"Explaining" a shape by estimating its generating skeleton

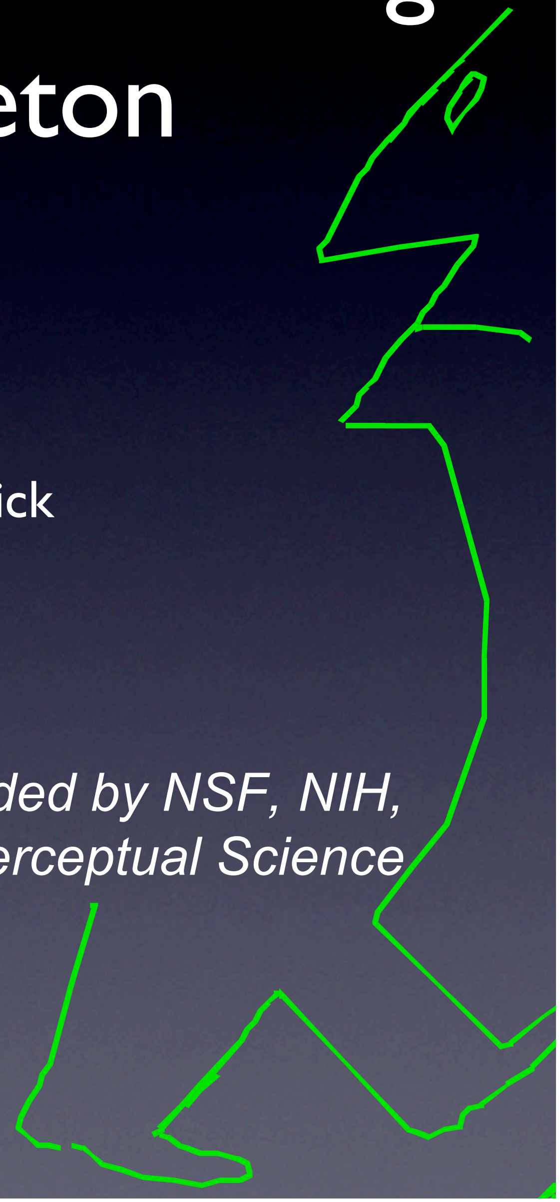

Jacob Feldman
Rutgers University - New Brunswick
USA

Joint work with Manish Singh

Funded by NSF, NIH,
Rutgers IGERT in Perceptual Science

Motivation

- We seek an effective *part-based* representation for shape
- Skeletal and medial axis representations are appealing…

(apologies to Marr & Nishihara, 1978)

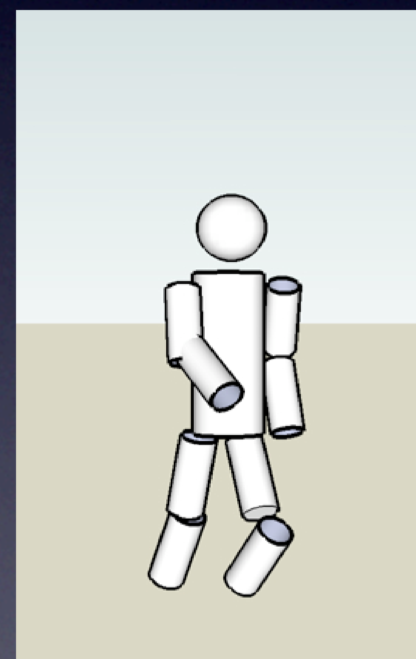

“Geon-guy” (apologies to Biederman, 1987)

...but problematic

Medial axis computation often gives counterintuitive results (forking)

...and is very sensitive to noise on the contour

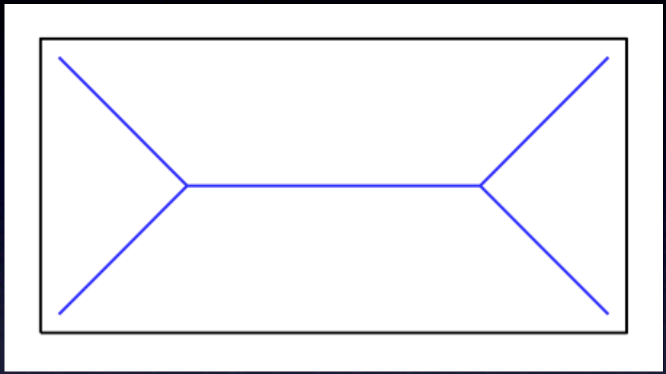

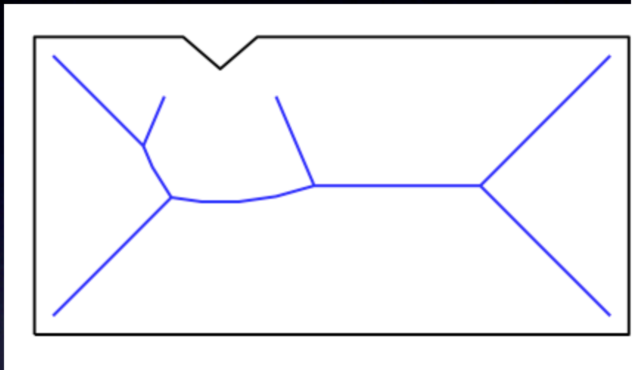

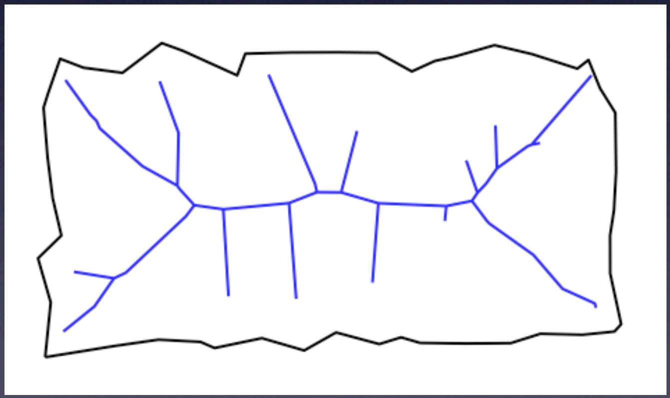

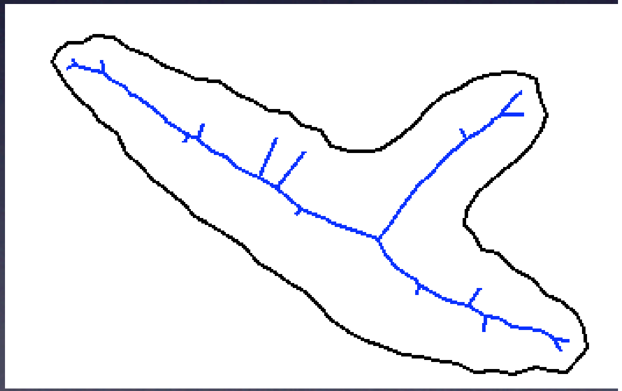

These problems ruin what would otherwise be an isomorphism between axes and parts

A different approach

- Computing medial axis representations is usually regarded as a “geometry problem”
- We view it as a *probabilistic inference* problem
- The goal is to estimate the shape’s *generative skeleton*—the skeleton from which it “grew”.

Bayesian estimation of the shape skeleton

- Define a prior on skeletons $p(\text{skel})$
- Define a *generative model* for shape given a skeleton, which defines the likelihood $p(\text{shape}|\text{skel})$
- Then we simply maximize the posterior
  $$p(\text{skel}|\text{shape}) \propto p(\text{skel})p(\text{shape}|\text{skel})$$
- Equivalent to minimizing the description length
  $$-\log p(\text{skel}|\text{shape}) \propto -\log p(\text{skel}) + -\log p(\text{shape}|\text{skel})$$
  $= \text{DL}(\text{skel}) + \text{DL}(\text{shape given skel})$ in MDL-speak

Forward model / generative model / likelihood function

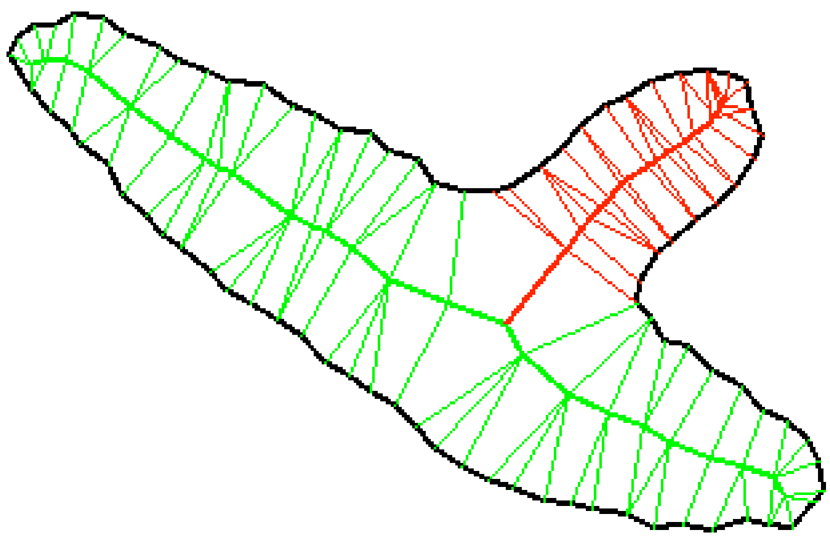

A skeleton…

…sprouts "ribs" in random directi(
(centered on normal)

…of random lengths

…whose endpoints join to become the shape.

