

SOLUTIONS 6

Solution 6.1. We know from Example 3.6 that

$$p(x|y_{1:n}) = \mathcal{N}(x; \mu_p, \sigma_p^2), \quad (1)$$

with

$$\mu_p = \frac{\sigma_0^2 \sum_{i=1}^n y_i + \sigma^2 \mu_0}{\sigma_0^2 n + \sigma^2}, \quad (2)$$

$$\sigma_p^2 = \frac{\sigma_0^2 \sigma^2}{\sigma_0^2 n + \sigma^2}. \quad (3)$$

Below, we will zset $\mu_0 = 0$, $\sigma_0^2 = 1$ and $\sigma^2 = 1$ which simplifies example, i.e., we have the true mean

$$\mu_p = \frac{\sum_{i=1}^n y_i}{M + 1}, \quad (4)$$

which we will use for checking our code. In this exercise, I will give different parts of the solution separately to improve clarity.

1. In this part, we need to simulate our data.

```
1 # simulate M data points
2 M = 100
3 rng = np.random.default_rng(25)
4 x = rng.normal(0, 1)
5 y = rng.normal(x, 1, M)
```

2. Secondly, we need to compute our true mean estimate as derived in (4):

```
1 # analytic posterior mean
2 mean_true = np.sum(y)/(M+1)
3 print("Analytic posterior mean: ", mean_true)
```

3. (This is for parts 3-4 together) We will then implement SNIS with $\mu_q = 0$ and $\sigma_q^2 = 1$. For this, we directly define log densities.

```
1 # define log prior
2 def logp(x):
3     return -x**2/2 - np.log(np.sqrt(2*np.pi))
4
5 # define log likelihood
6 def loglik(x, y):
7     return -(x-y)**2/2 - np.log(np.sqrt(2*np.pi))
8
9 # define log proposal
10 def logq(x):
11     return -x**2/2 - np.log(np.sqrt(2*np.pi))
12
13 def ESS(w):
14     return 1/np.sum(w**2)
15
16 N = 10000
17
18 x = rng.normal(0, 1, N) # sample from q(x)
```

```

19
20 logW = np.zeros(N)
21 for i in range(N):
22     logW[i] = np.sum(loglik(x[i], y)) + logp(x[i]) - logq(x[i])
23
24 log_hat_W = logW - np.max(logW)
25
26 w = np.exp(log_hat_W)/np.sum(np.exp(log_hat_W)) # weights with
27                                           log-trick
28 w2 = np.exp(logW)/np.sum(np.exp(logW)) # weights without log-
29                                           trick
30
31 # mean estimate
32 mean = np.sum(w*x)
33 mean2 = np.sum(w2*x)
34
35 print("Mean estimate (stable): ", mean)
36 print("ESS: ", ESS(w))
37 print("Mean estimate (unstable): ", mean2)
38 print("ESS: ", ESS(w2))

```

Solution 6.2. This is discussed in the lecture, so hopefully you will have more intuition about it. The first part of this exercise implements the SNIS.

```

1 import numpy as np
2 import matplotlib.pyplot as plt
3
4 def bar_p(x): # implementing the density just for visualisation!
5     return np.exp(-x[0]**2/10 - x[1]**2/10 - 2 * (x[1] - x[0]**2)**2)
6
7 def q(x):
8     return np.exp(- x[0]**2/2 - x[1]**2/2) / (2 * np.pi)
9
10 def logbar_p(x):
11     return - x[0]**2/10 - x[1]**2/10 - 2 * (x[1] - x[0]**2)**2
12
13 def loglik(y, x, sig):
14     H = [1, 0]
15     return -(y - H @ x)**2/(2 * sig**2) - np.log(sig * np.sqrt(2 * np.
16                                           pi))
17
18 def logq(x):
19     return - x[0]**2/2 - x[1]**2/2 - np.log(2 * np.pi)
20
21 def ESS(w):
22     return 1/np.sum(w**2)
23
24 y = 1
25 sig = 0.05
26
27 N = 10000
28 rng = np.random.default_rng(25)
29 # sample from q
30 x = rng.normal(0, 1, (2, N)) # 2 x N matrix (2 dimensional, N samples)
31
32 # compute logW
33 logW = np.zeros(N)
34 for i in range(N):

```

```

34     logW[i] = (loglik(y, x[:, i], sig)) + logbar_p(x[:, i]) - logq(x[:,
35                                     , i])
36 # compute log_hat_W
37 log_hat_W = logW - np.max(logW)
38 w = np.exp(log_hat_W)/np.sum(np.exp(log_hat_W))
39
40 # compute mean estimate
41 mean = np.sum(w*x, axis=1)
42
43 # compute ESS
44 print("ESS: ", ESS(w))

```

Having obtained the weights w and the samples x , we can now resample. We will use the following code for resampling and plot the result.

```

1 # resample N samples
2 x_resampled = np.zeros((2, N))
3 for i in range(N):
4     x_resampled[:, i] = x[:, rng.choice(N, p=w)]
5     # rng.choice chooses an index from 0 to N-1 with probability w
6
7 # plot resampled samples
8 x_bb = np.linspace(-4, 4, 100)
9 y_bb = np.linspace(-2, 6, 100)
10 X_bb, Y_bb = np.meshgrid(x_bb, y_bb)
11 Z_bb = np.zeros((100, 100))
12 for i in range(100):
13     for j in range(100):
14         Z_bb[i, j] = bar_p([X_bb[i, j], Y_bb[i, j]])
15 plt.contourf(X_bb, Y_bb, Z_bb, 100, cmap='RdBu')
16 plt.scatter(x_resampled[0, :], x_resampled[1, :], s=10, c='white')
17 plt.show()

```

Note that, as explained in the lecture, the result makes sense. We had a 2D banana prior but only observed x_1 dimension with some noise as $y = Hx + \sigma W$ where $W \sim \mathcal{N}(0, 1)$ and since $H = [1, 0]$, this equals to $y = x_1 + \sigma W$ with small σ . It means that we can only know that the object in 2D will reside parallel to x_1 axis, according to our prior. Hence the samples from the posterior taking a vertical shape along this axis (would get even more vertical with smaller σ – watch the discussion in the lecture).