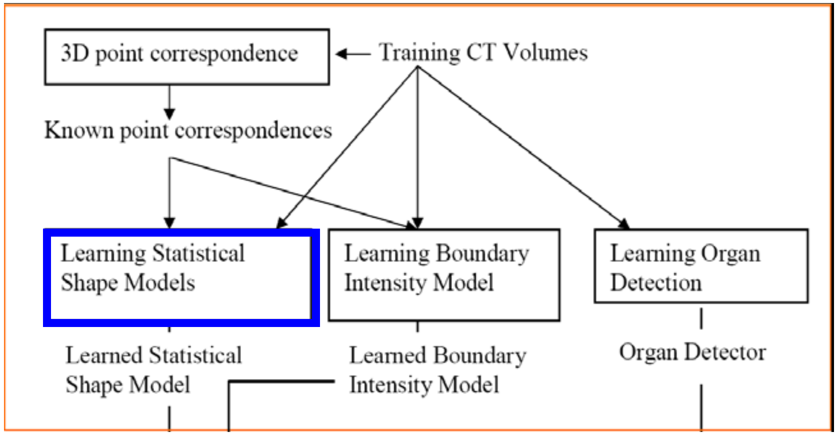

Input

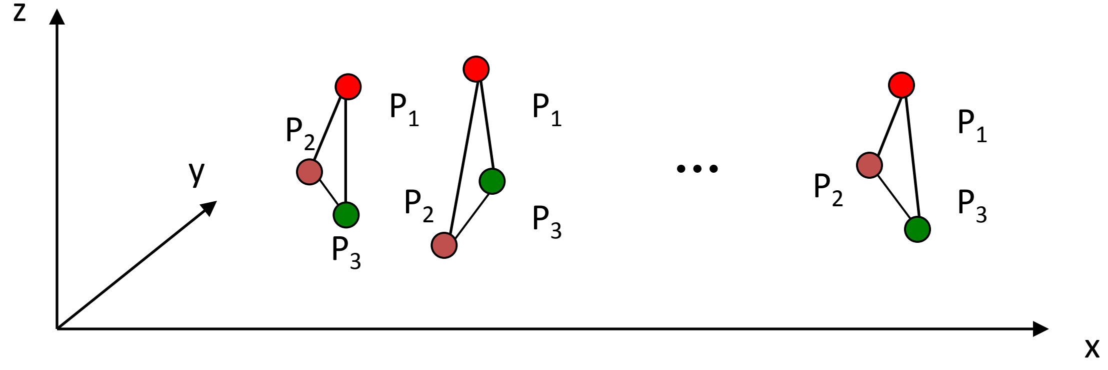

Output

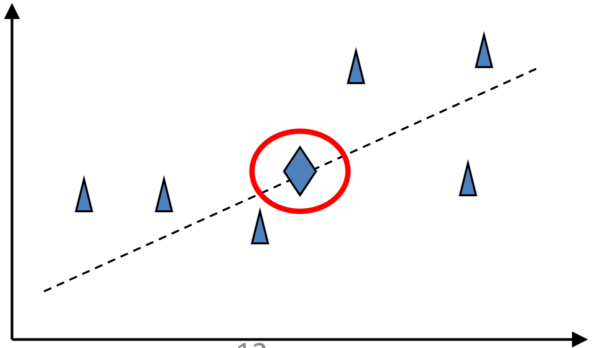

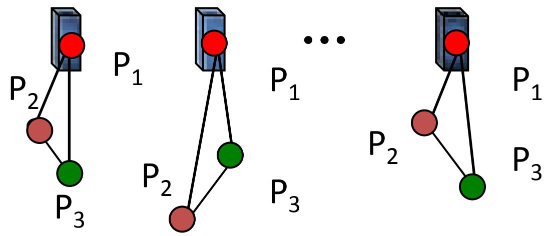

Intensity profiles

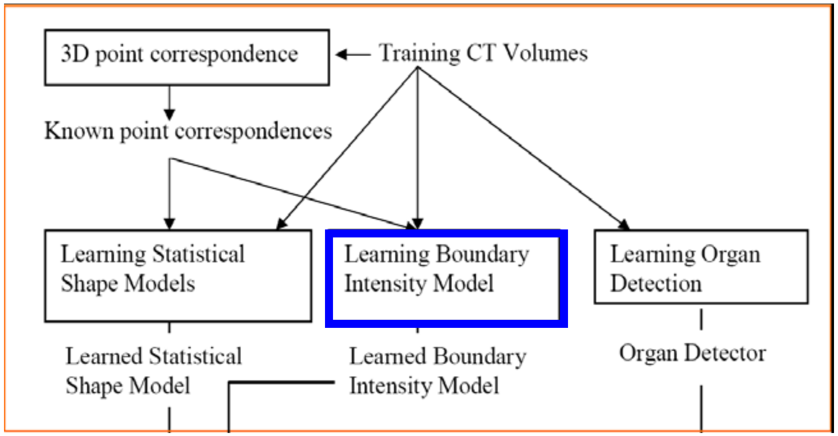

Input ●

Output ●

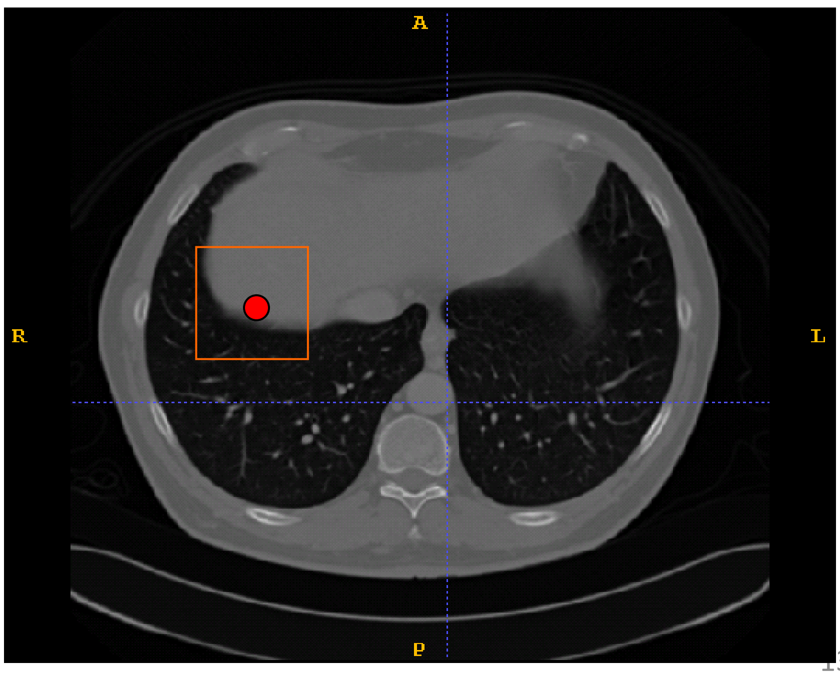

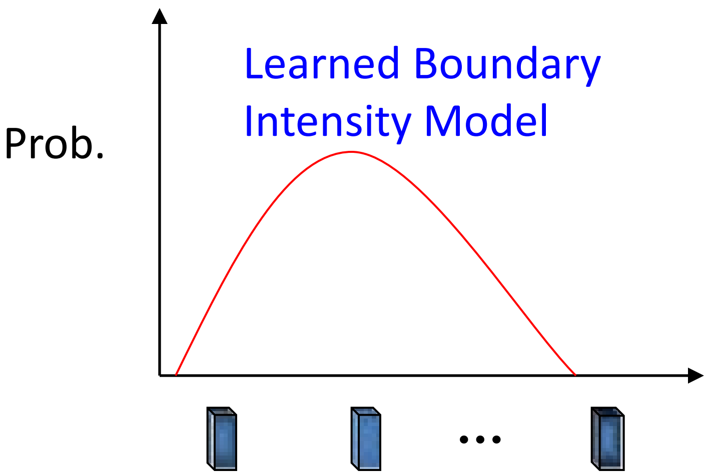

Intensity profiles

- Given a bounding box and the CT slices inside it, a classifier learns to decide if *everything* inside the box is liver or not.

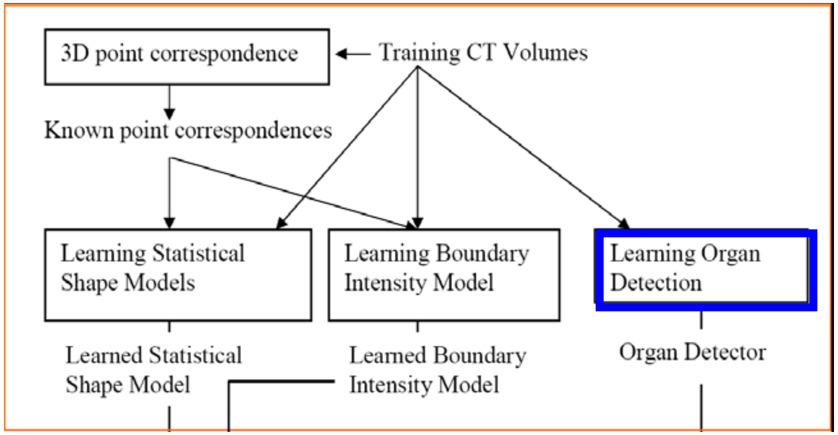

