30 Years of SLAM:

John Leonard
Samuel C. Collins Professor of Mechanical and Ocean Engineering
Massachusetts Institute of Technology

With thanks to many many people…..

30 Years of SLAM:

A Personal Historical Perspective on 3 Decades of Mobile Robotics Research

John Leonard
Samuel C. Collins Professor of Mechanical and Ocean Engineering
Massachusetts Institute of Technology

With thanks to many many people…..

Outline

- 1985-2015
  - 30 Years of papers
  - Are the "old" questions answered?
  - How do we measure progress?
- Is SLAM "solved?"
  - If yes, how do we know it is solved?
  - If no, what are the open questions?

Q: Where am I?

Q: Where am I? A: Jenkin Building, Oxford

Q: Where am I? A: Jenkin Building, Oxford

Hartley and Zisserman, Cambridge University Press

Jenkin Building Basement, Circa 1989

Why is SLAM Difficult?

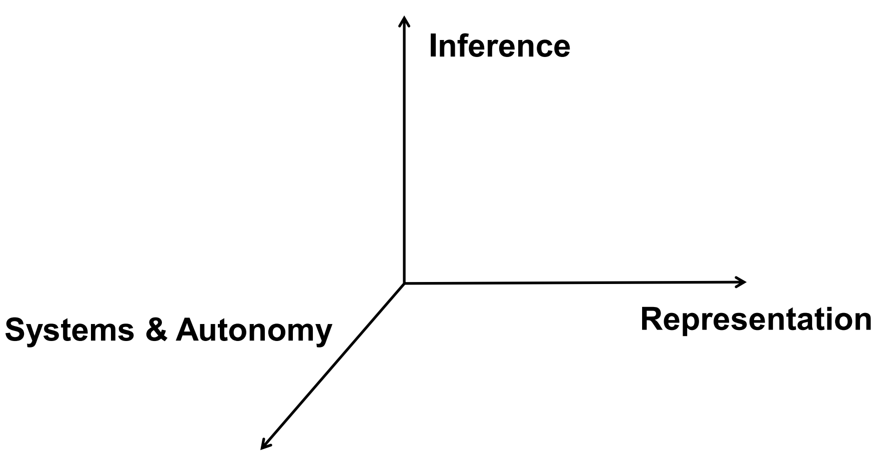

1985

Occupancy Grids
ICRA 1985

High Resolution Maps from Wide Angle Sonar

Hans P. Moravec Alberto Elfes
The Robotics Institute
Carnegie-Mellon University

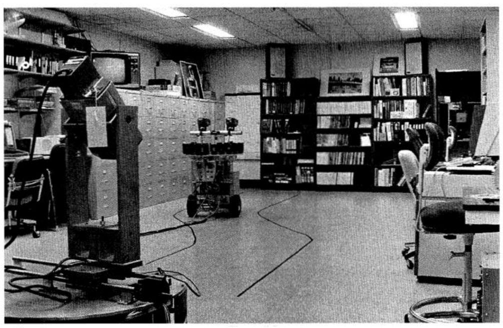

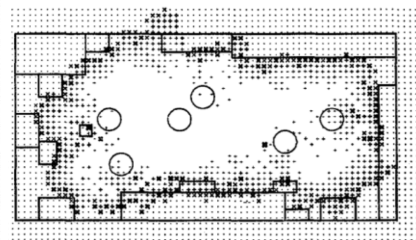

Figure 6: The Two-Dimensional Sonar Map After Thresholding.

1985

Visual Map Making for a Mobile Robot
Rodney Brooks, ICRA 1985

Visual Map Making for a Mobile Robot

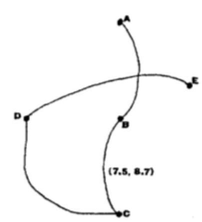

Rodney A. Brooks
MIT Artificial Intelligence Lab
545 Technology Square, Cambridge, Mass 02173.

Abstract. Mobile robots sense their environment and receive error laden readings. They try to move a certain distance and direction, and do so only approximately. Rather than try to engineer these problems away it may be possible, and may be necessary, to develop map making and navigation algorithms which explicitly represent these uncertainties, but still provide robust performance. The key idea is to use a relational map, which is rubbery and stretchy, rather than try to place observations in a 2-d coordinate system.

1. Introduction

We are interested in building mobile robot control systems useful for cheap robots (i.e., on the order of the price of an automobile) working in unstructured domains such as the home, material handling in factories, street cleaning, office and hotel cleaning, mining and agriculture. The same capabilities can be useful for robots, which do not have to be so cheap to be economically feasible, and which do tasks like: planetary exploration, space station maintenance and construction, asteroid mining, nuclear reactor operations, military reconaissance and general military operations.

2. Almost all mobile robot projects have had as one of their underlying assumptions that it is desirable to produce a world model in an absolute coordinate system. However all sensors and control systems have both systematic and random errors. The former can be dealt with by calibration techniques (although these are often time consuming and are confounded on mobile robots by the fact that the robot itself is not fixed to any coordinate system). The latter are always present. It is usual to model some worse case bounds on such errors but this will not always suffice (e.g. mismatches in stereo vision can produce depth measurements with error magnitude the full range of depths which can be measured). In any case the bounded errors at least must be dealt with in building models of the world and using them. A number of approaches have been taken to this problem:

a. Ignore it. This has only been successful in the most toylike of worlds.

b. Use fixed reference beacons. This implies that the environment is either structured for the robot's benefit in the case that beacons are explicitly installed, or that the environment has been pre-surveyed for the robot's benefit in the case that known positions of existing beacons (e.g. power outlets) are used.

"The key idea is to use a relational map, which is rubbery and stretchy, rather than to try to place observations in a 2-D coordinate frame.

Position Referencing and Consistent World Modeling for Mobile Robots, ICRA 1985

1985

POSITION REFERENCING AND CONSISTENT WORLD MODELING FOR MOBILE ROBOTS *

Raja Chatila and Jean-Paul Laumond

Laboratoire d'Automatique et d'Analyse des Systemes du CNRS
7, Avenue du Colonel Roche 31077 Toulouse Cedex, France

French Television, 1982

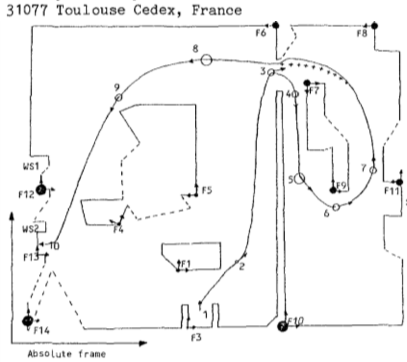

FIGURE 6: The final model with the associated errors on robot and frame position (the orientations errors are not represented)

Chatila and Laumond, ICRA 1985

A Stochastic Map for Spatial Relationships

1986

Smith, Self and Cheesemen

Proceedings of the Second Conference on Uncertainty in Artificial Intelligence, 1986

