Int J Comput Vis (2009) 83: 195–209
DOI 10.1007/s11263-008-0189-6

Stereo Matching Using Population-Based MCMC

Wonsik Kim · Joonyoung Park · Kyoung Mu Lee

Received: 7 February 2008 / Accepted: 14 October 2008 / Published online: 31 October 2008

Abstract In this paper, we propose a new stereo matching method using the population-based Markov Chain Monte Carlo (Pop-MCMC), which belongs to the sampling-based methods. Since the previous MCMC methods produce only one sample at a time, only local moves are available. In contrast, the proposed Pop-MCMC uses multiple chains in parallel and produces multiple samples at a time. It thereby enables global moves by exchanging information between samples, which in turn, leads to faster mixing rate. In the view of optimization, it means that we can reach a lower energy state rapidly. In order to apply Pop-MCMC to the stereo matching problem, we design two effective 2-D mutation and crossover moves among multiple chains to explore a high dimensional state space efficiently. The experimental results on real stereo images demonstrate that the proposed algorithm gives much faster convergence rate than conventional sampling-based methods including SA (Simulated Annealing) and SWC (Swendsen-Wang Cuts). And it also gives consistently lower energy solutions than BP (Belief Propagation) in our experiments. In addition, we also analyze the effect of each move in Pop-MCMC and examine the effect of parameters such as temperature and the number of the chains.

W. Kim · K.M. Lee (✉)
Department of EECS, ASRI, Seoul National University,
151-742 Seoul, Korea
e-mail: kyoungmu@snu.ac.kr

W. Kim
e-mail: ultra16@snu.ac.kr

J. Park
DM research Lab., LG Electronics Inc., 16 Woomyeon-Dong,
Seocho-Gu, 137-724 Seoul, Korea
e-mail: kzoome@lge.com

Keywords Stereo · Population-based MCMC · Energy minimization

1 Introduction

Stereo matching is one of the classical problems in computer vision (Scharstein and Szeliski 2002). The goal of stereo matching is to determine disparities, which are the differences in the horizontal coordinates of corresponding pixels in a pair of rectified images. If we obtain an accurate disparity map, we can recover 3-D scene information. However, it remains challenging problem because of occluded regions, noise of camera sensor, textureless regions, etc. Stereo matching algorithms can be classified into two approaches. One is the local approach, and the other is the global approach. In the local approach, disparities are determined by comparing the intensity values in local windows by some measures such as Sum of Absolute Differences (SAD), Sum of Squared Differences (SSD), and Birchfield-Tomasi measure (Birchfield and Tomasi 1998). Although local approaches are fast, they have difficulties in obtaining an accurate disparity map due to some intrinsic difficulties including the noise and choosing appropriate window size problems. To overcome those difficulties, in general, global approaches employ the smoothness constraint that reflects the smoothly varying surface assumption. Usually, an energy function that is composed of local (likelihood) and global (prior) constraint is defined and solved by various energy minimization techniques. Recently, deterministic methods such as dynamic programming (Ohta and Kanade 1985; Veksler 2005), Graph cuts (Boykov et al. 2001; Kolmogorov and Zabih 2001, 2002, 2004; Boykov and Kolmogorov 2004), message passing method and its variants (Yedidia

et al. 2000; Sun et al. 2003; Kolmogorov 2006), and linear programming-based methods (Komodak et al. 2007a, 2007b) are widely used for the global optimization and achieve good performances (Szeliski et al. 2006).

Graph cut methods are fast and provides very low energy solution with standard 4-neighborhood benchmark problems. However, it can be applied to a limited class of energy functions (Kolmogorov and Zabih 2004). Szeliski et al. (2006) showed that α-expansion move method was faster and slightly better than α-β-swap move method in all cases in their experiments. However, α-expansion move method can be applied to more limited class of energy functions than α-β-swap move. Belief Propagation (BP) is a message passing method originally developed for graphs without cycles. In general, although it is not guaranteed to converge, it has been successfully applied to loopy graphs (Sun et al. 2003). Tree-reweighted message passing (TRW) is also a message passing method (Kolmogorov 2006). It finds lower energy solution than Graph cuts in many problems. An important property of TRW is that it gives a lower bound on the energy which can be used to check how close our solution to the global minimum energy. All of the deterministic methods are approximation algorithms. Although Graph cuts provides global minimum for some restricted energy models, none of these methods guarantee to obtain the global minimum solution for a general stereo model in practical time since it is known to be an NP hard problem.

In contrast to the deterministic methods, stochastic approaches such as sampling-based methods can be used to find global optimum. Sampling-based methods were originally developed to generate samples from a given target distribution or to integrate functions in high dimensional space. These Sampling-based methods are also have been used for statistical estimation and optimization. In this paper, we use a sampling-based method for energy minimization to solve the stereo matching problem.

The Monte-Carlo method is the most primitive sampling-based method. In this method, a new sample is drawn depending on a pre-determined proposal distribution. This distribution is independent on previous samples. However, there are some difficulties in applying the Monte Carlo methods to vision problems as an optimizer. In general, we need to solve vision problems in very high-dimensional solution spaces. Even if it is assumed to be 100 pixels in the width and height, respectively, the dimension of the image space can be as high as 10^4. Monte Carlo methods would take infinitely long time since the acceptance rate would be almost zero in such a high-dimensional case. Moreover, we need to design a proper proposal distribution close to the target distribution. To resolve these problems, Markov Chain Monte Carlo (MCMC) methods had been tried. In MCMC, a new sample is drawn from the previous sample with a local transition probability, based on the Markov chain. Contrary to simple Monte Carlo methods, the acceptance rates of MCMC methods are high enough, and the proposal distributions are designable even in high-dimensional problems. Therefore, MCMC methods are more appropriate for the application to vision problems than the Monte Carlo methods. However, difficulties still remain in applying MCMC to vision problem as an optimizer. Since most MCMC methods allow only local moves in a large solution space, it still takes very long time to reach the global optimum.

To overcome the limitations of MCMC methods as an optimizer, recently Swendsen-Wang Cuts (SWC) was proposed (Barbu and Zhu 2004, 2005). In SWC, it is shown that bigger local moves are possible than in previous methods while maintaining the detailed balance. SWC uses Simulated Annealing (SA) (Kirkpatrick et al. 1983) to find the global optimum. Although SWC allows bigger local moves, a very slow annealing process is needed to approach the global optimum with probability 1. This is an apparent drawback of SWC. Therefore, we need a faster annealing process for real vision applications. However, fast annealing does not always guarantee the global optimum and the samples are often trapped in local optima.

In this paper, we propose a new MCMC method called Population-Based MCMC (Pop-MCMC) (Liang and Wong 2000; Jasra and Stephens 2007) that can overcome the drawbacks of SWC for stereo matching problem. Our goal is to obtain the lower energy state faster than other sampling methods including SWC which have been previously applied to this problem. In Pop-MCMC, two or more samples are drawn at the same time. Samples can exchange information with each other. This makes it possible to perform global moves of samples. It means that the mixing rate of drawn samples becomes faster. And in the view of optimization, the faster mixing rate means that it takes shorter time for the samples to approach the global optimum than conventional methods. A preliminary version of this paper has been appeared in Park et al. (2007).

This paper presents the design of Pop-MCMC for stereo matching and comparison of performance with the other methods such as SA, SWC and BP. The paper is organized as follows: In Sect. 2, SWC and Pop-MCMC are briefly introduced. Then, we present how Pop-MCMC is applied to stereo matching in Sect. 3. Section 4 gives the experimental results. In the final Sect. 5, we conclude the paper with discussions.

2 Related Works

In this section, we first describe the SWC, which has been applied to vision problems (Barbu and Zhu 2004, 2005). Then, we present the description of Pop-MCMC.

2.1 Swendsen-Wang Cuts (SWC)

SWC originated from the Swendsen-Wang (SW) method. Swendsen and Wang proposed SW method in 1987 (Swendsen and Wang 1987). It overcame the slow convergence of previous sampling-based methods such as Gibbs sampler. Let us explain SW briefly.

We consider a 2-D lattice graph $G = \langle V, E \rangle$, where V is the set of nodes and E is the set of edges connecting neighboring nodes. Each node $v_i \in V$ is assigned a label $x_i \in \{1, 2, \ldots, L\}$. The number of possible labels is L. In a 2-D lattice graph, each node has four edges. We assume that this graph follows the Potts model, which is often used in vision as a prior model. The formulation of Potts model is as follows.

$$\pi(\mathbf{X}) = \frac{1}{Z} \exp \beta \sum_{\langle i,j \rangle \in E} \mathbf{1}(x_i = x_j), \tag{1}$$

where $\mathbf{X}$ represents $(x_1, \ldots, x_N)$ and N is the number of nodes. Z and β are constants. $\mathbf{1}(\cdot)$ represents a Boolean function. When the graph follows the Potts model, a global minimum should be the states in which all the nodes have the same labels.

In Gibbs sampler, the label of only one node can be flipped to generate the next sample. So it needs a generation of $O(L^N)$ samples to reach the global optimum. In contrast, in SW the labels of a cluster of nodes are flipped at the same time.

However SW has several drawbacks. It assumes fixed number of labels, and does not create new labels in the case when the number of labels is unknown. And it is only applicable to Ising/Potts model. In addition, it does not consider the external field, such as the observed visual data in vision.

