# Efficient MCMC Sampling with Implicit Shape Representations

Jason Chang & John W. Fisher III

## Motivation

- Segmentation often formulated as energy minimization: $\arg\min E(\mathcal{C};I)$
- An "optimal" solution might not accurately represent how humans segment images

**Approach:** Instead of MAP estimation, characterize the entire distribution

- Marginal statistics are more robust to noise and can characterize multiple modes
- Calculate statistics from samples of $p(\mathcal{C}\,|\,I) \propto \exp(-E(\mathcal{C};I))$

**Problem:** How do we draw samples from this distribution?

## Level Set Representation

- Level set methods [10] represent a 2D curve, $\mathcal{C}$, implicitly with a 3D surface, $\varphi$
- Topological changes are easily incorporated by simply changing the surface

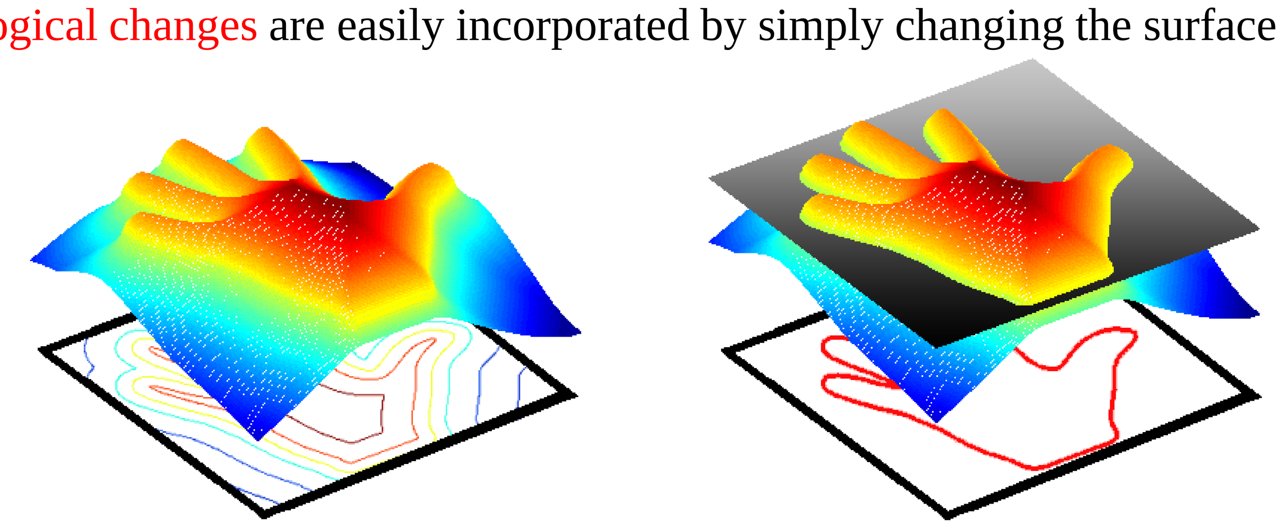

## Metropolis-Hastings

- Metropolis-Hastings MCMC [9] allows sampling from any distribution if it can be evaluated up to a constant scale
- Process: (1) Generate a sample from a proposal distribution $q(\hat{\varphi}^{(t+1)}\,|\,\varphi^{(t)})$
  (2) Accept the sample with the following probability

$$\Pr\left[\varphi^{(t+1)} = \hat{\varphi}^{(t+1)}\,|\,\varphi^{(t)}, I\right] = \min\left(\underbrace{\frac{p(\hat{\varphi}^{(t+1)}\,|\,I)}{p(\varphi^{(t)}\,|\,I)}}_{\text{Posterior Sample Ratio (PSR)}} \cdot \underbrace{\frac{q(\varphi^{(t)}\,|\,\hat{\varphi}^{(t+1)})}{q(\hat{\varphi}^{(t+1)}\,|\,\varphi^{(t)})}}_{\text{Forward-Backward Ratio (FBR)}},\ 1\right)$$