Inverse inference / estimation / posterior

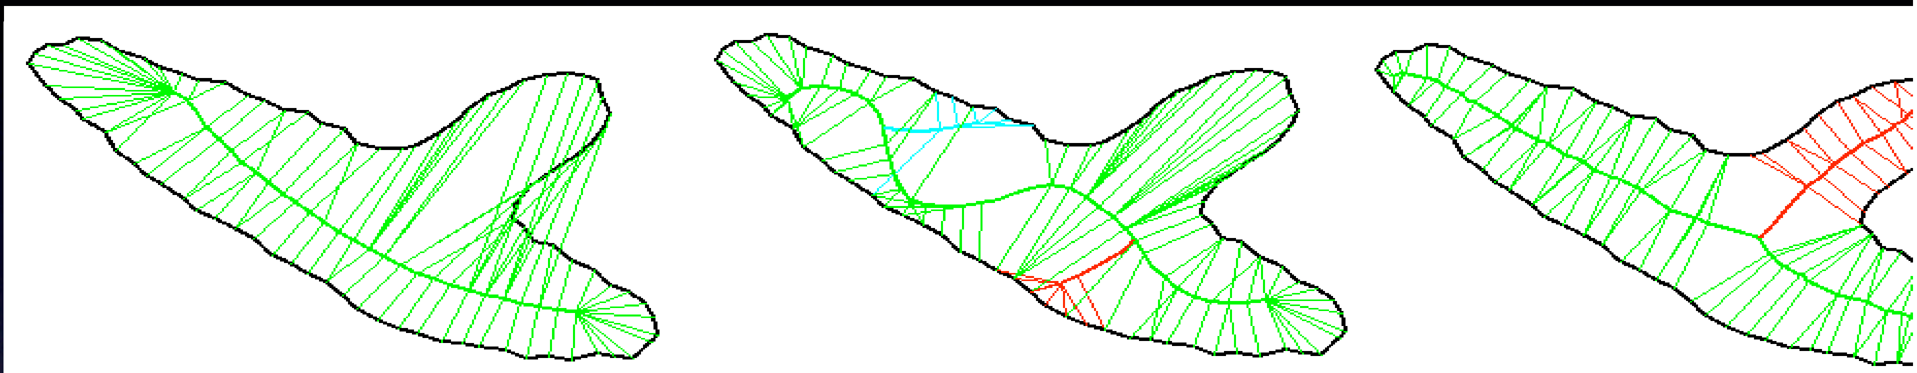

Too simple... *Too complex...* *Just right!*

- Goal: Over all skeletons, find the one with maximum posterior *p*(shape|skel)*p*(skel), called the maximum a posteriori or MAP skeleton

- This skeleton best "explains" the shape.

The prior on skeletons

A hierarchical generalization of our prior on contours $p(C)$:

A skeleton with n axes has prior $p(\text{skel})$ =

$$p_a p(C_1) \cdot p_a p(C_2) \cdot \ldots \cdot p_a p(C_n) = p_a^n \prod_i p(C)$$

and complexity $-\log p(\text{skel})$

$$n \log p_a + \sum_i \log p(C_i)$$

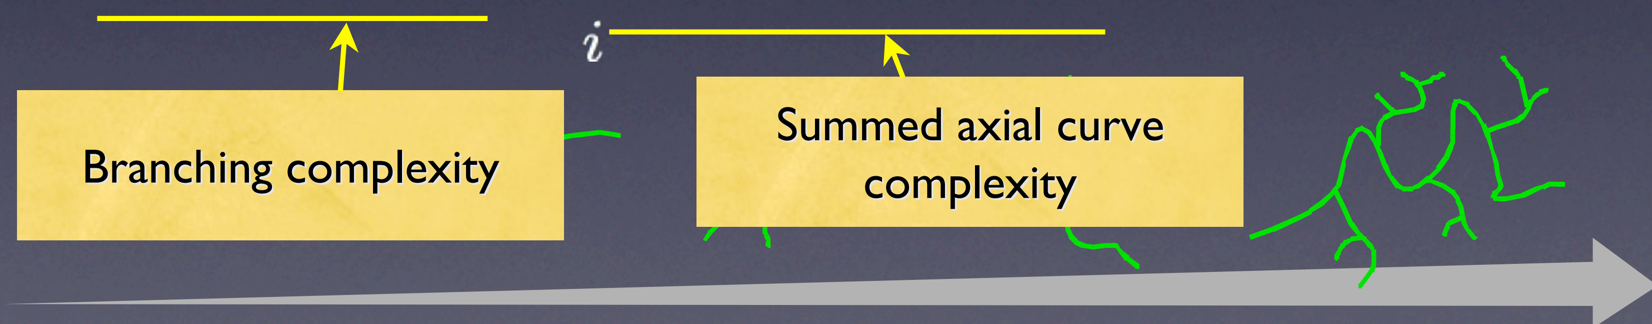

Increasing skeleton complexity / decreasing prior

Generative (likelihood) model for shapes

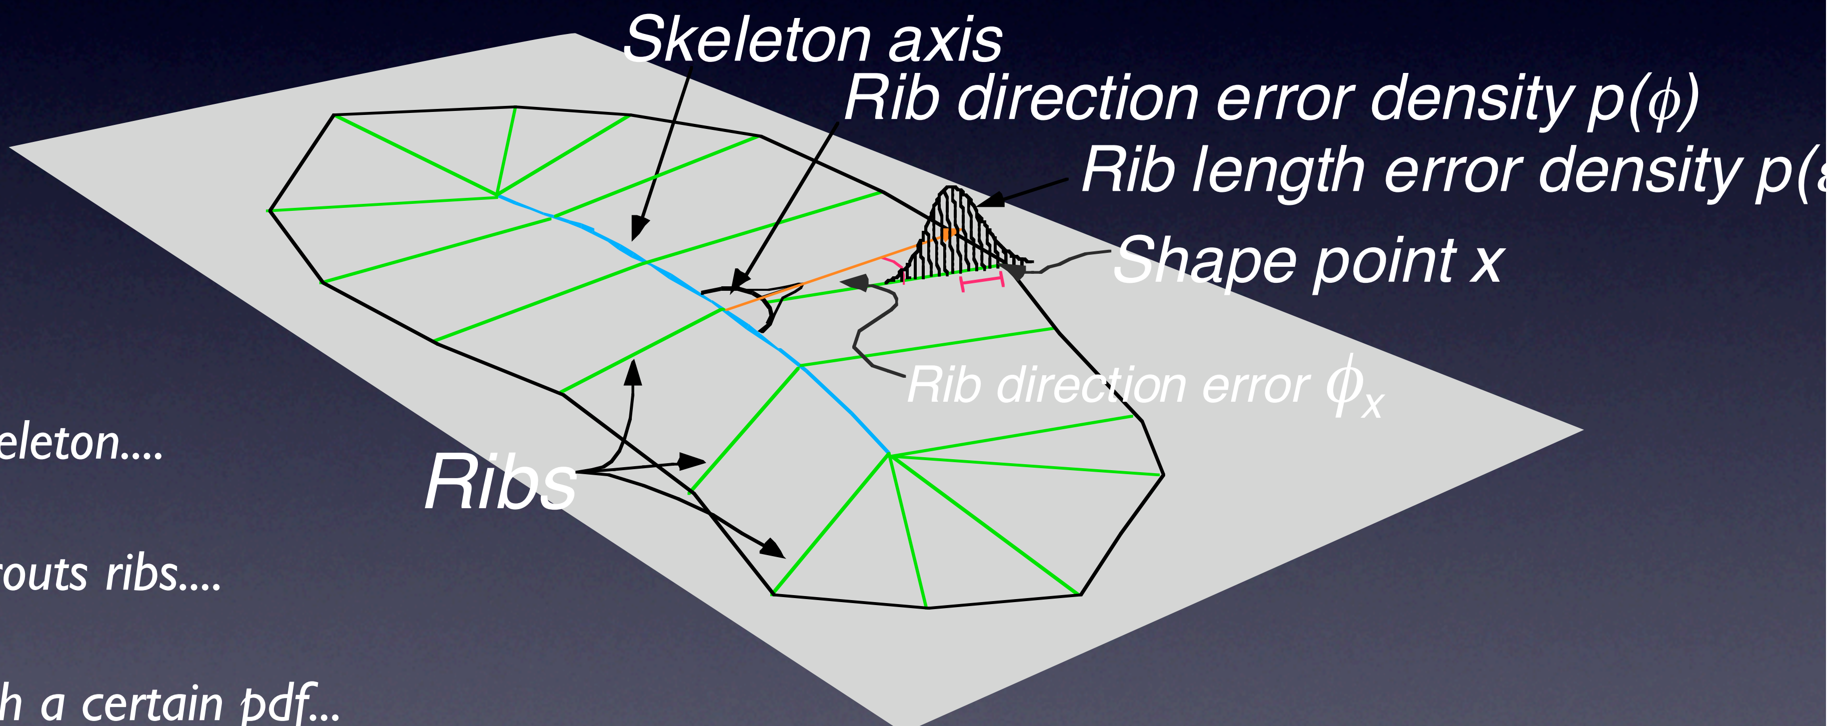

A skeleton....

