Convex Energy Minimization Over Multi-Region, Probabilistic Segmentation Spaces

Shawn Andrews
(with Chris McIntosh and Ghassan Hamarneh)

September 1^{st}, 2011

Medical Image Segmentation

- Segmentation is a fundamental task in medical image analysis
- Shape of organs and tissues crucial
- Enables analysis, diagnosis, and treatment

Medical Image Segmentation

- Segmentation is a fundamental task in medical image analysis
- Shape of organs and tissues crucial
- Enables analysis, diagnosis, and treatment
- Manual segmentation most accurate, but too expensive
- Semi and fully automatic methods greatly decrease time required by an expert
- Difficulties: noise, large image sizes, partial volume effects, and anatomical variability

Energy Functions

- Formulate methods as an energy minimization problem
- Domain of the energy function is a set of possible segmentations:

$$\text{Optimal Segmentation} = \arg\min E(\text{Possible Segs.})$$

Energy Functions

- Formulate methods as an energy minimization problem
- Domain of the energy function is a set of possible segmentations:

$$\text{Optimal Segmentation} = \arg\min E(\text{Possible Segs.})$$

- Energy function construction:
  1. User input
  2. Image information
  3. Prior knowledge

Energy Minimization

- Global minimum cannot always be found

Energy Minimization

- Global minimum cannot always be found

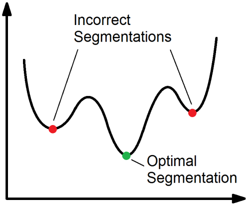

Nonconvex Function

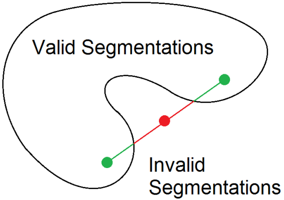

Nonconvex Domain

Convexity

- Nonconvex energies, local minimum found depends on
  1. Optimization scheme
  2. Initialization

Convexity

- Nonconvex energies, local minimum found depends on
  1. Optimization scheme
  2. Initialization
- But, convexity limits expressibility

Convexity

- Nonconvex energies, local minimum found depends on
  1. Optimization scheme
  2. Initialization
- But, convexity limits expressibility
- Question 1: What features can our function include while maintaining convexity?

Image Information

- Energy functions must incorporate image information to be meaningful
- Many convex terms exist

Image Information

- Energy functions must incorporate image information to be meaningful
- Many convex terms exist
- Image information often not sufficient...

Shape Priors

- Regions' shapes usually conform to a distribution
- Shape priors greatly increase accuracy

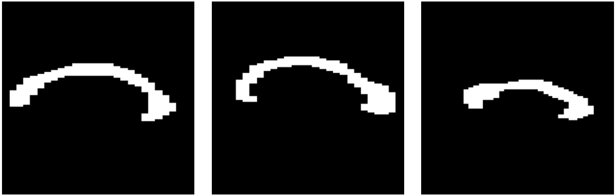

Corpus Callosum

Shape Priors

- Regions' shapes usually conform to a distribution
- Shape priors greatly increase accuracy

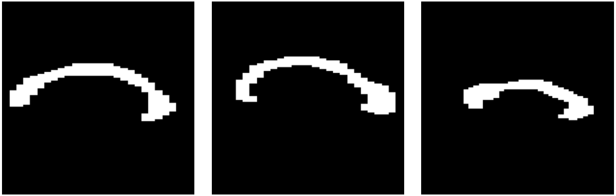

Corpus Callosum

- A complex shape prior may result in nonconvexity

Convex Domain

- Many segmentation representations exist
- Nonconvexity allows more descriptive, anatomically justified representations

Convex Domain

- Many segmentation representations exist
- Nonconvexity allows more descriptive, anatomically justified representations
- If domain nonconvex, global minimum may be unattainable

Convex Domain

- Many segmentation representations exist
- Nonconvexity allows more descriptive, anatomically justified representations
- If domain nonconvex, global minimum may be unattainable
- Question 2: What can we encode in a representation while maintaining convexity?

