Continuous Medial Models in Two-Sample Statistics of Shape

Timothy B. Terriberry

Outline

- Introduction
    - Motivation
    - Background
    - Contributions and Thesis Statement
- A New Continuous Medial Model
- Nonlinear Hypothesis Testing
- Conclusion

Motivation

- Do two groups of anatomical shapes show statistically significant differences?

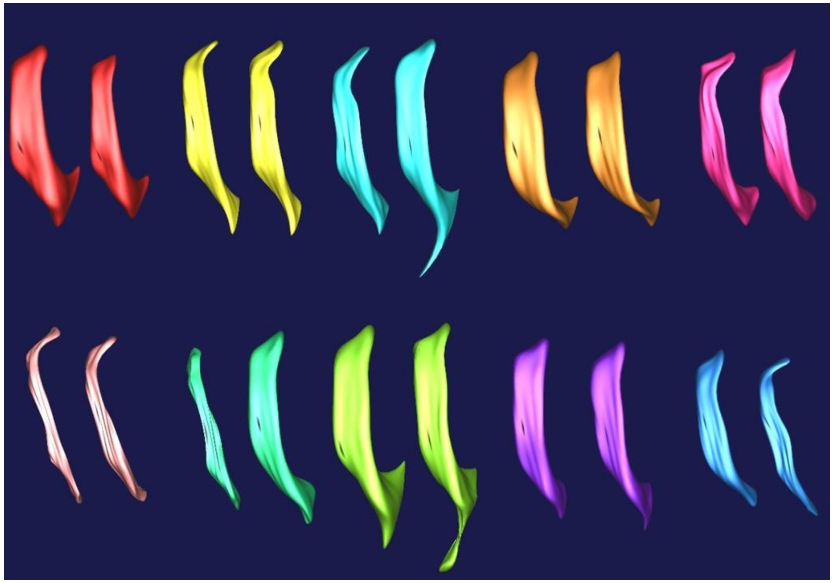

MZ

DZ

Motivation

- Requires answers to the questions:
- How should the shapes be represented?
- How can these representations be compared?

Which Shape Representation?

- The medial axis (Blum 1967)
  - Evidence suggests used by HVS
  - Naturally decomposes objects into articulated parts
  - Represents intuitive growth and deformation as local properties
    - Bending
    - Twisting
    - Thickening/thinning
  - Good coarse-scale shape representation with a small number of parameters

The Medial Axis in 2D

- **Definition: the closure of the centers of all maximally inscribed balls**

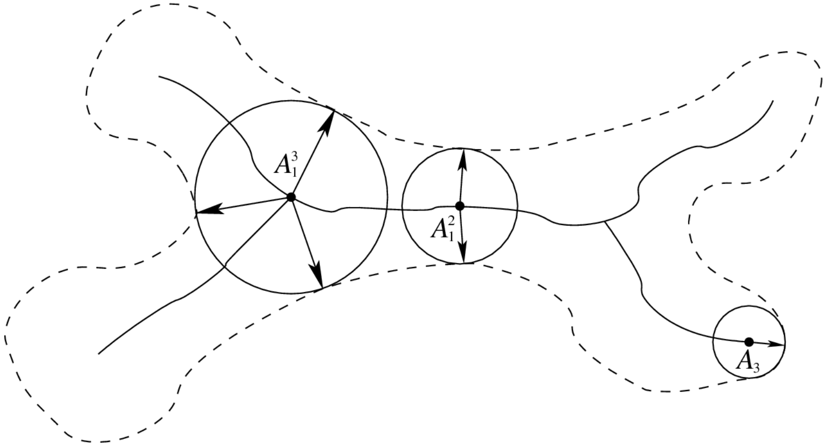

- **When coupled with the radius, gives a complete shape description**

Medial-Boundary Link

- Spoke vectors point from the medial axis to points of tangency on the boundary
  - Maps between the medial axis and the boundary
  - Can be used to analyze properties of the boundary in object-relative, medial coordinates
  - Two spokes per axis point, except at branch, end points
    - Arranged symmetrically around the axis

