Efficient Topology-Controlled Sampling of Implicit Shapes

Jason Chang & John W. Fisher III

Motivation

Image segmentation is often formulated as an energy minimization problem, where one tries to find the labeling, ℓ, that minimizes a surrogate energy functional $E(\ell;x)$. Due to the ill-posed nature of the problem, oftentimes multiple solutions exist and the optimal energy may not correspond to the best segmentations.

In this event, statistics over the distribution of segmentations may be more informative. Consequently, we look at sampling from the following distribution:

$$p_{L|X}(\ell \mid x) \propto \exp\left(-E(\ell;x)\right)$$

The work of [2] illustrates the improvement from sampling algorithms in boundary detection on the Berkeley Segmentation Dataset [8].

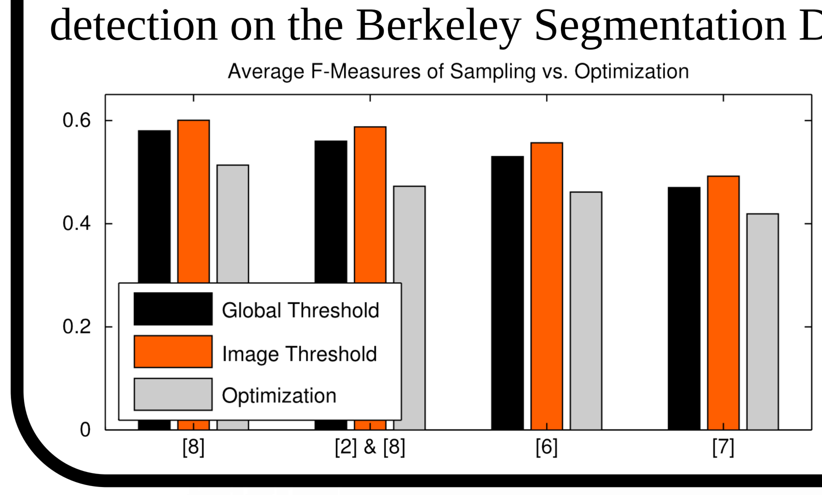

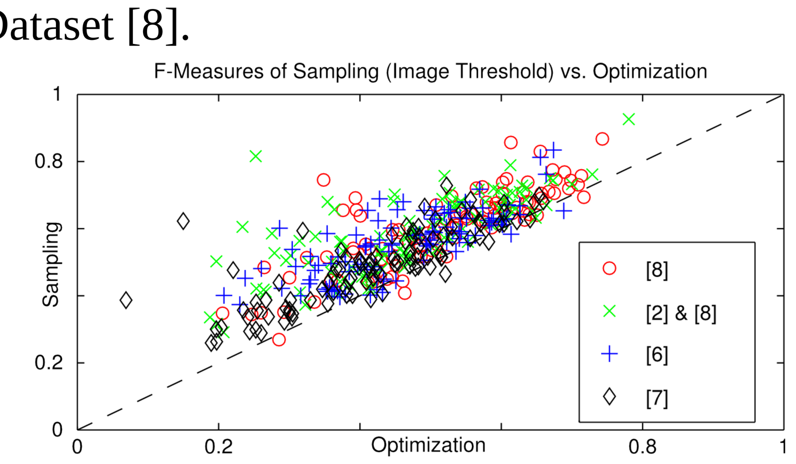

Level Set Representation

The level set representation [10] is widely used in image segmentation. A 2D labeling, ℓ, is implicitly represented by the sign of the level set of a 3D surface, φ. Evolving the curve is accomplished by changing the entire surface. This implicit representation easily represents topological changes.

