2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) October 25-29, 2020, Las Vegas, NV, USA (Virtual)

F1/10

Autonomous Racing

Simultaneous Localization And Mapping (SLAM)

otkey:	V 021+	Shift +	0.48+	Super +	
		Preview			
	30.97		Previewing 19	equines extra CF trame patent	NJ time
ize inc.	Prof.	Ma	dhu	r Beł	าไ
	-				
	Cor	mpute	er Sci	ence	
		CS 4	4501		
				art preview	
	Jnive	ersity	′ of ∖	′irgini	a
		Thread] Street (152 arted inco	amilis a mont ut		
-1-7281-	6211-9/20/\$3	1.00 ©2020 I	EEE		

Georbo	
	$\label{eq:construction} for a static property of the static product of the static prod$
- 1	File Edit View Search Serminal Help
ΩI <u>K</u> .	<pre>[18/0] [1546198341.471190, 1148.642000]: Loading controller: left_front_wheel_we locate controller.</pre>
	[18r0] [1546196341.496223, 1148.452000]: Loading controller: right_front_wheet_v
	elocity controller
	[Isro] [Iskeiseski.seriez, Iike.seskee]: Leasing controller: Cert_steering_winge position_controller
	[19F0] [1546190541.527393, 1148.686886]: Loading controller: right_steering_hing
DOM: UNITED OF	<pre>position_controlier [18F0] [1566196541.548764, 1148.7648860]: Loading controller: juint state control </pre>
	ler
	<pre>[INFO] [ISA4198341.SSE409, II48.FIE009]: Controller Spawmer: Loaded controllers: left_rear_wheel_velocity_controller, right_rear_wheel_velocity_controller, left _front_wheel_velocity_controller, right_front_wheel_velocity_controller, left_st eering_hinge_position_controller, right_steering_hinge_position_controller, join t_state_controller [INFO] [ISA4198341.SSE153, 1148.FIS6001: Started controllers: left_rear_wheel_ve</pre>
	locity controller, right_rear_wheel_velocity_controller, left_front_wheel_veloci
Contraction of the local division of the loc	ty_controller, right_front_wheel_velocity_controller, left_steering_hinge_positi
	[1870] [1140390342.029930836, 1149.121000000]: lookupTransform base_link to Las
	er timed out, Could not transform laser scan into base frame.
	1149.463000 seconds
	Contraction of the second seco
	Entered @ Now Camera III Select + Frank Camera Measure 20 Provide Insta
A 11	
1000 A	
and the second in	
Dente The Dente Art 15	
173	<u> </u>
4/3	Q Time .

Sun 1204

ROS Time: 1206.36 ROS Elapsed: 17.77 Wall Time: 1413.94 Wall Elapsed: 73.93 Experimental

Activities R Riz .				Wed 19:48		÷ 40 (5 •
			m	apping_demo.rviz* – RViz			ж
		(Blanding)	se Estimate 📝 2	2D Nav Goal 💡 Publish Point 🛟	= v		
		-				Hector	
- -						SLAM	
	-	-		8			
Topic	/map						
Alpha	0.7						
Color Scheme	map			1			
Draw Behind	0.05						_ ^
Width	2048						
Height	2048						
+ Position	-51,225: -51,225: 0						
+ Orientation	0: 0: 0: 1						
E 🖊 Path	1						
🗄 🖌 Status: Ok		×.			TARGAL		
Add	emove						
Time	The second second			a construction of the second s			×
ROS Time: 14585040	023.35 ROS Etapsed:	2.01 Wall	Time: 1460591289.42	Wall Elapsed: 22.83		Experiment	ntal
Reset Left-Click: Rot	ate. Middle-Click: Move X	(Y. Right-Click: Zoom. Shift	More options.			1	30 fps



- What is Simultaneous Localization and Mapping ?
- What are Occupancy Maps?
- How are SLAM and Scan Matching related ?
- SLAM in ROS Hector Mapping
- Google Cartographer SLAM Overview

Problem Setting

	4.4.5	 	▶.	Z
		4	duine -	
•	[10.000
	_			ç <u> </u>





Localization: given a map, use sensor data to estimate the current pose of the robot Mapping: given robot pose at each time (trajectory), use sensor data to build map



Simultaneous Localization and Mapping (SLAM): use sensor data to build map and estimate robot trajectory

A brief history of SLAM

Why do we need a map?