A Stochastic Map for Spatial Relationships

1986

Smith, Self and Cheesemen

Proceedings of the Second Conference on Uncertainty in Artificial Intelligence, 1986

"Rather than treat spatial uncertainty as a side issue in geometrical reasoning, we believe it must be an intrinsic part of spatial representations. In this paper, we describe a representation for spatial information, called the stochastic map, and associated procedures for building it, reading information from it, and **revising it incrementally** *as new information is obtained."*

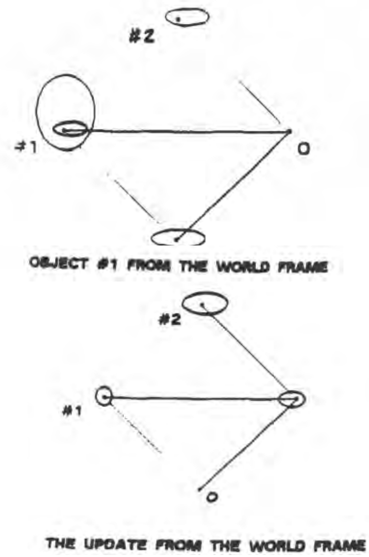

Consistent Integration and Propagation of Disparate Sensor Observations
Hugh Durrant-Whyte

1987

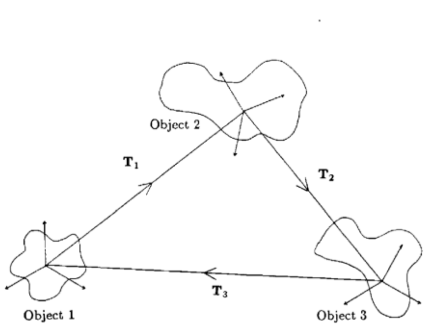

Fig. 1. Three objects in a relational loop.

The International journal of robotics research 6 (3), 3-24

Tracking and Data Association
Bar-Shalom and Fortmann, Academic Press, 1988

1988

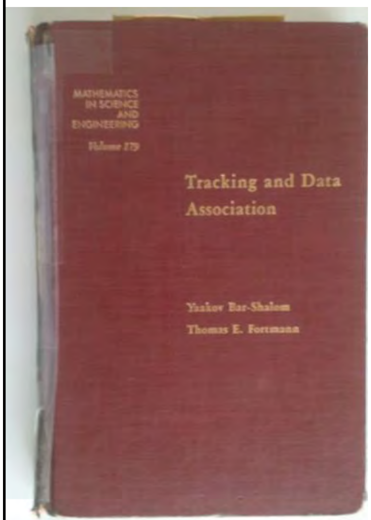

Tracking and Data Association
Bar-Shalom and Fortmann, Academic Press, 1988

1988

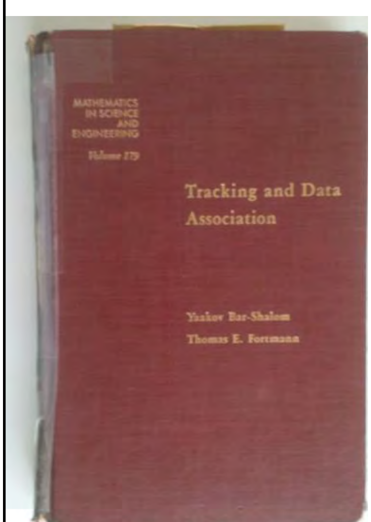

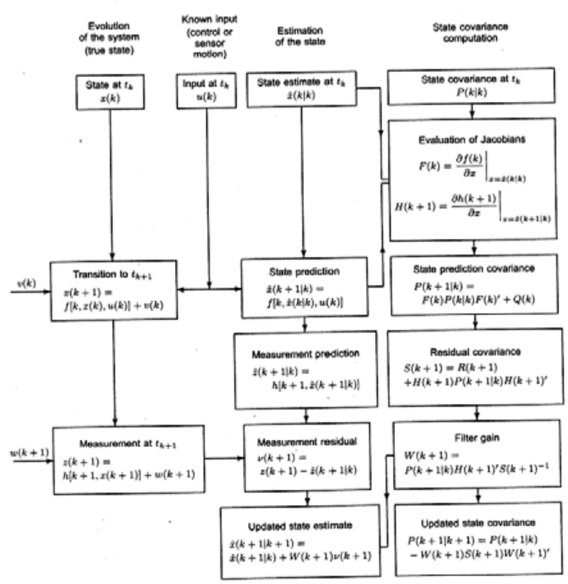

Localization with an *a priori* map

1990

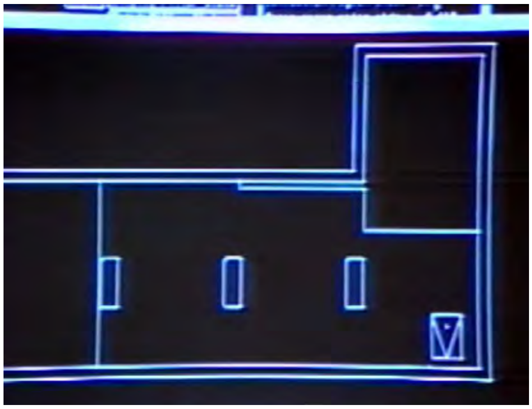

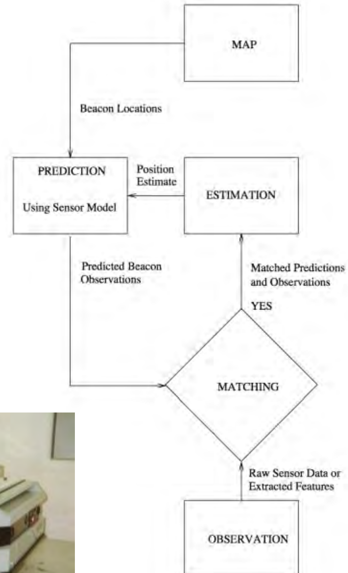

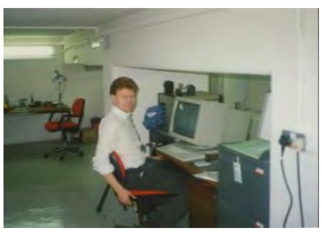

Jenkin Building Basement, Oxford, 1990

Mapping from Known Locations

1990

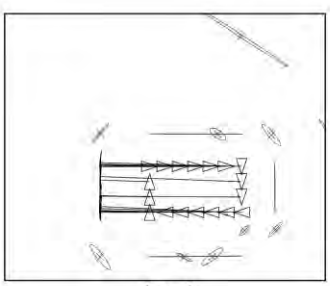

cluster 4

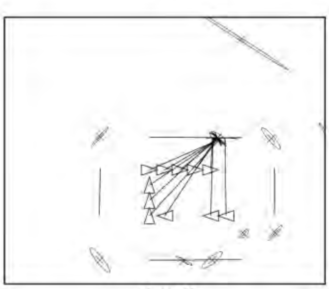

cluster 6

The First Complete SLAM Implementation!

