
Exercises 9

Gaussian processes in scikit-learn

Exercise 9.1 GP posterior

Repeat problem 8.2 using `GaussianProcessRegressor` from `sklearn.gaussian_process`, only plotting the credibility regions based on the posterior predictive variance (not drawing samples). Some useful hints:

- As all supervised machine learning methods in `sklearn`, you first have to construct an object from the model class (in this case `GaussianProcessRegressor`), and thereafter train it on data by using its member function `fit()`. To obtain predictions, use the member function `predict()`. To the latter, you will either have to pass `return_std=True` or `return_cov=True` in order to obtain information about the posterior predictive variance.
- When you construct the model, you have to define a kernel. The kernels are available in `sklearn.gaussian_process.kernels`, where the squared exponential/RBF kernel is available as `RBF`
- The function `fit()` automatically optimizes the hyperparameters. To turn that feature off, you have to pass the argument `optimizer=None` to `GaussianProcessRegressor`.
- To include the measurement noise, you can formulate it as part of the kernel by using the kernel `WhiteKernel`

Exercise 9.2 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. That is done automatically by the `fit()`-function in scikit-learn. Use, as before, the RBF kernel and measurement noise together, this time with the data

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

- (a) You still have to provide an initial value of the hyperparameters. Try $\ell = 1$ and $\sigma_n^2 = 0.1$. What hyperparameters do you get when optimizing? Plot the corresponding mean and credibility regions.
- (b) Try instead to initialize with $\ell = 10$ and $\sigma_n^2 = 1$. What do you get now?
- (c) Try to explain what happens by making a grid over different hyperparameter values, and inspect the marginal likelihood for each point in that grid. The `GaussianProcessRegressor` class has a member function `log_marginal_likelihood()` which you may use. (Do not forget to turn off the hyperparameter optimization!)

Exercise 9.3 Modeling CO₂ levels

The amount of carbon dioxide in the atmosphere has been measured continuously at the Mauna Loa observatory, Hawaii. In this problem, you should use a Gaussian process to model the data from 1958 to 2003, and see how well that model can be used for predicting the data from 2004-2019. They present their latest data at their homepage <https://www.esrl.noaa.gov/gmd/ccgg/trends/>, but for your convenience you can use the data in a more convenient format available here https://github.com/gpschool/labs/raw/2019/.resources/mauna_loa and the following code snippet

```
1 import pickle
2
3 with open("mauna_loa", "rb") as fid:
4     data = pickle.load(fid)
5
6 x = data['X'].flatten()
7 y = data['Y'].flatten()
8
9 xtest = data['Xtest'].flatten()
10 ytest = data['Ytest'].flatten()
```

Here, `x` and `y` contains your training data.

Start exploring some simple kernels, and thereafter you may have a look page 118-122 of [6], <http://www.gaussianprocess.org/gpml/chapters/RW5.pdf>, for some inspiration on how to design a more bespoke kernel for this problem.

