

Gaussian Processes

General Principles

Through varying intercepts and slopes, we have seen how to quantify some of the unique features that generate variation across clusters and covariance among the observations within each cluster. But through the covariance matrix that is used to account for correlation between clusters, we are inherently assuming linear relationships between clusters. What if we want to model the relationship between two variables that are not linearly related? In this case, we can use a Gaussian Process (GP) to model the relationship between two variables.

Considerations

Caution

- To capture complex, non-linear relationships in data where the underlying function is smooth but has an unknown functional form, GPs use a kernel .
- The choice of kernel hyperparameters can significantly impact results; thus, GPs require choosing an appropriate kernel function that captures the expected behavior of your data.
- Through kernel definition, we can incorporate domain knowledge.
- They scale poorly with dataset size ($O(n^3)$ complexity) due to matrix operations; thus, memory requirements can be substantial for large datasets, which has led to neural networks being used instead to resolve large non-linear problems.

Example

Below is an example code snippet demonstrating Gaussian Process regression using the BayesForge (BF) package. Data consist of a continuous dependent variable (*total_tools*), representing the number of tools invented in the islands, and a continuous independent variable (*population*), representing the population of the islands. The goal is to estimate the effect of population on the total tools. We use the distance matrix of the islands for the

kernel function in order to capture the spatial dependence of the relationship. This example is based on McElreath (2018).

Python

```
from BayesForge import bf
import jax.numpy as jnp
import pandas as pd
# Setup device-----
m = bf(platform='cpu')

# Import Data & Data Manipulation -----
# Import
from importlib.resources import files
data_path = m.load.kline2(only_path=True)
m.data(data_path, sep=';')

data_path2 = files('BayesForge.Resources') / 'islandsDistMatrix.csv'
islandsDistMatrix = pd.read_csv(data_path2, index_col=0)

m.data_to_model(['total_tools', 'population'])
m.data_on_model["society"] = jnp.arange(0,10)# index observations
m.data_on_model["Dmat"] = islandsDistMatrix.values # Distance matrix

def model(Dmat, population, society, total_tools):
    a = m.dist.exponential(1, name = 'a')
    b = m.dist.exponential(1, name = 'b')
    g = m.dist.exponential(1, name = 'g')

    # non-centered Gaussian Process prior
    etasq = m.dist.exponential(2, name = 'etasq')
    rhosq = m.dist.exponential(0.5, name = 'rhosq')
    SIGMA = etasq * jnp.exp(-rhosq * jnp.square(Dmat))
    SIGMA = SIGMA.at[jnp.diag_indices(Dmat.shape[0])].add(0.001)
    k = m.dist.multivariate_normal(0, SIGMA, name = 'k')

    lambda_ = a * population**b / g * jnp.exp(k[society])

    m.dist.poisson(lambda_, obs=total_tools)
```

```
# Run sampler -----
m.fit(model, progress_bar=False)
m.summary()
```

jax.local_device_count 32

	mean	sd	hdi_5.5%	hdi_94.5%	mcse_mean	mcse_sd	ess_bulk	ess_tail	r_hat
a	1.40	1.07	0.04	2.74	0.03	0.02	1311.55	1297.04	1.00
b	0.28	0.09	0.14	0.41	0.00	0.00	1162.80	1121.37	1.00
etasq	0.21	0.22	0.01	0.43	0.01	0.01	902.49	1219.58	1.00
g	0.60	0.57	0.01	1.25	0.02	0.01	1308.38	1491.80	1.00
k[0]	-0.14	0.32	-0.61	0.36	0.01	0.01	768.94	673.50	1.00
k[1]	-0.02	0.31	-0.48	0.46	0.01	0.01	715.49	648.26	1.00
k[2]	-0.05	0.30	-0.53	0.39	0.01	0.01	737.38	671.64	1.01
k[3]	0.37	0.28	-0.05	0.77	0.01	0.01	789.29	699.13	1.01
k[4]	0.10	0.28	-0.32	0.49	0.01	0.01	708.20	524.43	1.01
k[5]	-0.37	0.29	-0.84	0.03	0.01	0.01	830.21	731.78	1.01
k[6]	0.16	0.28	-0.25	0.55	0.01	0.01	742.21	644.47	1.01
k[7]	-0.19	0.28	-0.64	0.21	0.01	0.01	740.76	586.73	1.01
k[8]	0.27	0.27	-0.15	0.63	0.01	0.01	725.69	634.01	1.01
k[9]	-0.15	0.37	-0.77	0.36	0.01	0.01	858.31	743.44	1.01
rhosq	1.26	1.55	0.01	2.98	0.05	0.04	873.77	922.26	1.00