The Medial Axis in 3D

- Definition easy to extend to 3D, but geometry more complicated

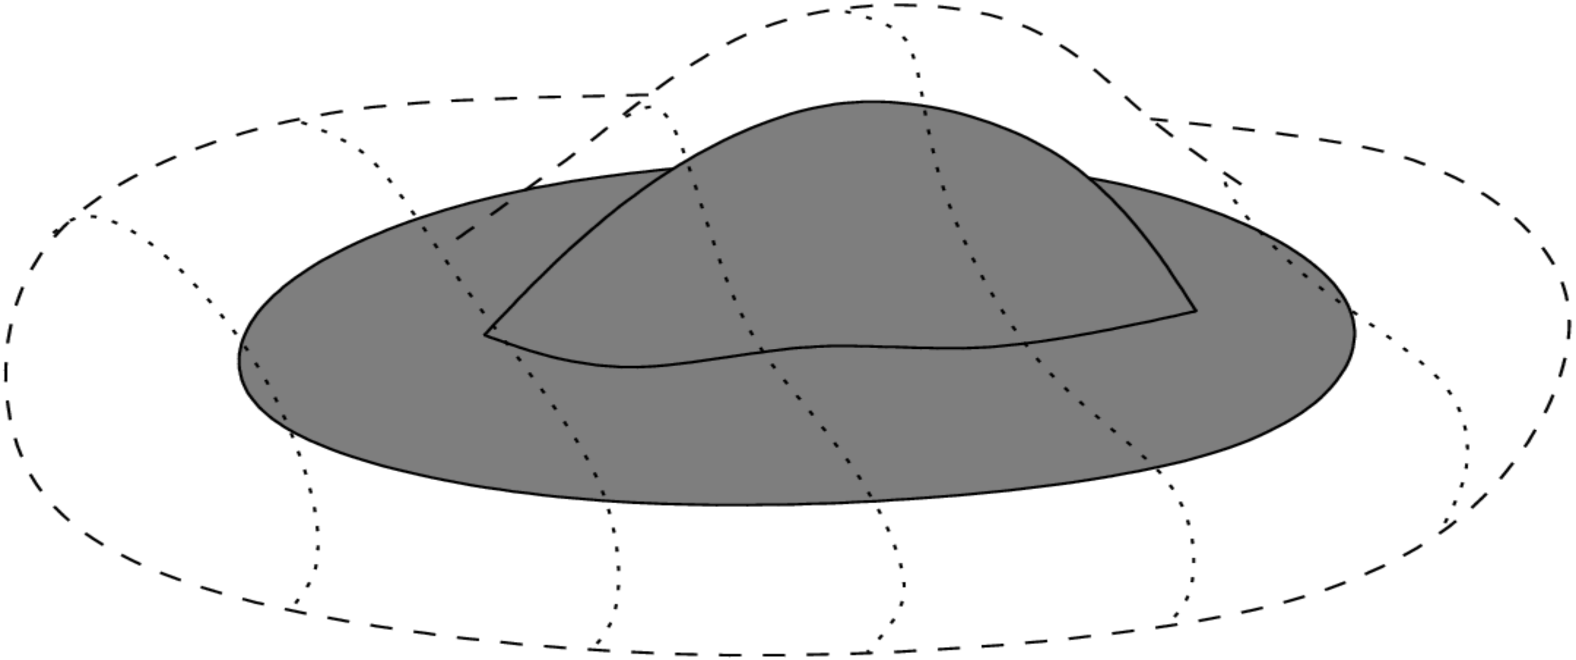

The Generic 3D Medial Axis

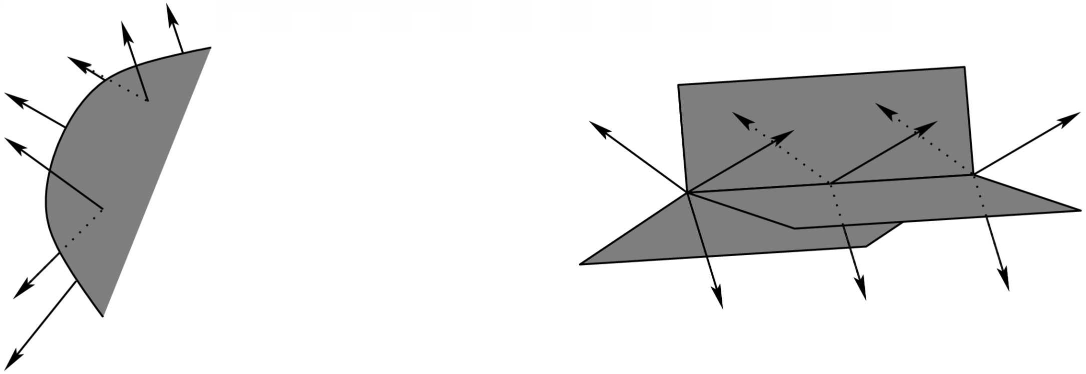

(a) edge curve

(b) branch curve

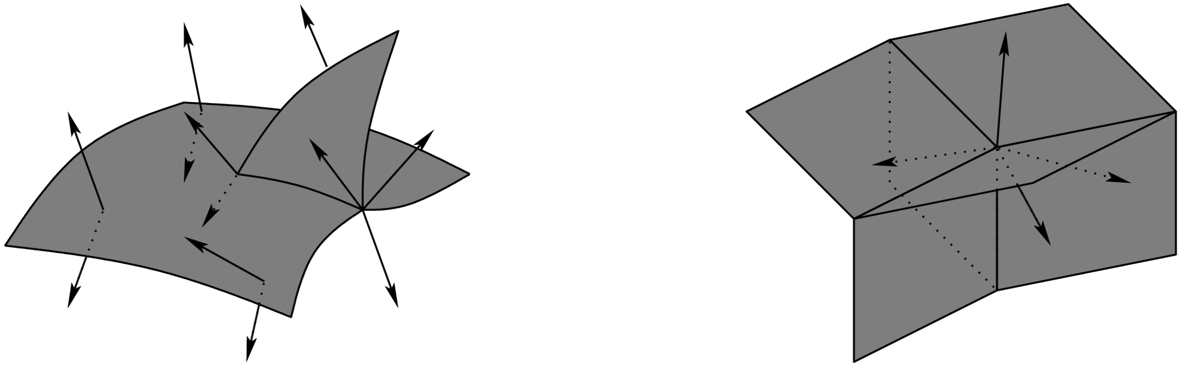

(c) fin point

(d) 6-junction

The Generic 3D Medial Axis

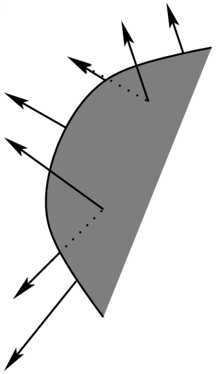

(a) edge curve

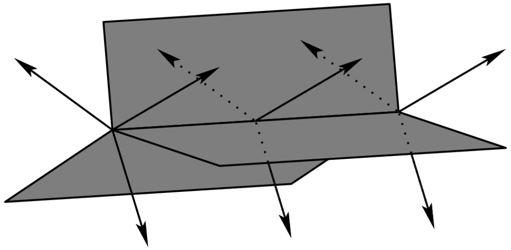

(b) branch curve

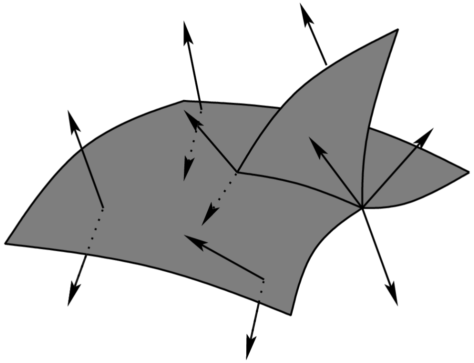

(c) fin point

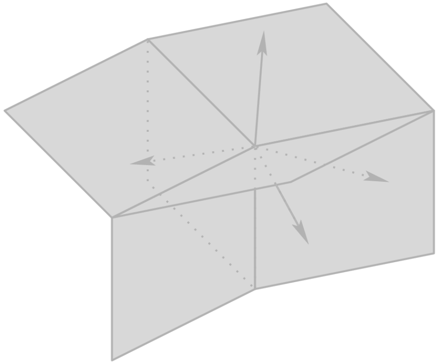

(d) 6-junction

Active Shape Models

- Want a set of models with common topology for statistics
- Medial topology is unstable to small perturbations
- Stable approach
  - Construct a template with a fixed topology
  - Deform the template to fit each target shape

Discrete m-reps

- **Coarsely sampled medial axis**
- **Medial-Boundary Link Tenuous**
  - **Boundary reconstructed by interpolating spoke ends (Thall, 2004), does not respect medial geometry**
  - **New interpolation of Han et al. (2006) relaxes medial constraints**
    - **Shape representation is not unique**
    - **Only approximates the discrete samples**
    - **Requires expensive numerical integration**

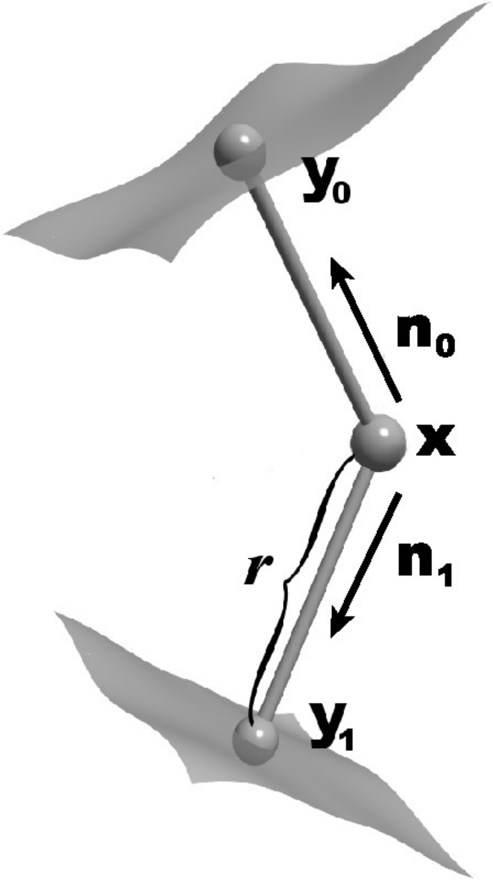

Continuous m-reps

- Interpolate the medial axis; interpolated boundary follows
  - Must ensure medial axis is valid
  - Easy in 2D: finite number of singular points, finite number of boundary conditions
  - Hard in 3D: infinitely many singular points, finite number of parameters with which to enforce boundary conditions

Previous Work

- Yushkevich et al. 2003
  - Interpolated axis, radius with B-splines
  - Solved implicitly for edge of axis
  - Limited to a single sheet
  - Every shape has a different domain, thus difficult to compare groups of shapes

Previous Work

- Yushkevich et al. 2005
  - Interpolated axis, potential function with Fourier basis functions
  - Solved PDE to produce a radius satisfying appropriate boundary conditions
  - Every shape has a fixed domain
  - Not enough free parameters in the PDE solver to allow branching

Thesis Contribution

- A novel "control curve" formulation capable of enforcing first- and second-order boundary conditions on the edges of subdivision surfaces
- A new continuous medial model based on this fomulation
  - ◆ Speed of cubic B-splines
  - ◆ Common domain for all shapes
  - ◆ Supports branching in 3D

How to Compare Shapes?

- **Must have a notion of "correspondence"**
    - **Which pieces of one shape match those on another?**
- **Must have a procedure for comparing points in nonlinear shape spaces**

Medial Correspondence

- Parameterization of a continuous m-rep gives an implicit correspondence
- But control points can wander during the model fitting process
  - A slowly-changing radius allows points to move without a large effect on the boundary
- Current approaches indirect
  - Add penalty terms for regularity, deformation from template, etc.

Enforcing Correspondence

- Correspondence cannot be derived from shape information alone
- Instead, take correspondence on the surface/volume as input
  - Compute with SPHARM, DetCov, MDL, symmetric fluid warping, image features, etc.

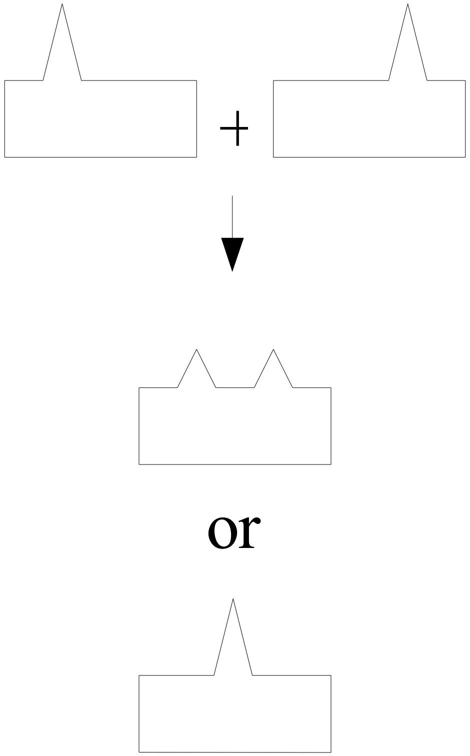

Thesis Contribution

- A correspondence optimization method that works in tandem with the model fitting process to produce a group of models with a common parameterization
- Utilizes the medial-boundary link to optimize implicit medial correspondence to match explicit boundary correspondence

Hypothesis Testing

- Comparing statistical distributions in shape spaces is challenging
  - High dimension, low sample size
  - Shape parameters are non-Euclidean
  - Different variables are not commensurate
  - Neighboring sites are correlated

Previous Work

- Previous studies focused on a single aspect of m-reps at a time
    - Position
    - Radius
- Ignored spoke direction, even though it is essential to describe twisting, bending

Thesis Contribution

- A novel nonlinear hypothesis test for m-reps, which simultaneously considers all of the shape parameters
  - ◆ Valid in any direct product of metric spaces
  - ◆ Reduces to an ordinary permutation test in the univariate case
  - ◆ Converges to the classic Hotelling's T^2 test in linear case

Thesis Statement

- ***Nonlinear growth and deformation of biological objects requires nonlinear shape models to effectively characterize these processes.***
- ***A nonlinear hypothesis test based on multivariate permutation tests provides an effective means of detecting and localizing their effects between groups.***
- ***A continuous medial model, which describes this variation in a natural way, can be used to effectively represent these shapes.***
- ***The geometric link between the medial axis and the boundary such a model provides can be used to optimize the correspondence across a population of objects, reducing parameterization noise to increase the power of statistical tests.***

Outline

- Introduction
- A New Continuous Medial Model
    - Boundary Reconstruction
    - Edge Curves
    - Branch Curves/Fin Points
    - Optimization
- Nonlinear Hypothesis Testing
- Conclusion

Boundary Reconstruction

- Defined by a continuous medial axis *m* and radius function *r*
- Boundary recovered via:

$$B = \boldsymbol{m} + r\vec{U}$$

$$\vec{U} = -\nabla r \pm \sqrt{1-\|\nabla r\|^{2}}\,\hat{n}$$

- Here ∇r is the Riemannian gradient:

$$\nabla r = \begin{bmatrix}\boldsymbol{m}_u & \boldsymbol{m}_v\end{bmatrix}\begin{bmatrix}\boldsymbol{m}_u \cdot \boldsymbol{m}_u & \boldsymbol{m}_u \cdot \boldsymbol{m}_v \\ \boldsymbol{m}_v \cdot \boldsymbol{m}_u & \boldsymbol{m}_v \cdot \boldsymbol{m}_v\end{bmatrix}^{-1}\begin{bmatrix}r_u \\ r_v\end{bmatrix}$$

Edge Curves

- Boundary condition:

$$\|\nabla r\| = 1$$

- Must hold at every point on the edge
- Solution:
  - Construct spline patches via Catmull-Clark
  - Interpolate curves in the edge direction
  - Convert to an interpolating basis to fix edge curve
  - Replace one curve with a control curve
  - Complete interpolation across the edge

