



上海科技大学  
ShanghaiTech University

## CS283: Robotics Fall 2020: SLAM I

---

Qingwen Xu    Sören Schwertfeger


ShanghaiTech University

# DEFINITION OF SLAM

---

# What is SLAM?

- **Localization**: inferring location given a map
- **Mapping**: inferring a map given locations
- **SLAM**: learning a map and locating the robot simultaneously
- SLAM has long been regarded as a chicken-and-egg problem:
  - a map is needed for localization and
  - a pose estimate is needed for mapping





Material derived from Wolfram Burgard:

<http://ais.informatik.uni-freiburg.de/teaching/ss20/robotics/slides/13-slam.pdf>

# SLAM Applications


- At home: vacuum cleaner, lawn mower
- Air: surveillance with unmanned air vehicles
- Underwater: reef monitoring
- Underground: exploration of mines
- Space: terrain mapping for localization
- .....






# Why is SLAM a Hard Problem?

- Robot path and map are both unknown
- Errors in map and pose estimates correlated



# Why is SLAM a Hard Problem?

- The mapping between observations and landmarks is unknown
- Picking wrong data associations can have catastrophic consequences (divergence)



# Overview of SLAM Methods

- Camera

- Feature-Based Methods
  - MonoSLAM
  - PTAM
  - ORB-SLAM
- Direct Methods
  - DTAM
  - LSD-SLAM
  - DSO
- Semi-Direct Methods
  - SVO
- Others
  - PoseNet
  - CNN-SLAM
  - ...

- Laser

- Pose Graph
  - Cartographer
  - Karto-SLAM
  - Hector-SLAM
  - BLAM
  - LIO
- Particle Filter
  - FastSLAM
  - Gmapping
- Extended Kalman Filter
  - EKF-SLAM
  - LINS
- Others
  - LOAM
  - IMLS-SLAM
  - ...

# SLAM Front-end & Back-end

- Front-end
  - calculate relative poses between several frames/ to map
    - scan matching
    - image registration
    - ...
  - estimate absolute poses
  - construct the local map
- Back-end
  - optimize the absolute poses and mapping
  - only if a loop was closed




# FRONT END - LASER

---


# Registration methods for Range Data

- ICP
- NDT
- Robust point matching (soft point correspondences)
- Coherent point drift
- Kernel correlation
- Approximations of the squared distance functions to curves and surfaces
- Feature extracting methods
  - Corners in point clouds
  - Lines
  - Planes
- Spectral methods



# ICP: Iterative Closest Points Algorithm

- Align two partially-overlapping point sets (2D or 3D)
- Given initial guess for relative transform
- Warning: Using 3D ICP for 2D data may mirror the data (e.g. 180 degree roll)!
  - Use 2D ICP!
- ROS: Point Cloud Library (PCL)




# Data Types

- Point sets
- Line segment sets (polylines)
- Implicit curves :  $f(x,y,z) = 0$
- Parametric curves :  $(x(u),y(u),z(u))$
- Triangle sets (meshes)
- Implicit surfaces :  $s(x,y,z) = 0$
- Parametric surfaces  $(x(u,v),y(u,v),z(u,v))$




# Motivation

- Scan Matching - Registration
- Shape inspection
- Motion estimation
- Appearance analysis
- Texture Mapping
- Tracking








# Aligning 3D Data

- Continuous lines or a set of points...




# Corresponding Point Set Alignment

- Let  $M$  be a model point set. (or map or previous scan)
- Let  $S$  be a scene point set. (current scan)

We assume :

1.  $N_M = N_S$ .
2. Each point  $S_i$  correspond to  $M_i$  .





# Corresponding Point Set Alignment

The Mean Squared Error (MSE) objective function :

$$f(R, T) = \frac{1}{N_S} \sum_{i=1}^{N_S} \|m_i - Rot(s_i) - Trans\|^2$$


$$f(q) = \frac{1}{N_S} \sum_{i=1}^{N_S} \|m_i - R(q_R)s_i - q_T\|^2$$

The alignment is :

$$(rot, trans, d_{mse}) = \Phi(M, S)$$

# Aligning 3D Data

- If correct correspondences are known, can find correct relative rotation/ translation as closed form solution:
  - **Horn's quaternion method**
  - SVD Arun et al.
  - Orthonormal matrices Horn et al.
  - Dual quaternions Walker et al.



See:

A. Lorusso, D. Eggert, and R. Fisher.


A Comparison of Four Algorithms for Estimating 3-D Rigid Transformations.

In *Proceedings of the 4th British Machine Vision Conference (BMVC '95)*, pages 237 - 246, Birmingham, England, September 1995.

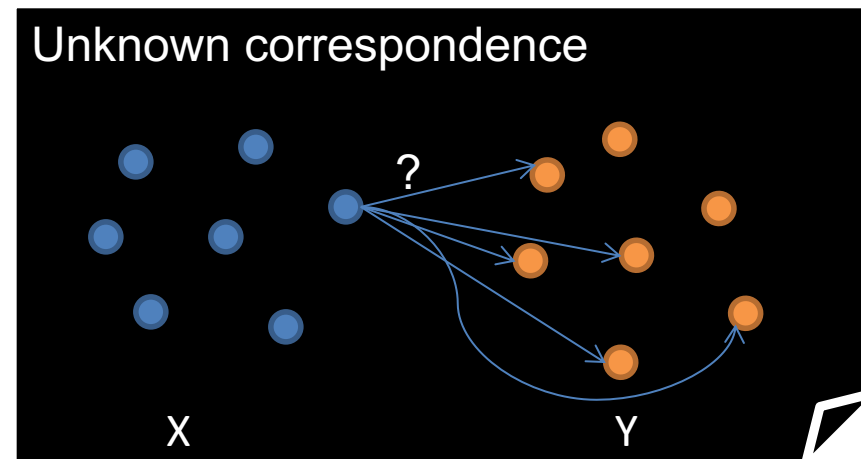
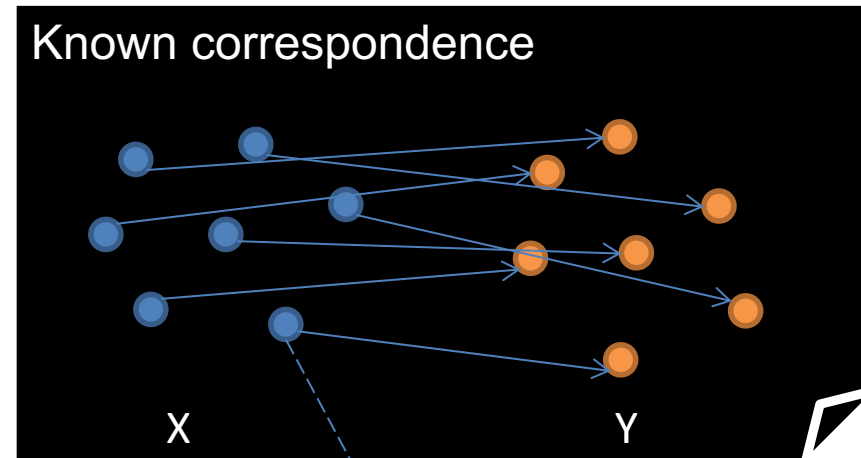

