

CSE 590A - Robotics

Homework 1 - EKF and Particle Filter for Localization

Due Sunday April 23rd @ 11:59pm

The key goal of this homework is to get an understanding of the properties of Kalman filters and Particle filters for state estimation. There are two parts to the homework - a written assignment and a programming assignment. There are two written questions which help you derive some useful quantities for the EKF implementation. In the programming assignment, you will implement an Extended Kalman Filter (EKF), a Particle Filter (PF) for landmark based localization, and a learned observation model from camera images. You will also analyze their performance under various conditions. The zip file containing the starter code for this homework can be found at <https://courses.cs.washington.edu/courses/csep590a/23sp/homeworks/assignment1.zip>.

Useful reading material: Lecture notes, Chapters 3,4,5,7 & 8 of Probabilistic Robotics, Thrun, Burgard and Fox (pdf shared with class).

Collaboration: Students can discuss questions, but each student MUST write up their own solution, and code their own solution. We will be checking code/PDFs for plagiarism.

Late Policy: This assignment may be handed in up to 5 days late. If you have used up your 6 late days this quarter, there will be a penalty of 10% of the maximum grade per day.

1 Writing assignments




Figure 1: Odometry-based motion model

1.1 Motion Model Jacobian [5 points]

Figure 1 describes a simple motion model. The state of the robot is its 2D position and orientation: $s = [x, y, \theta]$, where x and y are its location on the 2D ground plane, and θ is the angular direction the robot is facing. The

control to the robot is $u = [\delta_{rot1}, \delta_{trans}, \delta_{rot2}]$. This means that the robot controls itself at each time step by first rotating by δ_{rot1} , then driving forward some distance δ_{trans} , then rotating again by δ_{rot2} .

Given the current state $s_t = [x_t, y_t, \theta_t]$ and controls $u = [\delta_{rot1}, \delta_{trans}, \delta_{rot2}]$, we can use this motion model to describe an estimate for the next state $s_{t+1} = [x_{t+1}, y_{t+1}, \theta_{t+1}]$ using the following equations:

$$\begin{aligned} x_{t+1} &= x_t + \delta_{trans} * \cos(\theta_t + \delta_{rot1}) \\ y_{t+1} &= y_t + \delta_{trans} * \sin(\theta_t + \delta_{rot1}) \\ \theta_{t+1} &= \theta_t + \delta_{rot1} + \delta_{rot2} \end{aligned}$$

We refer to these functions as the motion model g such that $g(u, s_t) = s_{t+1}$. Derive the Jacobians of g with respect to the state $G = \frac{\partial g}{\partial s_t}$ and control $V = \frac{\partial g}{\partial u}$:

$$G = \begin{pmatrix} \frac{\partial x_{t+1}}{\partial x_t} & \frac{\partial x_{t+1}}{\partial y_t} & \frac{\partial x_{t+1}}{\partial \theta_t} \\ \frac{\partial y_{t+1}}{\partial x_t} & \frac{\partial y_{t+1}}{\partial y_t} & \frac{\partial y_{t+1}}{\partial \theta_t} \\ \frac{\partial \theta_{t+1}}{\partial x_t} & \frac{\partial \theta_{t+1}}{\partial y_t} & \frac{\partial \theta_{t+1}}{\partial \theta_t} \end{pmatrix} \quad V = \begin{pmatrix} \frac{\partial x_{t+1}}{\partial \delta_{rot1}} & \frac{\partial x_{t+1}}{\partial \delta_{trans}} & \frac{\partial x_{t+1}}{\partial \delta_{rot2}} \\ \frac{\partial y_{t+1}}{\partial \delta_{rot1}} & \frac{\partial y_{t+1}}{\partial \delta_{trans}} & \frac{\partial y_{t+1}}{\partial \delta_{rot2}} \\ \frac{\partial \theta_{t+1}}{\partial \delta_{rot1}} & \frac{\partial \theta_{t+1}}{\partial \delta_{trans}} & \frac{\partial \theta_{t+1}}{\partial \delta_{rot2}} \end{pmatrix}$$

1.2 Observation Model Jacobian [5 points]

Assume there is a landmark l at position (l_x, l_y) .