1991

The First Complete SLAM Implementation!
Philippe Moutarlier, LAAS

1991

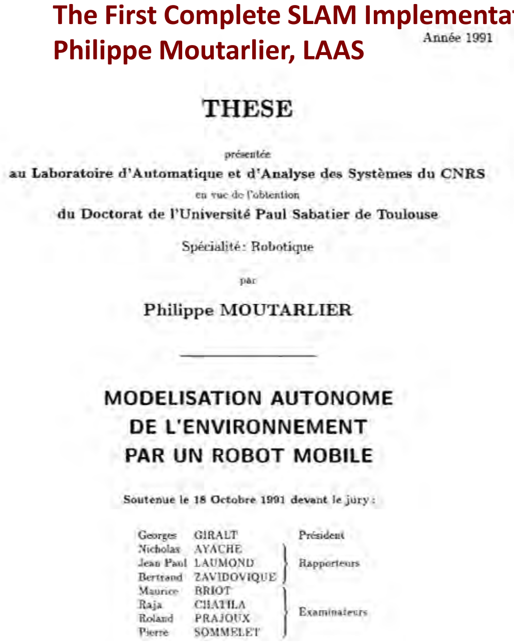

Année 1991

THESE

présentée

au Laboratoire d'Automatique et d'Analyse des Systèmes du CNRS

en vue de l'obtention

du Doctorat de l'Université Paul Sabatier de Toulouse

Spécialité: Robotique

par

Philippe MOUTARLIER

MODELISATION AUTONOME
DE L'ENVIRONNEMENT
PAR UN ROBOT MOBILE

Soutenue le 18 Octobre 1991 devant le jury:

Georges	GIRALT	Président
Nicholas	AYACHE	Rapporteurs
Jean Paul	LAUMOND	
Bertrand	ZAVIDOVIQUE	
Maurice	BRIOT	Examinateurs
Raja	CHATILA	
Roland	PRAJOUX	
Pierre	SOMMELET	

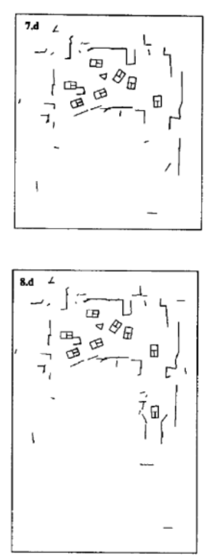

Figure 3.21 : A partir de maintenant les résultats vont différer suivant les méthodes. Nous finissons d'abord l'exploration et nous reviendrons sur ces différences.(Yaka)

Dynamic Environments (?)

1992

Dynamic Map Building for an Autonomous Mobile Robot

John J. Leonard
Hugh F. Durrant-Whyte
Department of Engineering Science
University of Oxford
Parks Road, Oxford OX1 3PJ
England

Ingemar J. Cox
NEC Research Institute
Princeton, New Jersey 08540

Abstract

This article presents an algorithm for autonomous map building and maintenance for a mobile robot. We believe that mobile robot navigation can be treated as a problem of tracking geometric features that occur naturally in the environment. We represent each feature in the map by a location estimate (the feature state vector) and two distinct measures of uncertainty: a covariance matrix to represent uncertainty in feature location, and a credibility measure to represent our belief in the validity of the feature. During each position update cycle, predicted measurements are generated for each geometric feature in the map and compared with actual sensor observations. Successful matches cause a feature's credibility to be increased. Unpredicted observations are used to initialize new geometric features, while unobserved predictions result in a geometric feature's credibility being decreased. We describe experimental results obtained with the algorithm that demonstrate successful map building using real sonar data.

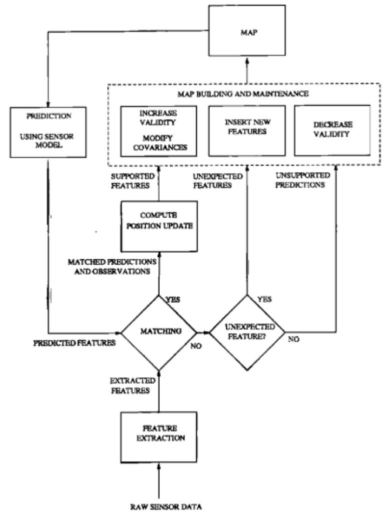

IJRR, 1992

1995

4 Simultaneous Localization and Map Building

The ideal localization system would allow the robot to start at a fixed location with no map of the surrounding environment, and from here to incrementally both build a map and use this map to locate its position relative to the start point. This is known as the simultaneous localization and map building (SLAM) problem. The SLAM problem is complex both as a mathematical exercise and in practical realization. It has received considerable attention over the years from a number of researchers [14, 5]. The essential problems are reasonably well known, however solutions have remained elusive, and a *significant* implementation has still not been attempted. We describe here the essential elements of the SLAM problem and some recent work in obtaining solutions to this problem.

Hugh Durrant-Whyte
Grasp Lab PhD 1986

Advisor: Lou Paul
Mentors: Ruzena Bacjsy and Max Mintz

Durrant-Whyte, et al., "Localization of Autonomous Guided Vehicles", ISRR, 1995

1996

New approach to simultaneous localization and dynamic map building. Csorba, Uhlmann, and Durrant-Whyte

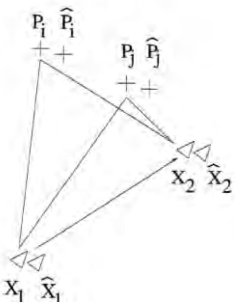

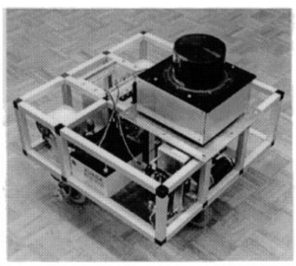

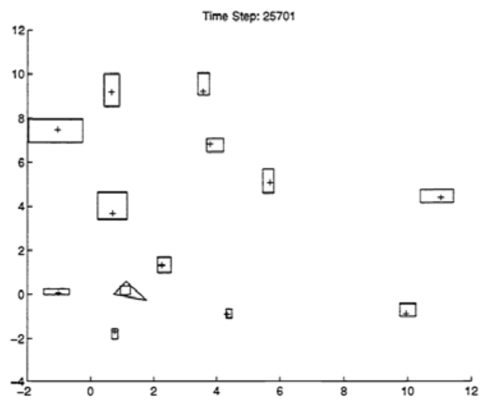

SPIE 1996

1997

Consistent Pose Estimation (Lu and Milios)

Our approach is to maintain all the local frames of data as well as the relative spatial relationships between local frames. These spatial relationships are modeled as random variables and are derived from matching pairwise scans or from odometry. Then we formulate a procedure based on the maximum likelihood criterion to optimally combine all the spatial relations. Consistency is achieved by using all the spatial relations as constraints to solve for the data frame poses simultaneously.