...sprouts ribs....

...with a certain pdf...

...whose endpoints join to form the shape.

Bayesian (posterior ratio) criterion for the "significance" of an axis

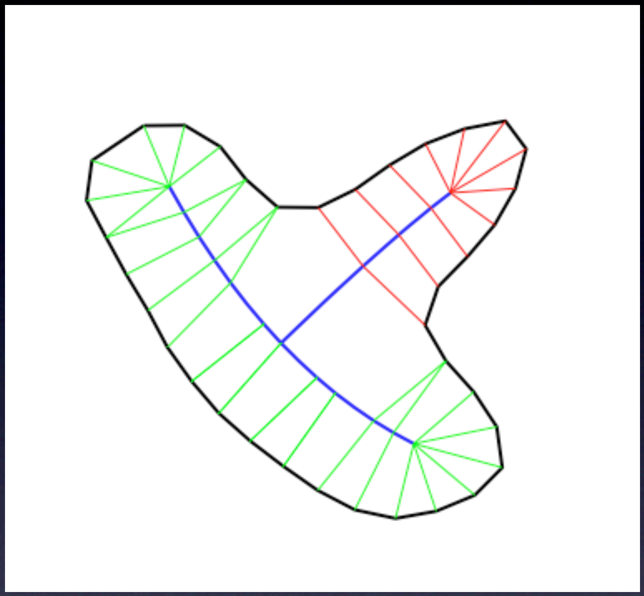

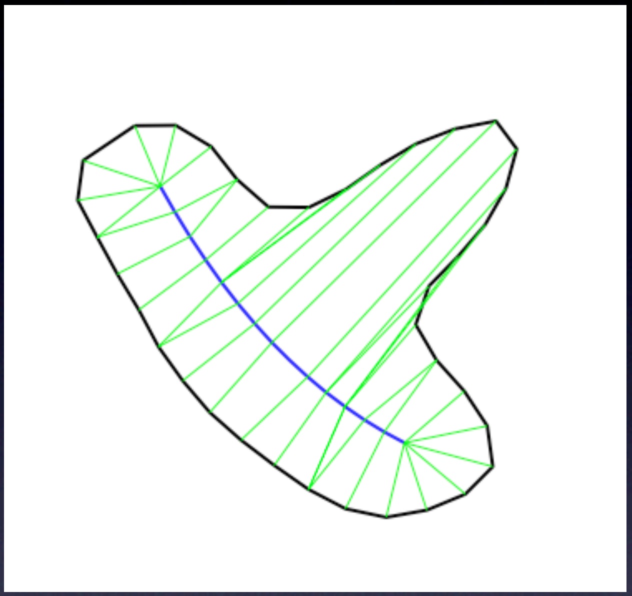

Compare posteriors *with the axis* vs. *without the axis*

If $\dfrac{p(\text{skel+axis}|\text{shape})}{p(\text{skel}|\text{shape})} > 1$, *keep the axis.*

Estimating the skeleton

the fine print

- Initialize skeleton
  (we use a conventional Voronoi-based medial axis)
- Prune "nonsignificant" axes using posterior ratio rule
- Parameterize skeleton estimate
  (we use a piecewise cubic spline approximation)
- Begin gradient descent in skeleton parameter space
  (we use a home-brewed variant of Expectation-Maximization)
- Many other details I'm not mentioning

Results

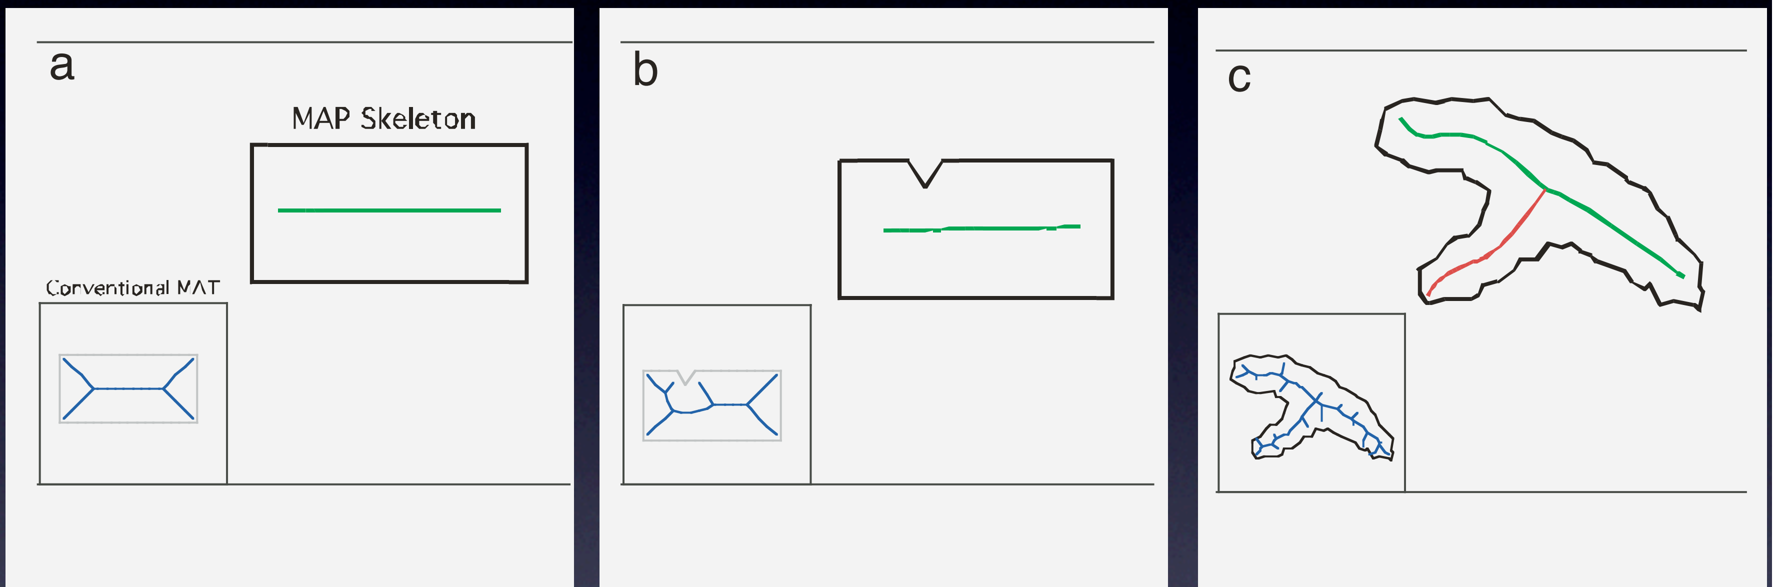

Robustness against contour noise

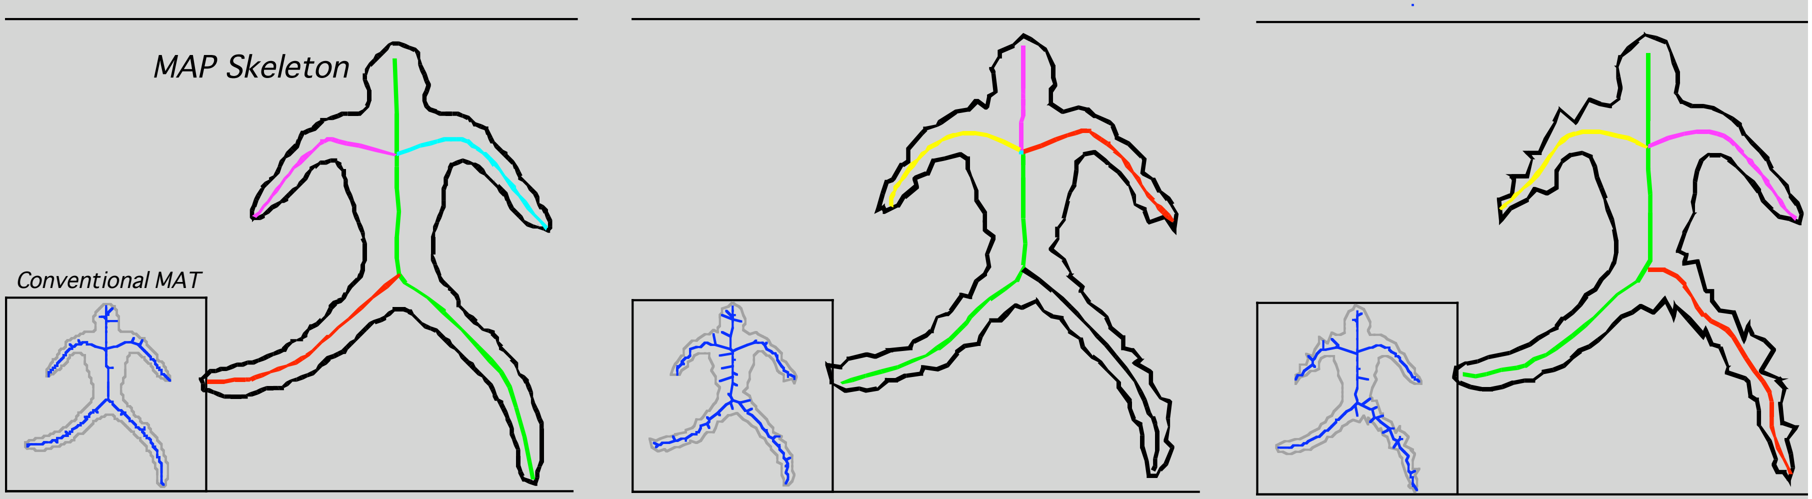

"Dude7"

