A Sample of Monte Carlo Methods in Robotics and Vision

Frank Dellaert
College of Computing
Georgia Institute of Technology

Microsoft Research May 27 2004

Credits

- Zia Khan
- Tucker Balch
- Michael Kaess
- Rafal Zboinski
- Ananth Ranganathan

Outline

The CPL and BORG Labs

Computational Perception Lab @Georgia Tech

Aaron Bobick, Frank Dellaert, Irfan Essa, Jim Rehg, andThad Starner

Other vision faculty In ECE: Ramesh Jain, Allen Tennenbaum, Tony Yezzi

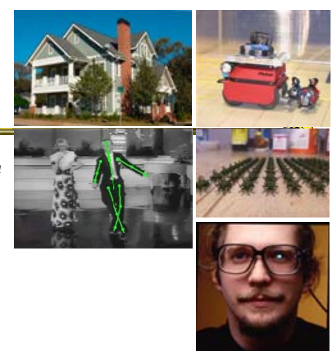

Structure from Motion…

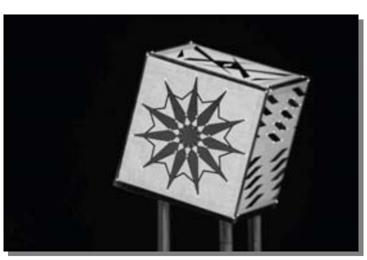

…without Correspondences

CVPR 2000, NIPS, Machine Learning Journal

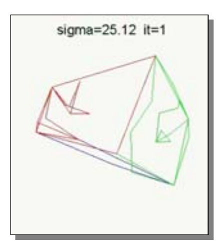

Current Main Effort: 4D Atlanta

4D Atlanta

- Idea:
  - Take 10000 images over 100 years
  - Build a 3D model with a time slider
- 2 PhD Students
- 2 MSc Students

- Assumptions about urban scenes (Manhattan), Symmetry (a la Yi Ma), Grammar-based inference, Markov chain Monte Carlo

Atlanta World

- CVPR 2004 Poster, with Grant Schindler

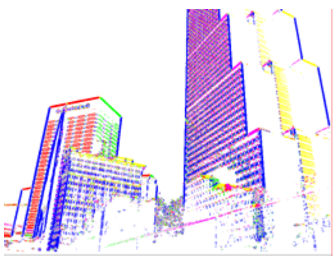

The BORG lab

- With Tucker Balch, Thad Starner

Real-Time Urban Reconstruction

•4D Atlanta, only real time, multiple cameras ☺
•Large scale SFM: closing the loop

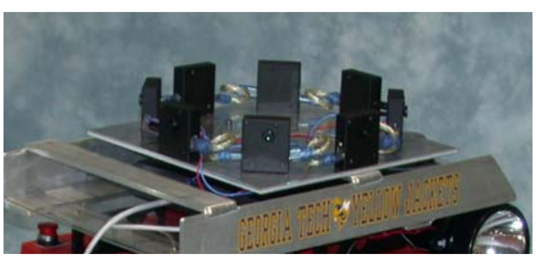

The Biotracking Project: Tracking Social Insects

Overview

- Influx of probabilistic modeling and inference…

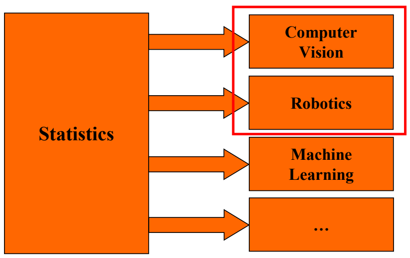

A sample of Methods

- Particle Filtering (Bootstrap Filter)
  - Monte Carlo Localization
- MCMC
  - Multi-Target Tracking
- Rao-Blackwellization
  - EigenTracking
- MCMC + RB
  - Piecewise Continuous Curve Fitting
  - Probabilistic Topological Maps

