IROS 2005 Advanced Tutorial on

SLAM - Getting it Working in Real World Applications

Tuesday, 2. August 2005.

Shaw Conference Centre, Salons 17/18

Organised by: Martin Adams

Contributors: **Martin Adams, Juan Nieto, Eduardo Nebot, Jose Neira, Carlos Estrada, Juan D. Tardos, Matthew Walter, John Leonard, Cyrill Stachniss, Wolfram Burgard, Agostino Martinelli, Roland Siegwart, Jonghyuk Kim, Salah Sukkarieh.**

IROS 2005 Advanced Tutorial on

SLAM - Getting it Working in Real World Applications

Schedule: Tuesday, 2. August 2005.

Shaw Conference Centre, Salons 17/18

08.45 Tutorial presenters meet.

09.00 – 09.40 “An Introduction to SLAM – Current Issues” (Martin Adams)

09.40 – 10.20 “Representation of Dense Information in Large Areas” (Juan Nieto, Eduardo Nebot)

10.20 – 10.40 Coffee Break

10.40 – 11.20 “Multivehicle mapping in Large Environments” (Jose Neira, Carlos Estrada and Juan D. Tardos)

11.20 – 12.00 “Sparse Information Based Representations” (M, Walter, John Leonard)

12.00 – 13.30 Lunch

13.30 – 14.10 “Rao-Blackwellized Particle Filters and Loop Closing” (Cyrill Stachniss, Wolfram Burgard)

14.10 – 14.50 “Combining Metric and Topological maps” (A. Martinelli, Roland Siegwart)

14.50 – 15.10 Coffee Break

15.10 – 15.50 “UAV Navigation: Airborne Inertial SLAM” (J. Kim, Salah Sukkareih)

15.50 – 16.30 “Rich Representations with RADAR and LADAR Based SLAM” (Martin Adams)

16.30 End.

An Introduction to SLAM

Martin D. Adams

School of Electrical and Electronic Engineering,

NTU,

Singapore

Email : *eadams@ntu.edu.sg*

Overview

- Brief History of SLAM.
- Estimation Techniques Used in SLAM.
- Corresponding SLAM Methods.
- Problems in SLAM – Convergence Scaling.
- Some Provocative Questions!

Brief History – A Classical Problem

- Researchers (1984 onwards):
  - Brooks, Chatila & Laumond, Moravec & Elfes, Crowley, Faugeras & Ayache, Smith & Cheeseman, Leonard & Durrant-Whyte, Moutarlier & Chatila, Cox, Simmons, Castellanos, Tardos & Neira, Fox, Thrun, Burgard, Gutman, Christensen, Jensfelt, Nebot, Newman, ACFR Group, Dissanayake, Julier & Uhlmann, Davison, Adams,Lacroix, Burgard …..

Brief History – Smith, Self & Cheeseman

- Seminal Paper "Estimating Uncertain Spatial Relationships in Robotics":
  - "This paper applies state estimation theory to the problem of estimating parameters of an entire spatial configuration of objects, with the ability to transform estimates into any frame of interest."
  - ASSUMPTIONS:
    - "…we linearize inherently non-linear relationships." (Angular errors are small).
    - "Estimating only two moments of the probability density functions of the uncertain spatial relationships is adequate for decision making."

Estimation Techniques

- The Extended Kalman Filter (EKF).
- Particle Filters.
- Distribution Approximation Filters (DAF) - UKF.
- Covariance Intersection (CI).

The Kalman Filter

KF is popular due to well known properties:

- **Optimality** – *the* optimal linear mean-squared error estimator.
- **Recursiveness** – In each iteration, only the current observations are used.
- **Linearity** – Estimate is obtained from a weighted linear sum of information from predictions and observations.
- **Covariance modelling** – Uncertainties associated with system behaviour and observations explicitly modelled using random variables with specified means and covariances.

The Extended Kalman Filter

- Success of Kalman filter depends on predictions.
- For Non-Linear system models, problem – how to predict state and its covariance.
- Potentially unbounded number of parameters required to model the entire distribution of the state errors.
- All filters use approximations – The simplest is the EKF.
- EKF predicts future state assuming that process and observation models linear on the scale of the error.

