This assignment consists of two questions.

Please use Python for implementation using only standard python libraries (e.g., numpy, matplotlib etc.). Please do not use any third party libaries/implementations for algorithms.

The revised submission date is 5pm on March 24, 2021. No late submissions.

This assignment will carry 20% of the total grade. Points in this assignment will be awarded out of 100.

This assignment is to be done individually or in pairs. Code implementation must be from your own original efforts. Do not use an existing or previous implementations done by others. Do not violate the honor code (refer to the course policy on honor code violations discussed in class).

Please submit your implementation and the accompanying report in a single .zip file in the format {EntryNumber}.zip on Moodle. The report should contain responses to questions and the desired plots/graphs. Briefly describe the key findings/insights for the graphs. Ensure the reproduction of graphs (modulo probabilistic execution) for your submission.

Additional readings: Artificial Intelligence: A Modern Approach (Ch. 15) and Probabilistic Robotics (Ch. 2 and Ch. 3).

Teaching Assistant Incharge: Mr. Vikas Upadhyay (vikas.upadhyay@cse.iitd.ac.in).

1. (40 points) Consider the problem of localizing an aerial vehicle that is flying at a constant height while surveying an area. Please see figure below. Assume a uniform discretization for the region with grid cells of size $1m \times 1m$. At each time step, the robot can execute one of four actions: *up*, *down*, *left* or *right* to an adjacent grid cell, selected with likelihoods 0.4, 0.1, 0.2 and 0.3 respectively. The area is instrumented with four sensors positioned at the indicated grid cells (see figure). Each sensor reports the discrete *presence* or *absence* of a target without reporting any other information such as the target's position in the grid. Sensor observations are stochastic. The likelihood of reporting the presence of a target is shown in the figure (right). Sensor observations are generated independently at each time step and each sensor reports independent of all other sensors. Our goal is to estimate the position of the agent in the grid given the assumptions above.

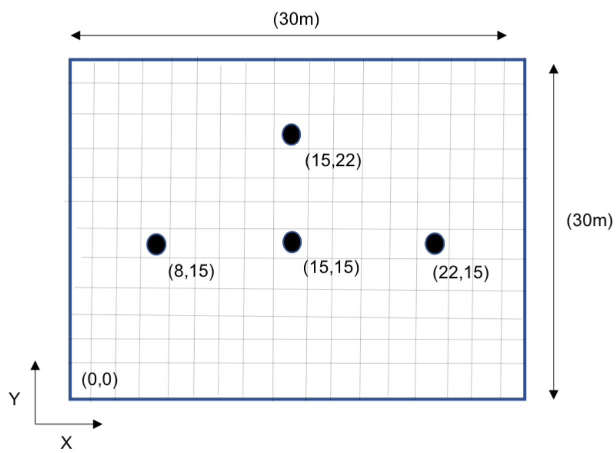


| | | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| 0.5 | 0.5 | 0.5 | 0.5 | 0.5 | 0.5 | 0.5 | 0.5 | 0.5 |
| 0.5 | 0.6 | 0.6 | 0.6 | 0.6 | 0.6 | 0.6 | 0.6 | 0.5 |
| 0.5 | 0.6 | 0.7 | 0.7 | 0.7 | 0.7 | 0.7 | 0.6 | 0.5 |
| 0.5 | 0.6 | 0.7 | 0.8 | 0.8 | 0.8 | 0.7 | 0.6 | 0.5 |
| 0.5 | 0.6 | 0.7 | 0.8 | **0.9** | 0.8 | 0.7 | 0.6 | 0.5 |
| 0.5 | 0.6 | 0.7 | 0.8 | 0.8 | 0.8 | 0.7 | 0.6 | 0.5 |
| 0.5 | 0.6 | 0.7 | 0.7 | 0.7 | 0.7 | 0.7 | 0.6 | 0.5 |
| 0.5 | 0.6 | 0.6 | 0.6 | 0.6 | 0.6 | 0.6 | 0.6 | 0.5 |
| 0.5 | 0.5 | 0.5 | 0.5 | 0.5 | 0.5 | 0.5 | 0.5 | 0.5 |

Figure 1: *Left*: The robot is moving in a $30m \times 30m$ area. The area is instrumented with noisy sensors that report the presence of absence of a target. The sensors are located at grid cells: $\{(8, 15), (15, 15), (22, 15), (15, 22)\}$ where $(i, j)$ indicates the index along $x$ and $y$ axes. The figure (including the the overlaid grid cells) is indicative and not to scale. *Right*: Sensors at each time step report the presence or absence of the target. The spatial likelihood of detecting the robot centered on the sensor origin (cell with dark borders) is shown. The likelihood reported in cell $(i, j)$ indicates the probability that the sensor will report a detection if the robot is present in cell $(i, j)$. The detection likelihood is assumed zero for other cells not shown.

(a) Simulate the robot's motion in the grid from any starting position and record the sequence of sensor observations (given the the assumptions above) for $T = 25$ time steps.

