

Gaussian Processes

Kaiwen Cai

1. Multivariate Gaussian Distribution

Let us look at the multivariate case of Gaussian Distribution (aka., Normal Distribution). Each random variable is distributed randomly and their joint distribution is also a Gaussian distribution. A multivariate Gaussian distribution is defined by its mean $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$.

$$\mathbf{X} = \begin{bmatrix} X_1 \\ X_2 \\ \vdots \\ X_n \end{bmatrix} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$$

where

$$\boldsymbol{\Sigma} = \text{Cov}(X_i, X_j) = \mathbf{E} \left[(X_i - \mu_i)(X_j - \mu_j)^T \right]$$

2. Gaussian Process Regression

Suppose we observe a training set $\mathcal{D} = \{(\mathbf{x}_i, f_i), i = 1, 2, \dots, N\}$, where $f_i = f(\mathbf{x}_i)$. For a test set $\mathcal{D}_* = \{(\mathbf{x}_{i,*}), i = 1, 2, \dots, N_*\}$, we want to predict the corresponding output $f_{i,*}$. We regard the combination of the training set and the test set as a multivariate Gaussian distribution:

$$\begin{bmatrix} \mathbf{f} \\ \mathbf{f}_* \end{bmatrix} \sim \mathcal{N} \left(\begin{bmatrix} \boldsymbol{\mu} \\ \boldsymbol{\mu}_* \end{bmatrix}, \begin{bmatrix} \mathbf{K} & \mathbf{K}_* \\ \mathbf{K}_*^T & \mathbf{K}_{**} \end{bmatrix} \right)$$

where $\mathbf{K} = \kappa(\mathbf{X}, \mathbf{X}) \in \mathbb{R}^{N \times N}$, $\mathbf{K}_* = \kappa(\mathbf{X}, \mathbf{X}_*) \in \mathbb{R}^{N \times N_*}$, $\mathbf{K}_{**} = \kappa(\mathbf{X}_*, \mathbf{X}_*) \in \mathbb{R}^{N_* \times N_*}$, κ is a predefined kernel function (here we adopt a RBF kernel):

$$\kappa(x, x') = \sigma_f^2 \exp \left(-\frac{1}{2\ell^2} (x - x')^2 \right)$$

With the observed training set, we update the test set

$$p(\mathbf{f}_* | \mathbf{X}_*, \mathbf{X}, \mathbf{f}) \sim \mathcal{N}(\boldsymbol{\mu}_*, \boldsymbol{\Sigma}_*)$$

where

$$\begin{aligned} \boldsymbol{\mu}_* &= \boldsymbol{\mu}(\mathbf{X}_*) + \mathbf{K}_*^T \mathbf{K}^{-1} (\mathbf{f} - \boldsymbol{\mu}(\mathbf{X})) \\ \boldsymbol{\Sigma}_* &= \mathbf{K}_{**} - \mathbf{K}_*^T \mathbf{K}^{-1} \mathbf{K}_* \end{aligned}$$

Now we sample from the multivariate distribution $p(\mathbf{f}_* | \mathbf{X}_*, \mathbf{X}, \mathbf{f}) \sim \mathcal{N}(\boldsymbol{\mu}_*, \boldsymbol{\Sigma}_*)$. Recall when we have a univariate Gaussian distribution $x \sim \mathcal{N}(\mu, \sigma^2)$, we sample in a way $x \sim \mu + \sigma \cdot \mathcal{N}(0, 1)$. The equivalent way of sampling from a multivariate distribution is: $\mathbf{f}_* \sim \boldsymbol{\mu} + \mathbf{B} \cdot \mathcal{N}(\mathbf{0}, \mathbf{I})$, where $\mathbf{B}\mathbf{B}^T = \boldsymbol{\Sigma}_*$.