The Extended Kalman Filter

State evolves as:

$$\mathbf{x}(k+1) = f(\mathbf{x}(k), \mathbf{u}(k), \mathbf{v}(k), k)$$

Error Covariance results from linearisation of $f(..)$

$$\mathbf{P}(k+1|k) = \nabla_{\mathbf{x}} f \mathbf{P}(k|k) \nabla_{\mathbf{x}}^T f + \nabla_{\mathbf{v}} f \mathbf{Q}(k) \nabla_{\mathbf{v}}^T f$$

Linearisation justified if filter estimate close to true estimate.

The Extended Kalman Filter

EKF Problems:

- Necessary to evaluate Jacobian matrices for poccess and observation models.
- Linearisation can introduce significant errors.
- Good example: Julier & Uhlmann – simple conversion from polar to Cartesian coordinates.

Particle Filters

Instead of transforming only means and variances – transform individually drawn samples (Gordon, Salmond, Smith).

1. Draw random samples $\{x_{k-1}(i) : i = 1, \ldots, N\}$ from the PDF $p(x_{k-1} \mid z_{k-1})$

2. Draw random samples $\{w_{k-1}(i) : i = 1, \ldots, N\}$ from the PDF $p(w_{k-1})$

3. *Predict:* Pass each sample pair through non-linear system model:

$$x_k^*(i) = f_{k-1}(x_{k-1}(i), w_{k-1}(i))$$

4. *Update:* After measurement, find likelihood of each prior sample and obtain normalised weight:

$$q_i = p(z_k \mid x_k^*(i)) / \Sigma_{j=1}^{N} p(z_k \mid x_k^*(i))$$

Particle Filters

5. Define distribution over $\{x_k^*(i) : i = 1,\dots,N\}$
   With prob. Mass q_i associated with element *i.*

6. Resample *N* times to generate samples $\{x_k(i) : i = 1,\dots,N\}$
   so that for any *j*

$$P(x_k(j)) = P(x_k^*(i)) = q_i$$

Problem: *How many particles are necessary for satisfactory performance?*

Distribution Approximation Filter

Also called Unscented Kalman Filter (UKF): (Julier, Uhlmann)

- As with Particle Filter – transform points individually.
- Calculate predicted mean and covariance from transformed samples.
- *Difference* – Initial points **not** chosen at random – they are symmetrically distributed points (*sigma points*) chosen to capture specific information about the distribution.

Distribution Approximation Filter

Advantages over EKF:

- Not necessary to find Jacobian matrices.
- Reaches the accuracy of the second order Gauss filter.

Remark:

- Choice of “Sigma Points” requires knowledge of prior distribution and a tuning parameter (kappa)

The SLAM Problem

- The state space for a SLAM Kalman filter stores all the vehicle and all the beacon states in a single state space [6]

$$\hat{\mathbf{x}}(k \mid k) = \left[\hat{\mathbf{x}}_v^T(k \mid k) \ldots \hat{\mathbf{p}}_N^T(k \mid k)\right]^T$$

$$\mathbf{P}(k \mid k) = \begin{pmatrix} \mathbf{P}_{vv}(k \mid k) & \mathbf{P}_{v1}(k \mid k) & \ldots & \mathbf{P}_{vN}(k \mid k) \\ \mathbf{P}_{1v}(k \mid k) & \mathbf{P}_{11}(k \mid k) & \ldots & \mathbf{P}_{1N}(k \mid k) \\ \mathbf{P}_{2v}(k \mid k) & \mathbf{P}_{21}(k \mid k) & \ldots & \mathbf{P}_{2N}(k \mid k) \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{P}_{Nv}(k \mid k) & \mathbf{P}_{N1}(k \mid k) & \ldots & \mathbf{P}_{NN}(k \mid k) \end{pmatrix}$$

- Storage costs are $O(N^2)$, computational costs are $O(N^3)$
- Therefore, the full Kalman can only be used for a small number of beacons

Suboptimal Fusion Methods

- Most of the computational and storage costs are taken up by storing and updating the cross terms
- These are required for the *optimal* solution
- Therefore, if most of the cross terms can be eliminated, the the computational and storage costs can be reduced or even made constant time
- However, the cross terms encode the informational relationship between the states: changing their values can cause the filter to appear to have more information than is truly available [2].
- Principled suboptimal data fusion algorithms must be used