1. Suppose the robot receives measurement of the landmark in terms of the bearing angle ϕ relative to the direction it is currently facing. Write down the estimation for ϕ in terms of the landmark's position (l_x, l_y) and the robot's state (x_t, y_t, θ_t) .
2. Derive Jacobian of the measurement with respect to the robot state:

$$H = \begin{bmatrix} \frac{\partial \phi}{\partial x_t} & \frac{\partial \phi}{\partial y_t} & \frac{\partial \phi}{\partial \theta_t} \end{bmatrix}$$

2 Programming assignments

Environment Requirement

We strongly encourage you to use a Linux machine for coding assignments. For training Neural Networks, we have provided colab notebook. Please see detailed instructions below.

(a) Linux / Mac

Virtual environment is strongly recommended! To set up a virtual environment, first install conda from <https://www.anaconda.com/>, then run:

```
conda create -n cse590a python=3.9
conda activate cse590a
pip install numpy
pip install pybullet
```

```
pip install matplotlib
pip install torch
```

(b) Windows

If you have a windows machine and would like to use Ubuntu, please go to: <https://www.vmware.com/products/workstation-player.html> and download VMware Workstation Player for Windows. You will also need to download both the [.vmdk](#) and the [.ovf](#) files. However, VM is not mandatory.

To set up the VM on your Windows machine, click the .exe file and follow installation steps. Once the VMware player is installed, open VMWare Workstation 17 Player, and choose "Open a Virtual Machine" and select "ros-noetic-22wi.ovf" and import. You should be good to go! Once you set up your VM, log in "robotics" with the password: "robotics".

Inside your VM machine, open a terminal, and please follow (a) Linux/Mac to set up the virtual environment on your VM machine.

(c) CoLab

CoLab is a free Jupyter notebook environment that runs entirely in the cloud hosted by Google. Please go to <https://colab.research.google.com> and click "Upload". Select your ".ipynb" jupyter notebook and you are all set!

Problem Overview

In the programming component of this assignment, you will implement an Extended Kalman Filter (EKF) and Particle Filter (PF) for localizing a robot based on landmarks. Then you will also implement a learned observation model which estimates landmark bearing angles from images. In our environment, the robot is driving in a path that would cause it to follow a rectangle. However, because it's motion is noisy, it's true path will deviate substantially from the intended path. The goal is to use the EKF and PF to find out where the robot actually is, given information about how the robot is moving and the landmark observations.

We will use the odometry-based motion model you derived in question 1.2. We assume that there are landmarks present in the robot's environment. The robot receives the bearings (angles) to the landmarks and the ID of the landmarks as observations: (bearing, landmark ID).

We assume a noise model for the odometry motion model with parameters α as shown below([1] Table 5.6), and a separate noise model for the bearing observations with parameter β (see [1], Section 6.6).

At each timestep, the robot starts from the current state and moves according to the control input plus noise. The robot then receives a noisy landmark observation from the world. You will use this information to localize the robot over the whole time sequence with an EKF and PF.

Code Overview

The starter code is written in Python and depends on matplotlib, numpy, pybullet and torch. This section gives a brief overview of each file.

- `localization.py` – This is your main entry point for running experiments.
- `train_localization_model.ipynb` – A jupyter notebook that will train your observation model once you have implemented it.

1:	Algorithm <code>sample_motion_model_odometry</code>(u_t, x_{t-1}):
2:	$\delta_{\text{rot1}} = \text{atan2}(\bar{y}' - \bar{y}, \bar{x}' - \bar{x}) - \bar{\theta}$
3:	$\delta_{\text{trans}} = \sqrt{(\bar{x} - \bar{x}')^2 + (\bar{y} - \bar{y}')^2}$
4:	$\delta_{\text{rot2}} = \bar{\theta}' - \bar{\theta} - \delta_{\text{rot1}}$
5:	$\hat{\delta}_{\text{rot1}} = \delta_{\text{rot1}} - \text{sample}(\alpha_1 \delta_{\text{rot1}} + \alpha_2 \delta_{\text{trans}})$
6:	$\hat{\delta}_{\text{trans}} = \delta_{\text{trans}} - \text{sample}(\alpha_3 \delta_{\text{trans}} + \alpha_4 (\delta_{\text{rot1}} + \delta_{\text{rot2}}))$
7:	$\hat{\delta}_{\text{rot2}} = \delta_{\text{rot2}} - \text{sample}(\alpha_1 \delta_{\text{rot2}} + \alpha_2 \delta_{\text{trans}})$
8:	$x' = x + \hat{\delta}_{\text{trans}} \cos(\theta + \hat{\delta}_{\text{rot1}})$
9:	$y' = y + \hat{\delta}_{\text{trans}} \sin(\theta + \hat{\delta}_{\text{rot1}})$
10:	$\theta' = \theta + \hat{\delta}_{\text{rot1}} + \hat{\delta}_{\text{rot2}}$
11:	return $x_t = (x', y', \theta')^T$

