A Latent Source Model for Patch-Based Image Segmentation

George H. Chen   Devavrat Shah   Polina Golland

Motivation

Patch-based segmentation methods popular now
→ circumvent nonrigid registration for label fusion
→ leverage fast approximate nearest-neighbor search

Goal: Develop theory to understand why these methods work

Patch-Based Segmentation

Variants of the following nearest-neighbor algorithm:

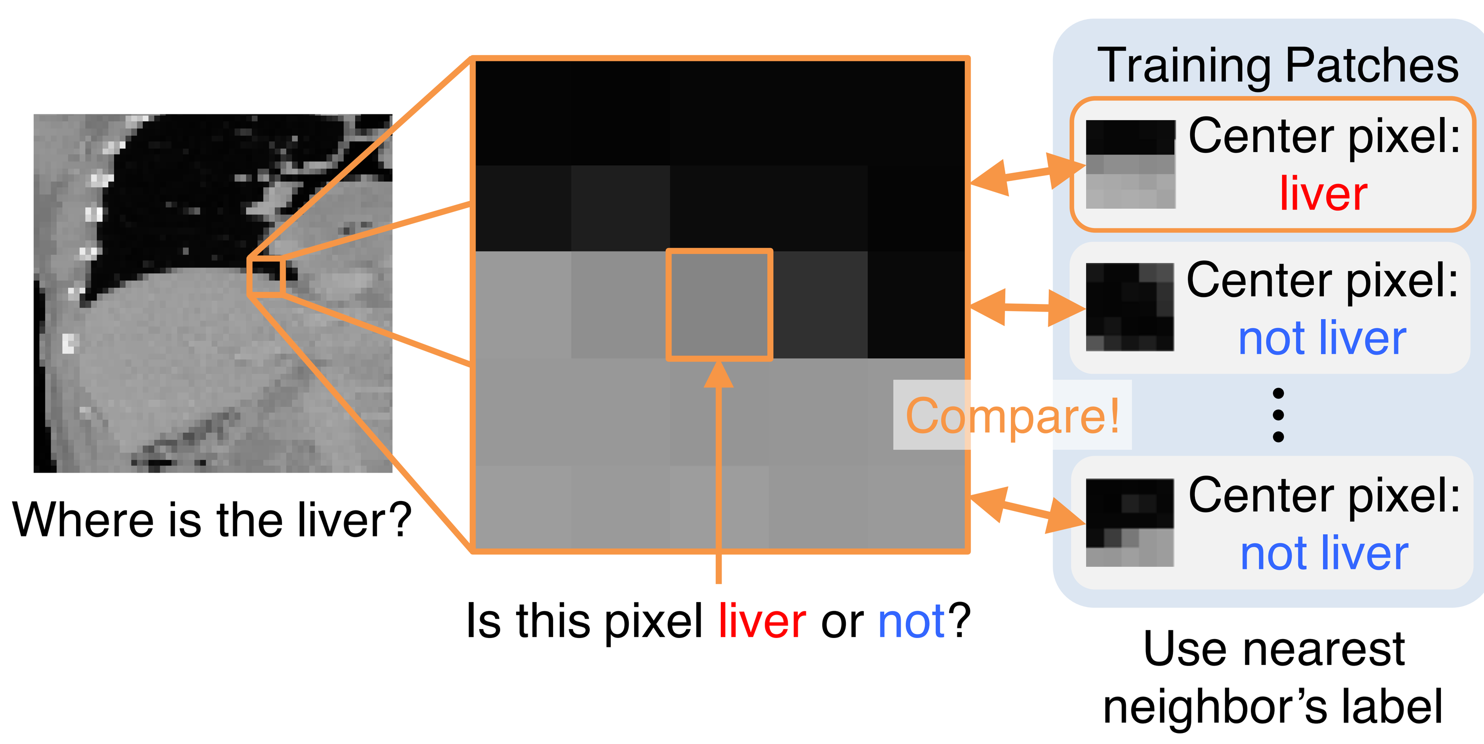

There are many ways to improve this algorithm (weighted majority voting, feature descriptors, smoothing, etc.)

Contributions

New theoretical guarantee
→ characterizes pixel mislabeling rate of nearest-neighbor, weighted majority voting patch-based segmentation

New probabilistic model for patch-based segmentation
→ leads to new iterative algorithm with many existing patch-based methods as special cases

Key Ideas

Patches Cluster

In real images, what do plausible small patches look like?