Monte Carlo Localization

1D Robot Localization

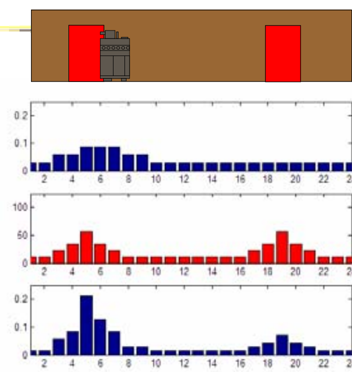

Prior P(X)

Likelihood L(X;Z)

Posterior P(X|Z)

Importance Sampling

- Densities are decidedly non-Gaussian
- Histogram approach does not scale
- Monte Carlo Approximation
- Sample from P(X|Z) by:
  - sample from prior P(x)
  - weight each sample $x^{(r)}$ using an importance weight equal to likelihood $L(x^{(r)};Z)$

1D Importance Sampling

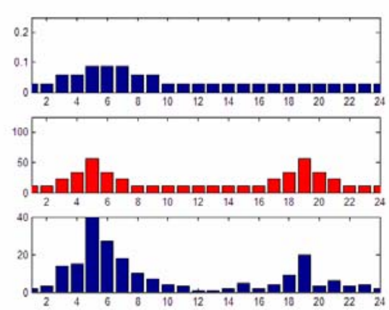

Sampling Advantages

- Arbitrary densities
- Memory = O(#samples)
- Only in "Typical Set"
- Great visualization tool !

- minus: Approximate
- Rejection and Importance Sampling do not scale to large spaces

Bayes Filter and Particle Filter

Recursive Bayes Filter Equation:

Motion Model

$$P(X_t|Z^t) = kP(Z_t|X_t)\int_{X_{t-1}} P(X_t|X_{t-1})P(X_{t-1}|Z^{t-1})$$

Predictive Density

Monte Carlo Approximation:

$$P(X_t|Z^t) \approx kP(Z_t|X_t)\sum_r \pi_{t-1}^{(r)} P(X_t|X_{t-1}^{(r)})$$

Particle Filter

Empirical predictive density = Mixture Model

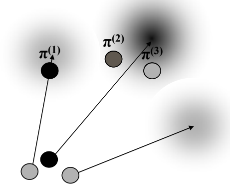

$$\pi_t^{(s)} = P(Z_t|X_t^{(s)})$$

First appeared in 70's, re-discovered by Kitagawa, Isard, …

3D Particle filter for robot pose: Monte Carlo Localization

Dellaert, Fox & Thrun ICRA 99

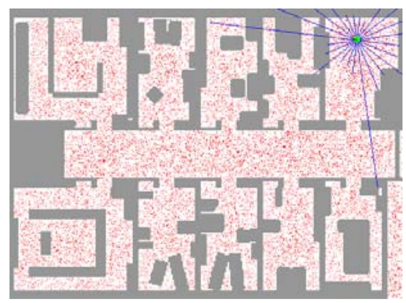

Multi-Target Tracking

An MCMC-Based Particle Filter for Tracking Multiple, Interacting Targets, ECCV 2004 Prague, With Zia Khan & Tucker Balch

Motivation

- How to track many INTERACTING targets ?

Traditional Multi-Target Tracking

- In essence: curve fitting !

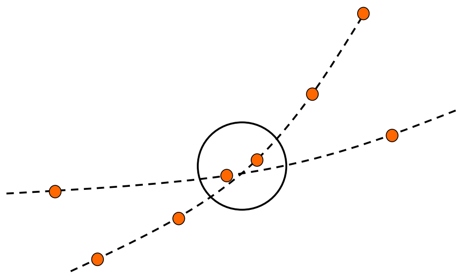

Ants are not Airplanes !

