AI for Autonomous Robots: Bridging Theory and Practice

Day 3: Localisation and Mapping

Dr. Timothy Wiley
School of Computing Technologies
RMIT University

ESSAI July 2026

RMIT UNIVERSITY

Navigation

Alternative to A*

ESSAI July 2026

Wall Following

As silly as it may sound, one of the most effective navigation techniques in a environment is wall-following. Method:

1. Stick you “hand” on a left/right wall
2. Only move around the space following the wall

To overcome “islands” fake walls are added using the path taken in exploration.

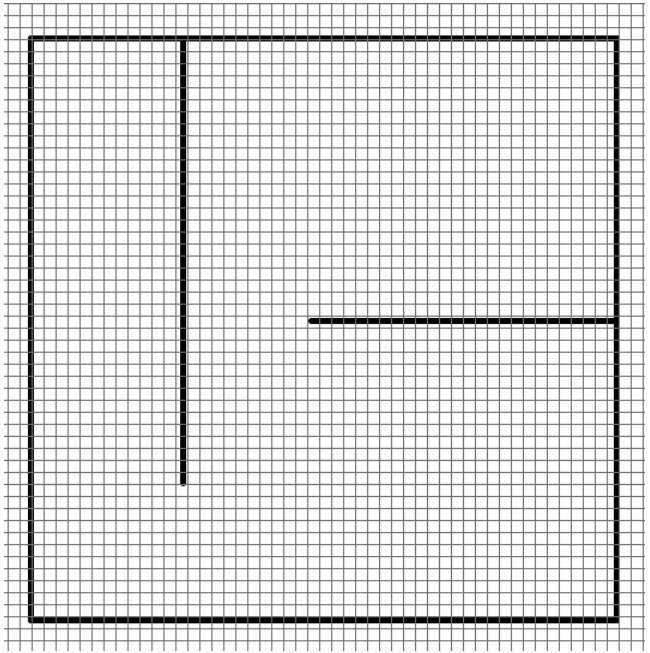

Rapidly-exploring Random Tree (RRT)

Both A* and D* have problems:

- Level of discretisation of an environment
- Computational overheads for large spaces
- Computation overhead for replanning
- Execution of navigation at the level of resolution suitable for the robot control

Rapidly-exploring Random Tree (RRT)

RRT is a sampling method for constructing global path plans. It builds a tree of points from the robot location to the goal. Algorithm:

1. Randomly choose a point in the map
2. Find the closest point in the tree to the chosen point
3. Add the chosen point to the tree iff:
    1. The point is within a configurable radius, and
    2. There is no obstacle between the chosen point and tree point
4. Repeat until the space is explored as desired

Rapidly-exploring Random Tree (RRT)

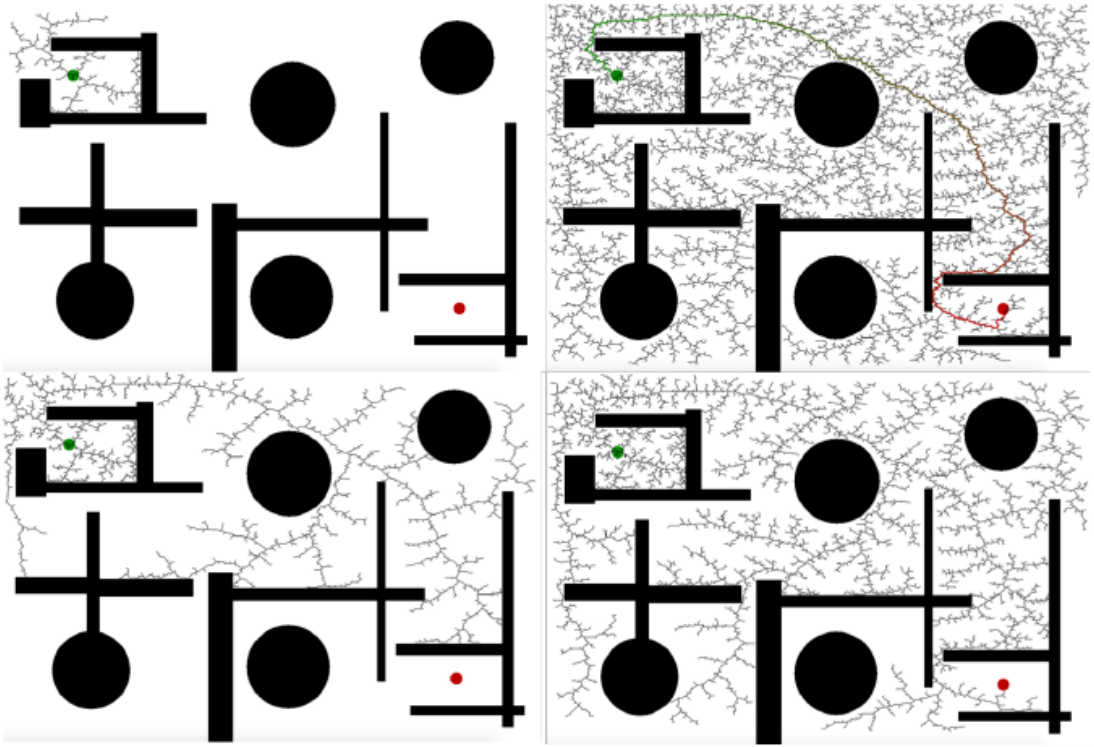

Image: Correll, 2022, introduction to Autonomous Robots

Rapidly-exploring Random Tree (RRT)

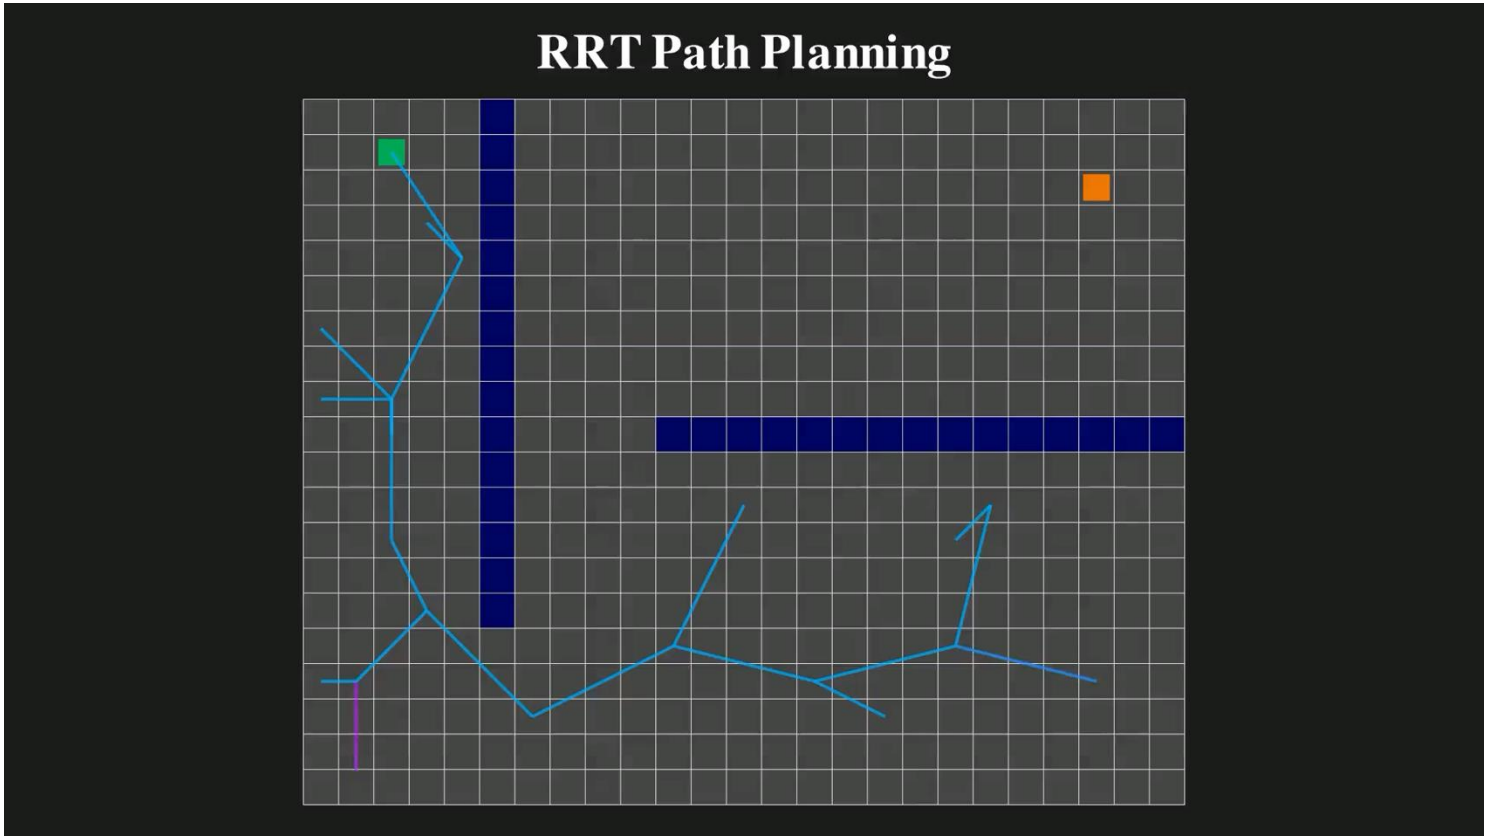

Dynamic Obstacles

A* (and RRT) path planning is not directly capable of re-planning when dynamic obstacles are encountered as navigation executes.

This issue can be “trivially” solved by re-planning with A* from the point a dynamic obstacle is detected.

However, this is computationally inefficient in the sense that existing information from the A* path is thrown-away and a new “global” path re-computed.