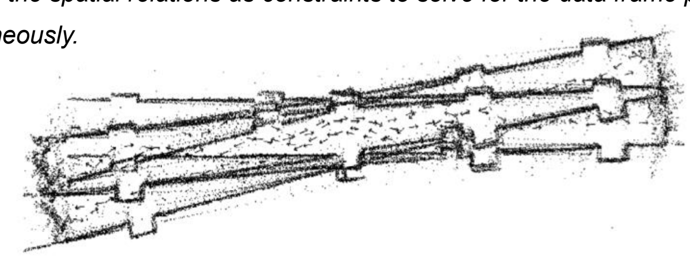

(a)

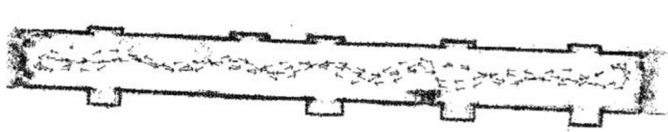

Visual SLAM – Andrew Davison

1998

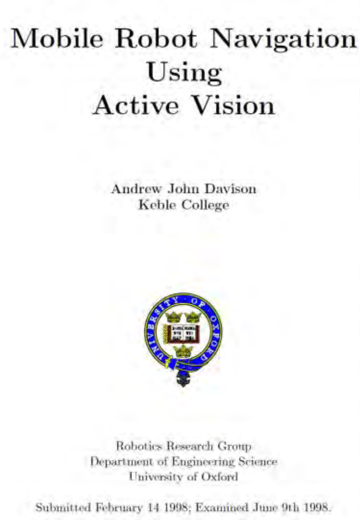
Mobile Robot Navigation Using Active Vision

Andrew John Davison
Keble College

Robotics Research Group
Department of Engineering Science
University of Oxford

Submitted February 14 1998; Examined June 9th 1998.

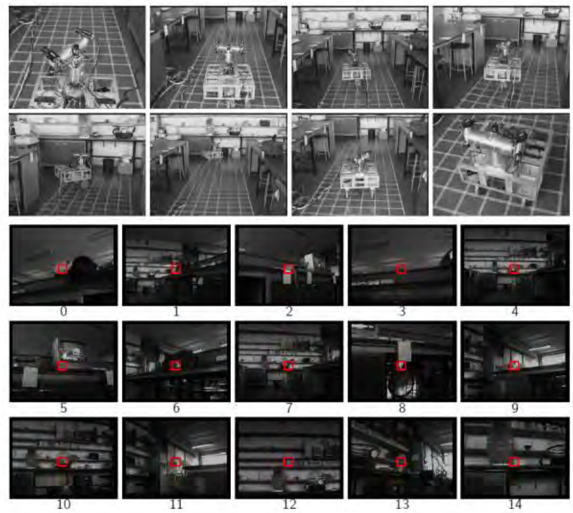

Visual SLAM – Andrew Davison

1998

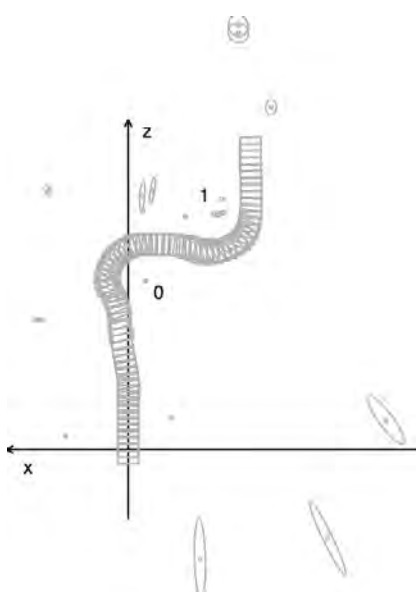

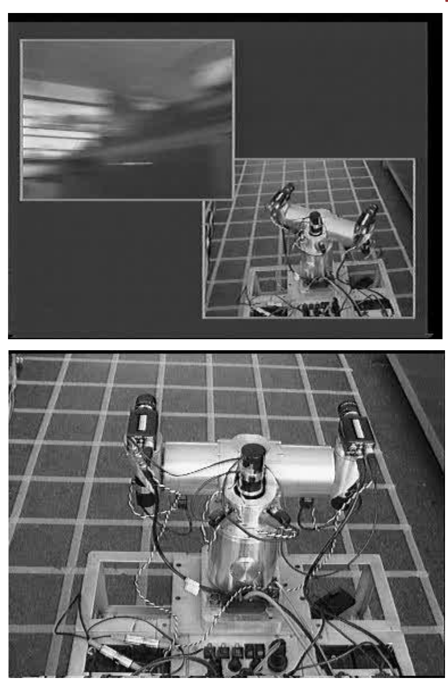

1999

Loop-Closing – Gutmann and Konolige

``A map is represented as an undirected graph: nodes are robot poses with associated scans and links are constraints between poses obtained from dead-reckoning, scan-matching, or correlation (Gutmann and Konolige, CIRA, 1999)"

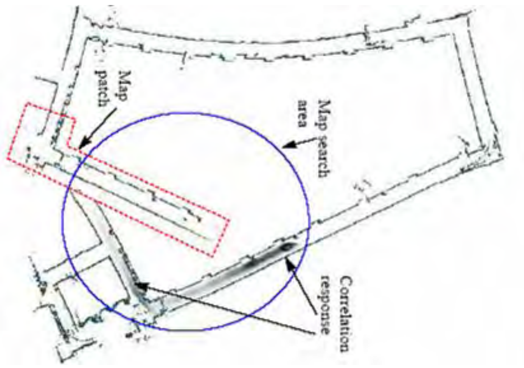

J.-S. Gutmann and K. Konolige. Incremental Mapping of Large Cyclic Environments, in: *International Symposium on Computational Intelligence in Robotics and Automation (CIRA'99)*, Monterey, November 1999.

2000

Probabilistic Algorithms and the Interactive Museum Tour-Guide Robot Minerva – Thrun et al.