To overcome the above limitations of SW, SWC has been proposed by extending SW from the Metropolis-Hastings perspective (Barbu and Zhu 2004, 2005). SWC can be applicable to arbitrary posterior probabilities, and can incorporate external data easily. The summary of SWC is described in the following:

Assume that a current state is A, repeat the process below.

1. If the labels of two neighboring nodes s and t are different, the edge connecting two nodes is removed. If the labels are the same, we determine whether the edge is retained or not with the probability q_e. If there exists external field, we consider it in designing the probability q_e. This process is repeated for all edge $e = \langle s, t \rangle \in E$. Then nodes connected by remaining edges are considered as a cluster.
2. One cluster V_0 is randomly selected.
3. New label l' of the chosen cluster V_0 is proposed with a proposal distribution $q(l'|V_0, A)$.
4. Determine whether we accept the newly generated sample (or state B) with acceptance probability α by the following Metropolis-Hastings rule

$$\alpha = \min\left(1, \frac{q(V_0 \mid B) q(l \mid V_0, B) \pi(B \mid I)}{q(V_0 \mid A) q(l' \mid V_0, A) \pi(A \mid I)}\right), \tag{2}$$

where I represents the external field, that is, the observed input image. No matter how q_e and the proposal distribution $q(l'|V_0, A)$ are designed, the detailed balance is maintained by Metropolis-Hastings kernel. Therefore, we can appropriately design q_e and the proposal distribution of the new label freely, so as to use the information of input image properly.

In order to reduce the complexity of SWC, a modified clustering method, SWC-2 has been proposed (Barbu and Zhu 2005). In SWC-2, a connected node cluster V_0 is determined by a recursive method as described in the following.

1. Select a seed node $v \in V$ randomly, and assign it to a cluster V_0.
2. Repeat until no more node is added to V_0.
   For any edge $e = \langle s, t \rangle \in E$ between the node $s \in V_0$ and its neighboring node $t \notin V_0$,
   (a) If the labels of the two nodes s and t are deferent, remove the edge. Otherwise, determine whether the edge e should be retained or not with probability q_e, same as in SWC.
   (a) If the edge e is not removed, add the node t to the cluster V_0.

Note that in constructing V_0, we need to calculate q_e only for the edges at the border of the cluster V_0. It leads to saving of computational costs. In our work, we adopt SWC-2 as a part of the proposed algorithm.

2.2 Population-Based MCMC (Pop-MCMC)

Pop-MCMC or evolutionary Monte Carlo is a stochastic simulation method that combines a population of Metropolis-Hastings samplers and Evolutionary Algorithm to improve the performance of MCMC samplers. Pop-MCMC generates multiple chains in parallel with different temperatures, and exchanges information among them to accelerate the mixing rate. This method can be considered as a variant of the Parallel Tempering (PT), that was proposed by Geyer (1991) and modified by others later (Hukushima and Nemoto 1996). PT aims to overcome the problems of traditional single process MCMC using a Metropolis-Hastings update, which has low mixing rate. The basic idea of PT is to simulate multiple replicas of the original system in parallel at a series of different temperatures, and swap the configurations with a Metropolis-Hastings criterion. The target

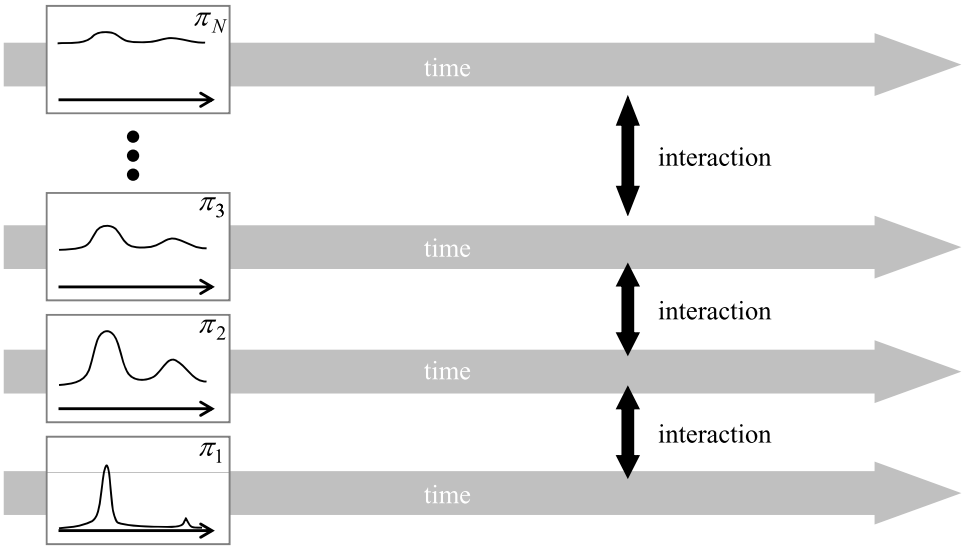

Fig. 1 Chains in Parallel Tempering

distribution of ith chain is defined as follows.

$$\pi_i(\mathbf{X}) = \pi(\mathbf{X})^{\frac{1}{T_i}}, \tag{3}$$

where $\pi(\mathbf{X})$ is an original target distribution, and T_i is the temperature of the ith chain. In the chain with high temperature, the target distribution is nearly flat as depicted in Fig. 1, where the heights of barriers between local optima are very low. Therefore, the samples in such chain can freely wander in contrast to the samples in a chain with low temperature. By exchanging these higher-temperature configurations with the configuration of a low temperature of our interest, we can allow the low temperature simulation to sample configurations much more efficiently than with local Metropolis updates only. This leads to a faster mixing rate between samples, and helps to escape from local minima.

Pop-MCMC allows chains to exchange information more actively than PT by introducing a new move called the *crossover* move. It originated from the genetic algorithm and then modified to fit the MCMC framework (Liang and Wong 2000). In Pop-MCMC, the Markov chain state is augmented as the population of all chains. Given an original target distribution $\pi(\mathbf{X})$, a new expanded target distribution is defined as follows.

$$\pi^*(\mathbf{X}_{1:N}) = \prod_{i=1}^{N} \pi_i(\mathbf{X}_i), \tag{4}$$

where N is the number of chains to use. We assume that $\pi_k \equiv \pi$ for at least one chain $k \in \{1, \ldots, N\}$. $\mathbf{X}_{1:N} = \{\mathbf{X}_1, \ldots, \mathbf{X}_N\}$ is a population composed of samples of N chains. Each component $\mathbf{X}_i$ in the vector $\mathbf{X}_{1:N}$ is called as a *chromosome*. The term chromosome is borrowed from genetic algorithm. The goal of Pop-MCMC is to generate samples $\mathbf{X}_{1:N}$ which follow the new target distribution π^*. And a collection of chromosomes from the kth chain, which has the target distribution $\pi_k = \pi$, is what we want to obtain finally.

Pop-MCMC has three different types of moves; mutation, exchange and crossover, which are described below in detail.

Mutation Move The mutation move updates a chromosome of a single chain using a Markov kernel, while other chains keep unchanged. We can use a conventional MCMC algorithm. Let us suppose that the current population is $\mathbf{X}_{1:N} = \{\mathbf{X}_1, \ldots, \mathbf{X}_i, \ldots, \mathbf{X}_N\}$. Among N chains, we randomly select ith chain and generate a new chromosome $\mathbf{Y}_i$ from the current chromosome $\mathbf{X}_i$ by an MCMC algorithm. Then, a new population $\mathbf{Y}_{1:N} = \{\mathbf{X}_1, \ldots, \mathbf{Y}_i, \ldots, \mathbf{X}_N\}$ is proposed, and is accepted according to the Metropolis-Hastings rule with probability

$$\alpha = \min(1, \gamma_m), \tag{5}$$

where

$$\begin{aligned}\gamma_m &= \frac{\pi^*(\mathbf{Y}_{1:N})}{\pi^*(\mathbf{X}_{1:N})} \cdot \frac{T(\mathbf{Y}_{1:N} \to \mathbf{X}_{1:N})}{T(\mathbf{X}_{1:N} \to \mathbf{Y}_{1:N})} \\ &= \frac{\pi_i(\mathbf{Y}_i)}{\pi_i(\mathbf{X}_i)} \cdot \frac{T(\mathbf{Y}_i \to \mathbf{X}_i)}{T(\mathbf{X}_i \to \mathbf{Y}_i)},\end{aligned} \tag{6}$$

where T denotes the transition probability between populations. In short, in the mutation move an MCMC move is performed at a specific chain independently, while other chains are kept unchanged. The irreducibility of Pop-MCMC is guaranteed by this mutation move.