Input

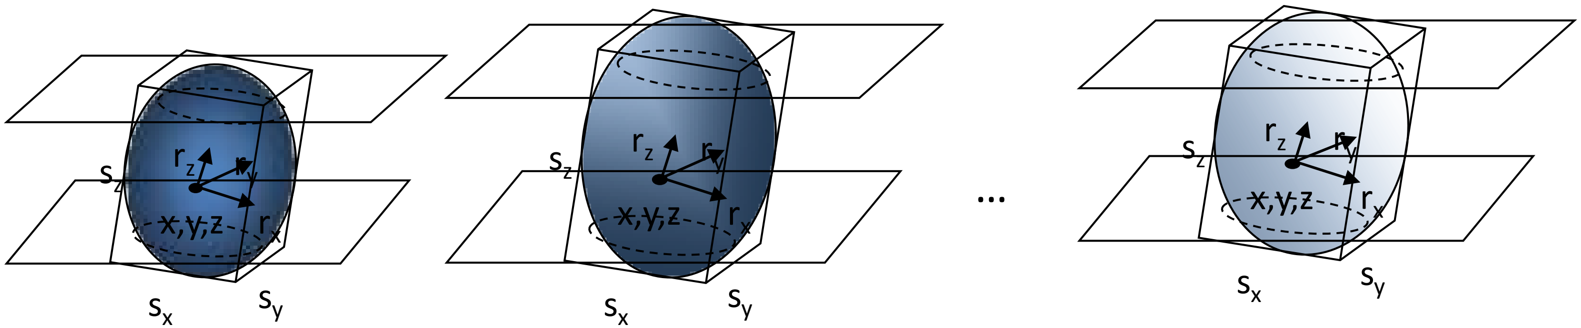

Output

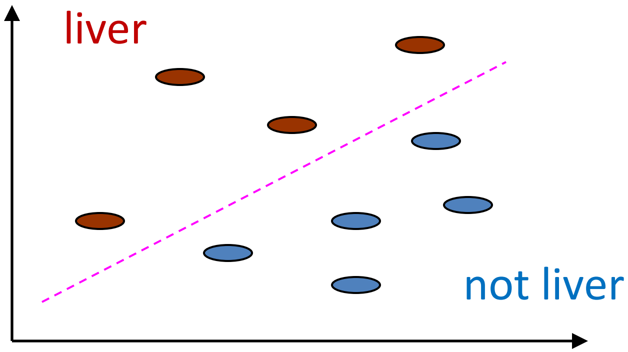

Active Shape Models: Testing

- There are 3 steps to the testing phase

  1. organ detection: use the learned organ detector to detect the organ in the testing volume and return a bounding box
  2. shape model initialization: initialize the learned statistical model based on the detected bounding box
  3. boundary refinement: use the learned boundary intensity model to estimate the refinement to the model for this shape