D*/D*-Lite

To overcome issues of replanning, the D*/D*-Lite family of algorithms by Anthony Stentz, Sven Koenig, and Maxim Likhachev perform incremental replanning (via A*) in a local space around the affected area where new obstacles are detected.

D*-Lite (Koenig & Likhachev 2002) is designed for dynamic re-updates as a robot moves and detects obstacles or updates to the map that affect it’s pre-computed path

D* Lite: Concept

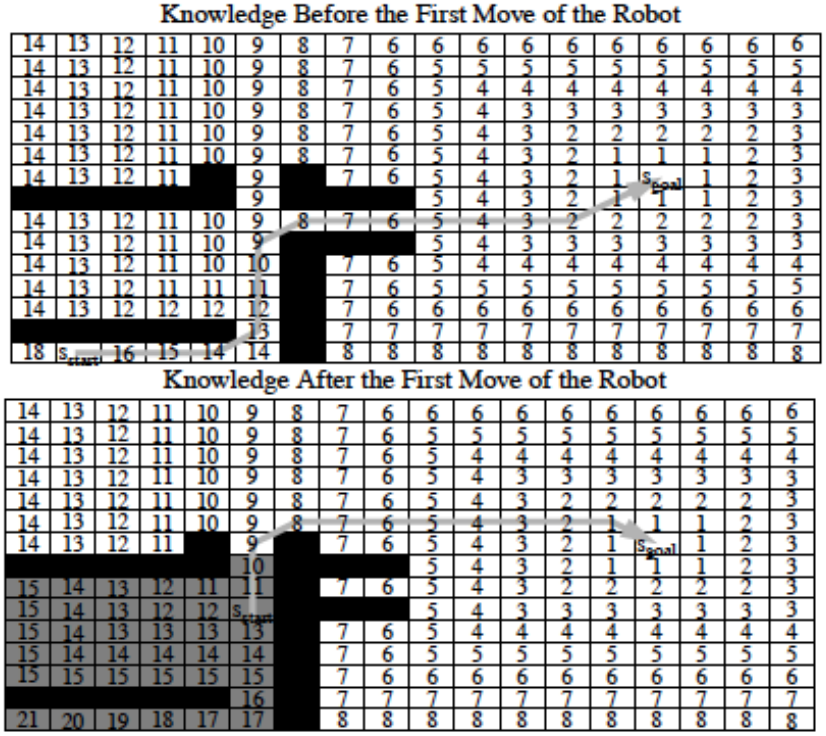

Image: Koenig, Sven, and Maxim Likhatchev. "D lite." Eighteenth national conference on Artificial intelligence. 2002.*

D* Lite: Comparison to A*

The broad differences of D*-Lite to A* are:

- Costs (and heuristics) are computed "in reverse" from the goal to the robot position. That is $g(s_{goal}) = h(s_{goal}) = 0$
- Costs of edges in the traversable graph are tracked for changes
  - A non-traversable edge has cost of ∞
- Maintains a secondary "1-step look ahead" value for each state

$$rhs(s) = \begin{cases} 0 & \text{if } s = s_{goal}, \\ \min_{s' \in Pred(s)} \left(g(s') + c(s', s)\right) & \text{otherwise .} \end{cases}$$

  - This is "more accurate" than $g(s)$ for changing edge costs
- Track k_m - an optimistic estimate of the cost to the goal

D* Lite: Overview

The general overview of D* Lite is:

1. Initialise with computing the shortest path (A*) to the start
2. If a change in the cost of a link is detected:
    1. Use a priority queue open list, U, for node evaluation order
    2. Modify the cost of the node s_{robot}
        1. Recalculate $rhs(s)$
        2. If $g(s) \neq rhs(s)$ add s to U using k_m to modify the priority of the nodes
    3. Recompute the shortest path for all nodes in U, always updating the cost of the node by (2)

D* Lite: Algorithm

procedure CalculateKey(s)
{01'} return $[\min(g(s), rhs(s)) + h(s_{start}, s) + k_m ; \min(g(s), rhs(s))]$;

procedure Initialize()
{02'} $U = \emptyset$;
{03'} $k_m = 0$;
{04'} for all $s \in S$ $rhs(s) = g(s) = \infty$;
{05'} $rhs(s_{goal}) = 0$;
{06'} U.Insert(s_{goal}, CalculateKey(s_{goal}));

procedure UpdateVertex(u)
{07'} if $(u \neq s_{goal})$ $rhs(u) = \min_{s' \in Succ(u)}(c(u, s') + g(s'))$;
{08'} if $(u \in U)$ U.Remove(u);
{09'} if $(g(u) \neq rhs(u))$ U.Insert$(u$, CalculateKey$(u))$;

procedure ComputeShortestPath()
{10'} while (U.TopKey() $\dot{<}$ CalculateKey(s_{start}) OR $rhs(s_{start}) \neq g(s_{start})$)
{11'} k_{old} = U.TopKey();
{12'} u = U.Pop();
{13'} if $(k_{old} \dot{<}$ CalculateKey$(u))$
{14'} U.Insert$(u$, CalculateKey$(u))$;
{15'} else if $(g(u) > rhs(u))$
{16'} $g(u) = rhs(u)$;
{17'} for all $s \in Pred(u)$ UpdateVertex(s);
{18'} else
{19'} $g(u) = \infty$;
{20'} for all $s \in Pred(u) \cup \{u\}$ UpdateVertex(s);

procedure Main()
{21'} $s_{last} = s_{start}$;
{22'} Initialize();
{23'} ComputeShortestPath();
{24'} while $(s_{start} \neq s_{goal})$
{25'} /* if $(g(s_{start}) = \infty)$ then there is no known path */
{26'} $s_{start} = \arg\min_{s' \in Succ(s_{start})}(c(s_{start}, s') + g(s'))$;
{27'} Move to s_{start};
{28'} Scan graph for changed edge costs;
{29'} if any edge costs changed
{30'} $k_m = k_m + h(s_{last}, s_{start})$;
{31'} $s_{last} = s_{start}$;
{32'} for all directed edges (u, v) with changed edge costs
{33'} Update the edge cost $c(u, v)$;
{34'} UpdateVertex(u);
{35'} ComputeShortestPath();

Image: Koenig, Sven, and Maxim Likhachev. "D lite." Eighteenth national conference on Artificial intelligence. 2002.*

D*/D*-Lite

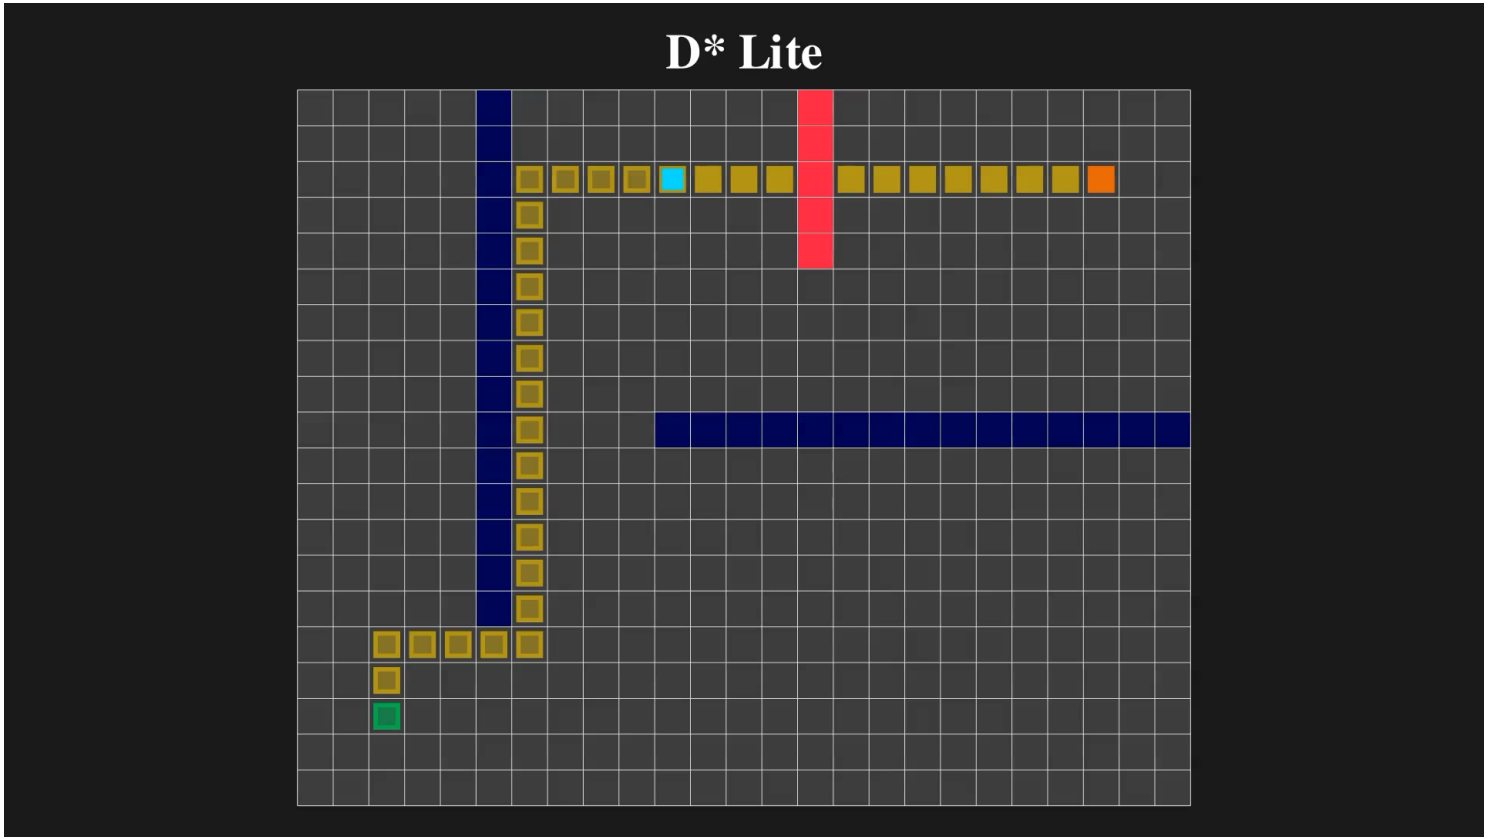

Navigation

Multi-Map planning & Re-Planning

ESSAI July 2026

Multiple Map resolutions

Planning in the global map is inefficient, especially for re-planning.

