
Image Segmentation with Pseudo-marginal MCMC Sampling and Nonparametric Shape Priors

Ertunc Erdil^{1,2*} Sinan Yıldırım¹ Tolga Tasdizen³ Müjdat Çetin^{1,4}

¹ Faculty of Engineering and Natural Sciences, Sabanci University, Istanbul, Turkey

² ARM Ltd., 1 Summerpool Road, Loughborough, Leicester, UK

³ Department of Electrical and Computer Engineering, University of Utah, Salt Lake City, UT, USA

⁴ Department of Electrical and Computer Engineering, University of Rochester, Rochester, NY, USA
{ertuncerdil, sinanyildirim, mcetin}@sabanciuniv.edu tolga@sci.utah.edu

Abstract

In this paper, we propose an efficient pseudo-marginal Markov chain Monte Carlo (MCMC) sampling approach to draw samples from posterior shape distributions for image segmentation. The computation time of the proposed approach is independent from the size of the training set used to learn the shape prior distribution nonparametrically. Therefore, it scales well for very large data sets. Our approach is able to characterize the posterior probability density in the space of shapes through its samples, and to return multiple solutions, potentially from different modes of a multimodal probability density, which would be encountered, e.g., in segmenting objects from multiple shape classes. Experimental results demonstrate the potential of the proposed approach.

1 Introduction

Incorporating prior shape density into the segmentation process has been widely studied in the literature [20], [29], [25], [2], [30], [18]. These methods can usually handle Gaussian-like, unimodal shape prior densities. Kim et al. [21] and Cremers et al. [8] incorporate nonparametric density estimation based shape priors into the segmentation process using level sets. Therefore, these methods and their variants can learn “multimodal” shape densities, which can be encountered in problems involving shape densities containing multiple classes of shapes [14], [6], [31], [28], [9], [23], [11], [24]. These methods minimize an energy function and find a solution at a local optimum. This does not provide any measure of the degree of confidence/uncertainty in that result and any information about the characteristics of the posterior density.

There are a limited number of Markov chain Monte Carlo (MCMC) based image segmentation methods in the literature. Most of these methods generate samples from the posterior density by assuming the prior density is uniform [13], [4], [5]. The only sampling-based segmentation approach in the literature that uses a shape prior is the one proposed by Chen et al. [7]. However, the method cannot handle topological changes in shapes.

Our major contribution in this paper is a pseudo-marginal Markov chain Monte Carlo (MCMC) sampling-based image segmentation approach that exploits nonparametric shape priors. The proposed approach is able to characterize the posterior density through its samples. Our approach can learn from very large data sets efficiently by using pseudo-marginal sampling. To the best of our knowledge, this is the first approach that performs pseudo-marginal MCMC shape sampling-based image segmentation through an energy functional that uses nonparametric shape priors and level

*This work has been supported by the Scientific and Technological Research Council of Turkey (TUBITAK) under Grant 113E603.

sets. Also, unlike other MCMC sampling-based segmentation approaches in the literature, the proposed approach perfectly satisfies the necessary conditions to implement MCMC sampling which is a crucial step for developing an MCMC sampler.

A precursor to this work was presented in [12]. This paper advances that prior work in two major ways: (1) while [12] approximately satisfies the necessary conditions of MCMC sampling, the approach presented in this paper perfectly satisfies these conditions; (2) unlike [12] we use pseudo-marginal sampling to be able to learn shape densities from very large data sets; the method in [12] becomes inefficient with large training data.

2 Model and problem definition

The image segmentation problem involves estimating an unknown segmenting curve for an object given an observed image $y \in \mathcal{Y}^{M \times N}$ where \mathcal{Y} is the set of the values that the pixels of y can take. We are particularly interested in problems in which the prior shape distribution has components corresponding to different object classes. We denote the class of the object by $s \in \{1, \dots, n\}$ where $n \geq 1$ is the total number of classes, which is known. For simplicity we assume that s has a uniform distribution over $\{1, \dots, n\}$ so that $p(s) = 1/n$, $s = 1, \dots, n$.