Binary vs. Multi-Region

- Often multiple regions of interest
- Multi-region representations explicitly encode regional interactions

Binary vs. Multi-Region

- Often multiple regions of interest
- Multi-region representations explicitly encode regional interactions

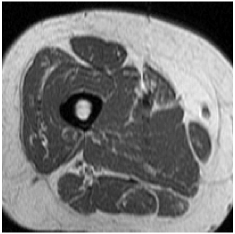
Thigh MRI

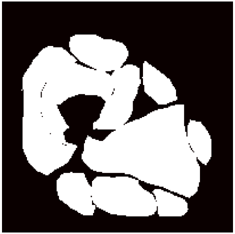
Binary

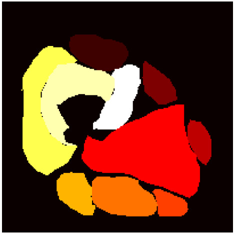
Multi-Region

Crisp vs. Probabilistic

- Probabilistic representations encode uncertainty
- Partial volume effect, probabilistic prior models, etc.

Crisp vs. Probabilistic

- Probabilistic representations encode uncertainty
- Partial volume effect, probabilistic prior models, etc.

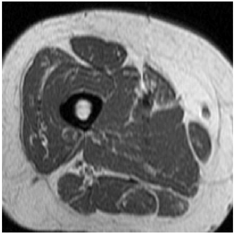

Thigh MRI

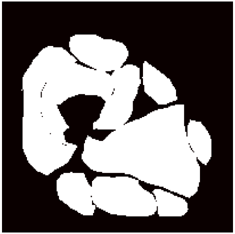

Crisp

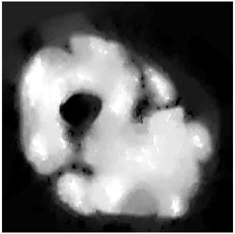

Probabilistic

Method

- Answering Question 1 & 2: We create an energy incorporating probabilistic, multi-region shape priors, while maintaining convexity
- First try: create shape priors using principal component analysis (PCA) on training segmentations (Cremers et al. '08)

Method

- Answering Question 1 & 2: We create an energy incorporating probabilistic, multi-region shape priors, while maintaining convexity
- First try: create shape priors using principal component analysis (PCA) on training segmentations (Cremers et al. '08)
- But this has limitations...

Notation

- Ω: image domain, with n pixels
- R: number of regions
- $\mathcal{S}^R$: the simplex of size R,

$$\mathcal{S}^R = \left\{ \{x_1, \ldots, x_R\} \in \mathbb{R}^R \,\middle|\, \sum_{r=1}^{R} x_r = 1 \right\}$$

- Probabilistic, multi-region segmentation:

$$q : \Omega \to \mathcal{S}^R$$

Training Data

- Enforce a statistically feasible segmentation space
- PCA on training data
- N ground truth (GT) segmentations:

$$\{q_1, \ldots, q_N\}$$

Principal Component Analysis

- q_0: mean of training GTs
- Ψ: an $(nR) \times k$ matrix of the k eigenmodes
- Statistically feasible segmentations parameterized by $\alpha \in \mathbb{R}^k$:

$$q(\alpha) = q_0 + \Psi\alpha$$

Principal Component Analysis

- q_0: mean of training GTs
- Ψ: an $(nR) \times k$ matrix of the k eigenmodes
- Statistically feasible segmentations parameterized by $\alpha \in \mathbb{R}^k$:

$$q(\alpha) = q_0 + \Psi\alpha$$

- Not all α give valid segmentations

Simplex PCA

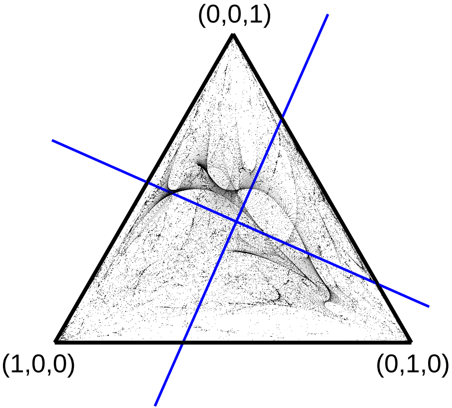

Simplicial Geometry

- Want: $T : \mathcal{S}^R \to \mathbb{R}^{R-1}$

Simplicial Geometry

- Want: $T : \mathcal{S}^R \rightarrow \mathbb{R}^{R-1}$
- For $p = \{p_1, \ldots, p_R\} \in \mathcal{S}^R$, LogOdds (Pohl et al. '08):

$$\text{LogOdds}(p) = \left\{ \log \frac{p_1}{p_R}, \ldots, \log \frac{p_{R-1}}{p_R} \right\} \in \mathbb{R}^{R-1}$$

Simplicial Geometry

- Want: $T : \mathcal{S}^R \rightarrow \mathbb{R}^{R-1}$
- For $p = \{p_1, \ldots, p_R\} \in \mathcal{S}^R$, LogOdds (Pohl et al. '08):

$$\text{LogOdds}(p) = \left\{ \log \frac{p_1}{p_R}, \ldots, \log \frac{p_{R-1}}{p_R} \right\} \in \mathbb{R}^{R-1}$$

- LogOdds not symmetric, but the isometric log-ratio (ILR) transform (Egozcue et al. '03) is:

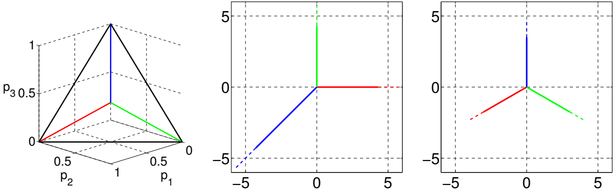

Simplex LogOdds ILR

Simplicial Geometry (cont'd)

- (Aitchison '86): a Hilbert space structure for $\mathcal{S}^R$

Simplicial Geometry (cont'd)

- (Aitchison '86): a Hilbert space structure for $\mathcal{S}^R$
- $p, q \in \mathcal{S}^R$, $\alpha \in \mathbb{R}$, $\mathcal{C}[\cdot]$: normalization function, $g(\cdot)$: geometric mean

$$p \oplus q = \mathcal{C}[p_1 q_1, \ldots, p_n q_n] \ ,$$

$$\alpha \odot p = \mathcal{C}[p_1^\alpha, p_2^\alpha, \ldots, p_n^\alpha] \ ,$$

$$\langle p, q \rangle_S = \sum_{i=1}^{n} \log \frac{p_i}{g(p)} \log \frac{q_i}{g(q)} \ ,$$

$$d_S(p, q) = \sqrt{\sum_{i=1}^{n} \left(\log \frac{p_i}{g(p)} - \log \frac{q_i}{g(q)} \right)^2} \ .$$

ILR Representation

- ILR maps operations to Euclidean counterparts

ILR Representation

- ILR maps operations to Euclidean counterparts
- Find basis for $\mathcal{S}^R$, $E = \{e_1, \ldots, e_{R-1}\}$
- ILR projects onto E:

$$\text{ILR}(p) = (\langle p, e_1 \rangle_S, \ldots \langle p, e_{R-1} \rangle_S) \in \mathbb{R}^{R-1}.$$

ILR Representation

- ILR maps operations to Euclidean counterparts
- Find basis for $\mathcal{S}^R$, $E = \{e_1, \ldots, e_{R-1}\}$
- ILR projects onto E:

$$\mathrm{ILR}(p) = (\langle p, e_1 \rangle_S, \ldots \langle p, e_{R-1} \rangle_S) \in \mathbb{R}^{R-1}.$$