Instead, practical path planning uses multiple maps with :

- Different levels of resolution
  - Global map: course map
  - Local map: fine grained
- Different dimensions
- Different recalculation time frames

CostMaps can also be generated at both local and global views.

Re-Planning

A simple solution to both…

- Dynamic obstacles
- Unobserved spaces

… is replanning, that is, re-calculating the path periodically

However, this can lead to inefficiencies.

Mapping Calculation and Re-Planning Hierarchy

For practical efficiency, at each level, both mapping and re-planning are computed periodically at different frequencies

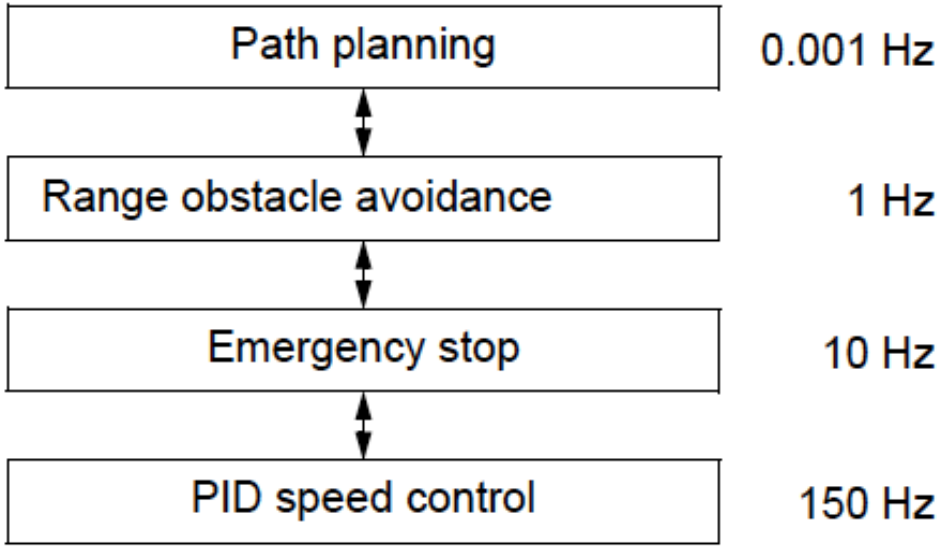

Image: Arkin, 2011, Introduction to Autonomous Mobile Robots

Map Generation at different levels

As noted earlier, the live generation of a “global map” is done through SLAM. However this is also computationally intense.

“Maps” at other levels are generally taken from direct sensor input.

For a typical robot with a laser, this is converting the laser scan directly into an occupancy grid

Obstacle Avoidance

With relatively quick re-planning at the local-map level, dynamic obstacles can be avoided as:

- The local occupancy map (and costmap) are quick to compute
- The dynamic obstacles are generally “small”
- Avoidance can be computed by re-planning in the local map

Worked example with ROSBot

SLAM

Map Generation

ESSAI July 2026

Localisation

… First this

ESSAI July 2026

Probabilistic Localisation

Localisation estimating the robot position in known map given:

- The estimate of the robots previous position
- How the robot moved
- What the robot has observed

This is computing the probability:

$$p(x_t|x_{t-1}, u_t, z_t)$$

- At time: t
- Pose: x_t
- Motion (odometry): u_t
- Measurement: z_t

Bayes Rule

$$P(x, y) = P(x \mid y)P(y) = P(y \mid x)P(x)$$

$$\Rightarrow$$

$$P(x \mid y) = \frac{P(y \mid x)\ P(x)}{P(y)} = \frac{\text{likelihood} \cdot \text{prior}}{\text{evidence}}$$

$$p(x_t | x_{t-1}, u_t, z_t) = \frac{p(x_{t-1}, u_t, z_t | x_t) p(x_t)}{p(x_{t-1}, u_t, z_t)}$$

Slide: Burgard, Fox & Thrun, 2006, Probabilistic Robotics

Markov Assumption

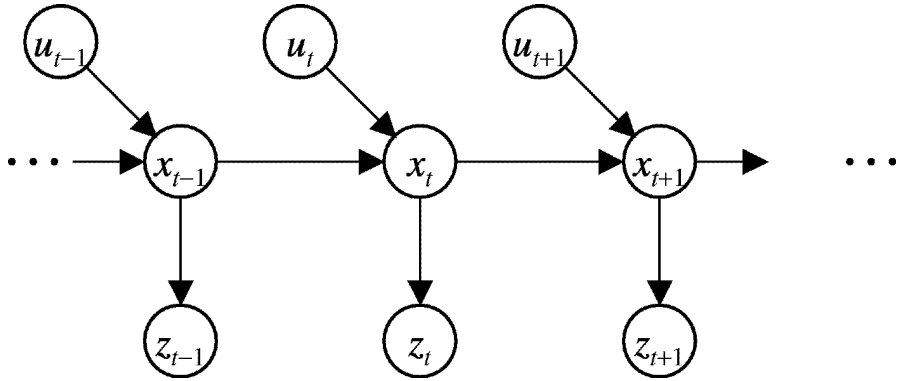

$$p(z_t \mid x_{0:t}, z_{1:t}, u_{1:t}) = p(z_t \mid x_t)$$

$$p(x_t \mid x_{1:t-1}, z_{1:t}, u_{1:t}) = p(x_t \mid x_{t-1}, u_t)$$

Underlying Assumptions

- Static world
- Independent noise
- Perfect model, no approximation errors

Slide: Burgard, Fox & Thrun, 2006, Probabilistic Robotics

Bayes Filters

$$Bel(x_t) = P(x_t \mid u_1, z_1 \ldots, u_t, z_t)$$

Bayes $$= \eta \; P(z_t \mid x_t, u_1, z_1, \ldots, u_t) \; P(x_t \mid u_1, z_1, \ldots, u_t)$$

Markov $$= \eta \; P(z_t \mid x_t) \; P(x_t \mid u_1, z_1, \ldots, u_t)$$

Total prob. $$= \eta \; P(z_t \mid x_t) \int P(x_t \mid u_1, z_1, \ldots, u_t, x_{t-1})$$

$$P(x_{t-1} \mid u_1, z_1, \ldots, u_t) \; dx_{t-1}$$

Markov $$= \eta \; P(z_t \mid x_t) \int P(x_t \mid u_t, x_{t-1}) \; P(x_{t-1} \mid u_1, z_1, \ldots, u_t) \; dx_{t-1}$$

Markov $$= \eta P(z_t \mid x_t) \int P(x_t \mid u_t, x_{t-1}) \, P(x_{t-1} \mid u_1, z_1, \ldots, z_{t-1}) \; dx_{t-1}$$

$$= \eta \; P(z_t \mid x_t) \int P(x_t \mid u_t, x_{t-1}) \; Bel(x_{t-1}) \; dx_{t-1}$$

Slide: Burgard, Fox & Thrun, 2006, Probabilistic Robotics

Bayes Rule

Markov Assumption and Conditional Independence allows the calculation of the "belief" in the robot's pose: $Bel(x_t)$

to a two-part method:

- Prediction: $\overline{Bel(x_t)} = \int p(x_t|u_t, x_{t-1}) Bel(x_{t-1}) dx$
- Correction: $Bel(x_t) = \eta p(z_t|x_t) \overline{Bel(x_t)}$
  - Where: η is a normalisation term
  - Note, the term belief is used as a synonym for the state probability

Bayes Rule: Decomposition

Why does it breakdown?

- Conditional independence gives that with complete information, two different probabilities are sufficient to give the full state
  - State Transition Probability:

$$p(x_t \mid x_{t-1}, u_t, z_t) = p(x_t \mid x_{t-1}, u_t)$$

  - Measurement Probability: $p(x_t \mid x_{t-1}, u_t, z_t) = p(z_t \mid x_t)$
- With incomplete information, we can combine these in the prediction and correction steps

The prediction and correction (measurement) steps are a core part of all localisation and mapping algorithms