# Horn's method

Material by Toru Tamaki, Miho Abe,  
Bisser Raytchev, Kazufumi Kaneda

- Input
  - Two point sets:  $X$  and  $Y$
- Output
  - Rotation matrix  $R$
  - Translation vector  $t$

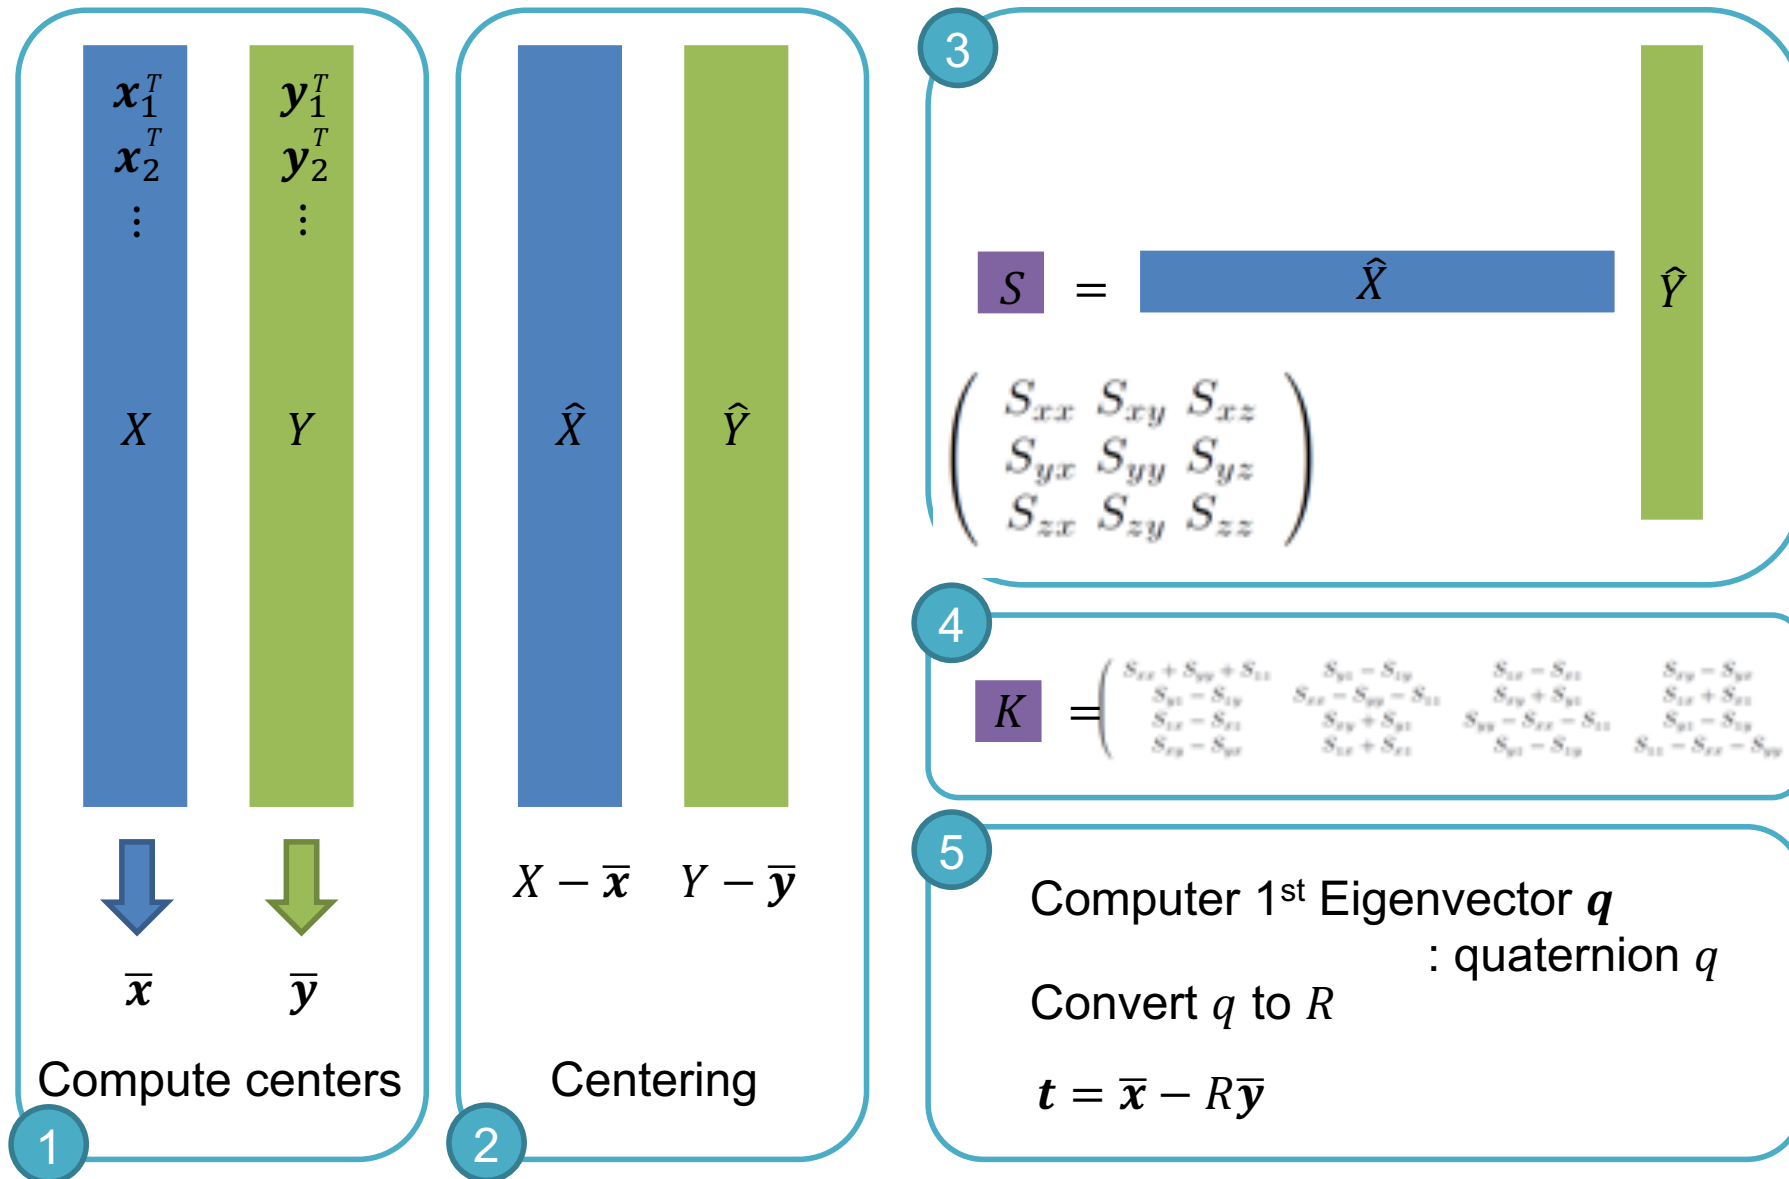


# Horn's method: correspondence is known.






# Horn's method: correspondence is known.




# Aligning 3D Data

- How to find correspondences: User input? Feature detection? Signatures?
- Alternative: assume **closest** points correspond




# Aligning 3D Data

- How to find correspondences: User input? Feature detection? Signatures?
- Alternative: assume **closest** points correspond



# Aligning 3D Data

- Converges if starting position “close enough“



# Closest Point

- Given 2 points  $r_1$  and  $r_2$ , the Euclidean distance is:

$$d(r_1, r_2) = \|r_1 - r_2\| = \sqrt{(x_1 - x_2)^2 + (y_1 - y_2)^2 + (z_1 - z_2)^2}$$

- Given a point  $r_1$  and set of points  $A$ , the Euclidean distance is:

$$d(r_1, A) = \min_{i \in 1..n} d(r_1, a_i)$$

# Finding Matches

- The scene shape  $S$  is aligned to be in the best alignment with the model shape  $M$ .
- The distance of each point  $s$  of the scene from the model is :

$$d(s, M) = \min_{m \in M} d \|m - s\|$$

# Finding Matches

$$d(s, M) = \min_{m \in M} d\|m - s\| = d(s, y)$$

$$y \in M$$

$$Y = C(S, M)$$

$$Y \subseteq M$$

C – the closest point operator

Y – the set of closest points to S

# Finding Matches

- Finding each match is performed in  $O(MN)$  worst case.
- Given  $Y$  we can calculate alignment


$$(rot, trans, d) = \Phi(S, Y)$$

- $S$  is updated to be :