Organ Detector

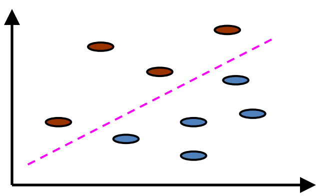

Learned Statistical

Shape Model

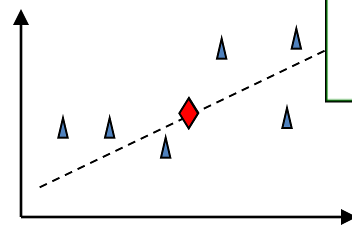

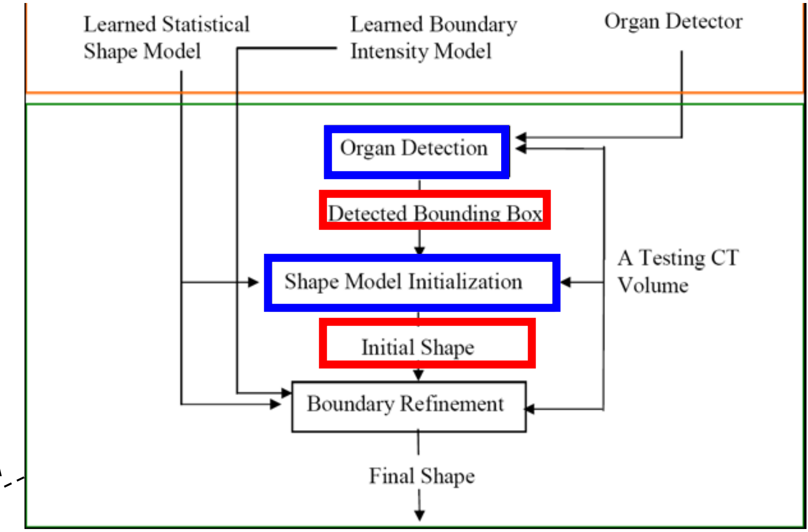

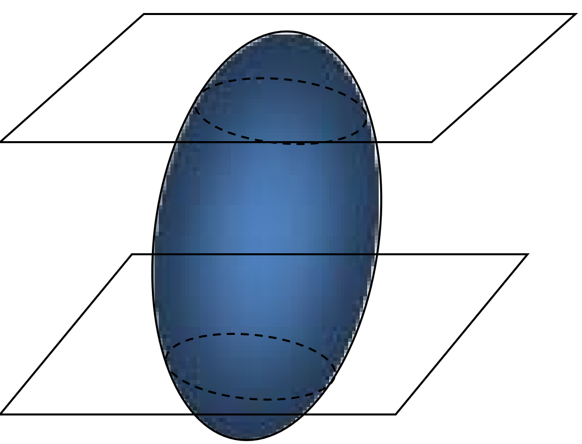

Test Organ

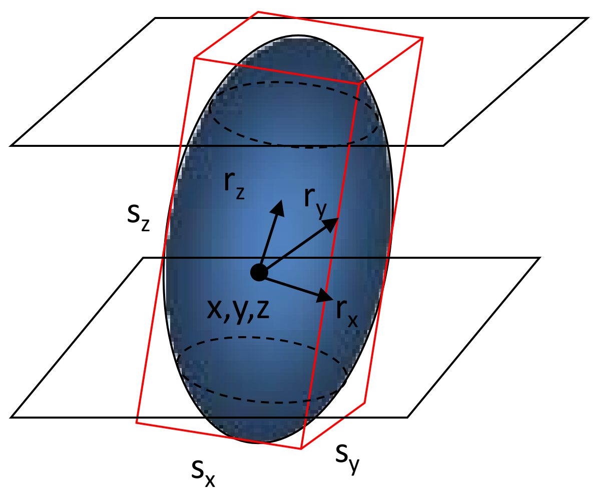

Find Bounding Box

Initialize Shape Model

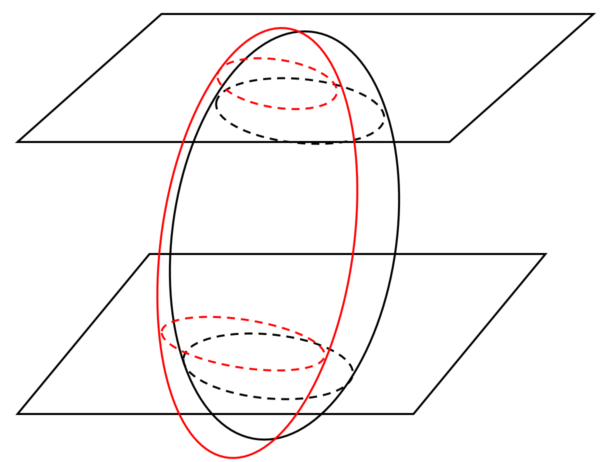

Boundary Refinement

Methods for Point Correspondences

1. Principal Component Analysis (PCA)

PCA takes in the points of each shape in the training set.
It produces a set of basis vectors (the components).

Each shape can then be represented as a linear combination of these components.

$$\tilde{x} = \overline{x} + \sum_{k=1}^{K} c_k b_k \text{ where } \overline{x} \text{ is the mean shape}$$

The optimal K projection axes b_k, k = 1 to K are the eigenvectors of the covariance matrix of the training set of points corresponding to the K largest eigenvalues.