Bayes Rule: Example

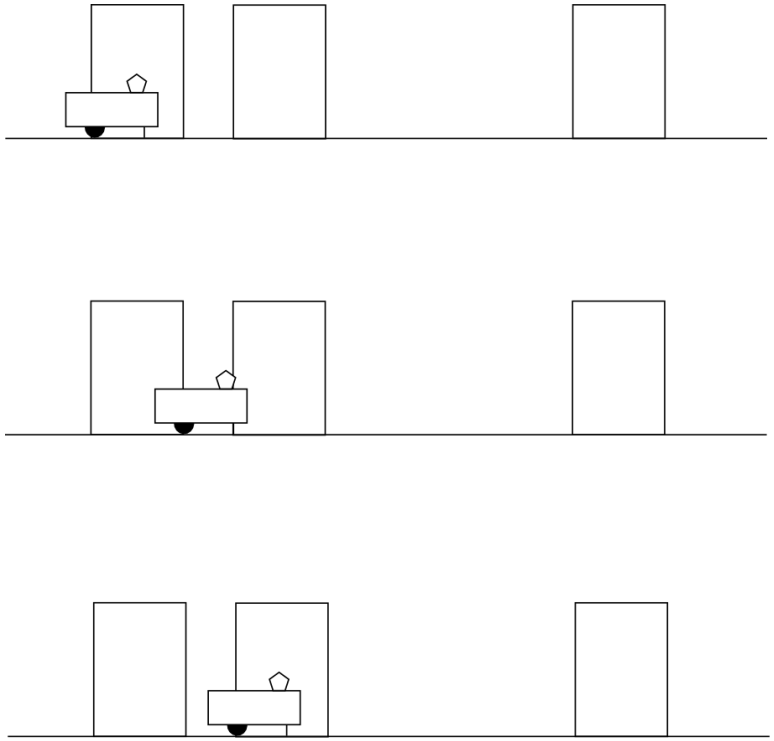

Image: Correll, 2022, introduction to Autonomous Robots

Bayes Rule: Example

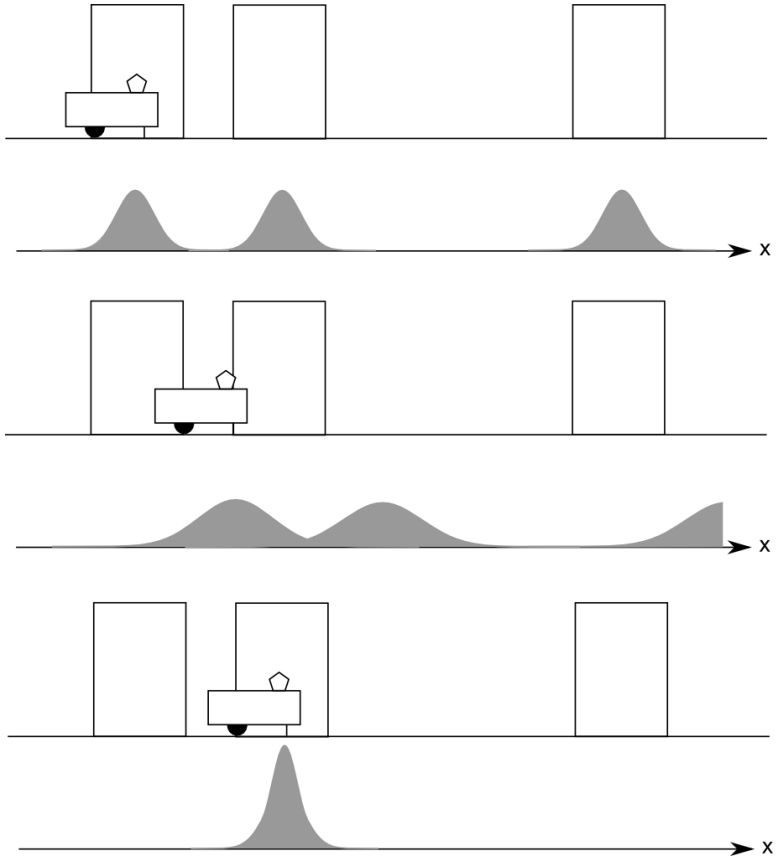

Image: Correll, 2022, introduction to Autonomous Robots

Diversions

Odometry &
ICP Matching

ESSAI July 2026

Odometry

Odometry is the estimation of the robot's pose based only on the actuators of the robot.

For example, on the ROSBot, odometery is computed using the encoder to estimate the "joint" position.

- Measures the rate at which the wheel turns
  - Estimate position of the robot from starting position
- Accurate, only if:
  - Encoder reliable
  - Time reliable
  - No slip
- Still drifts over time

Odometry

In ROS:

- If a robot supports Odometry, then the frame typically is /odom
- The ROSBot software launches with an odometry measurement

ICP: Iterative Closet Point Matching

For localisation (using laser scans and occupancy grids) a way to “match up” a laser scan to an occupancy grid is required.

ICP (Iterative Closest Point) laser scan matching is a technique to align two or more scans. The goal is to estimate the transformation (i.e., rotation and translation) that relates the scans to each other.

ICP: Iterative Closet Point Matching

ICP laser scan matching works by:

1. Identifying corresponding points between two scans, which are typically found by searching for the nearest point in one scan to each point in the other scan.
2. Iteratively updates the transformation estimate to minimize the distance between the corresponding points. There are many ways to transform the pose:
    1. Gradient descent
    2. Nonlinear optimization technique, such as the Gauss-Newton method, which adjusts the transformation estimate to reduce the overall distance between the points.