- Interaction changes behavior

Results: Vanilla Particle Filters

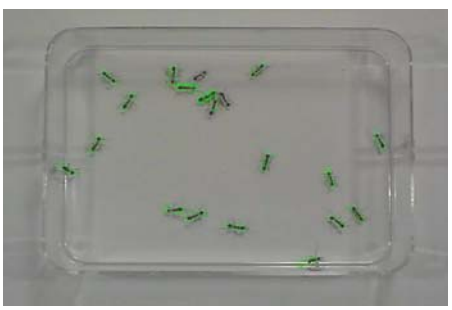

Our Solution: MRF Motion Model

- MRF = Markov Random Field, built on the fly

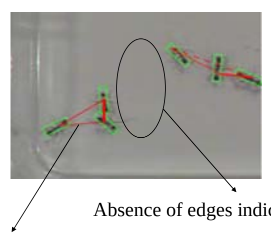

Absence of edges indicates **no interaction**

Edges indicate **interaction**

MRF Interaction Factor

- Pairwise MRF:

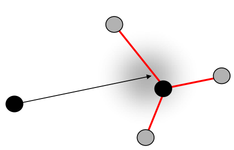

$$P(X_t|X_{t-1}) \propto \prod_i P(X_{it}|X_{i(t-1)}) \prod_{ij\in E} \psi(X_{it}, X_{jt})$$

Joint MRF Particle Filter

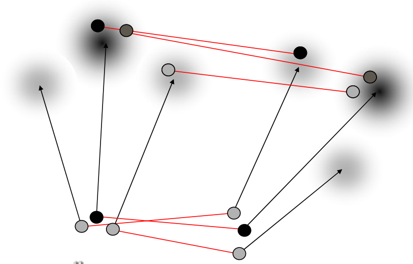

$$\pi_t^{(s)} = \prod_{i=1}^{n} P(Z_{it}|X_{it}^{(s)}) \prod_{ij\in E} \psi(X_{it}^{(s)}, X_{jt}^{(s)})$$

Results: Joint MRF Particle Filter

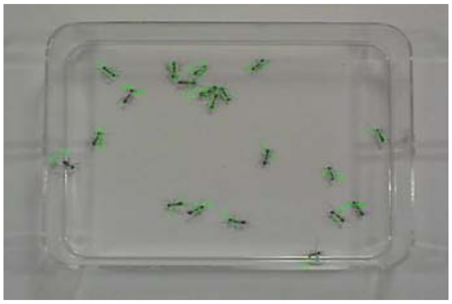

Solution: Marlov Chain Monte Carlo

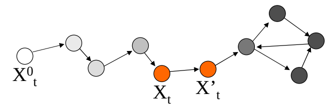

- Start at X^0_t
- Propose a move $Q(X'_t|X_t)$
- Calculate acceptance ratio

  a = $Q(X_t \mid X'_t)\, p(X_t) / Q(X'_t \mid X_t)\, p(X_t)$
- If a>=1, accept move

  otherwise only accept move with probability a

MCMC Particle Filter

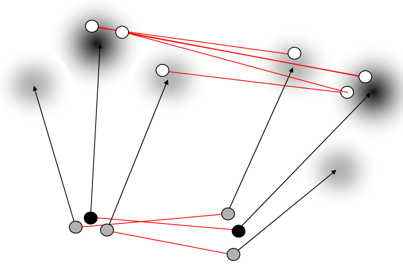

Results: MCMC

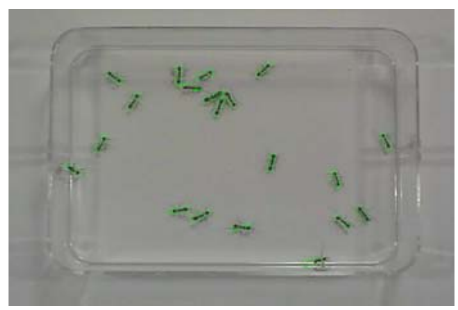

Quantitative Results (10K frames)

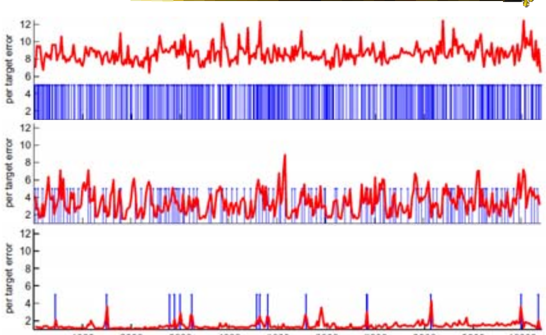

Rao-Blackwellized EigenTracking

Coming CVPR 2004 Talk,
With Zia Khan and Tucker Balch

Motivation

- Honeybees are more challenging ☺
- Eigenspace Representation:

- Generative PPCA Model (Tipping&Bishop)
- Learned using EM, from 146 color images of bees

Particle Filter

- Added dimensionality = problem
- Solution: integrate out PPCA coefficients

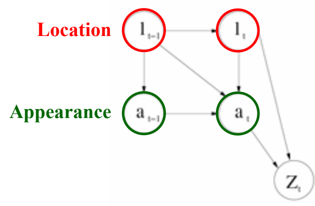

Marginal Bayes Filter

- Bayes filter for location and appearance

$$P(l_t, a_t|Z_t) = kP(Z_t|l_t, a_t)\times \int_{l_{t-1}}\int_{a_{t-1}} P(l_t, a_t|l_{t-1}, a_{t-1})P(l_{t-1}, a_{t-1})$$

- Marginalized to location only:

$$P(l_t|Z_t) = k\int_{a_t} P(Z_t|l_t, a_t)\times \int_{l_{t-1}}\int_{a_{t-1}} P(l_t, a_t|l_{t-1}, a_{t-1})P(l_{t-1}, a_{t-1})$$

Rao-Blackwellized Filter

- Hybrid approximation:

$$P(l_{t-1}, a_{t-1}) = P(l_{t-1})P(a_{t-1}|l_{t-1}) \approx \sum w_{t-1}^{(i)}\delta(l_{t-1}^{(i)})\alpha_{t-1}^{(i)}(a_{t-1})$$

$$\alpha_{t-1}^{(i)}(a_{t-1}) \triangleq P(a_{t-1}|l_{t-1}^{(i)}, Z^{t-1})$$

- Location is sampled
- Each sample carries a conditional Gaussian over the appearance coefficients
- Marginalization with PPCA is very efficient

Simplified Filter

- Dynamic Bayes Net:

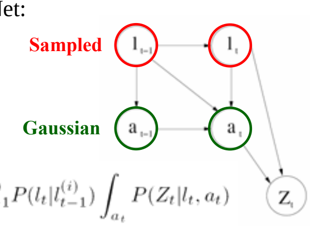

Approximation:

$$P(l_t|Z_t) \approx k\sum_i w_{t-1}^{(i)}P(l_t|l_{t-1}^{(i)})\int_{a_t} P(Z_t|l_t, a_t) \times \int_{a_{t-1}} P(a_t|l_t, l_{t-1}^{(i)}, a_{t-1})\alpha_{t-1}^{(i)}(a_{t-1})$$