Edge Curves

- **Start with a B-spline patch**

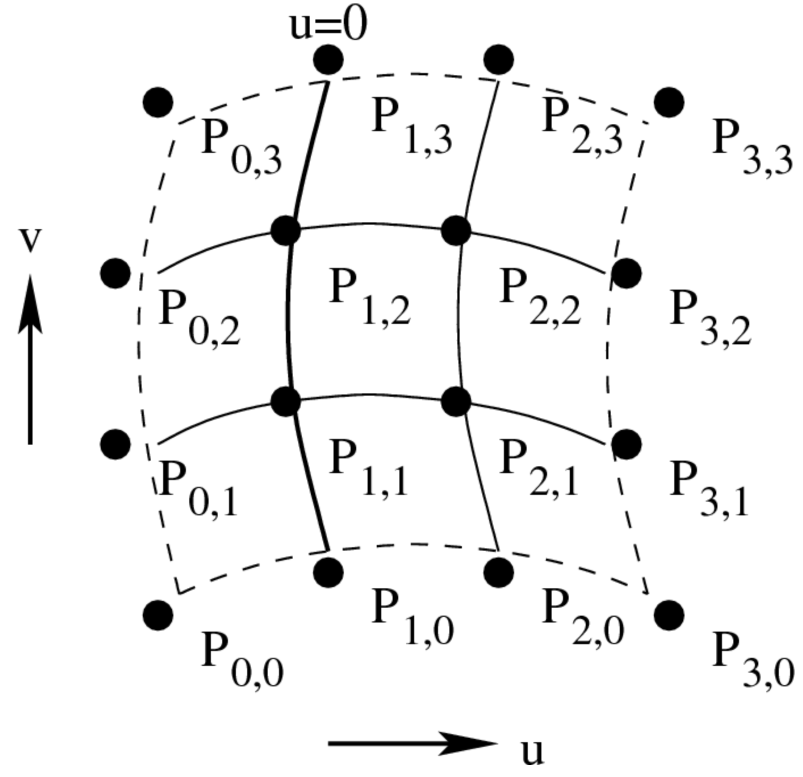

Edge Curves

- **Interpolate in the *v* direction**

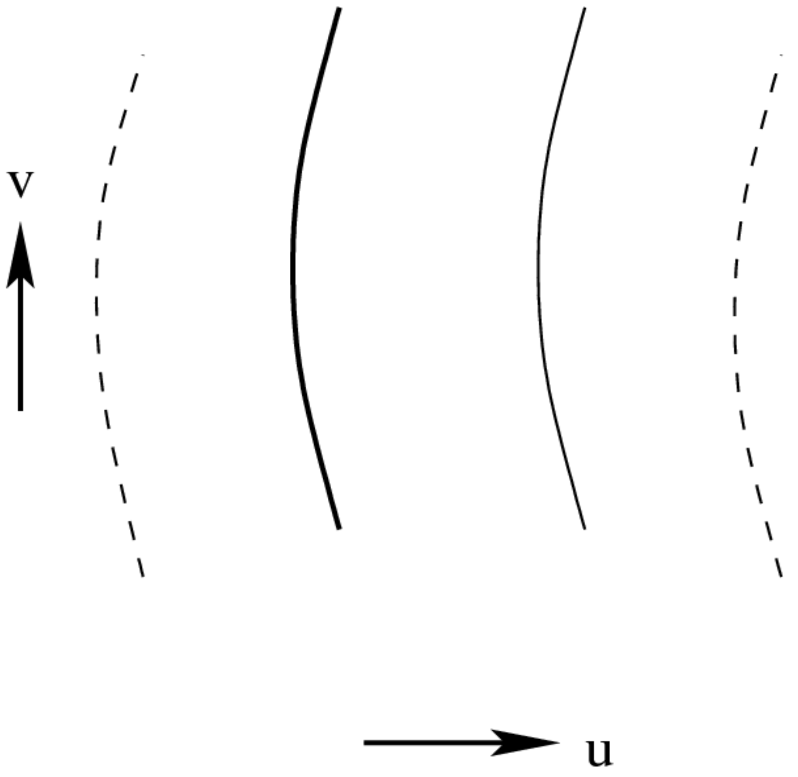

Edge Curves

- Convert to an interpolating basis in *u*

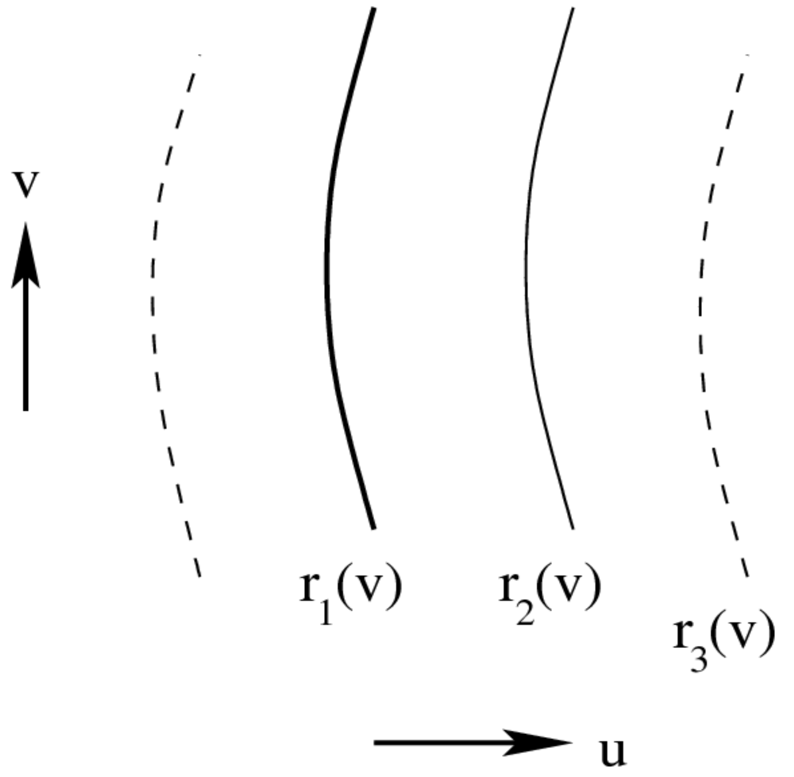

Edge Curves

- **Replace the curve on the left with a *control curve***
  - **Value of curve chosen to force $\|\nabla r\| = 1$ by setting r_u**

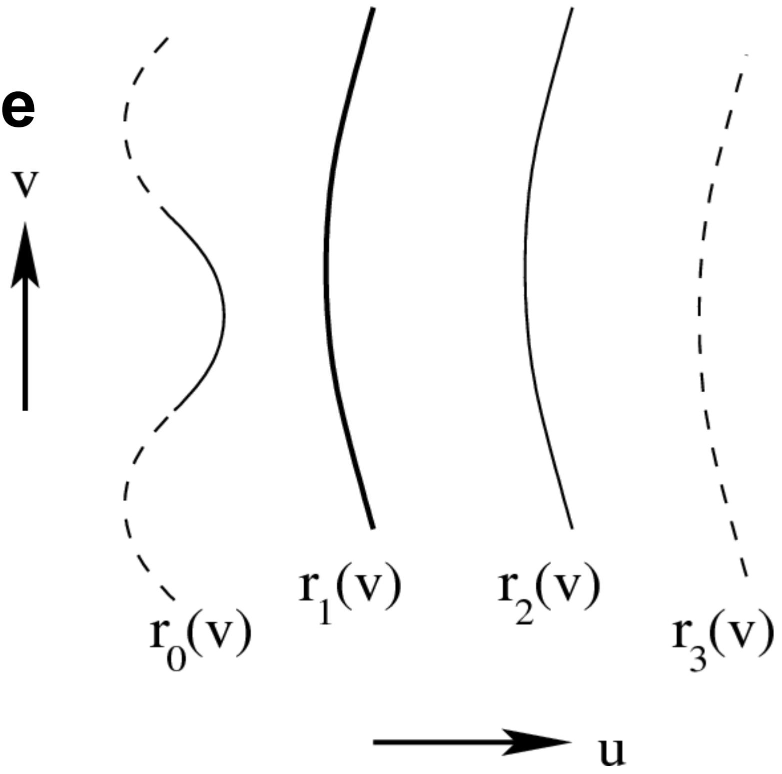

Edge Curves

- **Finish interpolating in the *u* direction**

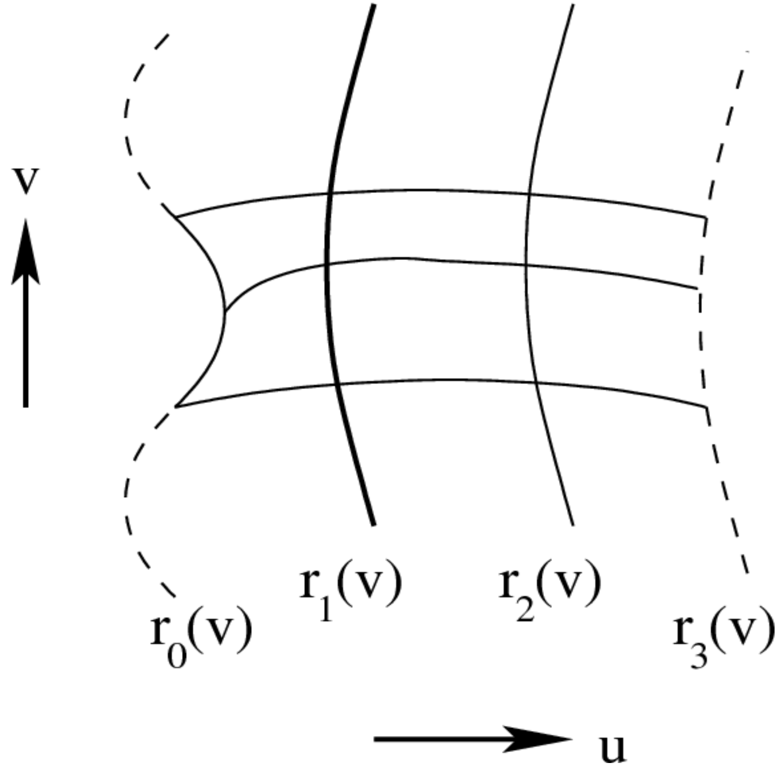

Edge Curves

- Final Result

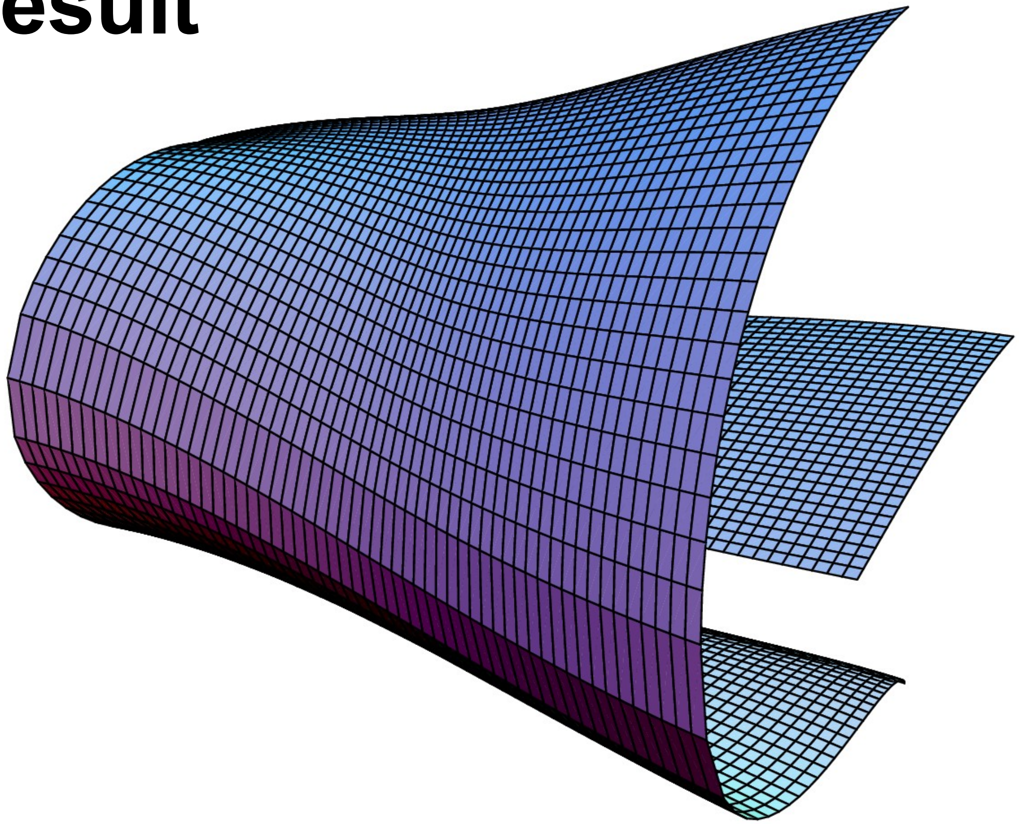

Edge Curves

- Interpolation is asymmetric in *u* and *v*
- Therefore the mesh must have no corners
  - All edge vertices must have valence 3
  - Either delete the corners or add an extra edge to make them valence 3