Exchange Move The exchange move is the same as that used in PT. In this move, two different chains are randomly chosen first. And then the chromosomes of those

two chains are exchanged to propose a new population. Let us suppose that the current population is $\mathbf{X}_{1:N} = \{\mathbf{X}_1, \ldots, \mathbf{X}_i, \ldots, \mathbf{X}_j, \ldots, \mathbf{X}_N\}$, and the ith and jth chains are selected. Then, the newly proposed population will be $\mathbf{Y}_{1:N} = \{\mathbf{X}_1, \ldots, \mathbf{X}_j, \ldots, \mathbf{X}_i, \ldots, \mathbf{X}_N\}$. Similar to the mutation move, the new population is accepted according to the acceptance probability:

$$\alpha = \min(1, \gamma_e), \tag{7}$$

$$\gamma_e = \frac{\pi^*(\mathbf{Y}_{1:N})}{\pi^*(\mathbf{X}_{1:N})} \cdot \frac{T(\mathbf{Y}_{1:N} \to \mathbf{X}_{1:N})}{T(\mathbf{X}_{1:N} \to \mathbf{Y}_{1:N})} = \frac{\pi_i(\mathbf{X}_j)\pi_j(\mathbf{X}_i)}{\pi_i(\mathbf{X}_i)\pi_j(\mathbf{X}_j)}. \tag{8}$$

The last equality holds due to the definition of the target distribution and the symmetry property of the transition probability.

In general, to obtain the higher acceptance rate, exchange moves are performed on chains that have similar target distributions with neighboring temperatures.

Crossover Move The crossover move is newly introduced in Pop-MCMC. The main concept of this move is borrowed from the genetic algorithm. The design of this move is the main contribution of Pop-MCMC. There are several variations of the crossover moves. One of the popular moves is the 1-point crossover move. The basic idea of it is as follows: As in the exchange move, two different chains, say ith and jth chains, are randomly selected. If the chromosome is a d-D vector, we randomly choose a natural number k between 1 and $(d-1)$. And new chromosomes $\mathbf{Y}_i$ and $\mathbf{Y}_j$ are proposed by swapping the same part of chromosomes $\mathbf{X}_i$ and $\mathbf{X}_j$ as follows.

$$\begin{bmatrix} \mathbf{X}_i = (x_{i1}, \ldots, x_{ik}, x_{i(k+1)}, \ldots, x_{id}) \\ \mathbf{X}_j = (x_{j1}, \ldots, x_{jk}, x_{j(k+1)}, \ldots, x_{jd}) \end{bmatrix} \to \begin{bmatrix} \mathbf{Y}_i = (x_{i1}, \ldots, x_{ik}, x_{j(k+1)}, \ldots, x_{jd}) \\ \mathbf{Y}_j = (x_{j1}, \ldots, x_{jk}, x_{i(k+1)}, \ldots, x_{id}) \end{bmatrix}. \tag{9}$$

In this case, the ratio of proposal distributions in the acceptance probability is canceled by symmetry. We only need to calculate the ratio of the target distributions. A new population is proposed as $\mathbf{Y}_{1:N} = \{\mathbf{X}_1, \ldots, \mathbf{Y}_i, \ldots, \mathbf{Y}_j, \ldots, \mathbf{X}_N\}$, and according to the Metropolis-Hastings rule, it is accepted with probability

$$\alpha = \min(1, \gamma_c), \tag{10}$$

where

$$\gamma_c = \frac{\pi^*(\mathbf{Y}_{1:N})}{\pi^*(\mathbf{X}_{1:N})} \cdot \frac{T(\mathbf{Y}_{1:N} \to \mathbf{X}_{1:N})}{T(\mathbf{X}_{1:N} \to \mathbf{Y}_{1:N})} = \frac{\pi_i(\mathbf{Y}_i)\pi_j(\mathbf{Y}_j)}{\pi_i(\mathbf{X}_i)\pi_j(\mathbf{X}_j)} \cdot \frac{q(\mathbf{X}_i, \mathbf{X}_j|\mathbf{Y}_i, \mathbf{Y}_j)}{q(\mathbf{Y}_i, \mathbf{Y}_j|\mathbf{X}_i, \mathbf{X}_j)}, \tag{11}$$

where $T(\mathbf{X}_{1:N} \to \mathbf{Y}_{1:N})$ is $p(i, j|\mathbf{X}_{1:N}) \cdot q(\mathbf{Y}_i, \mathbf{Y}_j|\mathbf{X}_i, \mathbf{X}_j)$. $p(i, j|\mathbf{X}_{1:N})$ denotes the probability that ith and jth chains are chosen and $q(\mathbf{Y}_i, \mathbf{Y}_j|\mathbf{X}_i, \mathbf{X}_j)$ indicates the probability that the chromosomes $\mathbf{Y}_i$ and $\mathbf{Y}_j$ are proposed, when current chromosomes $\mathbf{X}_i$ and $\mathbf{X}_j$ are given. Choosing chains is independent of the current state, so $p(i, j|\mathbf{X}_{1:N})$ and $p(i, j|\mathbf{Y}_{1:N})$ are canceled out in the second equality.

In order to include various ways of exchanging information between chromosomes, the 2-point crossover move, k-point crossover move, and adaptive crossover move were also proposed in the literature (Liang and Wong 2000, 2001).

3 Proposed Algorithm

In this paper, we apply the Pop-MCMC method to stereo matching. For this purpose, we design new effective 2-D mutation and crossover moves to explore the high dimensional state space efficiently. In order to improve the accuracy of the disparity map, various energy models have been newly proposed for the stereo problem. Among them, we choose the segment-based energy model since it is known as one of the best energy models and it is robust to noise (Tao et al. 2001; Bleyer and Gelautz 2005; Hong and Chen 2004; Klaus et al. 2006). This model assumes that each segment corresponds to a planar patch in the scene. In a segment-based energy model, the reference image is first over-segmented. This segment-based energy model also reduces running time since the number of nodes is much smaller than pixel-based model. Mean-shift algorithm is often used for the segmentation (Comanicu and Meer 2002).

3.1 Segment-Based Stereo Energy Model

Each segment is defined as a node $v \in V$, and neighboring nodes s and t are connected with edges $\langle s, t\rangle \in E$. Then we construct a graph $\mathbf{G} = (V, E)$. And the energy function is defined by

$$E(\mathbf{X}) = \sum_{v \in V} C_{\mathrm{SEG}}(f_v) + \sum_{\langle s,t\rangle \in N} \beta_{s,t} \mathbf{1}(f_s \neq f_t), \tag{12}$$

where $\mathbf{X}$ represents the current state of every segment, f_v is an estimated plane for each segment, $C_{\mathrm{SEG}}(f_v)$ is a matching cost, and $\beta_{s,t}$ is a penalty for different neighboring nodes of s and t, which are defined by

$$C_{\mathrm{SEG}}(f_v) = \sum_{(x,y) \in V} C(x, y, f_v(x, y)), \tag{13}$$

$$\beta_{s,t} = \gamma \cdot BL(s, t) \cdot \mathcal{S}(s, t), \tag{14}$$

where function $C(x, y, f_v(x, y))$ is the Birchfield-Tomasi cost, $BL(s, t)$ is the shared border length, and $\mathcal{S}(s, t)$ is the mean color similarity defined by

$$\mathcal{S}(s,t) = \frac{1}{2}\left(1 - \min\left(1, \frac{|R_{V_s} - R_{V_t}| + |G_{V_s} - G_{V_t}| + |B_{V_s} - B_{V_t}|}{255}\right)\right) + \frac{1}{2}, \tag{15}$$

where R_{V_s}, G_{V_s} and B_{V_s} are average intensity values of segment V_s, which are between 0 and 255. Mean color similarity has a value between $\frac{1}{2}$ and 1. When two neighboring segments have similar intensities, it becomes closer to 1. By varying γ, we can control the relative effect of matching cost and smoothness cost.

We first need to make a list of the planes for assigning each segment to a plane by examining segment by segment. For each pixel, we calculate the initial disparity by using SAD (Sum of Absolute Differences) and WTA (Winner Takes All) schemes. Using these initial disparities, we fit a plane for each segment. The equation of a plane in 3D-space can be written by

$$d(x, y) = c_1 x + c_2 y + c_3, \tag{16}$$

where x and y are the coordinates of a pixel, and $d(x, y)$ is its disparity. Based on the above equation, we construct the following algebraic equation for each segment.

$$\mathbf{A}\,[c_1, c_2, c_3]^{\mathrm{T}} = \mathbf{B}, \tag{17}$$

where the ith row of the matrix $\mathbf{A}$ is the coordinates $[x_i, y_i, 1]$ of the ith pixel, and the ith row of the matrix $\mathbf{B}$ is the disparity $d(x_i, y_i)$ of that pixel. Then, the values of c_1, c_2, c_3 are obtained as a least squares solution by solving (17). In this method, the outlier disparities are initially detected and removed by a disparity crosschecking method (Hong and Chen 2004). Once we find the plane parameters, we can further identify more outlier disparities that are not close to the fitted plane. For those pixels with outlier disparities, we re-estimate the correct disparities by confining the search range to be small near the fitted plane. Then, the least squares method is repeated to update parameters c_1, c_2, c_3 based on the modified disparities.

The above plane fitting process is repeated for each segment and newly found planes are added to a list. After that, each segment is assigned to a plane in the list that has lowest C_{SEG} value. Then we group the segments assigned to the same plane. And for each group, the above plane fitting is repeated in order to improve the accuracy of a plane. At last, we have the final list of the planes to use. Although, this plane-based model does not explicitly handle the occlusion, occluded pixels are likely to be detected as outlier through the crosscheck in plane estimation.