Intuitive Meaning of Principal Components

eigenvector corresponding to highest eigenvalue

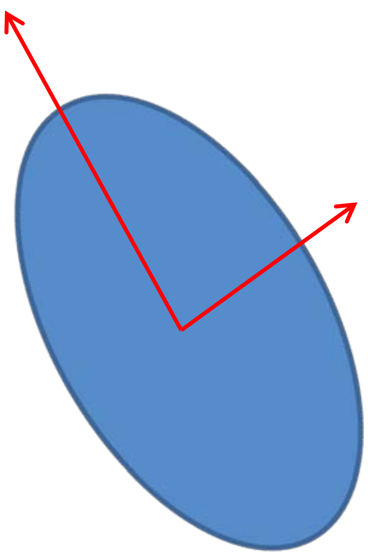

eigenvector corresponding to second eigenvalue

Eigenimages for Face Recognition

training images

mean image

3 eigen-images

linear approxi-mations

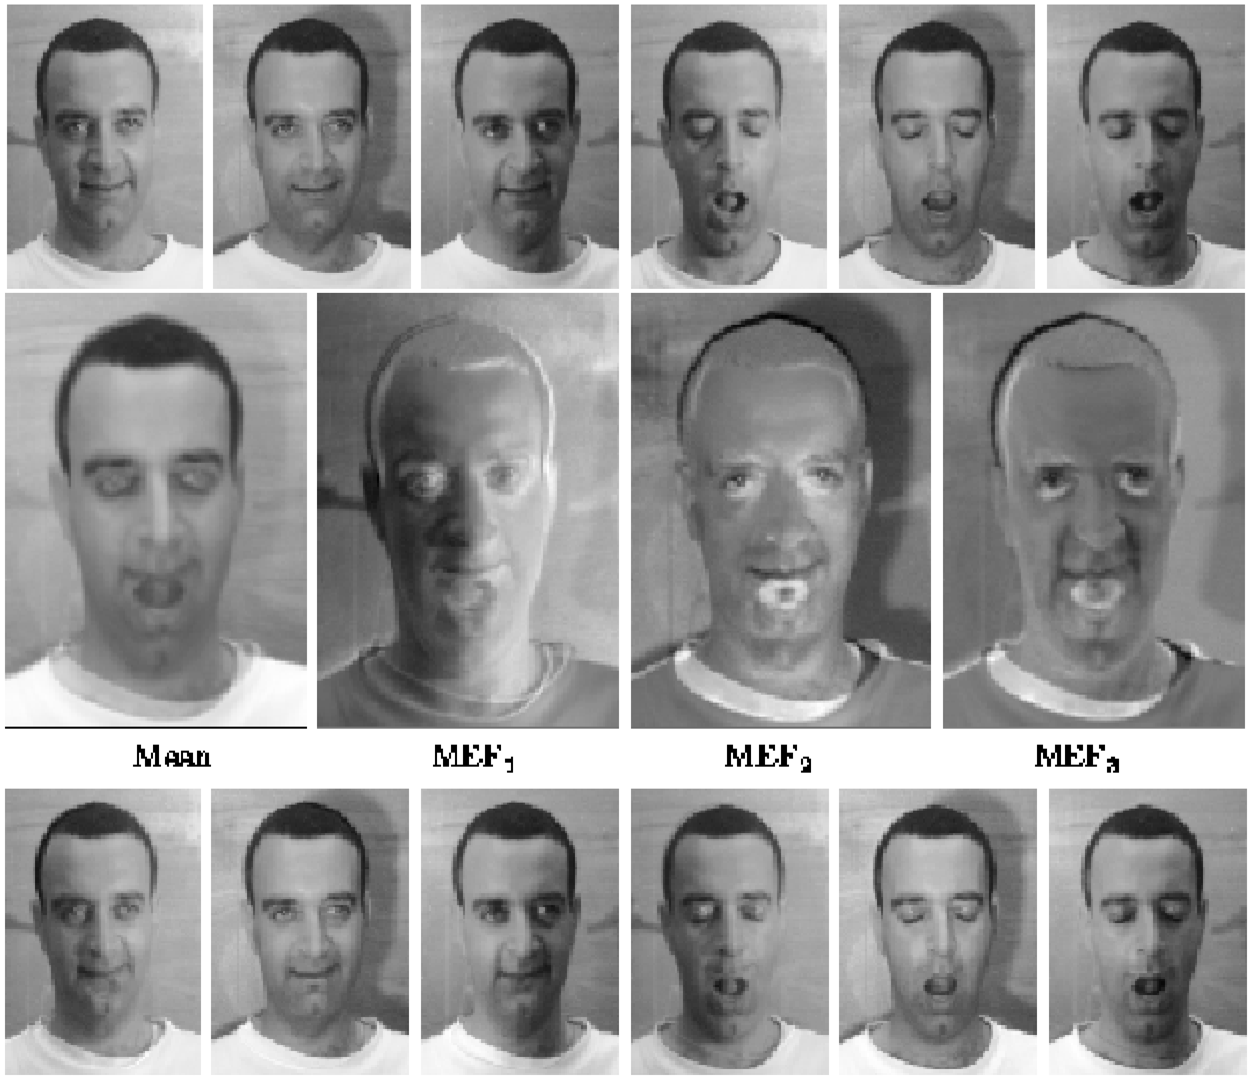

3D Point Correspondence (MDL)

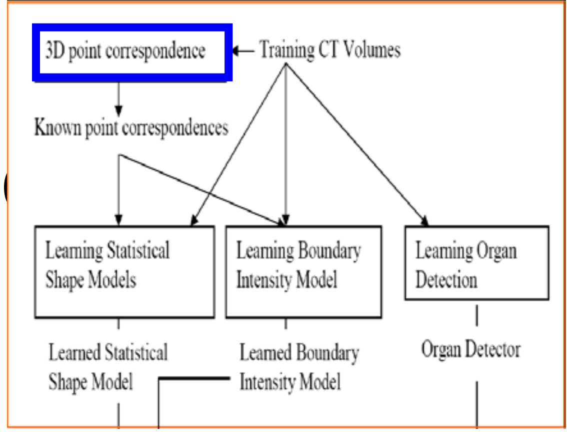

- Goal: Find 3D Point Correspondence
- Idea: Minimize MDL-based objective function
  - Evaluate the quality of the correspondence

$$F = \sum_{k=1}^{N} L_k \text{ with } L_k = \begin{cases} 1 + \log(\lambda_k / \lambda_{cut}), & \text{if } \lambda_k \geq \lambda_{cut} \\ \lambda_k / \lambda_{cut}, & \text{otherwise} \end{cases}$$

- The λ_ks are the eigenvalues from PCA.
- How: Gradient descent
  - Manipulate correspondences by parameterization and re-parameterization.

Davies et al. [IEEE TMI'02]

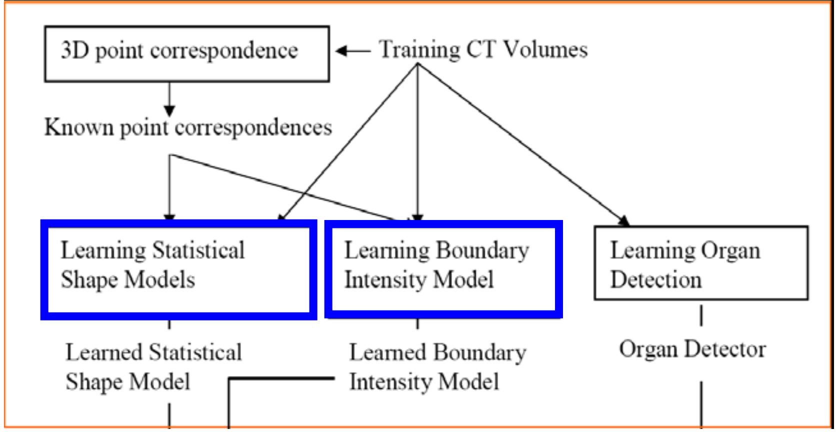

- Statistical Shape Models
  - Principal Component Analysis (PCA)
  - Kernel PCA

- Boundary Intensity Models
  - Gaussian distribution
  - AdaBoosted histogram classifiers
  - Heuristics

Cootes et al. [IEEE PAMI 01], Twining et al. [BMVC'01]

Cootes et al. [IEEE PAMI 01], Li [ICCV'05], Kainmuller et al. [MICCAI'07]

Organ Dete
(MSL)