Solutions 9

Gaussian processes in scikit-learn

Solution to Exercise 9.1

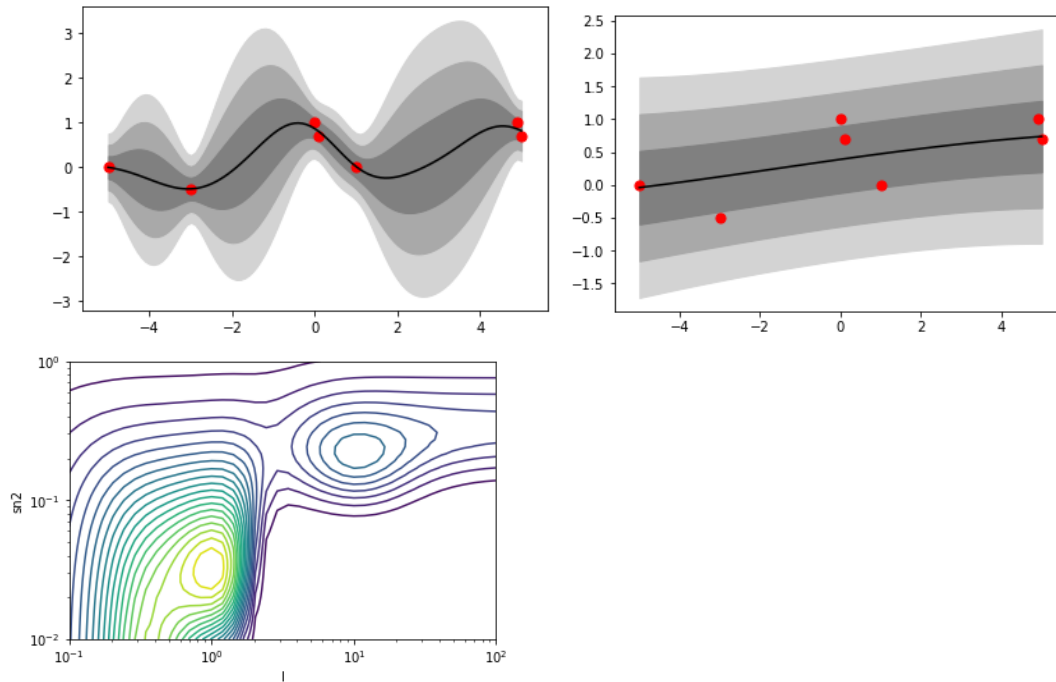
```
1 import numpy as np
2 from matplotlib import pyplot as plt
3
4 from sklearn.gaussian_process import GaussianProcessRegressor
5 from sklearn.gaussian_process.kernels import RBF, WhiteKernel
6
7 n = 5
8 x=np.array([-4,-3,-1,0,2]) # Observed inputs
9 f=np.array([-2,0,1,2,-1]) # Observed function values
10
11 m = 101
12 xs = np.linspace(-5,5,m) # Test input vector
13
14 # Without noise
15 kernel = RBF(length_scale=np.sqrt(2))
16 gp = GaussianProcessRegressor(kernel=kernel,optimizer=None).fit(x[:, np.newaxis], f)
17 f_mean, f_std = gp.predict(xs[:, np.newaxis], return_std=True)
18 plt.plot(xs, f_mean, 'k')
19 plt.fill_between(xs, f_mean - 3*f_std, f_mean + 3*f_std, color='lightgray')
20 plt.fill_between(xs, f_mean - 2*f_std, f_mean + 2*f_std, color='darkgray')
21 plt.fill_between(xs, f_mean - 1*f_std, f_mean + 1*f_std, color='gray')
22 plt.scatter(x, f, c='r', s=50)
23 plt.show()
24
25 # With noise
26 kernel = RBF(length_scale=np.sqrt(2)) + WhiteKernel(noise_level=0.1)
27 gp = GaussianProcessRegressor(kernel=kernel,optimizer=None).fit(x[:, np.newaxis], f)
28 f_mean, f_std = gp.predict(xs[:, np.newaxis], return_std=True)
29 plt.plot(xs, f_mean, 'k')
30 plt.fill_between(xs, f_mean - 3*f_std, f_mean + 3*f_std, color='lightgray')
31 plt.fill_between(xs, f_mean - 2*f_std, f_mean + 2*f_std, color='darkgray')
32 plt.fill_between(xs, f_mean - 1*f_std, f_mean + 1*f_std, color='gray')
33 plt.scatter(x, f, c='r', s=50)
34 plt.show()
35
36 # Different length scale
37 kernel = RBF(length_scale=0.5) + WhiteKernel(noise_level=0.1)
38 gp = GaussianProcessRegressor(kernel=kernel,optimizer=None).fit(x[:, np.newaxis], f)
39 f_mean, f_std = gp.predict(xs[:, np.newaxis], return_std=True)
40 plt.plot(xs, f_mean, 'k')
41 plt.fill_between(xs, f_mean - 3*f_std, f_mean + 3*f_std, color='lightgray')
42 plt.fill_between(xs, f_mean - 2*f_std, f_mean + 2*f_std, color='darkgray')
43 plt.fill_between(xs, f_mean - 1*f_std, f_mean + 1*f_std, color='gray')
44 plt.scatter(x, f, c='r', s=50)
45 plt.show()
```

Solution to Exercise 9.2

```

1 import numpy as np
2 from matplotlib import pyplot as plt
3
4 from sklearn.gaussian_process import GaussianProcessRegressor
5 from sklearn.gaussian_process.kernels import RBF, WhiteKernel
6
7 n = 5
8 x=np.array([-5,-3,0,.1,1,4.9,5]) # Observed inputs
9 y=np.array([0,-.5,1,.7,0,1,.7]) # Observed function values
10
11 m = 101
12 xs = np.linspace(-5,5,m) # Test input vector
13
14 # Optimize hyperparameters
15 kernel = RBF(length_scale=1) + WhiteKernel(noise_level=0.1)
16 gp = GaussianProcessRegressor(kernel=kernel).fit(x[:, np.newaxis], y)
17 print(gp.kernel_)
18 f_mean, f_std = gp.predict(xs[:, np.newaxis], return_std=True)
19 plt.plot(xs, f_mean, 'k')
20 plt.fill_between(xs, f_mean - 3*f_std, f_mean + 3*f_std, color='lightgray')
21 plt.fill_between(xs, f_mean - 2*f_std, f_mean + 2*f_std, color='darkgray')
22 plt.fill_between(xs, f_mean - 1*f_std, f_mean + 1*f_std, color='gray')
23 plt.scatter(x, y, c='r', s=50)
24 plt.show()
25
26 # Optimize hyperparameters with another initialization
27 kernel = RBF(length_scale=10) + WhiteKernel(noise_level=1)
28 gp = GaussianProcessRegressor(kernel=kernel).fit(x[:, np.newaxis], y)
29 print(gp.kernel_)
30 f_mean, f_std = gp.predict(xs[:, np.newaxis], return_std=True)
31 plt.plot(xs, f_mean, 'k')
32 plt.fill_between(xs, f_mean - 3*f_std, f_mean + 3*f_std, color='lightgray')
33 plt.fill_between(xs, f_mean - 2*f_std, f_mean + 2*f_std, color='darkgray')
34 plt.fill_between(xs, f_mean - 1*f_std, f_mean + 1*f_std, color='gray')
35 plt.scatter(x, y, c='r', s=50)
36 plt.show()
37
38
39 # Make a grid of different hyperparameter values to explore the marginal likelihood landscape
40 lv = np.logspace(-1,2,40)
41 sn2v = np.logspace(-2,0,40)
42
43 L,SN2 = np.meshgrid(lv,sn2v)
44
45 xv = np.linspace(-6,6,100)
46
47 margloglik = np.zeros(L.size)
48
49 for i in range(L.size):
50
51     l = L.flatten()[i]
52     sn2 = SN2.flatten()[i]
53
54     # With noise
55     kernel = RBF(length_scale=1) + WhiteKernel(noise_level=sn2)
56     gp = GaussianProcessRegressor(kernel=kernel,optimizer=None).fit(x[:, np.newaxis], y)
57
58     margloglik[i] = gp.log_marginal_likelihood()
59
60 plt.contour(lv,sn2v,np.exp(margloglik.reshape(L.shape))-np.max(margloglik)),levels=20)
61 plt.xscale('log')
62 plt.yscale('log')
63 plt.xlabel('l')
64 plt.ylabel('sn2')
65 plt.show()

