



1 Particle Filtering: Warmup

- (a) **True / False:** The particle filtering algorithm is consistent since it gives correct probabilities as the number of samples N tends to infinity.
- (b) **True / False:** The number of samples we use in the particle filtering algorithm increases from one time step to the next.
- (c) The following state space contains 10 particles. The left grid shows the prior belief distribution of the particles at time t , while the grid on the right shows the particles weighted by the observations $P(e_t|S_t)$.



Fill in the following grids to update the belief distribution. Each square in the “Belief” grid should correspond to $\hat{P}(S_t|e_{1:t-1})$, the estimated probability of a particle being in state S at time t . Each square in the “Unnormalized” grid should correspond to the probability $P(S_t, e_t|e_{1:t-1})$. The “Normalized” grid should contain our updated belief distribution $\hat{P}(S_t|e_t, e_{1:t-1})$.



2 Tracking the Jabberwock

Lewis' Jabberwock is in the wild: its position is in a two-dimensional discrete grid, but this time the grid is not bounded. In other words, the position of the Jabberwock is a pair of integers $z = (x, y) \in \mathbb{Z}^2 = \{\dots, -2, -1, 0, 1, 2, \dots\} \times \{\dots, -2, -1, 0, 1, 2, \dots\}$. At each time step $t = 1, 2, 3, \dots$, the Jabberwock is in some cell $Z_t = z \in \mathbb{Z}^2$, and it moves to cell Z_{t+1} randomly as follows: with probability $1/2$, it stays where it is; otherwise, it chooses one of its four neighboring cells uniformly at random (fortunately, no teleportation is allowed this week!).

- (a) Write a function for the transition probability $P(Z_{t+1} = (x', y') | Z_t = (x, y))$.

We will use the particle filtering algorithm to track the Jabberwock. As a source of randomness use values in order from the following sequence $\{a_i\}_{1 \leq i \leq 14}$. Use these values to sample from any discrete distribution of the form $P(X)$ where X takes values in $\{1, 2, \dots, N\}$. Given $a_i \sim U[0, 1]$, return j such that $\sum_{k=1}^{j-1} P(X = k) \leq a_i < \sum_{k=1}^j P(X = k)$. You have to fix an ordering of the elements for this procedure to make sense.

a_1	a_2	a_3	a_4	a_5	a_6	a_7	a_8	a_9	a_{10}	a_{11}	a_{12}
0.142	0.522	0.916	0.792	0.703	0.231	0.036	0.859	0.677	0.221	0.156	0.249

At each time step t you get an observation of the x coordinate R_t in which the Jabberwock sits, but it is a noisy observation. Given the true position $Z_t = (x, y)$, you observe the correct value according to the following probability:

$$P(R_t = r | Z_t = (x, y)) \propto (0.5)^{|x-r|}$$

- (b) Suppose that you know that half of the time, the Jabberwock starts at $z_1 = (0, 0)$, and the other half, at $z_1 = (1, 1)$. You get the following observations: $R_1 = 1, R_2 = 0, R_3 = 1$. Fill out the table for each time step using a particle filter with 2 particles to compute an approximation to $P(Z_1, Z_2, Z_3 | r_1, r_2, r_3)$. Sample transitions from the table below using the a_i 's as our source of randomness. The a_i 's you should use for each step have been indicated in the last row of each table. Note that going "left" decrements the x-coordinate by 1, and going "down" decrements the y-coordinate by 1.