- After enough iterations, we obtain one sample from the true distribution

## Previous Methods

- [4] alternates between an implicit and explicit representation
- Perturbs explicit points normal to the curve
- Must solve correspondence problem because of the FBR

$$\frac{p(\hat{\varphi}^{(t+1)}\,|\,I)}{p(\varphi^{(t)}\,|\,I)} \cdot \frac{q(\varphi^{(t)}\,|\,\hat{\varphi}^{(t+1)})}{q(\hat{\varphi}^{(t+1)}\,|\,\varphi^{(t)})}$$

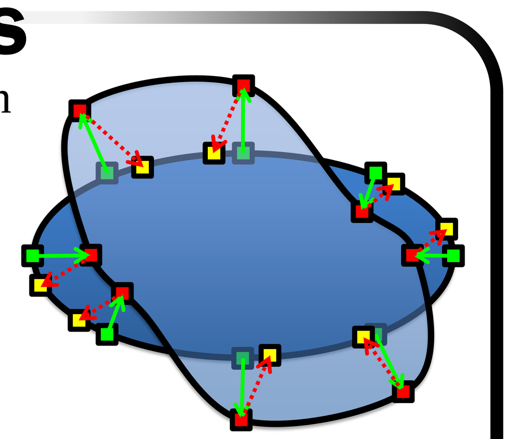

- [3] stays in the implicit representation
- If signed-distance function (SDF) imposed after perturbation, a multiple-to-multiple mapping exists in correspondence problem
- Smooths a perturbation to a single point that preserves the SDF
- Solves the correspondence problem, but slow convergence

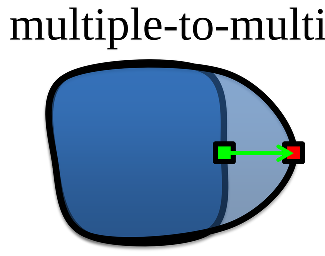

- Both methods restricted to binary segmentation of a single simply connected shape


[1] Brox, T., Weickert, J.: Level set segemtnation with multiple regions. In: *Image Processing, IEEE Transactions on*, 2006.
[2] Chang, J. and J. W. Fisher III. Analysis of Orientation and Scale in Smoothly Varying Textures. In: *ICCV*, 2009.
[3] Chen, S. and R. J. Radke. Markov Chaing Monte Carlo Shape Sampling using Level Sets. In: *NORDIA*, in conjunction with *ICCV*, 2009.
[4] Fan, A.C., Fisher III, J.W., Wells III, W.M., Levitt, J.J., Willsky, A.S.: MCMC curve sampling for image segmentation. In: *MICCAI*, 2007.
[5] Heiler M. and C. Schnorr. Natural Image Statistics for Natural Image Segmentation. In: *ICCV*, 2003.
[6] Houhou, N., Thiran, J.P., Bresson, X.: Fast texture segmentation model based on the shape operator and active contour. In: *CVPR*, 2008.
[7] Kim, J., Fisher III, J.W., Yezzi, A., Cetin, M., Willsky, A.: A nonparametric statistcal method for image segmentation using information theory and curve evolution. In: *Image Processing, IEEE Transactions on*, 2005.
[8] Martin, D., Fowlkes, C., Tal, D., Malik, J.: A database of human segmented natural images and its applications to evaluating segmentation algorithms and measuring ecological statistics. In: *ICCV*, 2001.
[9] Metropolis, N., Rosenbluth, A., Rosenbluth, M., Teller, A., Teller, E.: Equations of state calculations by fast computing machines. In: *Journal of Chemical Physics*, 1953.
[10] Osher, S., Fedkiw, R.: Level Set Methods and Dynamic Implicit Surfaces. Springer (2002).