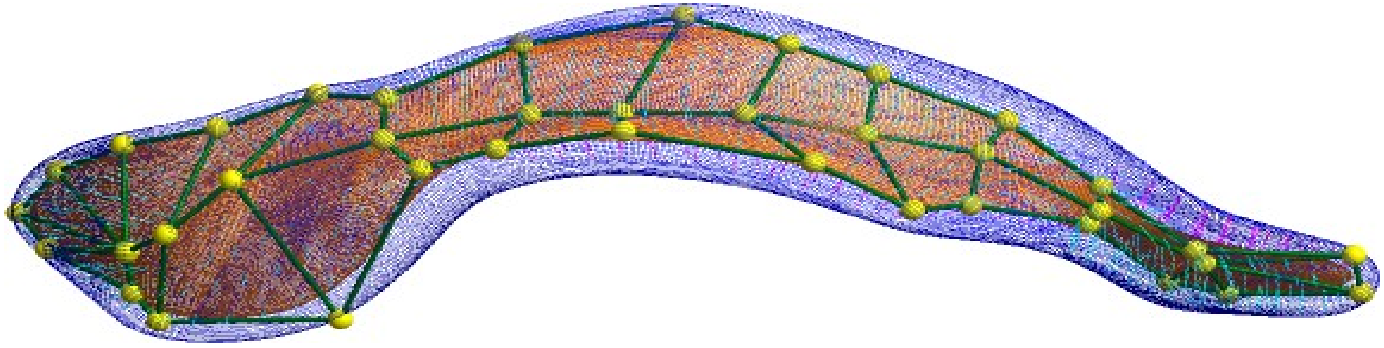

Branch Curves

- **Spoke directions must match neighbors to get a closed surface**
  - **Use a control curve to enforce**
- **Assume all 3 patches meet on the *u=0* curve (share control points)**

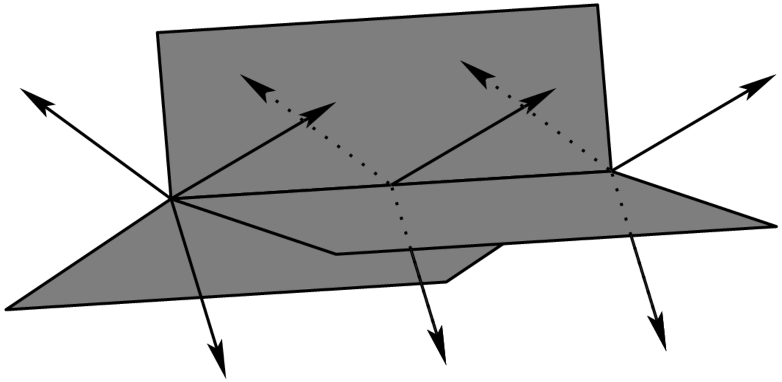

Branch Curves

- **The tips of -∇*r* and *U* lie in the same plane in all three patches**

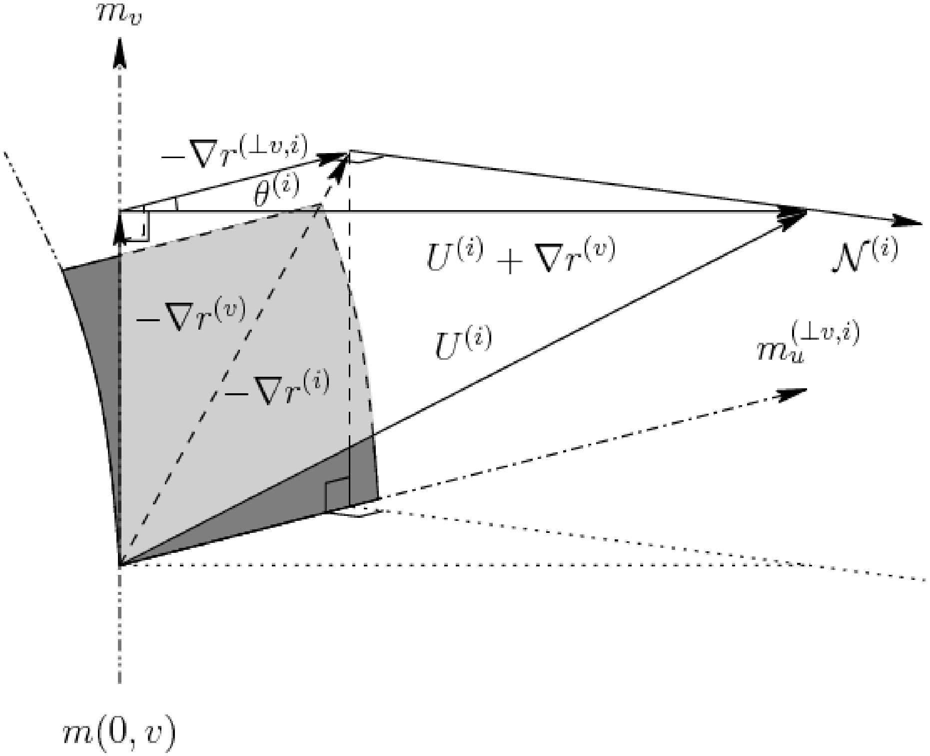

Branch Curves

- **Project everything into this plane**

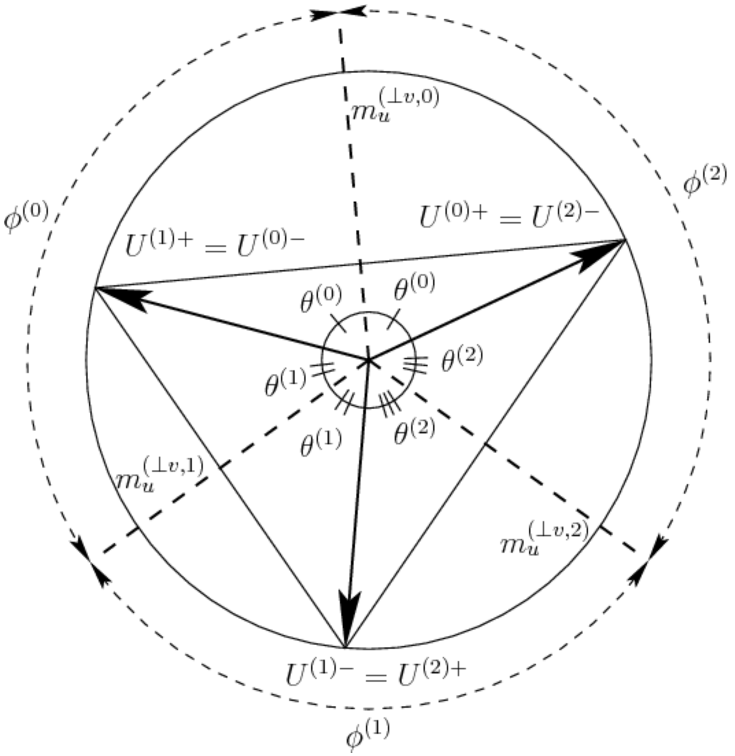

  - **Need**

$$\theta^{(0)}+\theta^{(1)}+\theta^{(2)}=\pi$$

  - **And**

$$\phi^{(i)}=\theta^{(i)}+\theta^{(i\oplus 1)}$$

  - **Can also be expressed in terms of ∇r, r_u**

Fin Points

- **Edge curve meets branch curve**
- **Enforce edge condition on patch** 0 **as normal, take** $\theta^{(0)}$ **as fixed**
- **Use a control curve on** r **to set** $\theta^{(1)}$ **and** $\theta^{(2)}$ **so they all sum to** π
- **Use another on** $\boldsymbol{m}$ **to set** $\boldsymbol{m}_u^{(\perp v,1)}$ **and** $\boldsymbol{m}_u^{(\perp v,2)}$ **to bisect them**
- **3 Additional constraints at fin point**

Fin Points

- Final result

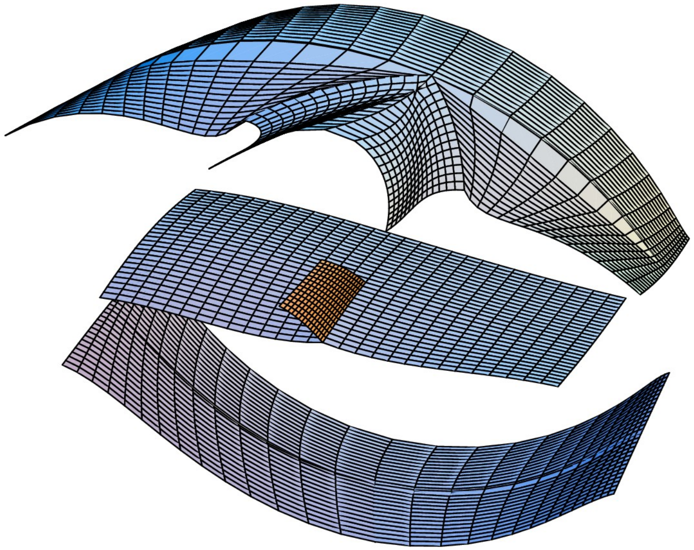

Optimization

- Objective functions constructed from medial integrals (Damon 2005)
  - Sample medial axis and approximate numerically
  - A “medial measure” weights each sample
  - Integrals over the boundary can be written and computed as medial integrals
  - Continuous representation affords exact, analytic gradients

Multiscale Optimization

- **Blur target binary image at successively smaller scales**
- **At each scale, match the boundary to a level set of the blurred image**

$$F_I^\sigma \triangleq \int_{\tilde{M}} \left(I_\sigma(m_0 + rU_0) - \ell_0\right)^2 \cdot \det\left(\mathbf{I} - rS_{\mathrm{rad}}\right)\, dM$$