- In order to support path planning
- Limiting the error in state estimates, by providing the opportunity to 'reset'
- Later... do we really need a map?

Historical Development (1986-2004): Probabilistic Foundations

- EKF (you will still find this in visual inertial odometry)
- Particle Filter (very efficient localization)
- We will cover these methods next class in the context of localization

Modern Era (2004-Now): Algorithmic Improvements

- Maximum a-posteriori Estimation
- Other names: factor graph optimization, graph-SLAM, smoothing and mapping (SAM), bundle adjustment

Limitations : Basic Path Planning

- High Level Path Assignments
 - 2nd right, 2nd right, 1st right, 1st left, 1st right



Race Lines





• Occupancy: binary R.V.

 $m_{x,y}:\{free,occupied\} \rightarrow \{\ 0\ ,1\ \}$

[Review – Into Probability] Given some probability space (Ω, P) , a **random variable** $X: \Omega \to R$ is a *function* that maps the sample *s*pace to the reals.

• Occupancy: binary R.V.

 $m_{x,y}:\{free, occupied\} \rightarrow \{0, 1\}$

Occupancy grid map

 fine-grained grid map where an occupancy variable associated
 with each cell
 x



- Occupancy grid mapping
 - : A Bayesian filtering to maintain a occupancy grid map.



Measurement









a range sensor





 $p(z = 1 | m_{x,y} = 1) : \text{True occupied measurement}$ $p(z = 0 | m_{x,y} = 1) : \text{False free measurement}$ $p(z = 1 | m_{x,y} = 0) : \text{False occupied measurement}$ $p(z = 0 | m_{x,y} = 0) : \text{True free measurement}$







$$Odd: = \frac{(X \ happens)}{(X \ not \ happens)} = \frac{p(X)}{p(X^c)}$$

More convenient when we use "Odd"

$$Odd((m_{x,y}=1) \ given \ z) = \frac{p(m_{x,y}=1|z)}{p(m_{x,y}=0|z)}$$

• Odd



• Odd

$$Odd = \frac{p(m_{x,y} = 1|z)}{p(m_{x,y} = 0|z)} = \frac{p(z|m_{x,y} = 1)p(m_{x,y} = 1)/p(z)}{p(z|m_{x,y} = 0)p(m_{x,y} = 0)/p(z)}$$
$$p(m_{x,y} = 0|z) = \frac{p(z|m_{x,y} = 0)p(m_{x,y} = 0)}{p(z)}$$
(Bayes' Rule)

• Take the log!

Odd:
$$\frac{p(m_{x,y} = 1|z)}{p(m_{x,y} = 0|z)} = \frac{p(z|m_{x,y} = 1)p(m_{x,y} = 1)}{p(z|m_{x,y} = 0)p(m_{x,y} = 0)}$$

Log-Odd:
$$\log \frac{p(m_{x,y} = 1|z)}{p(m_{x,y} = 0|z)} = \log \frac{p(z|m_{x,y} = 1)p(m_{x,y} = 1)}{p(z|m_{x,y} = 0)p(m_{x,y} = 0)}$$

 $= \log \frac{p(z|m_{x,y} = 1)}{p(z|m_{x,y} = 0)} + \log \frac{p(m_{x,y} = 1)}{p(m_{x,y} = 0)}$

 $\log odd^+ = \log odd_{492} meas + \log odd^-$

• Log-odd update



 $\log odd^+ = \log odd meas + \log odd^-$

• Log-odd update



 $\log odd^+ = \log odd meas + \log odd^-$

Measurement model in log-odd form

$$\log \frac{p(z|m_{x,y}=1)}{p(z|m_{x,y}=0)}$$

Two possible measurement:

(Trivial Case : cells not measured)

• Example

Constant Measurement Model

 $\log odd_occ \coloneqq 0.9$ $\log odd_free \coloneqq 0.7$

Initial Map:

 $\log odd = 0 \quad \text{for all (x,y)}$ $p(m_{x,y} = 1) = p(m_{x,y} = 0) = 0.5$

Update Rule:

 $\log odd += \log odd_meas$

	0	0	0	0	0	0	0	0
	0	0	0	0	0	0	Р	Ð
	0	0	0	0	0	0	Ø	0
t0	0	0	0	0	0	0	Œ	Ð

• Example

Constant Measurement Model

 $\log odd_occ \coloneqq 0.9$ $\log odd_free \coloneqq 0.7$

<u>Update</u>

Case I : cells with z=1

 $\log odd \leftarrow 0 + \log odd_occ$

Case II : cells with z=0

 $\log odd \leftarrow 0 - \log odd_free$

Update Rule:

 $\log odd += \log odd_meas$



• Example

Constant Measurement Model

 $\log odd_occ \coloneqq 0.9$ $\log odd_free \coloneqq 0.7$

<u>Update</u>

Case I : cells with z=1

 $\log odd \leftarrow 0 + \log odd_occ$

Case II : cells with z=0

 $\log odd \leftarrow 0 - \log odd_free$

Update Rule:

 $\log odd += \log odd_meas$







The Map

Body frame



















$$i = ceil(x/r)$$






Distance measurement: dKnown state: (x_1, x_2, θ)

$$\begin{bmatrix} x_{1,occ} \\ x_{2,occ} \end{bmatrix} = \begin{bmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{bmatrix} \begin{bmatrix} d \\ 0 \end{bmatrix} + \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}$$

$$\begin{bmatrix} i_{1,occ} \\ i_{2,occ} \end{bmatrix} = ceil \left(\frac{1}{r} \begin{bmatrix} x_{1,occ} \\ x_{2,occ} \end{bmatrix} \right)$$



The Map



The Map



The Robot



The Map

Distance measurement: $(d_1, d_2, d_3, d_4, d_5)$ Directions of rays: $(\alpha_1, \alpha_2, \alpha_3, \alpha_4, \alpha_5)$ Known state: (x_1, x_2, θ)

For *k*-th occupied cell: $\begin{bmatrix} x_{1k} \\ x_{2k} \end{bmatrix} = \begin{bmatrix} d_k \cos(\theta + \alpha_k) \\ -d_k \sin(\theta + \alpha_k) \end{bmatrix} + \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}$ $\begin{bmatrix} i_{1k} \\ i_{2k} \end{bmatrix} = ceil\left(\frac{1}{r} \begin{bmatrix} x_{1k} \\ x_{2k} \end{bmatrix}\right)$



The Map

Registering the first Scan



Registering the first Scan



Scan Matching

Pose of the Car at t = t1

Laser Scans w.r.t car at Time $t = t_1$



Laser Scans w.r.t car at Time $t = t_2$

Scan Matching



Laser Scans w.r.t car at Time $t = t_1$

Laser Scans w.r.t car at Time $t = t_2$







Multi-Resolution Map Representation



20 cm Grid Cell

10 cm Grid Cell

5 cm Grid Cell

Saving the map



Saving the map



- ROS Package called
 MAP Server
- Allows saving a map currently being published over /map topic
- The saved map can be loaded for future tasks.

System Tf tree



Parameters for Hector SLAM : ROS

- map resolution
- map_update_distance_thresh
- map_update_angle_thresh
- laser_max_dist
- update_factor_free
- update_factor_occupied

Google Cartographer Overview

The Problem

Why did Google (not Waymo) make a SLAM package?



What's different about Cartographer

The contribution of this paper is a novel method for **reducing the computational requirements of computing loop closure constraints** from laser range data.

Loop-closure

Two regions in the map are found to be the same region in the world even though their position is incompatible given the uncertainty estimate in the map — the classic loop closure problem.

The system must then be able to calculate the transformation needed to align these two regions to 'close the loop'.

529



System Overview: Sensor Inputs

Up to four types of sensor measurements: (1) Range (II) Odometry (III) IMU (IV) Fixed Frame Pose

The FI/IO vehicle supplies two sensors:

HOKUYO

2D LIDAR

(Range)



System Overview: Frontend

Cartographer consists of two separate subsystems: local SLAM (frontend or trajectory builder) and global SLAM (backend).

The job of local SLAM is to generate good submaps.

The job of global SLAM is to tie the submaps consistently together.