Optimization of the Simultaneous Localization and Map Building Algorithm for Real Time Implementation (Guivant and Nebot)

2001

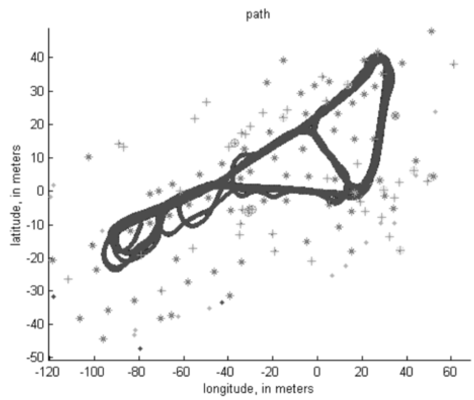

IEEE Trans R&A, 2001

Real-time SLAM using laser
Paul Newman

2002

Overhead view of scene (MIT Lobby 7)

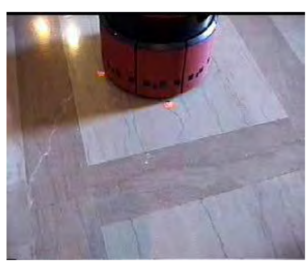

Starting position

View from the robot

Real-time software

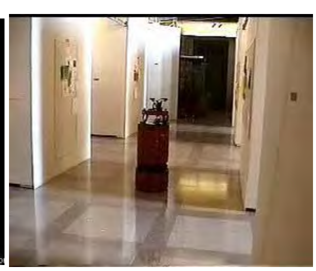

Return to home

Final adjustment

First SLAM Summer School

2002

Montemerlo and Thrun, FastSLAM

2002

AAAI 2002 Video courtesy of Cyrill Stachniss (from several years later)

An Atlas Framework for Scalable Mapping (Bosse)

2003

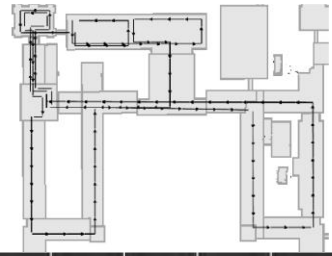

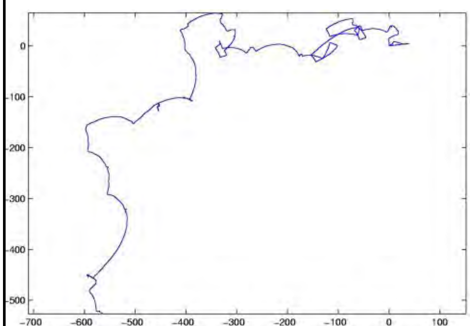

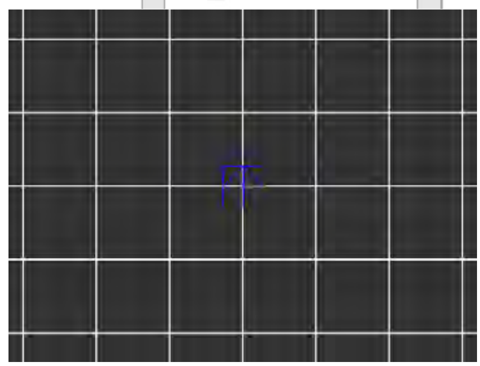

ICRA 2003

2005

Thrun, Burgard and Fox, MIT Press

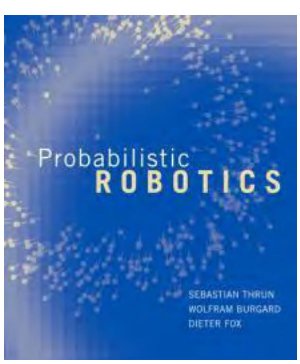

Wolfram Burgard, Dieter Fox and Sebastian Thrun (July, 2014)

Photo courtesy Wolfram Burgard

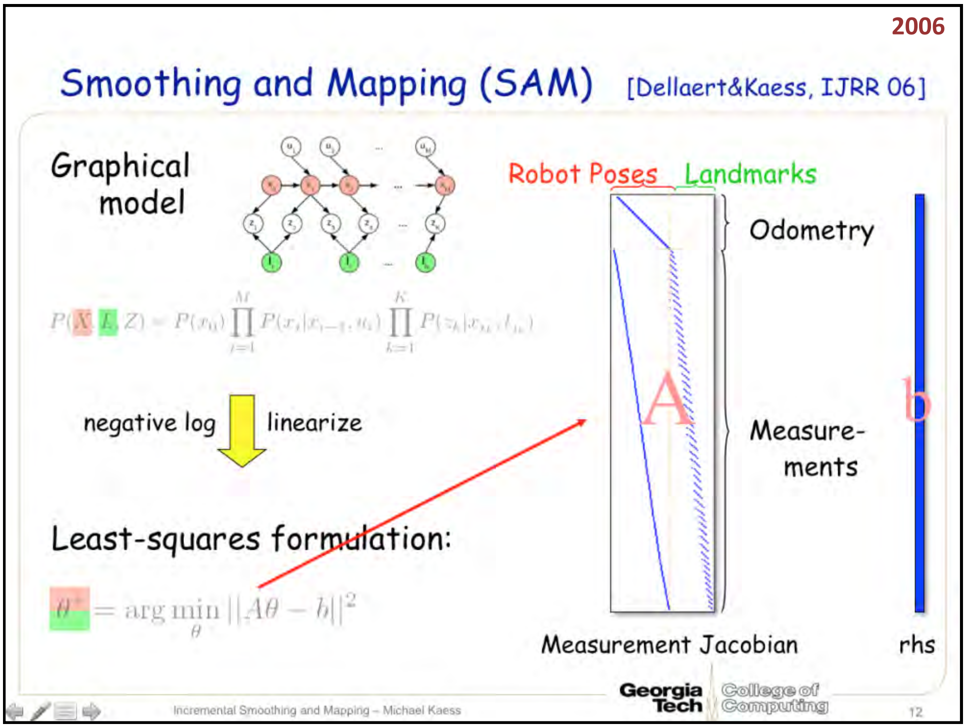

2007

G. Grisetti, C. Stachniss, S. Grzonka and W. Burgard
A Tree Parameterization for Efficiently Computing Maximum Likelihood Maps using Gradient Descent

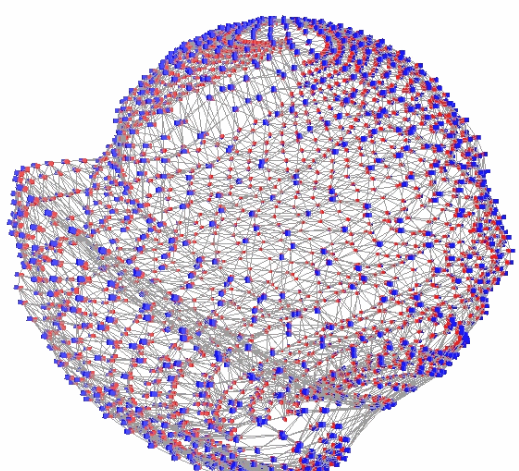

RSS 2007

2007

Parallel Tracking and Mapping (PTAM)
Klein and Murray

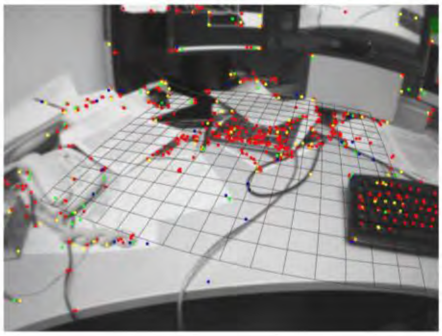

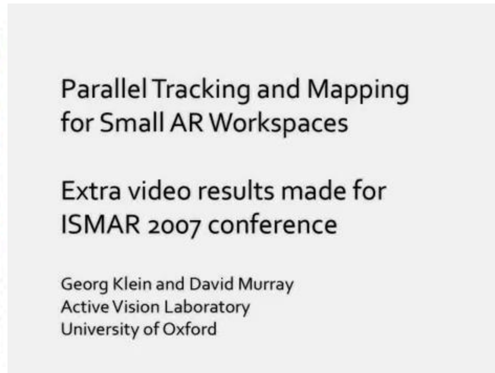

ISMAR, 2007 (Best Paper Award)