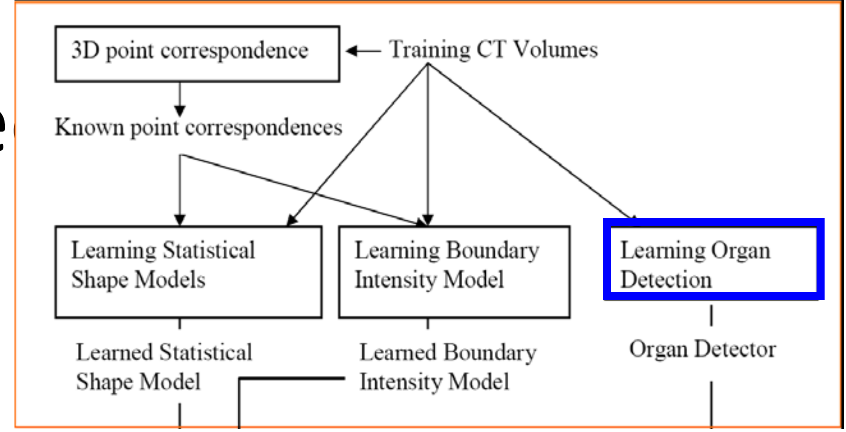

- Goal: Find the bounding box
  - The parameter space is 9D.
  - 3D positions, 3D scales and 3D orientations.
- Idea
  - Uniform and exhaustive search is unnecessary
- How: decompose the problem into three steps
  - position estimation, position-scale estimation and finally position-scale-orientation estimation.

Zheng et al. [ICCV'07]

Two ASM-based Systems

Kainmuller et al. [MICCAI'07]	Ling et al. [CVPR'08]
Statistical Shape Models	
– PCA – 43 CT volumes	– PCA, hierarchical shape pyramids – 75 volumes
Boundary Intensity Model	
– Heuristics	– A boundary classifier
Liver Detection	
– Lungs detection and DICOM info	– MSL (marginal space learning)
Performance	
– Ranked first in a recent liver segmentation competition. – 10 testing volumes – 1.1mm (the average symmetric surface distance) – 15 minutes.	– 5 fold cross validation – 1.59 mm (the average symmetric surface distance) – 1.38 mm (the median) – 12 seconds.

Experimental Setting

- Datasets:
  - 4 types of organs (livers, left kidneys, right kidneys, spleens)
  - 15-20 subjects
- Leave-one-out cross validation
- Measure the reconstruction error
- Metrics: Euclidean and Hausdorff distance

MDL-2D

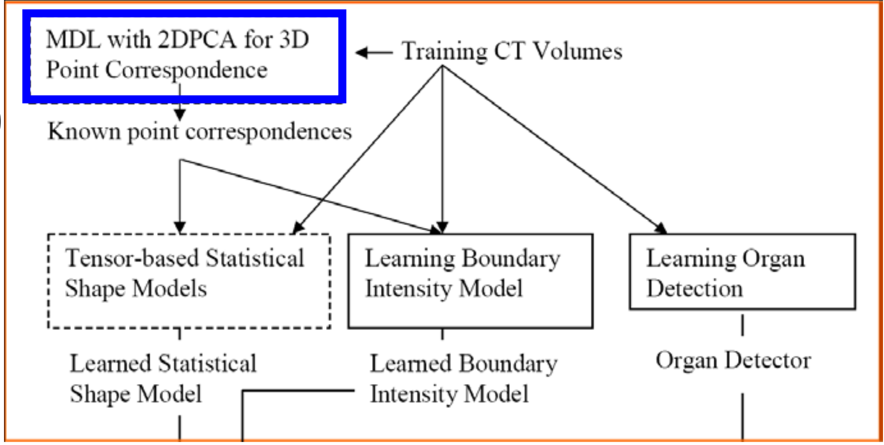

- MDL-based objective function

$$F = \sum_{k=1}^{N} L_k \text{ with } L_k = \begin{cases} 1 + \log(\lambda_k/\lambda_{cut}), & \text{if } \lambda_k \geq \lambda_{cut} \\ \lambda_k/\lambda_{cut}, & \text{otherwise} \end{cases}$$

- Idea: Generalize the objective function to 2DPCA space
  - Replace eigenvalues from PCA with from 2DPCA
  - How: Gradient descent
- Comparisons: original MDL vs. MDL-2DPCA

* Chen and Shapiro [EMBC'09]

Results (3D Point Correspondences)

