Uncertainties in tomographic reconstructions based on Bayesian geometric models

Ken Hanson, CCS-2, LANL
Greg Cunningham, DX-3, LANL

Los Alamos NATIONAL LABORATORY

Presentation available under
http://home.lanl.gov/kmh/

Overview

- Bayesian approach to model-based analysis
- Example - tomographic reconstruction from two views
- Deformable geometric models
- Bayes Inference Engine - a radiographic modeling tool
- Maximum a posteriori (MAP) reconstruction
- Sampling from probability density functions
  - ▹ Markov Chain Monte Carlo (MCMC) technique
  - ▹ probabilistic interpretation of priors
- Estimation of uncertainty in reconstructed shape
  - ▹ Use of MCMC to sample posterior
  - ▹ Hard truth approach - probe model stiffness

Bayesian approach to model-based analysis

- Models
  - ▹ used to analyze physical world
  - ▹ parameters inferred from data
- Bayesian analysis
  - ▹ uncertainties in parameters described by probability density functions (pdf)
  - ▹ prior knowledge may be incorporated
  - ▹ quantitatively and logically consistent methodology for making inferences
  - ▹ open ended approach
    - can incorporate new data
    - can extend models and choose between alternatives

Bayesian viewpoint

- Focus on probability distribution functions (pdf)
  - ▹ uncertainties in estimates more central than the estimates themselves
- Bayes law: $p(\mathbf{a}|\mathbf{d}) \sim p(\mathbf{a})\, p(\mathbf{d}|\mathbf{a})$
  - ▹ where **a** is parameter vector and **d** represents data
  - ▹ pdf before experiment, $p(\mathbf{a})$ (called *prior*)
  - ▹ modified by pdf describing experiments, $p(\mathbf{d}|\mathbf{a})$ (*likelihood*)
  - ▹ yields pdf summarizing what is known, $p(\mathbf{a}|\mathbf{d})$ (*posterior*)
- Experiment should provide decisive information
  - ▹ posterior much narrower than prior

Bayesian model building

- Steps in model building
  - ▹ choose how to model (represent) object
  - ▹ assign priors to parameters based on what is known beforehand
  - ▹ for given measurements, determine model with highest posterior probability (MAP)
  - ▹ assess uncertainties in model parameters
- Higher levels of inference
  - ▹ assess suitability of model to explain data
  - ▹ if necessary, try alternative models and decide among them

Example - tomographic reconstruction

- Problem - reconstruct object from two projections
  - ▹ 2 orthogonal, parallel projections (128 samples in each view)
  - ▹ Gaussian noise;
    rms-dev 5% of proj. max

Original object

Two orthogonal projections
with 5% rms noise

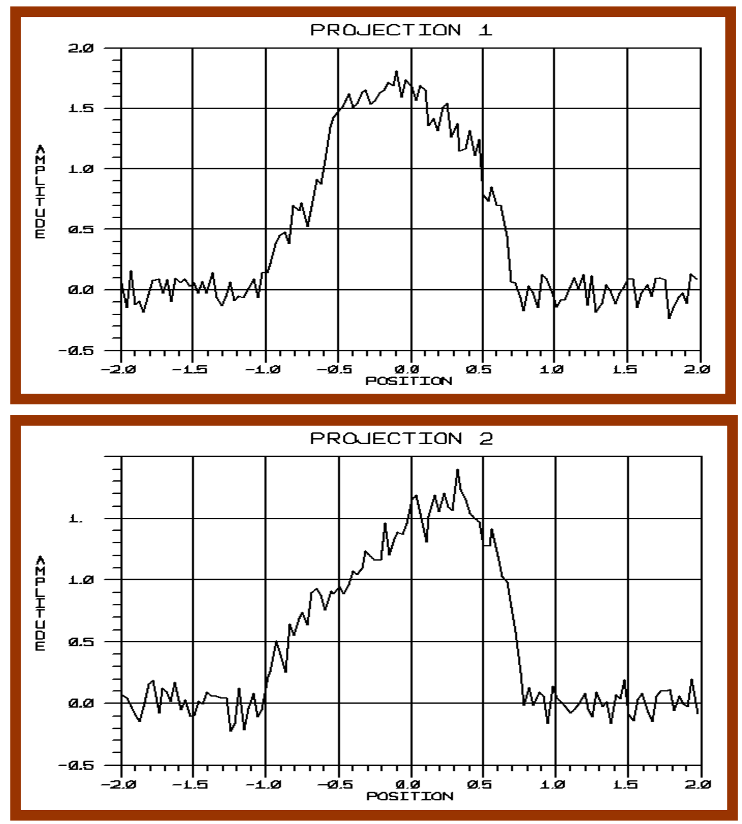

Prior information in reconstruction

- Assumptions about object
  - ▹ object density is uniform
  - ▹ abrupt change in density at edge
  - ▹ boundary is relatively smooth
- Object model
  - ▹ object boundary - deformable geometric model
    - relatively smooth
  - ▹ interior has uniform density (known)
  - ▹ exterior density is zero
  - ▹ only variables are those describing boundary

