A Markov-Chain Monte Carlo Approach to Simultaneous Localization and Mapping

Péter Torma
GusGus AB
Hungary

András György
Machine Learning Res. Group
MTA SZTAKI, Hungary

Csaba Szepesvári, Roshan Shariff
Dept. of Computing Sciences
University of Alberta, Canada

As is well-acknowledged in the literature, Monte Carlo methods are sometimes non-trivial to set up. In this case study, we apply Markov-chain Monte Carlo (MCMC) techniques to the Bayesian version of the Simultaneous Localization and Mapping (SLAM) problem from robotics. The SLAM posterior over robot trajectories and landmark positions in an unknown environment has high dimensionality and can be multimodal and/or highly peaked, which makes it challenging to sample from. Thus the use of a simple Gibbs sampler, though shown to be possible in the literature, is highly inefficient. Assuming a prior distribution over the landmark positions, we present an efficient, provably consistent, and rapidly converging algorithm to sample from the SLAM posterior (for arbitrary dynamics and observation models) by representing it as the stationary distribution of a Markov process.

Previous SLAM algorithms either assume linearized models and Gaussian noise (e.g. EKF-SLAM) or use particle filters (e.g. FastSLAM). Since real-world robots usually have non-linear dynamics, the former approach may fail to converge. Particle filters approximate the posterior distribution over robot trajectories with a set of samples drawn from it. This fails when the certainty in the robot's position increases sharply, which happens when the robot observes a landmark it has seen earlier ("a loop is closed") and is thus able to correct accumulated errors in its position. If no sample in the set lies close to the true trajectory, the particle set collapses to an inaccurate estimate.

The new algorithm constructs an *inference graph* in which the vertices represent robot and landmark locations. An assignment of values to the vertices induces a unique labeling of the edges, representing the geometric relationships between the vertex values. The many constraints, however, make most edge labelings inconsistent: the posterior is restricted to consistent labelings. By randomly selecting and resampling an edge label from a spanning tree of the graph, we are able to sample from a proposal distribution that maintains the consistency property. The robot dynamics and observation models, together with the observations and control inputs, assign a probability density to any labeling of the inference graph. We then use a generalized Metropolis-Hastings algorithm to sample a consistent edge labeling, from which the robot trajectory and landmark locations can be recovered. In the limit, the map and trajectory thus determined follow the SLAM posterior distribution. Since the graph stores the entire history of observations and controls, the edge resampling technique is able to correct errors arbitrarily far back in time, thereby successfully tracking the robot's position even after long loops have been closed.

We describe how the structure of the SLAM inference graph makes this edge resampling algorithm efficient, and suggest implementation strategies. Since the publication of the algorithm in [1], we have improved its running time with a caching technique that exploits the nature of the SLAM problem. We have also run more experiments demonstrating that the resulting Markov process converges rapidly to the correct map and trajectory estimates. Thus MCMC methods allow the algorithm to solve the problem without compromising correctness or losing convergence guarantees.

References

[1] P. Torma, A. György, and Cs. Szepesvári. A Markov-chain Monte Carlo approach to simultaneous localization and mapping. In *Proc. of the 21th International Conference on Artificial Intelligence and Statistics, JMLR: W&CP*, volume 9, pages 605–612, 2010.