Synthetic Ellipsoids

- Average Dice coefficient: 96.72%
- MAD over all cases: 0.318 voxels

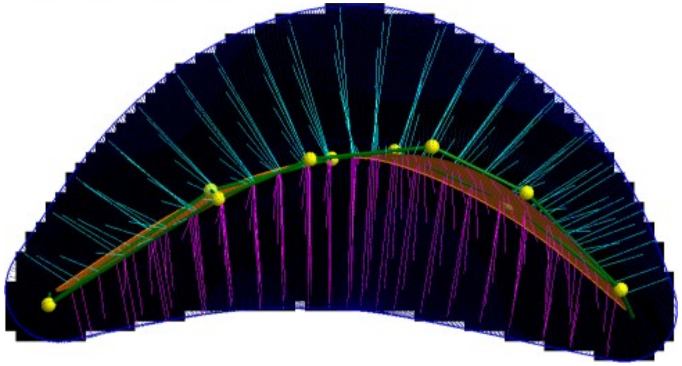

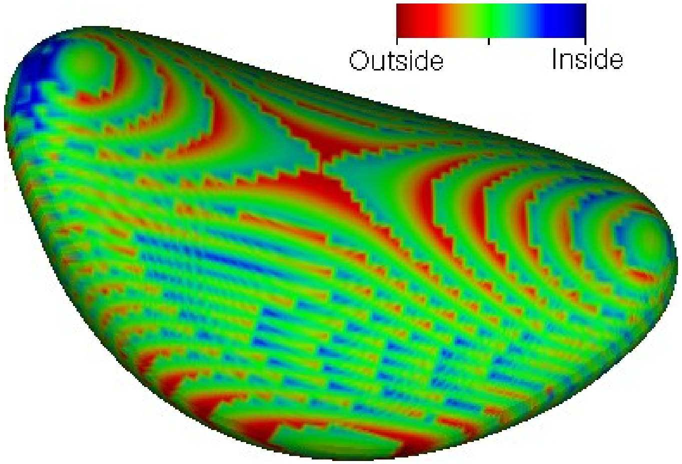

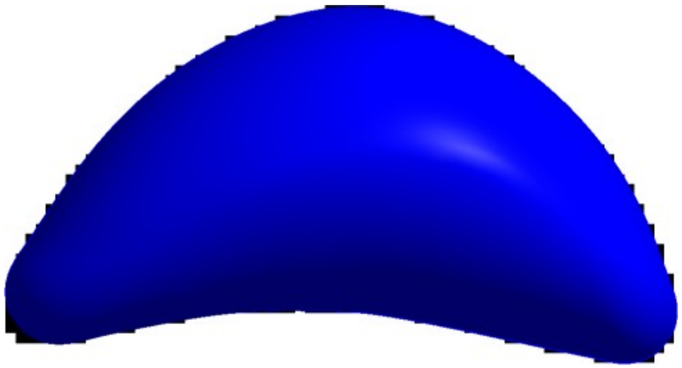

Autism Study

- **Average Dice coefficient: 93.92%**
- **MAD over all cases: 0.242 mm**

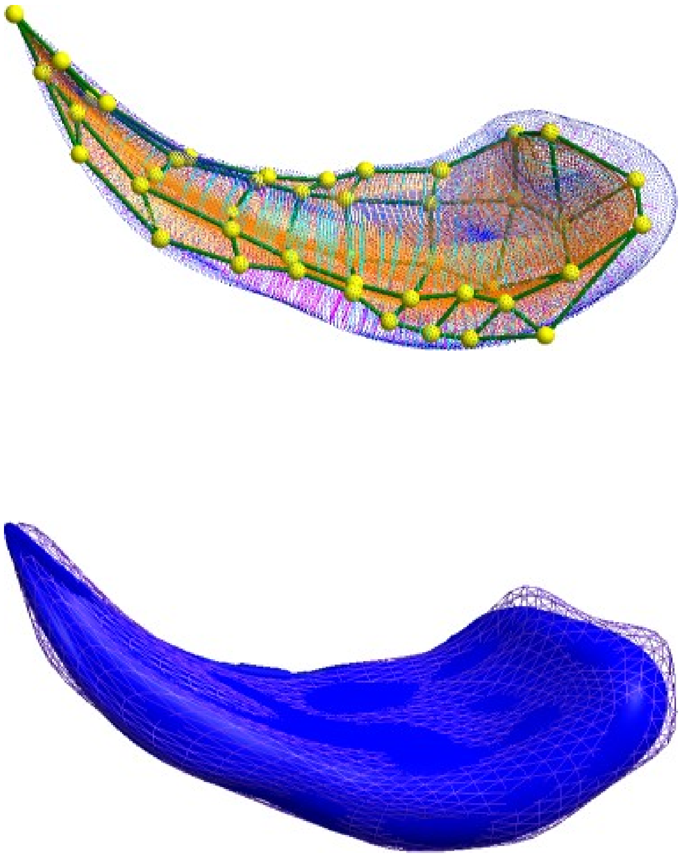

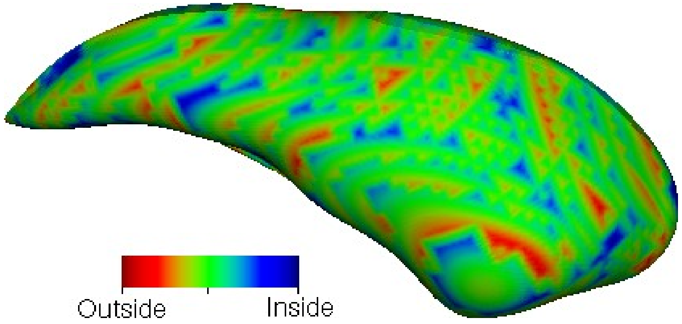

Optimizing Correspondence

- **Minimize mismatch between given correspondence and implicit medial correspondence**

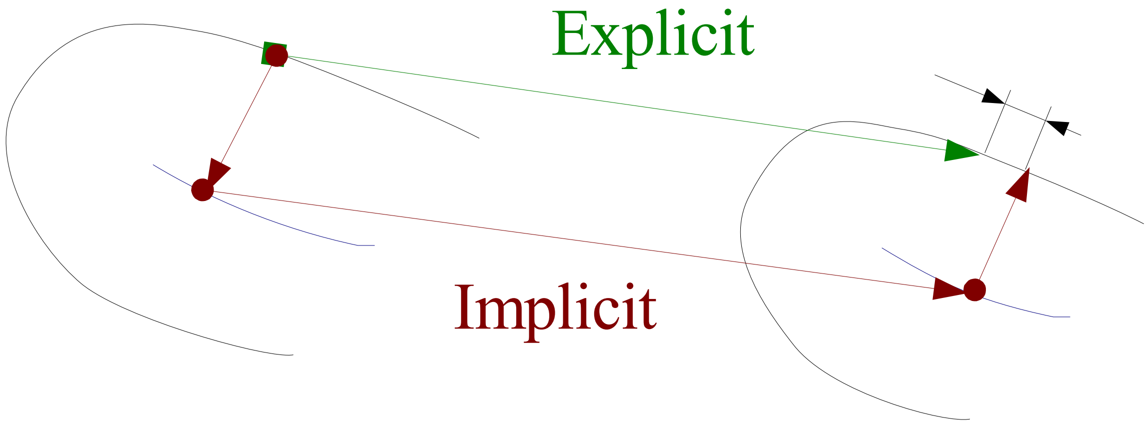

Optimizing Correspondence

- **Map implicitly corresponding points to a reference coordinate system using the explicit correspondence**

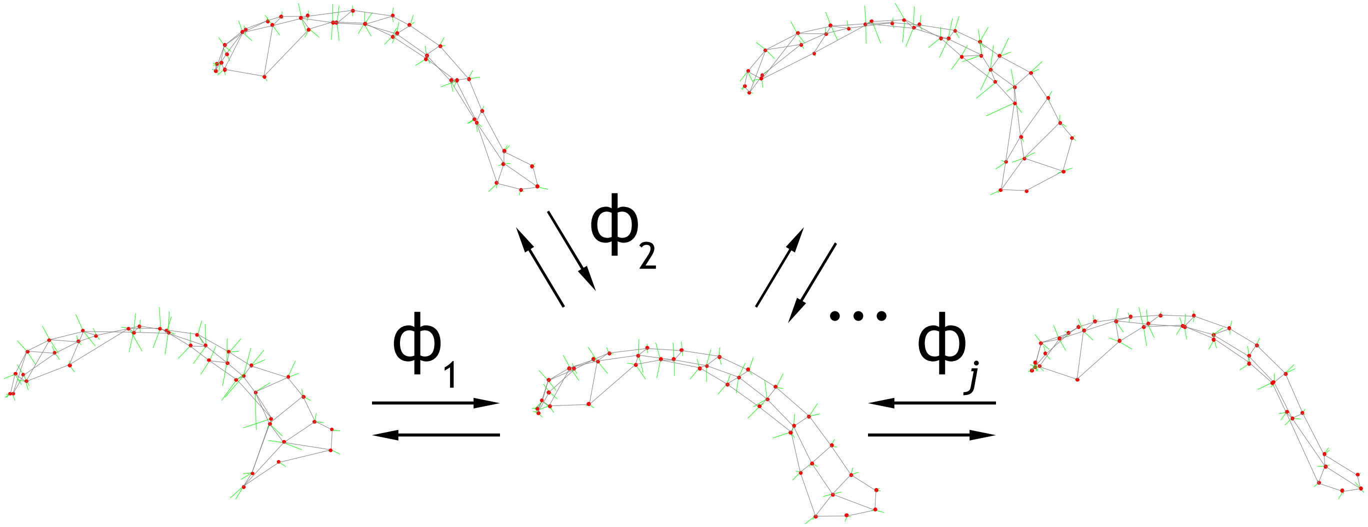

Optimizing Correspondence

- **Minimize size of clusters in reference coordinate system**
  - **Trivial minimizer: cluster mean**
  - **O(*n*) instead of O(n^2)**

$$F_C^j \triangleq \int_{\tilde{M}} \left(\phi_j(m_0 + rU_0) - \mu_{(m_0,U_0)}\right)^2 \cdot \det\left(\mathbf{I} - rS_{\text{rad}}\right)\, dM$$

Optimizing Correspondence

- How to trade off between model fit and correspondence match? Don’t.
- Take alternating steps of each one
- Optimize correspondence by sliding control points along model produced by last model fit step
  - Doesn’t change the model within the tolerance of the interpolation, only changes the parameterization

Synthetic Ellipsoids

- Ellipsoids deformed by known diffeomorphisms
- No correspondence optimization:

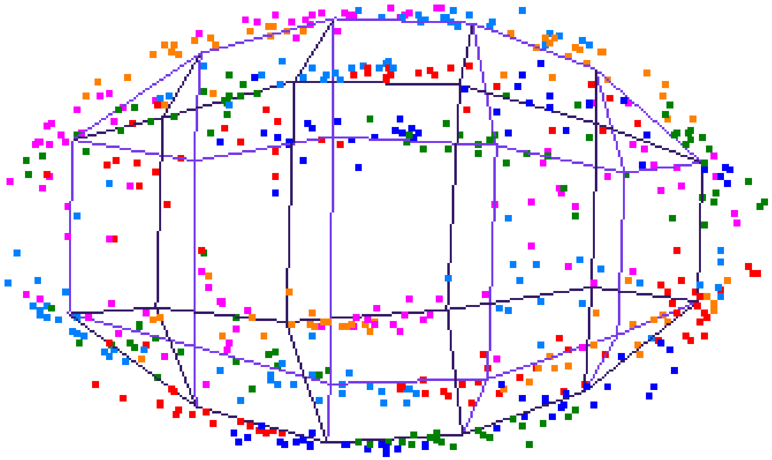

Synthetic Ellipsoids

- Ellipsoids deformed by known diffeomorphisms
- With correspondence optimization:

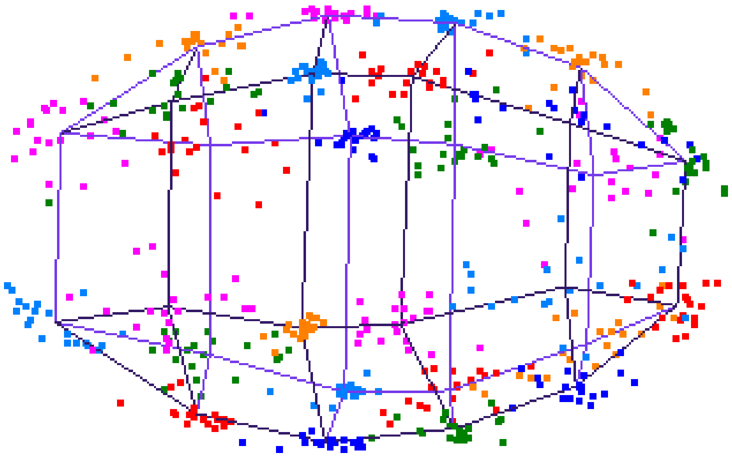

Autism Study

- **Left caudate, correspondence given by SPHARMs**
- **No correspondence optimization:**

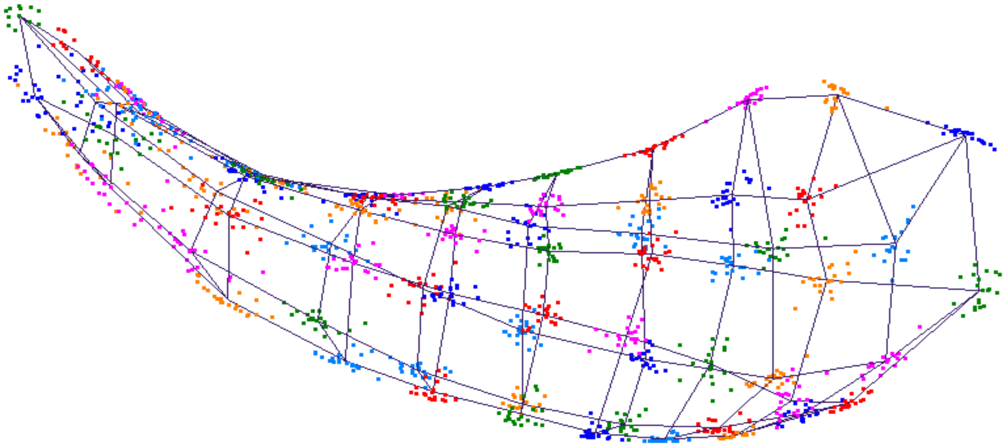

Autism Study

- **Left caudate, correspondence given by SPHARMs**
- **With correspondence optimization:**

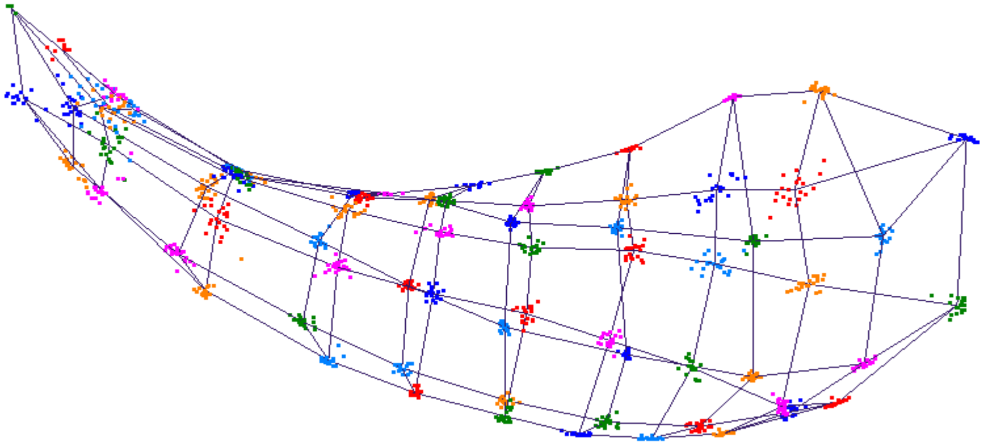

Outline

- Introduction
- A New Continuous Medial Model
- Nonlinear Hypothesis Testing
    - Metric for m-reps
    - Testing procedure
    - Twin Ventricle application
- Conclusion

Metric for M-reps

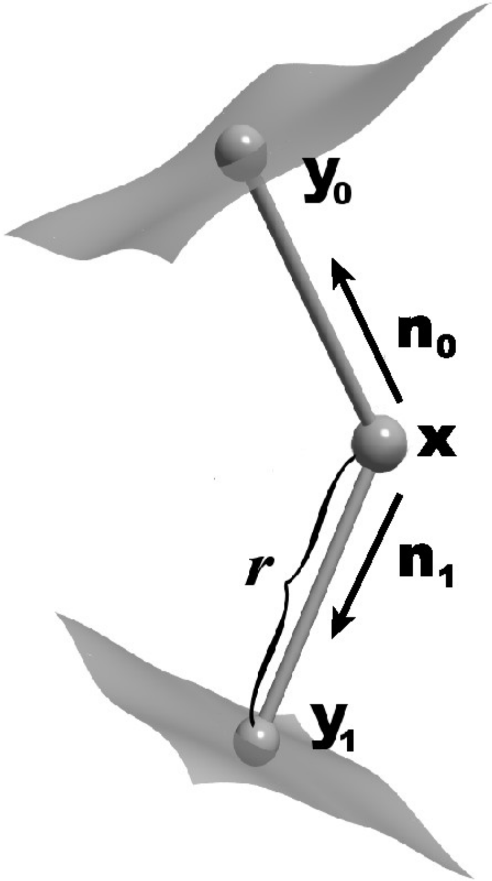

- 8 dimensions per medial atom
    - x (3), r (1), n_0 (2), n_1 (2)
- Riemannian symmetric space
    - $R^3 \times R^+ \times S^2 \times S^2$ (Fletcher et al. 2003)
    - Nonlinear, except R^3

Metric for M-reps

- **Each parameter has a metric invariant to geometric transformations**
  - **R^3 - Euclidean metric (invariant to translation)**
  - **R^+ - $|\log(r_1) - \log(r_2)|$ (invariant to scale)**
  - **S^2 - Distance on sphere (invariant to rotation)**
- **Can define the *Fréchet mean* via the metric**

$$\hat{\mu} = \underset{x \in M}{argmin} \sum_{i=1}^{n} d(x, x_i)^2$$

Testing Approaches

- Permutation tests:
  - Handles nonlinear, non-Gaussian
  - Usually univariate: requires a single, summarizing statistic
  - Assumes: isotropic distribution around mean
- Classic, linear: Hotelling's T^2 test
  - $T^2 \propto D^2 = (\hat{\mu}_1 - \hat{\mu}_2)^T \hat{\Sigma}^{-1} (\hat{\mu}_1 - \hat{\mu}_2)$
  - Handles correlation, scale
  - Assumes: populations have same (anisotropic) distribution around the mean

Our Approach

- Generalize permutation tests to capture desirable properties of Hotelling's test
  - Use a true multivariate permutation test framework (Pesarin 2001)
    - Perform partial tests on individual features
    - Combine the test results into a single score
  - Trivial example: Bonferroni correction
    - min p-value multiplied by number of tests
    - Too pessimistic for high-dimension data

Our Approach

- Marginal permutation tests on individual features generate *uniformly distributed* and *parameterization invariant* p-values
- Using a c.d.f., map the uniform distribution to a standard distribution, and perform tests there
- Gives an *unbiased* global test for equality of population distributions

Example

- Two data sets
  - Size $n_1 = n_2 = 10$
- M=2 dimensional feature vectors
  - Position, Scale
- Drawn from multivariate normal distributions (common covariance)
  - Second parameter exponentiated
  - Then both parameters scaled by 10

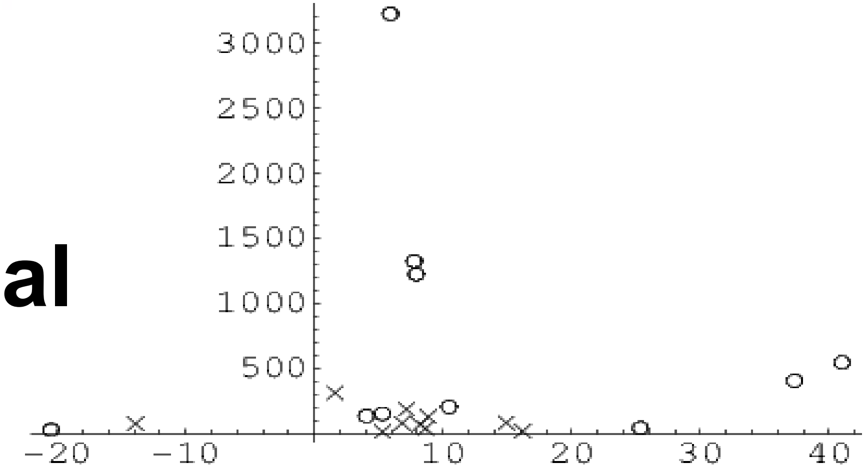

Step 1: Partial Tests

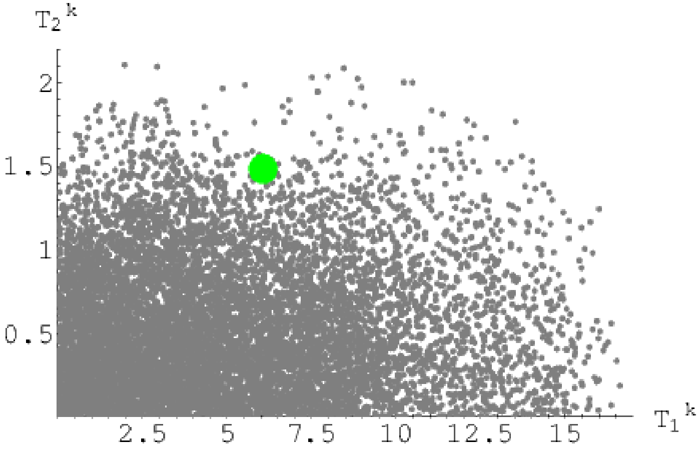

- **Choose N random assignments to group 1 or 2**
- **For each feature *j* and permutation *k***
  - **Compute a *test statistic* T_j^k, e.g. $d(\hat{\mu}_{1,j}^k, \hat{\mu}_{2,j}^k)$**
  - **Also compute T_j^o, the statistics for the observed data**

