Importance Sampling for Dependent Random Variables with Application to Contour Grouping

Anonymous CVPR submission

Paper ID ****

Abstract

1. Background

Our goal is to compute a posterior distribution $p(X_1, \ldots, X_m \mid Z)$, where $(X_1, \ldots, X_m)$ is a vector of random variables (RVs) and $Z = (z_1, \ldots, z_m)$ is a vector of observations. This will allow us to find value assignments $X_t = \widehat{x}_t$ for $t = 1, \ldots, m$ to RVs that maximize this posterior:

$$\widehat{x}_{1:m} = \underset{x_{1:m}}{\operatorname{argmax}}\, p(x_{1:m} \mid Z), \tag{1}$$

where $x_{1:m} = (x_1, \ldots, x_m) \in \mathcal{X}^m$ is a state space vector. Thus, the possible values of each RV range over the state space $\mathcal{X}$. We will achieve our goal by approximating the posterior distribution with a final number of samples in the framework of Bayesian Importance Sampling (BIS). Since it is usually difficult to draw samples from the pdf $p(x_{1:m}|Z)$, we will draw samples $x^{(i)}_{1:m} \sim q(x_{1:m}|Z)$ for $i = 1, \ldots, N$ from a proposal pdf q, form which samples are easily generated. Then approximation to the density p is given by

$$p(x_{1:m}|Z) \approx \sum_{i=1}^{N} w^{(i)} \delta(x_{1:m} - x^{(i)}_{1:m}), \tag{2}$$

where δ is the Dirac delta function and

$$w^{(i)} = \frac{p(x^{(i)}_{1:m}|Z)}{q(x^{(i)}_{1:m}|Z)} \tag{3}$$

are normalized weights (so that they sum to one).

Since it is still hard to draw samples from q due to high dimensionality of $x_{1:m}$, Sequential Importance Sampling (SIS) is usually utilized. Following the order of dimensions in the vector of RVs $X = (X_1, \ldots, X_t)$ samples are generated

$$x^{(i)}_t \sim q(x|x_{1:t-1}, z_{1:t}) \tag{4}$$

for $t = 1, \ldots m$, and the particles are built sequentially $x^{(i)}_{1:t} = (x^{(i)}_{1:t-1}, x^{(i)}_t)$ for $i = 1, \ldots, N$. The weights can also be recursive updated according to the formula

$$w(x^{(i)}_{1:t}) = w(x^{(i)}_{1:t-1}) \frac{p(z_t|x^{(i)}_{1:t}, z_{1:t-1}) p(x^{(i)}_t|x^{(i)}_{1:t-1})}{q(x^{(i)}_t|x^{(i)}_{1:t-1}, z_{1:t})} \tag{5}$$

such that $w^{(i)} = w(x^{(i)}_{1:m})$, which means that we obtain a set of weighted samples form $p(x_{1:m}|Z)$ once $t = m$.

A common assumption underlying the equations (4) and (5) is that there exist two known functions f and h such that

$$x_t = f(x_{t-1}) + u_t \tag{6}$$

$$z_t = h(x_t) + v_t \tag{7}$$

where both u_t and v_t are mutually independent and identically distributed sequences with known probability density functions (pdfs). Often they are assumed to be Gaussians that model state prediction noise u_t and observation noise v_t. (The assumption in (6) can be replied by a weaker one: $x_t = f(x_{1:t-1}) + u_t$, i.e., that we can determine x_t if we know all previous states $x_{1:t-1}$. For the simplicity of presentation we will use (6).

The equation (5) can be simplified by making a common assumption that $q(x^{(i)}_t|x^{(i)}_{1:t-1}, z_{1:t}) = p(x^{(i)}_t|x^{(i)}_{1:t-1})$, which yields

$$w(x^{(i)}_{1:t}) = w(x^{(i)}_{1:t-1}) p(z_t|x^{(i)}_{1:t}, z_{1:t-1}), \tag{8}$$

and, consequently, the samples are generated from

$$x^{(i)}_t \sim p(x|x^{(i)}_{1:t-1}), \tag{9}$$

i.e., the current observations $z_{1:t-1}$ have no influence on the samples.

