

1 NNGP Based Models

1.1 Spatial regression model

1.2 Random-effects and Response NNGP models

1.3 Construction of A , D

2 Code NNGP Based Model in Stan

2.1 Intro of simulation data

2.2 Data block in Stan

2.3 Parameter and model block in Stan

2.4 User-defined likelihood function for NNGP models

3 Simulation study

3.1 Response NNGP models in Stan

3.2 Random-effects NNGP models for simulation study

4 Results and Discussion

4.1 Response NNGP model for simulation study

4.2 Random-effects NNGP models for simulation study

4.3 Discussion

References

Nearest Neighbor Gaussian Processes (NNGP) based models in Stan

Lu Zhang

lu.zhang@ucla.edu (<mailto:lu.zhang@ucla.edu>)

11/5/2017

1 NNGP Based Models

Nearest Neighbor Gaussian Processes (NNGP) based models is a family of highly scalable Gaussian processes based models. In brief, NNGP extends the Vecchia's approximation (Vecchia (1988)) to a process using conditional independence given information from neighboring locations. In this section, I will briefly review response and random-effects NNGP models. For more details of NNGP, please refer to Datta et al. (2016).

1.1 Spatial regression model

We envision a spatial regression model at any location s

$$y(s) = m_{\theta}(s) + w(s) + \epsilon(s), \quad \epsilon(s) \stackrel{iid}{\sim} \mathcal{N}(0, \tau^2) \quad (1)$$

where, usually, $m_{\theta}(s) = x(s)^{\top} \beta$ and $w(s)$ is a latent spatial process capturing spatial dependence. Let S be the set of n observed locations. If we model the process $w(s)$ with a Gaussian process $w(s) \sim GP(0, C_{\theta}(\cdot, \cdot))$, then a customary Bayesian hierarchical models for observations on $S = \{s_1, \dots, s_N\}$ can be constructed as

$$p(\theta) \times \mathcal{N}(w(S) | 0, C_{\theta}(S, S)) \times \mathcal{N}(y(S) | m_{\theta}(S) + w(S), \tau^2 I_n) \quad (2)$$

Another implementation of GP is to marginalize over the latent process $w(s)$ and construct the outcome process $y(s)$ directly with a GP. By integrating out $w(s)$, we have a more parsimonious model whose parameters set collapses from $\{\theta, w\}$ to $\{\theta\}$. In a Bayesian setting, θ will be sampled from its posterior distribution

$$p(\theta | y(S)) \propto p(\theta) \times \mathcal{N}(y(S) | m_{\theta}(S), C_{\theta}(S, S) + \tau^2 I_n) \quad (3)$$

To distinguish these two models, we shall call the former as a random-effects model, and the latter as a response model. These two models are referred as the latent variable GP and the marginal likelihood GP in Stan reference manual (STAN DEVELOPMENT TEAM (2017)).

1.2 Random-effects and Response NNGP models

Nearest neighbor Gaussian process (NNGP) provides an alternative to the Gaussian process in the models discussed in the preceding subsection. The likelihoods of two models basing on NNGP derived from the original Gaussian process coincide with the Vecchia's approximation (Vecchia 1988) of the original models. In particular, a random-effects NNGP model has a posterior distribution proportional to

$$p(\theta) \times \mathcal{N}(w(S) \mid 0, C^*) \times \mathcal{N}(y(S) \mid m_\theta(S) + w(S), \tau^2 I_n) \quad (4)$$

where $C^{*-1} = (I - A^*)^\top D^{*-1} (I - A^*)$ is the precision matrix of the latent process $w(s)$ over S . Here A^* is sparse and strictly lower triangular with at most M ($M \ll N$) non-zero entries in each row, and D^* is a diagonal matrix. One can readily calculate the determinant of C^* by the product of the diagonal elements in D^* . The likelihood of $w(S)$ based on precision matrix C^{*-1} serves as a good approximation to the likelihood of $w(s)$ in (2), while the storage and computational burden of the former is linear in N .