ICP: Iterative Closet Point Matching

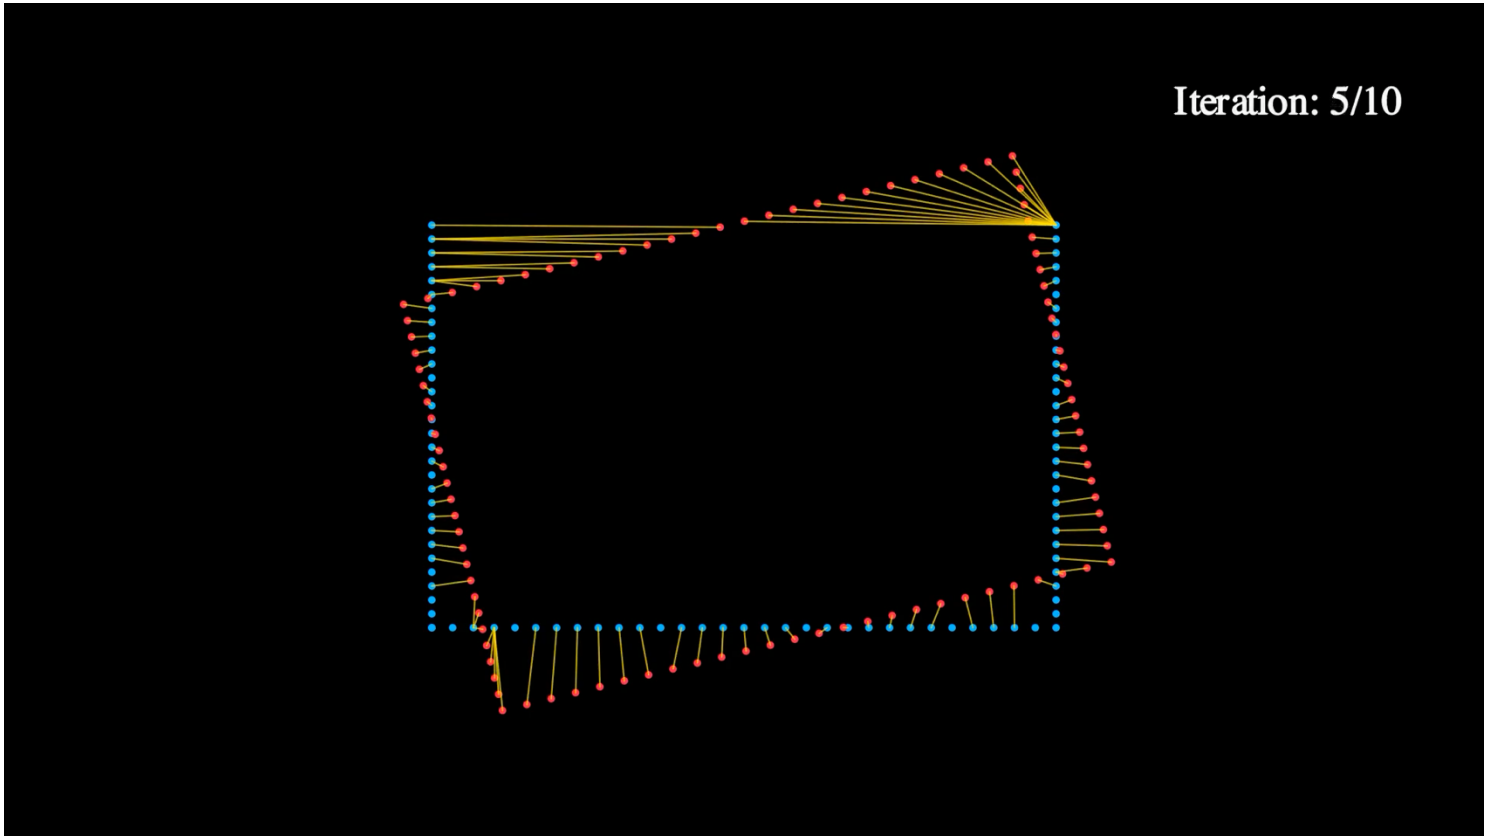

ICP: Iterative Closet Point Matching

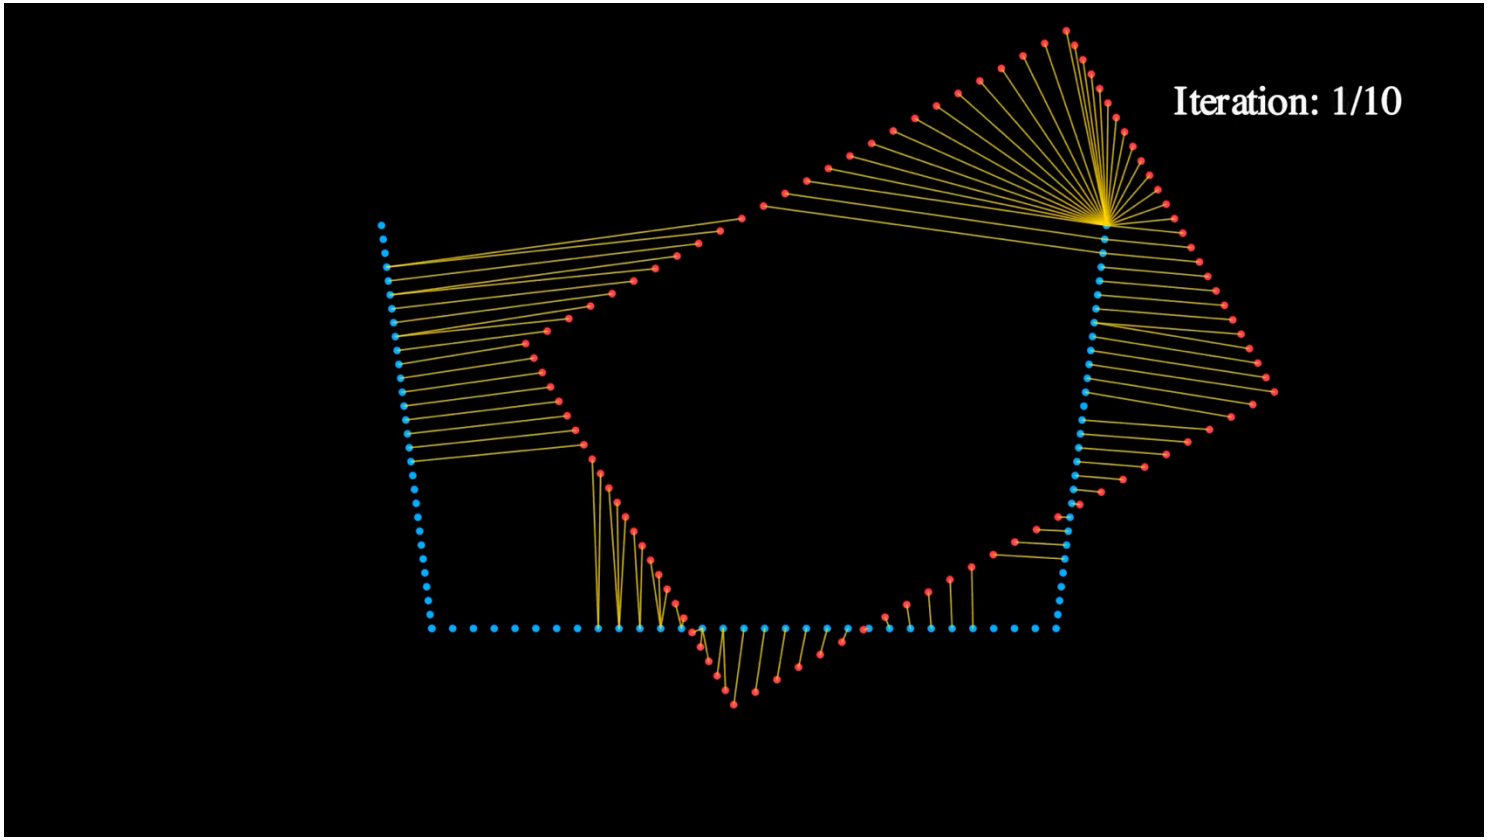

Localisation

Particle Filter

ESSAI July 2026

Particle Filter Localisation

The Particle Filter approach does away with trying to measure and calculate the probabilities.

Instead, it uses a random sampling approach via “particles”. The assumption is that with enough sampling the estimate of the robot’s pose can be found to within an acceptable error margin.

The particle filter follows the same two step process:

1. Predict where the robot has gone creating new particles
2. Correct and Update the most likely particles

Particle Swarm Algorithm

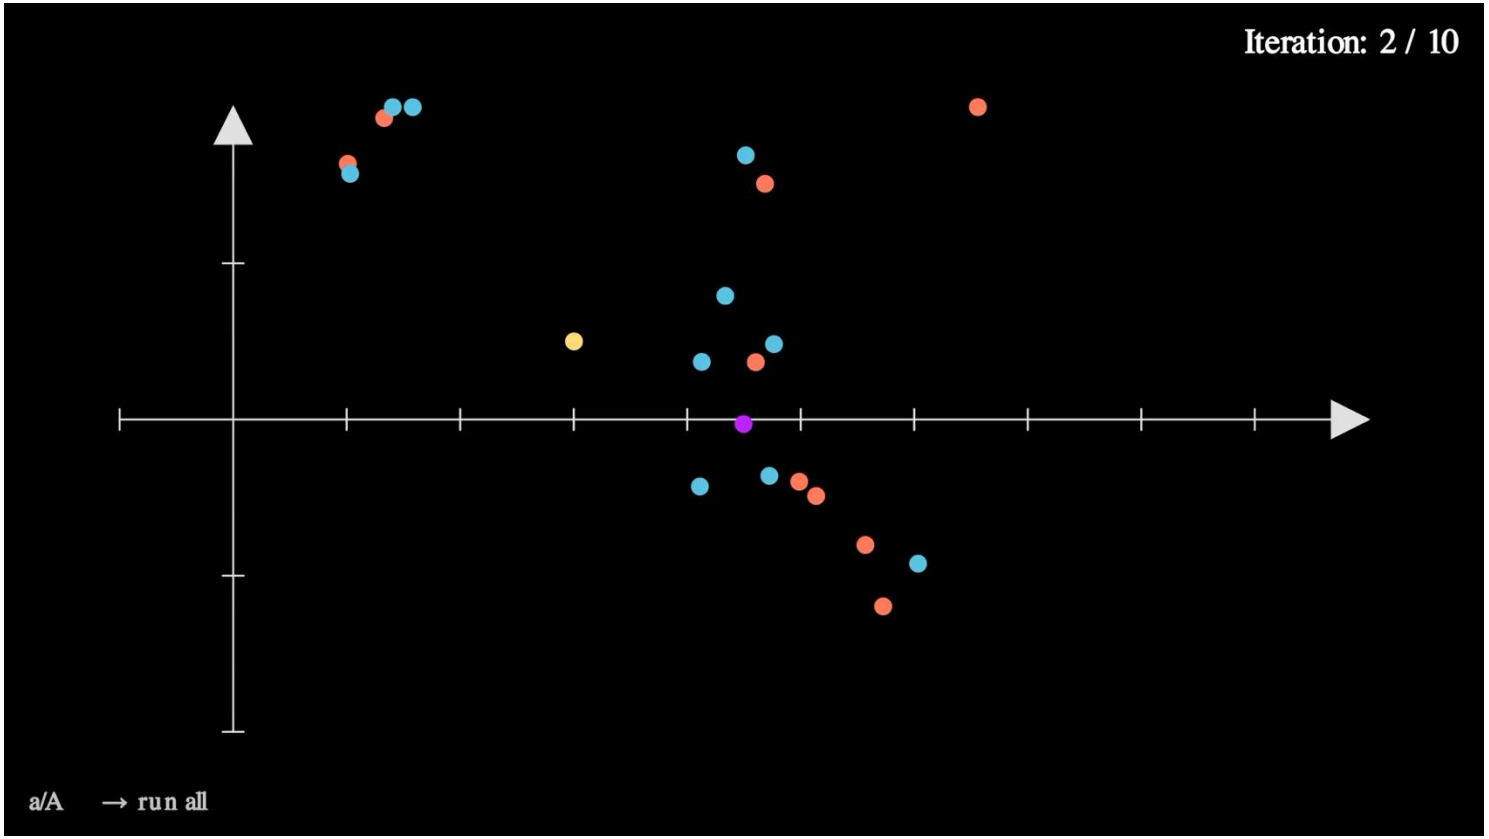

Particle Filter Localisation: Algorithm

1: **Algorithm Particle_filter**$(\mathcal{X}_{t-1}, u_t, z_t)$**:**
2: $\bar{\mathcal{X}}_t = \mathcal{X}_t = \emptyset$
3: for $m = 1$ to M do
4: sample $x_t^{[m]} \sim p(x_t \mid u_t, x_{t-1}^{[m]})$
5: $w_t^{[m]} = p(z_t \mid x_t^{[m]})$
6: $\bar{\mathcal{X}}_t = \bar{\mathcal{X}}_t + \langle x_t^{[m]}, w_t^{[m]} \rangle$
7: endfor
8: for $m = 1$ to M do
9: draw i with probability $\propto w_t^{[i]}$
10: add $x_t^{[i]}$ to $\mathcal{X}_t$
11: endfor
12: return $\mathcal{X}_t$

Image: Burgard, Fox & Thrun, 2006, Probabilistic Robotics

Particle Filter Localisation: Algorithm

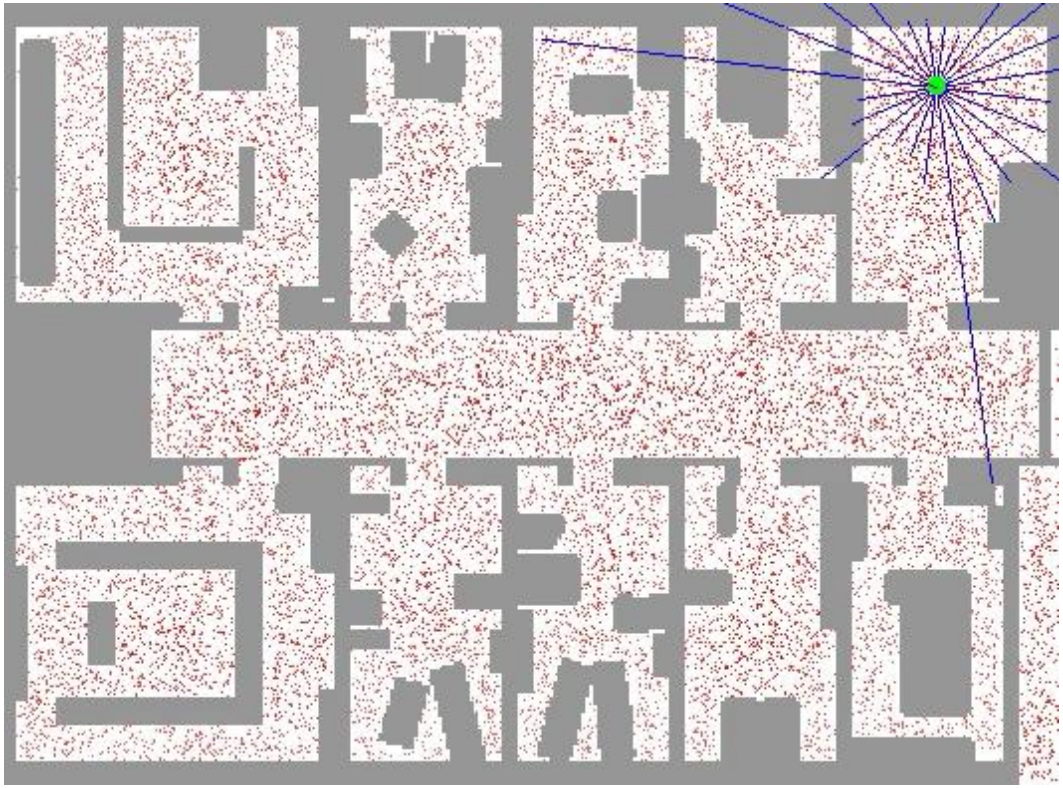

Image: Burgard, Fox & Thrun, 2006, Probabilistic Robotics

Particle Filter Localisation: Algorithm

An interesting property is that the particle filter maps to the bayes estimation:

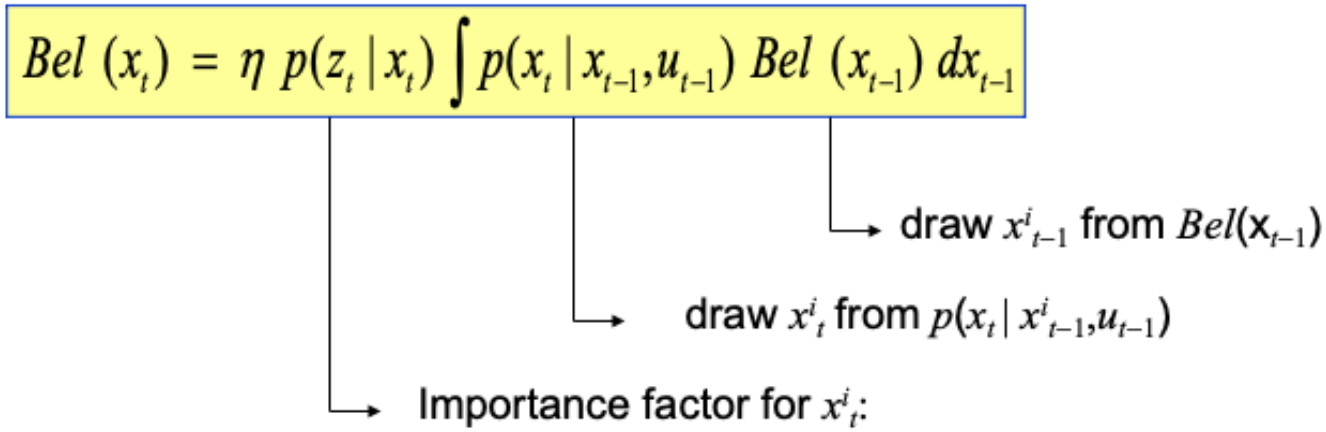

Image: Burgard, Fox & Thrun, 2006, Probabilistic Robotics

Localisation

Kalman Filter

ESSAI July 2026

Kalman Filter Localisation

The Kalman Filter is an estimator under two critical assumptions:

- Linear system
- Gaussian system

It is a generalised method to compute the state vector to predict in the case of localisation, the pose (position + orientation) of the robot. However, the Kalman filter can estimate more general state vectors if desired, including velocity and acceleration.

Kalman Filter

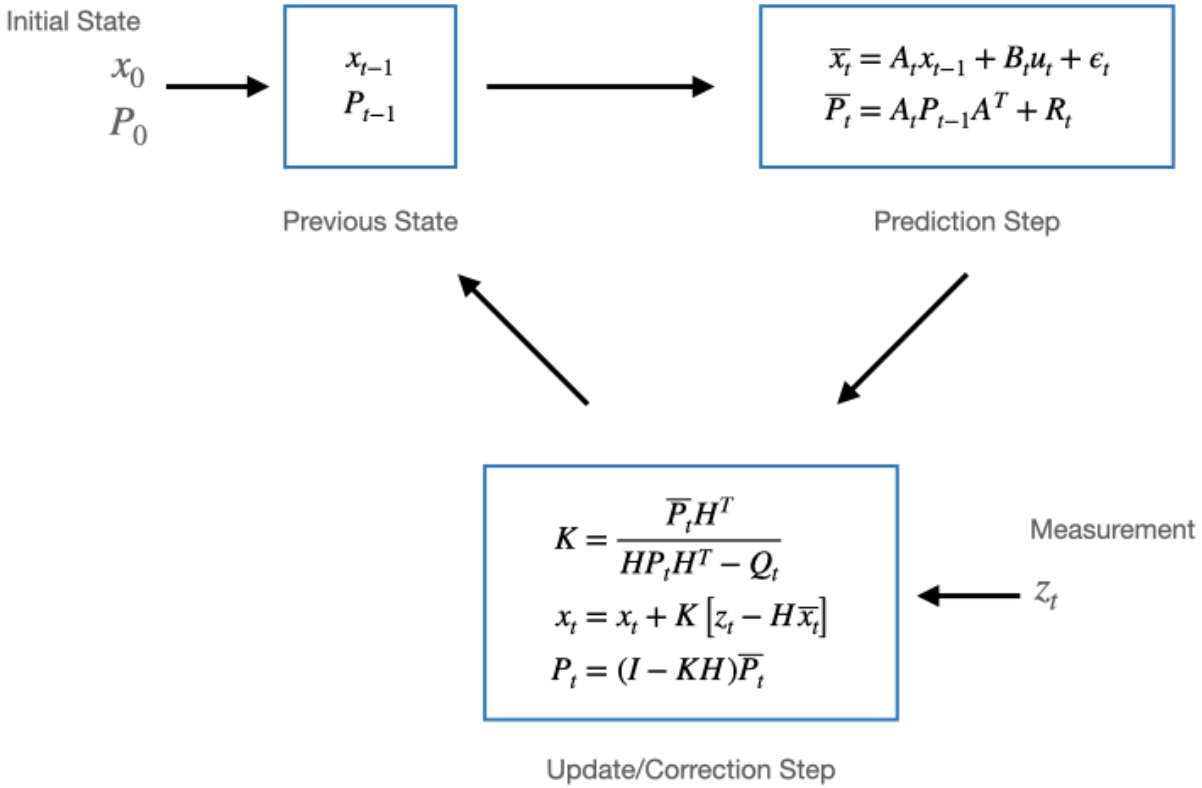

Kalman Filter

The Kalman Filter acts over a discrete time-step

- Matrices A, B, H Are all "state/representation transformation" matrices
- P_t models the prediction error:
  - ε - error of motion update
  - R_t - process error/noise, static increase in the uncertainty due to the motion update
  - Q_t - sensor error/noise, error in measurement of the state from the sensors
  - Forms a Covariance Matrix
- K is the Kalman Gain, which is the ratio of how much to trust the predication vs the measurement
- The Innovation, is the difference between the measurement and prediction