2009

FAB-MAP: Probabilistic Localization and Mapping in the Space of Appearance
Mark Cummins and Paul Newman

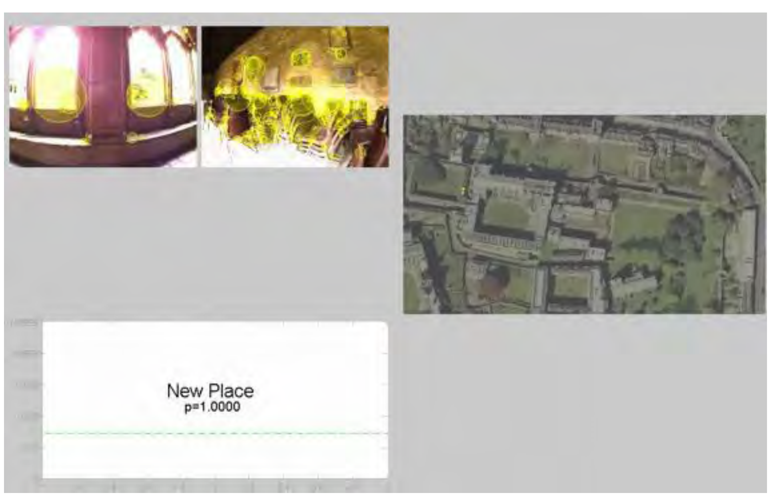

IJRR 2009

Sibley et al. – Relative Bundle Adjustment/VSLAM

2010

Planes, Trains and Automobiles – Autonomy for the Modern Robot

Gabe Sibley, Christopher Mei, Ian Reid and Paul Newman

Abstract—We are concerned with enabling truly large scale autonomous navigation in typical human environments. To this end we describe the acquisition and modeling of large urban spaces from data that reflects human sensory input. Over 200GB of image and inertial data are captured using head-mounted stereo cameras. This data is processed into a relative map covering 121 km of Southern England. We point out the numerous challenges we encounter, and highlight in particular the problem of undetected ego-motion, which occurs when the robot finds itself on-or-within a moving frame of reference. In contrast to global-frame representations, we find that the continuous relative representation naturally accommodates moving-reference-frames – without having to identify them first, and without inconsistency. Within a moving-reference-frame, and without drift-less global exteroceptive sensing, motion with respect to the global-frame is effectively unobservable. This underlying truth drives us towards relative topometric solutions like relative bundle adjustment (RBA), which has no problem representing distance and metric Euclidean structure, yet does not suffer inconsistency introduced by the attempt to solve in the global-frame.

I. Introduction

Autonomous navigation in human working environments is an important problem, and this paper is motivated by our attempt to make sense of the 121 km path between Oxford and London depicted in Figs. 1 and 2. The map begins in an office in Oxford, and proceeds with various forms of transport

Figure 1: 121 km path between Oxford in the upper left and London in the bottom right. We compute visual estimates for 89.4% of this distance. Using appearance-based place recognition and inertial dead reckoning, 100% is covered topologically, which is sufficient for path planning. The graph begins in an office in Oxford, and proceeds with various forms of transport including: foot, bicycle, train, subway, escalator, rickshaw, punting-boat and ferris wheel. Note that we cannot detect our true position in the global inertial frame – when we are traveling on the train or subway for instance, motion with respect to the global inertial frame becomes effectively unobservable in the presence of noise.

ICRA 2010

KinectFusion – Izadi, Necombe et al.

2011

SIGGRAPH Talks 2011

KinectFusion:

Real-Time Dynamic 3D Surface Reconstruction and Interaction

Shahram Izadi 1, Richard Newcombe 2, David Kim 1,3, Otmar Hilliges 1, David Molyneaux 1,4, Pushmeet Kohli 1, Jamie Shotton 1, Steve Hodges 1, Dustin Freeman 5, Andrew Davison 2, Andrew Fitzgibbon 1

1 Microsoft Research Cambridge 2 Imperial College London
3 Newcastle University 4 Lancaster University
5 University of Toronto

2012

Kintinuous (Whelan, McDonald et al.)

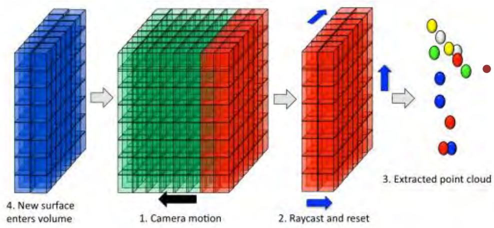

- Extension of KinectFusion (Newcombe, et al. ISMAR '11)
- Treat volumetric model as a cyclical buffer.
  - As region leaves the range of the buffer, extract surface data.
  - As region enters the range of the buffer, initialise and track the new data.

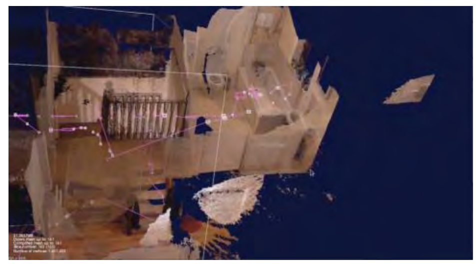

Whelan et al. RSS 2012 RGB-D Workshop

2013

SLAM++
Salas-Moreno et al.

