

Lab 8: Building Occupancy Grids with TurtleBot*

EECS/ME/BIOE C106A/206A Fall 2022

Goals

By the end of this lab, you should be able to:

- Use the ROS parameter server to set parameter values that can be shared across multiple nodes
 - Understand and explain how an occupancy grid works and when to use one
 - Map out the lab space using your own custom occupancy grid
 - Point out any important deficiencies in your implementation
-

Contents

1	The ROS parameter server	2
2	Generating and updating the occupancy grid	2
2.1	Updating with log-odds	2
2.2	Using the laser scan	3
3	Testing your occupancy grid	4

Introduction

In this lab, we will build and test one of the most useful datastructures in mobile robotics: the occupancy grid.¹ The key idea behind the occupancy grid is to represent space as — you guessed it — a grid, in which every cell, or *voxel*, is either occupied or free. Since nothing is ever really certain in life (i.e., measurements are noisy), occupancy grids actually keep track of the *probability* that each cell is occupied. When the robot receives a measurement of the environment, typically from a laser scanner, it updates these probabilities to incorporate the new information.

This lab is broken up into three phases:

1. Learn how to use the ROS parameter server.
2. Write the key steps in an occupancy grid update.
3. Test your implementation and identify any shortcomings.

*Developed by David Fridovich-Keil and Laura Hallock, Fall 2017. Updated by Valmik Prabhu, Nandita Iyer, Ravi Pandya, and Philipp Wu, Fall 2018.

¹A Google search for “occupancy grid” turns up lots of great references that go into more detail.

1 The ROS parameter server

We haven't really exposed you to the ROS parameter server before, but since it is one of the more useful features of ROS, we want you to get some practice using it.² ROS parameters are key-value pairs that ROS allows you to specify when launching a nodes (e.g., in a `launch` file) that may be queried by those nodes at run-time. This can be an extremely useful tool for writing flexible code and for enforcing that multiple nodes hold the same value for some particular variable.

Begin by creating a new workspace called `lab8` in your `ros_workspaces` directory. Next, create a package named `mapping` inside the `src` directory of your `lab8` workspace. It should depend on `rospy`, `visualization_msgs`, `geometry_msgs`, `sensor_msgs`, `std_msgs`, and `tf2_ros`.

Our starter code for this lab is on GitHub for you to clone so that you can easily access any updates we make to the starter code. It can be found at <https://github.com/ucb-ee106/106a-fa22-labs-starter/tree/main/Lab8>. We also highly recommend you make a **private** GitHub repository for each of your labs just in case.

Create a folder in the `mapping` package called `launch` and copy the `demo.launch` file from the starter code into the new folder. Don't forget to build and source your workspace!

Inside the file `demo.launch`, you'll see that a number of command-line arguments are declared (along with default values). These arguments are then mapped to specific parameters in a node called `mapper`. These parameters will need to be read in by that node at run-time.

From the starter code, move `occupancy_grid_2d.py` and `mapping_node.py` into the appropriate location in your `mapping` package and make the files executable. In the `occupancy_grid_2d.py` file, locate the `LoadParameters` function. We've loaded one parameter for you, but you'll need to finish this function by loading the rest. Note that two of the variables in `occupancy_grid_2d.py`, `x_res` and `y_res`, are not on the parameter server. How do you think you should generate these variables? (You should not be editing the launch file.)

Checkpoint 1

Submit a checkoff request at <https://tinyurl.com/fa22-106alab> for a staff member to come and check off your work. At this point you should be able to:

- Run `demo.launch` without any error messages complaining about failing to load parameters.
- Explain each parameter you have loaded in `LoadParameters`.

2 Generating and updating the occupancy grid

Now for the fun part! In the file `occupancy_grid_2d.py` file, locate the function `SensorCallback` and fill in the details. The main idea here is that each grid cell contains the *log-odds ratio of occupancy*.

2.1 Updating with log-odds

We briefly introduce the mathematics of the log-odd update rule. Let X_{ij} be a binary random variable that indicates the *true* occupancy of the cell at row i and column j , where $X_{ij} = 1$ corresponds to cell being occupied and $X_{ij} = 0$ being free. Let Y_{ij} be another binary random variable that represents the *measured* occupancy of that cell.

By Bayes' Rule, we have

$$\Pr(X_{ij} | Y_{ij}) := \frac{\Pr(Y_{ij} | X_{ij})\Pr(X_{ij})}{\Pr(Y_{ij})}. \quad (1)$$

Define the *odd* of a binary random variable Z as the ratio of the following two probability

$$\text{odd}(Z) = \frac{\Pr(Z = 1)}{\Pr(Z = 0)}.$$

²See http://wiki.ros.org/rospy_tutorials/Tutorials/Parameters for a more detailed description of the server's purpose and usage.