(b) Estimate the robot's current position at time $t$ given the sequence of observations generated by the simulation above till time $T$. Visualize and plot the estimated log-likelihood over the grid locations at each time step. Plot the estimated and ground truth locations at each time step.

(c) Estimate agent's current and past positions given the sequence of observations received till time $T$. Plot the estimated and the ground truth locations at each time step.

(d) Determine the error between actual path and estimated path obtained in parts (b) and (c) above using the Manhattan distance metric. Plot and describe your finding(s).

(e) Compute the *predictive* likelihood over the robot's future location. Plot the likelihood over the next (i) 10 time steps and (ii) 25 time steps. Describe your finding(s).

(f) Estimate the robot's *most-likely* path given the full sequence of observations till time $T$.

2. (60 points) Consider an airplane flying in a $(x, y)$-plane. A noisy sensor (e.g., a radar) provides measurements $z_t = (x'_t, y'_t)$. The plane is controlled by providing velocity increments $u_t = [\delta\dot{x}_t, \delta\dot{y}_t]$, which get added to the velocity components $\dot{x}_t$ and $\dot{y}_t$ at each time step. The uncertainty in the motions is characterized by a presence of Gaussian noise $\epsilon_t \sim \mathcal{N}(0, R)$. Similarly, the measurement uncertainty is characterized by presence of Gaussian noise $\delta_t \sim \mathcal{N}(0, Q)$. Our goal is to estimate its positions $[x_t, y_t]$ and velocities $[\dot{x}_t, \dot{y}_t]$ at time instant $t$ from noisy observations $[x'_t, y'_t]$.

   (a) Implement the motion model for this problem. Initially, assume that the control inputs are both zero. Assume the noise parameters as: $\sigma_{rx} = 1.0$, $\sigma_{ry} = 1.0$, $\sigma_{r\dot{x}} = 0.01$ and $\sigma_{r\dot{y}} = 0.01$ forming the covariance matrix $R$ as $diag(\sigma^2_{rx}, \sigma^2_{ry}, \sigma^2_{r\dot{x}}, \sigma^2_{r\dot{y}})$ where $diag$ denotes the diagonal elements of $R$ with off-diagonals as zero. Next, implement the observation model for this problem. Assume that the observation noise is distributed as an isotropic Gaussian with a standard deviation of 10. Simulate the motion and the sensor models for $T = 200$ time steps. Plot the actual trajectory and the observed trajectory of the vehicle.

   (b) Implement a Bayes filter for the problem to estimate the vehicle state given the assumptions above. Please formally write down the the model for estimation. Let $\hat{x}_t$ denote the estimated state at time $t$. We know that the vehicle is expected to start at location $(10, 10)$ with an initial velocity of 1.0 along both $x$ and $y$ axes. Further, our prior belief over the vehicle's initial state has a standard deviation of 0.01 for each state variable.

   (c) Plot the actual trajectory $[x_t, y_t]$, the noisy observations $[x'_t, y'_t]$ and the trajectory estimated by the filter $[\hat{x}_t, \hat{y}_t]$. Additionally, plot the uncertainty ellipses for the estimated trajectory $[\hat{x}_t, \hat{y}_t]$. An uncertainty ellipse denotes the locus of points that are one standard deviation away from the mean.

   (d) Implement a control policy where $\delta\dot{x}_t$ varies as a sine wave and $\delta\dot{y}_t$ varies as a cosine wave. Plot the true, observed and the estimated $[x_t, y_t]$ trajectories under the given control inputs. Compute and plot the error between the true and the estimated trajectory using the euclidean distance metric.

   (e) Increase and decrease the uncertainty in the sensor model in comparison with the uncertainty in the motion model. Plot the estimated trajectory and explain how the noise variation impacts the filter performance.

   (f) Next, model higher uncertainty in the initial belief over the vehicle's position assuming a standard deviation of 100 in both the $x$ and $y$ positions. As before, plot the true, observed and the estimated trajectories under the sine-cosine control policy.

   (g) Assume that the sensor observations drop out at time instants $t = 10$ and $t = 30$ for a period of 10 time steps and are re-acquired after that period. Simulate and show the evolution of uncertainty in the vehicle's $[x_t, y_t]$ position by plotting the uncertainty ellipses.

   (h) Plot the estimated velocities $[\hat{\dot{x}}_t, \hat{\dot{y}}_t]$ and the true velocities of the vehicle $[\dot{x}_t, \dot{y}_t]$. Briefly explain if the estimator can or cannot track the true values.

   (i) Simulate a second vehicle in the environments with linear-Gaussian motion and sensor models (as above) with different initial state estimates and noise characteristics. Let $a$ and $b$ index the two vehicles. The sensor receives two sets of measurements, $z^1_t = [x'^1_t, y'^1_t]$ and $z^2_t = [x'^2_t, y'^2_t]$. Estimate the latent states $x^a_t$ and $x^b_t$ at time $t$ for the two vehicles. The estimator requires a strategy for associating observations $z^1_t$ and $z^2_t$ with the latent states $x^a_t$ and $x^b_t$, known as the *data association* step. Implement and a data association strategy and study its behavior in your simulation. Can your solution scale to more (4-5) agents?