Reconstruction error

Livers

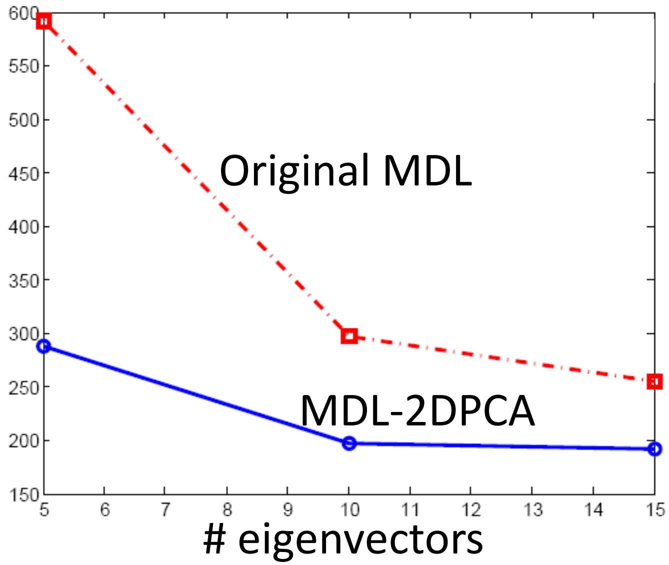

Left Kidneys

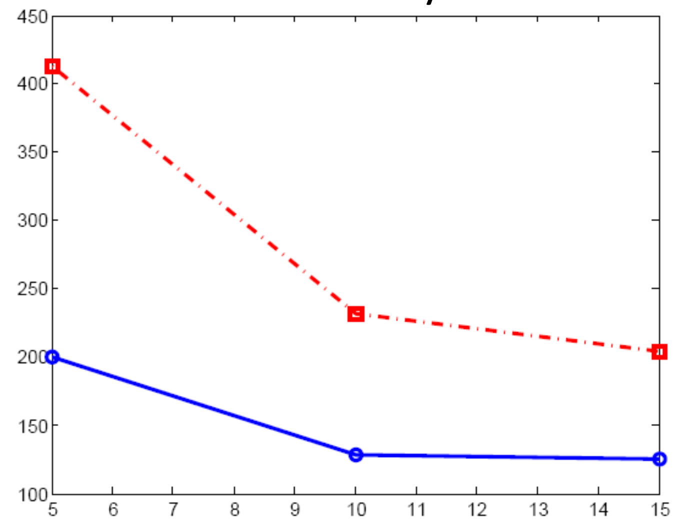

Right Kidneys

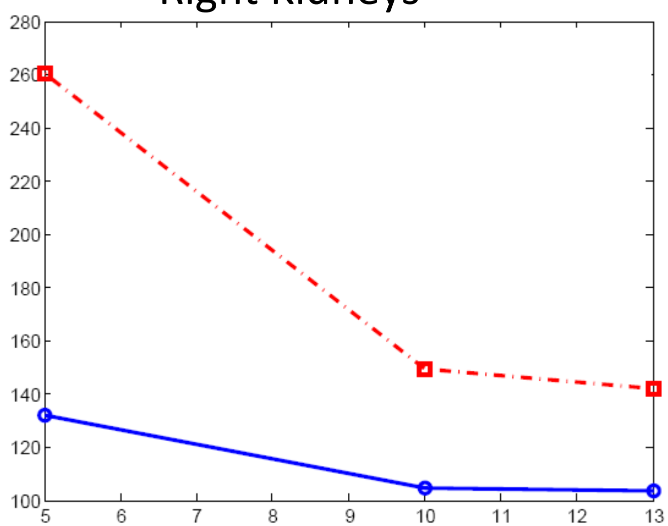

Spleens

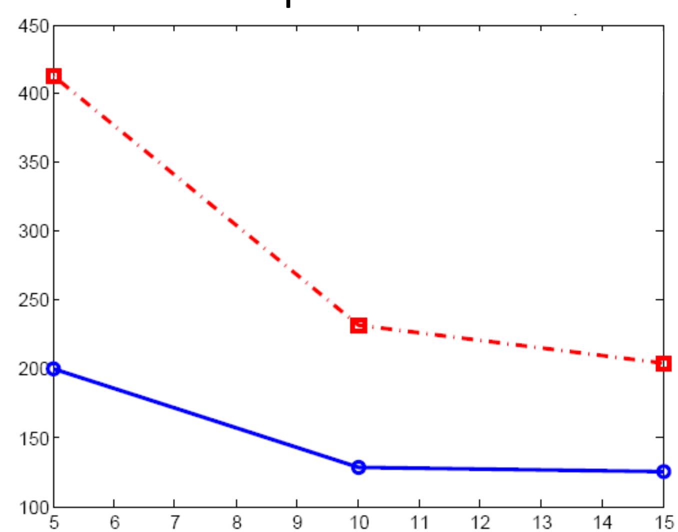

Tensor-base

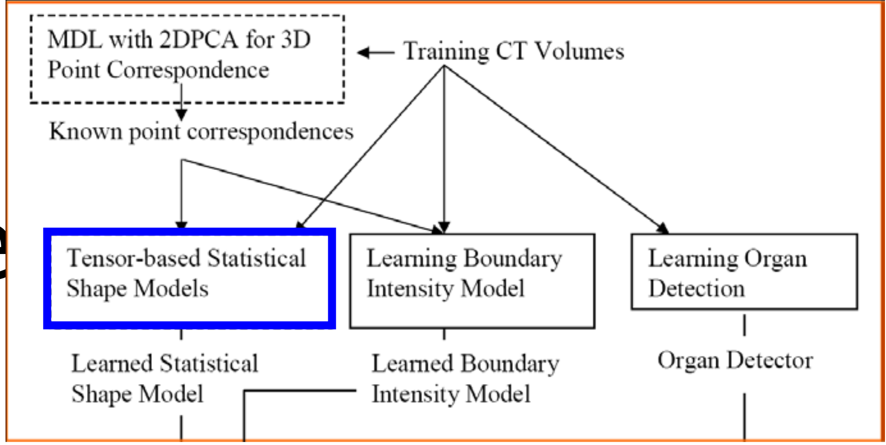

- Idea: Tensor-based dimension reduction methods
  - 2DPCA
  - Parafac model
  - Tucker decomposition
- Comparisons: PCA vs. Tensor-based dimension reduction

* Chen and Shapiro [to appear in EMBC'09]

Results (Statistical Shape Models)