It is convenient to rewrite the Bayes' Rule (1) with odds below.

$$\begin{aligned}
\frac{\Pr(X_{ij} = 1 | Y_{ij})}{\Pr(X_{ij} = 0 | Y_{ij})} &= \frac{\Pr(Y_{ij} | X_{ij} = 1)\Pr(X_{ij} = 1)/\Pr(Y_{ij})}{\Pr(Y_{ij} | X_{ij} = 0)\Pr(X_{ij} = 0)/\Pr(Y_{ij})}, \\
\frac{\Pr(X_{ij} = 1 | Y_{ij})}{\Pr(X_{ij} = 0 | Y_{ij})} &= \frac{\Pr(Y_{ij} | X_{ij} = 1)\Pr(X_{ij} = 1)}{\Pr(Y_{ij} | X_{ij} = 0)\Pr(X_{ij} = 0)}, \\
\log \frac{\Pr(X_{ij} = 1 | Y_{ij})}{\Pr(X_{ij} = 0 | Y_{ij})} &= \log \frac{\Pr(Y_{ij} | X_{ij} = 1)}{\Pr(Y_{ij} | X_{ij} = 0)} + \log \frac{\Pr(X_{ij} = 1)}{\Pr(X_{ij} = 0)}, \\
\text{logodd}(X_{ij} | Y_{ij}) &= \log \frac{\Pr(Y_{ij} | X_{ij} = 1)}{\Pr(Y_{ij} | X_{ij} = 0)} + \text{logodd}(X_{ij}).
\end{aligned} \tag{2}$$

Without loss of generality, rename the following log odd ratios

$$\Delta_{occ} := \log \frac{\Pr(Y_{ij} = 1 | X_{ij} = 1)}{\Pr(Y_{ij} = 1 | X_{ij} = 0)}, \quad \Delta_{free} := \log \frac{\Pr(Y_{ij} = 0 | X_{ij} = 1)}{\Pr(Y_{ij} = 0 | X_{ij} = 0)}. \tag{3}$$

Substituting (3) into (2), we have

$$\begin{aligned}
\text{logodd}(X_{ij} | Y_{ij} = 1) &= \Delta_{occ} + \text{logodd}(X_{ij}), \\
\text{logodd}(X_{ij} | Y_{ij} = 0) &= \Delta_{free} + \text{logodd}(X_{ij}),
\end{aligned} \tag{4}$$

where $\text{logodd}(X_{ij})$ is the existing log odd ratio at cell (i, j) , $\text{logodd}(X_{ij} | Y_{ij} = 1)$ the updated log odd ratio after observing the cell (i, j) being occupied, and $\text{logodd}(X_{ij} | Y_{ij} = 0)$ the updated log odd ratio after observing it being unoccupied.

This may seem like an unnecessary mathematical complication, but it's actually very useful: conversion from probability to log-odds transforms the range of possible values from $[0, 1]$, which is bounded and centered around 0.5, to the range $(-\infty, \infty)$, which is unbounded and centered around 0, making things easier for analysis. The symmetry of log-odds around 0 is illustrated in the following table:

Probability	Odds	Log-Odds
0.1	0.111	-2.197
0.2	0.250	-1.386
0.3	0.428	-0.847
0.4	0.667	-0.405
0.5	1.000	0.000
0.6	1.500	0.405
0.7	2.333	0.847
0.8	4.000	1.386
0.9	9.000	2.197

When a scan ray terminates at a particular cell, that cell's log-odds ratio is incremented by some small amount — i.e., $\text{logodd}(X_{ij} | Y_{ij} = 1) = \Delta_{occ} + \text{logodd}(X_{ij})$ — and then upper bounded by a maximum threshold to ensure numerical stability. Likewise, when the ray passes through a cell (and does not terminate there), that cell's log-odds ratio is decremented by some other amount — i.e., $\text{logodd}(X_{ij} | Y_{ij} = 0) = \Delta_{free} + \text{logodd}(X_{ij})$, where by convention Δ_{free} is negative — and similarly lower bounded by a minimum threshold.

For a video introduction of occupancy grid map and the log-odd update rule, please refer to the following tutorial videos by UPenn: <https://tinyurl.com/ogintroduction> and <https://tinyurl.com/oglogodd>. The videos explain the rationale behind log-odd update but does not cover the threshold saturation technique.