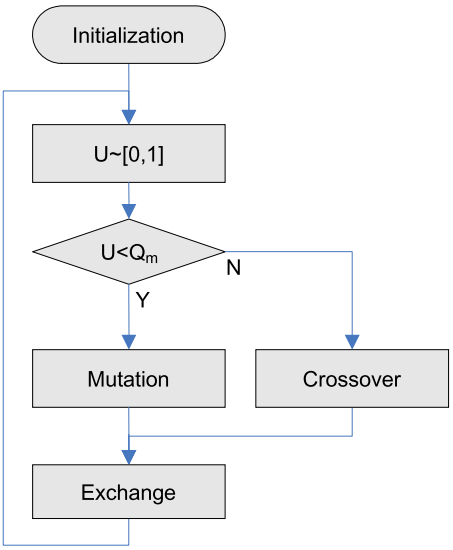

Fig. 2 The overall flow chart of the proposed Pop-MCMC algorithm applied to stereo matching

3.2 Design of Pop-MCMC

Given a target probability distribution $\pi(\mathbf{X}) \propto \exp\{-E(\mathbf{X})\}$, our aim is to find the state $\mathbf{X}$ where the probability is maximized. In Pop-MCMC, we draw multiple samples from multiple chains at the same time with respect to the following distributions.

$$\pi_i(\mathbf{X}_i) = \pi(\mathbf{X}_i)^{\frac{1}{T_i}} \propto \exp\left\{-\frac{E(\mathbf{X}_i)}{T_i}\right\}, \tag{18}$$

where T_i is the temperature of ith chain. The appropriate sequence of the temperatures can be designed empirically according to the target distribution. Each sample from each chain is a chromosome, and chromosomes interact with each other, which helps perform global moves.

The overall flow of Pop-MCMC is illustrated in Fig. 2. The three moves, mutation, crossover, and exchange moves are repeatedly performed and samples are generated at each iteration. In this process, we first choose a random number U between 0 and 1, and compare U with the mutation rate Q_m. Depending on the value of U, we choose one move between mutation and crossover. So, by varying Q_m, we can control the rates between the global move (crossover) and local move (mutation) easily. This means that Q_m adjusts the trade-off between exploration and convergence of the algorithm (Spears 1992). A proper value of Q_m can be chosen according to the given problem, the model, or the number of chains. For example, if a large number of chains are used, Q_m is usually set to a small value for faster convergence. Let us describe the detailed design of each move for the stereo problem.

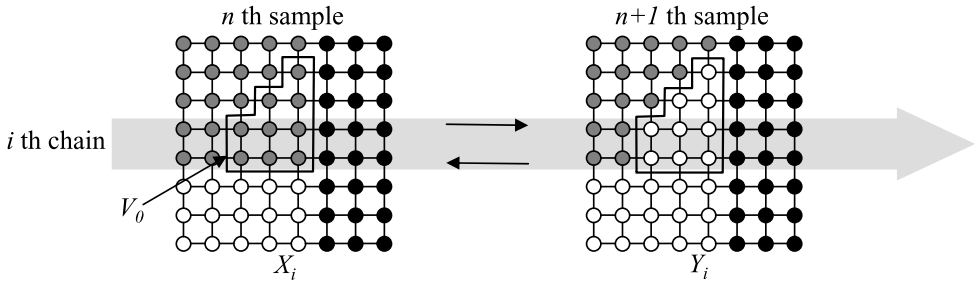

Fig. 3 An example of mutation move

Mutation Move In this paper, we employ the MCMC kernel of SWC-2 for the mutation move of a randomly selected chain. At first, we construct a random cluster V_0 as in SWC-2 for a selected chain. For clustering, we need to design edge probability q_e, which determines whether the edge should be retained or not. We define the edge probability,

$$q_e = 1 - \exp\left(-\frac{K_i \cdot \mathcal{S}(s,t)}{\frac{C_{\mathrm{SEG}}(f_{v_1})}{N(v_1)} + \frac{C_{\mathrm{SEG}}(f_{v_2})}{N(v_2)} + 2}\right), \tag{19}$$

where v_1 and v_2 represent neighboring nodes, $N(v)$ is the number of the pixels in the node (segment) v, and K_i represents a weighting factor for the chosen ith chain. The more similar the intensities of the connected nodes and the lower the matching costs are, the higher the probability that the edge remains. Note that the matching costs are normalized by the sizes of the corresponding segments. By varying K_i, we can control the average size of clusters. A bigger K_i tends to generate bigger clusters. We set K_i to increase as i increases. Consequently, clusters are likely to be small in lower-temperature chains and big in higher-temperature chains. It helps more effective exploration and also prevents chromosomes from correlating with each other.

The new label l' for the selected cluster V_0 is proposed according to the following proposal distribution.

$$q(l' \mid V_0, \mathbf{X}_i) = \exp\left[-\left\{\frac{\sum_{v \in V_0} C_{\mathrm{SEG}}(f_v = l')}{\sum_{v \in V_0} N(v)} + 1 - \prod_{\langle v_1, v_2 \rangle \in N, v_1 \in V_0, v_2 \notin V_0} \mathbf{1}(l' = f_{v_2})\right\}\right], \tag{20}$$

where l' is the newly proposed label for V_0, and $\mathbf{X}_i$ is the current state of selected ith chain. When the nodes in the cluster V_0 have low matching costs and there exist neighboring nodes of same label, the value of $q(l' \mid V_0, \mathbf{X}_i)$ becomes high. After a new label is proposed, it is accepted according to the Metropolis-Hastings rule. By substituting (18) and the transition probability in (2) into (6), we can calculate the acceptance probability:

$$\begin{aligned}\alpha &= \min(1, \gamma_m) \\ &= \min\left(1, \frac{\pi_i(\mathbf{Y}_i)}{\pi_i(\mathbf{X}_i)} \cdot \frac{T(\mathbf{Y}_i \to \mathbf{X}_i)}{T(\mathbf{X}_i \to \mathbf{Y}_i)}\right) \\ &= \min\left(1, \exp\left\{\frac{E(\mathbf{X}_i) - E(\mathbf{Y}_i)}{T_i}\right\} \cdot \frac{q(V_0 \mid \mathbf{Y}_i) q(l \mid V_0, \mathbf{Y}_i)}{q(V_0 \mid \mathbf{X}_i) q(l' \mid V_0, \mathbf{X}_i)}\right),\end{aligned} \tag{21}$$

where $\mathbf{Y}_i$ is the proposed state of the ith chain, and $q(V_0 \mid \mathbf{X}_i)$ is the probability for selecting cluster V_0 when current state is $\mathbf{X}_i$. Figure 3 illustrates an example of mutation move on the ith chain.

Exchange Move In this move, we choose two chains and propose to exchange the chromosomes of two chains. The proposal is accepted or not by the Metropolis-Hastings rule. Figure 4 shows an example of exchange move. Note that for the exchange move, there is no need for a special design for stereo matching problem. So, when the ith and jth chains are selected, by substituting (18) into (8), we can obtain the acceptance probability by

$$\begin{aligned}\alpha &= \min(1, \gamma_e) \\ &= \min\left(1, \frac{\pi_i(\mathbf{X}_j)\pi_j(\mathbf{X}_i)}{\pi_i(\mathbf{X}_i)\pi_j(\mathbf{X}_j)}\right) \\ &= \min\left(1, \exp\left[\{E(\mathbf{X}_i) - E(\mathbf{X}_j)\} \cdot \left(\frac{1}{T_i} - \frac{1}{T_j}\right)\right]\right),\end{aligned} \tag{22}$$

where $\mathbf{X}_i$ and T_i are the current state and temperature of the ith chain. In order to achieve faster mixing rate, we need to raise the acceptance rate, and this can be accomplished by choosing two neighboring chains that have similar temperatures. Then, from the above equation, the Metropolis-Hastings ratio tends to get bigger.

Crossover Move Typical crossover moves commonly used in conventional Pop-MCMC are the 1-point crossover and

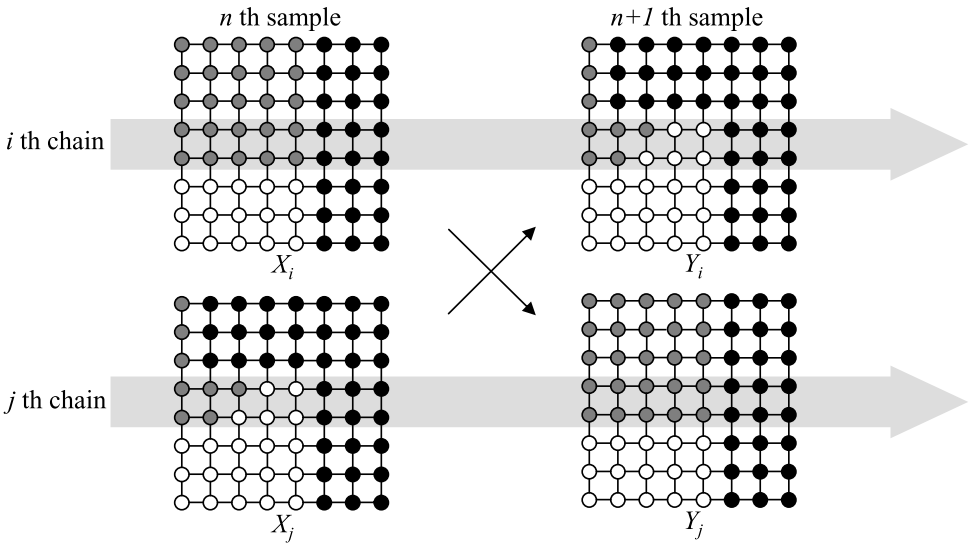

Fig. 4 An example of exchange move

2-point crossover moves. However, since these methods are designed for the chromosomes of 1-D vectors, it is inappropriate to apply them directly to the stereo matching problem, in which the chromosomes are 2-D state configurations. Nonetheless, the 1-point and 2-point crossover moves have an advantage of low computational complexity because the most of terms in the Metropolis-Hastings ratio cancel out each other. Therefore, in this work, we introduce a new 2-D crossover move that maintains this advantage. Detailed description of the proposed crossover move is as follows.