+ noise

+ noise on one arm and leg

Dog, cow

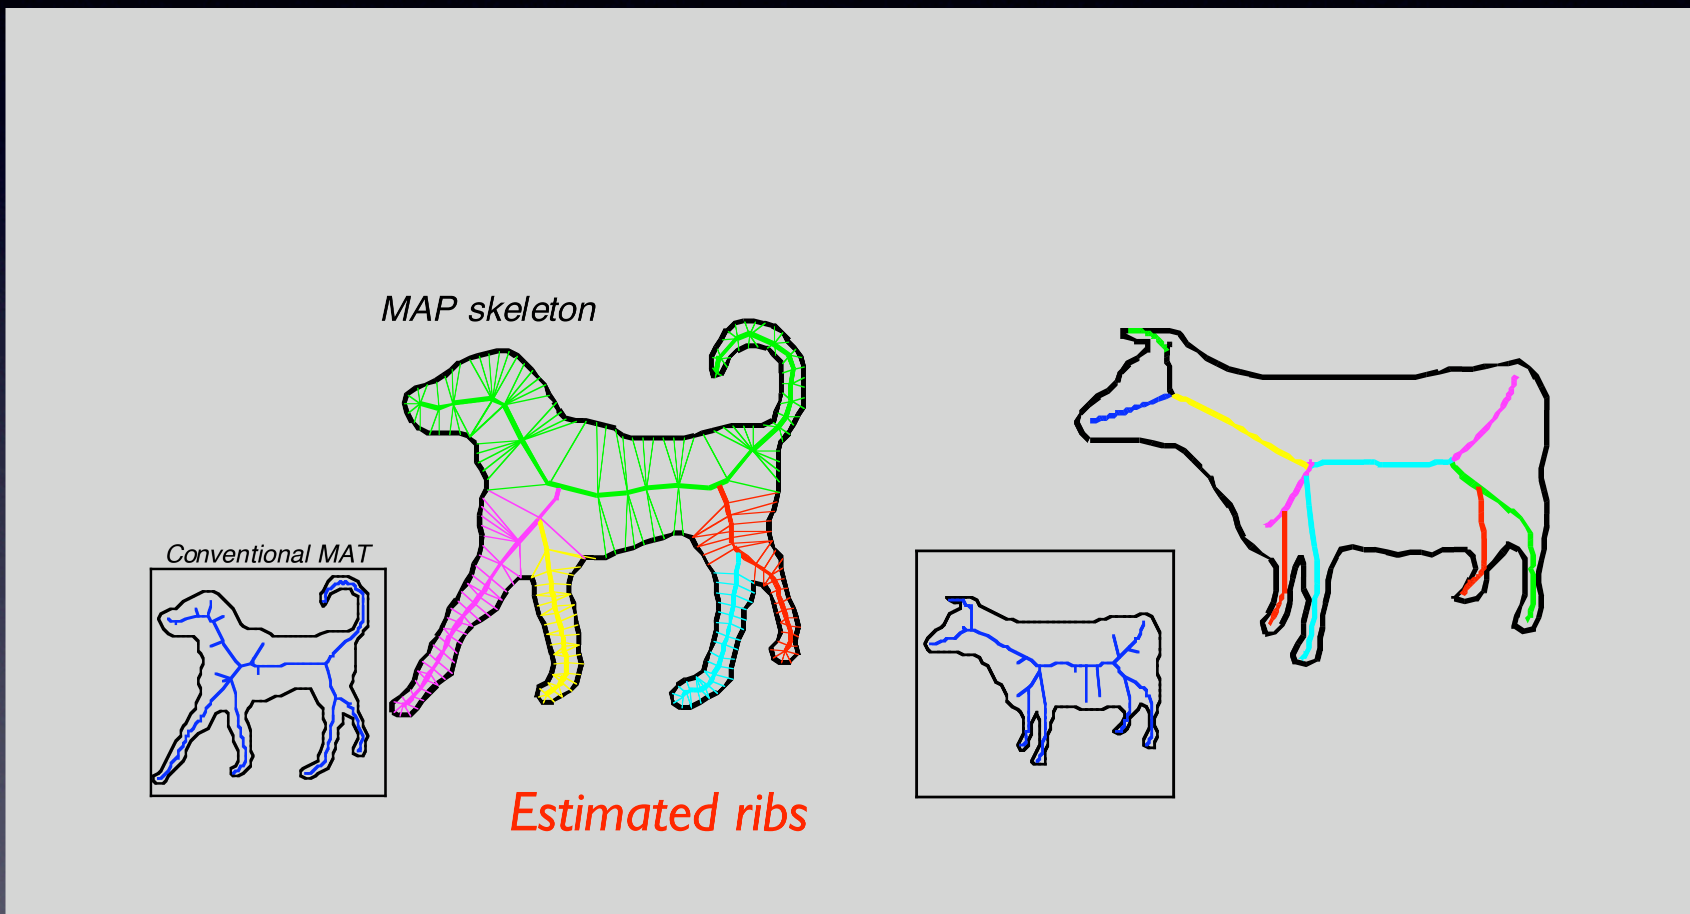

Camel, elephant

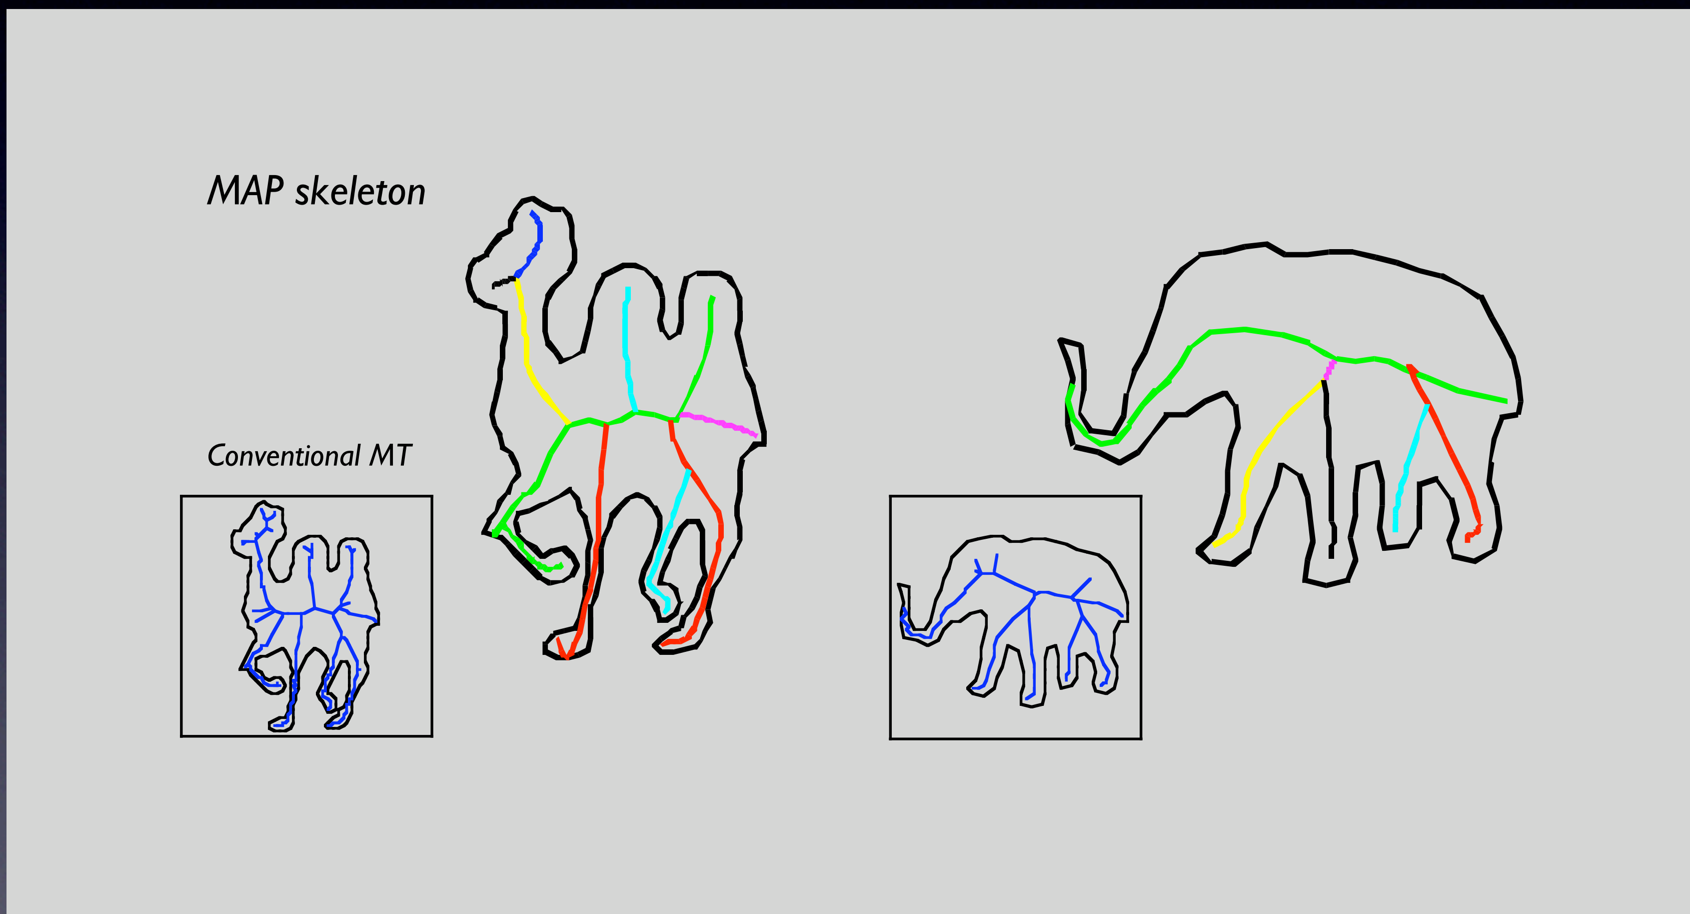

Donkey, fish

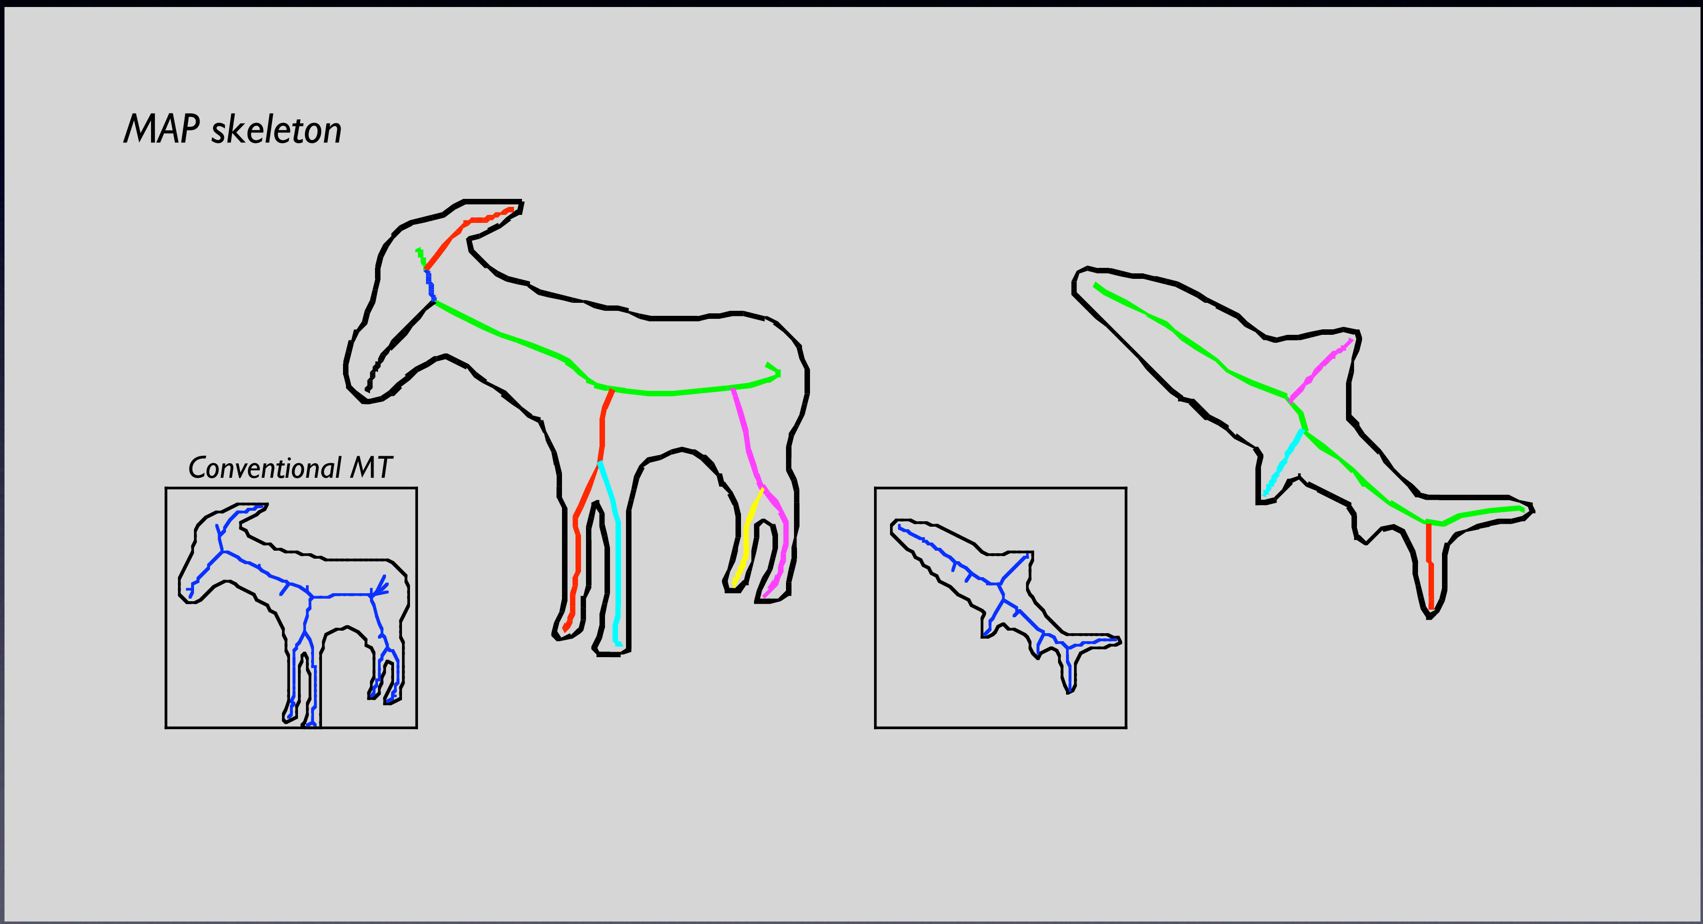

Skeletons and parts

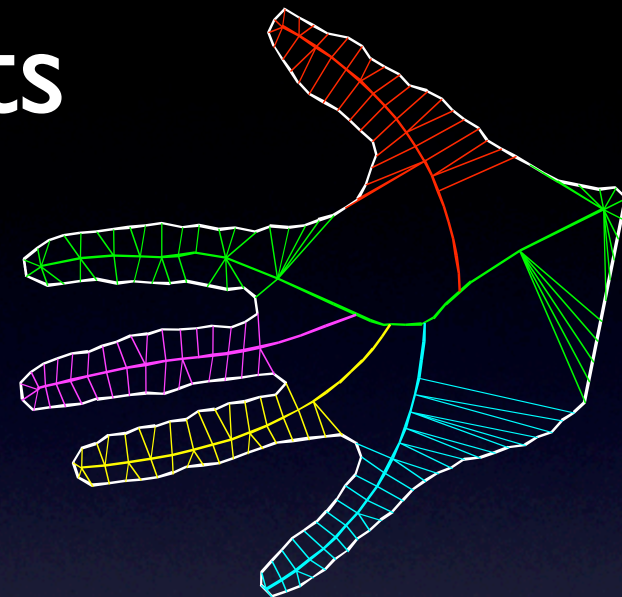

• Distinct axes in the estimated skeleton "own" (explain) points on the contour

• Many known principles of part decomposition approximately "fall out" of MAP skeleton estimation

  • minima rule (Hoffman & Richards, 1984)
  • short cuts (Singh, Seyranian, & Hoffman, 1999)
  • maximization of convexity (Rosin, 2000)

