
Exercises 8

Gaussian processes in numpy

Exercise 8.1 Sample from GP prior

In this exercises we will write the code needed to draw and plot samples of f from a Gaussian process prior with squared exponential (or, equivalently, RBF) kernel

$$f \sim \mathcal{GP}(m, \kappa), \quad \text{whith } m(x) = 0 \text{ and } \kappa(x, x') = \sigma_f^2 e^{-\frac{1}{2\ell^2} \|x-x'\|^2}$$

To implement this, we choose a vector of m test input points \mathbf{x}_* . We will choose \mathbf{x}_* to contain sufficiently many points, such that it will *appear* as a continuous line on the screen. We then evaluate the $m \times m$ covariance matrix $\kappa(\mathbf{x}_*, \mathbf{x}_*)$ and thereafter generate samples from the multivariate normal distribution

$$f(\mathbf{x}_*) \sim \mathcal{N}(m(\mathbf{x}_*), \kappa(\mathbf{x}_*, \mathbf{x}_*)).$$

- Use `numpy.linspace` to construct a vector \mathbf{x}_* with $m = 101$ elements equally spaced from -5 to 5.
- Construct a mean vector $m(\mathbf{x}_*)$ with 101 elements all equal to zero and the 101×101 covariance matrix $\kappa(x_*, x_*)$. The expression for $\kappa(\cdot, \cdot)$ is given above. Let the hyperparameters be $\ell = 2$ and $\sigma_f^2 = 1$.
- Use `scipy.stats.multivariate_normal` (you might need to use the option `allow_singular=True`) to draw 25 samples $f^{(1)}(\mathbf{x}_*), \dots, f^{(25)}(\mathbf{x}_*)$ from the multivariate normal distribution $f(\mathbf{x}_*) \sim \mathcal{N}(m(\mathbf{x}_*), \kappa(\mathbf{x}_*, \mathbf{x}_*))$.
- Plot the samples $f^{(1)}(\mathbf{x}_*), \dots, f^{(25)}(\mathbf{x}_*)$ versus the input vector \mathbf{x}_* .
- Try another value of ℓ and repeat (b)-(d). How do the two plots differ, and why?

Exercise 8.2 GP posterior

In this exercise we will perform Gaussian process regression. That means, based on the n observations $\mathcal{D} = \{x_i, f(x_i)\}_{i=1}^n$ and the prior belief $f \sim \mathcal{GP}(0, \kappa(x, x'))$, we want to find the posterior $p(f|\mathcal{D})$. (In the previous problem, we were only concerned with the prior $p(f)$, not conditioned on having observed the data \mathcal{D} .) We consider the same Gaussian process prior (same mean $m(x)$ and $\kappa(x, x')$ and hyperparameters) as in the previous exercise.

- Construct two vectors $\mathbf{x} = [-4, -3, -1, 0, 2]^\top$ and $\mathbf{f} = [-2, 0, 1, 2, -1]^\top$, which will be our training data (that is, $n = 5$).
- Keep \mathbf{x}_* as in the previous problem. In addition to the $m \times m$ matrix $\kappa(\mathbf{x}_*, \mathbf{x}_*)$, now also compute the $n \times m$ matrix $\kappa(\mathbf{x}, \mathbf{x}_*)$ and the $n \times n$ matrix $\kappa(\mathbf{x}, \mathbf{x})$. Hint: You might find it useful to define a function that returns $\kappa(x, x')$, taking x and x' as arguments.
- Use the training data \mathbf{x}, \mathbf{y} and the matrices constructed in (b) to compute the posterior mean $\boldsymbol{\mu}_{\text{posterior}}$ and the posterior covariance $\mathbf{K}_{\text{posterior}}$ for \mathbf{x}_* , by using the equations for conditional multivariate normal distributions.
- In a similar manner as in (c) and (d) in the previous problem, draw 25 samples from the multivariate distribution $f(\mathbf{x}_*) \sim \mathcal{N}(\boldsymbol{\mu}_{\text{posterior}}, \mathbf{K}_{\text{posterior}})$ and plot these samples ($f^{(j)}(\mathbf{x}_*)$ vs. \mathbf{x}_*) together with the posterior mean ($\boldsymbol{\mu}_{\text{posterior}}$ vs. \mathbf{x}_*) and the actual measurements (\mathbf{y} vs. \mathbf{x}). How do the samples in this plot differ from the prior samples in the previous problem?