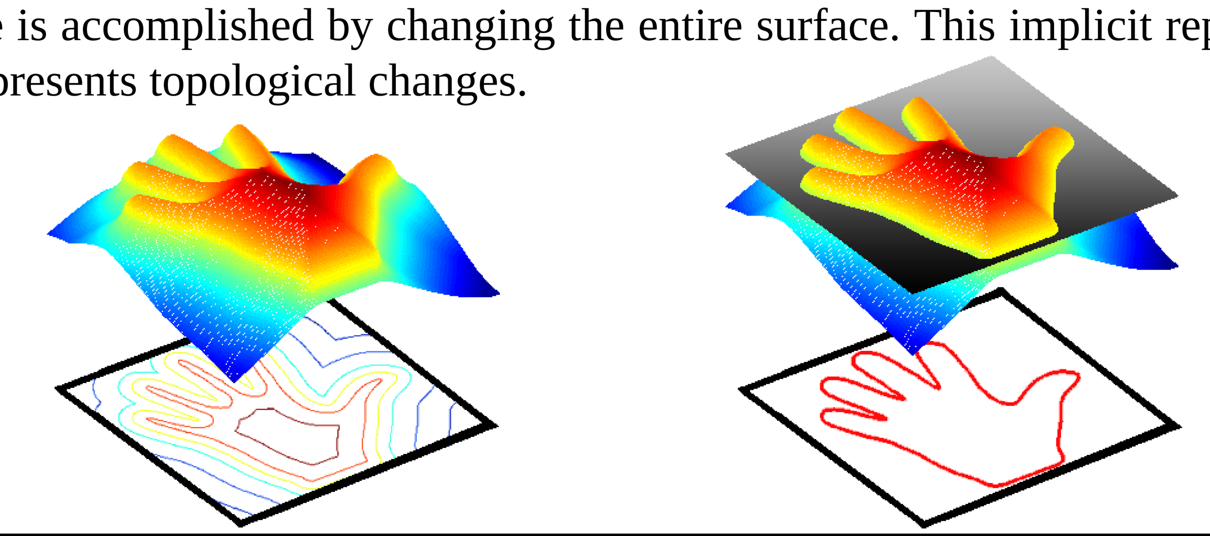

Metropolis-Hastings & Gibbs

Metropolis-Hastings (MH) MCMC methods [9] provide a way to sample from a distribution as long as the density can be evaluated to a constant scale factor. The use of Metropolis-Hastings for sampling curves was first proposed by [4]. This method was very slow and was restricted to a single simply connected component.

Given a previous sample, $\varphi^{(t)}$, one generates a new sample, $\hat{\varphi}^{(t+1)}$ from a proposal distribution, $q\left(\hat{\varphi}^{(t+1)} \mid \varphi^{(t)}\right)$, and accepts the sample with probability:

$$\Pr\left[\varphi^{(t+1)} = \hat{\varphi}^{(t+1)} \mid \varphi^{(t)}, x\right] = \min\left(\underbrace{\frac{p_{L|X}\left(\ell\left(\hat{\varphi}^{(t+1)}\right) \mid x\right)}{p_{L|X}\left(\ell\left(\varphi^{(t)}\right) \mid x\right)} \cdot \frac{q\left(\varphi^{(t)} \mid \hat{\varphi}^{(t+1)}\right)}{q\left(\hat{\varphi}^{(t+1)} \mid \varphi^{(t)}\right)}}_{\text{Hastings Ratio}}, 1\right)$$

Given enough time, it is guaranteed that this procedure ultimately produces a single sample from the desired distribution as long as the Markov chain is ergodic. However, the chain can typically take upwards of 100,000 samples to converge. One must take particular care in designing proposal distributions that have high acceptance probability and that can be generated and evaluated quickly.

Gibbs sampling is a special case of MH-MCMC where the proposal is chosen to be proportional to the target distribution:

$$q\left(\hat{\varphi}^{(t+1)} \mid \varphi^{(t)}\right) \propto p_{L|X}\left(\ell\left(\hat{\varphi}^{(t+1)}\right) \mid x\right)$$

In this case, the proposal is always accepted because the Hastings ratio evaluates to 1.

Gibbs-Inspired Proposal

A proposal is generated from the following procedure:

1) Generate a random mask, m, that selects a subset of pixels
2) Add a random constant value, f, to all pixels within the mask

$$\hat{\varphi}^{(t+1)}(m, f) = \varphi^{(t)} + f \cdot m$$

Because the target distribution only depends on the sign of the height at each pixel, any f in a range, r_i, has the same probability under the target distribution.