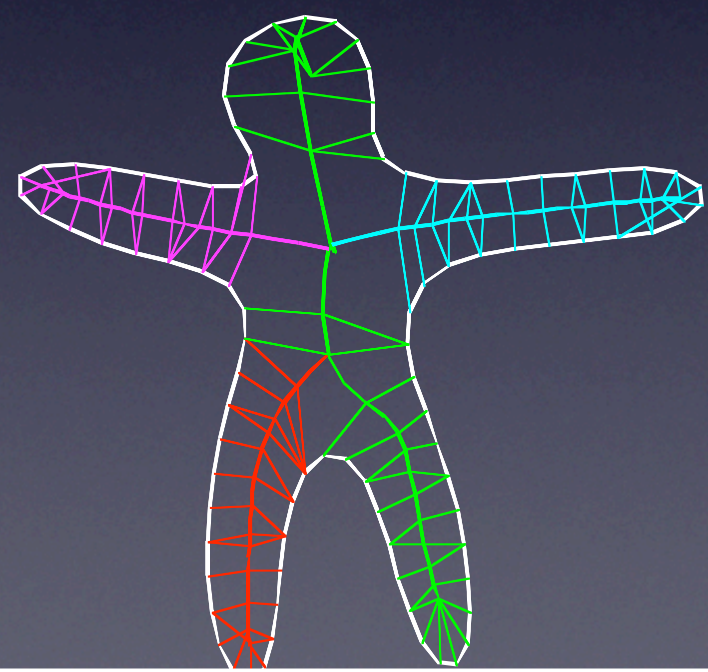

• Unifies theory of part-boundaries with theory of part-cuts

Extending “ownership” to interior regions

- For each interior point *x*, determine axial ownership by

$$p(x \text{ owned by } A_i) = p(A_i)p(x|A_i)$$

$$\propto \frac{1}{1+\text{depth}(A_i)} f(d(x,$$

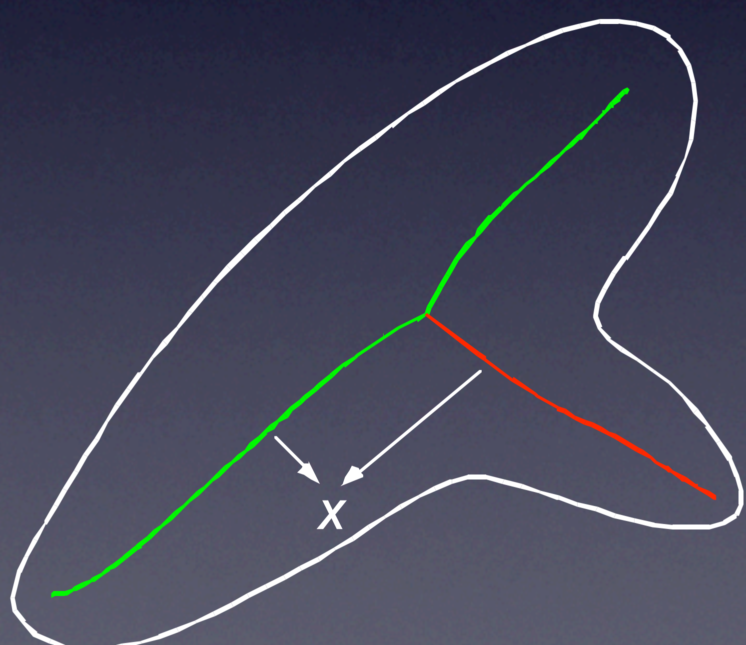

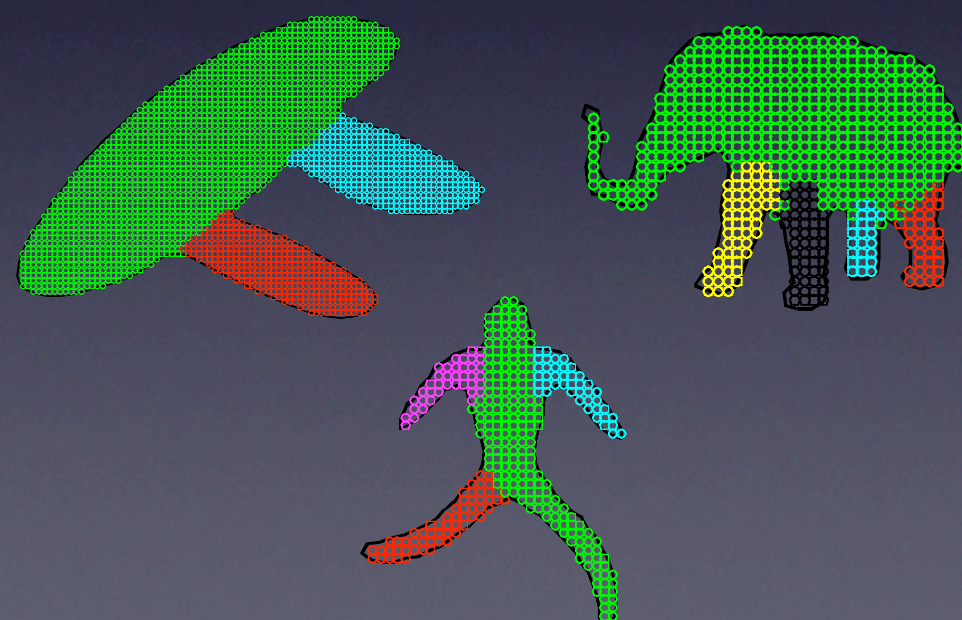

with Erica Briscoe

Shape similarity

- Given shapes *x* and *y*, what is sim(*x*,*y*)?

- Conceptualization: shapes are similar to the degree that they see to *share common generative origins*

- Operationalization: similarity is given by the "cross-likelihoods"

$$\text{sim}(x,y) \sim [p(\text{shape}_x|\text{skel}_y) + p(\text{shape}_y|\text{skel}_x)]/2$$

- Experiments: Similarity ratings on all pairs of 25 shapes, various types: "metric" differences, part-structure differences, non-axial shapes…