2. New Approach

We illustrate our key ideas on an example of multi robot localization. A team of m robots obtained m observations

$Z = \{z_1, \ldots, z_m\}$ by exploring their environment, where observation z_t comes from robot t. For example, if the robots are equipped with laser range scanners, then each z_t could be a vector of laser range readings representing distances to the closest obstacles from the robot t. Each RV X_t describes robot poses, i.e., its values x_t represent coordinates and the heading direction of the robot t. Our goal is to determine the state vector of robot poses $(x_1, \ldots, x_m)$ in a given top view map of the environment that maximizes the posterior distribution $p(X_1, \ldots, X_m|Z)$. In other words, $(x_1, \ldots, x_m)$ is a vector of most likely robot poses given the measurements Z.

It is possible to apply the classical PF robot localization by utilizing the order of the observations $Z = \{z_1, \ldots, z_m\}$, which follows the numbering of the robots. Hence the follower for each particle (i) is determined by importance sampling from the proposal distribution, i.e., sample $x_t^{(i)} \sim p(x|x_{1:t-1}^{(i)})$ and the particle weight is updated based on the evaluation of the observation z_t according to the recursive formula in Eq. 8. However, by doing so, we would have selected an arbitrary order, and in particular, if the robot localization task fails, it could be due to the selected order. Would we have selected a different order, the localization task could have been successful. Moreover, the observations $Z = \{z_1, \ldots, z_m\}$ are collected simultaneously at the same time. i.e., after each robot completed its exploration. Consequently, there is no reason to favor any particular order without utilizing further information.

In the proposed approach, the order of the observations is not predetermined, in particular, we do not follow the order of indices of the observations in Z. Our key idea is to utilize the PF framework to determine the most informative order of the observations. This way we are able to simultaneously find the most informative order and to utilize the observations in the order of their informativeness. Intuitively, it make sense, for example, if the first robot took its laser readings in the middle of a along a long corridor and the second robot at the entrance to a room with many distinctive features, then our approach will first process the laser readings obtained by the second robot, since they are more informative.

We stress that the SIS in Eq. 4 and particle evaluation in Eq. 5 utilize the sequential order of the RVs reflected in the order of dimensions in the state space $(x_1, \ldots, x_m)$. In many applications, this order is determined naturally by the time stamp of the observations, e.g., a single robot is collecting laser measurements at consecutive time points, in which case x_t denotes the robot pose at time t. The goal of our work is to extend SIS to applications in which there is no natural order of observations like the case of multi robot localization.

The key idea of the proposed approach is not to utilize the fix order of the dimensions, but instead compute the best possible order of the dimensions $(x_{i_1}, \ldots, x_{i_m})$ (or equivalently RVs) so that the corresponding sequence of observations $Z = (z_{i_1}, \ldots, z_{i_m})$ is most informative. For this we extend the underlying assumption (6) to

$$x_s = f(x_k, s) + u_k \quad \text{forall} \quad s \in \{1, \ldots, m\}, \tag{10}$$

which means that all RVs depend on each other, and we can determine the value of variable s for any given value x_k of variable k. In our multi robot mapping example, it means that each robot knows the relative pose of the other robots. We observe that now the sequence of states visited before time t is not a sequence of consecutive numbers $(1, \ldots, t-1)$ but any subsequence $(i_1, \ldots, i_{t-1})$ formed by $t-1$ different numbers in $\{1, \ldots, m\}$, and the function f allows us to determine not only the value of the next state but values of all remaining states. Due to noise factor u_k, each value x_s is an estimate or prediction of a true unknown value. Since a given state x_s can be determined based on all the states in the current sequence, we can combine all the predictions $x_s = f(x_{i_k}, s) + u_{i_k}$ for $k = 1, \ldots, t-1$ and improve the accuracy of x_s. The assumption (7) remains unchanged.