Python (Build in function)

```
from BayesForge import bf
import jax.numpy as jnp
import pandas as pd

# Setup device-----
m = bf(platform="cpu")

# Import Data & Data Manipulation -----
# Import
from importlib.resources import files

data_path = m.load.kline2(only_path=True)
m.data(data_path, sep=";")
```

```

islandsDistMatrix = m.load.islands_dist_matrix(frame=False)["data"]

m.data_to_model(["total_tools", "population"])
m.data_on_model["society"] = jnp.arange(0, 10) # index observations
m.data_on_model["Dmat"] = islandsDistMatrix # Distance matrix

def model(Dmat, population, society, total_tools):
    a = m.dist.exponential(1, name="a")
    b = m.dist.exponential(1, name="b")
    g = m.dist.exponential(1, name="g")

    k = m.gaussian.gaussian_process(Dmat, etasq=2, rhosq=0.5, sigmaq=0.001)

    lambda_ = a * population**b / g * jnp.exp(k[society])

    m.dist.poisson(lambda_, obs=total_tools)

# Run sampler -----
m.fit(model)
m.summary()

```

```
jax.local_device_count 32
```

```
0%|          | 0/2000 [00:00<?, ?it/s]
```

```
0%|          | 0/2000 [00:00<?, ?it/s]
```

```
0%|          | 0/2000 [00:00<?, ?it/s]
```

```
0%|          | 0/2000 [00:00<?, ?it/s]
```

	mean	sd	hdi_5.5%	hdi_94.5%	mcse_mean	mcse_sd	ess_bulk	ess_tail	r_hat
a	1.19	1.06	0.01	2.45	0.03	0.02	1110.13	1047.64	1.0
b	0.32	0.15	0.09	0.56	0.01	0.00	765.61	756.87	1.0
g	0.84	0.87	0.00	1.93	0.03	0.02	756.89	641.80	1.0
kernel[0]	-0.15	0.68	-1.24	0.96	0.03	0.02	392.96	721.41	1.0
kernel[1]	0.14	0.64	-0.87	1.20	0.03	0.02	364.97	578.05	1.0

	mean	sd	hdi_5.5%	hdi_94.5%	mcse_mean	mcse_sd	ess_bulk	ess_tail	r_hat
kernel[2]	0.07	0.63	-0.95	1.07	0.03	0.02	363.15	544.45	1.0
kernel[3]	0.52	0.62	-0.45	1.52	0.03	0.02	349.09	569.19	1.0
kernel[4]	0.13	0.63	-0.87	1.11	0.03	0.02	348.83	489.39	1.0
kernel[5]	-0.46	0.64	-1.53	0.51	0.03	0.02	375.52	507.80	1.0
kernel[6]	0.24	0.63	-0.78	1.21	0.03	0.02	345.04	462.64	1.0
kernel[7]	-0.22	0.65	-1.31	0.74	0.03	0.02	353.64	462.89	1.0
kernel[8]	0.35	0.64	-0.62	1.42	0.03	0.02	347.00	485.94	1.0
kernel[9]	-0.27	0.85	-1.74	0.98	0.04	0.03	442.01	619.11	1.0

R

```

library(BayesForge)
jnp = reticulate::import('jax.numpy')
pd = reticulate::import('pandas')
# setup platform-----
m=importBF(platform='cpu')

# Import data -----
m$data(m$load$kline2(only_path=T), sep=';')
islandsDistMatrix = m$load$islands_dist_matrix(frame = FALSE)$data
m$data_to_model(list('total_tools', 'population'))
m$data_on_model$society = jnp$arange(0,10, dtype='int64')
m$data_on_model$Dmat = jnp$array(islandsDistMatrix)

# Define model -----
model <- function(Dmat, population, society, total_tools){
  a = bf.dist.exponential(1, name = 'a')
  b = bf.dist.exponential(1, name = 'b')
  g = bf.dist.exponential(1, name = 'g')

  # non-centered Gaussian Process prior
  etasq = bf.dist.exponential(2, name = 'etasq')
  rhosq = bf.dist.exponential(0.5, name = 'rhosq')
  k = m$gaussian$gaussian_process(Dmat, etasq, rhosq, 0.01)

  lambda_ = a * population**b / g * jnp$exp(k[society])
  m$dist$poisson(lambda_, obs=total_tools)
}

```

```

# Run MCMC -----
m$fit(model) # Optimize model parameters through MCMC sampling

# Summary -----
m$summary() # Get posterior distribution

