

Machine Learning for Computer Vision

July 17, 2018

Topic: HMMs, Variational Inference and Sampling methods

Exercise 1: Gibbs sampling

Show that the Gibbs sampling algorithm satisfies detailed balance:

$$p(z)T(z, z') = p(z')T(z', z)$$

This follows from the fact that in Gibbs sampling, we sample a single variable, z_k at a time, while all other variables, $z_{-k} = \{z_i\}_{i \neq k}$, remain unchanged ($z'_{-k} = z_{-k}$). We denote as $T(z, z')$ the transition probability from z to z' and we get

$$\begin{aligned} p(z)T(z, z') &= p(z_k, z_{-k})p(z'_k|z_{-k}) && \text{(Joint probability)} \\ &= p(z_k|z_{-k})p(z_{-k})p(z'_k|z_{-k}) && \text{(Product Rule)} \\ &= p(z_k|z'_{-k})p(z'_{-k})p(z'_k|z'_{-k}) && (z_{-k} = z'_{-k}) \\ &= p(z_k|z'_{-k})p(z'_k, z'_{-k}) && \text{(Product Rule)} \\ &= T(z', z)p(z') && \text{(Joint probability),} \end{aligned}$$

where we have used the transition probability in Gibbs sampling $T(z, z') = p(z'_k|z_{-k})$.

Exercise 2: Particle Filter

a) *What kind of spaces can we explore with a particle filter?*

With particle filters we can explore continuous state spaces.

b) *What kind of distributions can we approximate with a particle filter?*

Particle filter is non-parametric, meaning we can approximate arbitrary (multi-modal) distributions. Given enough particles we can approximate any function.

c) *In a Monte Carlo localization problem what do the particles and the particle weights correspond to?*

The particles themselves correspond to the motion model as they represent the state after motion with noise. The particle weights are computed according to the measurement model so they represent the likelihood of a measurement.

d) **Programming** : *Implement a particle filter for global localization.*

See code.