Table 5.6 Algorithm for sampling from $p(x_t | u_t, x_{t-1})$ based on odometry information. Here the pose at time t is represented by $x_{t-1} = (x \ y \ \theta)^T$. The control is a differentiable set of two pose estimates obtained by the robot's odometer, $u_t = (\bar{x}_{t-1} \ \bar{x}_t)^T$, with $x_{t-1} = (\bar{x} \ \bar{y} \ \bar{\theta})$ and $x_t = (\bar{x}' \ \bar{y}' \ \bar{\theta}')$.

Figure 2: Table 5.6 from [1]. The **sample** function samples a single random value with mean 0 and variance equal to the function argument.




Figure 3: The test environment in pybullet. The blue line traces out the robot's true position, which is a result of noisy actions. The green line traces the robot's position assuming that actions weren't noisy. The black line indicates the bearing angle to the current landmark.

- `soccer_field.py` – This implements the dynamics and observation functions, as well as the noise models for both. It also contains the interface to display the robot in pybullet. Add your Jacobian implementations here!
- `utils.py` – This contains a useful function for normalizing an angle between $[-\pi, \pi]$.
- `policies.py` – This contains a simple policy, which you can safely ignore.
- `ekf.py` – Add your extended Kalman filter implementation here!
- `pf.py` – Add your particle filter implementation here!
- `observation_model.py` – Add your observation model implementation here!
- `racecar.urdf` – The robot description of the racecar.
- `meshes/` – Meshes used by the racecar.

- `make_data.py` – A script which generates a dataset to use for training your learned observation model.
- `requirements.txt` – A list of pip packages necessary to run the assignment code.

Command-Line Interface To visualize the robot in the soccer field environment, run

```
$ python localization.py --plot none
```

The blue line traces out the robot’s position, which is a result of noisy actions. The green line traces the robot’s position assuming that actions weren’t noisy. The black line indicates the bearing angle to the current landmark. After you implement a filter, the filter’s estimate of the robot’s position will be drawn in red.

To check the implementation of ekf and pf, run

```
$ python localization.py --plot ekf
$ python localization.py --plot pf
```

To see other command-line flags available to you, run

```
$ python localization.py -h
```

Data Format

- state: $[x, y, \theta]$
- control: $[\delta_{rot1}, \delta_{trans}, \delta_{rot2}]$
- observation: $[\theta_{bearing}]$

Hints

- Make sure to call `utils.minimized_angle` any time an angle or angle difference could exceed $[-\pi, \pi]$.
- Make sure to use the low-variance systematic sampler from lecture. It gives you a smoother particle distribution and also requires only a single random number per resampling step.
- Switch `pybullet` from GUI to DIRECT (headless) for a significant speedup. This can be done by omitting the `--plot` flag when running `localization.py`

2.1 EKF Implementation [30 points]

Implement the extended Kalman filter algorithm in `ekf.py`. You will need to fill in `ExtendedKalmanFilter.update` and the `Field` methods `G`, `V`, and `H`. Your results from a successful EKF implementation should be comparable to the following results.

```
$ python localization.py ekf --seed 1
...
Mean position error: 5.4683282563859095
Mean Mahalanobis error: 2.28460917211501
ANEES: 0.7615363907050033
```

Note that these numbers are included for reference and you are not expected to exactly replicate these numbers. Depending on your implementation, you may get different results than this.

- (a) Plot the real robot path and the filter path under the default (provided) parameters. For this you can use a screen capture of the pybullet environment.
- (b) Plot the position error and ANEES as the noise factors for both the data and the filter range over $r = [1/64, 1/16, 1/4, 4, 16, 64]$. You should run 10 trials per value of r with different random seeds and plot the mean (optionally standard deviation). One run might look something like:

```
$ python localization.py ekf --data-factor 4 --filter-factor 4
```

Note that the value for `--data-factor` should be the same as `--filter-factor`. Discuss anything interesting you observe.