The proposed sampling is as follows. In the first iteration $(t = 1)$ we generate m samples

$$x_s^{(i)} \sim p_s(x) \quad \text{for} \quad s \in \{1, \ldots, m\}. \tag{11}$$

At iteration t, the assumption (10) allows us to generate $m - t + 1$ samples from

$$x_{i_s}^{(i)} \sim p_{i_s}(x|x_{i_{1:t-1}}^{(i)}) \quad \text{for} \quad s \in \{t, \ldots, m\}. \tag{12}$$

Hence at iteration t particle (i) has $m - t + 1$ followers. For example, if $m = 5$ and $x_{i_{1:3}}^{(i)} = x_{1,3,5}^{(i)}$, then particle (i) will have two followers $x_2^{(i)}$ and $x_4^{(i)}$. With reference to our multi robot example, when we determined the locations of robots $1, 3, 5$, we consider the two remaining robots 2 and 4 and the next robot whose position we want to determine by particle (i). Of course, we repeat an analogous process for each particle (i) for $i = 1, \ldots, N$.

In contrast, in the standard application of rule (9), at each iteration t particle (i) has one follower. Even when sometimes each particle (i) has many followers, all followers are in the same dimension, which means that we only determine possible locations of say robot 2 by the followers and do not consider locations of robot 4 for particle (i), since a strict order of the state dimensions is followed in the classical setting.

Then we compute weights for each $x_{i_s}^{(i)}$, which is now indexed with the particle index (i) for $i = 1, \ldots, N$ and the follower index i_s for $s \in \{t, \ldots, m\}$. The formula for the iterative weight computation is derived below.

Finally analog to the standard PF, we perform resampling. The resampling plays in our framework an additional

and a very crucial role. It performs the selection of the most informative followers of the particles. As is the case in the standard PF, the weight of each particle is based on the evaluation how the predicted observation $h(x_{i_s})$ differs from the current observation z_{i_s}. Consequently, in the proposed approach, the weights are evaluated with respect to different observations. Thus, in the proposed approach the value of the weight depends on two factors,

- how descriptive a given observation z_{i_s} is and
- how good the prediction $h(x_{i_s})$ of observation z_{i_s} is.

This allows resampling to select the most informative followers of current particles as new particles in each iteration. Since all observations Z are given, the order of RVs of particles is heavily determined by Z, and this order may be different for each particle. This is in strong contrast to the classical PF, where observations Z have no influence on the order of RVs, which is fixed.

3. Particle Filter with Static Observations

Our derivation is analog to the PF derivation, but it differs fundamentally, since unlike the standard PF framework, the observations Z do not arrive sequentially, but are available at once. To simplify the notation we replace the double indexing of the state variables $x_{i_s}^{(i)}$ with a bijection (onto and one-to-one function) $< \cdot >^{(i)}: \{1,\dots,m\} \to \{1,\dots,m\}$. Although we may have a different bijection for each particle, we will drop the index (i) from $< 1 : t >^{(i)}$, since the state variables carry the particle index. For example, we will denote $(x_{i_1}^{(i)}, x_{i_2}^{(i)}, x_{i_3}^{(i)}) = x_{i_{1:3}}^{(i)}$ as $x_{<1:3>}^{(i)}$.

We derive now the recursive weight update formula for the static observations Z. For every t from 2 to m, we have