We ultimately aim to estimate a binary segmenting curve $c \in \{0, 1\}^{M \times N}$ where 0's indicate the background and 1's indicate the object. The conditional density of y given c is independent from s and is denoted by $p_{Y|C}(y|c)$. We construct this density based on the piecewise-constant version of the Mumford-Shah functional [26], [3] which is a very common data fidelity term for image segmentation. To present the curves, we use level sets, which we define as a mapping $\phi : \{0, 1\}^{M \times N} \rightarrow \mathbb{R}^{M \times N}$. Our choice of level sets is to handle topological changes and use gradient flows effectively in our methodology. We denote the level set of c as $x = \phi(c)$.

We also have a training set of binary curves that are grouped into classes, $\mathbf{C} = \{\mathcal{C}_1, \dots, \mathcal{C}_n\}$, where each $\mathcal{C}_i = \{c_{i,1}, \dots, c_{i,m_i}\}$ is the collection of $m_i \geq 1$ segmented curves for class i . Let us also define the level set representation of the training set as $\mathbf{X} = \{\mathcal{X}_1, \dots, \mathcal{X}_n\}$, where each $\mathcal{X}_i = \{x_{i,1}, \dots, x_{i,m_i}\}$ with $x_{i,j} = \phi(c_{i,j})$, the level set representation of $c_{i,j}$. Now we can define the prior distribution for x given the class s as $p(x|s) = \frac{1}{m_s} \sum_{i=1}^{m_s} \mathcal{N}(x; x_{s,i}, \sigma^2 I)$, where $\mathcal{N}(x; \mu, \Sigma)$ is a Gaussian density with mean μ and covariance matrix Σ . This prior corresponds to a mixture of kernels with centers $x_{s,1}, \dots, x_{s,m_s}$ with kernel size σ [21]. To determine the kernel size σ , we use an ML kernel with leave-one-out [27]. Let us also define $\bar{\phi}$ as the pseudo-inverse² of ϕ such that $\bar{\phi}(\phi(c)) = c$. We use $\bar{\phi}$ to rewrite the conditional density of the data in terms of x as $p(y|x) = p_{Y|C}(y|\bar{\phi}(x))$. We can also write the joint density of s , x , and y as $p(s, x, y) = p(s)p(x|s)p(y|x)$.

The Bayesian image segmentation problem can be formulated as finding the posterior distribution of x given y which is $p(x|y) \propto p(y|x)p(x) = p(y|x) \sum_{s=1}^n p(s)p(x|s)$. However, estimating $p(x|y)$ can be difficult since the summation over classes makes the distribution hard to infer, e.g., using Monte Carlo sampling methods. Alternatively, we aim for the joint posterior distribution of s and x given y , $p(s, x|y) \propto p(s, x, y)$, whose marginal is still the desired posterior $p(x|y)$.

3 Methodology

3.1 Metropolis-Hastings within Gibbs

We aim to sample from $p(s, x|y)$ using Gibbs sampling by sampling from the conditional densities $p(s|y, x)$ and $p(x|y, s)$ in an alternating fashion [16], [15]. However, since the full conditional $p(x|s, y)$ is hard to sample from, we update x by using a Metropolis-Hastings (MH) move, which leads to the well known Metropolis-Hastings within Gibbs (MHwG) algorithm [17].

Both conditional densities in MHwG involve $p(x|s)$ which needs to be evaluated during MH updates. This can be too costly when m_s is large, which occurs when we have a big training set. Therefore, towards a more computationally efficient MCMC algorithm that scales with the training data size, we consider the following unbiased estimator of $p(x|s)$ via subsampling: $\hat{p}(x|s) = \frac{1}{\hat{m}_s} \sum_{j=1}^{\hat{m}_s} \mathcal{N}(x; x_{s,u_j}, \sigma^2 I)$ where $\{u_1, u_2, \dots, u_{\hat{m}_s}\} \subset \{1, 2, \dots, m_s\}$ is a set of sub-