We first choose two chains randomly and construct a cluster V_0 in a similar way as in SWC-2 (or the mutation move). However, there are two differences in constructing V_0 compared with SWC-2. First, q_e is set constant, not adaptively determined with the matching costs or the intensities of the input image, since there is no need for the nodes of the cluster V_0 to be homogeneous in this case. It is also computationally efficient to use q_e as a constant value because the proposal distribution part in the Metropolis-Hastings ratio is canceled out. Second, when we calculate the probability q_e, we do not have to check whether the labels of the nodes are the same or not, so the resulting cluster V_0 can have nodes with different labels. Therefore, compared with the mutation move that requires the identifying and removing processes of all the edges connecting the nodes with different labels, the selecting scheme and the calculation of the acceptance probability of V_0 in the crossover move is much simpler. Eventually this property enables high efficiency in computation, and also the freedom in the construction of V_0 helps to achieve faster convergence.

The process after constructing a cluster V_0 is similar to the 1-point crossover move. From the chromosomes $\mathbf{X}_i$ and $\mathbf{X}_j$ of two selected chains, new chromosomes $\mathbf{Y}_i$ and $\mathbf{Y}_j$ are proposed by exchanging the labels of the nodes which belong to the cluster V_0 as shown in Fig. 5. The acceptance probability $\alpha = \min(1, \gamma_c)$ of the newly proposed chromosomes is calculated, and the next population of samples is determined. By substituting (18) into the Metropolis-Hastings rule in (11), we can obtain γ_c as follows.

$$\begin{aligned}
\alpha &= \min(1, \gamma_c) \\
&= \min\left(1, \frac{\pi_i(\mathbf{Y}_i)\pi_j(\mathbf{Y}_j)}{\pi_i(\mathbf{X}_i)\pi_j(\mathbf{X}_j)} \cdot \frac{q(\mathbf{X}_i, \mathbf{X}_j|\mathbf{Y}_i, \mathbf{Y}_j)}{q(\mathbf{Y}_i, \mathbf{Y}_j|\mathbf{X}_i, \mathbf{X}_j)}\right) \\
&= \min\left(1, \frac{\pi_i(\mathbf{Y}_i)\pi_j(\mathbf{Y}_j)}{\pi_i(\mathbf{X}_i)\pi_j(\mathbf{X}_j)}\right) \\
&= \min\left(1, \exp\left[\frac{E(\mathbf{X}_i) - E(\mathbf{Y}_i)}{T_i} \right.\right. \\
&\qquad \left.\left. + \frac{E(\mathbf{X}_j) - E(\mathbf{Y}_j)}{T_j}\right]\right),
\end{aligned} \tag{23}$$

where we used the symmetric property of the proposal distribution $q(\mathbf{X}_i, \mathbf{X}_j|\mathbf{Y}_i, \mathbf{Y}_j)$.

The proposed Pop-MCMC algorithm is summarized in Algorithm 1.

4 Experimental Results

We have implemented the proposed algorithm on a 2.8 GHz Pentium IV PC platform. In this section, we evaluate the performance of the proposed algorithm by comparing with other conventional methods such as SWC-2, SA, BP, and Graph cuts. In addition, we illustrate the effects of each move, temperature parameter, and the number of chains. We tested the proposed algorithm on several benchmark images in the Middlebury datasets (http://vision.middlebury.edu/stereo). Figure 6 shows the reference images and the ground truth maps of the test images. We used the segment-based

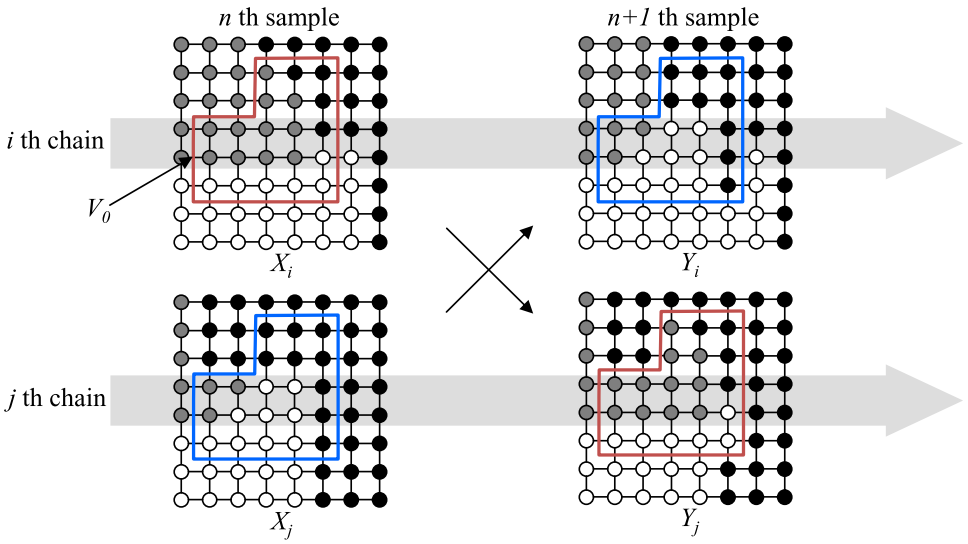

Fig. 5 An example of crossover move

Algorithm 1 Proposed Pop-MCMC algorithm

(Initialize)
Initialize the population $X_{1:N}$ by Winner-Takes-All manner with data cost.
Set the temperatures $T_1 < T_2 < \cdots < T_N$.
repeat
  if $U \sim [0, 1] < Q_m$ **then**
    for $i = 1$ to N **do**
      (Mutation)
      Select a random node v in ith chain.
      Draw a cluster from a node v with SWC-2.
      Propose a new label for the cluster and determine whether accept it or not with Metropolis-Hastings rule.
    end for
  else
    for $i = 1$ to $\lfloor \frac{N}{5} \rfloor$ **do**
      (Crossover)
      Select two random chains and a random node v.
      Draw a cluster from node v with modified SWC-2.
      Determine whether swap the cluster or not with Metropolis-Hastings rule.
    end for
  end if
  for $i = N - 1$ to 1 **do**
    (Exchange)
    Perform the exchange move onto ith and $i + 1$th chains with Metropolis-Hastings rule.
  end for
until The algorithm converges.

energy model in (12) for the test. Pop-MCMC, SWC-2, and SA methods were repeated ten times on each test stereo image pair since they are stochastic methods, and the averages and standard deviations of the resulting energies were compared.

We fixed the parameter values of Pop-MCMC for all the test sets. Empirically, the temperatures were set to be decreasing linearly in the range of predefined maximum and minimum temperature values. The maximum and minimum temperatures were set to 1.0 and 0.0001, respectively, and the number of chains was set to five. Q_m was set to be 0.25. For the edge probability of the ith chain, we set $K_i = 3i + 1$. This helped the chromosomes not to correlate with each other.

Figure 7 presents the comparison of the energy plots against running time in second for Pop-MCMC, SWC-2,

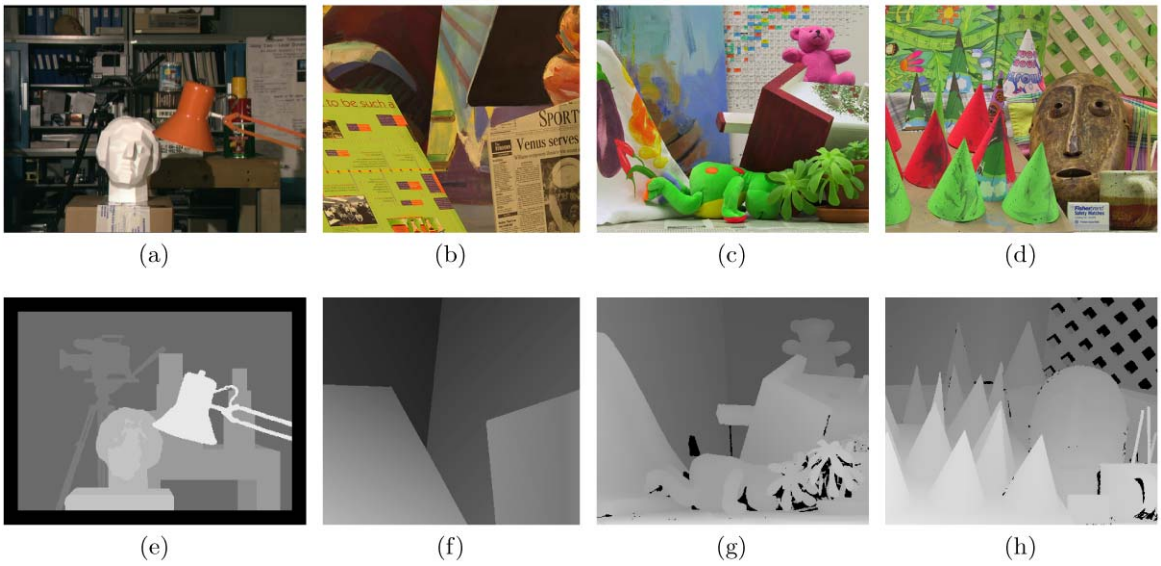

Fig. 6 Test stereo images: (**a**)–(**d**) reference images, (**e**)–(**h**) ground truth disparity maps. (**a**, **e**) Tsukuba, (**b**, **f**) Venus, (**c**, **g**) Teddy, and (**d**, **h**) Cones