$$\begin{aligned} w(x_{<1:t>}^{(i)}) &= \frac{p(x_{<1:t>}^{(i)} \mid Z)}{q(x_{<1:t>}^{(i)} \mid Z)} \\ &= \frac{p(x_{<t>}^{(i)}|x_{<1:t-1>}^{(i)}, Z)\, p(x_{<1:t-1>}^{(i)}|Z)}{q(x_{<t>}^{(i)}|x_{<1:t-1>}^{(i)}, Z)\, q(x_{<1:t-1>}^{(i)}|Z)} \\ &= \frac{p(x_{<t>}^{(i)}|x_{<1:t-1>}^{(i)}, Z)}{q(x_{<t>}^{(i)}|x_{<1:t-1>}^{(i)}, Z)}\, w(x_{<1:t-1>}^{(i)}) \\ &= \frac{p(Z|x_{<1:t>}^{(i)}) p(x_{<t>}^{(i)}|x_{<1:t-1>}^{(i)})}{p(Z|x_{<1:t-1>}^{(i)}) q(x_{<t>}^{(i)}|x_{<1:t-1>}^{(i)}, Z)}\, w(x_{<1:t-1>}^{(i)}) \end{aligned} \tag{13}$$

To obtain the last equation, we apply Bayes rule to decompose $p(x_{<t>}^{(i)}|x_{<1:t-1>}^{(i)}, Z)$ that interchanges $x_{<t>}^{(i)}$ and Z.

As it is often the case in PF applications, we assume that $q(x_{<t>}^{(i)}|x_{<1:t-1>}^{(i)}, Z) = p(x_{<t>}^{(i)}|x_{<1:t-1>}^{(i)})$. Using this simple exploration based proposal the weight recursion in (13) becomes:

$$w(x_{<1:t>}^{(i)}) = \quad = w(x_{<1:t-1>}^{(i)}) \frac{p(Z|x_{<1:t>}^{(i)})}{p(Z|x_{<1:t-1>}^{(i)})} \tag{14}$$

By recursive substitution of weights in (14), i.e., by applying (14) to $w(x_{<1:t-1>}^{(i)}), w(x_{<1:t-2>}^{(i)}), \dots, w(x_{<1:2>}^{(i)})$, we obtain

$$\begin{aligned} w(x_{<1:t>}^{(i)}) &= w(x_{<1:t-2>}^{(i)}) \frac{\cancel{p(Z|x_{<1:t-1>}^{(i)})}}{p(Z|x_{<1:t-2>}^{(i)})} \frac{p(Z|x_{<1:t>}^{(i)})}{\cancel{p(Z|x_{<1:t-1>}^{(i)})}} \\ &= \dots = w(x_{<1>}^{(i)}) \frac{p(Z|x_{<1:t-1>}^{(i)}, x_t)}{p(Z|x_{<1>}^{(i)})} \end{aligned} \tag{15}$$

Under the assumption that all particles have the same initial weight $w(x_{<1>}^{(i)})$ and the same initial observation probability $p(Z|x_{<1>}^{(i)})$ for $i = 1,\dots,N$, we obtain

$$w(x_{<1:t>}^{(i)}) = p(Z|x_{<1:t>}^{(i)}) \tag{16}$$

Since for $t = m$ we have $w(x_{<1:m>}^{(i)}) = w(x_{1:m}^{(i)})$, we obtain that the weights computed by the recursive formulas (14) (16) are equal to the weights in Eq. (3).

For comparison, the corresponding weight update in the standard PF framework ([4]) is

$$w(x_{1:t}^{(i)}) = w(x_{1:t-1}^{(i)})\, p(z_t|x_{1:t-1}^{(i)}, x_t), \tag{17}$$

where z_t denotes the new observations obtained at time t. Because our observations Z do not have any natural order, Z cannot be expressed as a sequence of observations. We do not make any Markov assumption in the proposed formula (16), i.e., the new state $x_{<t>}^{(i)}$ is dependent on all previous states $x_{<1:t-1>}^{(i)}$ for each particle (i).

The approximation of Eq. 2 is computed in a framework of Sequential Importance Resampling (SIR). We outline now our PF algorithm, which in each iteration, i.e., at every time step t, and for each particle $i = 1,\dots,N$ executes the 3 steps:

1) **Importance sampling / proposal:** Sample $m - t + 1$ followers of particle (i)

$$x_s^{(i)} \sim p(x|x_{<1:t-1>}^{(i)}) \text{ for } s \in \{t,\dots,m\}. \tag{18}$$

and set $x_{<1:t-1>,s}^{(i)} = (x_{<1:t-1>}^{(i)}, x_s^{(i)})$ for $s \in \{t,\dots,m\}$. This step is possible by assumption (10).