## Key Ideas

(1) Ignore SDF constraint – solves correspondence problem with large perturbations

$$\hat{\varphi}^{(t+1)} = \varphi^{(t)} + \mathbf{f}^{(t)}$$

(2) Bias proposal towards likely samples – increase Markov chain mixing time

- Biased proposals increase the PSR but decrease the FBR

$$\text{FBR} = \frac{\Pr[\text{red}]}{\Pr[\text{green}]} \ll 1$$

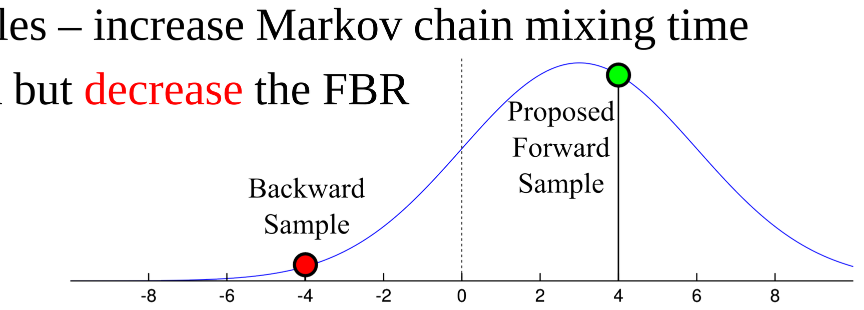


- FBR decreases with bias strength and number of random variables

(3) Take advantage of spatial correlations to only decrease the FBR a little

## Biased Filtered Point Sampling

- Generate a proposal perturbation as follows:

$$\mathbf{f}^{(t)} = \mathbf{h}^{(t)} * \left(\mathbf{c}^{(t)} \circ \mathbf{n}^{(t)}\right)$$

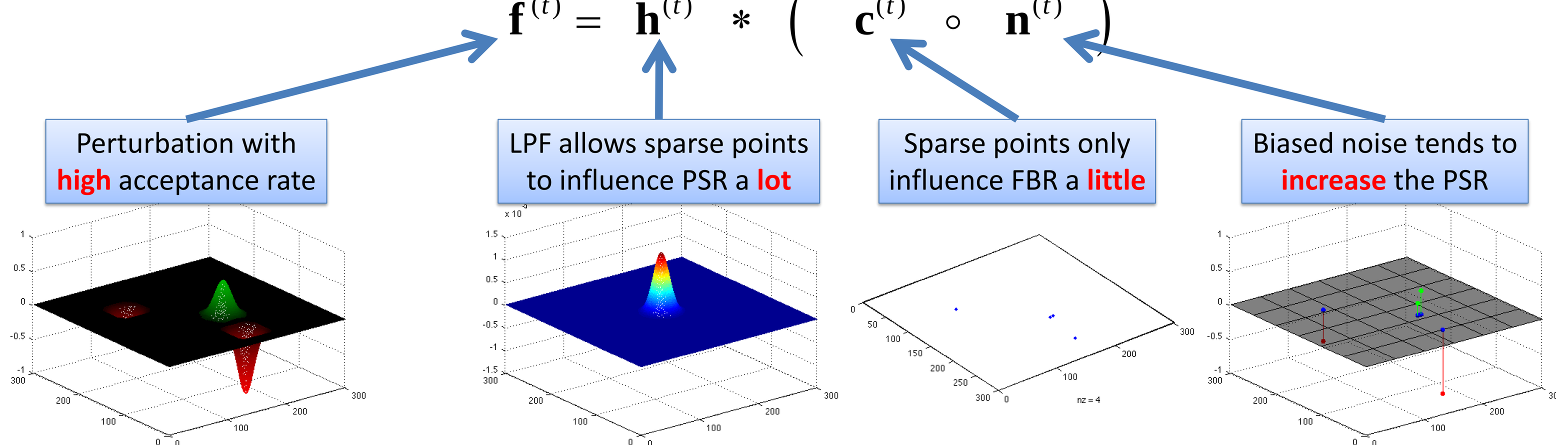


$c_i^{(t)} \sim \text{Bernoulli}\left(\alpha_c \exp\left[-v_i^{(t)} \cdot \text{sign}\left(\varphi_i^{(t)}\right)\right]\right)$ $\quad n_i^{(t)} \sim \mathcal{N}\left(\alpha_n v_i^{(t)}, \sigma^2\right)$ $\quad v_i^{(t)} = -\nabla E(\mathcal{C};I)$