²Note that ϕ is not invertible, so we define a pseudo-inverse.

samples generated via sampling without replacement and $\hat{m}_s \ll m_s$. This approximation of the prior leads to the approximation of the conditional posterior densities $\hat{p}(s|x, y) \propto p(s)\hat{p}(x|s)$ and $\hat{p}(x|s, y) \propto \hat{p}(x|s)p(y|x)$. Using this approximation does not generally guarantee that the Markov Chain has an equilibrium distribution that is exactly $p(x, s|y)$. To achieve a correct MH algorithm which uses the approximation $\hat{p}(x|s)$ and still targets $p(x, s|y)$, we adopt the pseudo-marginal MH algorithm of [1]; the details are described in the following section.

3.2 Pseudo-marginal MHwG

Assume that we have a non-negative random variable z such that given x and s , its conditional density given s and x , $g_{s,x}(z)$, satisfies $\int_0^\infty g_{s,x}(z)zdz = p(x|s)$. We choose z as an approximation to $p(x|s)$, in particular $z = \hat{p}(x|s)$ where $\hat{p}(x|s)$ is defined in Section 3.1, and its probability density $g_{s,x}(z)$ corresponds to the generation process of this approximation. (It will become clear that in fact we do not have to calculate $g_{s,x}(z)$ at all but we should be able to sample from it.) Note that z is a random variable since we generate $\{u_1, \dots, u_{\hat{m}_s}\}$ randomly when computing $\hat{p}(x|s)$. We can define the extended posterior density with the new variable z added as $p(x, s, z|y) \propto p(s)zg_{s,x}(z)p(y|x)$. When we integrate z out, we see that samples for s and x from $p(x, s, z|y)$ will admit the desired posterior $p(s, x|y): p(s)p(y|x) \int zg_{s,x}(z)dz = p(s)p(x|s)p(y|x)$. Now, the problem of generating samples from $p(s, x|y)$ can be replaced with the problem of generating samples from $p(x, s, z|y)$.

We propose a pseudo-marginal MHwG sampling procedure to generate samples from $p(x, s, z|y)$. Note the important remark that this algorithm also targets $p(x, s|y)$, hence $p(x|y)$ exactly. There are two major steps of the proposed approach which we explain in detail in the following: (1) condition the posterior on x and update s and z , (2) update x and z by conditioning the posterior on s .

Update step for the class s and z : The distribution from which we sample s and z in Metropolis-Hastings is $p(s, z|y, x)$. We can write this distribution as $p(s, z|y, x) = p(s, z|x) \propto p(s)zg_{s,x}(z)$. Since we regard the proposal mechanism as a joint update of s and z , the proposal generates (s', z') from the density $q(s'|s)g_{s',x}(z')$. Note that (s', z') denotes the candidate samples generated from the proposal distribution. We take $q(s'|s)$ as a uniform distribution $\mathcal{U}\{1, n\}$ and $z' = \hat{p}(x|s')$ [1]. Once s' and z' are sampled, they are either accepted with probability

$$\min \left\{ 1, \frac{p(s')z'g_{s',x}(z')q(s|s')g_{s,x}(z)}{p(s)zg_{s,x}(z)q(s'|s)g_{s',x}(z')} \right\} = \min \left\{ 1, \frac{p(s')z'q(s|s')}{p(s)zq(s'|s)} \right\}, \quad (1)$$

or the current values of s and z are kept. The probabilities in the MH ratio in (1) can be computed exactly which guarantees satisfying the necessary conditions to implement MCMC sampling.