Deformable geometric models

- Natural to describe objects in terms of their boundaries
- In data analysis aim is to balance
  - ▹ internal energy ε: measure of deformation
  - ▹ external energy, e.g. χ^2: measure of mismatch to data
- Constrain smoothness based on curvature κ
  - ▹ deformation energy, e.g., $\varepsilon \sim \int \kappa^2 \, ds$, for curve
  - ▹ controls number of degrees of freedom of curve
- Analogy to elastic materials - rods, sheets

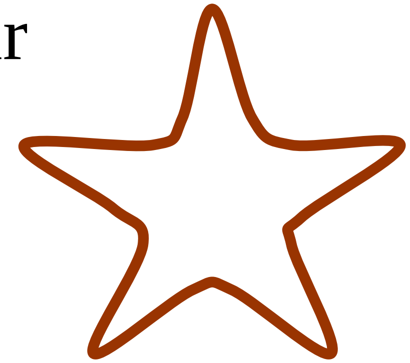

Tomographic reconstruction from two views

- Data consist of two orthogonal views
  - ▹ parallel projections, each containing 128 samples
  - ▹ Gaussian noise; rms-dev 5% of proj. max
- Object model
  - ▹ boundary is 50-sided polygon
  - ▹ smoothness achieved by prior on curvature κ
  - ▹ uniform (known) density inside boundary
- φ = - log posterior = $\frac{1}{2}\chi^2 + \frac{\alpha S}{(2\pi)^2}\oint \kappa^2 ds$,
  - ▹ where S is total perimeter,
  - ▹ χ^2 is sum of squares of residuals divided by noise variance

The Bayes Inference Engine

- Flexible modeling tool developed at LANL
  - ▹ object described as composite geometric and density model
  - ▹ measurement process (principally radiography)
- User interface via data-flow diagram
- Full interactivity with every aspect of model
- Provides
  - ▹ MAP estimate by optimization (gradient by ADICT)
  - ▹ samples of posterior by MCMC
  - ▹ uncertainty estimates

The Bayes Inference Engine

- BIE data-flow diagram to find MAP solution

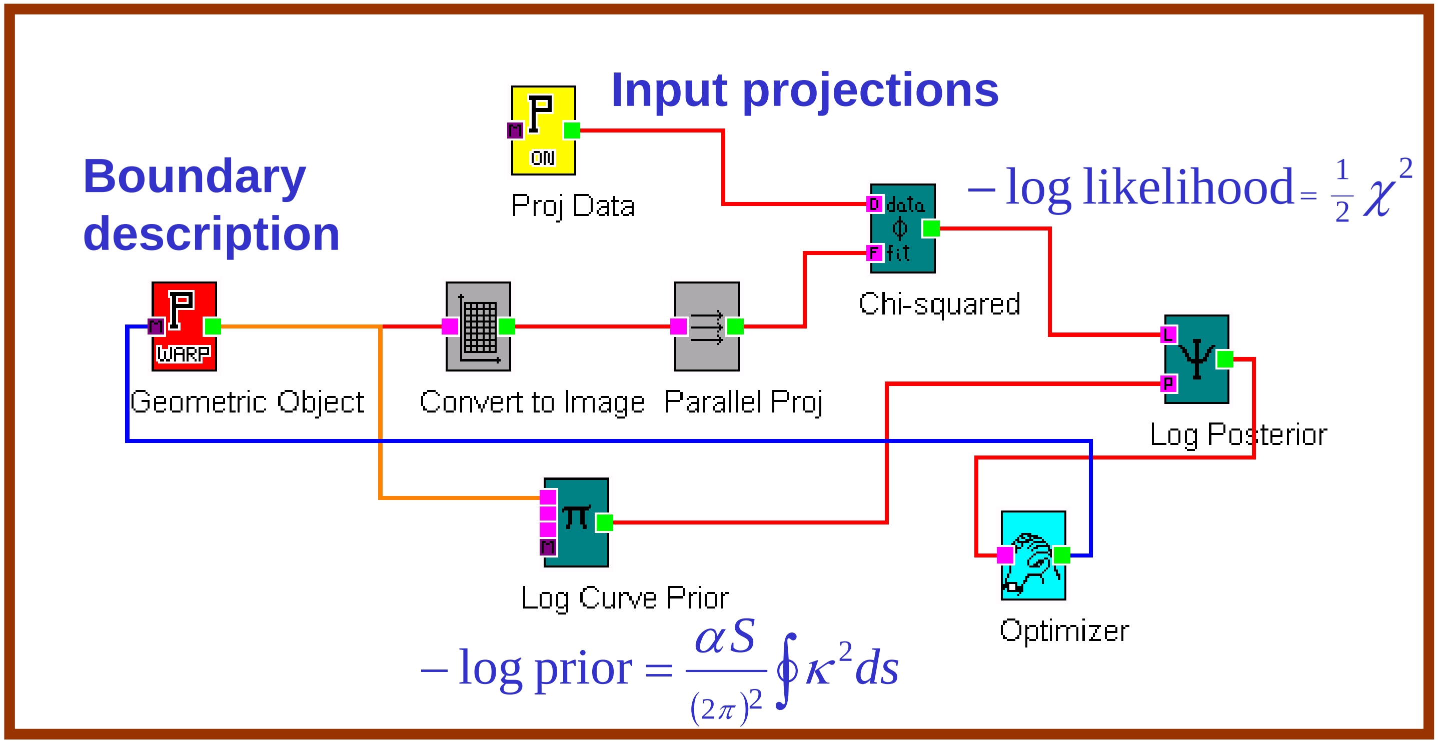

  - ▹ Optimizer uses gradients that are efficiently calculated by adjoint differentiation in code technique (ADICT)