2) **Importance weighting/evaluation:** An individual importance weight is assigned to each follower of each particle according to Eq. 16.

3) **Resampling:** At the sampling step we sample more followers than the number of particles. Thus we have a larger set of particles $x_{<1:t-1>,s}^{(i)}$ for $i = 1,\dots,N$ and

$s = t, \ldots, m$ from which we sub-sample N particles and assign qual weights to all of them as in the standard Sampling Importance Resampling (SIR) approach. While SIR requires a justification that it still computes (2) due to the presence of the old weights in (17), which are reset to be be equal to $1/N$ after resampling in SIR, this fact is obvious in the proposed approach, since the old weights are not present in (16).

The fact that we can consider more than one follower of each particle and reduce the number of followers by resampling is known in the PF literature and is referred to as prior boosting [2, 1]. It is used to capture multi-modal likelihood regions. As stated above, in the proposed approach the resampling has an additional role of selecting the the most informative random variables as followers of particles.

In order to be able to execute the proposed algorithm, we need to define the proposal distribution $q(x^{(i)}_{<t>}|x^{(i)}_{<1:t-1>}, Z) = p(x^{(i)}_{<t>}|x^{(i)}_{<1:t-1>})$ and the posterior distribution $p(Z|x^{(i)}_{<1:t>})$. We describe their constructions below.

Since the proposed PF framework performs sequential filtering, there are two important issues that need to be addressed: setting the initial particles $x^{(i)}_{<1>}$ for each $i = 1, \ldots, N$ (Section **??**) and the number of particles, which will be determined experimentally. We only mention here that the proposed approach is in some sense robust to the initial correspondences, since it does not matter with which correspondence we start as long as we start at some element of the optimal set of correspondences $\widehat{x}_{1:m} = (\widehat{x}_1, \ldots, \widehat{x}_m)$.

4. Improved Importance Sampling from MRFs

In this section we will develop the key motivational story behind our work i.e. sampling from the posterior distribution of a Markov Random Field. As is well known in computer vision community many low level problems have been modeled using MRFs. A limited work is done in modeling mid-level and high-level problems using MRFs not because they cannot be modeled in such a way but the inference process becomes computationally hard. In fact for a general MRF, inference becomes NP-hard. This is because inference in MRFs is closely related to assignment problems. Except for simple costs of assignments the assignment problem quickly turns out to be intractable. Our contribution is to be able to infer to a reasonable degree on the problem instances of contour grouping and object recognition where we can exploit certain results from importance sampling theory to effectively navigate the exponential space of assignments. this has applications in clustering, object recognition, contour grouping and any general applications where inferences can be modeled using Markov Random Fields or Conditional Random Fields.

5. Markov Random Fields

Let $S = \{s_1, s_2, \ldots, s_m\}$ be a family of random variables (RVs), which take the values $e_i \in E = \{e_1, e_2, \ldots, e_n\}$. E is called a set of labels. $s_i = e_i$ denotes the random event that the RV s_i gets e_i assigned. If f is a short hand for $(s_1 = e_1, s_2 = e_2, \ldots, s_m = e_m)$ and S forms a MRF as per definitions in [**?**]:

$$p(f) = \frac{1}{Z} e^{-\frac{1}{T}U(f)} \tag{19}$$

where Z is the normalizing constant and T is the temperature parameter determining the sharpness of the distribution: high-temperature makes all configurations equally likely since the effect of $U(f)$ goes down. Simply put $e^{-\frac{1}{T}U(f)} \rightarrow 1$ as $T \rightarrow \infty$.

At given T and particular U we can sample "patterns" of assignments by sampling from p.