- (e) Instead of plotting samples, plot a credibility region. Here, a credibility region is based on the (marginal) posterior variance. The 68% credibility region, for example, is the area between $\boldsymbol{\mu}_{\text{posterior}} - \sqrt{\mathbf{K}_{\text{posterior}}^d}$ and $\boldsymbol{\mu}_{\text{posterior}} + \sqrt{\mathbf{K}_{\text{posterior}}^d}$, where $\mathbf{K}_{\text{posterior}}^d$ is a vector with the diagonal elements of $\mathbf{K}_{\text{posterior}}$. What is the connection between the credibility regions and the samples you drew previously?
- (f) Now, consider the setting where the measurements are corrupted with noise, $y_i = f(\mathbf{x}_i) + \varepsilon$, $\varepsilon \sim \mathcal{N}(0, \sigma^2)$. Use $\sigma = 0.1$ and repeat (c)-(e) with this modification of the model. What is the difference in comparison to the previous plot? What is the interpretation?
- (g) Explore what happens with another length scale ℓ .

Exercise 8.3 Other covariance functions/kernels

The squared exponential kernel/covariance function gives samples which are smooth and infinitely continuously differentiable. Other kernels make other assumptions. Now try the previous problems using the exponential kernel instead,

$$\kappa(x, x') = e^{-\frac{1}{\ell}|x-x'|}.$$

Exercise 8.4 Learning hyperparameters

Until now, we have made GP regression using predefined hyperparameters, such as the lengthscale ℓ and noise variance σ^2 . In this exercise, we will estimate ℓ and σ^2 from the data by maximizing the marginal likelihood. The logarithm of the marginal likelihood for a Gaussian process observed with Gaussian noise is

$$p\{\mathbf{y}|\mathbf{x}, \ell\} = -\frac{1}{2}\mathbf{y}^T \mathbf{K}_y^{-1} \mathbf{y} - \frac{1}{2} \log |\mathbf{K}_y| - \frac{n}{2} \log 2\pi$$

where $\mathbf{K}_y = \kappa(\mathbf{x}, \mathbf{x}) + \sigma^2 \mathbf{I}$.

- (a) Write a function that takes \mathbf{x} , \mathbf{y} , ℓ and σ^2 as inputs and produces the marginal likelihood as output for the squared exponential covariance function.
- (b) Consider the same data as before. Use $\sigma^2 = 0$ and compute the marginal likelihood for values of ℓ between 0.1 and 1 and plot it. What seems to be the maximal value of the marginal likelihood on this interval? Do GP regression based on this value of ℓ .

Exercise 8.5 Learning hyperparameters II

In this exercise we investigate a setting where the marginal likelihood has multiple local minima.

- (a) Now, consider the following data

$$\mathbf{x} = [-5 \quad -3 \quad 0 \quad 0.1 \quad 1 \quad 4.9 \quad 5]^T, \quad \mathbf{y} = [0 \quad -0.5 \quad 1 \quad 0.7 \quad 0 \quad 1 \quad 0.7]^T$$

and compute the log marginal likelihood for both ℓ and σ . Use a logarithmic 2D-grid for values of ℓ spanning from 10^{-1} to 10^2 and for σ^2 spanning from 10^{-2} to 10^0 . Visualize the marginal likelihood on that grid with a contour plot.



- (b) Find the hyperparameters ℓ and σ^2 that correspond to the maximal marginal likelihood. Perform GP regression on the data using these hyperparameters.
- (c) Perform GP regression for the hyperparameters that correspond to other possible local optima of the marginal likelihood. What differences do you see in your posterior?