```



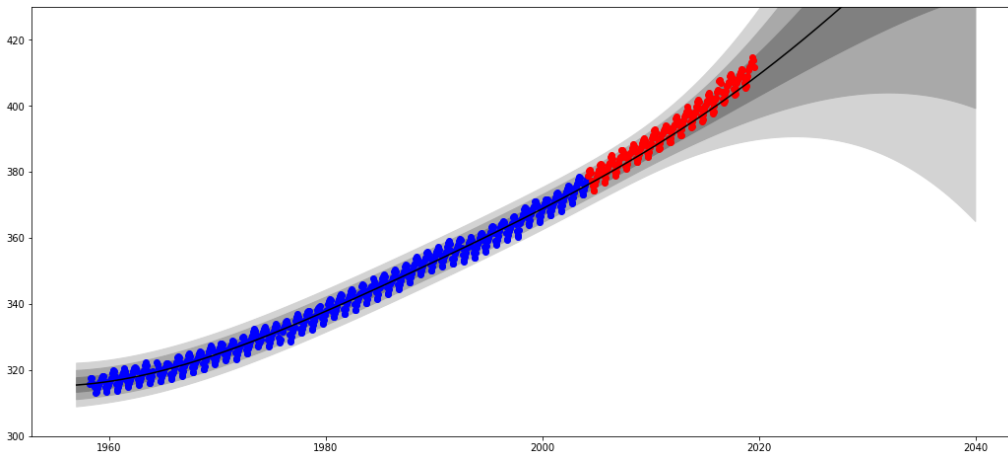
Solution to Exercise 9.3

A first modeling attempt is given by

```

1 import pickle
2
3 with open("mauna_loa", "rb") as fid:
4     data = pickle.load(fid)
5
6 x = data['X'].flatten()
7 y = data['Y'].flatten()
8
9 xtest = data['Xtest'].flatten()
10 ytest = data['Ytest'].flatten()
11
12
13 import numpy as np
14 from matplotlib import pyplot as plt
15
16 from sklearn.gaussian_process import GaussianProcessRegressor
17 from sklearn.gaussian_process.kernels import RBF, WhiteKernel
18
19
20 xs = np.linspace(1957,2040,200)
21
22 kernel = 100*RBF(length_scale=100) + WhiteKernel(noise_level=1)
23 gp = GaussianProcessRegressor(kernel=kernel).fit(x[:,np.newaxis], y)
24 print(gp.kernel_)
25
26 f_mean, f_std = gp.predict(xs[:,np.newaxis], return_std=True)
27
28 plt.figure(figsize=(18,8))
29 plt.plot(xs, f_mean, 'k')
30 plt.fill_between(xs, f_mean - 3*f_std, f_mean + 3*f_std, color='lightgray')
31 plt.fill_between(xs, f_mean - 2*f_std, f_mean + 2*f_std, color='darkgray')
32 plt.fill_between(xs, f_mean - 1*f_std, f_mean + 1*f_std, color='gray')
33 plt.scatter(x, y, c='b')
34 plt.scatter(xtest, ytest, c='r')
35 plt.ylim((300,430))
36 plt.show()

```



More elaborate models are implemented here: https://scikit-learn.org/stable/auto_examples/gaussian_process/plot_gpr_co2.html
and exercise 7 here: https://nbviewer.jupyter.org/github/gpschool/labs/blob/2019/2019/.answers/lab_1.ipynb (however using another GP package)

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.
- [6] Carl E. Rasmussen. *Gaussian Processes for machine learning*. MIT press, 2006.