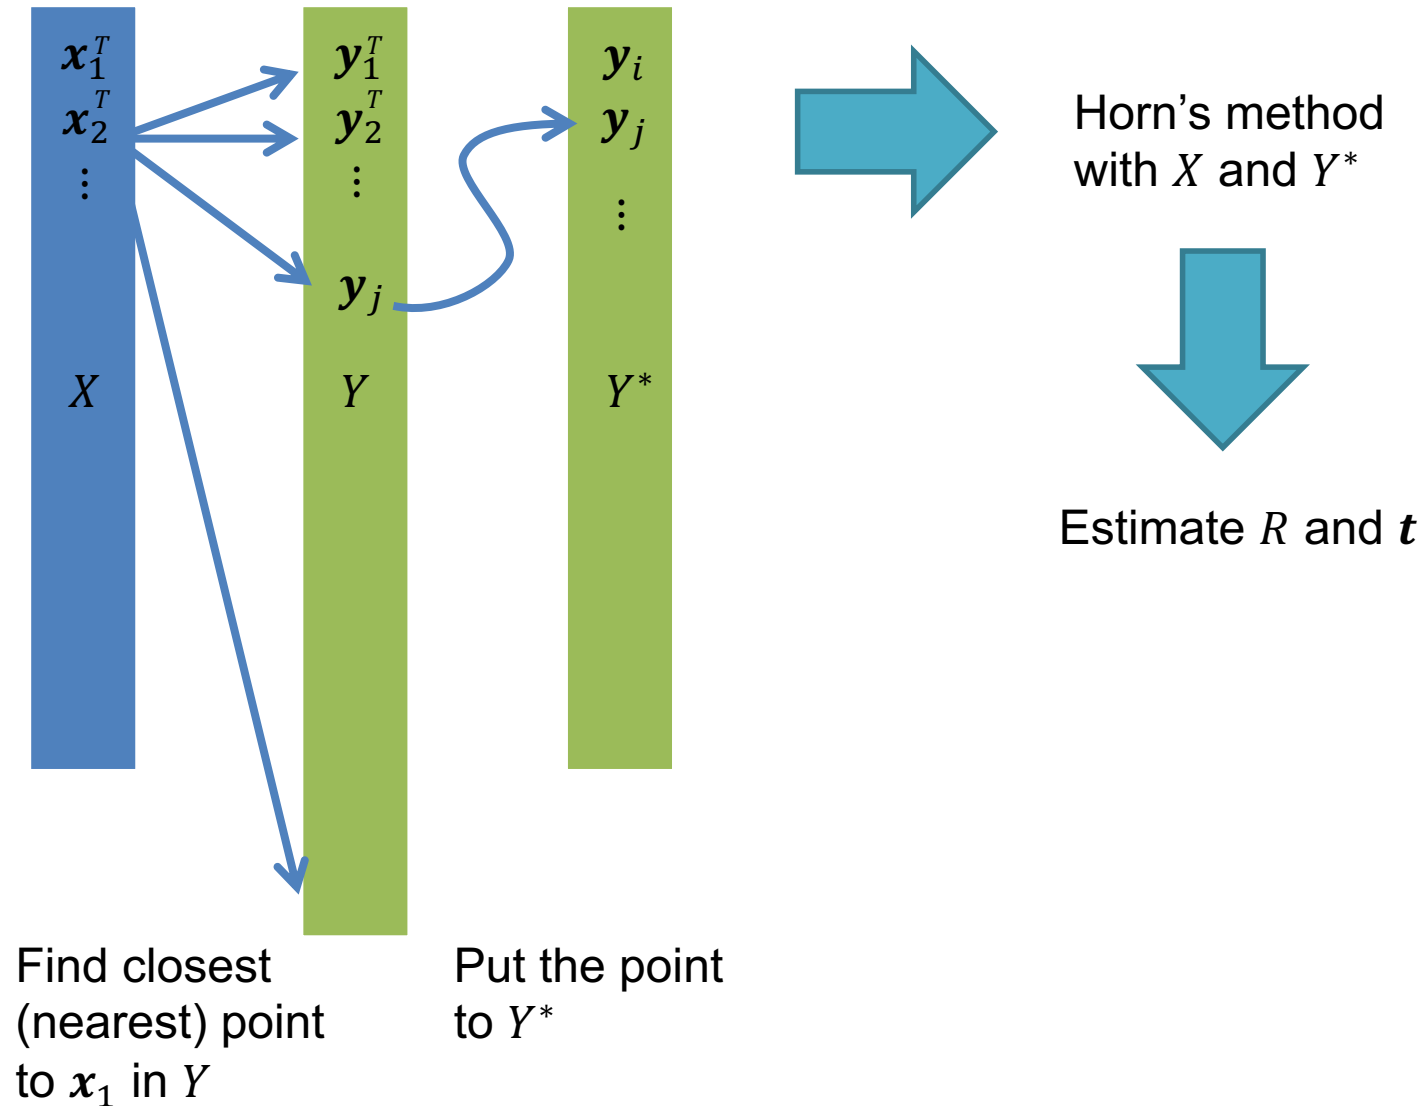
$$S_{new} = rot(S) + trans$$