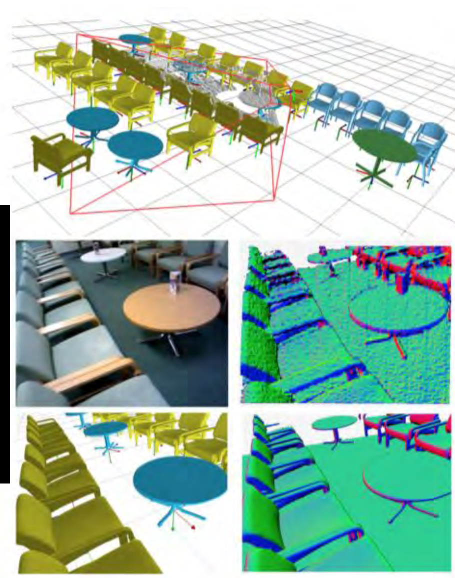

CVPR 2013

2013

Kintinous Processing Pipeline ("Cloud Slices" connected to pose graph SLAM optimization)

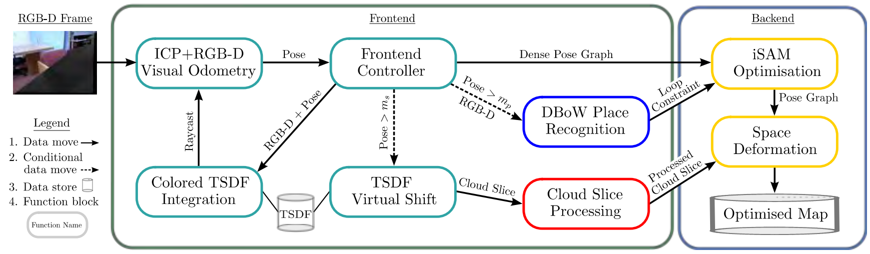

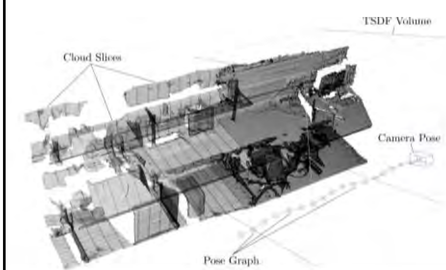

"Deformation-based Loop Closure for Large Scale Dense RGB-D SLAM" by T. Whelan, M. Kaess, J. Leonard and J. McDonald, IROS 2013

2014

Google Tango – Journey

www.youtube.com/watch?v=44vppay5UDc

Progress in SLAM online pose graph optimization (courtesy of Luca Carlone)

2015

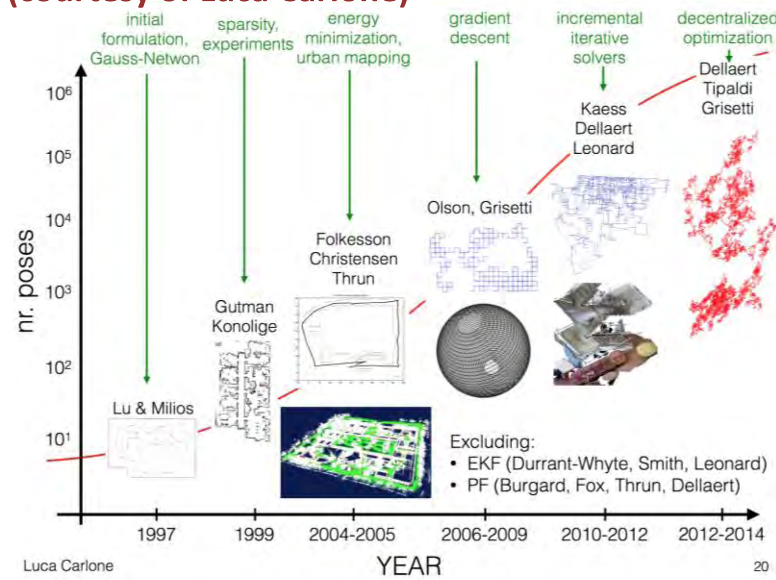

Check out the slides from his recent MIT talk at http://www.lucacarlone.com/

Kintinuous with Stereo – Walking over Stairs

2015

Kintinuous with Dense Stereo on MIT DRC Atlas
Maurice Fallon. Scott Kuidersma, Tom Whelan, and Russ Tedrake.

Conclusion and Future Research Challenges

Goals:

- My dream is to achieve *persistent autonomy* and *lifelong map learning* in highly dynamic environments
- Can we robustly integrate mapping and localization with real-time planning and control?

Open Questions:

- Robustness – we would love to have guarantees of performance, but we do not have them for most approaches
- Representation – how can we integrate many different types?
- We need dynamic scene understanding and robust vision (recent work in computer vision is very exciting, but current precision-recall curves indicate we have a long way to go)

Is SLAM "Solved?"

Its an Exciting Time to Work in Mobile Sensing!

Postdocs and PhD students that can build real-time 3D perception, navigation and motion planning systems are in high demand:

- Virtual Reality
- Mobile Devices
- Self-Driving Vehicles
- Drones

Big tech companies such as Google, Apple, Facebook and Uber

Small startups such as skydio

Traditional companies in transition, such as Ford, Delphi, Continental, Bosch…

Vision for Mobile Robotics: A Research Agenda

- We need an *object-based* understanding of the environment that facilitates life-long learning
- Let's build rich representations that leverage knowledge of location to better understand about objects, and concurrently uses information about objects to better understand location
  - Sudeep Pillai: Monocular SLAM Supported Object Recognition (presented at RSS2015 on Tuesday)
  - Ross Finman: Automatic Discovery of Objects in lifelong Dense RGB-D maps

    Key Idea: can we learn about objects through observing changes in the world?

Pillai and Leonard, RSS 2015

MONOCULAR SLAM SUPPORTED
OBJECT RECOGNITION

Sudeep Pillai & John J. Leonard
Computer Science and Artificial Intelligence Lab
Massachusetts Institute of Technology

Why is SLAM Difficult?

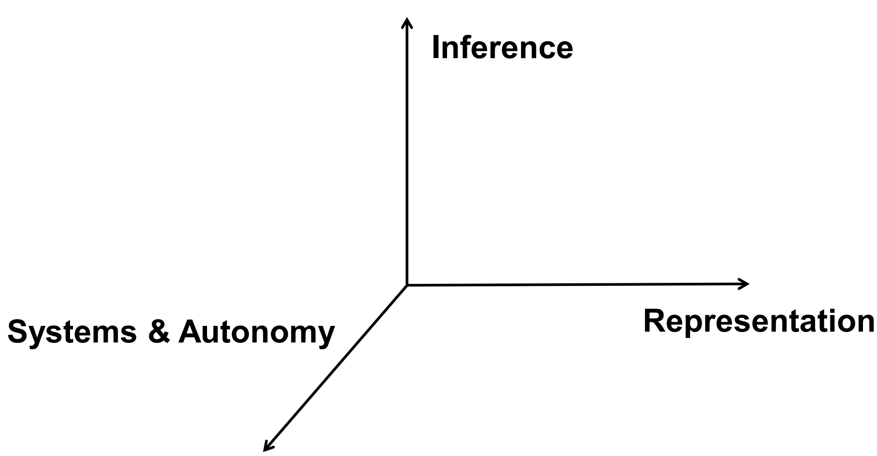

Why is SLAM Difficult?