SA, BP, and Graph cut methods (α-expansion move and α-β-swap move). The same energy model was applied to each method. For the implementation of SWC-2, we followed the work of Barbu and Zhu (2004, 2005), and for Graph cuts, we used the source code from (http://vision.middlebury.edu/MRF). α-expansion move method showed the best performance in all the tests. Although the attained energy values of Pop-MCMC are slightly higher than those of two graph cut algorithms, since the differences are very small, the proposed algorithm is comparable to them. And note that in contrast to the graph cut algorithms that can be applied only to submodular functions with pairwise priors, Pop-MCMC can be applied to wider class of energies including higher-order MRFs and highly complicated MRFs. Moreover, the convergence rate of Pop-MCMC can be substantially speeded up by using parallel implementation. Therefore Pop-MCMC can be a good alternative to α-expansion and α-β-swap move methods. Note also that Pop-MCMC algorithm reached much lower energy states than SA and SWC-2 on all the cases except Tsukuba. And, it even showed better performance than BP for all the test images. On the Tsukuba images, each method obtained relatively good result since the dimension of the solution space is low and thus the energy model is relatively simple. Note that on all test images, the convergence rates of Pop-MCMC are much faster and its standard deviations are consistently smaller than those of SA and SWC-2. From these results we can argue that conventional sampling-based methods like SA and SWC-2 are easily trapped at local minima, while Pop-MCMC is more likely to approach the global minimum due to the global moves in Pop-MCMC, such as exchange and crossover moves.

The disparity error rates of the Pop-MCMC were compared with those of other algorithms and shown in Table 1, and the resulting disparity maps of the proposed algorithm are shown in Fig. 8. Note that there are some limitations of the segment-based energy model. When real world objects are piecewise planar, the results are quite good. However, for the cases of Teddy and Cones that include objects with curved surfaces, the performance seems not satisfactory. And also, for a fronto-parallel plane, a non-segment based energy model can be superior to the segment-based energy model due to the smaller number of labels. In addition, since occlusion or visibility was not considered in our stereo model, the error rates at the vicinity of discontinuity were relatively large.

Figures 9 and 10 report the contribution of each move in Pop-MCMC. Figure 9 shows the statistics of each move while Pop-MCMC is running on the Venus images. We counted the number of accepted moves every ten seconds. As shown in the graphs, the mutation move occurred most frequently. In the beginning, all the three moves frequently occurred but as time went on, they tended to decrease since they were approaching the optimum. The exchange move occasionally occurred when higher-temperature chromosomes had lower energy states than those of lower-temperature chromosomes. Figure 10 compares the energy convergence rates for different combinations of moves. We illustrated the energy curves and the boxplots of the final state energies. We performed the experiment on the Venus images. The exchange move contributed larger amount than the crossover move. Obviously, when we combined all the three moves, they together helped each other to achieve fast convergence. Boxplots of the final state energies show not

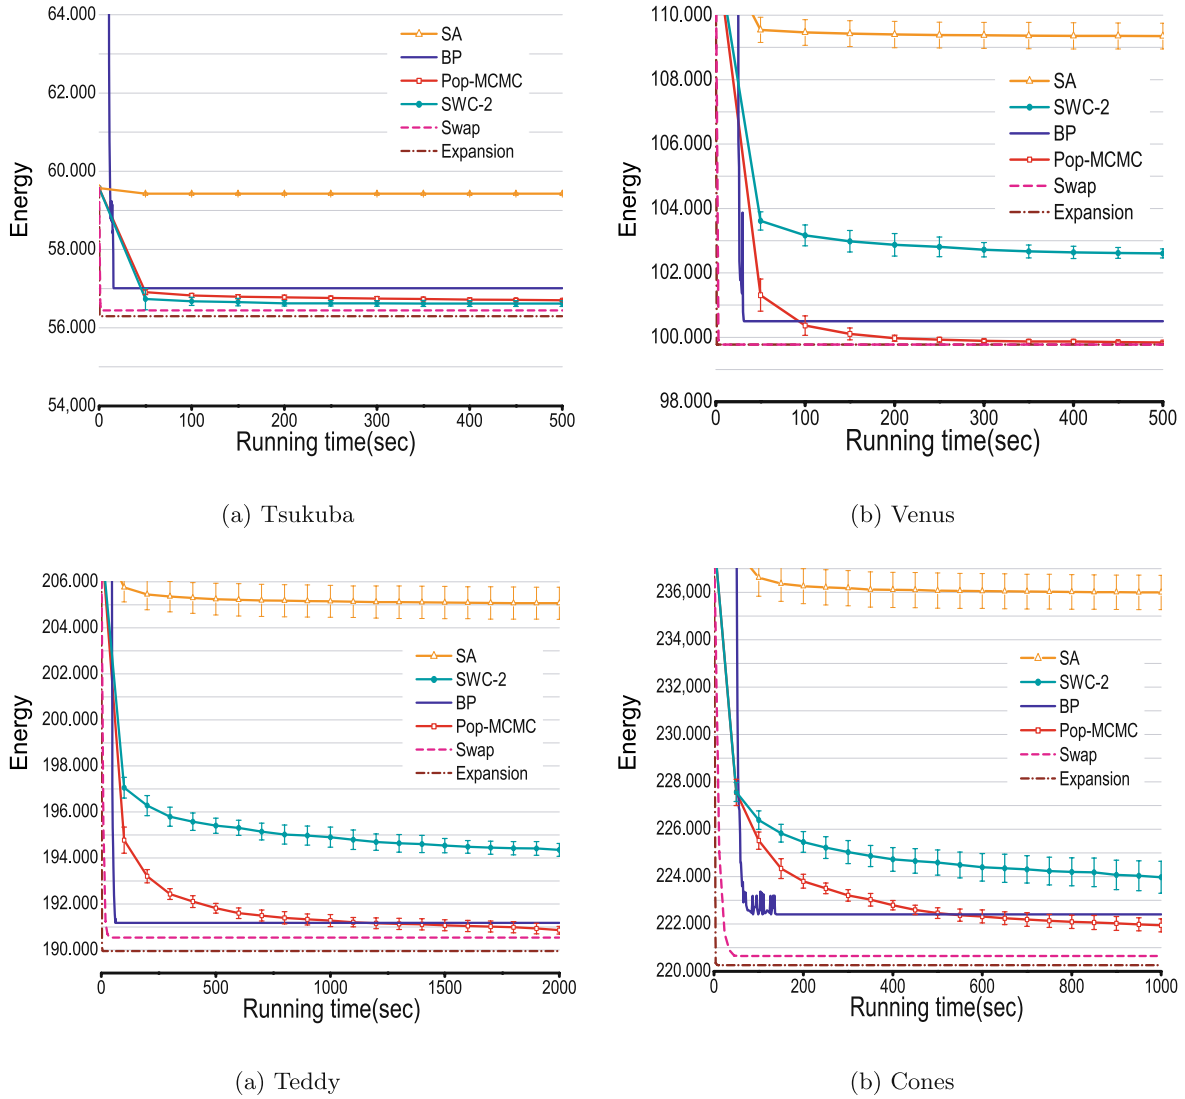

(a) Tsukuba (b) Venus

(a) Teddy (b) Cones

Fig. 7 Performance (energy vs. running time) comparison of Pop-MCMC, SA, SWC-2, α-expansion, α-β-swap, and BP on (**a**) Tsukuba, (**b**) Venus, (**c**) Teddy, and (**d**) Cones. Pop-MCMC obtains lower energy results than other methods except on Tsukuba

only that the three moves together could reach lower energy state, but also that they decreased the standard deviation, and in turn made the algorithm quite stable. It took about 190 seconds to minimize the energy to be 100,000 using all moves. However, if one of the moves was missing, it became much slower. For example, without the crossover move, it took about 440 seconds, and without the exchange move, it could not reach that state until 500 seconds.

Figures 11 and 12 exhibit the performance for differing parameters. Both experiments were carried on the Venus image. Figure 11 shows the energy convergence plots according to the variation of the max-temperature. The min-temperature was set to 0.0001. We observed that if the max-temperature was too low, it quickly moved to the nearest minimum but easily got stuck in local minima. While if it was too high, the algorithm was rarely trapped in local minima but the convergence speed became too slow. Figure 12 shows the energy convergence plots by varying the number of chains. If the population size was large, it helped each other to reach the global minimum by exchanging information. However, a large size of population usually increases redundancy in the algorithm. We found empirically that for our segment-based stereo energy model, the optimal max-temperature was 1.0 and the optimal number of chains was five.