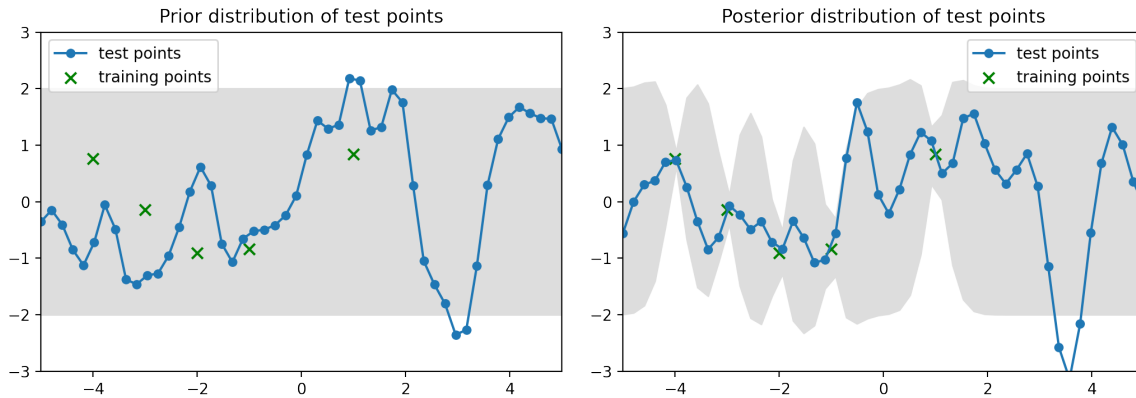


Figure 1. Prior and posterior distribution of Test set

```

1  %%
2  import numpy as np
3  import matplotlib.pyplot as plt
4  np.set_printoptions(precision=3)
5  # plt.style.use('ggplot')
6
7
8  # KERNEL FUNCTION ===== #
9  def kernel(a, b, param):
10     """
11     RBF kernel
12     """
13     sqdist = np.sum(a**2,1).reshape(-1,1) + np.sum(b**2,1) - 2*np.dot(a, b.T)
14     return np.exp(-.5 * (1/param) * sqdist)
15 param = 0.1
16
17
18 # TEST DATA ===== #
19 # PRIOR ===== #
20 n = 50
21 Xtest = np.linspace(-5, 5, n).reshape(-1, 1)
22 K_ss = kernel(Xtest, Xtest, param)
23 stdv = np.diag(K_ss)
24 L = np.linalg.cholesky(K_ss)
25 f_prior = np.dot(L, np.random.normal(size=(n, 1)))
26
27
28 # TRAINING DATA ===== #
29 Xtrain = np.array([-4, -3, -2, -1, 1]).reshape(5, 1)
30 ytrain = np.sin(Xtrain)
31 plt.plot(Xtest, f_prior)
32 plt.gca().fill_between(Xtest.flat,
33                        0 - 2 * stdv,
34                        0 + 2 * stdv,
35                        color="#d4d4d4")
36 plt.scatter(Xtrain, ytrain, s=50, marker='x', c='green')
37 plt.axis([-5, 5, -3, 3])

```

```

38 plt.title('Funtions from the GP prior')
39 plt.savefig('prior.png', bbox_inches='tight', dpi=200)
40 plt.show()
41
42 # POSTERIOR ===== #
43 K = kernel(Xtrain, Xtrain, param)
44 L = np.linalg.cholesky(K)
45
46 K_s = kernel(Xtrain, Xtest, param)
47 Lk = np.linalg.solve(L, K_s) # L^-1 * Ks
48 mu = np.dot(Lk.T, np.linalg.solve(L, ytrain)).reshape((n, ))
49 # = (L^-1 * Ks)^T * L^-1 * f
50 # = Ks^T * L^-T * L^-1 * f
51 # = Ks^T * (L * L^T)^-1 * f
52 # = Ks^T * K^-1 * f
53
54 s2 = np.diag(K_ss) - np.sum(Lk**2, axis=0)
55 # = Kss - (L^-1 * Ks). **2 NOTE How?
56 stdv = np.sqrt(s2)
57
58 L = np.linalg.cholesky(K_ss - np.dot(Lk.T, Lk))
59 # = Kss - (L^-1 * Ks)^T * (L^-1 * Ks)
60 # = Kss - Ks^T * L^-T * L^-1 * Ks
61 # = Kss - Ks^T * (L * L^T)^-1 * Ks
62
63 f_post = mu.reshape(-1, 1) + np.dot(L, np.random.normal(size=(n, 1)))
64
65 plt.plot(Xtest, f_post)
66 plt.gca().fill_between(Xtest.flat,
67                        mu - 2 * stdv,
68                        mu + 2 * stdv,
69                        color="#d4d4d4")
70 plt.scatter(Xtrain, ytrain, s=50, marker='x', c='green')
71 plt.axis([-5, 5, -3, 3])
72 plt.title('Funtions from the GP posterior')
73 plt.savefig('post.png', bbox_inches='tight', dpi=200)
74 plt.show()

```