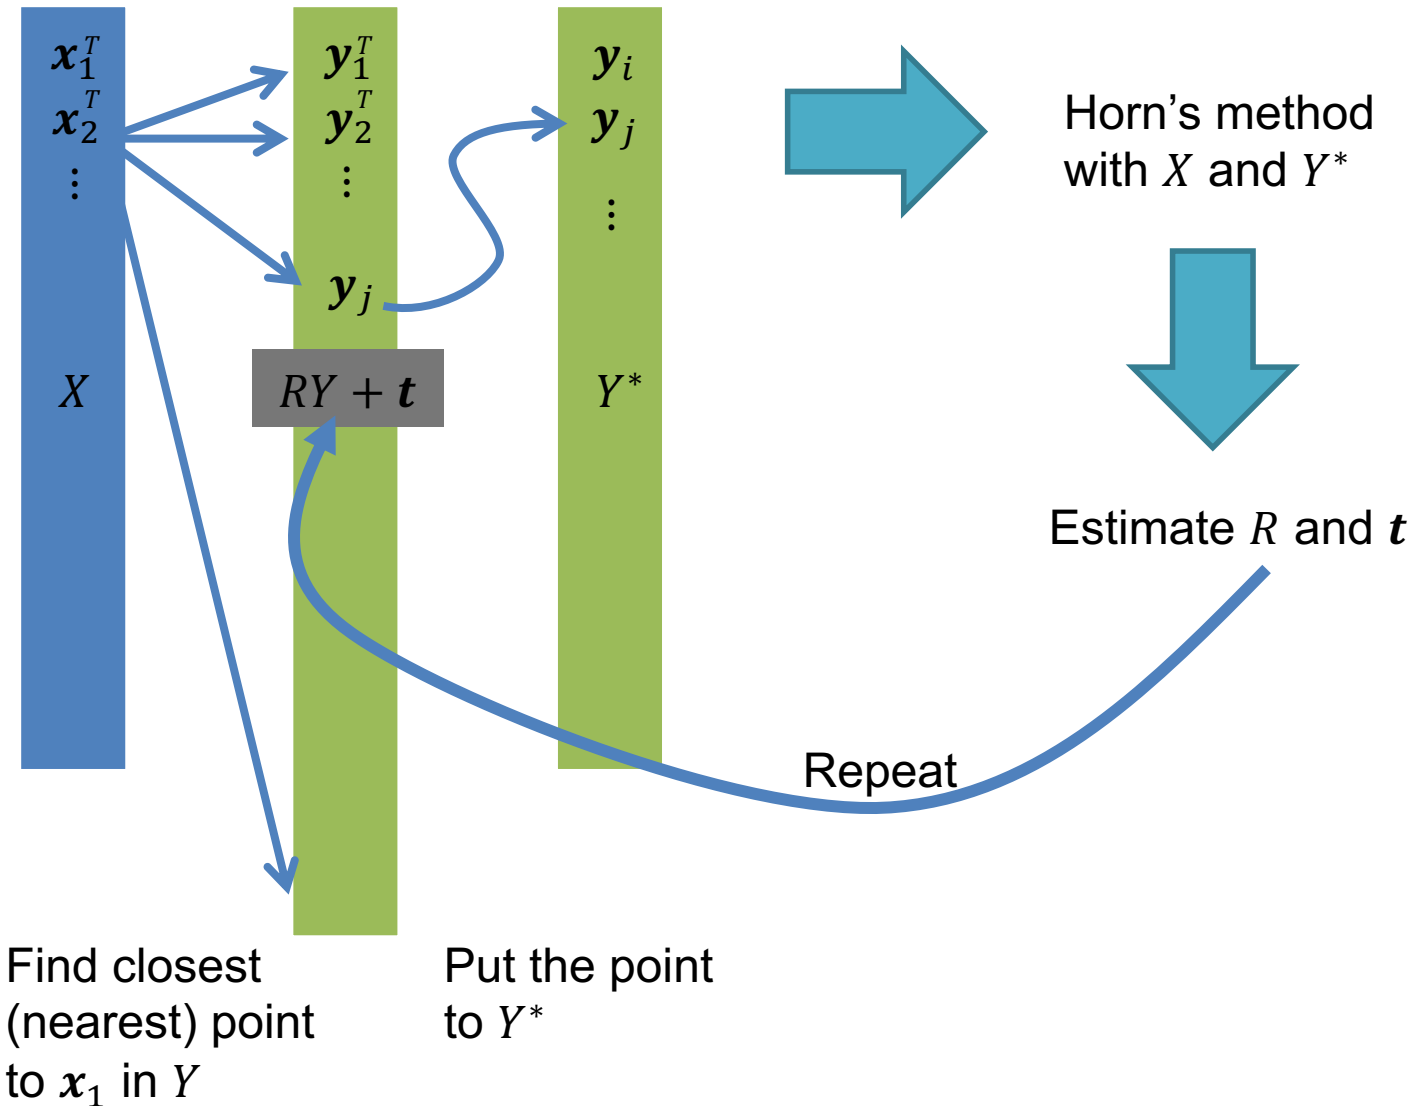
# ICP: correspondence is unknown.