- (c) Plot the position error and ANEES (average normalized estimation error squared) as the filter noise factors α, ϕ vary over r (as above) while the data noise factor is kept as default. You should run 10 trials per value of r and plot the mean (optionally standard deviation). Note that you should not pass the `--data-factor` argument here. How does the results compare with what you get using the default parameters? Discuss anything interesting you observe.
- (d) Discuss any challenges you faced and describe your debugging process.

2.2 PF Implementation [30 points]

Implement the particle filter algorithm in `pf.py`. You will need to fill in `ParticleFilter.update` and `ParticleFilter.resample`.

```
python localization.py pf --seed 1
...
Mean position error: 6.346563990607579
Mean Mahalanobis error: 6.251697842587203
ANEES: 2.083899280862401
```

Note that these numbers are included for reference and you are not expected to exactly replicate these numbers. Depending on your implementation, you may get different results than this.

- (a) Plot the real robot path and the filter path under the default parameters.
- (b) Plot the position error and ANEES as both the data and the filter noise factors range over $r = [1/64, 1/16, 1/4, 4, 16, 64]$. You should run 10 trials per value of r and plot the mean (optionally standard deviation). Discuss anything interesting you observe.
- (c) Plot the position error and ANEES as the filter noise factors α, ϕ vary over r (as above) while the data noise factor is kept as default. You should run 10 trials per value of r and plot the mean (optionally standard deviation). How does the results compare with what you get using the default parameters? Discuss anything interesting you observe.
- (d) Plot the mean position error and ANEES as α, ϕ range over r (as in (c)) and the number of particles varies over $[20, 50, 500]$. You should run 10 trials per value of r and plot the mean (optionally standard deviation). You should include three plots, one per each particle quantity. Discuss anything interesting you observe.
- (e) Discuss any challenges you faced and describe your debugging process.

2.3 Learned Observation Model [30 points]

In this section you will learn a model which takes panoramic images of the scene around the racecar, and predicts the bearing angles to each landmark. The racecar has four evenly spaced cameras that have been stitched together into a single panoramic image that is 32 pixels tall and 128 pixels wide that look like Figure 4:

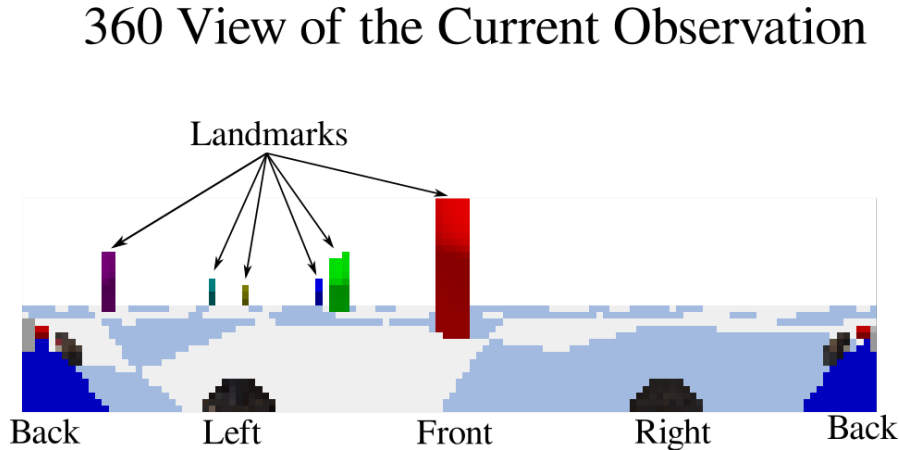


Figure 4: An illustration on the observation input of the car

- (a) **Generate Datasets:** The first step in any learning problem is to acquire your dataset. We have provided the script `make_data.py` to generate data for you. This will randomly place the car in the scene, capture a panoramic image, and record the ground-truth bearing angle to each landmark. You can generate the train and test dataset using the following code:

```
$ python make_data.py --split train --seed 1234 --size 10000
$ python make_data.py --split test --seed 12345 --size 1000
```

This will generate the folders `hw1_train_dataset` and `hw1_test_dataset` containing `rgb_XXXXXX.png` and `label_XXXXXX.npy`. The rgb images in each should look like 4. Zip these folders as `hw1_train_dataset.zip` and `hw1_test_dataset.zip`.