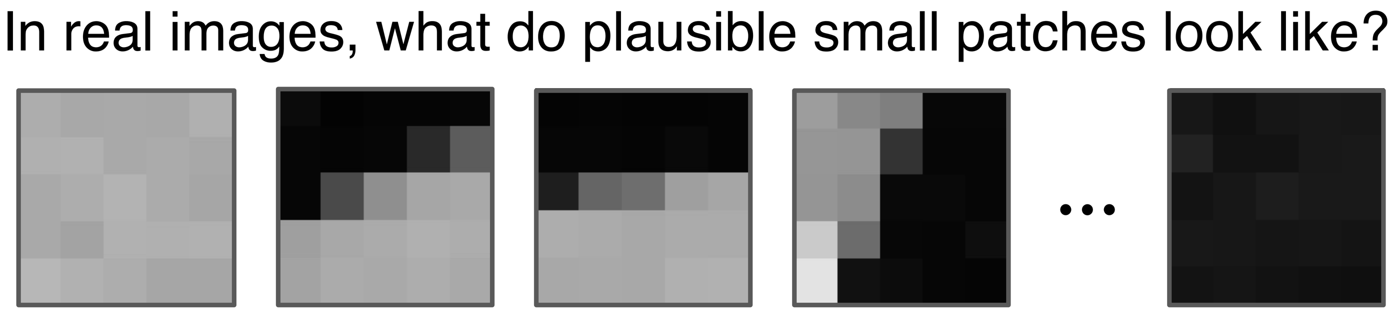

Can cluster these into "canonical" patches!

→ model image patches as noisy versions of few canonical patches ("latent sources" that generate patches)

Nearby Patches Appear Similar

Spatially nearby patches (within subject & across subjects) can be explained by same canonical patch