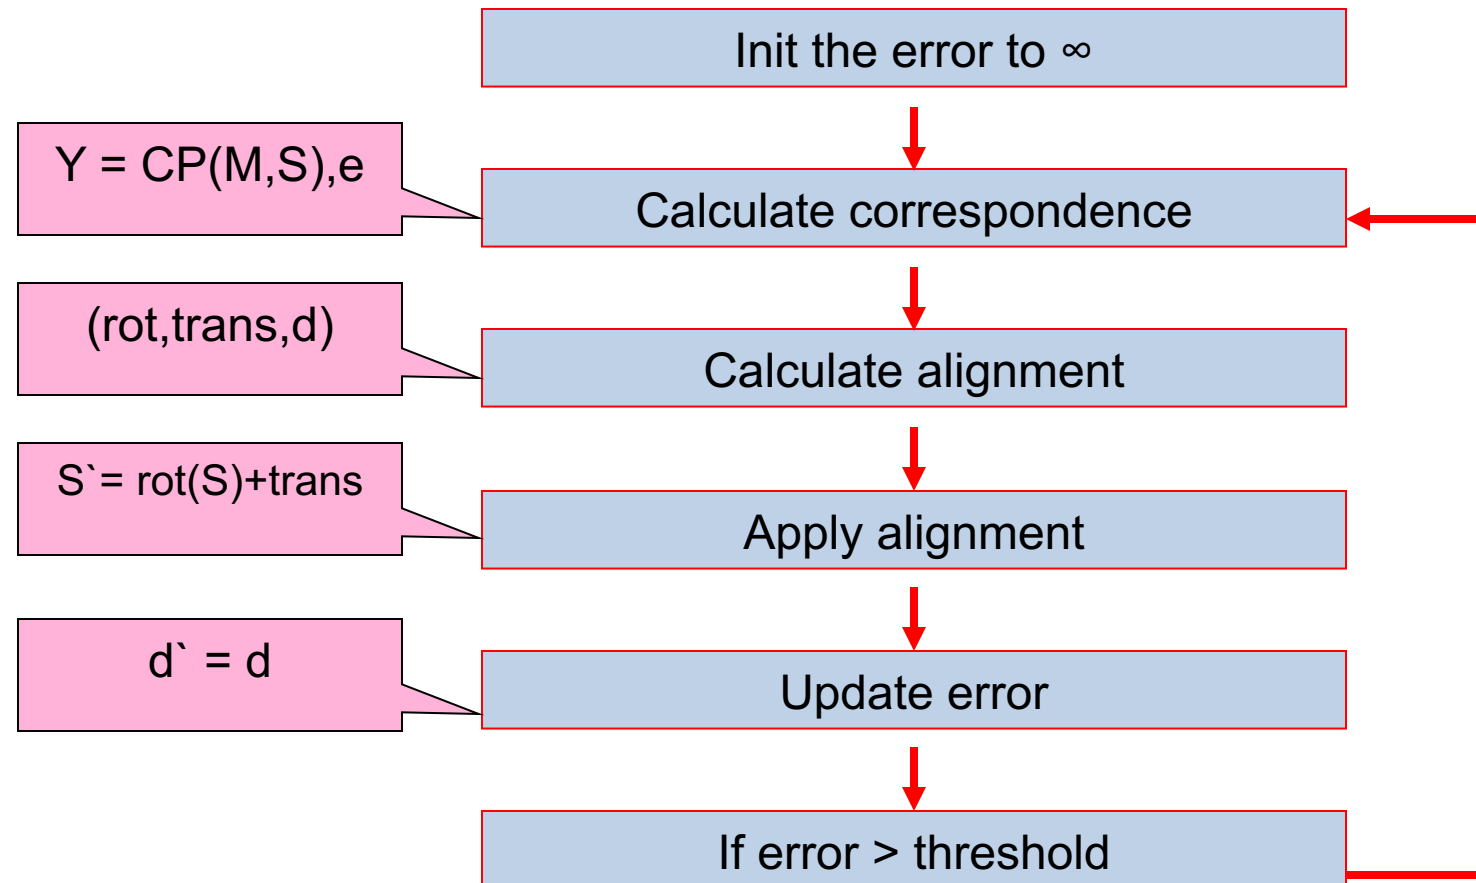
# ICP: correspondence is unknown.



# ICP: correspondence is unknown.



# The Algorithm



# The Algorithm

```
function ICP(Scene,Model)
begin
  E` ← + ∞;
  (Rot,Trans) ← In Initialize-Alignment(Scene,Model);
  repeat
    E ← E`;
    Aligned-Scene ← Apply-Alignment(Scene,Rot,Trans);
    Pairs ← Return-Closest-Pairs(Aligned-Scene,Model);
    (Rot,Trans,E`) ← Update-Alignment(Scene,Model,Pairs,Rot,Trans);
  Until |E` - E| < Threshold
  return (Rot,Trans);
end
```

# Convergence Theorem

- The ICP algorithm always converges monotonically to a local minimum with respect to the MSE distance objective function.

# Time analysis

Each iteration includes 3 main steps

A. Finding the closest points :

$O(N_M)$  per each point

$O(N_M * N_S)$  total.

B. Calculating the alignment:  $O(N_S)$

C. Updating the scene:  $O(N_S)$


# Optimizing the Algorithm

- K-D Tree :

Construction time:  $O(kn \log n)$

Space:  $O(n)$

Region Query :  $O(n^{1-1/k} + k)$





# Time analysis

Each iteration includes 3 main steps

A. Finding the closest points :

$O(N_M)$  per each point

$O(N_M \log N_S)$  total.