## M-ary Extension

- [1] uses $M$ level sets to represent $M$ regions in an optimization framework
  - This leads to problems in generating proper proposals
- Use $M$-1 level sets to represent $M$ regions
  - Must ensure no overlap or vacuum conditions

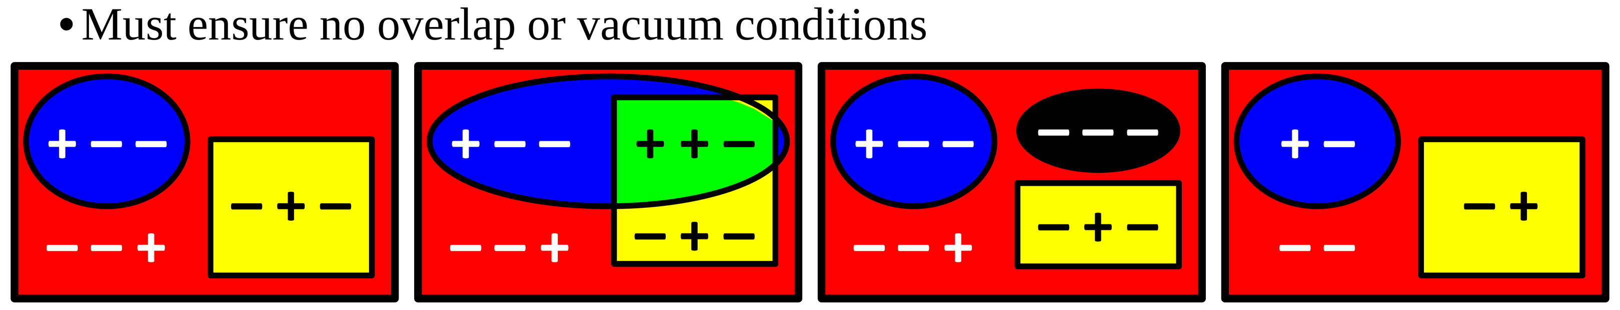


- Implicitly define a "null" region:

$$R_\ell = \{\varphi_\ell \geq 0\}\ ,\quad \forall \ell \in \mathcal{L} = \{1,2,\ldots,M\} \qquad R_0 = \bigcap_{\ell \in \mathcal{L}} \{\varphi_\ell < 0\}$$

- M-ary Proposal
  (1) Choose a random level set, $\varphi_\ell$
  (2) Generate perturbation that only affects $\varphi_\ell$: $\mathbf{f}^{(t)} = \mathbf{h}^{(t)} * \left(\mathbf{c} \circ \mathbf{n}^{(t)}\right) \circ \mathbf{1}\{R_\ell \cup R_0\}$

## Computation Times

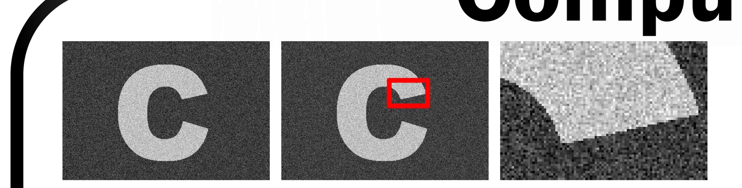

- Comparison to sampling methods [3] and [4]
- Incorporate a gradient bias in all methods to view its affects
- Results shown after 100,000 its.

| Alg. | Iterations to Convergence | Secs. Per Iteration | Total Gain |
|---|---|---|---|
| BFPS | 150 | 0.03 | x1 |
| UFPS | 40,000 | 0.03 | x267 |
| B[3] | 254,000* | 0.30 | x16,933 |
| U[3] | 896,000* | 0.26 | x51,769 |
| B[4] | 321,000* | 5.00 | x356,667 |
| U[4] | 336,000* | 5.00 | x373,333 |