- Probabilistic, multi-region segmentation:

$$\eta = \mathrm{ILR}(q) : \Omega \to \mathbb{R}^{R-1}$$

ILR Visualization

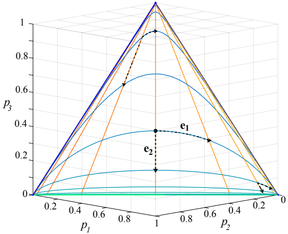

$\mathcal{S}^3$

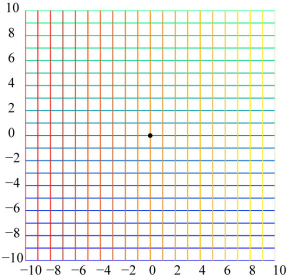

$\mathbb{R}^2 = \text{ILR}(\mathcal{S}^3)$

PCA Revisited

- GTs mapped to ILR space
- η_0: mean of training GTs
- Ψ': an $(n(R-1)) \times k$ matrix of the k eigenmodes

$$\eta(\alpha) = \eta_0 + \Psi' \alpha$$

PCA Revisited

- GTs mapped to ILR space
- η_0: mean of training GTs
- Ψ': an $(n(R-1)) \times k$ matrix of the k eigenmodes

$$\eta(\alpha) = \eta_0 + \Psi' \alpha$$

- Every $\alpha \in \mathbb{R}^k$ is valid
- Vector space representation achieved

ILR PCA

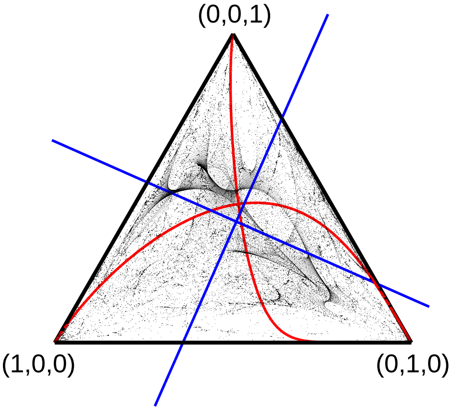

Energy Formulation

- Energy functional $E(\eta)$ is constructed with 3 terms:

$$E(\eta) = w_1 E_{\text{Intensity}}(\eta) + w_2 E_{\text{Gradient}}(\eta) + w_3 E_{\text{Shape}}(\eta)$$

- Each term $\Leftrightarrow$ different segmentation property

Intensity Energy Term

- Construct regional intensity distributions
- Intensity-based probabilistic segmentation:

$$p : \Omega \to \mathcal{S}^R$$

Intensity Energy Term

- Construct regional intensity distributions
- Intensity-based probabilistic segmentation:

$$p : \Omega \to \mathcal{S}^R$$

- Use squared distance:

$$E_{\text{Intensity}}(\eta) = \int_x \left\| \eta(x) - \text{ILR}(p(x)) \right\|^2 d\Omega$$

Gradient Energy Term

- $b(x)$: boundary indicator at pixel x
- $h(x)$: measure of the rate of segmentation change:

$$\eta(x) = \{\eta_1(x), \ldots, \eta_{R-1}(x)\} \in \mathbb{R}^{R-1}$$
$$h(x) = \sum_{r=1}^{R-1} \|\nabla_x \eta_r(x)\|^2$$

Gradient Energy Term

- $b(x)$: boundary indicator at pixel x
- $h(x)$: measure of the rate of segmentation change:

$$\eta(x) = \{\eta_1(x), \ldots, \eta_{R-1}(x)\} \in \mathbb{R}^{R-1}$$

$$h(x) = \sum_{r=1}^{R-1} \|\nabla_x \eta_r(x)\|^2$$

- Gradient-based term:

$$E_{\text{Gradient}}(\eta) = \int_x (1 - b(x))h(x)d\Omega$$

Shape Energy Term

- Modes with greater variance should have greater freedom
- Λ: $k \times k$ diagonal matrix with variances the diagonal