Solutions 8

Gaussian processes in numpy

Solution to Exercise 8.1

```
1 import numpy as np
2 import matplotlib.pyplot as plt
3 from scipy.stats import multivariate_normal
4
5 m = 101
6 xs = np.linspace(-5,5,m) # Test input vector
7 mxs = np.zeros(m) # Zero mean vector
8
9 l = np.sqrt(2) # hyperparameters
10 sf2 = 1
11 Kss = sf2*np.exp(-1/(2*l**2)*np.abs(xs[:,np.newaxis]-xs[:,np.newaxis].T)**2) # Covariance matrix
12
13 s = 25 # Draw samples from the prior
14 fs = multivariate_normal(mean=mxs,cov=Kss,allow_singular=True).rvs(s).T
15
16 plt.plot(xs,fs,'gray') # Plot the samples
17 plt.title('l = sqrt(2)')
18 plt.show()
19
20 # Use another length scale
21 l = 0.5
22 Kss = sf2*np.exp(-1/(2*l)*np.abs(xs[:,np.newaxis]-xs[:,np.newaxis].T)**2)
23 fs = multivariate_normal(mean=mxs,cov=Kss,allow_singular=True).rvs(s).T
24 plt.plot(xs,fs,'gray')
25 plt.title('l = 0.5')
26 plt.show()
```



Solution to Exercise 8.2 (d) All posterior samples passes through the observed data points (the prior samples do not necessarily do that). This is natural, since the posterior distribution of $f(\mathbf{x}_*)$ is conditioned on the observations, and must the posterior distribution pass through them.

(e) The 68% credibility region, for example, contains 68% of the posterior samples.

(f) The posterior distribution does not pass exactly through all observations any more, but the observations are not to some extent “explained” as noise.

```

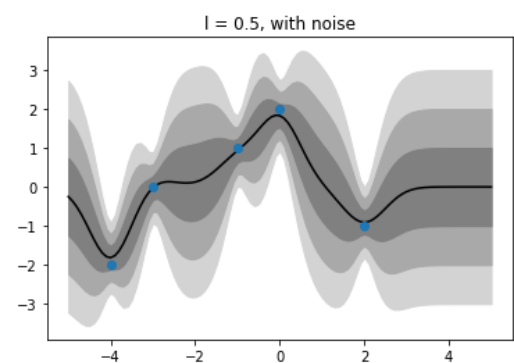

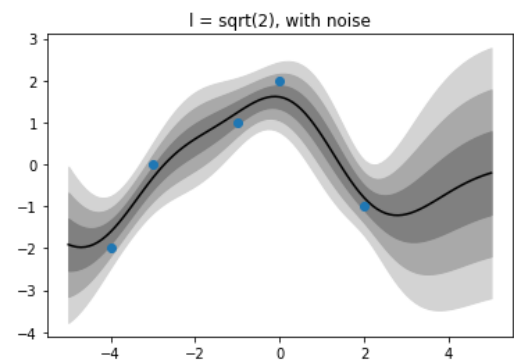

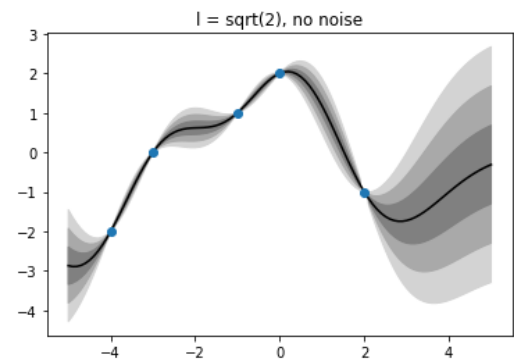
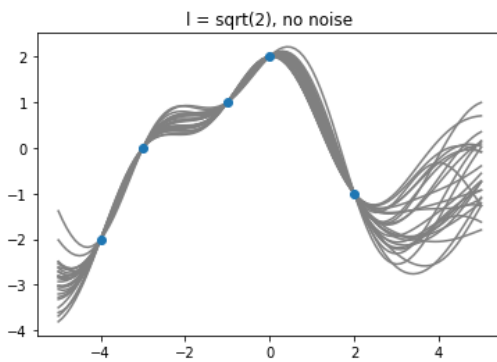
1 import numpy as np
2 import matplotlib.pyplot as plt
3 from scipy.stats import multivariate_normal
4
5 n = 5
6 x=np.array([-4,-3,-1,0,2]) # Observed inputs
7 f=np.array([-2,0,1,2,-1]) # Observed function values
8
9 m = 101
10 xs = np.linspace(-5,5,m) # Test input vector
11
12 l = np.sqrt(2) # hyperparameters
13 sf2 = 1
14 def k(x, xp, l):
15     return sf2*np.exp(-1/(2*l**2)*np.abs(x[:, np.newaxis]-xp[:, np.newaxis].T)**2)
16 Kss = k(xs, xs, l)
17 Ks = k(x, xs, l)
18 K = k(x, x, l)
19
20 mu_post = (Ks.T@np.linalg.inv(K))@f
21 K_post = Kss - Ks.T@np.linalg.inv(K)@Ks
22
23 s = 25 # Draw samples from the posterior
24 fs = multivariate_normal(mean=mu_post, cov=K_post, allow_singular=True).rvs(s).T
25
26 plt.plot(xs, fs, 'gray') # Plot the samples
27 plt.scatter(x, f, zorder=3)
28 plt.title('l = sqrt(2), no noise')
29 plt.show()
30
31 plt.plot(xs, mu_post, 'black') # Plot credibility regions
32 plt.fill_between(xs, mu_post + 3*np.sqrt(np.diag(K_post)), mu_post - 3*np.sqrt(np.diag(K_post)), color=
    'lightgray')
33 plt.fill_between(xs, mu_post + 2*np.sqrt(np.diag(K_post)), mu_post - 2*np.sqrt(np.diag(K_post)), color=
    'darkgray')
34 plt.fill_between(xs, mu_post + 1*np.sqrt(np.diag(K_post)), mu_post - 1*np.sqrt(np.diag(K_post)), color=
    'gray')
35 plt.scatter(x, f, zorder=3)
36 plt.title('l = sqrt(2), no noise')
37 plt.show()
38
39 # Include measurement noise
40 K = k(x, x, l) + 0.1*np.eye(n)
41 mu_post = (Ks.T@np.linalg.inv(K))@f
42 K_post = Kss - Ks.T@np.linalg.inv(K)@Ks
43 fs = multivariate_normal(mean=mu_post, cov=K_post, allow_singular=True).rvs(s).T
44
45 plt.plot(xs, fs, 'gray') # Plot the samples
46 plt.scatter(x, f, zorder=3)
47 plt.title('l = sqrt(2), with noise')
48 plt.show()
49
50 plt.plot(xs, mu_post, 'black') # Plot credibility regions
51 plt.fill_between(xs, mu_post + 3*np.sqrt(np.diag(K_post)), mu_post - 3*np.sqrt(np.diag(K_post)), color=
    'lightgray')
52 plt.fill_between(xs, mu_post + 2*np.sqrt(np.diag(K_post)), mu_post - 2*np.sqrt(np.diag(K_post)), color=
    'darkgray')
53 plt.fill_between(xs, mu_post + 1*np.sqrt(np.diag(K_post)), mu_post - 1*np.sqrt(np.diag(K_post)), color=
    'gray')
54 plt.scatter(x, f, zorder=3)
55 plt.title('l = sqrt(2), with noise')
56 plt.show()
57