B. Calculating the alignment:  $O(N_S)$

C. Updating the scene:  $O(N_S)$


# ICP Variants

- Variants on the following stages of ICP have been proposed:
  1. **Selecting** sample points (from one or both point clouds)
  2. **Matching** to points to a plane or mesh
  3. **Weighting** the correspondences
  4. **Rejecting** certain (outlier) point pairs
  5. Assigning an **error metric** to the current transform
  6. **Minimizing** the error metric w.r.t. transformation


# Performance of Variants

- Can analyze various aspects of performance:
  - Speed
  - Stability (local minima)
  - Tolerance of noise and/or outliers
  - Maximum initial misalignment


# ICP Variants

- 
1. Selecting sample points (from one or both meshes).
  2. Matching to points to a plane or mesh
  3. Weighting the correspondences.
  4. Rejecting certain (outlier) point pairs.
  5. Assigning an error metric to the current transform.
  6. Minimizing the error metric w.r.t. transformation.


# ICP Variants

1. Selecting sample points (from one or both meshes).
-  2. **Matching** to points to a plane or mesh.
3. Weighting the correspondences.
4. Rejecting certain (outlier) point pairs.
5. Assigning an error metric to the current transform.
6. Minimizing the error metric w.r.t. transformation.

# ICP Variants

1. Selecting sample points (from one or both meshes).
2. Matching to points to a plane or mesh.
-  3. **Weighting the correspondences.**
4. Rejecting certain (outlier) point pairs.
5. Assigning an error metric to the current transform.
6. Minimizing the error metric w.r.t. transformation.

# ICP Variants

1. Selecting sample points (from one or both meshes).
2. Matching to points to a plane or mesh.
3. Weighting the correspondences.
-  4. Rejecting certain (outlier) point pairs.
5. Assigning an error metric to the current transform.
6. Minimizing the error metric w.r.t. transformation.

# Rejecting Pairs

- Corresponding points with point to point distance higher than a given threshold.
- Rejection of worst  $n\%$  pairs based on some metric.
- Pairs containing points on end vertices.
- Rejection of pairs whose point to point distance is higher than  $n^*\sigma$ .
- Rejection of pairs that are not consistent with their neighboring pairs [Dorai 98] :


$(p_1, q_1)$  ,  $(p_2, q_2)$  are inconsistent iff

$$\left| \text{Dist}(p_1, p_2) - \text{Dist}(q_1, q_2) \right| > \text{threshold}$$




# Rejecting Pairs

Distance thresholding




# Rejecting Pairs

Points on end vertices





# Rejecting Pairs


## Inconsistent Pairs



# ICP Variants

1. Selecting sample points (from one or both meshes).
2. Matching to points to a plane or mesh.
3. Weighting the correspondences.
4. Rejecting certain (outlier) point pairs.
-  5. Assigning an error metric to the current transform.
-  6. Minimizing the error metric w.r.t. transformation.

