CS 4758/6758: Robot Learning: Homework 2

Ashutosh Saxena. Due: Feb 24 (Tuesday), 5pm

1 Cost Functions

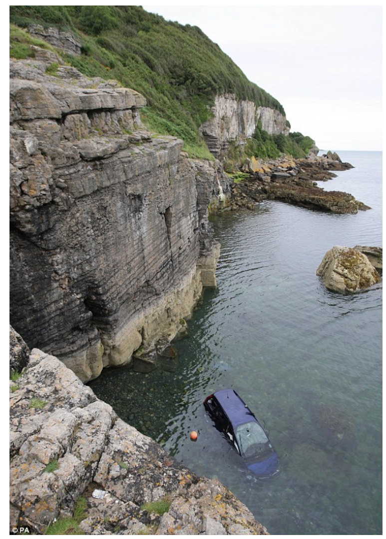

You are trying to train an autonomous vehicle to navigate mountain roads. One particularly treacherous stretch winds along the side of a cliff, with its face to the left and a steep drop-off to the right. Denoting the lateral position of the car by x, design a cost function $J(x) \geq 0$ describing how (sub)optimal its course is. Draw a graph of J and explain your choices.[1]

Some examples of the properties that your cost function should follow are: $\arg\min_x J(x)$ should give the most desired position of the car on the road. It also be not be discontinous, otherwise it becomes harder to run algorithms such as gradient descent. Does your optimization problem $\arg\min_x J(x)$ has local minima?

Draw or define $J_2(x)$ for a road which has a divider in between. Does $\arg\min_x J_2(x)$ have local minima? (10 points, plus bonus)

2 Sensor Noise

A robot (e.g., the ground robot or aerial robot) is equipped with three types of ranging sensors: laser, ultrasound, and infrared. To test the sensors, you place them in front of obstacles at known distances and measure their output. The file *sensors.train* contains the resulting data.

Call the vector of true distances y, and the matrix of sensor readings X (each row corresponds to a single reading, each column to a sensor). You want to figure out how to best combine the sensor data to accurately estimate y. You could compute a weighted average: using a vector of weights β, $\hat{y} = X\theta$ gives the estimated position. But what weights to use?

A) Find θ that minimizes $\|y - \hat{y}\|_2^2$. [Note: $\|v\|_2^2 = v^T v$] Report it.

[1] You need not give a formal mathematical description of J to get full credit.

Now lets take a different approach. Suppose for each sensor i ($i = \{1, 2, 3\}$), its readings x_i are related to the correct distance by $x_i = y + \epsilon_i$. Then ϵ_i represents the sensor's error. Plotting histograms of ϵ_i yields:

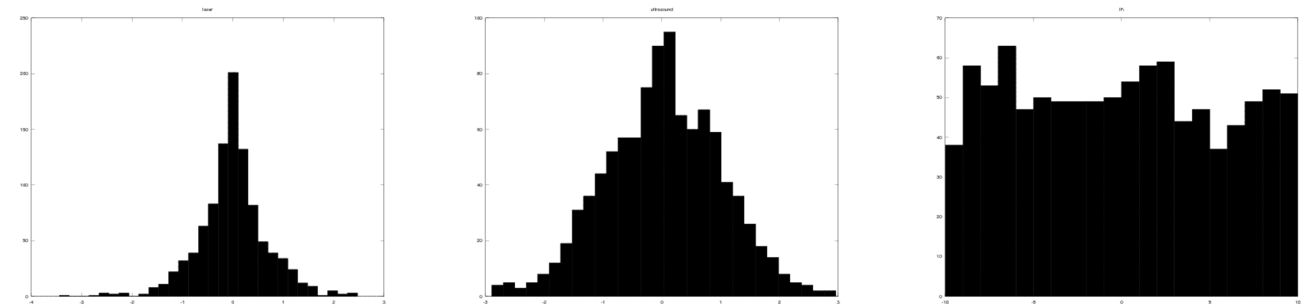

B) For each error signal, decide whether it is best modelled by a Gaussian, Laplacian or uniform distribution. For Laplace and Gaussian distributions, compute maximum likelihood estimates of their parameters $((\mu_g, \sigma)$ or $(\mu_l, b))$ from the data in *sensors.train*, and report them.

C) Assuming the probability distributions of **B)** to be correct, and also assuming that each sensor makes independent errors. You take a new triplet of measurements–one from each sensor (x_1, x_2, x_3). Write an expression for the likelihood of a position y; that is $\Pr(y|x_1, x_2, x_3; \mu_g, \mu_l, \sigma, b)$.