```

```

58 # Try another length scale
59 l = 0.5
60 Kss = k(xs,xs,l)
61 Ks = k(x,xs,l)
62 K = k(x,x,l) + 0.1*np.eye(n)
63 mu_post = (Ks.T@np.linalg.inv(K))@(f)
64 K_post = Kss - Ks.T@np.linalg.inv(K)@Ks
65 fs = multivariate_normal(mean=mu_post,cov=K_post,allow_singular=True).rvs(s).T
66 plt.plot(xs,fs,'gray')
67 plt.scatter(x,f,zorder=3)
68 plt.title('l = 0.5, with noise')
69 plt.show()
70 plt.plot(xs,mu_post,'black')
71 plt.fill_between(xs,mu_post + 3*np.sqrt(np.diag(K_post)),mu_post - 3*np.sqrt(np.diag(K_post)),color='lightgray')
72 plt.fill_between(xs,mu_post + 2*np.sqrt(np.diag(K_post)),mu_post - 2*np.sqrt(np.diag(K_post)),color='darkgray')
73 plt.fill_between(xs,mu_post + 1*np.sqrt(np.diag(K_post)),mu_post - 1*np.sqrt(np.diag(K_post)),color='gray')
74 plt.scatter(x,f,zorder=3)
75 plt.title('l = 0.5, with noise')
76 plt.show()






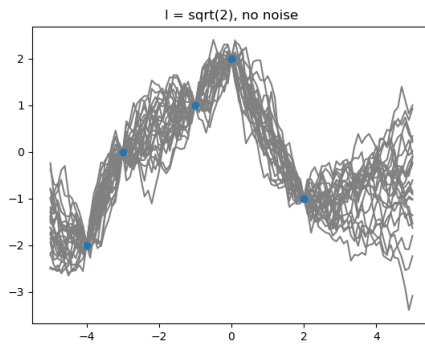
```