Update step for the level set x and z : The next step is to sample x and z from the conditional density $p(x, z|y, s)$. To achieve this, we perform Metropolis-Hastings sampling as we use for sampling s' and z' . The conditional density $p(x, z|y, s)$ can be written as $p(x, z|s, y) \propto zg_{s,x}(z)p(y|x)$. Also, for joint sampling of candidates (x', z') we can write the proposal density as $q_{s,j}(x', z'|y) = q_{s,j}(x'|x, y)g_{s,x'}(z')$ where j is sampled uniformly from $\{1, m_s\}$, and $z' = \hat{p}(x'|s)$ is generated using subsampling. Then, the Metropolis-Hastings ratio can be computed as follows:

$$\min \left\{ 1, \frac{z'g_{s',x}(z')p(y|x')q_{s,j}(x|x', y)g_{s,x}(z)}{zg_{s,x}(z)p(y|x)q_{s,j}(x'|x, y)g_{s',x}(z')} \right\} = \min \left\{ 1, \frac{z'p(y|x')q_{s,j}(x|x', y)}{zp(y|x)q_{s,j}(x'|x, y)} \right\} \quad (2)$$

Design of the proposal distribution: The crucial part in (2) is designing the proposal distribution to generate a candidate curve x' from x . Let us define an energy function whose gradient drives the curve to the desired location using the data and the training images in a particular class s as $E_s(x) := \log p(x|s) + \log p(y|x, s)$. When the training data set is too large, calculating $\log p(x|s)$ may be too expensive as discussed earlier. An unbiased estimator of $E_s(x)$ would be obtained as $E_{s,j}(x) := \log \mathcal{N}(x; x_{s,j}, \sigma^2 I) + \log p(y|x, s)$ where $j \sim \mathcal{U}(1, \dots, m_s)$. The proposal distribution after sampling the j^{th} training image in class s is then constructed as $q_{s,j}(x'|x, y) = \mathcal{N}(x'; x - \widehat{\nabla E}_{s,j}(x), \Sigma)$. Here, the shift term $\widehat{\nabla E}_{s,j}(x)$ is an approximation³ to the gradient $\nabla E_{s,j}(x)$ w.r.t. x and is given by $\widehat{\nabla E}_{s,j}(x) = \frac{1}{\sigma^2}(x - x_{s,j}) + [(y - \mu_{\text{in}})^{\odot 2} - (y - \mu_{\text{out}})^{\odot 2}]$ where $(\cdot)^{\odot k}$ is the element-wise power operation.

³Note that the term $(y - \mu_{\text{in}})^{\odot 2} - (y - \mu_{\text{out}})^{\odot 2}$ is a discrete approximation w.r.t. the level set representation x [3], so that the whole expression is an approximation to $\nabla E_{s,j}(x)$.

In the design of $q_{s,j}(x'|x, y)$, the most important part is selecting a covariance matrix, Σ , that generates smooth perturbations since smoother curves are more likely [21], [13]. In the proposed approach, we compute a positive semi-definite covariance matrix Σ such that it generates smooth random perturbations. Given an $M \times N$ image, we first generate an $MN \times MN$ matrix Z by drawing i.i.d. samples from a unit Gaussian distribution. Then, we construct another $MN \times MN$ matrix F where each column of F is a smoothed version of each row in Z . By assuming F is constructed by multiplying Z by a matrix \hat{A} , we can find the matrix \hat{A} as $\hat{A} = Z^{-1}F$ since Z is generally invertible. Given \hat{A} , a covariance matrix $\hat{\Sigma}$ can be computed by $\hat{\Sigma} = \hat{A}\hat{A}^T$. However, generally $\hat{\Sigma}$ is not positive semi-definite since \hat{A} is not generally a full rank matrix. Therefore, we find the closest positive semi-definite matrix to $\hat{\Sigma}$ using the approach in [19] which we take it as Σ .

4 Experimental results

We compare running time of the proposed pseudo-marginal sampling approach and conventional sampling approach which uses all training examples to estimate the prior shape density. We perform experiments on the MNIST data set [22] which contains 60,000 training examples. We construct training sets with various sizes from 1K to 50K by randomly selecting equal number of samples from each digit class. We generate 1000 samples using both the pseudo-marginal sampling and conventional sampling on a test image by using training sets with different sizes. The plot in Figure 1 shows average running time as a function of training set size for both pseudo-marginal sampling and conventional MCMC sampling. The average single sample generation time of the proposed approach does not change as the training set size increases since we choose $\hat{m}_s = 10$ in all experiments. We also measure the segmentation accuracy of all samples with the ground truth using Dice score [10]. Average Dice score⁴ result for pseudo-marginal sampling is 0.7718 whereas it is 0.7777 for the conventional sampling. The very slight decrease in Dice results of pseudo-marginal sampling can be acceptable in many applications when the huge gain in running time is considered.