MAP reconstruction

- Determine boundary that maximizes posterior probability

Reconstructed boundary (gray-scale) compared with original object (red line)

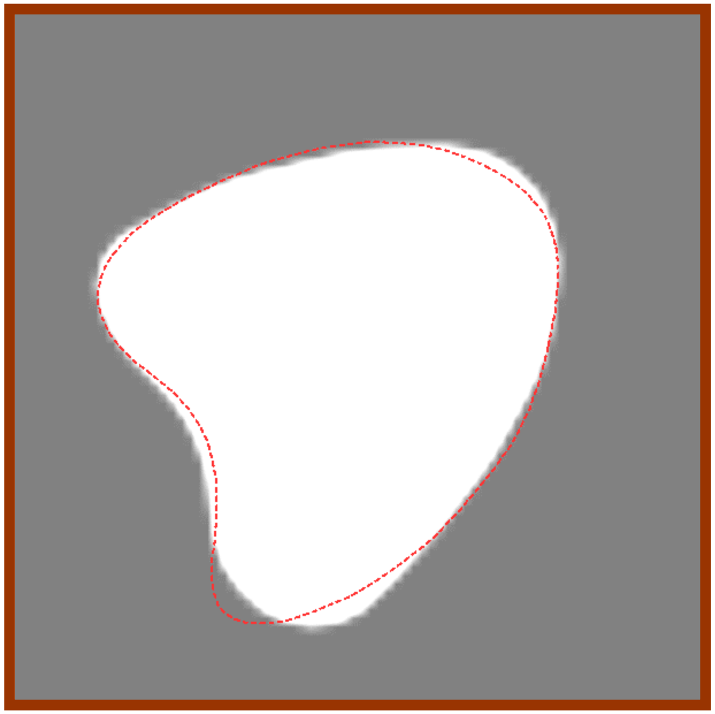

MCMC
Markov Chain Monte Carlo

- Generate sequence of random samples from specified probability density function
  - ▹ represent pdf with finite number of samples
- Markov chain - probability of *k*th state in sequence depends only on (*k*-*1*)th state
- Monte Carlo procedure
  - ▹ based on pseudo-random numbers generated by computer
  - ▹ estimated quantities always uncertain because of event statistics

MCMC - Metropolis algorithm

Generates sequence of random samples from an arbitrary target probability density function, q(**x**)

- Metropolis algorithm:
  - draw trial step from symmetric pdf, i.e., $T(\Delta\mathbf{x}) = T(-\Delta\mathbf{x})$
  - accept or reject trial step based on $q(\mathbf{x}_k + \Delta\mathbf{x})/q(\mathbf{x}_k)$
  - relies only on calculation of target pdf q(**x**)
  - simple and generally applicable
  - works well for small number of parameters; inefficient for many