*Optimistic lower bound on number of iterations obtained from linearly interpolating the sub-linear growth of the average energy

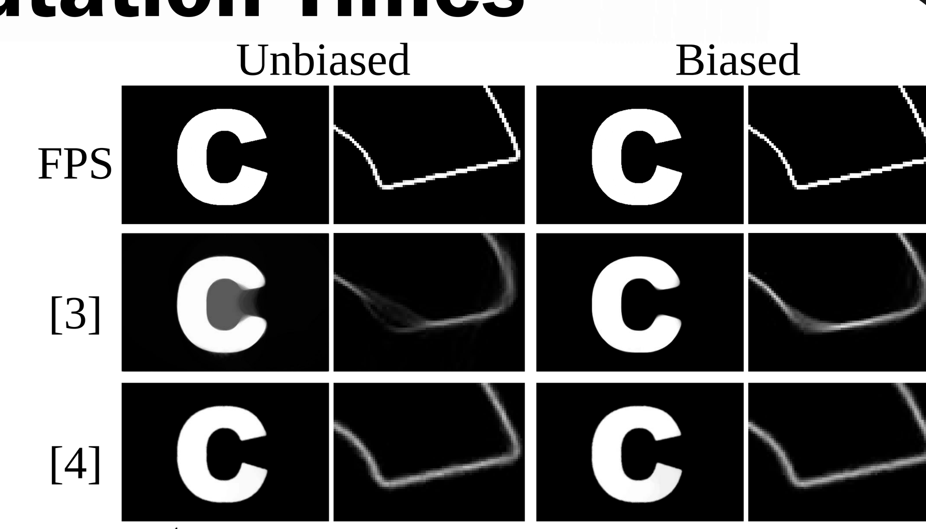


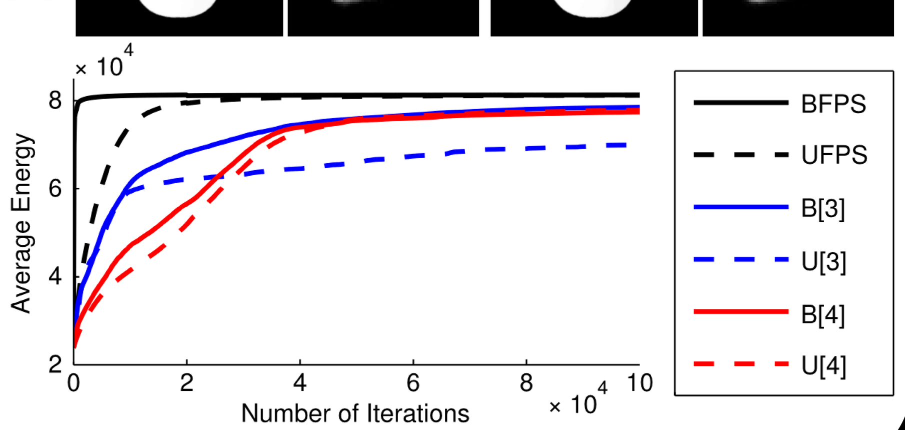


## Synthetic Results

- This example shows the importance of allowing topological changes. Methods [3] and [4] are unable to capture both the inside and outside of the circle.