Figure 1: Average running time for producing a single sample as a function of training set size.

We also run the proposed approach on the test image in Figure 2(a) and generate 1000 samples. The algorithm generates samples from digit classes 4, 7, and 9 as shown in Figure 2(c). The optimization-based approach of Kim et al. [21] finds a single segmentation solution shown in Figure 2(b). This experiment shows that our algorithm can produce samples from different modes.

5 Conclusion

We propose a pseudo-marginal Markov chain Monte Carlo (MCMC) sampling-based image segmentation approach that exploits nonparametric shape priors. The proposed approach generates samples from the posterior distribution $p(x|y)$ to avoid shortcomings of the optimization-based approaches which include getting stuck at local optima and being unable to characterize the posterior density. The proposed MCMC sampling approach deals with all these problem while being computationally efficient, unlike the conventional MCMC approaches, by using pseudo-marginal sampling principles. Moreover, our pseudo-marginal shape sampler perfectly satisfies the necessary conditions to implement MCMC sampling which is crucial for ensuring the generated samples come from the desired distribution. Existing methods in the literature only approximately satisfy these conditions.

⁴Note that Dice score takes values in $[0, 1]$ where 1 indicates the perfect match with the ground truth.



(a) Test image (b) Kim et al. [21] (c) The proposed approach

Figure 2: Marginal confidence bounds obtained by samples on a test image. Note that in (c) red indicates the least and blue indicates the most confidence boundaries.

References

- [1] Andrieu, C., Roberts, G.O.: The pseudo-marginal approach for efficient Monte Carlo computations. *The Annals of Statistics* pp. 697–725 (2009)
- [2] de Bruijne, M., van Ginneken, B., Viergever, M.A., Niessen, W.J.: Adapting active shape models for 3D segmentation of tubular structures in medical images. In: *Information Processing in Medical Imaging*, pp. 136–147. Springer (2003)
- [3] Chan, T.F., Vese, L., et al.: Active contours without edges. *IEEE transactions on Image processing* **10**(2), 266–277 (2001)
- [4] Chang, J., Fisher, J.: Efficient MCMC sampling with implicit shape representations. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2081–2088. IEEE (2011)
- [5] Chang, J., Fisher III, J.W.: Efficient topology-controlled sampling of implicit shapes. In: *Proceedings of the IEEE International Conference on Image Processing (ICIP)* (2012)
- [6] Chen, S., Radke, R.J.: Level set segmentation with both shape and intensity priors. In: *International Conference on Computer Vision*, pp. 763–770. IEEE (2009)
- [7] Chen, S., Radke, R.J.: Markov chain Monte Carlo shape sampling using level sets. In: *IEEE 12th International Conference on Computer Vision Workshops (ICCV Workshops)*, pp. 296–303. IEEE (2009)
- [8] Cremers, D., Osher, S.J., Soatto, S.: Kernel density estimation and intrinsic alignment for shape priors in level set segmentation. *International Journal of Computer Vision* **69**(3), 335–351 (2006)
- [9] Cremers, D., Rousson, M., Deriche, R.: A review of statistical approaches to level set segmentation: integrating color, texture, motion and shape. *International journal of computer vision* **72**(2), 195–215 (2007)
- [10] Dice, L.R.: Measures of the amount of ecologic association between species. *Ecology* **26**(3), 297–302 (1945)
- [11] Erdil, E., Ghani, M.U., Rada, L., Argunsah, A.O., Unay, D., Tasdizen, T., Cetin, M.: Nonparametric joint shape and feature priors for image segmentation. *IEEE Transactions on Image Processing* (2017)
- [12] Erdil, E., Yildirim, S., Cetin, M., Tasdizen, T.: MCMC shape sampling for image segmentation with nonparametric shape priors. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 411–419 (2016)
- [13] Fan, A.C., Fisher III, J.W., Wells III, W.M., Levitt, J.J., Willsky, A.S.: MCMC curve sampling for image segmentation. In: *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pp. 477–485. Springer (2007)
- [14] Foulonneau, A., Charbonnier, P., Heitz, F.: Multi-reference shape priors for active contours. *International journal of computer vision* **81**(1), 68–81 (2009)
- [15] Gelfand, A.E., Smith, A.F.: Sampling-based approaches to calculating marginal densities. *Journal of the American statistical association* **85**(410), 398–409 (1990)
- [16] Geman, S., Geman, D.: Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Transactions on pattern analysis and machine intelligence* (6), 721–741 (1984)
- [17] Geweke, J., Tanizaki, H.: Bayesian estimation of state-space models using the Metropolis–Hastings algorithm within Gibbs sampling. *Computational Statistics & Data Analysis* **37**(2), 151–170 (2001)
- [18] Heimann, T., Meinzer, H.P.: Statistical shape models for 3D medical image segmentation: a review. *Medical Image Analysis* **13**(4), 543–563 (2009)
- [19] Higham, N.J.: Computing a nearest symmetric positive semidefinite matrix. *Linear Algebra and its Applications* **103**, 103 – 118 (1988)
- [20] Kass, M., Witkin, A., Terzopoulos, D.: Snakes: Active contour models. *International Journal of Computer Vision* **1**(4), 321–331 (1988)
- [21] Kim, J., Çetin, M., Willsky, A.S.: Nonparametric shape priors for active contour-based image segmentation. *Signal Processing* **87**(12), 3021–3044 (2007)