$$E_{\text{Shape}}(\eta(\alpha)) = \alpha^T \Lambda^{-1} \alpha$$

Shape Energy Term

- Modes with greater variance should have greater freedom
- Λ: $k \times k$ diagonal matrix with variances the diagonal

$$E_{\text{Shape}}(\eta(\alpha)) = \alpha^T \Lambda^{-1} \alpha$$

- Strictly convex $\Rightarrow E$ is strictly convex

Final Segmentation

- $\eta(\alpha)$ is linear in $\alpha \Rightarrow E(\eta(\alpha))$ is convex in α
- Final segmentation q^*:

$$\alpha^* = \arg\min_{\alpha} E(\eta(\alpha)), \quad q^* = \mathrm{ILR}^{-1}(\eta(\alpha^*))$$

- Convex, multi-region, and includes a general shape prior

Feature Table

	Features			
Methods	**Convex Energy**	**Shape Prior**	**Multi-Region**	**Probab-ilistic**
Chan '04	√	X	X	√ *
Ishikawa '03	√	X	√ †	X
Veksler '08	√	√ ‡	X	X
Vu '08	X	√	√	X
Pock '08, Brown '09, Delong '09	√	X	√	X
Lellmann '09, Pock '09, Zach '08	√	X	√	√ *
Song '10	√	√ ‡	√	X
Pohl '07	X	√	√	√
Grady '05	√	X	√	√
Cremers '08	√	√	X	√
Our Method	√	√	√	√

∗ Relaxed 0-1 segmentations could be informally treated as probabilities.

† Allows only limited regional interaction terms.

‡ Only applicable to restricted classes of shapes.

Implementation Details

- Simple affine registration scheme sufficiently aligned images
- Energy term weights determined from the training data
- Validation using a leave-one-out heuristic

Thigh MRI Data

- 40 volumetric MRI thigh scans of size $175 \times 175 \times 85$ had all 11 knee extensor and flexor muscles segmented, for 12 regions
- Despite very poor intensity priors and borders, our method achieved an average DSC of 0.92 ± 0.03 with the GT

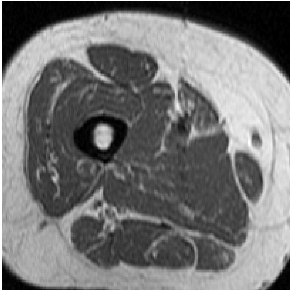

Thigh Muscle Segmentation

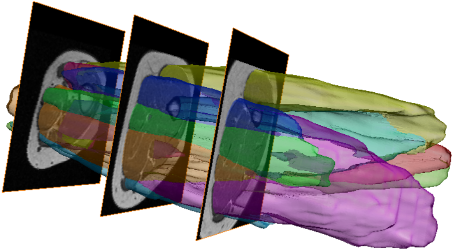

- A resulting segmentation of our method overlaid on several image slices

Thigh Muscle Segmentation

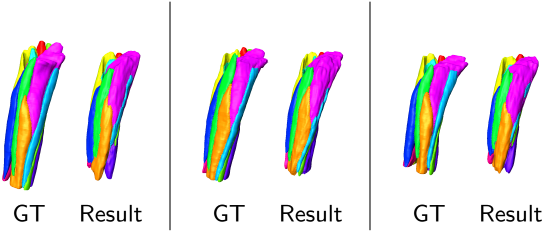

Conclusion

- Multi-region shape priors while maintaining convexity

Conclusion

- Multi-region shape priors while maintaining convexity
- Strictly convex unconstrained energy

Conclusion

- Multi-region shape priors while maintaining convexity
- Strictly convex unconstrained energy
- ILR-based segmentation simplifies statistical analysis

Future Work

- Include pose estimation

Future Work

- Include pose estimation
- Segmentation tasks with weak image information but strong shape priors

Future Work

- Include pose estimation
- Segmentation tasks with weak image information but strong shape priors
- More flexible shape spaces

Questions

- ?