2.2 Using the laser scan

Before starting any edits, read through the inline comments and try to understand what the function is doing at each step. This callback function receives a `sensor_msgs/LaserScan` message, which represents a single line depth scan around the robot (as would be generated by a LIDAR). The scan begins at some angle, gathers range information at a certain angular increment, and ends at some second angle. Use `rosmmsg show` or the online ROS documentation to see the contents of this message.

The callback function iterates through each ray of the scan using the `enumerate` function (look up the documentation for this function if you don't understand what it's doing). The first thing you'll be implementing is finding the




Figure 1: Example of occupancy grid

angle of the ray in the *fixed frame*. A quick look at `demo.launch` shows that the fixed frame is called `odom`, while the sensor frame is `base_footprint`. The frame `odom`, which stands for odometry³, is coincident with `base_footprint` when the robot is turned on. As the robot moves, `odom` moves in the opposite direction relative to `base_footprint`. Thus, if you set your fixed frame in RViz as `odom`, you'll see `base_footprint` moving as the turtlebot moves in the real world. Note: if you move the turtlebot manually (say by picking it up), the odometry won't be able to detect it and the `odom` frame will be wrong. If you do this, restart the bringup sequence on your turtlebot to reset the `odom` frame.

The next thing you'll be doing is "walking" backwards along the ray from the scan point to the sensor, updating the log-odds in each voxel the ray passes through. The `numpy.arange` function can be helpful in defining your loop. The function `PointToVoxel`, defined below `SensorCallback`, may be useful as well. If a voxel is occupied, you should increase the log odds at that voxel by your occupied update value, thresholding it at your occupied threshold value. If a voxel is free, you should increase the log odds at that voxel by your free update value, thresholding it at your free threshold value. Remember that you should only be updating each voxel once per ray.

When you're done, try running the launch file again, and make sure you don't get any error messages.

3 Testing your occupancy grid

Look back at the launch file again. You'll notice that the node's main source file is `mapping_node.py`, not `occupancy_grid_2d.py`. (Although this project is small by most standards, it is generally good practice to separate the actual executable node file from other files implementing different classes that your node uses.) Examine how the `mapping_node.py` file creates an occupancy grid, initializes it, and on success just idles. If you trace that initialization call into the `OccupancyGrid2d` class, you'll see that initialization loads all parameters, registers publishers and subscribers, and sets up any other class variables. If any of that fails, it returns `False`, which causes the whole node to crash. This is a very safe way to build your system because it minimizes the chance that your code crashes mid-operation. We strongly encourage you to use this sort of architecture in your projects.

Bringup the TurtleBot with a minimal launch, and then run the launch file from the `mapping` package. Next, we want to publish a transform between the base frame and the laser scan at the base so that RVIZ can visualize our laser scan. We will do this by running:

```
roslaunch tf2_ros static_transform_publisher 0.0 0.0 0.0 0.0 0.0 0.0 base_footprint base_scan
```

Now open RViz. Find and visualize the topic on which the occupancy grid is being published. You'll also need to change the fixed frame to `odom` and add `tf` to the display so you can see where the TurtleBot is. Add the `/scan/Laserscan` topic to the display to show the laser scanner's output in real time.

³Odometry is the use of data from motion sensors to estimate change in position over time. It is used in robotics by some legged or wheeled robots to estimate their position relative to a starting location.

Drive the TurtleBot around with the `turtlebot_teleop` node we used in Lab 4. You should see the floorplan begin to emerge as you drive around. Do you notice any systematic errors? Where are they coming from, and how would you address them?

Next, experiment with changing some of the parameters defined in your parameter server. While you can simply change the values in your launch file, it's cleanest (and most convenient) to set them via the command line so you can experiment with many different values without changing the defaults. (Hint: You've actually done this before using Baxter/Sawyer — `electric_gripper` is a parameter value!)

Experiment with changing the downsampling rate parameter. What is the downsampling rate's function, and why is it important? (The comments in the `occupancy_grid_2d.py` file might be helpful here.)

Lastly, experiment with changing the resolution of the map. (Note that the length of each cell isn't explicitly defined in the parameter server but can be calculated from the values there; which parameters do you need to modify to make the cells larger and smaller?) How does your map behave differently? Do you notice any change in error patterns?

Checkpoint 2

Submit a checkoff request at <https://tinyurl.com/fa22-106alab> for a staff member to come and check off your work. At this point you should be able to:

- Demonstrate the odometry-based localization and the associated map
- Describe any shortcomings you notice in the result, and hypothesize why they exist
- Explain any bottlenecks in the code — what's the slowest part of the computation?
- Change a parameter of the launch file from the command line