- [22] LeCun, Y., Bottou, L., Bengio, Y., Haffner, P.: Gradient-based learning applied to document recognition. *Proceedings of the IEEE* **86**(11), 2278–2324 (1998)
- [23] Mesadi, F., Cetin, M., Tasdizen, T.: Disjunctive normal parametric level set with application to image segmentation. *IEEE Transactions on Image Processing* **26**(6), 2618–2631 (2017)
- [24] Mesadi, F., Erdil, E., Cetin, M., Tasdizen, T.: Image segmentation using disjunctive normal Bayesian shape and appearance models. *IEEE Transactions on Medical Imaging* **PP**(99), 1–1 (2017)
- [25] Milborrow, S., Nicolls, F.: Locating facial features with an extended active shape model. In: *European Conference on Computer Vision (ECCV)*, pp. 504–513. Springer (2008)
- [26] Mumford, D., Shah, J.: Optimal approximations by piecewise smooth functions and associated variational problems. *Communications on pure and applied mathematics* **42**(5), 577–685 (1989)
- [27] Silverman, B.W.: *Density estimation for statistics and data analysis*, vol. 26. CRC press (1986)
- [28] Soğanlı, A., Uzunbaş, M.G., Çetin, M.: Combining learning-based intensity distributions with nonparametric shape priors for image segmentation. *Signal, Image and Video Processing* **8**(4), 789–798 (2014)
- [29] Van Ginneken, B., Frangi, A.F., Staal, J.J., Romeny, B.M., Viergever, M.A.: Active shape model segmentation with optimal features. *IEEE Transactions on Medical Imaging* **21**(8), 924–933 (2002)
- [30] Wang, W., Shan, S., Gao, W., Cao, B., Yin, B.: An improved active shape model for face alignment. In: *IEEE International Conference on Multimodal Interfaces*, p. 523 (2002)
- [31] Yang, R., Mirmehdi, M., Xie, X., Hall, D.: Shape and appearance priors for level set-based left ventricle segmentation. *IET Computer Vision* **7**(3), 170–183 (2013)