System Overview: Backend

Cartographer consists of two separate subsystems: local SLAM (frontend or trajectory builder) and global SLAM (backend).

The job of local SLAM is to generate good submaps.

The job of global SLAM is to tie the submaps consistently together.



Frontend: Local SLAM

What is a submap?

A submap is considered as complete when the local SLAM has received a given amount of range data.

Submaps must be small enough so that the drift inside them is below the resolution of the occupancy grid, so that they are locally correct.

On the other hand, they should be large enough to be distinct for loop closure to work properly. More on this later



What is a submap?

Submaps each have their own static transform and contain a collection of registered range-measurements.

Consecutive measurements are connected by constraints which are 'local'.

Here local means derived from odometry or recent scan overlaps and resultant scan matches.



What is a submap?

Submaps each have their own static transform and contain a collection of registered range-measurements.

Consecutive measurements are connected by constraints which are 'local'.

Here local means derived from odometry or recent scan overlaps and resultant scan matches. Once a motion between two scans is found by the scan matcher, it goes through a motion filter. A scan is inserted into the current submap only if its motion is above a certain distance, angle or time threshold.



Submaps can **store their range data in a probability grid**. For 2D a signed distance field representation is also possible.

Probability grids are a 2D table where each **cell has a fixed size and contains the odds** of being obstructed.

Odds are updated according to **"hits"** (where the range data is measured) and **"misses"** (the free space between the sensor and the measured points).



Updating the submap

- I. For every hit, we **insert the closest grid point** into the hit set.
- 2. For every miss, we insert the grid point associated with each pixel that intersects one of the rays between the scan origin and each scan point, excluding grid points which are already in the hit set.
- If a grid point has yet to be observed it is assigned a value p_{hit} or p_{miss} depending on which set it is in.

Define the submap as:

 $M: r\mathbf{Z} \times r\mathbf{Z} \to [p_{min}, p_{max}]$

Then the map is updated according to the following operation:

 $M_{new}(x) = \mathbf{clamp}\left(odds^{-1}\left(odds(M_{old}(x)) \, odds(p_{hit})\right)\right)$



Recall the definition of the odds function: $odds(p) = \frac{p}{1-p}$ Same compute new p from odds...



Scan Matching:

How do we know that 'hits' and misses map to a particular cell in the probability grid?

- The collection of scan points relative to a moving reference frame attached to the robot are, $H = \{h_k\}$ for k=1...K where h_k is a point in \mathbb{R}^2 .
- These points are placed in a submap at pose ξ in the global reference frame by applying a rigid body transform to each (rotation and translation).
- Submap construction is the iterative process of repeatedly aligning scan and submap coordinate frames. *Sound familiar?*

So we can understand why scan matching alone is insufficient for robust localization and explore the basics.

Does Cartographer use ICP?

No. Cartographer needs to support 3D LIDARs and the correspondence problem is insidious.

- Scan matching variants:
 - Iterative closest point (ICP)
 - Scan-to-scan
 - Scan-to-map
 - Map-to-map
 - Feature-based
 - RANSAC for outlier rejection
 - Correlative matching

Cartographer uses the Ceres-solver to formulate a nonlinear leastsquares correlative scan matching problem.

Closest-Point Matching

Find closest point in other the point set



Closest-point matching generally stable, but slow and requires preprocessing

Correlation-based Scan Matching

- In ICP, correspondences between two scans are explicitly computed, allowing a rigid-body transformation to be computed.
- Susceptible to local minima; poor initial estimates lead to incorrect data associations and divergence.
- Correlation based methods search for a rigidbody transformation (without computing correspondences) that projects one set of LIDAR points (the query scan) on top of a reference map.
- The reference map is generally implemented as a look-up table that assigns a cost for every projected query point.
Correlation-based Scan Matching



Backend: Global SLAM

Closing Loops

Goal: is our current scan in one of the submaps we have already seen?

Constraints take the form of relative poses ξ_{ij} and associated covariance matrices Σ_{ij} . Relative poses now include both submap and scan poses.

For a pair of submap i and scan j, the pose ξ_{ij} describes where in the submap coordinate frame the scan was matched.