Similarity results

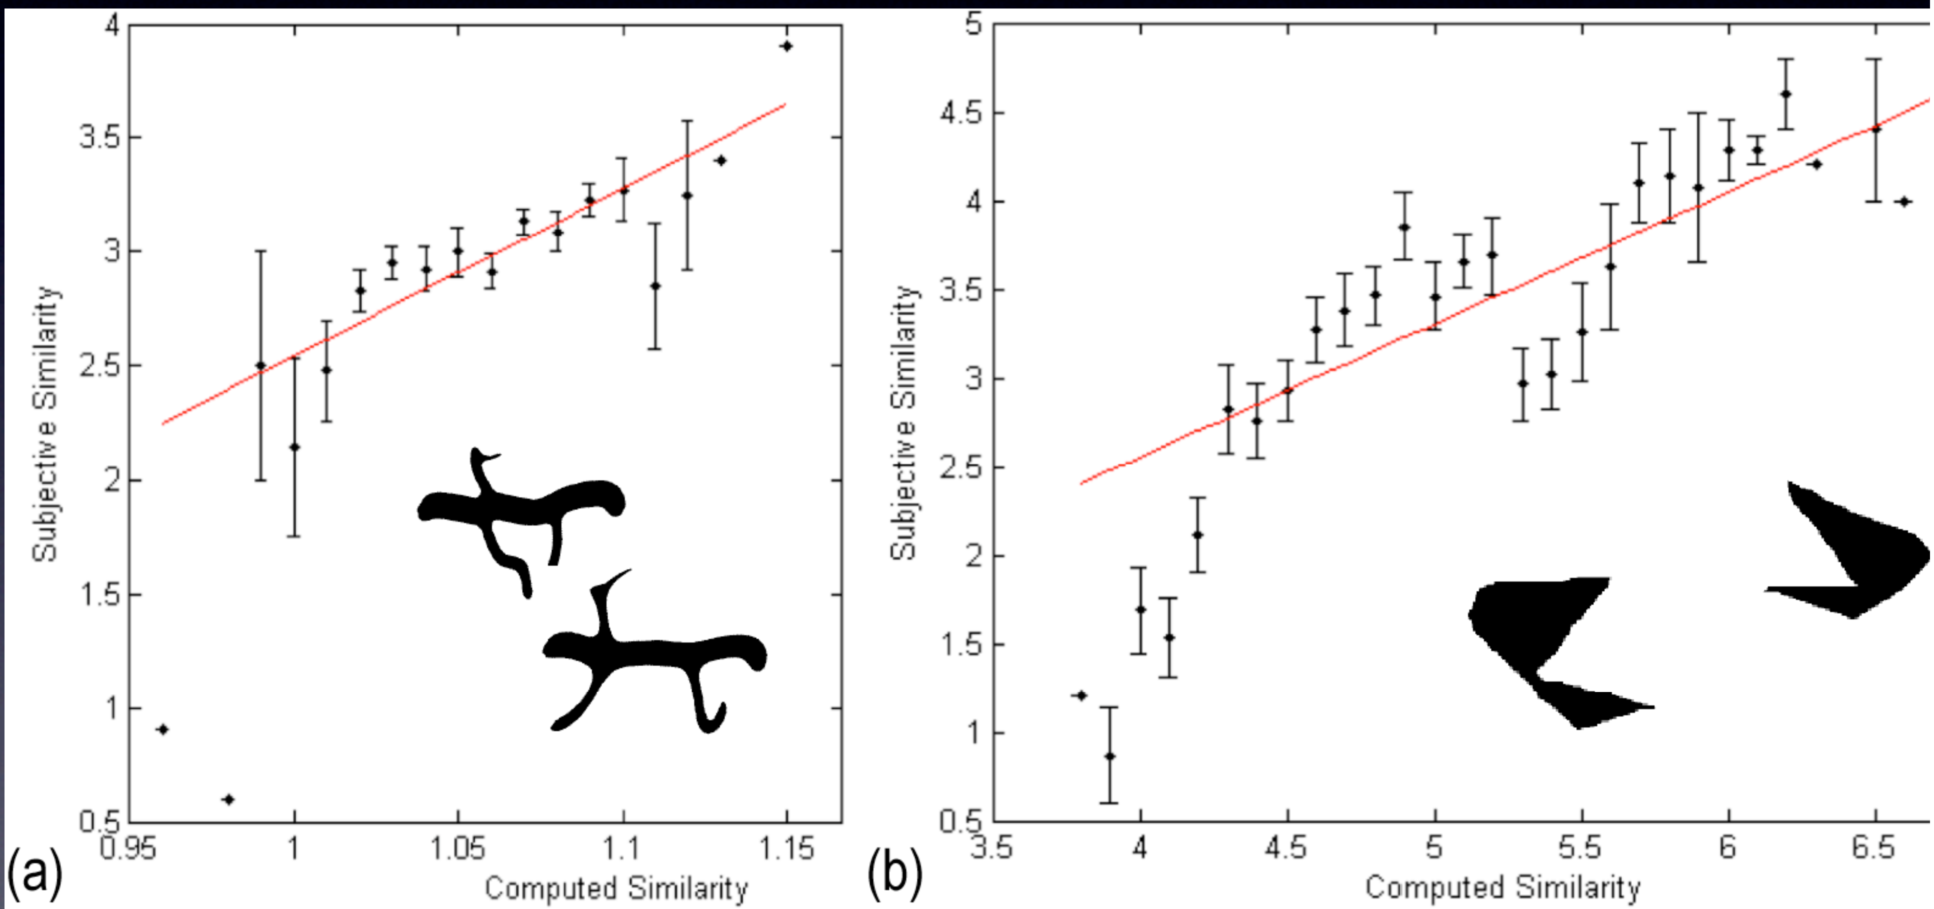

3D shape from the skeleton

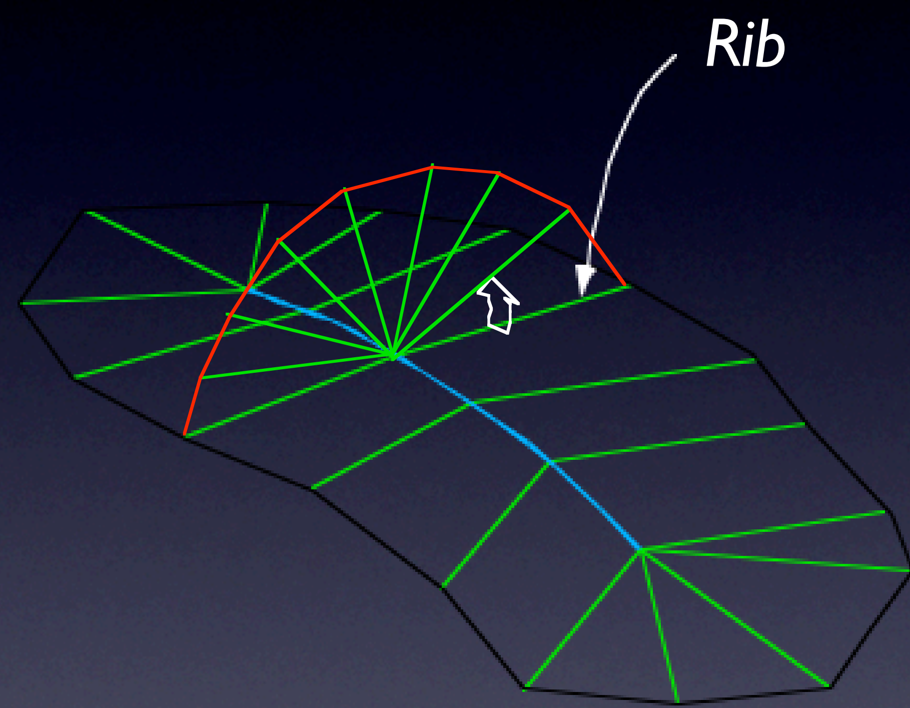

- The generative model can be easily extended to 3D
- "Inflate" the shape to produce a complete 3D model from the 2D skeleton

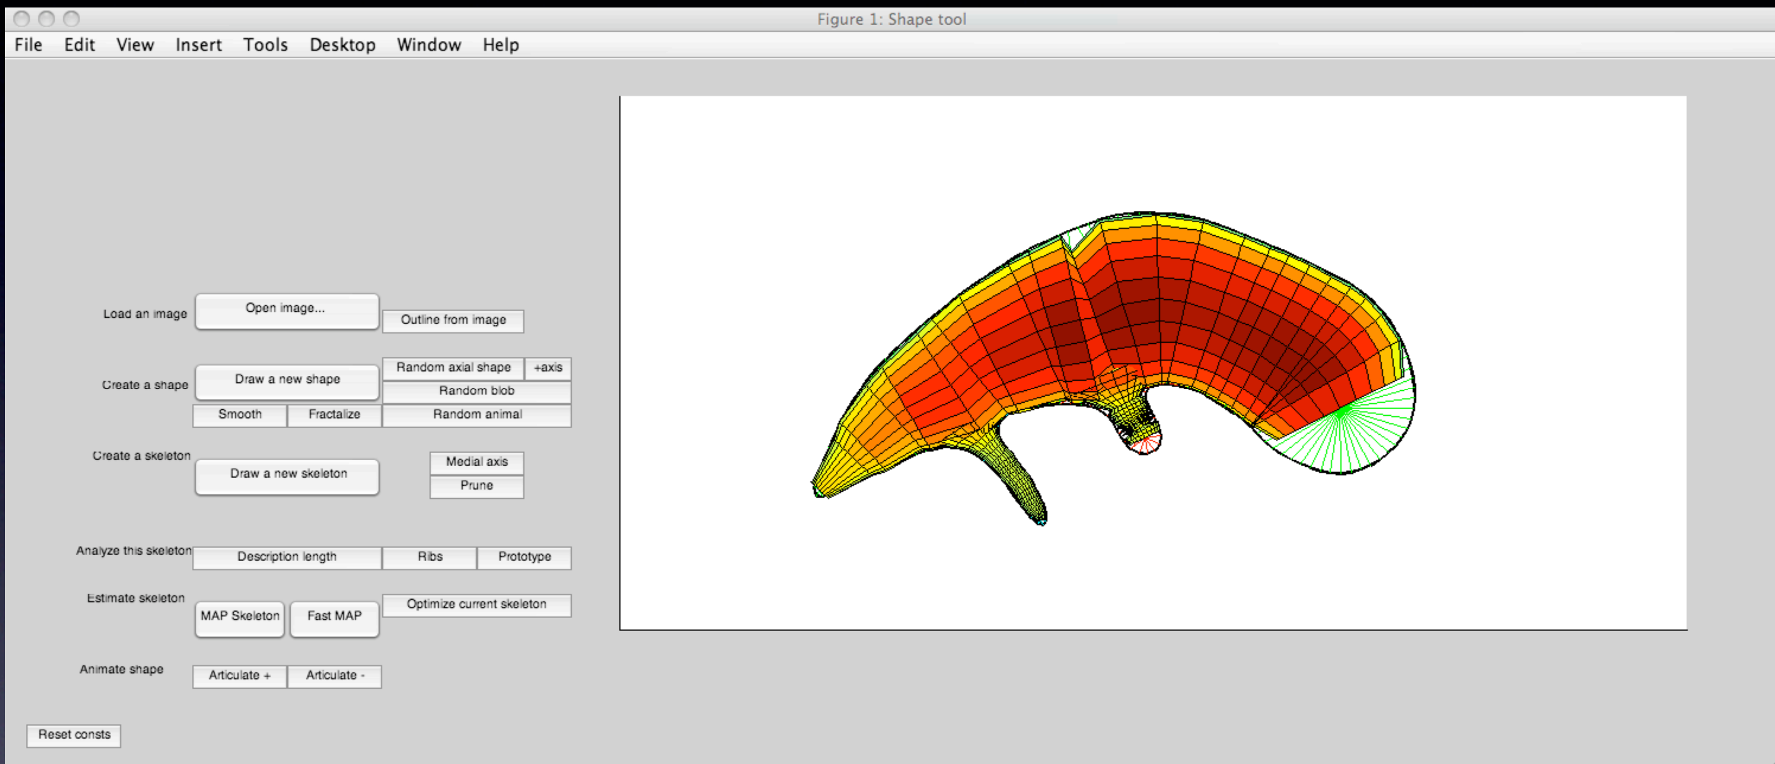

Figure 1: Shape tool
File
Edit
View
Insert
Tools
Desktop
Window
Help
Load an image
Open image...
Outline from image
Create a shape
Draw a new shape
Random axial shape
+axis
Random blob
Smooth
Fractalize
Random animal
Create a skeleton
Draw a new skeleton
Medial axis
Prune
Analyze this skeleton
Description length
Ribs
Prototype
Estimate skeleton
MAP Skeleton
Fast MAP
Optimize current skeleton
Animate shape
Articulate +
Articulate -
Reset consts

Statistics of natural shapes

- "Naturalizing the prior"