The Geometry of Kalman Filter Updates

- Given an estimate $\mathbf{a}$ with covariance $\mathbf{P}_{aa}$ and observation $\mathbf{b}$ with covariance $\mathbf{P}_{bb}$
- For any (known) cross correlation, the covariance ellipse of the updated estimate lies within the *intersection region* of the two covariance ellipses [7]
- This result applies even if $\dim(\mathbf{a}) \neq \dim(\mathbf{b})$

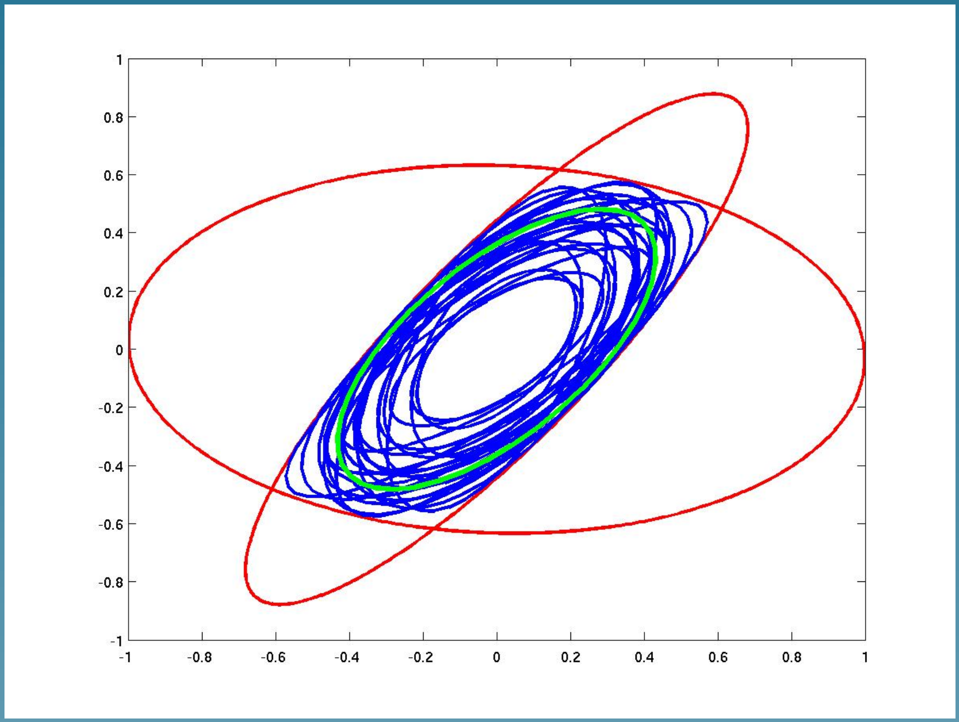

Covariance ellipses for $\mathbf{a}$ and $\mathbf{b}$ (red) and updates with various cross correlations (blue) and when independent (green)

Covariance Intersection (CI)

- The intersection region is given by

$$\mathbf{P}_{cc}^{-1} = \omega \mathbf{P}_{aa}^{-1} + (1-\omega)\mathbf{P}_{bb}^{-1}$$

  where $\omega \in [0, 1]$

- The degree of freedom defines a set of ellipses that pass through the intersection points