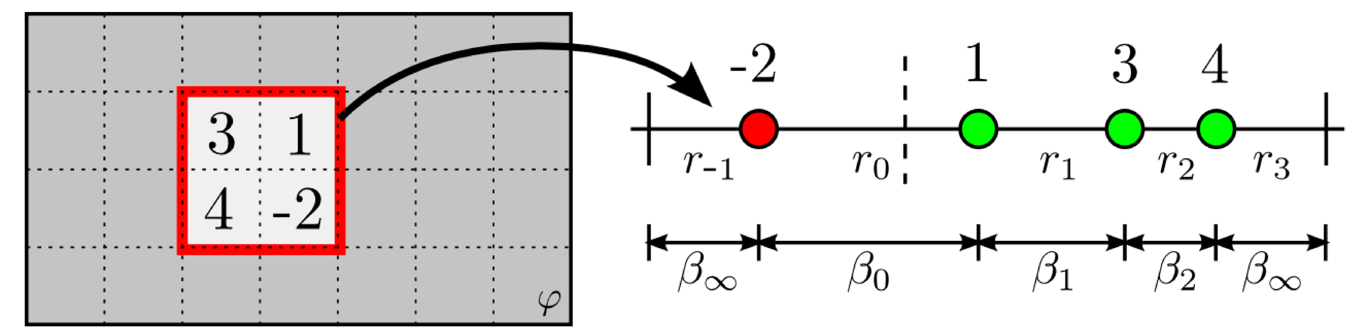

The perturbation is chosen by selecting a range followed by uniformly sampling a value in that range.

$$p_{F|M\Phi}\left(f \mid m, \varphi^{(t)}\right) = p_{R|M\Phi}\left(r \mid m, \varphi^{(t)}\right) p_{F|RM\Phi}\left(f \mid r, m, \varphi^{(t)}\right)$$

$$= p_{R|M\Phi}\left(r \mid m, \varphi^{(t)}\right) \frac{1}{\beta_r}$$

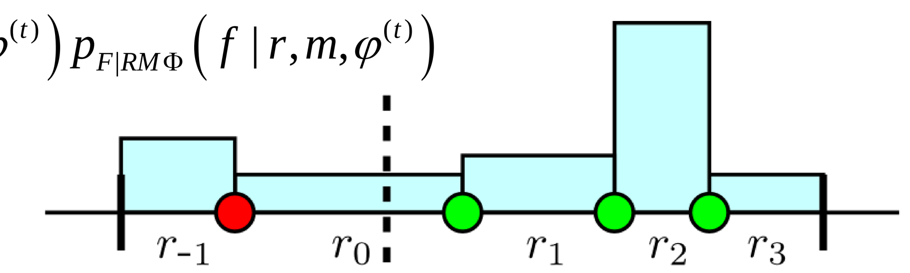

Because the value of f within a range does not affect the sign of the resulting level-set, the proposed labeling can be expressed as

$$\hat{\ell}^{(t+1)}(m, r) = \operatorname{sgn}\left(\hat{\varphi}^{(t+1)}\left(m, f \sim p_{F|R}(f \mid r, m, \varphi)\right)\right)$$

The following range proposal distribution results in a Hasting's ratio of 1:

$$p\left(r \mid m, \varphi^{(t)}\right) \propto \beta_r \, p_{L|X}\left(\hat{\ell}^{(t+1)}(m, r)\right)$$

Computation Times

We compare the running time of our GIMH sampler with three other binary shape samplers [2], [3], and [4]. The following shows the average energy across samples plotted against computation time. GIMH is approximately 100x faster than the next fastest algorithm, BFPS.