Step 2: Partial Test p-values

- **For each feature *j* and permutation *k***
  - **Compute a p-value using that feature's cumulative distribution:**

$$p(T_j^k) = \frac{1}{N}\sum_{l=1}^{N} H(T_j^l, T_j^k)$$

$$H(T_j^l, T_j^k) = \begin{cases} 1, & T_j^k \geq T_j^l \\ 0, & T_j^k < T_j^l \end{cases}$$

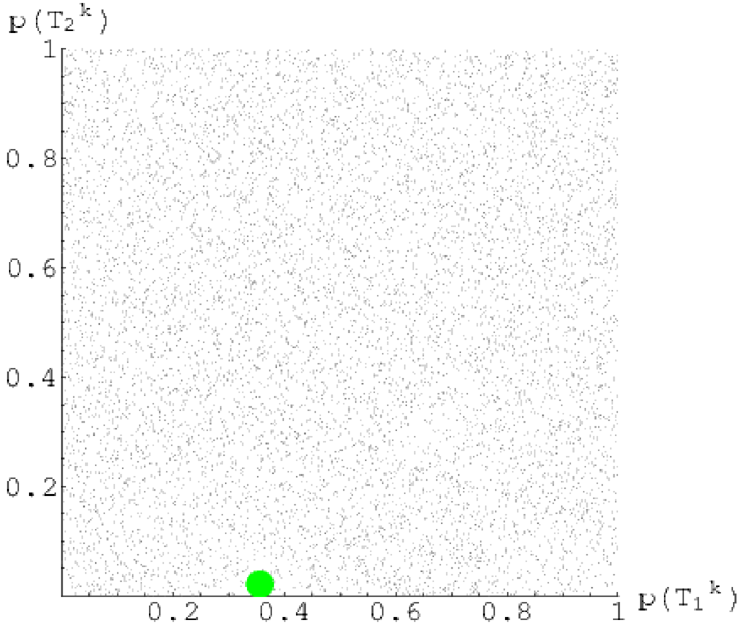

- **The marginal distributions are *uniform*, and invariant to scale**

Step 3: Combined Test

- **If the partial tests are**
  - **Significant for large values**
  - **Consistent**
  - **Marginally unbiased (unbiased regardless of whether or not other tests are true)**
- **And we choose a combining function $T'(p(\mathbf{T}^k))$ such that it is**
  - **Monotonically non-increasing in each p-value**
  - **Obtains its supremum T^* when any p-value is 0**
  - **Has finite critical values strictly smaller than T^***

Step 3: Combined Test

- Theorem: Then $T'(p(T^k))$ is an unbiased *global* test for equality of distributions (Pesarin 2001)

- What function should we use?
- One asymptotically equivalent to Hotelling's T^2 test (in linear case)
  - Uniformly most powerful, and affine invariant

Step 3: Combined Test (2-sided)

- With signed distances, T_j^k is significant for large *and* small values
- Map p-values for each feature to a standard normal distribution
  - $U_j^k = \Phi^{-1}(p(T_j^k) - \frac{1}{2\mathrm{N}})$, $\Phi =$ Gaussian c.d.f.
- Compute samp. covariance $\Sigma_U = \frac{1}{N} U^T U$
  - Full rank even for small samples: *N* is large
- Then $T'^k = (U^k)^T \hat{\Sigma}_U^{-1} U^k$

Step 3: Combined Test (1-sided)

- **Unsigned distances:**
  - Use positive half of c.d.f.

$$U_j^k = \Phi^{-1}\left(1-\frac{1}{2}\left(p(T_j^k)-\frac{1}{2N}\right)\right)$$

  - Assume U^k symmetric
  - Then covariance is zero when $j_1 \neq j_2 \rightarrow \Sigma_U = I$
  - So $T'^k = (U^k)^{\mathrm{T}} \cdot (U^k)$

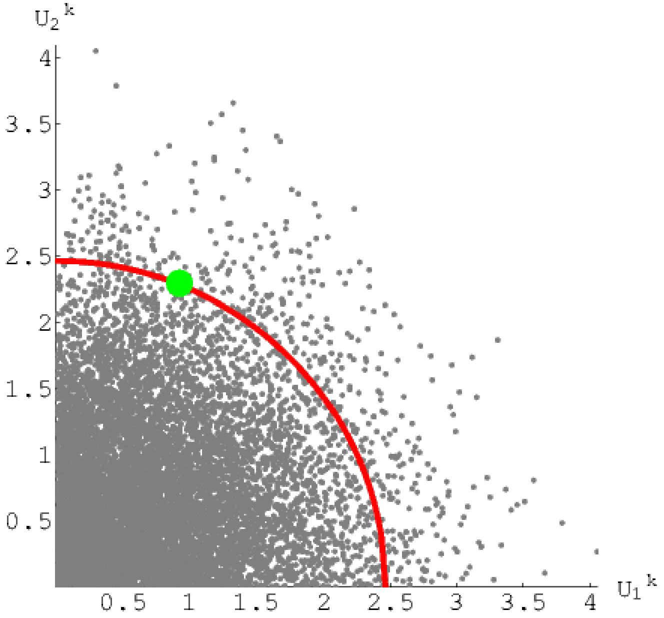

- **Use empirical distribution of T'^k to select a critical value**

Acceptance Region

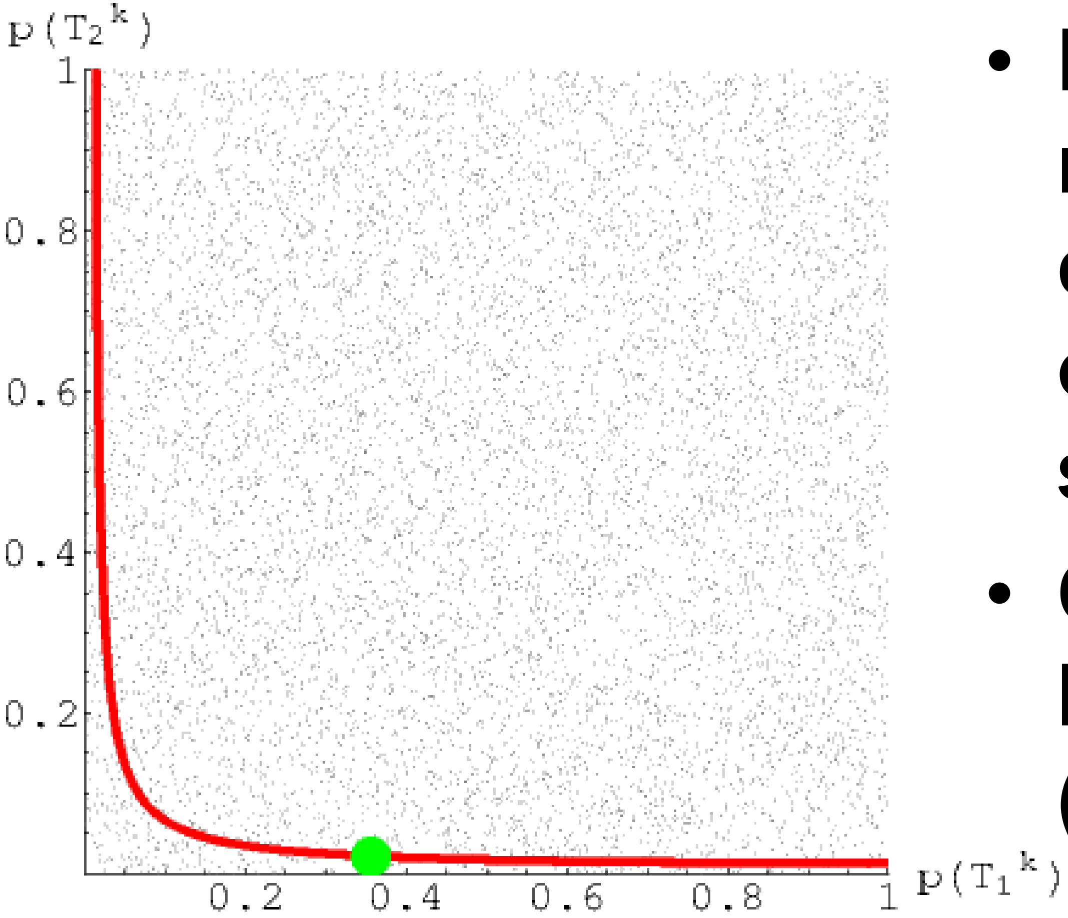

- Map critical region via c.d.f. to original space
- Contains both axes (p-value = 0)