- ω should be chosen to minimise the uncertainty in $\mathbf{P}_{cc}$

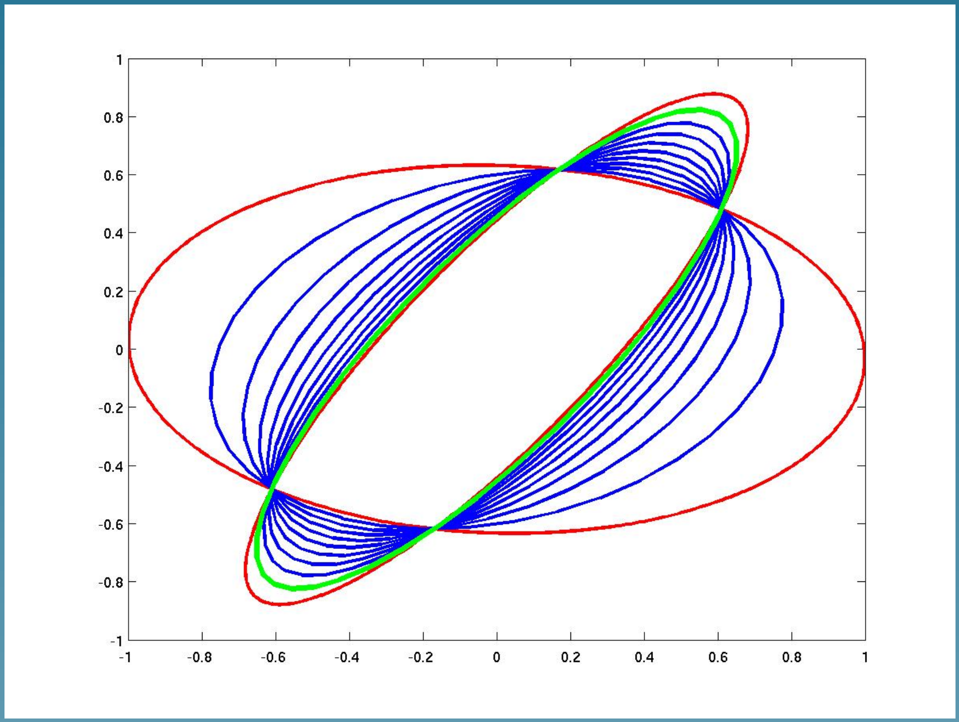

Covariance ellipses for CI with $\omega = [0 : 0.1 : 1]$ (blue) and minimum determinant value $\omega = 0.42$ (green)

Corresponding SLAM Methods

- ***EKF Based SLAM*** – Non-Linear Vehicle Kinematic and Sensor Measurement Models used to predict vehicle state and future observations. *Models linearised* to predict vehicle state and measurement uncertainties. (Chatila & Laumond, Crowley, Smith, Self & Cheeseman, Leonard & Durrant-Whyte, Cox)

- ***DAF/UKF Based SLAM*** – High order kinematic and dynamic vehicle models for outdoor, high speed vehicles used to implement UKF based SLAM. Model predictions carried out with the use of "Sigma Points" (Julier & Uhlmann).

- ***FAST SLAM*** – Particles (samples from a distribution) are used to estimate entire vehicle trajectories. Each particle has an *independent* EKF running for each land mark to estimate its position. *"Landmark estimates are conditionally independent given the robot's path."* (Thrun, Montemerlo, Koller, Wegbreit)

Other SLAM Related Methods

- ***Use of Relative Maps, ATLAS Framework***
  (P. Newman, J. Leonard).

- ***Symmetries & Purturbations (SP) Map***
  (J. Castellanos, J. Montiel, J. Neira, J. Tardos).

- ***Rao-Blackwellized Particle Filters***
  (C. Stachniss, W. Burgard).

- ***Scan Matching***
  (F. Lu, E. Milios).

- ***Sparse Extended Information Filters***
  (R. Eustice, H. Singh, J. Leonard)

Problems in SLAM

- ***Costs -*** Computational/Storage, Use of Information form based estimators.
- ***Sensing Issues*** – Interpretation of Sensor Data/Association.
- ***Environmental Representations*** – Feature based, grids, Sum-of-Gaussians, Scan Matching.
- ***Loop Closing*** – Recognising where you have been before.
- ***Scaling*** – Large Environments.
- ***Observability*** – Estimating a high dimensional state with low dimensional measurements.

Provocative Questions

1. “If SLAM is correctly formulated – Why is Loop Closing an issue at all?”
    - When a loop is closed – its simply another iteration of the SLAM filter – at any iteration, estimates can become inconsistent, if the filtering is incorrect – what makes “Loop Closing” special?

2. SLAM is an unobservable problem. Control engineers say that even trying to solve an unobservable problem is a waste of time!!! Why even try to solve SLAM?
    - Use or relative maps has shown that only by “anchoring” some features can SLAM become tractable – is that the “be all and end all” of SLAM”?