  In place of the naïve complexity prior, draw prior densities from *statistics of natural shapes*

- The goal is to find "meaningful" shape parameters and tune the representation to the environment

Sample domains: animals and leaves

with John Wilder

We gathered skeleton statistics from two shape databases…

Animals

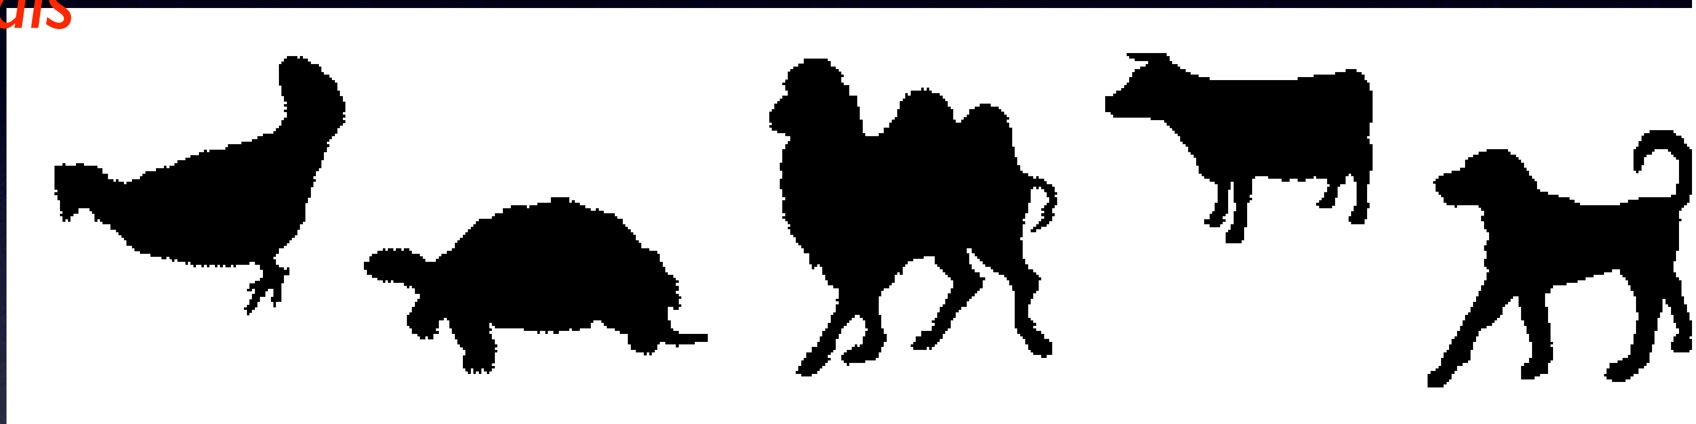

Leaves

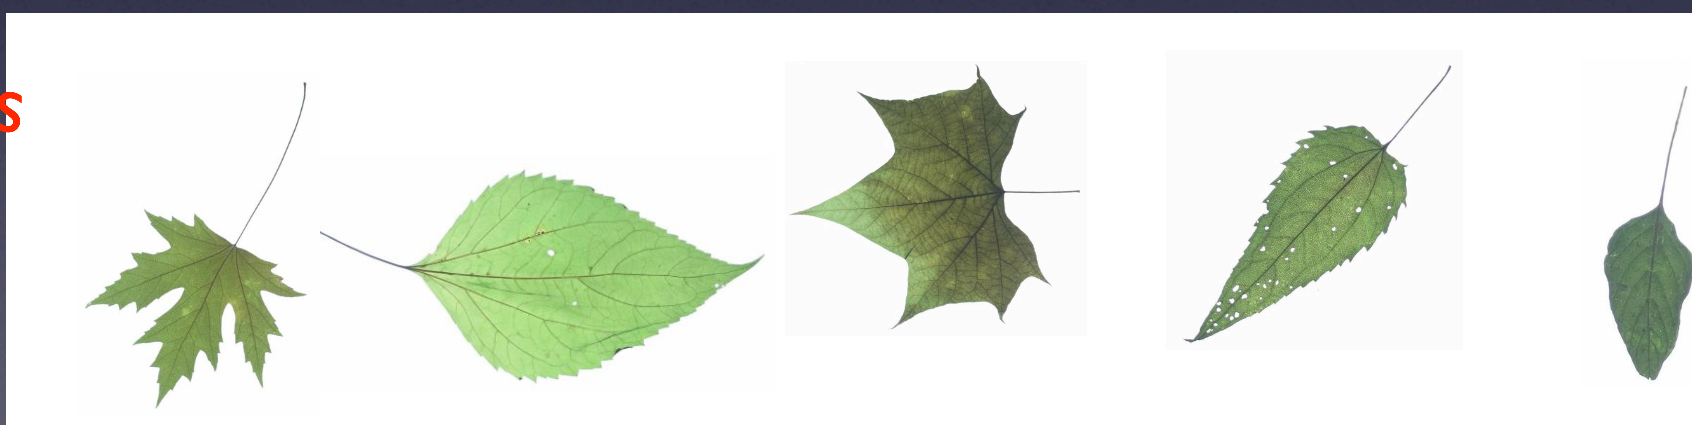

Empirical distributions of skeleton parameters

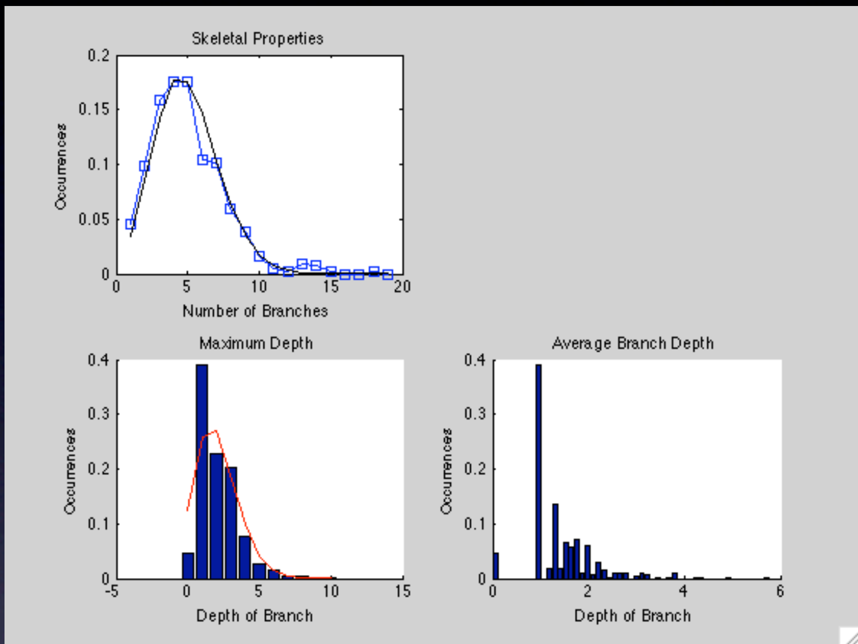

Animals

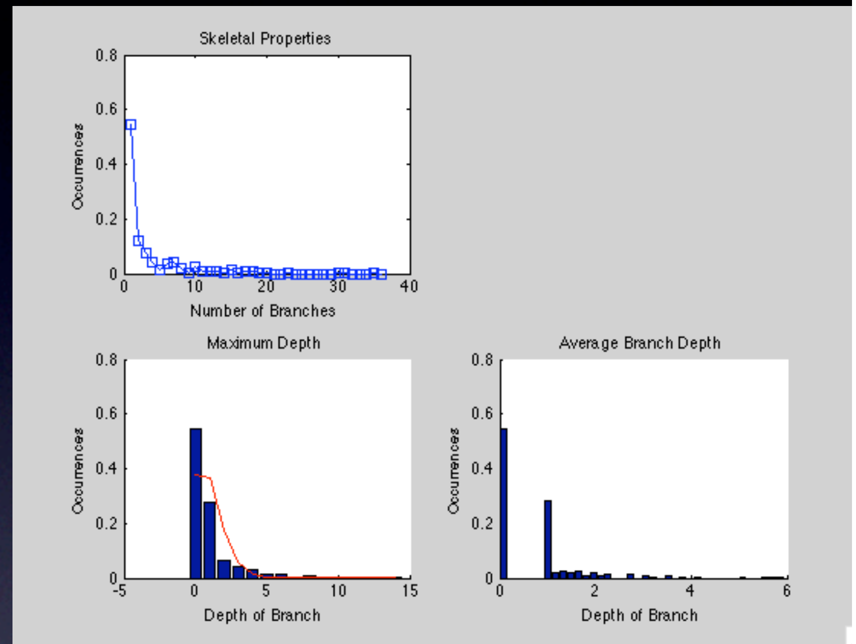

Leaves

Objectively quantifying "natural kinds"

Summary and conclusions

- Shape is poorly understood, even in the 2D case
  - skeletons are important
- The generating skeleton as a unifying conceptualization of shape

  Principled theoretical framework based on the idea of "explaining" the shape

  Bayesian estimation of the MAP skeleton yields part decomposition, similarity measur structure, etc.
- Many other extensions just beginning to be pursued

The end