Covariance Matrix of Error

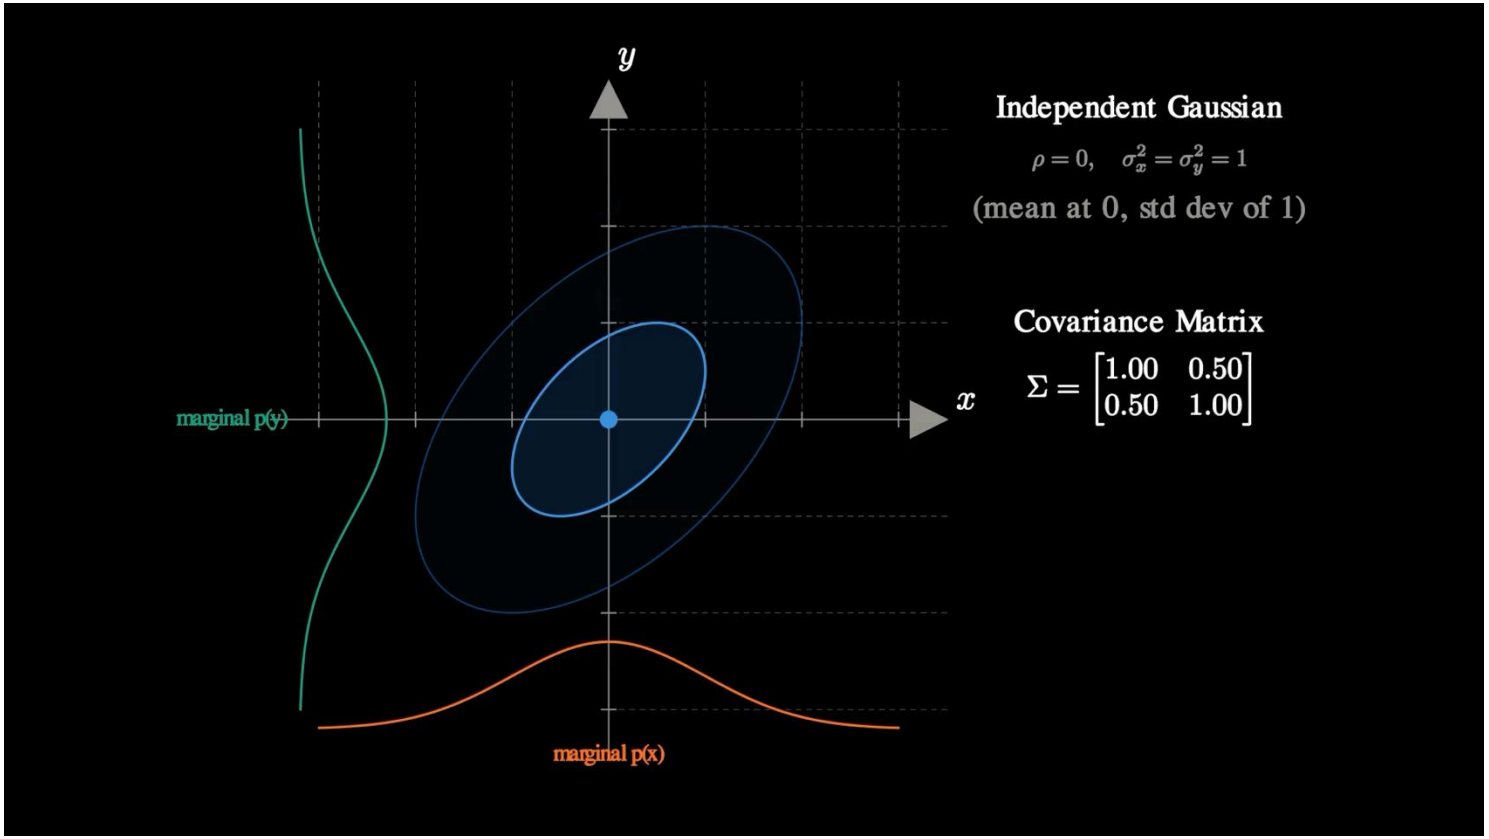

Kalman Filter Linear Assumption

The Kalman Filter assumes a linear system to ensure convergence

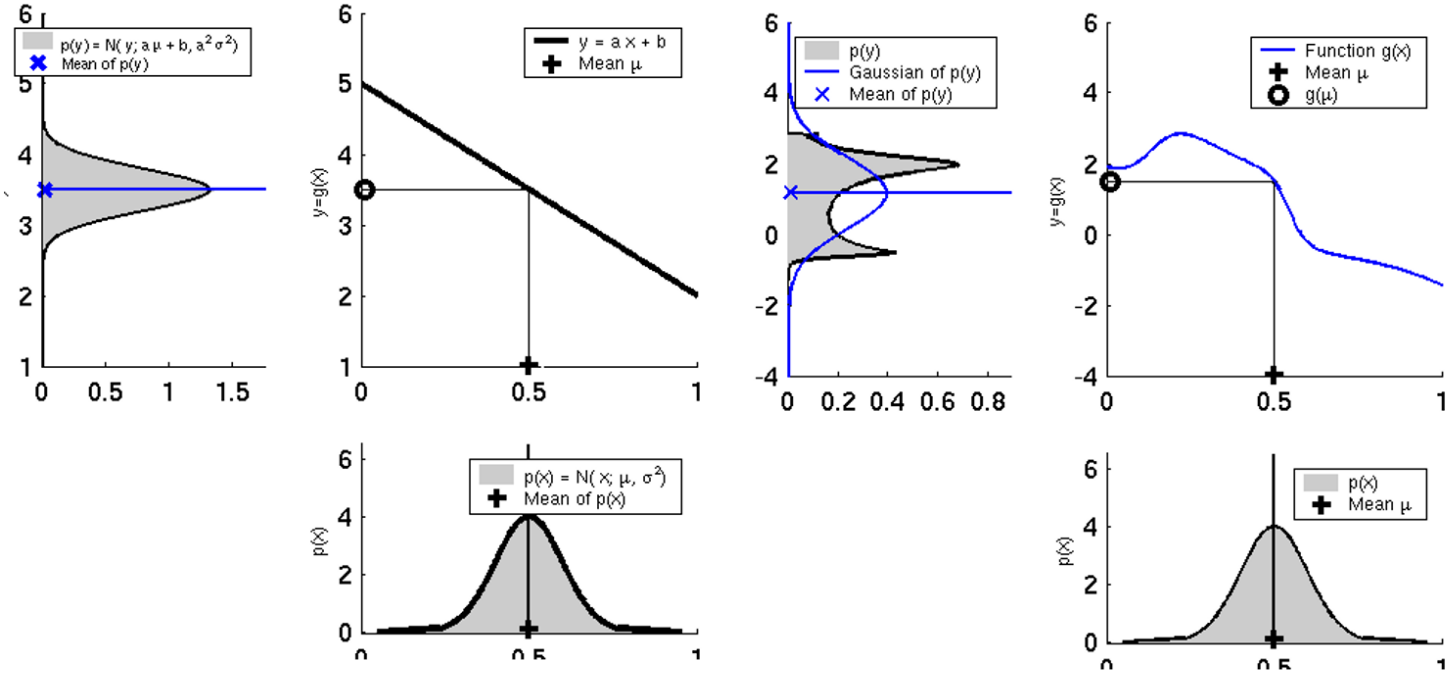

Extended Kalman Filter

The Extended Kalman Filter, makes a local linearisation assumption of a non-linear system at "the mean", that is, the position estimate for localisation.

Mathematic revision to ensure local linearisation:

- Transform the pose through the robot kinematics
- Replace the covariance matrix with the Jacobian matrix at the local point, essentially the partial derivates of all dimensions at "the mean"

SLAM

Map Generation

ESSAI July 2026

Simultaneous Localisation and Mapping

It is as the name suggests SLAM creates a map of the environment, at the same time as determining the robot's position within the map

The two cannot be separated, since generating the map from the robot's sensor observations requires knowledge of where the robot is within said map!

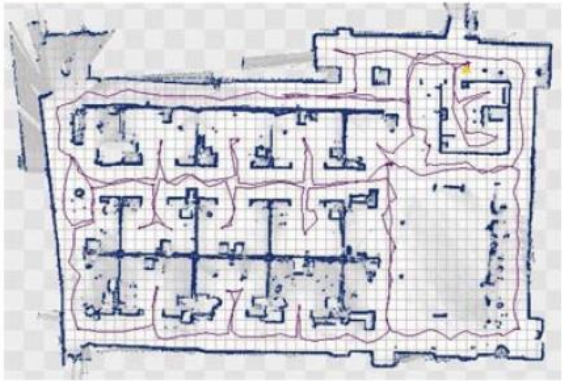

Simultaneous Localisation and Mapping

Thus SLAM estimates what the map should look like. SLAM is an extension to the localisation equation:

$$p(x_t, m | x_{t-1}, u_t, z_t)$$

- where m is the map

FastSLAM

FastSLAM is a particle filter-based SLAM approach. In FastSLAM:

- Each particle represents a possible location and map of the environment.
- As the robot moves, each particle is updated to refine its estimates of the location and map
- Particles are resampled based on an error measurement of both the location and map step

FastSLAM

FastSLAM uses "landmarks" in the environment. Landmarks are stored in a tree structure

- Each node in the tree represents a cluster of landmarks that are close together
- The root node represents the entire environment
- Each level of the tree represents a smaller region of space
- When a new landmark is detected, it is assigned to the node that corresponds to its location.
  - If the node does not yet, a new node is created.
  - As the robot moves and observes more landmarks, the nodes in the tree are updated.

FastSLAM

The key advantages of the landmark tree is that only a small part of the map needs to be updated without having to recompute the entire map every time new laser data is received. Only the affected nodes need to be updated.

FastSLAM

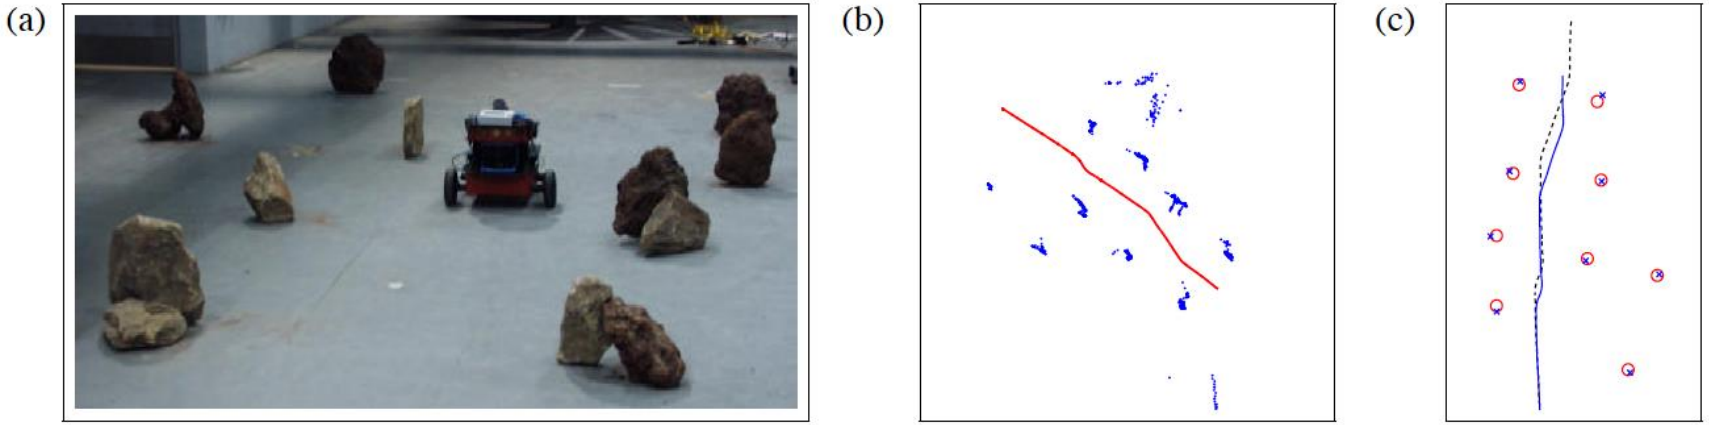

Figure 4: (a) Physical robot mapping rocks, in a testbed developed for Mars Rover research. (b) Raw range and path data. (c) Map generated using FastSLAM (dots), and locations of rocks determined manually (circles).