Loop Closure



Huber Loss

Uses a Huber loss to make the objective less susceptible to spurious constraints which are often added in symmetric environments... False positive is a disaster



$$L_{\delta}(a) = egin{cases} rac{1}{2}a^2 & ext{for } |a| \leq \delta, \ \delta(|a| - rac{1}{2}\delta), & ext{otherwise.} \end{cases}$$

Activities RIAZ -			Wed 19:48			÷ •0) Ů ↓	
			m	apping_demo.rviz* – RViz			ж
		(BELIER)	se Estimate 📝 2	D Nav Goal 🂡 Publish Point 🕂	— v		
80						Hector	
	C.					SLAM	
€ ≪ • F.	2						
Topic	/map						
Alpha	0.7						
Color Scheme	map						5
Pesolution	0.05						- C
Width	2048			2			
Height	2048						
Position	-51.225; -51.225; 0						
Orientation	0; 0; 0; 1						
- 🥓 Path	1						
🗄 🖌 Status: Ok					AND CALCE		
Add	emove Rename			Þ.	1		
Time				Q			×
ROS Time: 14585040	ROS Elapsed:	2.01 Wall	Time: 1460591289.42	Wall Elap\$20: 22.83		Experiment	ital
Reset Left-Click: Rot	tate. Middle-Click: Move)	(Y. Right-Click: Zoom. Shif	t: More options.			3	30 fps

Problem Setting





Input

I. Laser Scan





Input

- I. Laser Scan
- 2. Control Input



Overview

Input

- I. Laser Scan
- 2. Control Input
- 3. Map



Overview

Input

- I. Laser Scan
- 2. Control Input
- 3. Map

Output I. Pose



Overview

Input

- I. Laser Scan
- 2. Control Input
- 3. Map

Output

- I. Pose
- 2. Particles



Particle Filter Localization



The key idea of Bayes filtering is to estimate a probability density over the state-space conditioned on the data. This posterior is typically called the belief and is denoted:

 $\operatorname{Bel}(x_t) = p(x_t \mid d_{0\ldots t})$ Belief or posterior At time t Robot state The data from

86

time 0 to t.



- I. How would you describe the robot state in the localization problem?
- 2. Is the belief a probability distribution, what kind, how do you write it down?
- 3. How would you describe the data from the FI tenth car?

 $\operatorname{Bel}(x_t) = p(x_t \mid d_{0\ldots t})$ Belief or posterior At time t Robot state The data from time 0 to t.

87

For mobile robots, we distinguish two types of data: perceptual data such as laser range measurements, and odometry data, which carries information about robot motion. Denoting the former by o (for observation) and the latter by a (for action), we have:

$$Belief or posterior Robot state = p(x_t \mid o_t, a_{t-1}, o_{t-1}, \dots)$$

Recursion

After some simplification (see additional resources) we have:

Integrate out over previous beliefs. In practice finite approximation

$$Bel(x_t) = \eta \ p(o_t \mid x_t) \int p(x_t \mid x_{t-1}, a_{t-1}) Bel(x_{t-1}) dx_{t-1}$$

Normalization Constant Make sure everything adds up to 1! **Sensor Model** Compute how likely your measurements were given updated particles.

Motion Model

Simulate noisy dynamics of particles based on control input.

Previous Belief

Draw your particles with replacement according to importance weight.

Read left to right

In practice we represent the distributions non-parametrically using particles (a finite set of samples). More in the next slides.

Initialization

Now let's look at particle filters in three dimensions and the process of using them in our car. Here we see a part of a map previously generated using Cartographer. The black pixels on the map denote the walls and the grey pixels denote free space.



Initialization

The red arrow indicates our initial pose obtained from user input. Lets use particle filters to solve the pose tracking problem and localize more accurately in the map.



Initialization

First we need to generate a set of hypothesis for our first position. These the discrete particles drawn from a Gaussian distribution of mean being the initial guess and with a small covariance.



Motion Model

Applying the measured control input to the motion model (with noise), yields an updated set of possible poses.



For each particle we can create a 'fake' laser scan by raycasting against the map at the particles pose.



Questions

How hard is creating the fake laser scan relative to computing the motion model update?

What is special about the sensor model?



Questions

How hard is creating the fake laser scan relative to computing the motion model update? Way harder.