# BLAM: ICP in action





# FRONT END - CAMERA

---

# Methods

- **Feature-based Methods**

- SIFT
- ORB (ORB-SLAM)
- BRISK
- AKAZE



- **Direct Methods**

- Optical Flow
- Inverse Depth (LSD-SLAM)
- Fourier-Mellin Transform

- **Semi-Direct Methods**

- SVO

more details in the next lectures



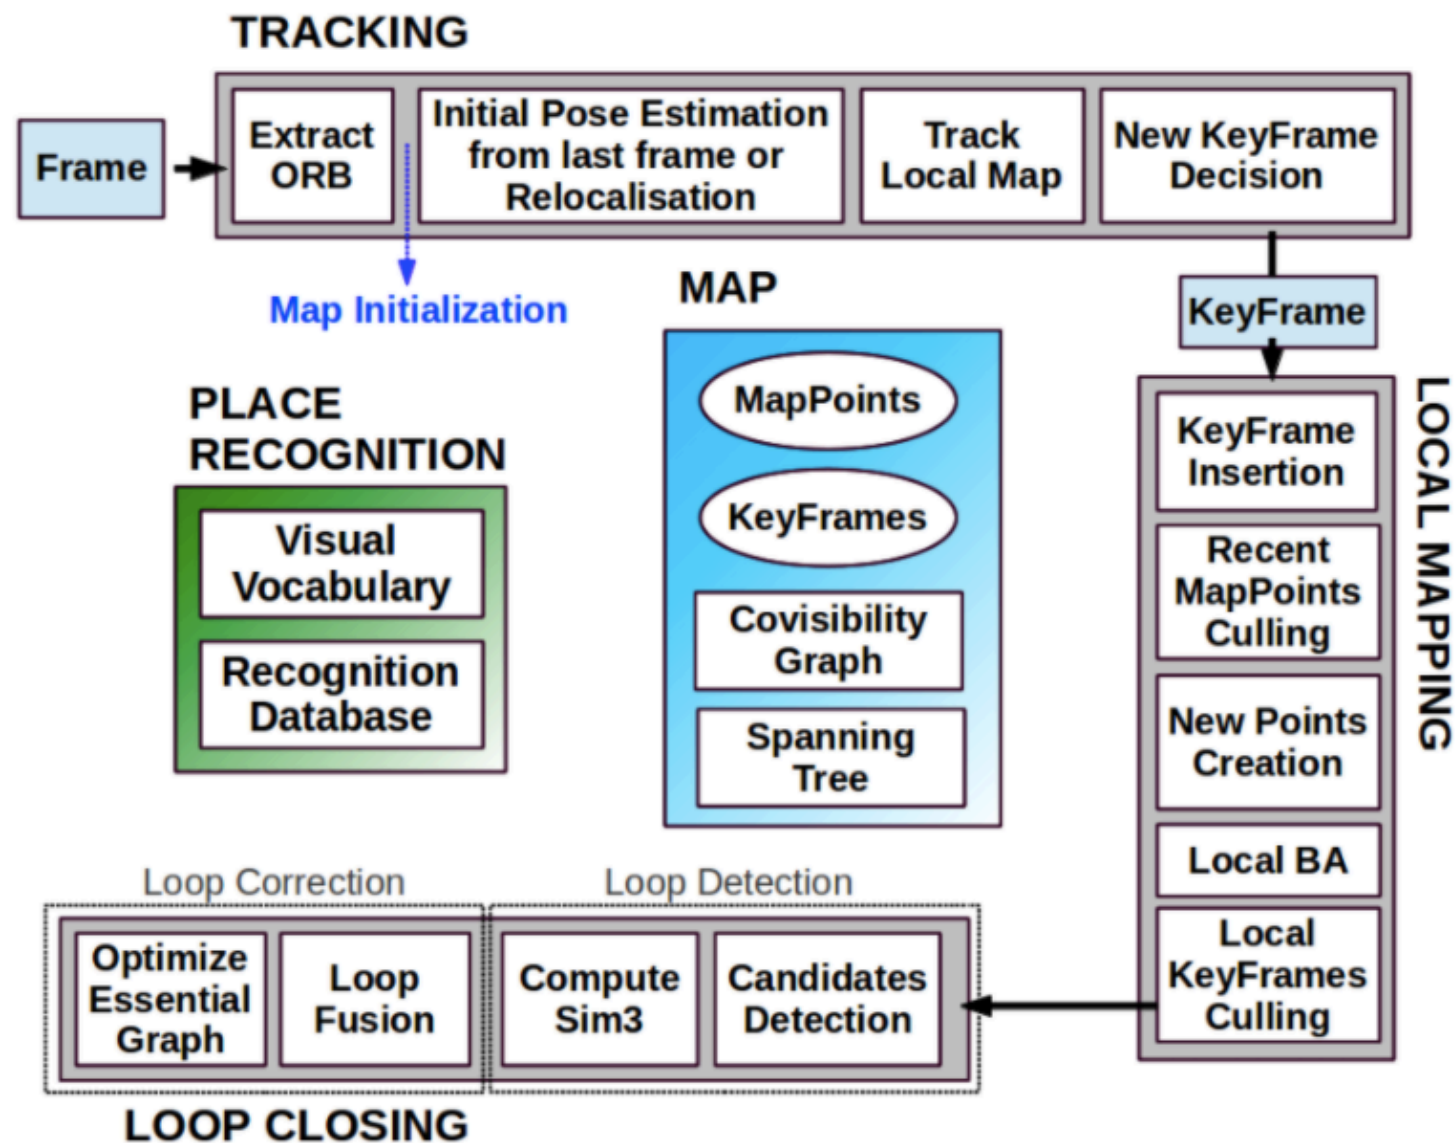


# Feature-based Methods

- Feature Extraction
  - Feature Detectors & Feature Descriptor; more in vision lectures
    - ORB, SIFT, AKAZE, BRISK, etc ...
- Feature Matching
  - BFM, KNN, etc ...
- Relative Pose Calculation
  - 5-pt, 7-pt, 8-pt, PnP, etc ...


# Feature-based Method: ORB-SLAM

*Mur-Artal R, Montiel J M M, Tardos J D. ORB-SLAM: a versatile and accurate monocular SLAM system[J]. IEEE transactions on robotics, 2015, 31(5): 1147-1163.*



# Direct Method: LSD-SLAM




- Construct Photometric Error




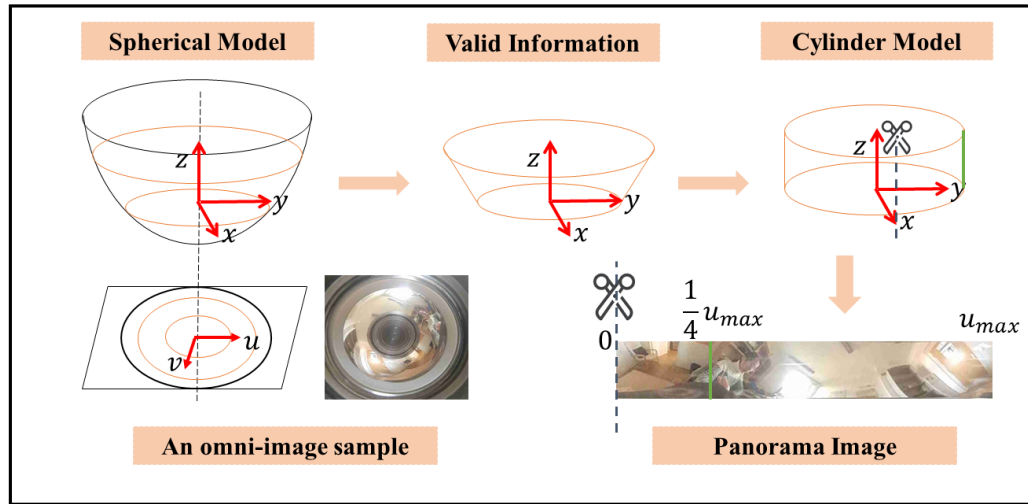
- Construct Depth Error
- Minimize Objective Error Function

# Direct Method: Fourier Mellin Transform

- Spectral based registration: detection of scaling, rotation and translation in 2 subsequent frames
- Processing spectrum magnitude decouples translation from affine transformations
  - Detection of signal shift between 2 signals by phase information
  - Resampling to polar coordinates → Rotation turns into signal shift !
  - Resampling the radial axis from linear to logarithmic presentation → Scaling turns into signal shift !
  - Calculate a Phase Only Match Filter (POMF) on the resampled magnitude spectra




# Pose Estimation for Omni-directional Cameras using Sinusoid Fitting



(1a)

(1b)

(1c)



$$y = B + A \sin(\omega x + \phi)$$

$$\Delta v(u_p) = \lambda_p t_z + \gamma \left\| P_{xy}(R) \cdot \sin(\gamma u_p + \frac{P_{xy}(R)}{\|P_{xy}(R)\|}) \right\|$$

$$\Delta u(u_p) = \gamma P_z(R) + \lambda_p \|P_{xy}(t)\| \cdot \sin(\gamma u_p + \hat{t}_{xy})$$

# BACK END


---

# Overview of Back-end

- Loop Detection
  - Find candidates of scan pairs/ scan with old map
  - E.g. based on global pose estimated (chain rule) OR image similarity (bag of words)
- Loop Closure
  - E.g. use scan matching to find the transform AND its uncertainty
- Optimization
  - Pose Graph optimization (e.g. minimize error of poses, based on uncertainty)
  - Bundle Adjustment
- Map Rendering
  - E.g. generate grid map based on optimized graph