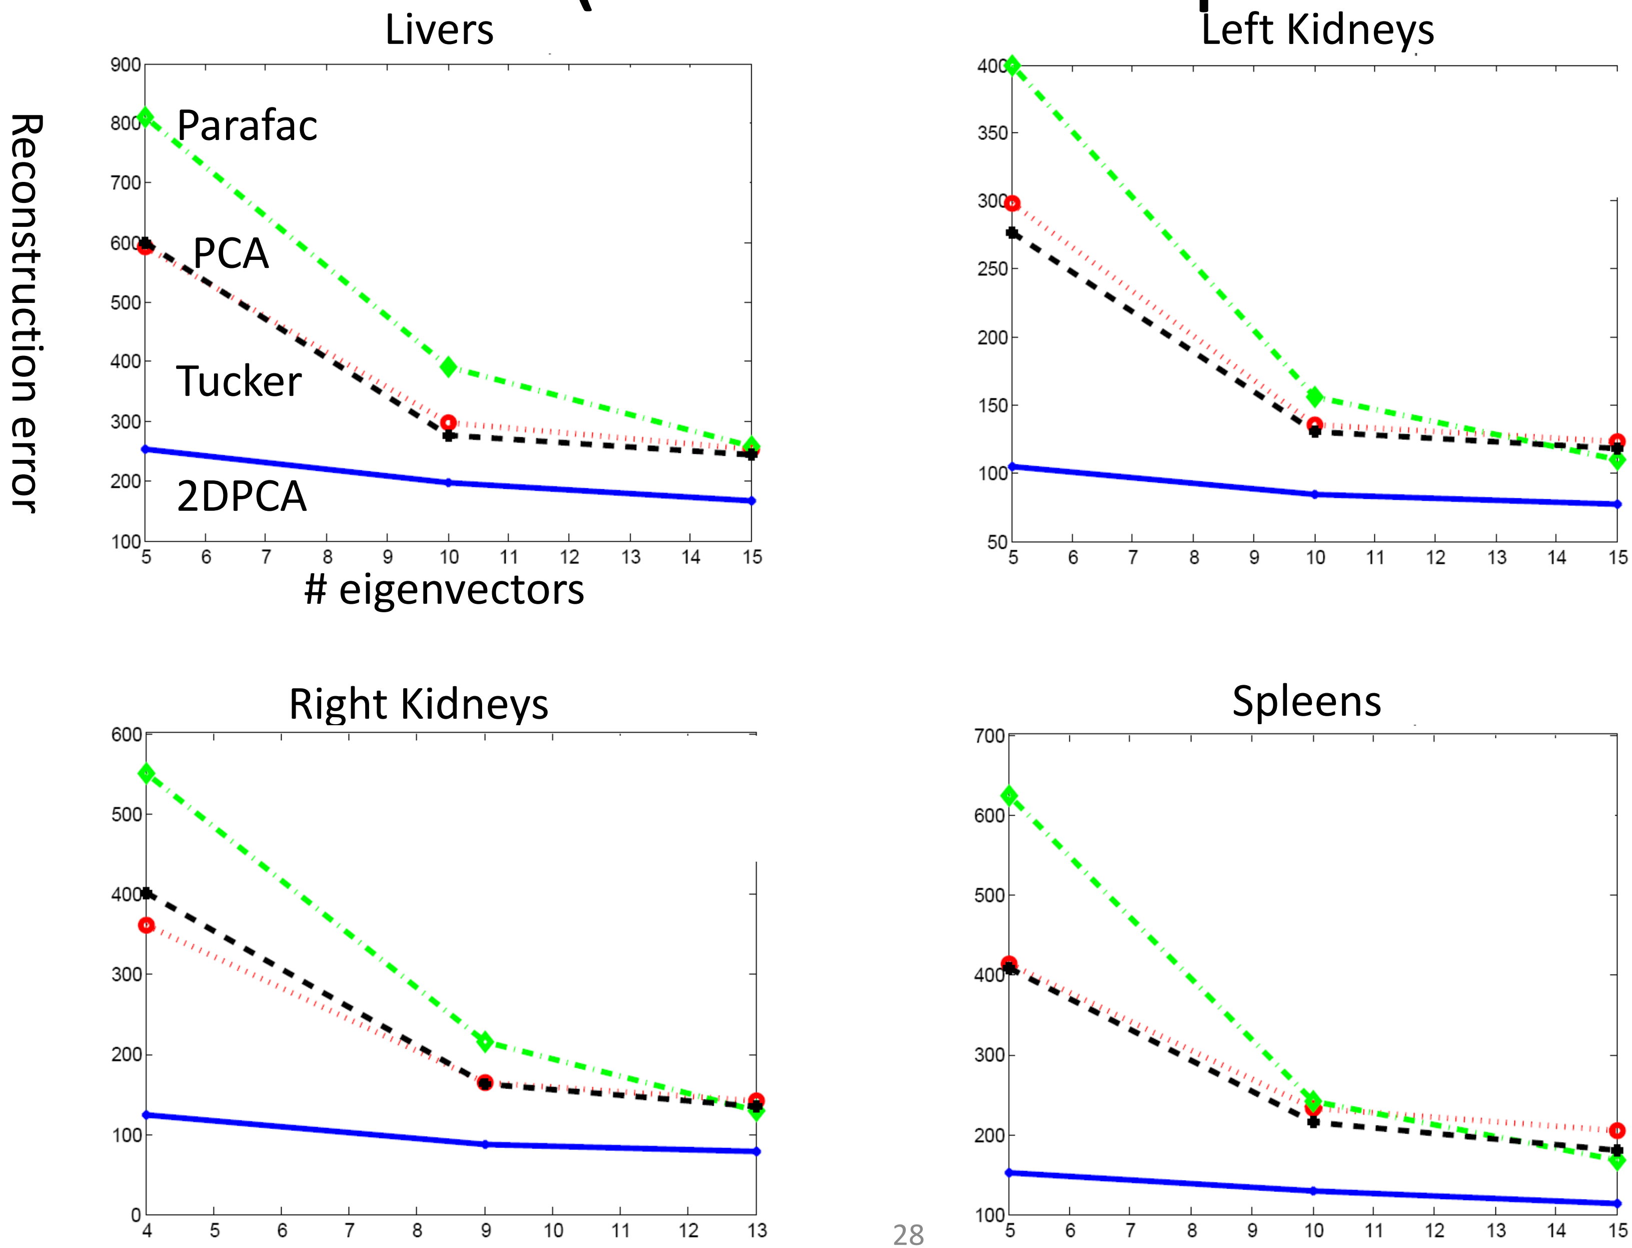

Organ Det (Boosting Ap

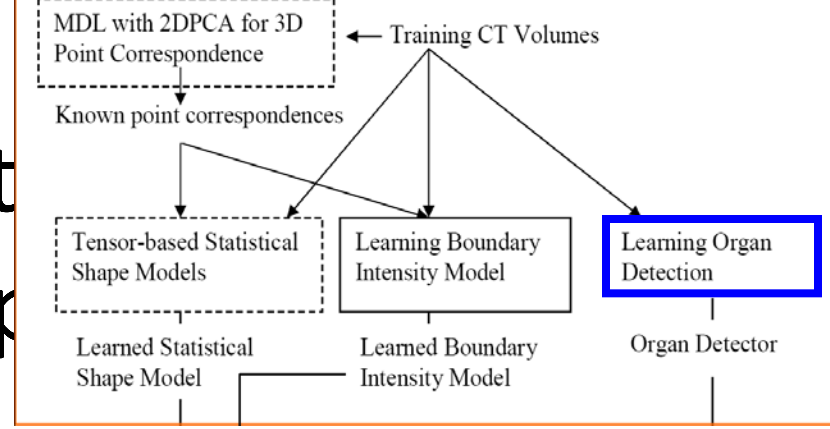

- Idea: Classify whether an image block contains an organ of interest
- How:
  - Partition slices into non-overlapping 32x32 blocks
  - Global features: gray-tone histogram of the image slice and its slice index
  - Local features: the position of a block, the mean and variance of its intensity values, and its intensity histogram.
  - 20,000 SVM linear classifiers + Adaboosting
- Comparisons: Manual vs. Adaboosting

Results (Organ Detection)

Livers (Training)

	Positive (predicted)	Negative (predicted)
Positive (actual)	96.23%	3.77%
Negative (actual)	4.57%	95.43%

Livers (Testing)

	Positive (predicted)	Negative (predicted)
Positive (actual)	91.23%	8.77%
Negative (actual)	6.57%	93.43%