What is special about the sensor model? It's embarrassingly



Raycasting

Efficient methods for raycasting such as Breshenham's Line method can be deployed. Alternately a signed distance field can be precomputed.



Computing Particle Weights

Recall, we have the 'real' laser scan which the car observed (shown in green). Note we don't know where the green dot is but we do know the range measurements.



Computing Particle Weights

We can compute a score for the fake laser scan:

$$p(o_t \mid x_t^i) = \frac{1}{\sigma\sqrt{2\pi}} e^{\frac{(o_t - r^i)^2}{2\sigma^2}}$$



The Algorithm

```
# a single recursive update step of the MCL algorithm
def MCL(X_{t-1}, a_{t-1}, o):
    X_t = \{\}
    for i in range(m):
         # sample a pose from the old particles according to old weights. Samples
         # implicitly represent the prior prob. dist. Bel(x \{t-1\}) in eqn. (3)
         x_{t-1}^{i} \sim X_{t-1}
         # update the sampled pose according to the motion model
         x_{t}^{i} \sim p(x_{t} | x_{t-1}, a_{t-1})
         # weight the updated pose according to the sensor model
         w_{t}^{i} = p(o | x_{t}^{i})
         # add the new pose and weight to the new distribution
         X_t = X_t \cup \{ (x_t^i, w_t^i) \}
    # normalize weights, should sum to 1
    X_{t} = normalize(X_{t})
    return X_{i}
# iterative application of the MCL algorithm
def particle_filter():
    X = Bel(x_o) \leftarrow initial particles
    while true:
         a = get_last_odometry()
         o = get last sensor readings()
         X = MCL(X, a, o)
         # inferred pose ← expected value over particle distribution
         pose = Ex[X]
                                                   572
```

Particle Filter without Resampling



Resampling

Resampling **Original Particles** A COMPANY OF A COMPANY ******** ******* WHERE AND CONTRACTOR OF THE OWNER. After N iterations



Particles



Particle Filters in ROS

• Adaptive Monte Carlo Localization Package

• Localization for a robot moving in a 2D space

• Localizes against a pre-existing map

AMCL Parameters

min_particles

Default: 100

The minimum number of particles to be used for calculating correlation

max_particles

Default: 500

The maximum number of particles to be used for calculating correlation

AMCL Parameters

update_min_d

Default: 0.2m

The minimum translation movement required by the vehicle before an pose update is published

update_min_a

Default: $\pi/_6$ radians

The minimum angular movement required by the vehicle before an pose update is published
AMCL Parameters

initial_pose_x Default:0
initial_pose_y Default:0
initial_pose_a Default:0

The initial mean position of the particles to initialize the particle filter

AMCL Parameters



The covariance of particles distributed around the mean

Tf tree – Where does AMCL fit in







References

S. Thrun, W. Burgard. "Probabilistic Robotics." Chapter 4 and Chapter 8. http://www.probabilistic-robotics.org/

S. Thrun. "Artificial Intelligence for Robotics, Lesson 3." Udacity. https://www.udacity.com/course/artificial-intelligence-for-robotics--cs373

S. Thrun, D. Fox, W. Burgard and F. Dellaert. "Robust Monte Carlo Localization for Mobile Robots." Artificial Intelligence Journal. 2001. <u>http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.71.6016&rep=rep1&type=pdf</u>

D. Fox, W. Burgard, and S. Thrun. "Markov localization for mobile robots in dynamic environments," Journal of Artificial Intelligence Research, vol. 11, pp. 391427, 1999. http://www.jair.org/media/616/live-616-1819-jair.pdf

D. Fox. "KLD-sampling: Adaptive particle filters," Advances in Neural Information Processing Systems 14 (NIPS), Cambridge, MA, 2002. MIT Press. <u>https://papers.nips.cc/paper/1998-kld-sampling-adaptive-particle-filters.pdf</u>

D. Bagnell "Particle Filters: The Good, The Bad, The Ugly" http://www.cs.cmu.edu/~16831-f12/notes/F14/16831_lecture05_gseyfarth_zbatts.pdf

C. Walsh, S. Karaman. "CDDT: Fast Approximate 2D Ray Casting for Accelerated Localization." Arxiv, 2017. http://arxiv.org/abs/1705.01167