Solution to Exercise 8.3

Change the kernel definition in the code for exercise 8.2 with the following

```
14 def k(x, xp, l):
15     return sf2*np.exp(-1/(2*l**2)*np.abs(x[:,np.newaxis]-xp[:,np.newaxis].T))
```



**Solution to Exercise 8.4 (a)** We use `scipy.optimize.minimize`

```
1 import numpy as np
2 import matplotlib.pyplot as plt
3 from scipy.stats import multivariate_normal
4 from scipy.optimize import minimize
5
6 n = 5
7 x=np.array([-4,-3,-1,0,2]) # Observed inputs
8 f=np.array([-2,0,1,2,-1]) # Observed function values
9
10 m = 101
11 xs = np.linspace(-5,5,m) # Test input vector
12
13
14 # Kernel function
15 def k(x, xp, l, sf2=1):
16     return sf2*np.exp(-1/(2*l**2)*np.abs(x[:,np.newaxis]-xp[:,np.newaxis].T)**2)
17
```

```

18 # marginal likelihood
19 def neg_log_marg_lik(y, x, l, s2): #
20     K = k(x, x, l) + s2*np.eye(len(x))
21     return 0.5*(y.T@np.linalg.inv(K))@y + np.log(np.linalg.det(K)) + len(x)*np.log(2*np.pi)
22
23 # Optimize the function
24 def optimal_hyperparameters(y, x, l=np.sqrt(2), s2=0):
25     result = minimize(lambda hyper: neg_log_marg_lik(y, x, hyper[0], hyper[1]), [l, s2])
26     return result.x[0], result.x[1]
27
28 # Create optimal kernel
29 l, s2 = optimal_hyperparameters(f, x)
30 K = k(x, x, l=1) + s2*np.eye(n)
31 Ks = k(x, xs, l=1)
32 Kss = k(xs, xs, l=1)
33
34 mu_post = (Ks.T@np.linalg.inv(K))@f
35 K_post = Kss - Ks.T@np.linalg.inv(K)@Ks
36
37 s = 25 # Draw samples from the posterior
38 fs = multivariate_normal(mean=mu_post, cov=K_post, allow_singular=True).rvs(s).T
39
40 mu_post = (Ks.T@np.linalg.inv(K))@f
41 K_post = Kss - Ks.T@np.linalg.inv(K)@Ks
42 fs = multivariate_normal(mean=mu_post, cov=K_post, allow_singular=True).rvs(s).T
43
44 plt.plot(xs, fs, 'gray') # Plot the samples
45 plt.scatter(x, f, zorder=3)
46 plt.title(f'l = {l:.3f}, with noise s2 = {s2:.3f}')
47 plt.show()
48
49 plt.plot(xs, mu_post, 'black') # Plot credibility regions
50 plt.fill_between(xs, mu_post + 3*np.sqrt(np.diag(K_post)), mu_post - 3*np.sqrt(np.diag(K_post)),
51                 color='lightgray')
51 plt.fill_between(xs, mu_post + 2*np.sqrt(np.diag(K_post)), mu_post - 2*np.sqrt(np.diag(K_post)),
52                 color='darkgray')
52 plt.fill_between(xs, mu_post + 1*np.sqrt(np.diag(K_post)), mu_post - 1*np.sqrt(np.diag(K_post)),
53                 color='gray')
53 plt.scatter(x, f, zorder=3)
54 plt.title(f'l = {l:.3f}, with noise s2 = {s2:.3f}')
55 plt.show()
56
57 # Same data, kernel, and marginal likelihood function as before

```



(b)

```




1 # Same data, kernel, and marginal likelihood function as before
2
3 # Create grid for l
4 nl = 101
5 cl = np.linspace(0.1, 1, nl)
6 marg_lik = np.zeros(nl)
7
8 # Compute marginal likelihood over the grid
9 for i in range(nl):
10     marg_lik[i] = -neg_log_marg_lik(f, x, l=cl[i], s2=0)