Probability $q(x_1, x_2)$

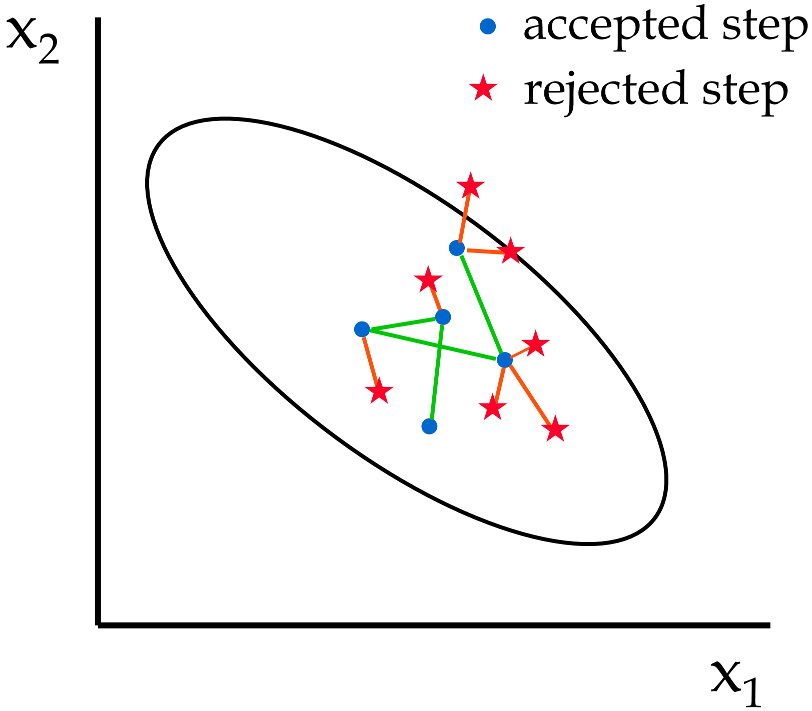

MCMC - Metropolis algorithm

- Generate sequence of random samples from probability density function $q(\mathbf{x})$, where $\mathbf{x}$ is vector of parameters
- Start with arbitrary $\mathbf{x}_0$
- Recursive loop to generate sequence: at point $\mathbf{x}_k$
  - ▹ pick new trial vector $\mathbf{x}^* = \mathbf{x}_k + \Delta\mathbf{x}$,
    where $\Delta\mathbf{x}$ drawn from symmetric p.d.f.
  - ▹ calculate: $r = q(\mathbf{x}^*)/q(\mathbf{x}_k)$
    - if greater than 1, accept step; $\mathbf{x}_{k+1} = \mathbf{x}^*$
    - if less than 1, accept step with probability r
    - otherwise, repeat previous point; $\mathbf{x}_{k+1} = \mathbf{x}_k$