Our Approach In Pictures

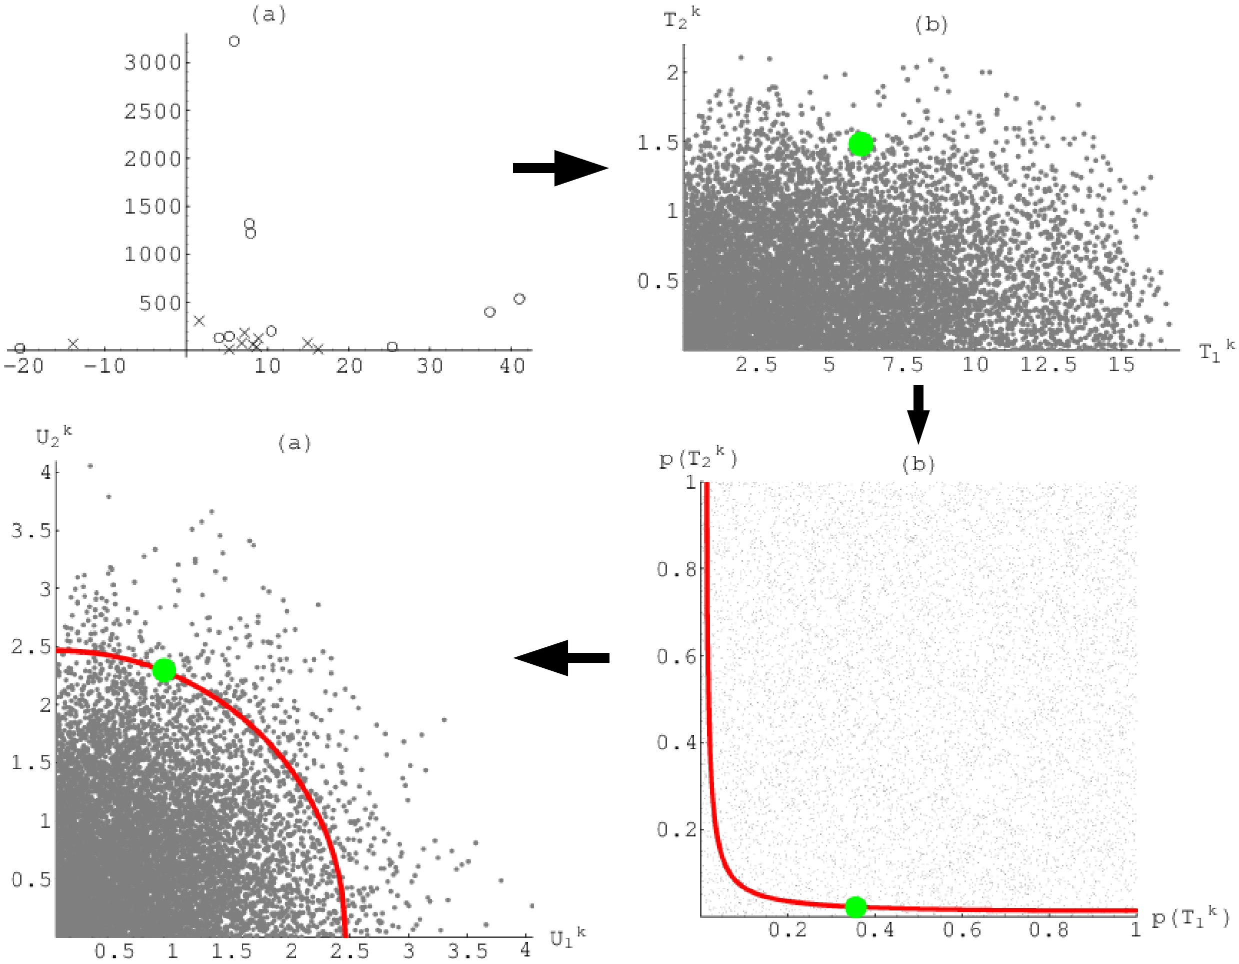

Application: Twin Ventricles

- MRI data of lateral ventricles from twin pairs
  - MZ - Healthy monozygotic: 9 pairs
  - DS - Monozygotic and discordant for schizophrenia: 9 pairs
  - DZ - Healthy dizygotic: 10 pairs
  - NR - Healthy non-related pairs: 10 pairs drawn from other healthy subjects

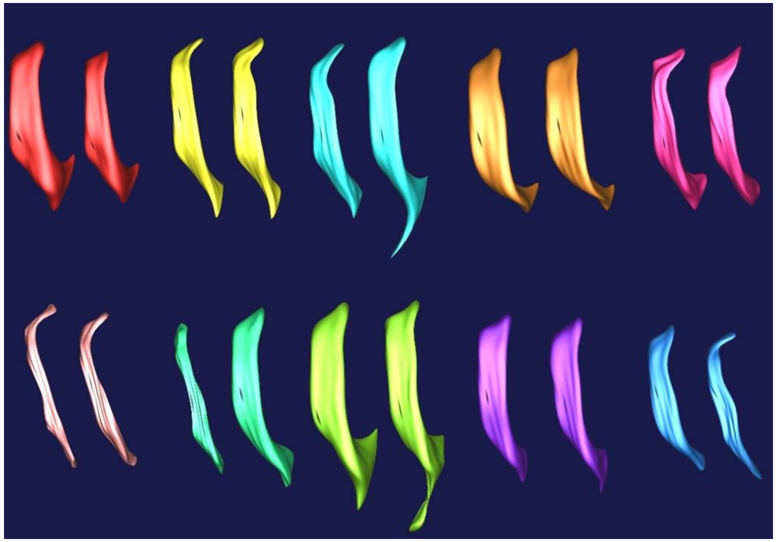

Global Results

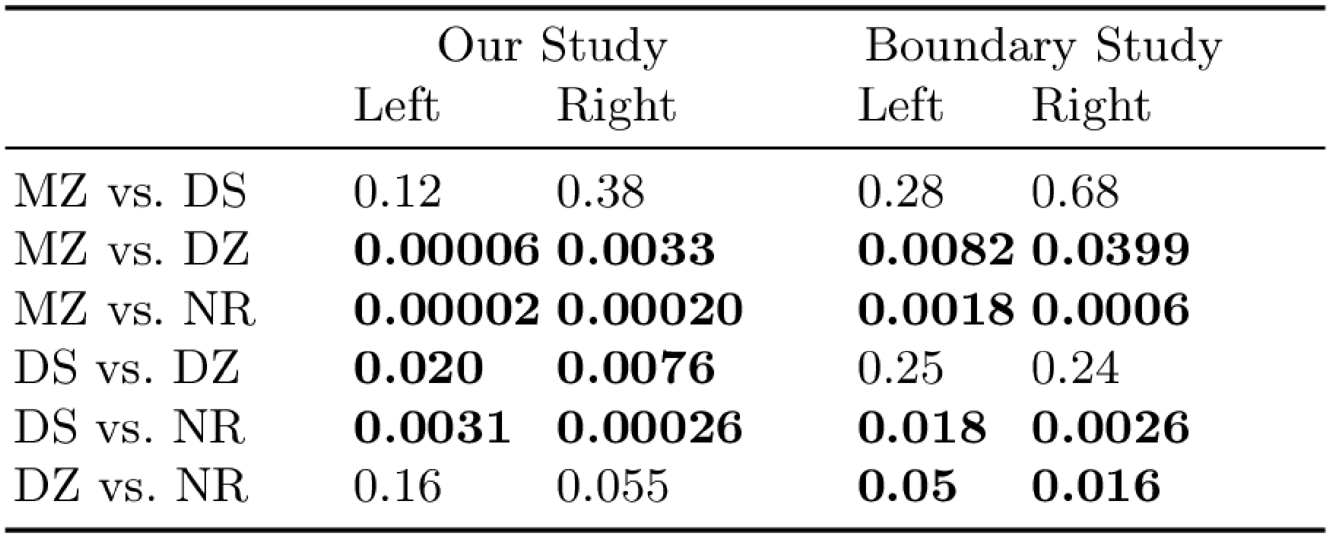

	Our Study		Boundary Study	
	Left	Right	Left	Right
MZ vs. DS	0.12	0.38	0.28	0.68
MZ vs. DZ	**0.00006**	**0.0033**	**0.0082**	**0.0399**
MZ vs. NR	**0.00002**	**0.00020**	**0.0018**	**0.0006**
DS vs. DZ	**0.020**	**0.0076**	0.25	0.24
DS vs. NR	**0.0031**	**0.00026**	**0.018**	**0.0026**
DZ vs. NR	0.16	0.055	**0.05**	**0.016**

- Comparison of our results with an earlier study on the PDMs (Styner et al. 2002)
  - Tests significant at 0.05 level in bold

Local Tests

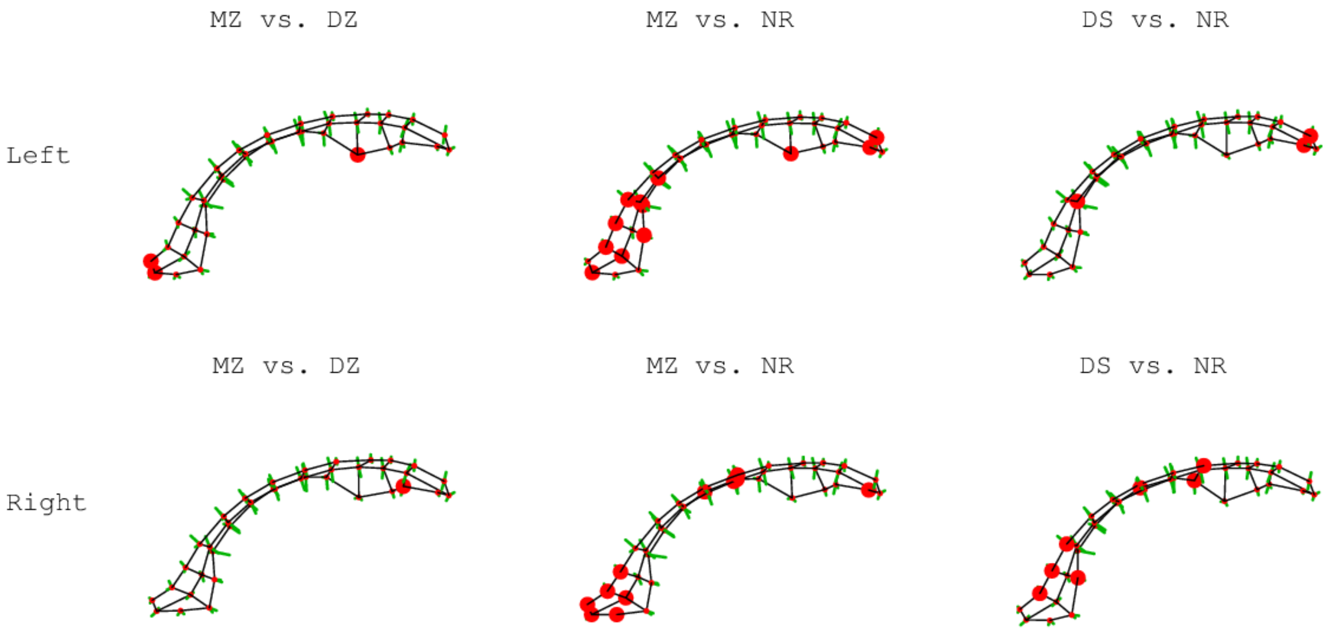

- **Local tests (*M* = 6 partial tests per atom, correction for multiple tests applied across atoms)**

Conclusion: Contributions

- First continuous 3D medial model that supports branching
- Multiscale model fitting procedure with explicit correspondence optimization
- Multivariate nonlinear hypothesis testing procedure for shape models

Questions?