```

```

11
12 idx = np.argmax(marg_lik)
13 lmax = cl[idx]
14
15 plt.plot(cl, marg_lik)
16 plt.scatter(cl[idx], marg_lik[idx])
17 plt.title(f'Marginal likelihood (max at {lmax:.3f})')
18 plt.show()
19
20 # Create optimal kernel
21 K = k(x, x, l=lmax)
22 Ks = k(x, xs, l=lmax)
23 Kss = k(xs, xs, l=lmax)
24
25 mu_post = (Ks.T@np.linalg.inv(K))@f
26 K_post = Kss - Ks.T@np.linalg.inv(K)@Ks
27
28 s = 25 # Draw samples from the posterior
29 fs = multivariate_normal(mean=mu_post, cov=K_post, allow_singular=True).rvs(s).T
30
31 mu_post = (Ks.T@np.linalg.inv(K))@f
32 K_post = Kss - Ks.T@np.linalg.inv(K)@Ks
33 fs = multivariate_normal(mean=mu_post, cov=K_post, allow_singular=True).rvs(s).T
34
35 plt.plot(xs, fs, 'gray') # Plot the samples
36 plt.scatter(x, f, zorder=3)
37 plt.title(f'l = {lmax:.3f}, no noise')
38 plt.show()
39
40 plt.plot(xs, mu_post, 'black') # Plot credibility regions
41 # Add a small value to prevent negative variances
42 plt.fill_between(xs, mu_post + 3*np.sqrt(np.diag(K_post) + 0.0001), mu_post - 3*np.sqrt(np.diag(
    K_post)+0.0001), color='lightgray')
43 plt.fill_between(xs, mu_post + 2*np.sqrt(np.diag(K_post) + 0.0001), mu_post - 2*np.sqrt(np.diag(
    K_post)+0.0001), color='darkgray')
44 plt.fill_between(xs, mu_post + 1*np.sqrt(np.diag(K_post) + 0.0001), mu_post - 1*np.sqrt(np.diag(
    K_post)+0.0001), color='gray')
45 plt.scatter(x, f, zorder=3)
46 plt.title(f'l = {lmax:.3f}, no noise')
47 plt.show()

```



Solution to Exercise 8.5 (a)

```

1 import numpy as np
2 import matplotlib.pyplot as plt
3 from scipy.stats import multivariate_normal
4 from scipy.optimize import minimize
5
6 x=np.array([-5, -3, 0, 0.1, 1, 4.9, 5])
7 y=np.array([0, -0.5, 1, 0.7, 0, 1, 0.7])
8 n = len(x)
9
10 m = 101
11 xs = np.linspace(-6, 6, m) # Test input vector
12
13
14 # Kernel function
15 def k(x, xp, l, sf2=1):


```



```

16     return sf2*np.exp(-1/(2*l**2)*np.abs(x[:,np.newaxis]-xp[:,np.newaxis].T)**2)
17
18 # marginal likelihood
19 def neg_log_marg_lik(y, x, l, s2): #
20     K = k(x, x, l) + s2*np.eye(len(x))
21     return 0.5*(y.T@np.linalg.inv(K))@y + np.log(np.linalg.det(K)) + len(x)*np.log(2*np.pi)
22
23 nl = 100
24 ns = 100
25 cl = np.logspace(-1, 1.1, nl)
26 cs = np.logspace(-2, -0.5, ns)
27
28 L, S = np.meshgrid(cl, cs)
29 Z = np.zeros((nl, ns))
30
31 for i in range(nl):
32     for j in range(ns):
33         Z[i,j] = - neg_log_marg_lik(y, x, L[i,j], S[i,j])
34
35
36 plt.contour(L, S, Z, 500, vmin=-7)
37 plt.show()