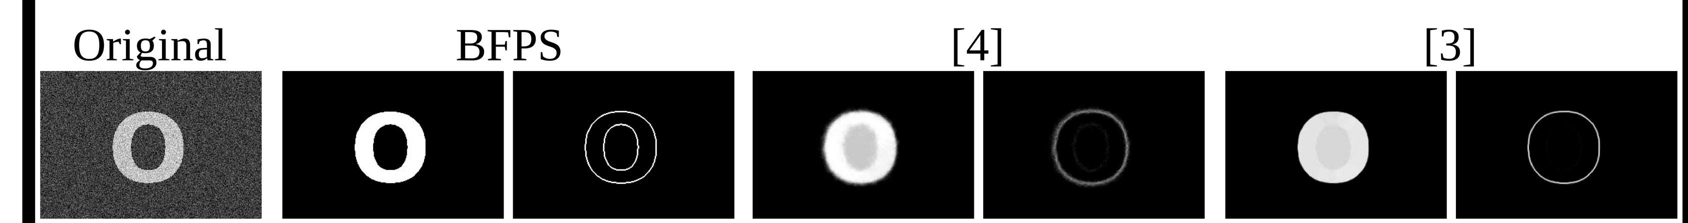


- This example shows the benefit of marginal statistics in the presence of noise
- In high SNR (single-mode) cases, sampling does not give much improvement because the distribution is very peaked
- In low SNR cases, the improvement from sampling is more pronounced

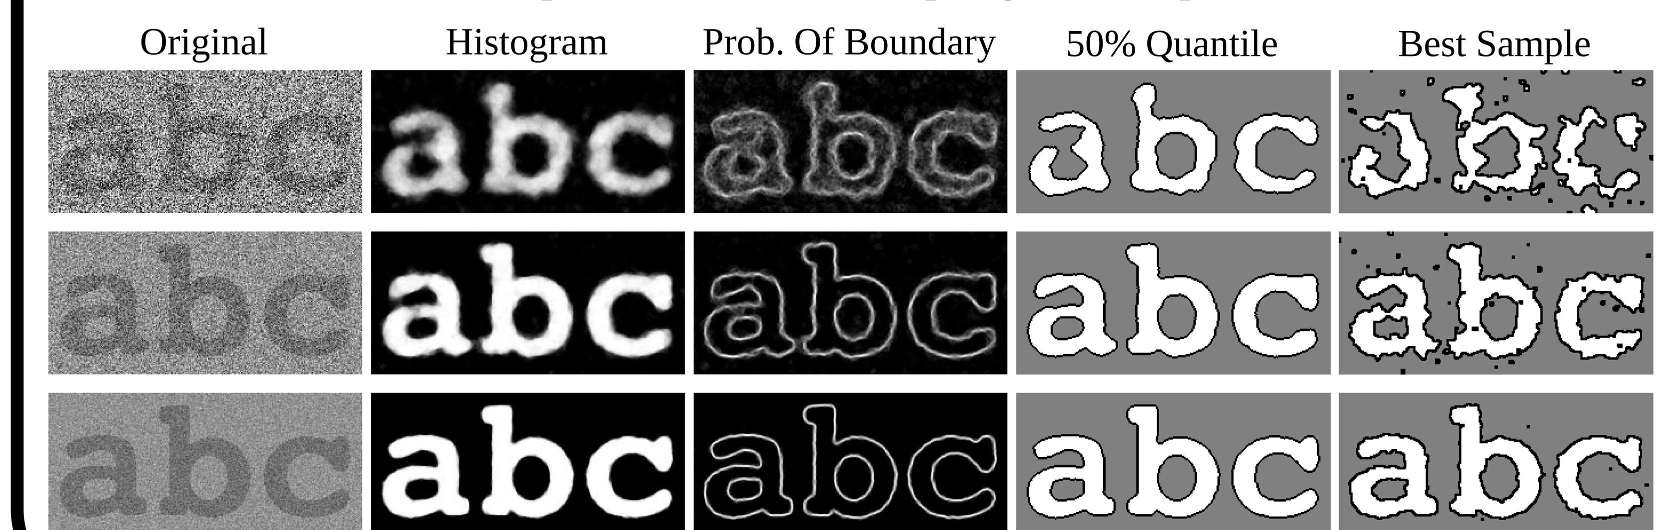


## Natural Image Results

The following are results obtained on [8]. The '**X**' is the F-measure using the sampling approach, and the '**+**' is the F-measure using gradient descent.