5 Conclusion

In this paper, we proposed a new stereo matching algorithm based on Pop-MCMC. We showed that the proposed sampling-based Pop-MCMC was a good optimizer

Table 1 The error rates for each test image (http://vision.middlebury.edu/stereo). For the sampling-based methods, we denote the average and standard deviation for ten trials. *nonocc*, *all*, and *disc* represent the error rate within non-occluded region, the whole image, and the vicinity of discontinuity, respectively

Method	Tsukuba			Venus			Teddy			Cones		
	nonocc	*all*	*disc*	*nonocc*	*all*	*disc*	*nonocc*	*all*	*disc*	*nonocc*	*all*	*disc*
Pop-MCMC	3.35 (±0.42)	3.88 (±0.42)	10.3 (±0.92)	0.22 (±0.01)	0.35 (±0.02)	2.89 (±0.17)	12.0 (±0.56)	17.9 (±0.69)	21.7 (±0.63)	13.3 (±0.37)	19.2 (±0.54)	23.7 (±0.57)
SWC	3.69 (±1.24)	4.28 (±1.23)	10.4 (±1.21)	0.9 (±0.27)	1.1 (±0.26)	5.57 (±0.11)	11.6 (±0.72)	17.8 (±0.86)	22.2 (±0.99)	13.5 (±0.72)	20.3 (±0.97)	23.4 (±0.76)
SA	3.5 (±0.28)	4.09 (±0.3)	9.58 (±0.48)	0.94 (±0.16)	1.34 (±0.21)	7.67 (±0.78)	14.8 (±0.61)	21.4 (±0.59)	24.3 (±0.79)	15.6 (±0.69)	22.9 (±0.78)	25.1 (±0.43)
BP	3.12	3.76	10.5	0.21	0.34	2.81	10.5	16.5	20.4	12.9	19.2	23.3
α-expansion	4.12	4.73	12.2	0.21	0.34	2.81	10.9	12.4	19.1	12.5	18.6	23.1
α-β-swap	2.56	3.09	9.15	0.21	0.34	2.81	10.5	12.0	19.7	13.0	19.0	23.6

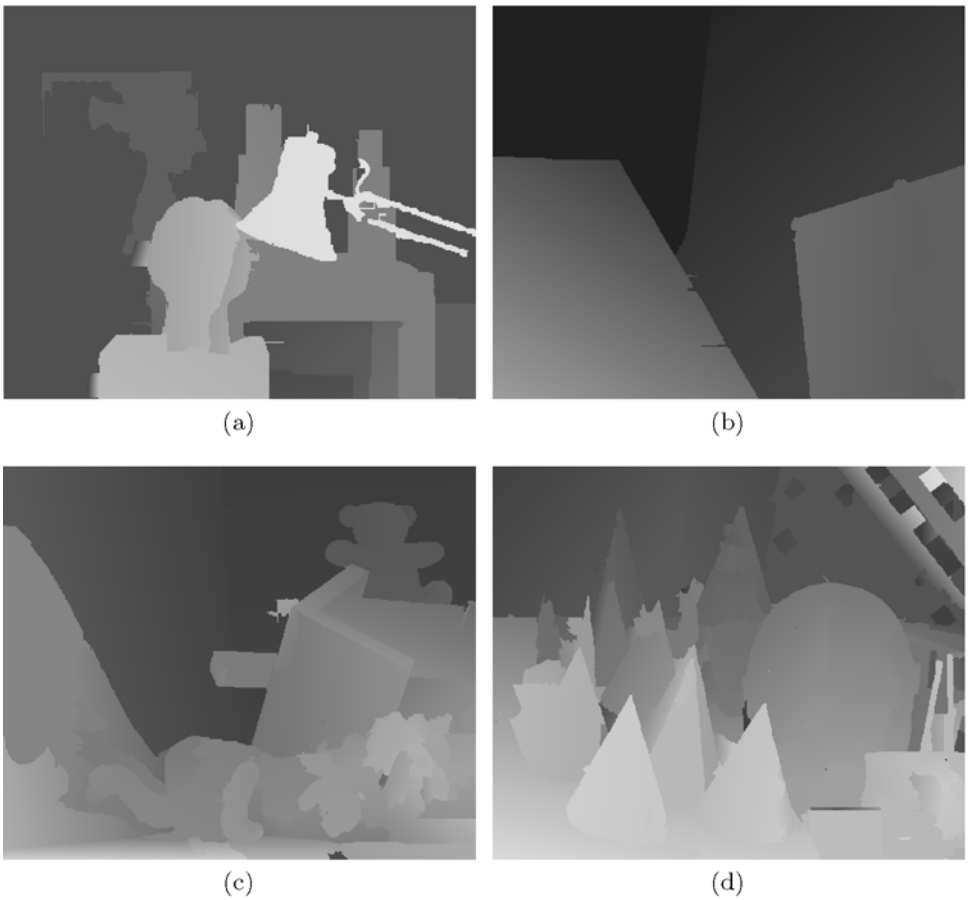

Fig. 8 Results of the proposed algorithm: the disparity maps of (**a**) Tsukuba, (**b**) Venus, (**c**) Teddy, and (**d**) Cones

for stereo problem. Pop-MCMC uses multiple chains in parallel, and establishes faster mixing rate by exchanging information between chromosomes. In this work, we designed new effective 2-D mutation and crossover moves for stereo matching based on cluster sampling technique. Consequently, it is shown that the proposed algorithm provides much faster convergence rate than conventional sampling-based methods including SA and SWC, and gives lower energy states than BP. We also investigated the contribution of each move. Combining all the three moves together made the algorithm more stable. In addition, we analyzed the effect of parameters such as temperature and the number of

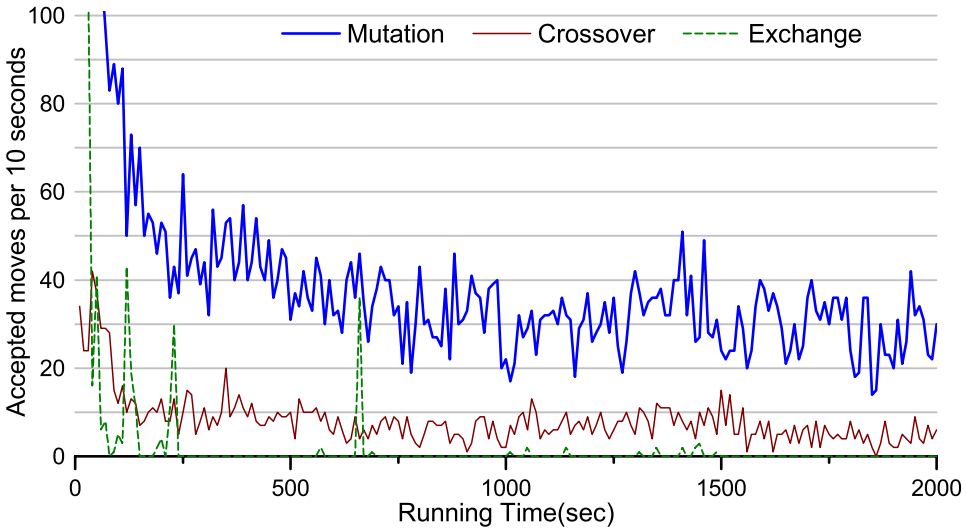

Fig. 9 The statistics of each move while running the algorithm on the Venus images. Initially, all the three moves are quite active, and then tend to decrease as time goes on. While mutation and crossover moves consistently occur, exchange move occurs occasionally

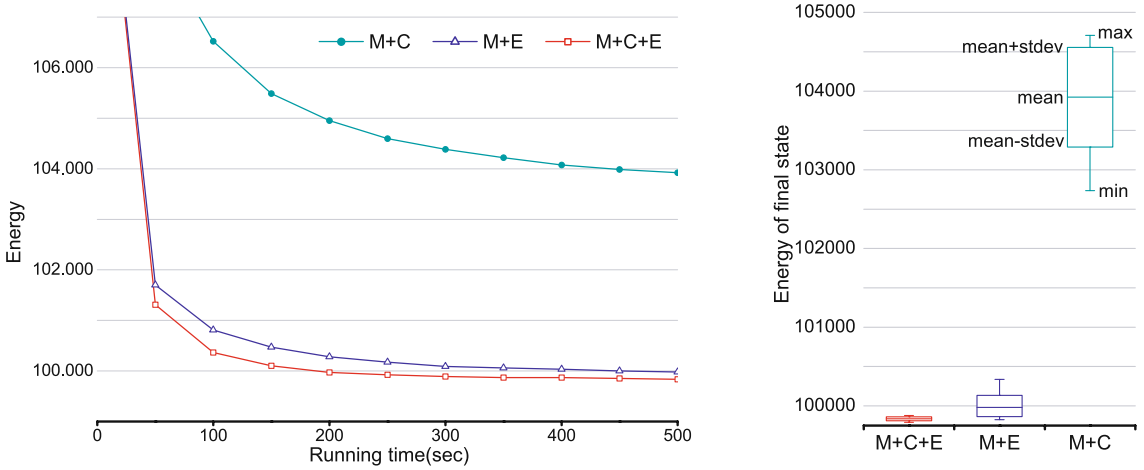

(a) (b)

Fig. 10 The performance of the Pop-MCMC for different combinations of moves: (**a**) Energy curves, (**b**) boxplots of the final states