D) Assign y its most likely value. That is, let the estimated position $\hat{y} = \arg\max_y \Pr(y|x_1, x_2, x_3; \mu_g, \mu_l, \sigma, b)$. If possible, give a closed form expression for $\hat{y}$. Otherwise, please specify how the maximization problem could be solved. Intuitively, how does this filter work? Is it linear?

E) Which algorithm is better? The first algorithm (linear regression) or the second (statistical modeling)? To decide, measure their performance on a test set. The file *sensors.test* contains different sample of identically distributed data. Report the RMS error of each algorithm. [Note : RMS error or root mean square error is $\frac{\|\hat{Y}-Y\|_2}{\sqrt{n}}$ where n is the number of observations.] Explain the results.

F) For this part assume that all $P(x_i|y)$ are Gaussian distributions. Argue (or use a result derived in the class) that $\hat{y}$ is a linear function of x_i's.

G) For this part assume all $P(x_i|y)$ is are Laplacian distributions. Prove that $\hat{y}$ is a median filter on values x_i's.

(50 points)

3 PID Control (10 pts)

3.1 Tuning

Each of the following step-responses was generated by a badly-tuned PID controller. For each, suggest a change to ***one*** of the coefficents that will produce the 'ideal' response shown. We are looking for answers like 'decrease the proportional term.' Explain your answers.

Fix each these controllers (by changing only one parameter each)

Figure 1: The desired response.

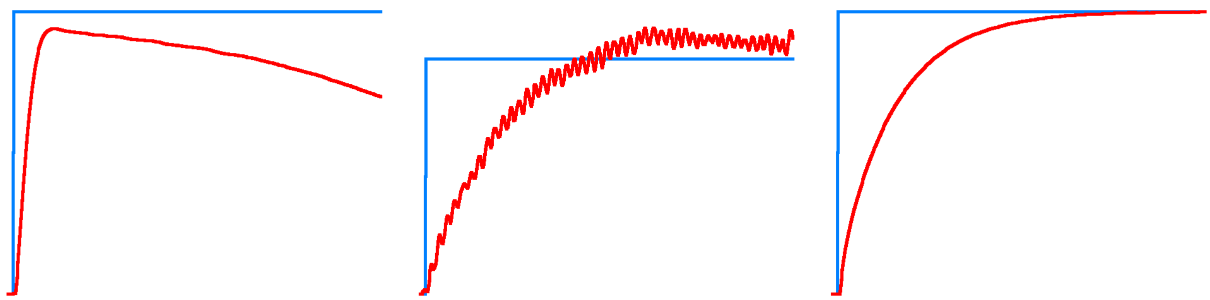

3.2 Stability

Given the one-dimensional, discrete-time system:

$$x_{t+1} = ax_t + u_t$$

and the controller:

$$u_t = -k_P(x_t - x_{set})$$

If $x_{set} = 0$, for what range of k_P is the controlled system stable? I.e., specify the range of k_P (only as a function of a) for which $x_{t\to\infty} \to x_{set}$ in the controlled system.

4 Perception for your Robot (30 points)

In this question, *you* have to take an image from *your* robot in the project, and you have to get openCV working on *your* computer.[2]

You are free to take any image(s) relevant to your project. You would be graded on the relevancy of the image to your project, and creativity in this question.

Demonstrate:

1. Results of applying a mean and median filter on the images you collect.

2. Results of applying edge detection filters on the images you collect.

3. Localize an object by doing either Option A or Option B. It should work on atleast three distinct RGB/RGBD images you took.

Option A: Given an image, goal is to find a colored object (of known color, e.g., a pingpong ball). Build an algorithm that outputs the (x,y) location of the object in the image. You can use color segmentation, edge detection, image filters, etc. Describe in a brief bulleted list what approach you applied. The approach should be reasonably robust to multiple images taken of the same object in different locations in a normal environment. Bonus points for more robustness.

Option B: Given the depthmap from Kinect sensor (or the point-cloud), find the location ((x, y, d) or (X, Y, Z)) of an object lying on the top of a flat table. It should work on *any* generic object lying on a flat surface. You can use depth segmentation, or other approaches. While you can assume known height of the flat surface, bonus points for using robust approaches such as Ransac to detect flat surfaces.

Please give your code printout, description of your approach, and image results to get credit.

[2]You'd get no credit if these rules are not followed. Using images collected by someone else, even your project partner, would be a violation of academic integrity.