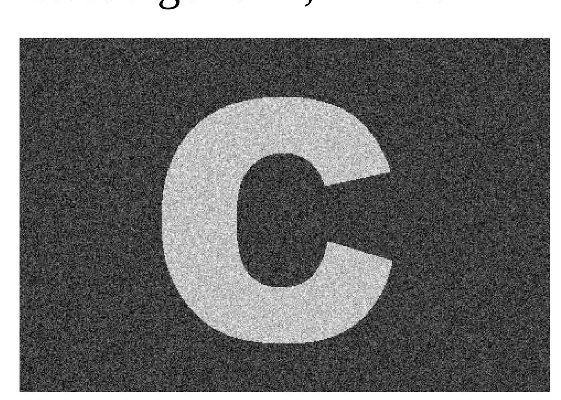

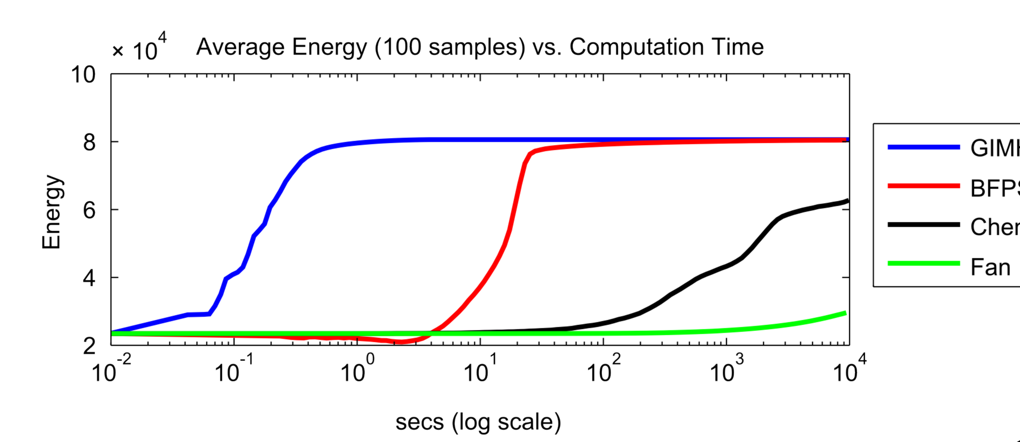

Gibbs Sampling Comparison

The formulation of the energy allows the use of a traditional Gibbs sampler. However, the small moves of a typical Gibbs sampler tend to converge to local extrema, such as the case below.

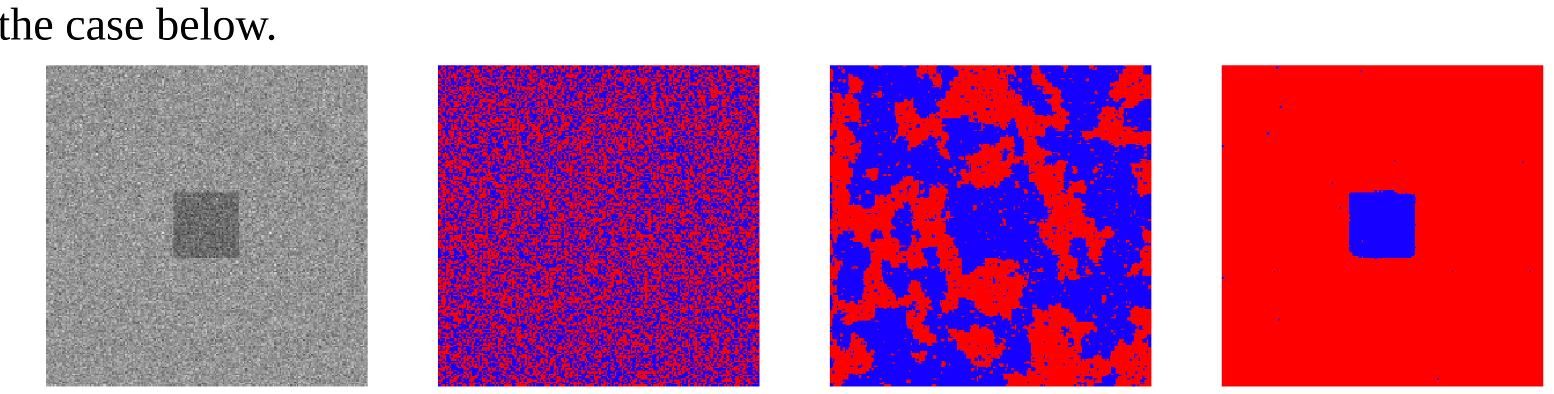

Image | Initialization | Gibbs Sample | GIMH Sample

GIMH is, however, related to block Gibbs sampling. In block Gibbs, a block of size M requires one to evaluate 2^M number of different configurations. In GIMH, the level-set implicitly orders the block of pixels, resulting in only M+1 configurations. Ergodicity of the chain is ensured because the ordering varies with time.