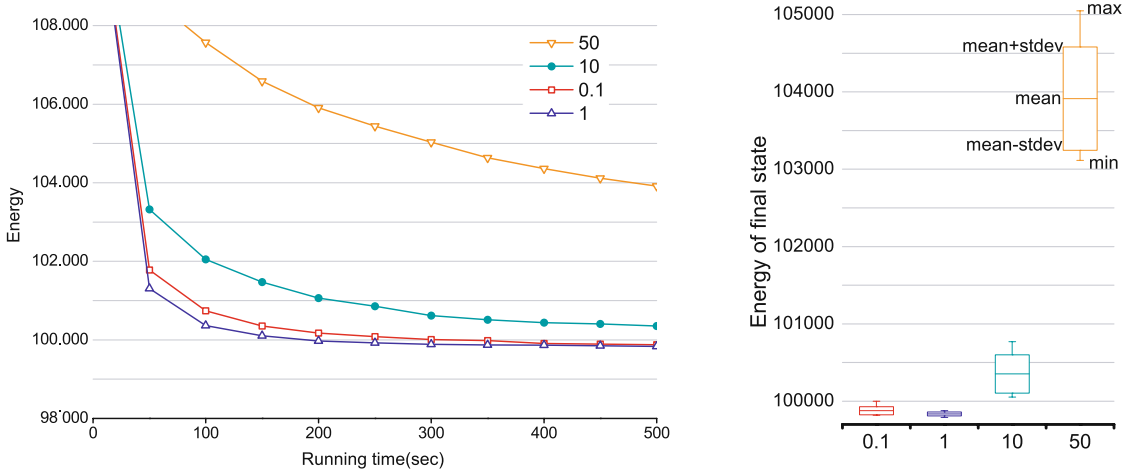

(a) (b)

Fig. 11 The performance of the Pop-MCMC for different max temperature values: (**a**) Energy curves, (**b**) boxplots of the final states

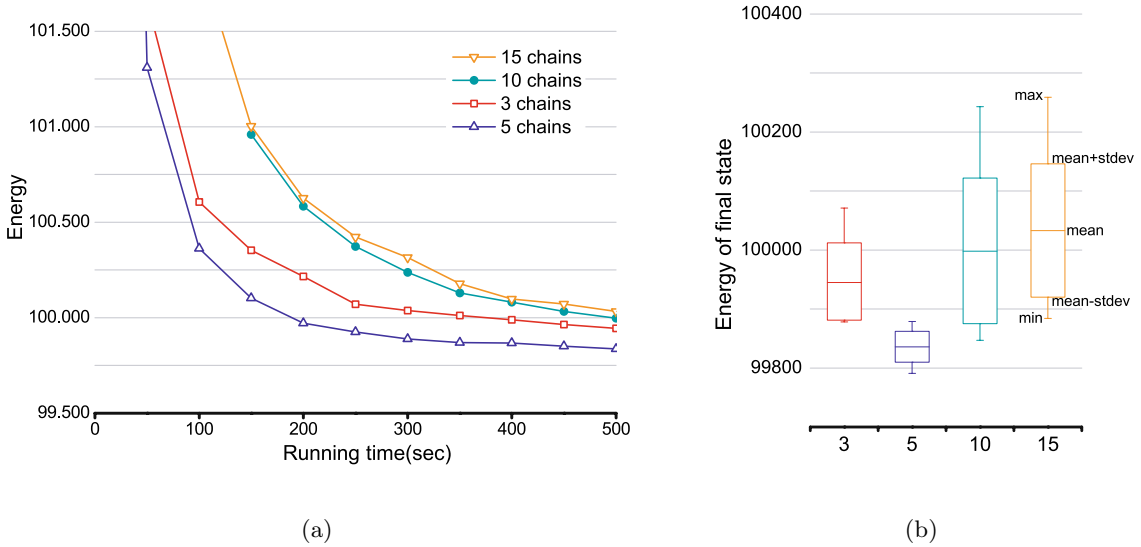

Fig. 12 The performance of the Pop-MCMC for different number of chains: (**a**) Energy curves, (**b**) boxplots of the final states

chains, and found the optimal parameters for our problem. We have a plan to apply and analyze the performance of the proposed method to more sophisticated stereo energy models including occlusion handling and visibility terms as well as the segmentation problem.

Acknowledgements This work was supported in part by the ITRC program by Ministry of Information and Communication and in part by Defense Acquisition Program Administration and Agency for Defense Development, through the Image Information Research Center, Korea.

References

Barbu, A., & Zhu, S. C. (2004). Multigrid and multi-level Swendsen-Wang cuts for hierarchic graph partition. In *Conference on computer vision and pattern recognition*, 731–738.

Barbu, A., & Zhu, S. C. (2005). Generalizing Swendsen-Wang to sampling arbitrary posterior probabilities. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1239–1253

Birchfield, S., & Tomasi, C. (1998). A pixel dissimilarity measure that is insensitive to image sampling. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.

Bleyer, M., & Gelautz, M. (2005). Graph-based surface reconstruction from stereo pairs using image segmentation. In *Proceedings of society of photo-optical instrumentation engineers: Vol. 5665. Videometrics VIII* (pp. 288–299).

Boykov, Y., & Kolmogorov, V. (2004). An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *26*(9), 1124–1137.

Boykov, Y., Veksler, O., & Zabih, R. (2001). Fast approximate energy minimization via graph cuts. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *23*(11), 1222–1239.

Comanicu, D., & Meer, P. (2002). Mean shift: A robust approach toward feature space analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *24*, 603–619.

Geyer, C. J. (1991). Markov chain Monte Carlo maximum likelihood. *Computing Science and Statistics*, 156–163.

Hong, L., & Chen, G. (2004). Segment-based stereo matching using graph cuts. In *Conference on computer vision and pattern recognition* (pp. 74–81).

Hukushima, K., & Nemoto, K. (1996). Exchange Monte Carlo method and application to spin glass simulations. *Journal of the Physical Society of Japan*, *65*, 1604–1608.

Jasra, A., & Stephens, D. A. (2007). On population-based simulation for static inference. *Statistics and Computing*, *17*(3), 263–279.

Kirkpatrick, S., Gelatt, C., & Vecchi, M. (1983). Optimization by simulated annealing. *Science*, *220*(4598), 671–680.

Klaus, A., Sormann, M., & Karner, K. (2006). Segment-based stereo matching using belief propagation and a self-adapting dissimilarity measure. In *International conference on pattern recognition* (pp. 15–18).

Kolmogorov, V. (2006). Convergent tree-reweighted message passing for energy minimization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *28*(10), 1568–1583.

Kolmogorov, V., & Zabih, R. (2001). Computing visual correspondence with occlusions using graph cuts. In *Proceedings of international conference on computer vision* (pp. 508–515).

Kolmogorov, V., & Zabih, R. (2002). Multi-camera scene reconstruction via graph cuts. In *Proceedings of European conference on computer vision* (pp. 82–96).

Kolmogorov, V., & Zabih, R. (2004). What energy functions can be minimized via graph cuts? *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *26*(2), 147–159.

Komodakis, N., Tziritas, G., & Paragios, N. (2007a). Fast, approximately optimal solutions for single and dynamic MRFs. In *Conference on computer vision and pattern recognition*.

Komodakis, N., Paragios, N., & Tziritas, G. (2007b). MRF optimization via dual decomposition: Message-passing revisited. In *Proceedings of international conference on computer vision*.

Liang, F., & Wong, W. H. (2000). Evolutionary Monte Carlo: Applications to C_p model sampling and change point problem. *Statistica Sinica*, *10*, 317–342.

Liang, F., & Wong, W. H. (2001). Real parameter evolutionary Monte Carlo with applications to Bayesian mixture models. *Journal of the American Statistical Association*, 653–666.

Ohta, Y., & Kanade, T. (1985). Stereo by intra- and inter-scanline search. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *2*, 449–470.

Park, J., Kim, W., & Lee, K. M. (2007). Stereo matching using population-based MCMC. In *Proceedings of Asian conference on computer vision*.

Scharstein, D., & Szeliski, R. (2002). A taxonomy and evaluation of dense two-frame stereo correspondence algorithms. *International Journal of Computer Vision*, *47*(1–3), 7–42.

Spears, W. M. (1992). Crossover or mutation? In *Foundations of genetic algorithms* (Vol. 2). San Mateo: Morgan Kaufmann.

Sun, J., Zheng, N. N., & Shum, H. Y. (2003). Stereo matching using belief propagation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *25*(7), 787–800.

Swendsen, R. H., & Wang, J. S. (1987). Nonuniversal critical dynamics in MC simulations. *Physical Review Letters*, 86–88.

Szeliski, R., Zabih, R., Scharstein, D., Veksler, O., Kolmogorov, V., Agarwala, A., Tappen, M., & Rother, C. (2006). A comparative study of energy minimization methods for Markov random fields. In *Proceedings of European conference on computer vision*.

Tao, H., Sawhney, H. S., & Kumar, R. (2001). A global matching framework for stereo computation. In *Proceedings of international conference on computer vision* (Vol. 1, pp. 532–539).

Veksler, O. (2005). Stereo correspondence by dynamic programming on a tree. In *Conference on computer vision and pattern recognition*.

Yedidia, J. S., Freeman, W. T., & Weiss, Y. (2000). Generalized belief propagation. In *Proceedings of neural information processing systems* (pp. 689–695).