The Bayes Inference Engine

- BIE data-flow diagram to produce MCMC sequence

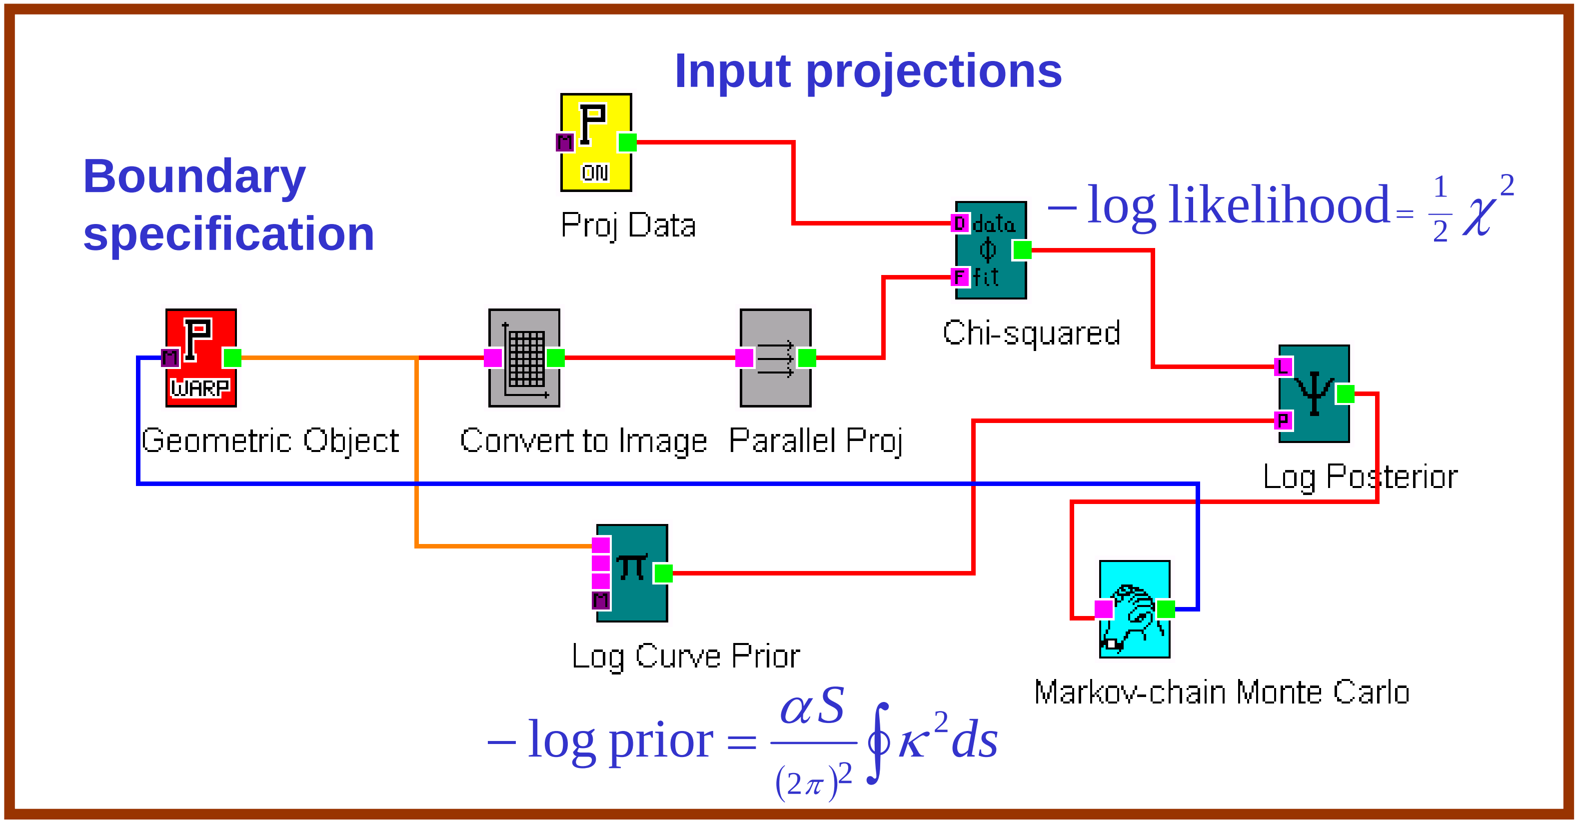

Probabilistic interpretation of prior for deformable model

- Probability of shape: $\sim \exp\left[-\frac{\alpha S}{(2\pi)^2} \oint \kappa^2 ds \right]$
- Sample prior pdf using MCMC
  - ▹ shows variety of shapes deemed admissible before experiment
  - ▹ decide on $\alpha = 5$ on basis of appearance of shapes

Plausible shapes drawn from prior for α = 5

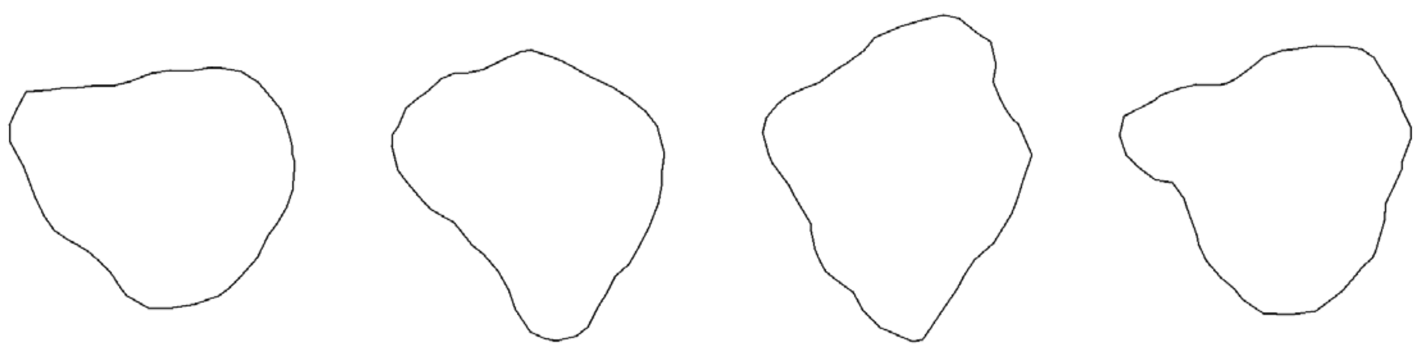

Visualization of uncertainty

- Problem inherently difficult for numerous parameters
  - ▹ wish to see correlations among uncertainties in parameters
- View MCMC sequence as video loop
  - ▹ advantage is one directly observes model in normal way
- View several plausible realizations from MCMC sequence
- Marginalized uncertainties (one parameter at a time)
  - ▹ rms uncertainty (or variance) for each parameter
  - ▹ credible intervals

Uncertainties in two-view reconstruction

- From MCMC samples from posterior with 150,000 steps, display three selected boundaries
  - ▹ these represent alternative plausible solutions