[0; 0.5)	Stay
[0.5; 0.625)	Up
[0.625; 0.75)	Left
[0.75; 0.875)	Right
[0.875; 1)	Down

Initial	Belief $\hat{P}(z_1)$	Weights $P(r_1 z_1)$	Unnormalized $\hat{P}(z_1, r_1)$	Normalized $\hat{P}(z_1 r_1)$	Resampling
$p_1 = (0, 0)$ $p_2 = (1, 1)$ a_1, a_2	1/2 1/2				$p_1 = (,)$ $p_2 = (,)$ a_3, a_4

Transition $P(z_2 z_1)$	Belief $\hat{P}(z_2 r_1)$	Weights $P(r_2 z_2)$	Unnormalized $\hat{P}(z_2, r_2 r_1)$	Normalized $\hat{P}(z_2 r_1, r_2)$	Resampling
$p_1 = (,)$ $p_2 = (,)$ a_5, a_6					$p_1 = (,)$ $p_2 = (,)$ a_7, a_8

Transition $P(z_3 z_2)$	Belief $\hat{P}(z_3 r_1, r_2)$	Weights $P(r_3 z_3)$	Unnormalized $\hat{P}(z_3, r_3 r_1, r_2)$	Normalized $\hat{P}(z_3 r_1, r_2, r_3)$	Resampling
$p_1 = (,)$ $p_2 = (,)$ a_9, a_{10}					$p_1 = (,)$ $p_2 = (,)$ a_{11}, a_{12}

- (c) Use your samples (the unweighted particles in the last step) to evaluate the posterior probability that the x-coordinate of Z_3 is different than the column of Z_3 , i.e. $X_3 \neq Y_3$.

- (d) What is the problem of using variable elimination instead of a particle filter for tracking Jabberwock?

3 Game Theory: Equilibrium Warm-Up

- (a) What is a Nash Equilibrium?
- (b) What is another example of a solution concept that might be useful? What is the difference between this concept and Nash Equilibrium?
- (c) What are some drawbacks to following the strategy of Nash equilibrium as a solution concept?
- (d) What is the difference between a weakly dominant and strictly dominant strategy?
- (e) What is the difference between pure and mixed strategies?
- (f) Consider rock paper scissors where Player 1's strategy is to always play rock, and Player 2's strategy is to play scissors or paper with equal probability. Is this a Nash Equilibrium?
- (g) Recursively remove dominated strategies to find the Nash Equilibrium of the following game. The order of utilities in each cell is the roman numeral player then the alphabet player.

	A	B	C
i	3,0	0,-5	0,-4
ii	1,-1	3,3	-2,4
iii	2,4	4,1	-1,8

4 Game Theory: Pacman Hunt!

Let us define the Pacman Hunt game for Aarathi and Maia (derivative of the well-known Stag Hunt game!). Suppose Aarathi and Maia are both hunters, they can choose to hunt Pacman or ghosts. If they hunt a ghost, they will always be successful and gain the modest payoff of 1. If they hunt Pacman, they are successful only if they both chose to hunt Pacman because then they can cooperate. In this case, they gain payoff of 2 each. However, if one hunts Pacman and one hunts ghosts, the one that hunted Pacman gets a payoff of 0. We can formulate the game in the below utility table:

Aarathi, Maia	Pacman	Ghost
Pacman	2,2	0,1
Ghost	1,0	1,1

- (a) What are the pure Nash Equilibria of this problem?
- (b) We will now investigate the possibility of a mixed Nash equilibrium. Recall that in a mixed Nash Equilibrium, the utilities of the weighted actions are equal. Let p be the probability that Maia picks Pacman.
- (i) What is the expected value of action Pacman for Aarathi?
- (ii) What is the expected value of action Ghost for Aarathi?
- (iii) What value of p makes these two expected values the same?
- (iv) Since the table is symmetric, the probability that equalizes the value of action Pacman and Ghost for both players is the same. What is the expected utility for both Maia and Aarathi if they play according to the mixed Nash Equilibria? How does this utility compare to the equilibria from (a)?
- (c) What if the game changed and somehow the Pacmen in the wild got larger or more profitable so the utility table then became

Aarathi, Maia	Pacman	Ghost
Pacman	3,3	0,1
Ghost	1,0	1,1

Calculate the mixed Nash Equilibrium for this game. Are the results surprising?