Code available at http://people.csail.mit.edu/jchang7/

Topology Control

Topological numbers [5] allow one to identify if changing a single pixel will cause a topology change. These numbers count the number of 4- and 8-connected components in a 3x3 neighborhood. Topology changes occur unless $T_4 = T_8 = 1$.

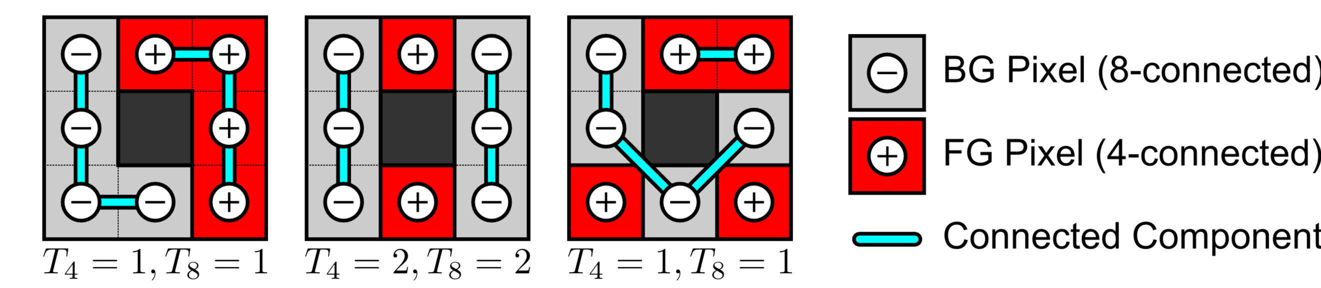

Extended topological numbers [12], can *identify* the actual topological change that occurs. These can be computationally prohibitive in 3D, but are easily found in 2D.

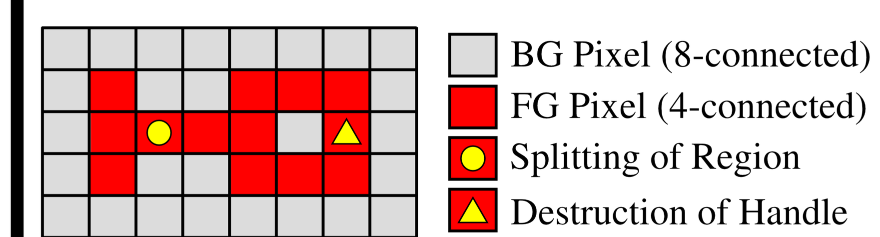

$$T_4(\circ) = 2 \quad = \quad T_4(\triangle) = 2$$
$$T_8(\circ) = 1 \quad = \quad T_8(\triangle) = 2$$
$$T_4^+(\circ) = 2 \quad \neq \quad T_4^+(\triangle) = 1$$
$$T_8^+(\circ) = 1 \quad \neq \quad T_8^+(\triangle) = 2$$

T_4	T_4^+	T_8	T_8^+	Add to FG		Add to BG	
				FG	BG	FG	BG
0	0	1	1	CR	CH	DR	DH
1	1	0	0	DH	DR	CH	CR
1	1	1	1	-	-	-	-
≥2	<T_4	X	X	CH	SR	X	X
≥2	≥2	X	X	MR	DH	X	X
X	X	≥2	<T_8	X	X	SR	CH
X	X	≥2	≥2	X	X	DH	MR

'C' – Create, 'D' – Destroy, 'S' – Split, 'M' – Merge
'H' – Handle(s), 'R' – Region(s), 'X' – any value;

Imposing Different Constraints

With the mapping of topological numbers to topology changes, one can control the evolution of the shape to have any topology. We show four different topology constraints: unconstrained, topology preserving, genus preserving, and foreground connected component preserving. A sample from each constraint is shown below.