The state space to be explored to understand the patterns of the posterior p is n^m. At lower Ts it finding Maximum A Posteriori (MAP) estimates of p has important applications in high-level vision problems. Thus one of the goals of sampling from MRF is:

$$\widehat{f} = \underset{f \in E^S}{\operatorname{argmax}}\, p(f) \tag{20}$$

There have been many sampling algorithms like Gibbs sampler, Hot Coupling ([3]), Tree sampling, Swendsen-Wang sampling etc. But most of them assume restrictive conditional independences. Recently Hamze et. al. proposed a very generic importance sampling method called Large Flip Importance Sampling (LFIS) to sample from the posterior [**?**]. The main motivation for their approach comes from N-Fold Way (NFW, [**?**]) and Tabu search ([**?**]) where they use heuristics to improve the sampling of the exponential state space using memory and heuristics to design good moves in the state space. Since the moves are no-longer MCMC in the traditional sense they introduce importance weights to the distinct states visited by N copies of the sampler. Independently there has been an application of similar strategy using particle filters with static observations in [**?**]. In this paper we combine the strengths of both the approaches and present an improved sampler that employs better weigthing scheme and navigational strategy to explore state space so as to compute MAP in an efficient way.

We first present a brief overview of both the approaches and then combine both into an integrated approach.

5.1. Large Flip Importance Sampling

Sampling algorithms used in computer vision often tend to use different terminology which might makes things a bit difficult to understand convergence complexity issues. Hence, even though the Gibbs sampler is the simplest

MCMC sampler for MRF we would like to present the algorithm below so as to make the connections between different approaches explicit. where f_i indicates the random

Algorithm 1 Gibbs sampler

1: Initialize a sample: $f^{(0)} = \left(s_1^{(0)} = e_1^{(0)}, s_2^{(0)} = e_2^{(0)}, \ldots, s_m^{(0)} = e_m^{(0)}\right)$.
2: **for** $i = 1$ to T **do**
3: Draw $f_1^{(i)} \sim p(f_1 | f_2^{(i-1)}, f_3^{(i-1)}, \ldots, f_m^{(i-1)})$.
4: Draw $f_2^{(i)} \sim p(f_2 | f_1^{(i)}, f_3^{(i-1)}, \ldots, f_m^{(i-1)})$. ⋮
5: Draw $f_m^{(i)} \sim p(f_m | f_1^{(i)}, f_2^{(i)}, \ldots, f_{m-1}^{(i)})$.
6: **end for**
7: Return $\{f^{(1)}, f^{(2)}, \ldots, f^{(T)}\}$.

event $s_i = e_i$. The above algorithm shows that the invariance kernel at a particular iteration above, for MCMC type convergence is defined as:

$$K(f, f') = \frac{1}{m} \sum_{i=1}^{m} p(f_i | f_{M \setminus i}) \tag{21}$$

The convergence condition says

$$\| K^{(T)}(f^{(0)}, f') - p(f') \| \to 0, \text{ as } T \to \infty \tag{22}$$

As can be seen from above there are two computational complexity issues in a sampling algorithm viz. (1) simulation complexity, (2) optimization complexity. Simulation complexity is because of the computations needed in drawing the samples while the optimization complexity involves T. The longer one runs the algorithm the closer one gets to the optimum. Of course in the most general case T can be exponentially large. Most practical algorithm designs involve in reducing T so as to visit non-trivial and important states of E^S as soon as possible without wasting computing resources. Such design algorithms are called *event driven* MCMC approaches. Motivated from NFW (algorithm below) LFIS was developed to address the "cycling" problem of NFW. where $\widehat{f}^{(i)} = f^{(\Theta_i \ldots \Theta_{i+1}-1)}$ i.e. the above algorithm avoids visiting duplicate states and the actual mixing time of the equivalent Gibbs sampler is Θ_T. Hence NFW effectively reduced the mixing time by cleverly covering more states in less time. Geometric$(p_{\text{flip}}(\widehat{f}_1^{(i-1)}, \widehat{f}_2^{(i-1)}, \ldots, \widehat{f}_m^{(i-1)}))$ is the probability that the $\widehat{f}^{(i-1)}$ changes over the next iteration. $\nu(j, e, \widehat{f}^{(i-1)})$ is the discrete posterior probability of the joint event "variable j assumes value e" given that a change in state occurred, i.e. $\widehat{f}^{(i)} \neq \widehat{f}^{(i-1)}$. Essentially the NFW simulates $T - 1$ *flips*.