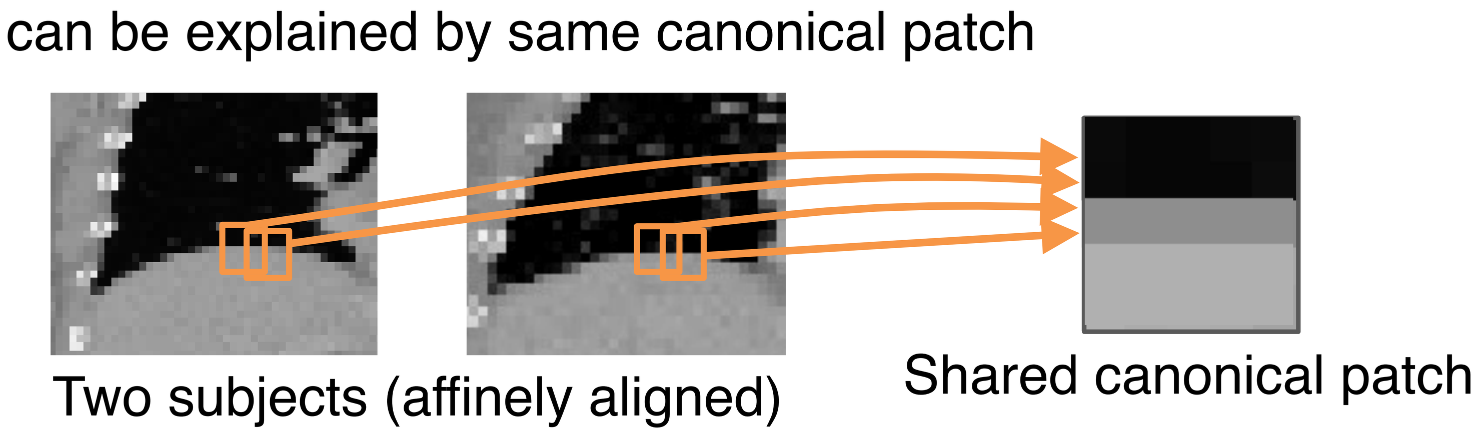

→ model nearby patches to share which canonical patches they are generated from

Point-wise Segmentation

Predict each pixel's label separately using patches

Generative Model for Each Patch

Each pixel has its own canonical patches in two bags

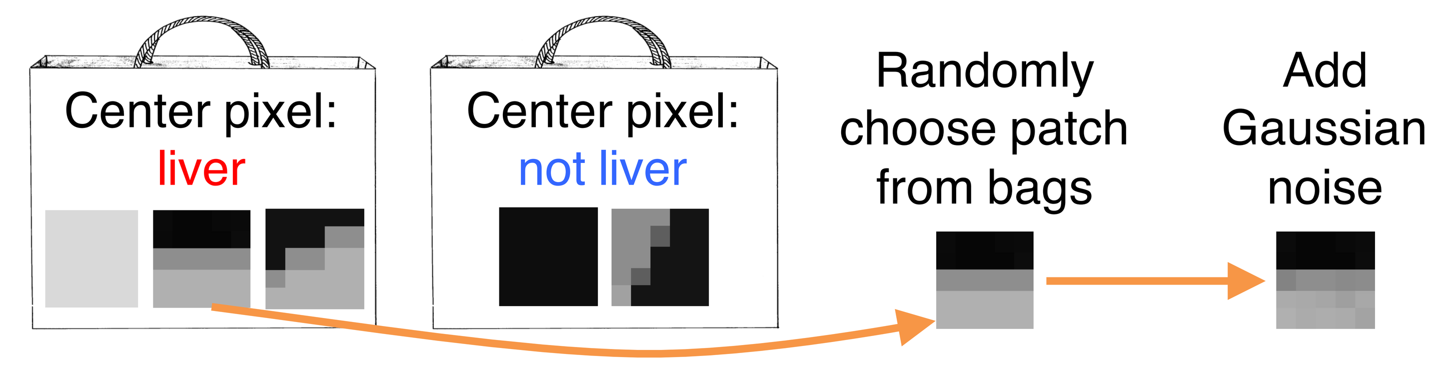

Theoretical Guarantee

Nearest-neighbor/weighted majority voting

Assume:
- Nearby pixels share enough canonical patches
- Canonical patches with opposite labels different enough

Then can make average pixel mislabeling rate (0 to 1) → 0 with # training subjects = $\Theta(k \log k)$

k = max # canonical patches a pixel has

Interpretation
- # training subjects sufficient: enough to see all canonical patches per pixel
- error doesn't → 0: assumptions don't hold across image

Multi-point Segmentation

Predict label *patches* and merge label patch estimates

Probabilistic Model

Two additions to point-wise segmentation model:

More Local Structure
Look at label patch, not just center pixel

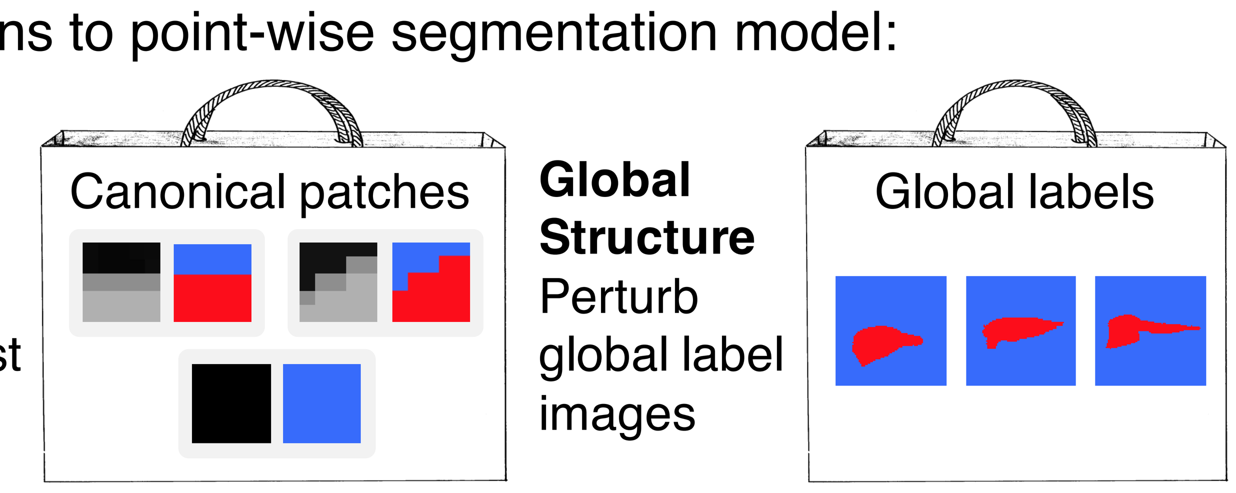

Global Structure
Perturb global label images

Inference: ADMM Algorithm

Iterate between:
- Predict each label patch separately – parallelizable (e.g., weighted majority voting)
- Merge label patches (uniform global prior → average each pixel's label predictions)

New ADMM algorithm

Multi-point/in-painting
(Coupe et al '11, Rousseau et al '11, Wachinger et al '14, Zoran & Weiss '11)

Point-wise
(Bai et al '13, Coupe et al '11, Rousseau et al '11, Wachinger et al '14)

Results