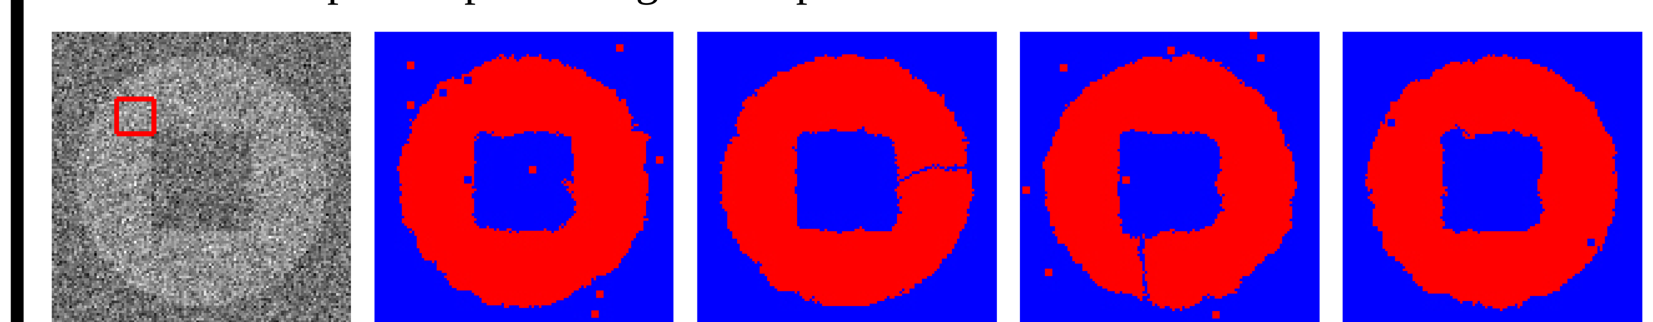

Initialization | Unconstrained | Topology Preserving | Genus Preserving | FG Connected Component Preserving

Synthetic Results

The following examples illustrate when topology constraints can help. The additional prior knowledge can improve robustness. However, incorrect, overly-restrictive constraints can cause problems.