Images: Montemerlo, Michael, et al. "FastSLAM: A factored solution to the simultaneous localization and mapping problem." AAAI (2002)

FastSLAM

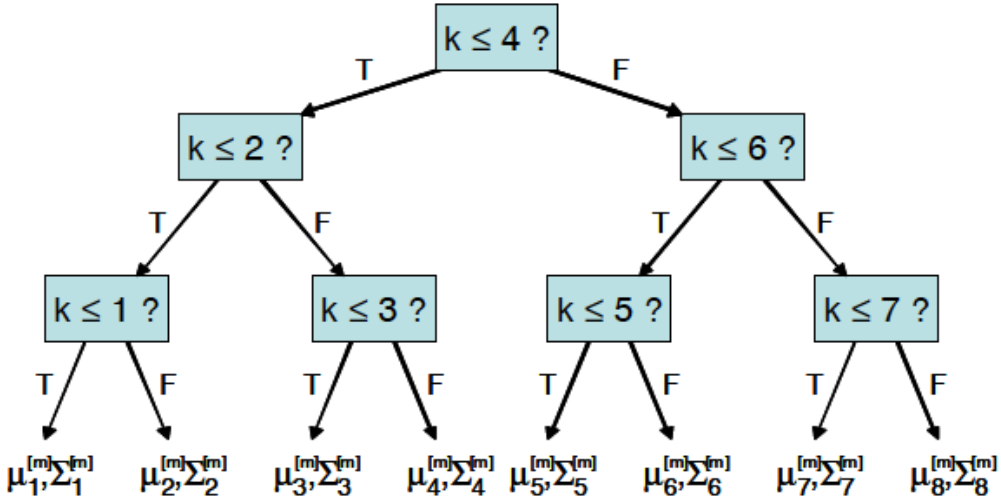

Figure 2: A tree representing $K = 8$ landmark estimates within a single particle.

Images: Montemerlo, Michael, et al. "FastSLAM: A factored solution to the simultaneous localization and mapping problem." AAAI (2002)

Loop Closure

Loop closure occurs when the robot revisits the same part of the environment twice, and the map should "close the loop" of map, that is, ensure the walls (or features) of the map align at the location where the robot revisited.

Loop closure is challenging because:

- Odometry errors from robot motor encoders
- Estimation errors in map computation
- Mapping algorithms that perform partial map updates

Loop Closure

Generally, most SLAM methods fail at proper loop closure unless it is explicitly handled by the algorithm:

The landmark tree in FastSLAM allows errors in landmarks when loop closure is detected to be "distributed" among the landmark tree.

FastSLAM: Loop Closure

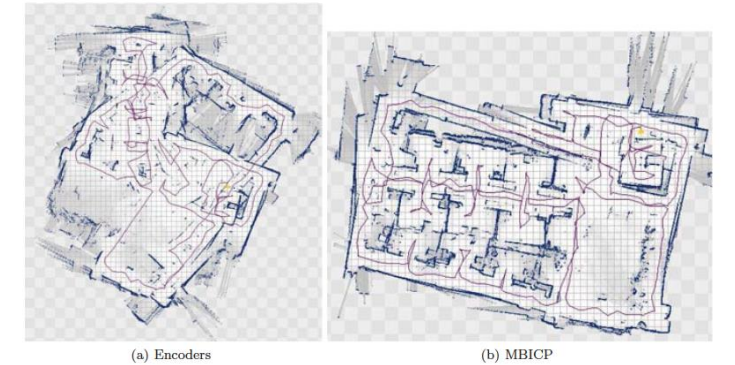

(a) Encoders (b) MBICP

(c) OG-MBICP (d) FastSLAM

Figure 3. Full map of office environment.

Images: Milstein, A., McGill, M., Wiley, T., Salleh, R. & Sammut, C. A Method for Fast Encoder-Free Mapping in Unstructured Environments. Journal of Fields Robotics, 28, 817–831 (2011)

FastSLAM: Loop Closure

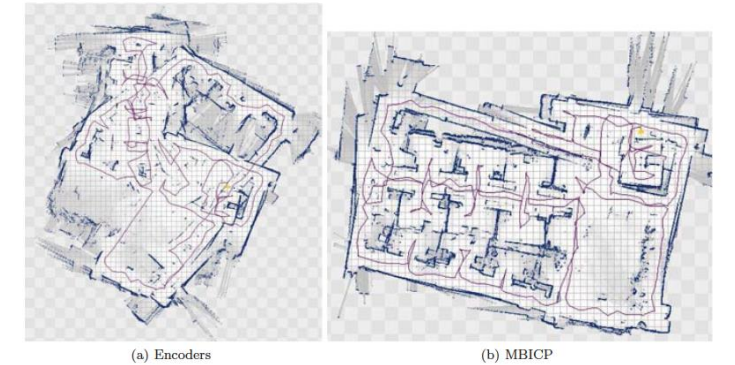

(a) Encoders (b) MBICP

(c) OG-MBICP (d) FastSLAM

Figure 3. Full map of office environment.

Images: Milstein, A., McGill, M., Wiley, T., Salleh, R. & Sammut, C. A Method for Fast Encoder-Free Mapping in Unstructured Environments. Journal of Fields Robotics, 28, 817–831 (2011)

GraphSLAM

GraphSLAM is a graph-based approach to SLAM that uses a graph representation of the map, which is then optimised as loops are detected. In GraphSLAM:

- Nodes of the graph are small localmaps that include a pose of the map, and map landmarks
- Edges represent kinematic displacements between the localmaps.
- In each localmap, ICP can be used to align new laser scans, and build-up the local map
- A new localmap (and node) is created if:
  - the robot moves too far to be outside a localmap
  - The ICP alignment error is too great

GraphSLAM

Loop closure is specifically detected by comparing the current localmaps to all other localmaps in the graph:

- When a loop is detected, a loop constraint is added to the graph
- A pose error is computed between the localmaps (for both displacement and rotation)
- This error is then "smoothed-out" or distributed between all edges of the graph to refine the kinematic transform between each localmap
- This results in a more accurate map

However, GraphSLAM is reliant on loop-closure to correct for mapping errors.

GraphSLAM

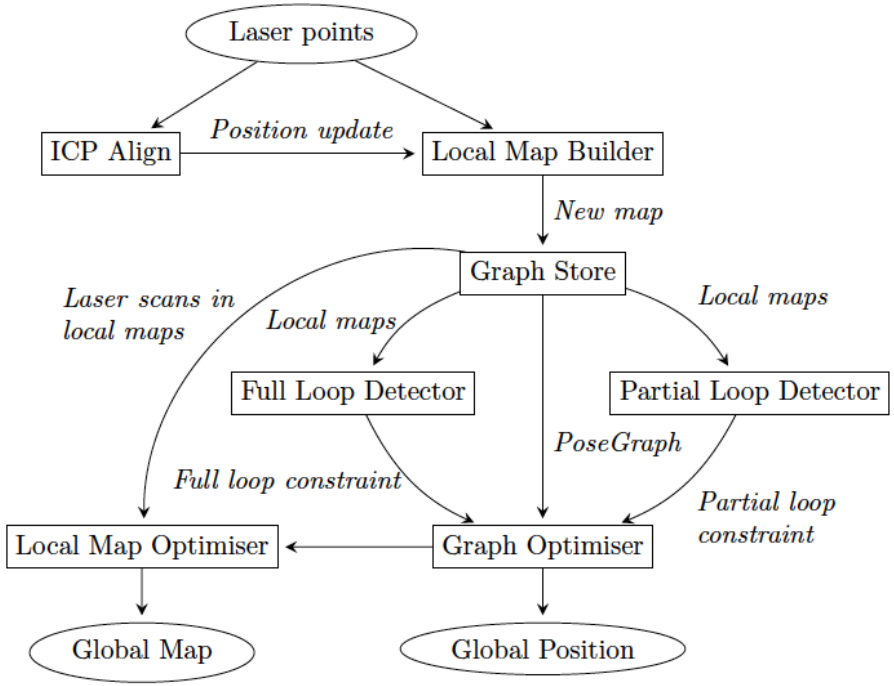

Images: Ratter, A. & Sammut, C. Local Map Based Graph SLAM with Hierarchical Loop Closure and Optimisation. in (2015).

GraphSLAM

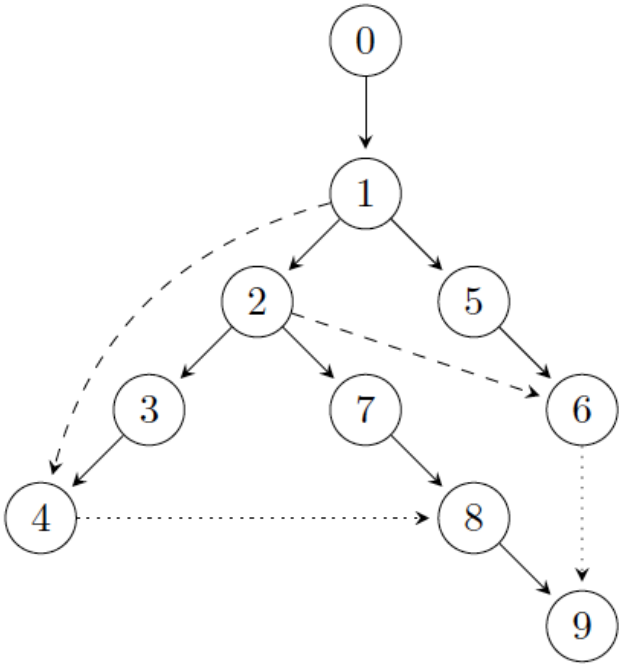

Images: Ratter, A. & Sammut, C. Local Map Based Graph SLAM with Hierarchical Loop Closure and Optimisation. in (2015).

Noon Gudgin

Thank you

Day 4: Task Planning

ESSAI July 2026