Algorithm 2 N-Fold Way

1: Initialize a sample: $\widehat{f}^{(0)} = \left(s_1^{(0)} = e_1^{(0)}, s_2^{(0)} = e_2^{(0)}, \ldots, s_m^{(0)} = e_m^{(0)}\right)$.
2: **for** $i = 1$ to T **do**
3: Draw $\tau \sim$ Geometric$(p_{\text{flip}}(\widehat{f}_1^{(i-1)}, \widehat{f}_2^{(i-1)}, \ldots, \widehat{f}_m^{(i-1)}))$.
4: Draw $\widehat{f}^{(i)} \sim \nu(j, e, \widehat{f}^{(i-1)})$.
5: Set $\Theta_i = \Theta_i + \tau$.
6: **end for**
7: Return $\{\widehat{f}^{(1)}, \widehat{f}^{(2)}, \ldots, \widehat{f}^{(T)}\}$.

One reason for longer mixing times for some posteriors is the conditionals involved in the kernel can make the state not change i.e. $f = f'$ for a long time. NFW exploits the clever perspective of "simulating" the long waiting time if the conditionals for change of state space visited can be computed. The main problem involved is in computing this "change conditionals" which if not computed using exhaustive flips can result in "cycling" of the states i.e. same states get visited over and over again instead of visiting new states. LFIS avoids this by having explicit memory and running N copies of the sampler. After the run it computes importance weights to each of the unique states visited as follows:

$$w^{(i)} = \frac{p(\widetilde{f}^{(i)})}{\sum_{i=1}^{S_L} p(\widetilde{f}^{(i)})} \tag{23}$$

where $S_L \subset E^S$. These importance samples represent the posterior $p(f|S_L)$ instead of $p(f)$. Obviously if the algorithm is run long enough $S_L \to E^S$ and the samples represent the posterior $p(f)$.

5.2. Particle Filters with Static Prior

Another recent work that uses importance sampling for sampling from MRF is in [**?**]. Owing to the similarity to the notations used in robot mapping, they call it particle filters with static prior. There the idea is to sample states *sequentially* and in such a way that samples build-up the assignments bottom-up using clever conditionals. Below we summarize that approach using the terminology used in this paper. $f_{<1:t>}$ denotes t random variables are instantiated not necessarily s_1 through s_t. Also $p(f_{<t>}|f_{<1:t-1>})$ denotes the conditional probability distribution of instantiating a random variable which has not been instantiated so far.

Assume there are N particles. Let's take a closer look at the journey of one particle. By definition $f_{<1:m>} \equiv f$ since f has m random variables. Since there are only m random variables the equivalent of number of iterations (T) in Gibbs sampler needed is only m. The importance samples $\{(f_{<1:m>}^{(i)}, w_{<1:m>}^{(i)})\}$ represent the posterior $p(f_{<m>}|f_{<1:m-1>}^{(i)})$. But there are three key observations:

- If N is large enough and the conditionals are designed

Algorithm 3 Large Flip Importance Sampling

Initialize a sample: $\widehat{f}^{(0)} = \left(s_1^{(0)} = e_1^{(0)}, s_2^{(0)} = e_2^{(0)}, \ldots, s_m^{(0)} = e_m^{(0)}\right)$.
$k = 1$.
$\Gamma_k \sim \mathcal{U}[\Gamma_{\min}, \Gamma_{\max}]$
$n = 0$.
$F_0^k = \phi$.
for $i = 1$ to T **do**
  Draw $\widehat{f}^{(i)} \sim \nu(j, e, \widehat{f}^{(i-1)})$.
  $\mathcal{F}_n^k = \mathcal{F}_{n-1}^k \cup \widehat{f}^{(i)}$.
  $n = n + 1$.
  if $n = \Gamma_k$ **then**
    $k = k + 1$.
    $\Gamma_k \sim \mathcal{U}[\Gamma_{\min}, \Gamma_{\max}]$
    $n = 0$.
    $F_0^k = \phi$.
  end if
end for
Pick all the *unique* states, $\{\widetilde{f}^{(i)}\}_{i=1}^{T}$, from $\{\mathcal{F}_{\Gamma_k}^k\}$ for all available k.
Assign importance weights to those unique states as per Eq. (23).
Return $\{(\widetilde{f}^{(1)}, w^{(1)}), (\widetilde{f}^{(2)}, w^{(2)}), \ldots, (\widetilde{f}^{(T)}, w^{(T)})\}$.

Algorithm 4 Particle filter with static prior

for $i = 1$ to N **do**
  Draw $f_{<1>}^{(i)} \sim p(f_{<1>})$.
  Set $w_{<1>}^{(i)} = p(f_{<1>}^{(i)})$.
  Draw $f_{<2>}^{(i)} \sim p(f_{<2>}|f_{<1>}(i))$.
  Set $f_{<1:2>}^{(i)} = \{f_{<2>}^{(i)}, f_{<1>}^{(i)}\}$.
  Set $w_{<1:2>}^{(i)} = w_{<1>}^{(i)} * p(f_{<2>}^{(i)}|f_{<1>}(i))$.
  Draw $f_{<m>}^{(i)} \sim p(f_{<m>}|f_{<1:m-1>}^{(i)})$.
  Set $f_{<1:m>}^{(i)} = \{f_{<m>}^{(i)}, f_{<1:m-1>}^{(i)}\}$.
  Set $w_{<1:m>}^{(i)} = w_{<1:m-1>}^{(i)} * p(f_{<m>}^{(i)}|f_{<1:m-1>}^{(i)})$.
end for
Return $\{(f_{<1:m>}^{(1)}, w_{<1:m>}^{(1)}), (f_{<1:m>}^{(2)}, w_{<1:m>}^{(2)}), \ldots, (f_{<1:m>}^{(N)}, w_{<1:m>}^{(N)})\}$.

properly the samples would represent non-trivial patterns in the posterior and since in [**?**] the goal was to identify good matching between model and image contours the samples were able to serve the purpose.

- The samples are weighted *incrementally* using recursive importance weighting unlike LFIS where the samples are weighted at the end of the sampling process. For derivations of the recursive importance weighting please refer to [**?**].

- Although the samples are not generated in a standard MCMC fashion like Gibbs sampler the importance weighting and the conditional distributions result in useful samples. This is the key basis for LFIS also.

Acknowledgments

This work was supported in part by NSF IIS-0812118 and by DOE DE-FG52-06NA27508 grants. N. Adluru is supported by CIBM and MIR at the UW-Madison.

References

[1] J. Carpenter, P. Clifford, and P. Fearnhead. Building robust simulation-based filters for evolving data sets. Technical report, Dept. of Statistics, University of Oxford, 1999. 4

[2] N. Gordon, D. Salmond, and A. Smith. Novel approach to nonlinear/non-gaussian bayesian state estimation. In *Radar and Signal Processing, IEE Proceedings of*, volume 140, pages 107–113, April 1993. 4

[3] F. Hamze and N. de Freitas. Hot coupling: A particle approach to inference and normalization on pairwise undirected graphs of arbitrary topology. In *NIPS*, pages 1–8, 2005. 4

[4] S. Thrun, W. Burgard, and D. Fox. *Probabilistic Robotics*. The MIT Press Cambridge, 2005. 3