Quantiles

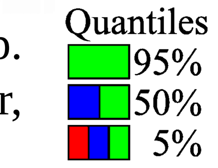

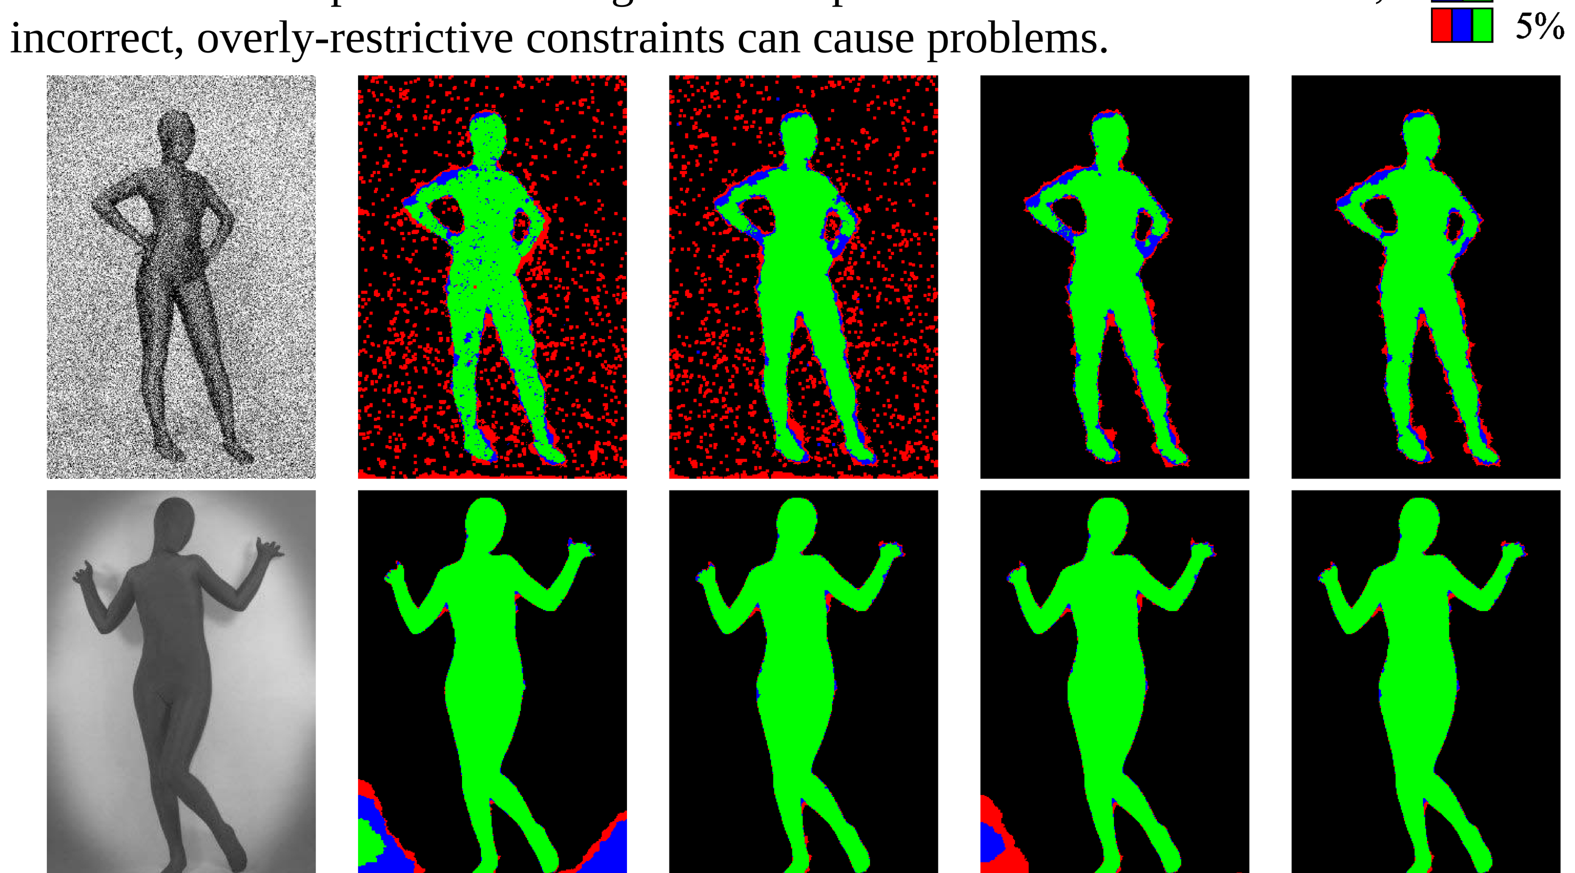

[1] Chang, J. and J. W. Fisher III. Analysis of Orientation and Scale in Smoothly Varying Textures. In: *ICCV*, 2009.
[2] Chang, J. and J. W. Fisher III. Efficient MCMC Sampling with Implicit Shape Representations. In: *CVPR*, 2011.
[3] Chen, S. and R. J. Radke. Markov Chaing Monte Carlo Shape Sampling 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] Han, X., C. Xu, and J. L. Prince. A Topology Preserving Level Set Method for Geometric Deformable Models. In: *IEEE Transacations on PAMI*, 2003.
[6] Heiler M. and C. Schnorr. Natural Image Statistics for Natural Image Segmentation. In: *ICCV*, 2003.
[7] Houhou, N., Thiran, J.P., Bresson, X.: Fast texture segmentation model based on the shape operator and active contour. In: *CVPR*, 2008.
[8] 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: *IEEE Transactions on IP*, 2005.
[9] 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.
[10] 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).
[12] Segonne, F. Active ContoursUnder Topology Control – Genus Preserving Level Sets. In: *Int. Journal of Computer Vision*, 2008.