```



(b)

```

1 # Local maximum
2 local_max = minimize(lambda hyper: neg_log_marg_lik(y, x, hyper[0], hyper[1]), [10, 0.2])
3 print(local_max.message)
4
5 # Create first optimal kernel
6 l, s2 = local_max.x
7 K = k(x, x, l=1) + s2*np.eye(n)
8 Ks = k(x, xs, l=1)
9 Kss = k(xs, xs, l=1)
10
11 mu_post = (Ks.T@np.linalg.inv(K))@y
12 K_post = Kss - Ks.T@np.linalg.inv(K)@Ks
13
14 plt.plot(xs,mu_post,'black') # Plot credibility regions
15 plt.fill_between(xs,mu_post + 3*np.sqrt(np.diag(K_post)),mu_post - 3*np.sqrt(np.diag(K_post)),
16                 color='lightgray')
17 plt.fill_between(xs,mu_post + 2*np.sqrt(np.diag(K_post)),mu_post - 2*np.sqrt(np.diag(K_post)),
18                 color='darkgray')
19 plt.fill_between(xs,mu_post + 1*np.sqrt(np.diag(K_post)),mu_post - 1*np.sqrt(np.diag(K_post)),
20                 color='gray')
21 plt.scatter(x,y,zorder=3)
22 plt.title(f'Local: l = {l:.3f}, with noise s2 = {s2:.3f}, logp(y) = {-local_max.fun:.3f}')
23 plt.show()

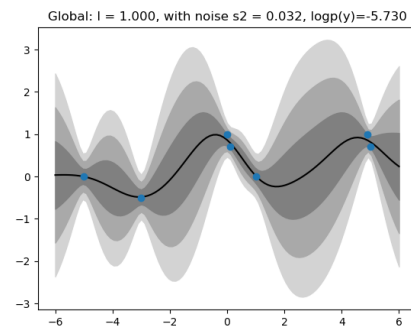
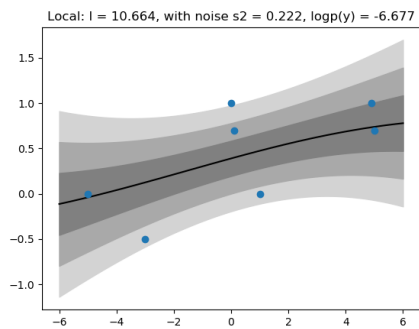
```

(c)

```

1 # Constrained optimization to find the global maximum
2 from scipy.optimize import LinearConstraint
3 global_max = minimize(lambda hyper: neg_log_marg_lik(y, x, hyper[0], hyper[1]), [1, 0.002],
4                       constraints=[LinearConstraint(np.eye(2), [0,0], [1, 0.05])])
5 print(global_max.message)
6
7 # Create second optimal kernel
8 l, s2 = global_max.x
9 K = k(x, x, l=1) + s2*np.eye(n)
10 Ks = k(x, xs, l=1)
11 Kss = k(xs, xs, l=1)
12
13 mu_post = (Ks.T@np.linalg.inv(K))@y
14 K_post = Kss - Ks.T@np.linalg.inv(K)@Ks
15
16 plt.plot(xs,mu_post,'black') # Plot credibility regions
17 plt.fill_between(xs,mu_post + 3*np.sqrt(np.diag(K_post)),mu_post - 3*np.sqrt(np.diag(K_post)),
18                 color='lightgray')
19 plt.fill_between(xs,mu_post + 2*np.sqrt(np.diag(K_post)),mu_post - 2*np.sqrt(np.diag(K_post)),
20                 color='darkgray')
21 plt.fill_between(xs,mu_post + 1*np.sqrt(np.diag(K_post)),mu_post - 1*np.sqrt(np.diag(K_post)),
22                 color='gray')
23 plt.scatter(x,y,zorder=3)
24 plt.title(f'Global: l = {l:.3f}, with noise s2 = {s2:.3f}, logp(y)={-global_max.fun:.3f}')
25 plt.show()

```



Bibliography

- [1] David Barber. *Bayesian reasoning and machine learning*. Cambridge University Press, 2012.
- [2] Christopher M Bishop. *Pattern recognition and machine learning*. springer, 2006.
- [3] Kevin B Korb and Ann E Nicholson. *Bayesian artificial intelligence*. CRC press, 2010.
- [4] Steffen L Lauritzen and David J Spiegelhalter. Local computations with probabilities on graphical structures and their application to expert systems. *Journal of the Royal Statistical Society: Series B (Methodological)*, 50(2):157–194, 1988.
- [5] Kevin P Murphy. *Machine learning: a probabilistic perspective*. MIT press, 2012.