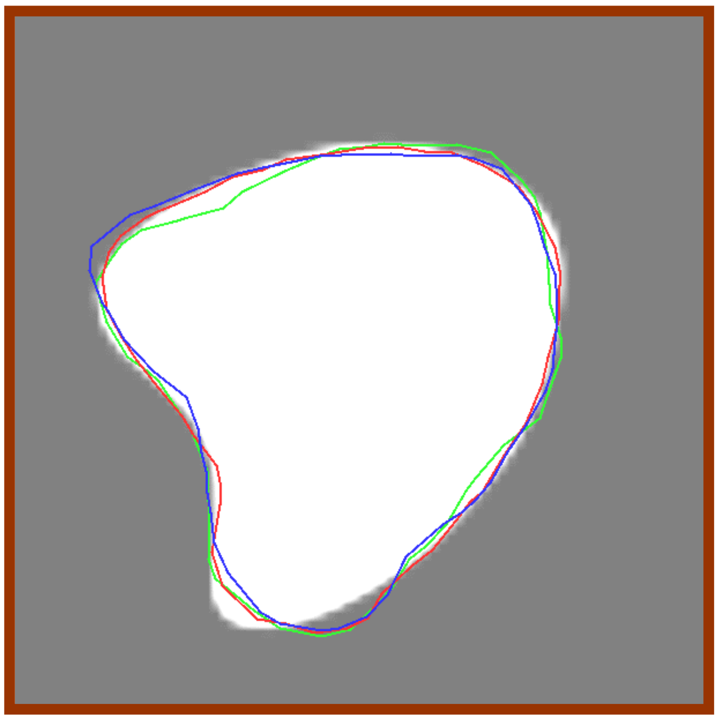

compared to original object

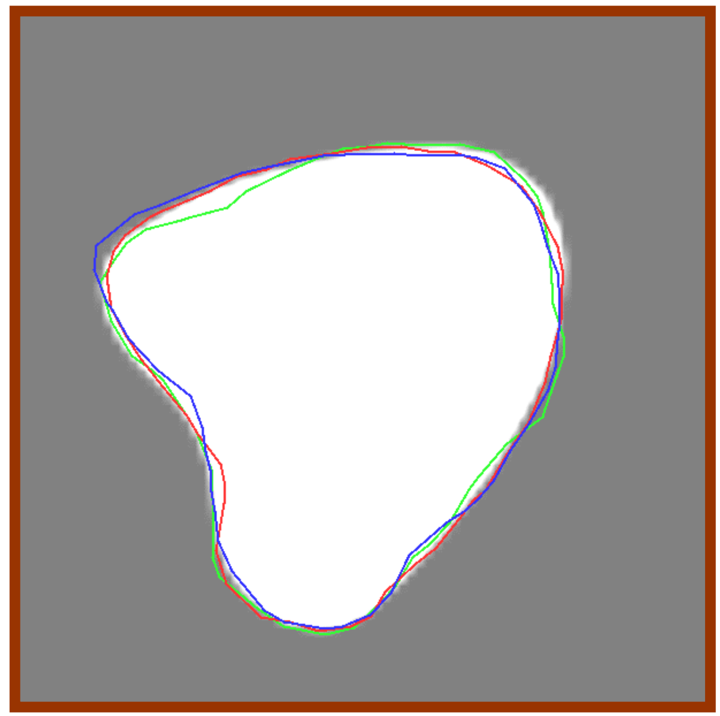

compared to MAP estimated object

Posterior mean of gray-scale image

- Average gray-scale images over MCMC samples from posterior
- Value of pixel is probability it lies inside object boundary
- Amount of blur in edge is related to magnitude of uncertainty in edge localization
- Observe that posterior median nearly same as MAP boundary
  - ▹ implies posterior probability distribution symmetric about MAP parameter set

Posterior mean of gray-scale image

- Pixels in posterior mean image with value 0.5 represent posterior median boundary position
  - ▹ similar to MAP boundary for two-view problem

Posterior mean image compared to MAP boundary (red line)

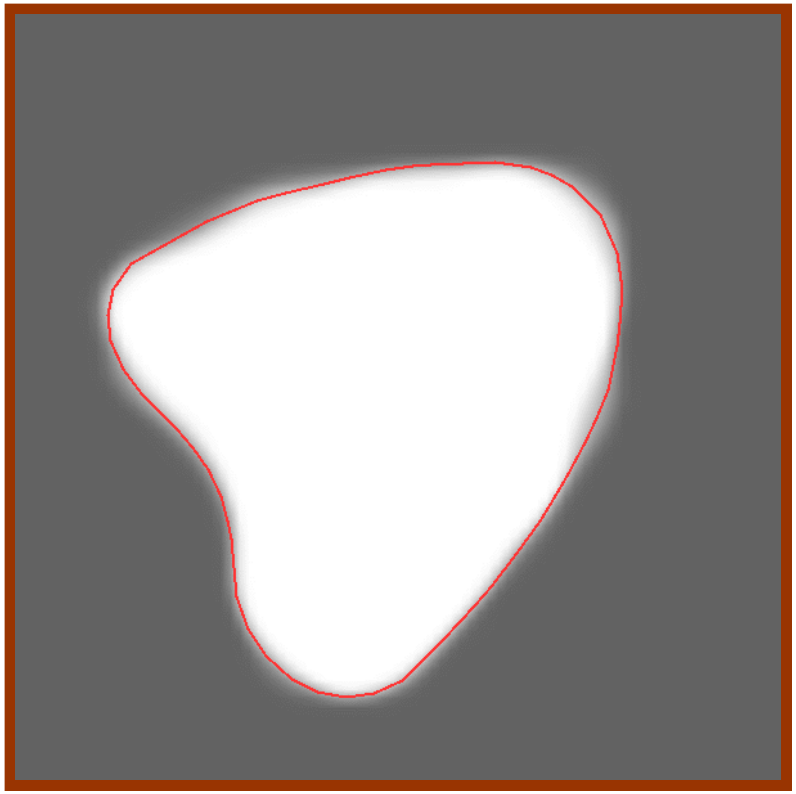

Uncertainty in edge localization

- Steepness of edge profile of posterior mean image indicates uncertainty in edge localization
  - ▹ uncertainty is nonstationary; varies with position

At top of reconstruction (tangent point)