True Positive : Green, False Positive: Blue

False Negative: Red, True Negative: Cyan

Sub. 1

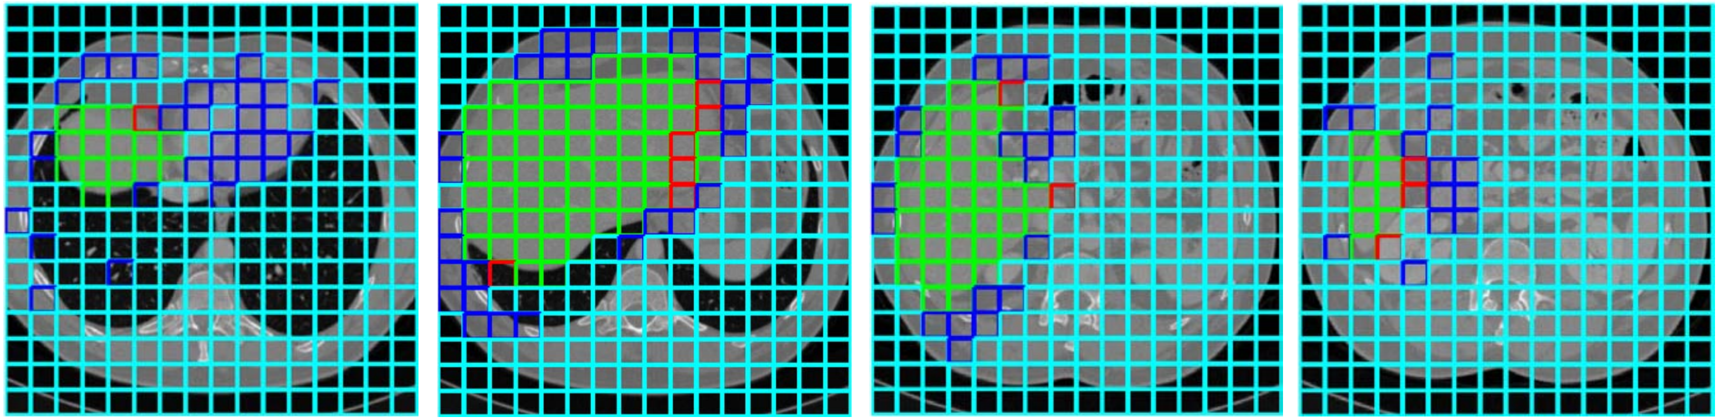

Sub. 2

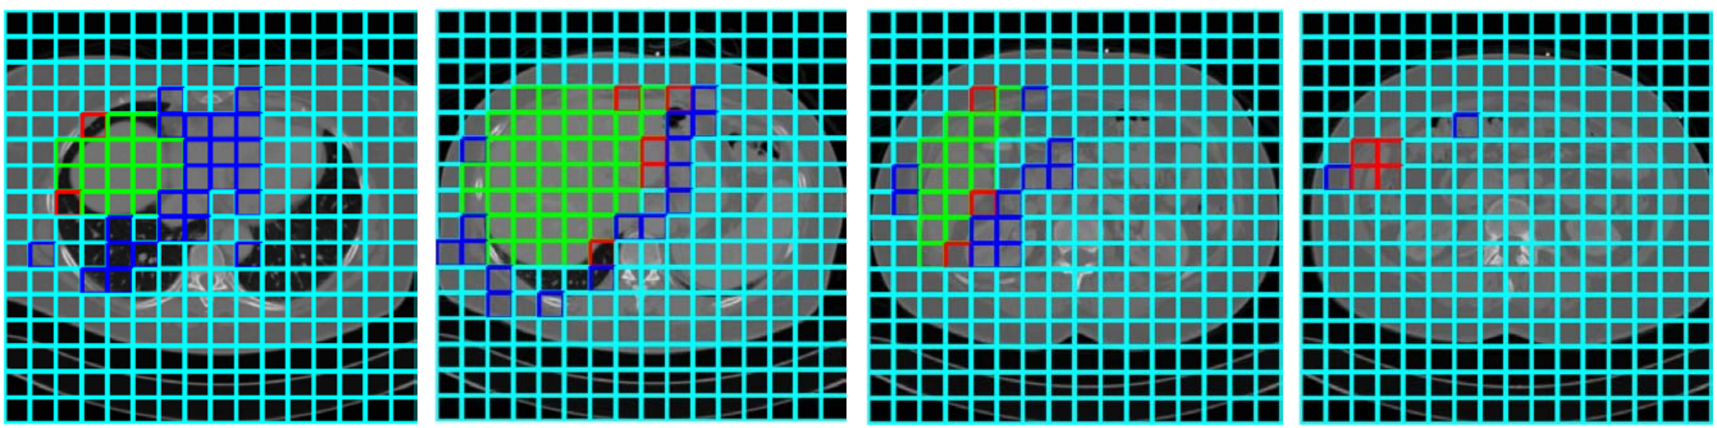

Graph Cuts Bas
Boundary Refine

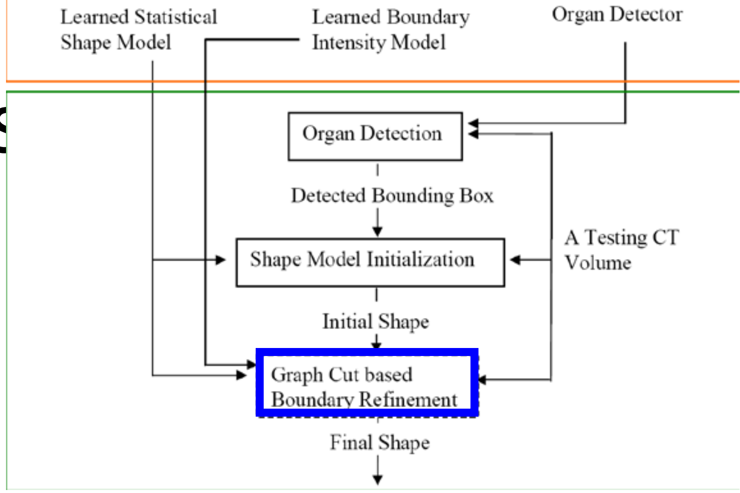

- Idea: Adding hard constraints to min s-t cuts
- Min s-t cuts with side constraints
  - NP-hard in general cases
  - Approximation algorithm: standard rounding algorithm
- Comparisons: with constraints vs. without constraints

Chen and Shapiro [ICPR'08]

Results (Boundary Refinement)