- (b) **Implement Neural Network:** Next you will build a neural network to predict landmark heading directions from visual observations. To do this, it will be easiest to work in the provided colab notebook. Go to colab and upload `train_observation_model.ipynb`. Then click the folder icon next to the notebook shown in Figure 5 to open the file browser of your colab runtime. From here, upload `observation_model.py`, `hw1_train_dataset.zip` and `hw1_test_dataset.zip`.

Implement your neural network initialization and forward functions in the uploaded `observation_model.py`. Your network should take a single `output_channels` argument that allows you to change how many values the network predicts (we will use this to change the prediction format in the next section). The forward function should take a batch of images in as a `[batch_size x 3 x 32 x 128]` tensor and output a `batch_size x output_channels` tensor. We recommend starting with a simple convolution network, this assignment should not require a complicated model.

Hint: Try using a network similar to the one shown in Tutorial 2. Because the image sizes are different here, you may need a different number of Convolution and MaxPool layers than in that example.

Feel free to use external sources for network ideas, but please mention which ones you used in your writeup.

Please write a short description of the networks you tried, what worked, and what didn't.

WARNING: Make sure you download your `observation_model.py` from colab frequently, and include it in your writeup when you are finished! If your runtime disconnects, you may lose the changes you made to this file!

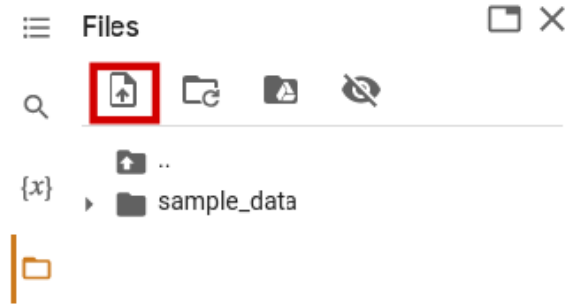


Figure 5: Click the file upload icon to upload your files.

(c) **Train Models:** In this section you will train two different versions of your model, one that predicts the bearing angles ϕ directly, and the other which predicts $\cos(\phi)$ and $\sin(\phi)$ instead.

- (i) At the top of the notebook, a variable `supervision_mode` is defined. Set this to "phi" and run all the cells in the training notebook. This will train your network to predict the bearing angles of each landmark directly (6 channels, one for each landmark). After each epoch of training, the script will generate a set of debugging images that look like Figure 6. The first row of lines under the image shows the ground truth heading of the landmarks, while the second row shows the network's current predictions.




Figure 6: Debug images produced by the evaluation.

Report the error mean and std of your model.

- (ii) Change `supervision_mode` to "xy". This will train your network to predict $\cos(\phi)$ and $\sin(\phi)$ instead of the bearing angles directly (12 channels, two for each landmark). Report the mean and std of your model.
- (iii) Did one of these work better than the other? Why do you think one might be superior?

For reference, you should be able to get a model with less than 0.07 mean error for at least one of the supervision modes. Figure 7 shows the loss curve for a fairly simple convolutional neural network that achieves less than 0.05 mean error after 20 epochs:

Error Mean: 0.041499
 Error Std: 0.066491

(d) **System Integration:** The training runs above should generate a `checkpoint.pt` file, which stores the weights of your network. Download your `observation_model.py` and the `checkpoint.pt` that corresponds to the best `supervision_mode` back to the assignment directory. Run the following, replacing `MODE` with the best `supervision_mode` you found:

```
python localization.py pf --use-learned-observation-model checkpoint.pt --supervision-mode MODE
```





Figure 7: Loss curve for a simple convolutional network.

Report the position error output in the writeup across five different random seeds.

- (e) Discuss any challenges you faced and describe your debugging process.
- (f) **Extra Credit: Analysis [10 points]:** Train your network again, with (1) an increased training dataset of 50000 examples (2) make the network deeper and/or wider (3) adjust training parameters such as learning rate (4) more epochs. What modifications are most helpful? Summarize your observation in the writeup.

3 Submission

You will be using Gradescope <https://www.gradescope.com/> to submit the homework. Please submit the written assignment answers as a PDF. For the code, submit a zip file of the entire working directory.

References

- [1] Thrun, Burgard and Fox (2005), Probabilistic Robotics.