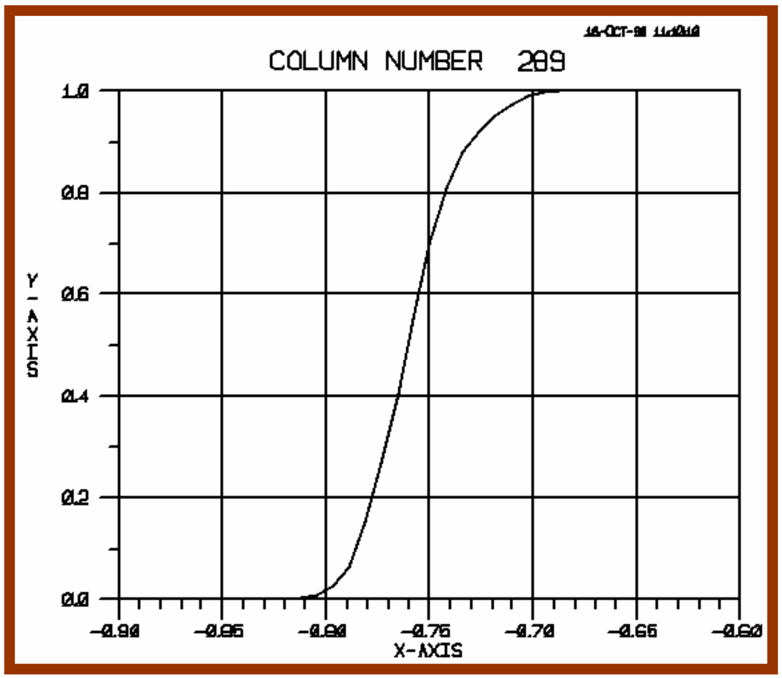

Top, left of center, less well determined

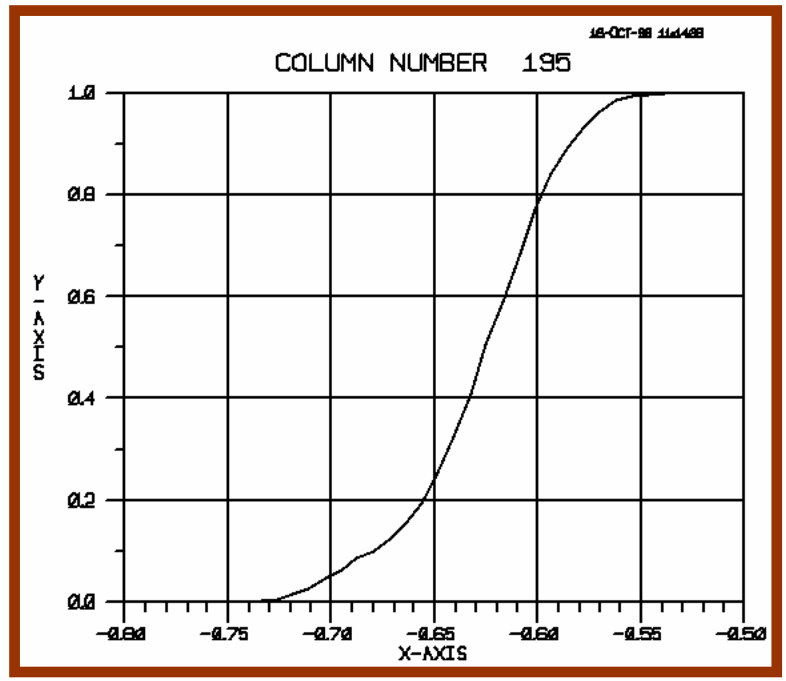

Credible interval

- Bayesian "confidence interval"
  - ▹ probability that actual parameter lies within interval
  - ▹ different from standard definition of confidence interval, which is based on (hypothetical) repeated experiments
- For MCMC posterior mean image, determines credible interval for boundary position
  - ▹ 95% credible interval is region of posterior mean image whose pixel values lie between 0.025 and 0.975.

Credible interval

- 95% credible interval of boundary localization for two-view reconstruction compared with original object boundary (red line)
  - ▹ narrower at tangent points
  - ▹ 92% of original boundary lies inside 95% credible interval
- Marginalized measure of uncertainty - ignores correlations among different positions