Sampled

q=0

q=20

Dancers, q=10, n=500

Piecewise Continuous Curve-Fitting

ECCV 2004 Prague, with Michael Kaess and Rafal Zboinski

Reconstructing Objects with Jagged Edges

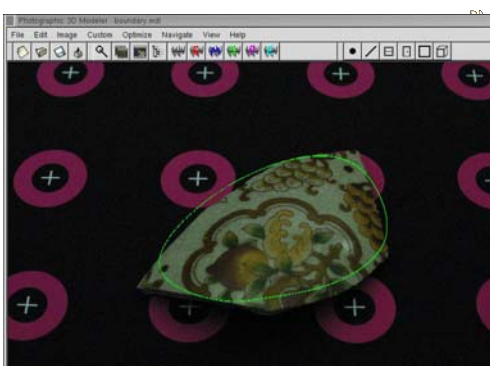

Subdivision Curves

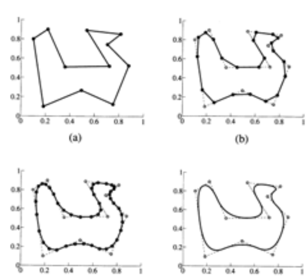

Tagged Subdivision Curves

- Hughes Hoppe paper: piecewise smooth surface fitting

- In this context:
  - 3D tagged subdivision curves

Tagged Curve Example

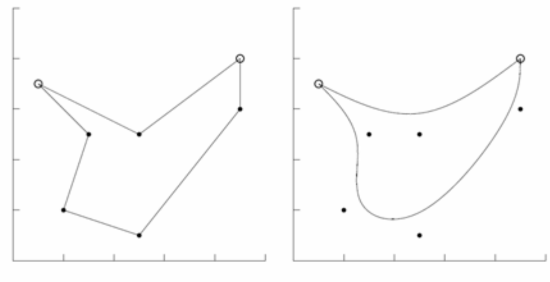

Rao-Blackwellized Sampling

- MCMC sampling over discrete tag configurations
- For each sample: optimize over control points
- Approximate mode by a Gaussian
- Marginalize Analytically

$$P(T|Z) = \int P(\Theta, T|Z) d\Theta$$

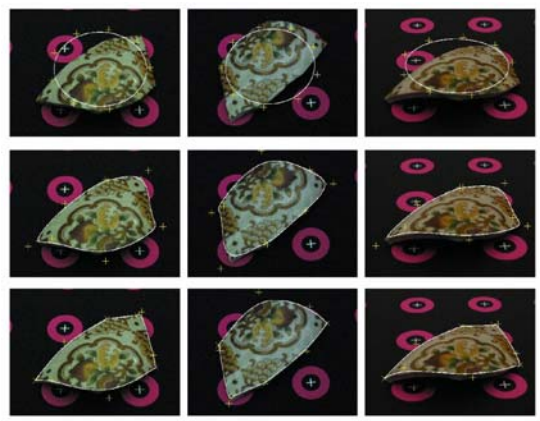

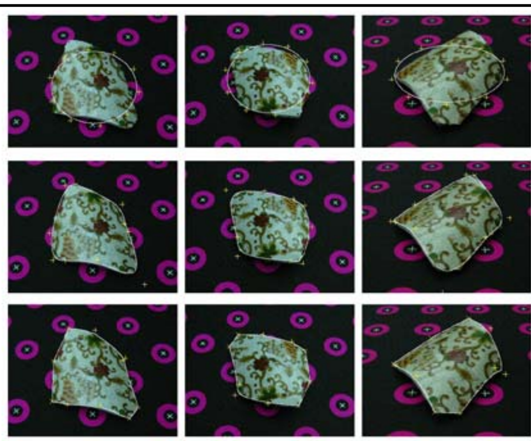

Marginals

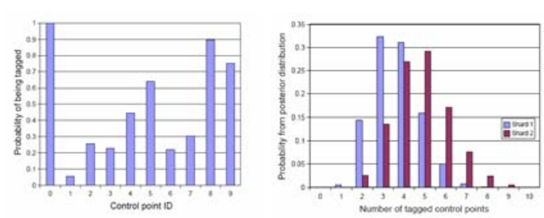

Probabilistic Topological Maps

Submissions to IROS, NIPS,
With Ananth Ranganathan

Motivation

- Metric Maps
- Topological Maps
- How to reason about topology given incomplete or noisy observations ?

Problem

Odometry measurements are noisy:

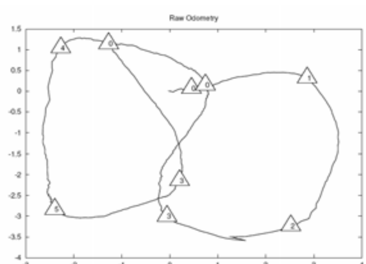

Correct Topology and ML Path

Given ground truth topology, calculate ML path:

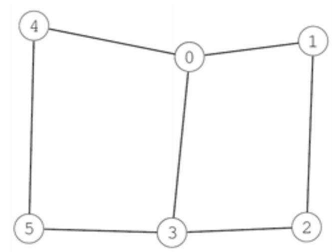

Probabilistic Topological Maps

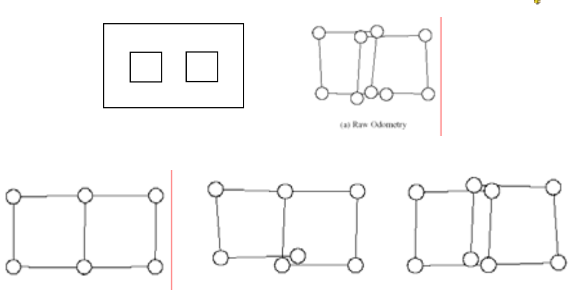

(a) Raw Odometry

(b) 28.8 %

(c) 23.0 %

(d) 10.8 %

Set Partitions

- Topologies ⇔ Set Partitions

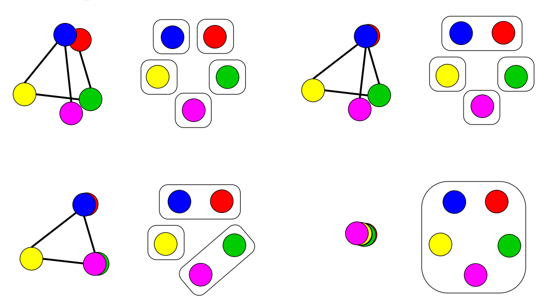

Bell numbers

- 1, 2, 5, 15, 52, 203, 877, 4140, 21147, 115975
- Combinatorial explosion !
- Idea: use MCMC Sampling over topologies

MCMC Proposal

- Pick k at random, assign it to group t in 1..m
- Some possibilities:

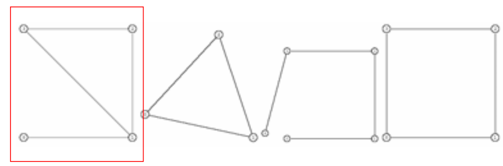

Acceptance Ratio

- Pick k at random, assign it to group t in 1..m

1. k is in a singleton
   (a) $|S_t| = 0$:
      i. k is put in a group by itself again, $n' = n$
      ii. $q(T'|T) = \frac{m-n+1}{m}$
      iii. Reverse case is 1(a), with $q(T|T') = \frac{m-n'+1}{m}$, and hence $r = \frac{m-n'+1}{m-n+1} = 1$
   (b) $|S_t| \neq 0$:
      i. k joins a non-empty group, $n' = n - 1$
      ii. $q(T'|T) = \frac{1}{m}$
      iii. Reverse case is 2(a), with $q(T|T') = \frac{m-n'}{m}$, and hence $r = \frac{m-n'}{1} = \frac{m-n+1}{1}$
2. k is *not* in a singleton
   (a) $|S_t| = 0$:
      i. k is put in a group by itself, $n' = n + 1$
      ii. $q(T'|T) = \frac{m-n}{m}$
      iii. Reverse case is 1(b), with $q(T|T') = \frac{1}{m}$, and hence $r = \frac{1}{m-n} = \frac{1}{m-n}$
   (b) $|S_t| \neq 0$:
      i. e joins another non-empty group, $n' = n$
      ii. $q(x'|x) = \frac{1}{m}$
      iii. Reverse case is 2(b), with $q(x|x') = \frac{1}{m}$, and hence $r = \frac{1}{1} = 1$

Rao-Blackwellized Sampling

- MCMC sampling over discrete tag configurations
- For each sample: optimize over robot trajectory
- Approximate mode by a Gaussian
- Marginalize Analytically

$$P(T|Z) = \int P(\Theta, T|Z)d\Theta$$

Results

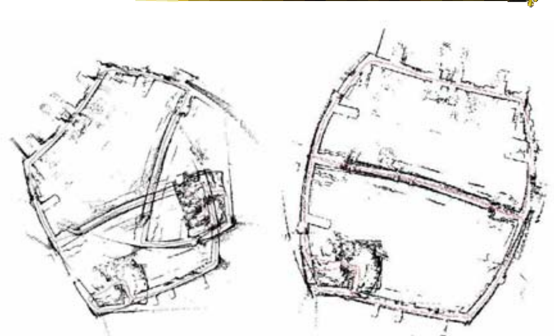

The End