```

Julia

```

using BayesForge

# Setup device-----
m = importBF(platform="cpu")

# Import Data & Data Manipulation -----
# Import
data_path = m.load.kline2(only_path = true)
m.data(data_path, sep=";")

islandsDistMatrix = m.load.islands_dist_matrix(frame = false)["data"]
m.data_to_model(["total_tools", "population"])
m.data_on_model["society"] = jnp.arange(0,10)# index observations
m.data_on_model["Dmat"] = jnp.array(islandsDistMatrix) # Distance matrix

# Define model -----
@BF function model(Dmat, population, society, total_tools)
    a = m.dist.exponential(1, name = "a")
    b = m.dist.exponential(1, name = "b")
    g = m.dist.exponential(1, name = "g")

    # non-centered Gaussian Process prior
    etasq = m.dist.exponential(2, name = "etasq")
    rhosq = m.dist.exponential(0.5, name = "rhosq")
    SIGMA = etasq * jnp.exp(-rhosq * jnp.square(Dmat))
    SIGMA = SIGMA.at[jnp.diag_indices(Dmat.shape[0])].add(etasq)
    k = m.dist.multivariate_normal(0, SIGMA, name = "k")

    lambda_ = a * population^b / g * jnp.exp(k[society])

```

```

    m.dist.poisson(lambda_, obs=total_tools)

end

# Run mcmc -----
m.fit(model) # Optimize model parameters through MCMC sampling

# Summary -----
m.summary() # Get posterior distributions

```

Mathematical Details

Formula

The following equation allows us to evaluate the relationship between the dependent variable Y distributed normal, and the independent variable X while incorporating a GP for the effect of variable Q :

$$Y_{[i]} \sim \text{Normal}(\alpha + \beta X_{[i]} + \gamma_{[Q_{[i]}]}, \sigma)$$

where:

- $Y_{[i]}$ is the i -th value for the dependent variable Y .
- α is the intercept term.
- β is the regression coefficient term.
- $X_{[i]}$ is the i -th value for the independent variable X .
- $Q_{[i]}$ is an integer-valued independent variable (e.g., year-of-birth, age, year) for observation i .
- γ is a vector output from a Gaussian process:

$$\gamma \sim \text{MVNormal}(Z, \varsigma\Omega\varsigma)$$

where:

- Z represents the mean vector of the multivariate normal distribution and set to zero .
- ς is a diagonal matrix of standard deviations.
- Ω is a correlation matrix.

- Multiple kernel functions for Ω exist and will be discussed in the [Note\(s\)](#) section. But the most common one is the quadratic kernel:

$$\Omega_{[i,j]} = \eta \exp(-\phi^2 D_{[i,j]}^2)$$

Where:

- η is the maximal correlation.
- ϕ determines the rate of decline.
- $D_{[i,j]}$ is the distance between the i -th and j -th categories.

Bayesian model

In the Bayesian formulation, we define each parameter with priors . We can express a Bayesian version of this GP using the following model:

$$Y_i = \alpha + \beta X_i + \gamma_{Z_i}$$

$$\gamma \sim \text{MVNormal} \left(\begin{pmatrix} 0 \\ \vdots \\ 0 \end{pmatrix}, K \right)$$

$$K_{i,j} = \eta^2 \exp(-p^2 D_{i,j}^2) + \delta_{i,j} \sigma^2$$

$$\alpha \sim \text{Normal}(0, 1)$$

$$\eta^2 \sim \text{HalfCauchy}(0, 1)$$

$$p^2 \sim \text{HalfCauchy}(0, 1)$$

where:

- Y_i is the i -th value for the dependent variable Y .
- α is the intercept term with a prior of $\text{Normal}(0, 1)$.
- β is the regression coefficient term with a prior of $\text{Normal}(0, 1)$.

- X_i is the i-th value for the independent variable X .
- γ_{Z_i} is the Gaussian process i-th value for the independent variable Z .
- γ is the latent function modeled by the GP.
- K_{ij} is the kernel function evaluated at the corresponding points, $K_{ij} = k(Z_i, Z_j)$, with priors of HalfCauchy(0,1) for η^2 and p^2 to ensure positive values.

Notes

i Note

Common kernel functions include:

- *Radial Basis Function* (RBF) or Squared Exponential Kernel:

$$k(x, x') = \sigma^2 \exp\left(-\frac{\|x - x'\|^2}{2l^2}\right)$$

- *Rational Quadratic Kernel*, this kernel is equivalent to adding together many RBF kernels with different length scales:

$$k(x, x') = \sigma^2 \left(1 + \frac{\|x - x'\|^2}{2l^2}\right)^{-\alpha}$$

- *Periodic kernel* allows for modeling functions that repeat themselves exactly:

$$k(x, x') = \sigma^2 \exp\left(-\frac{2 \sin^2(\pi\|x - x'\|/p)}{l^2}\right)$$

- *Locally Periodic Kernel*:

$$k(x, x') = \sigma^2 \exp\left(-\frac{2 \sin^2(\pi\|x - x'\|/p)}{l^2}\right) \exp\left(-\frac{\|x - x'\|^2}{2l^2}\right)$$

- Any slope or intercept in your model can be defined using a Gaussian Process.

Reference(s)

McElreath, Richard. 2018. *Statistical Rethinking: A Bayesian course with examples in R and Stan*. Chapman; Hall/CRC.

<https://www.cs.toronto.edu/~duvenaud/cookbook/>