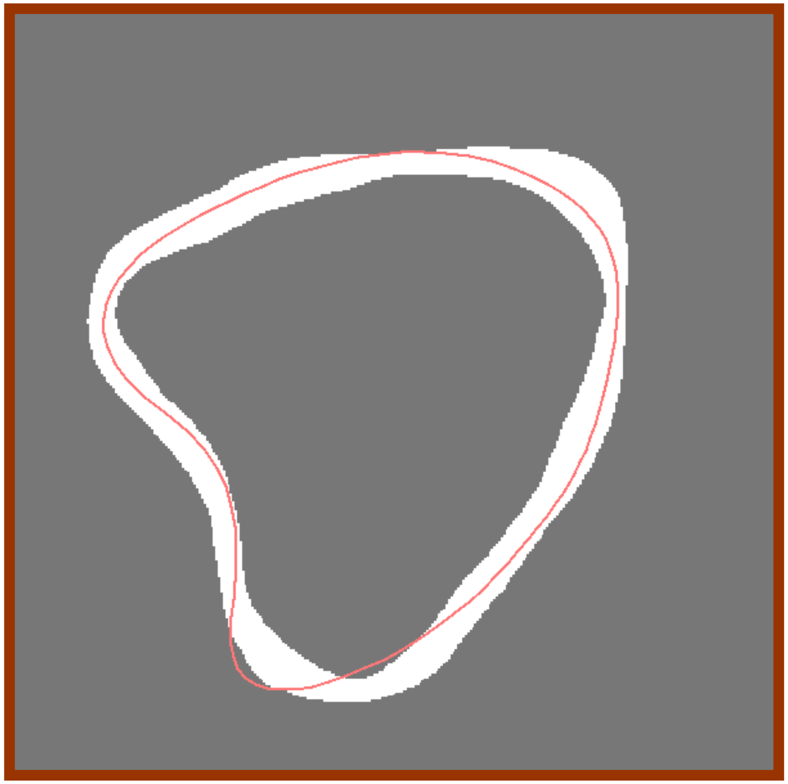

Important issues

- Bayesian vs. frequentist approach to uncertainty assessment
  - ▹ MCMC sampling of posterior
    - single data set, single object
  - ▹ Monte Carlo simulation of repeated experiments to determine characteristics of the estimator used
    - variety of data sets (variety of objects)
- Advantages of Bayesian approach
  - ▹ applies to specific data set supplied
  - ▹ illuminates null space; multiple solutions that yield exactly same measurements

Important issues

- Markov Chain Monte Carlo
  - ▹ efficiency - number of function evaluations required to obtain given level of accuracy in posterior characterization
    - choice of trial step distribution
    - account for correlations among different parameters
    - for Metropolis algorithm, efficiency ~ (number parameters)$^{-1}$
  - ▹ burn in period at beginning of MCMC sequence to reach equilibrium with target pdf
    - how long should burn in be?
  - ▹ need algorithms to improve efficiency
    - hybrid method, based on Hamiltonian dynamics (needs gradient)

Hard truth method

- Interpret φ = - log probability as potential function; sum of
  - ▹ deformation energy
  - ▹ $\frac{1}{2}\chi^2$
- Stiffness of model proportional to curvature of φ
- Row of covariance matrix found by applying a force to parameters at MAP solution and reminimizing φ

Applying force (white bar) to MAP boundary (red) moves it to new location (yellow-dashed)

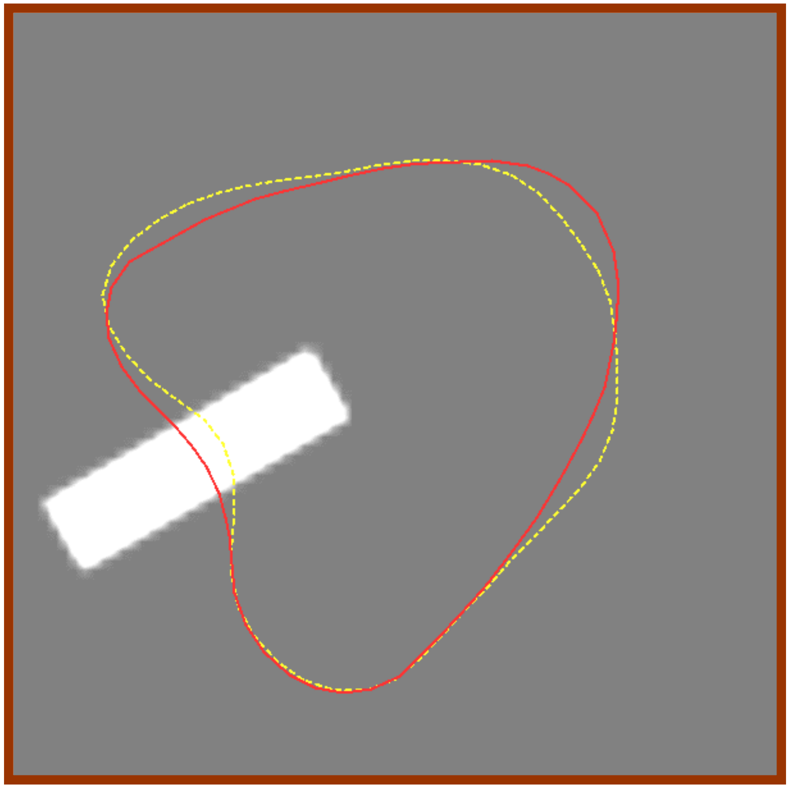

Bibliography

- "Uncertainty assessment for reconstructions based on deformable models," K. M. Hanson et al., *Int. J. Imaging Syst. Technol.* **8**, pp. 506-512 (1997)
- "Posterior sampling with improved efficiency," K. M. Hanson et al., in *Medical Imaging: Image Processing*, *Proc. SPIE* **3338**, pp. 371-382 (1998)
- "The hard truth," K. M. Hanson et al., in *Maximum Entropy and Bayesian Methods*, pp. 157-164 (Kluwer, 1996)
- "Operation of the Bayes Inference Engine," K. M. Hanson et al., in *Maximum Entropy and Bayesian Methods*, pp. 309-318 (Kluwer, 1999)

These and other papers available under http://home.lanl.gov/kmh/