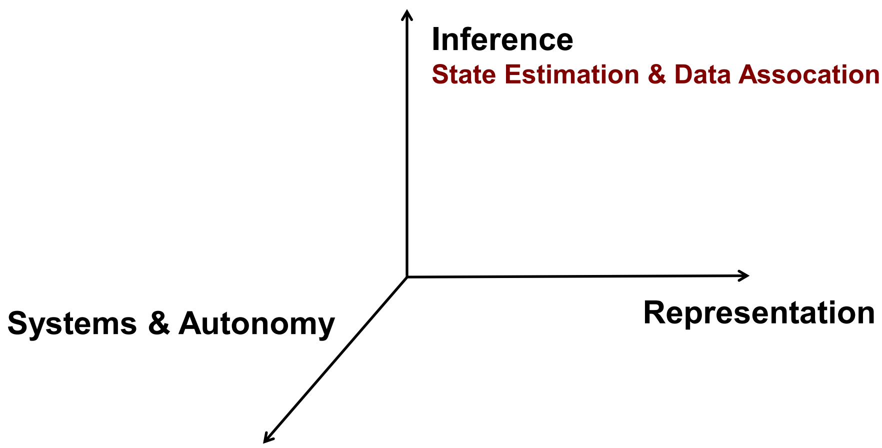

Why is SLAM Difficult?

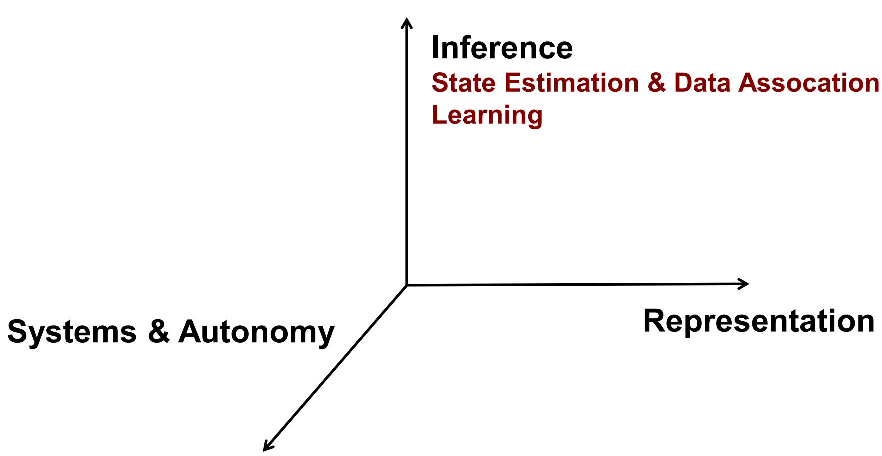

Why is SLAM Difficult?

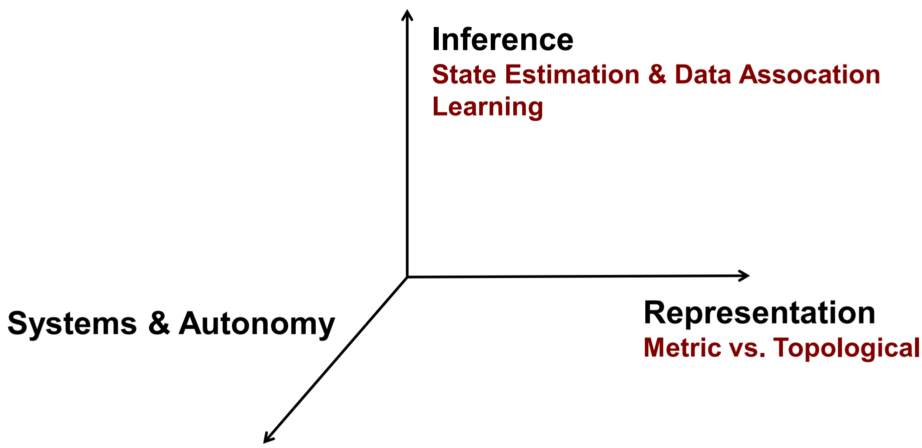

Why is SLAM Difficult?

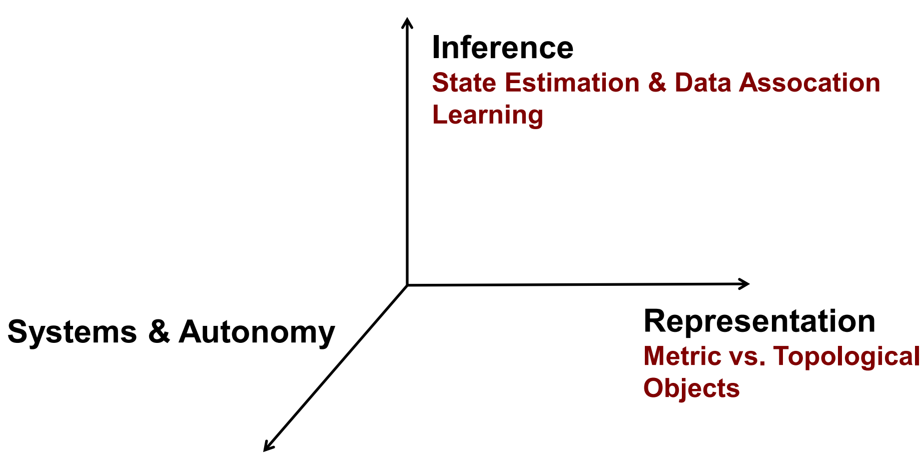

Why is SLAM Difficult?

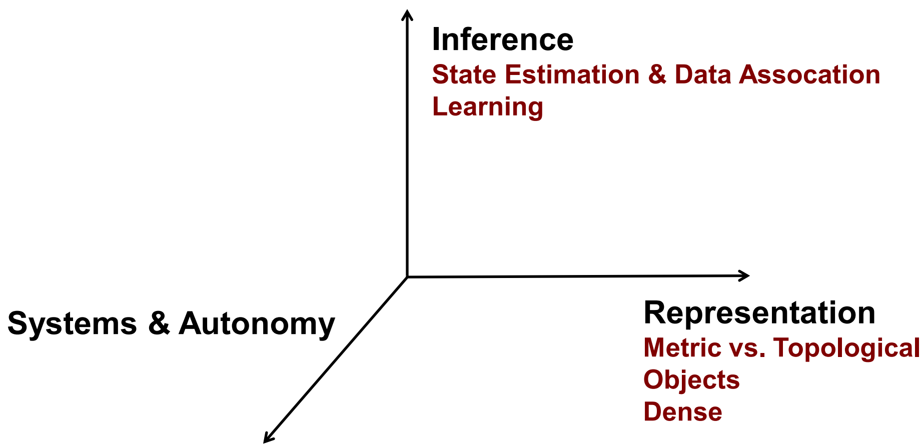

Why is SLAM Difficult?