# Loop Closure


- Before loop closure





# Loop Closure

- After loop closure



# Loop Closure

- Recognizing an already mapped area, typically after a long exploration path (the robot “closes a loop”)
- Structurally identical to data association, but
  - high levels of ambiguity
  - possibly useless validation gates
  - environment symmetries
- Uncertainties collapse after a loop closure (whether the closure was correct or not)

# Loop Closure

- By revisiting already mapped areas, uncertainties in robot and landmark estimates can be reduced
- This can be exploited when exploring an environment for the sake of better (e.g. more accurate) maps
- Exploration: the problem of where to acquire new information

# Robust Loop Closing over Time for Pose Graph SLAM

---

Instituto de Investigación en  
Ingeniería de Aragón (I3A)

---

Yasir Latif: [ylatif@unizar.es](mailto:ylatif@unizar.es)  
José Neira: [jneira@unizar.es](mailto:jneira@unizar.es)

Volgenau School of Engineering  
George Mason University

---

César Cadena: [ccadenal@gmu.edu](mailto:ccadenal@gmu.edu)

*This work was supported by Spanish DPI2009-13710 and DPI2009-07130,  
by DGA-FSE (group T04), and by the US Army Research Office  
(W911NF-1110476)*

# THREE SLAM PARADIGMS

---

# The Three SLAM Paradigms

- Most of the SLAM algorithms are based on the following three different approaches:
  - Extended Kalman Filter SLAM: (called EKF SLAM)
  - Particle Filter SLAM: (called FAST SLAM)
  - Graph-Based SLAM

# EKF SLAM: overview

- **Extended state vector**  $y_t$  : robot pose  $x_t$  + position of all the features  $m_i$  in the map:

$$y_t = [x_t, m_0, \dots, m_{n-1}]^T$$

- Example: 2D line-landmarks, size of  $y_t = 3 + 2n$  : three variables to represent the robot pose +  $2n$  variables for the  $n$  line-landmarks having vector components

$$(\alpha_i, r_i)$$

$$y_t = [x_t, y_t, \theta_t, \alpha_0, r_0, \dots, \alpha_{n-1}, r_{n-1}]^T$$

- As the robot moves and takes measurements, the state vector and covariance matrix are updated using the standard equations of the extended Kalman filter
- Drawback: EKF SLAM is computationally very expensive.

# Particle Filter SLAM: FastSLAM

- **FastSLAM approach**

- Using particle filters.
- Particle filters: mathematical models that represent probability distribution as a set of discrete particles that occupy the state space.

- **Particle filter update**


- Generate new particle distribution using motion model and controls

- a) For each particle:

1. Compare particle's prediction of measurements with actual measurements
2. Particles whose predictions match the measurements are given a high weight

- b) Filter resample:

- Resample particles based on weight
- Filter resample
  - Assign each particle a weight depending on how well its estimate of the state agrees with the measurements and randomly draw particles from previous distribution based on weights creating a new distribution.




probability distribution (ellipse) as particle set (red dots)




# Graph-Based SLAM (1/3)

- SLAM problem can be interpreted as a sparse graph of nodes and constraints between nodes.
- The nodes of the graph are the robot locations and the features in the map.
- Constraints: relative position between consecutive robot poses, (given by the odometry input  $\mathbf{u}$ ) and the relative position between the robot locations and the features observed from those locations.



# Graph-Based SLAM (2/3)

- Constraints are not rigid but soft constraints!
- Relaxation: compute the solution to the full SLAM problem =>
  - Compute best estimate of the robot path and the environment map.
  - Graph-based SLAM represents robot locations and features as the nodes of an elastic net. The SLAM solution can then be found by computing the state of minimal energy of this net



# Graph-Based SLAM (3/3)


- Significant advantage of graph-based SLAM techniques over EKF SLAM:
  - EKF SLAM: computation and memory for to update and store the covariance matrix is quadratic with the number of features.
  - Graph-based SLAM: update time of the graph is constant and the required memory is linear in the number of features.
- However, the final graph optimization can become computationally costly if the robot path is long.
- Libraries for graph-based slam: g2o, ceres

# SLAM EXAMPLES

---


# 3D Scan matching using Spectral Method

Flood Gate (sonar)



Crashed car park

Disaster City (3D LRF)



# Jacobs 3D Mapping – Plane Mapping


Pose Graph

3D Range Sensing

Plane Extraction

Planar Scan Matching

Relax Loop-Closing Errors



# Plane Extraction from 3D Point Clouds

- **Plane Fitting**

- Assumes 3D sensor has radial Gaussian noise dependent on range
- Uses Approximate Least Squares solution to find the best fit.
- Estimates covariance matrix of the plane parameters

- **Range Image Segmentation**

- Is based on region growing algorithm
- Uses incremental formulas, therefore is fast
- Has linear computational complexity


**Given a range image, returns a polygonal model i.e. a set of planar features and boundaries.**

# Plane Registration (Scan Matching)

- Determining the correspondence set maximizing the global rigid body motion constraint.
- Finding the optimal decoupled rotations (Wahba's problem) and translations (closed form least squares) with related uncertainties.
- No motion estimates from any other source are needed.
- Very fast
- MUMC: Finding Minimally Uncertain Maximal Consensus
  - Of matched planes
- Idea: select two non-parallel plane matches  $\Rightarrow$  fixes rotation and only leaves one degree of translation!




# Relaxation of Errors (Translation)



- **Only translation errors are relaxed**
  - Good rotation estimates from the plane matching
  - Non-linear optimization can be exchanged with linear if rotation is assumed to be known precisely.
  - This leads to a fast relaxation method

Experiment Lab Run: 29 3D point-clouds; size of each:  $541 \times 361 = 195,301$





**Universidad**  
Zaragoza



Instituto Universitario de Investigación  
**en Ingeniería de Aragón**  
**Universidad Zaragoza**

# ORB-SLAM2: an Open-Source SLAM System for Monocular, Stereo and RGB-D Cameras

Raúl Mur-Artal and Juan D. Tardós

raulmur@unizar.es

tardos@unizar.es

# LOAM: Lidar Odometry and Mapping in Real-time

(Open Source Code and Open Datasets)

The Field Robotics Center  
At the Robotics Institute of Carnegie Mellon University