A response NNGP model yields posterior distribution:

$$p(\theta \mid y(S)) \propto p(\theta) \times \mathcal{N}(y(S) \mid m_\theta(S), \{C_\theta + \tau^2 I\}^*) \quad (5)$$

where $\{C_\theta + \tau^2 I\}^{*-1} = (I - A)^\top D^{-1} (I - A)$, analogous to random-effects NNGP model, can be treated as an approximation of $\{C_\theta(S, S) + \tau^2 I_n\}^{-1}$. Notice that although one can obtain the response model by integrating out latent process in a random-effects model, the corresponding NNGP model doesn't have this property.

1.3 Construction of A, D

The details of Matrices A, D and two models can be found in Finley et al. (2017). Here we use the response NNGP model to show how to construct matrix A and D . Let $N(s_i)$ be at most M closest points to s_i among the locations indexed less than i . The i th row ($i > 1$) of A has nonzero entries on positions indexed by $N(s_i)$, and the nonzero entries are calculated by

$$A(i, N(s_i)) = C_\theta(s_i, N(s_i))(C_\theta(N(s_i), N(s_i)) + \tau^2 I)^{-1} \quad (6)$$

And the i th element on the diagonal of D satisfies

$$D(i, i) = C_\theta(s_i, s_i) + \tau^2 - C_\theta(s_i, N(s_i))(C_\theta(N(s_i), N(s_i)) + \tau^2 I)^{-1} C_\theta(N(s_i), s_i)$$

These equations are derived from the distribution of $E[y(s_i) \mid y(N(s_i))]$. The nonzero entries in each row of A are precisely the weights obtained by predicting $y(s_i)$, or "kriging", based upon the values of $y(s)$ at neighboring locations, i.e.,

$N(s_i)$. And the diagonal elements in D are the variance of $y(s_i)$ conditioning on its' neighbors in the "past" $y(N(s_i))$.

2 Code NNGP Based Model in Stan

In this section, I will use a simulation data to show how to code NNGP based models efficiently in Stan.

2.1 Intro of simulation data

We generated response Y along with a covariate x at $n = 500$ randomly sited locations in a unit square domain by the following model:

$$y(s) = \beta_0 + x(s)\beta_1 + w(s) + \epsilon(s), \quad \epsilon(s) \sim N(0, \tau^2) \quad (8)$$

where the zero-centered spatial random effect $w(s)$ were sampled from a Gaussian process with a covariance function C_θ specified by exponential:

$$C_\theta(s_i, s_j) = \sigma^2 \exp(-\phi \|s_i - s_j\|), \quad s_i, s_j \in S \quad (9)$$

The predictor x were generated from $N(0, 1)$. The setting of parameters is listed in the code.

```
rmvn <- function(N, mu = 0, V = matrix(1)){
  P <- length(mu)
  if(any(is.na(match(dim(V), P))))
    stop("Dimension problem!")
  D <- chol(V)
  t(matrix(rnorm(N * P), ncol = P) %*% D + rep(mu, rep(N,
P)))
}

set.seed(1234)
N <- 500
coords <- cbind(runif(N), runif(N))
X <- as.matrix(cbind(1, rnorm(N)))
B <- as.matrix(c(1, 5))
sigma.sq <- 2
tau.sq <- 0.1
phi <- 3 / 0.5

D <- as.matrix(dist(coords))
R <- exp(-phi*D)
w <- rmvn(1, rep(0, N), sigma.sq*R)
Y <- rnorm(N, X %*% B + w, sqrt(tau.sq))
```

2.2 Data block in Stan

The following block shows the elements needed for NNGP based models

```

data {
  int<lower=1> N;
  int<lower=1> M;
  int<lower=1> P;
  vector[N] Y;
  matrix[N, P + 1] X;
  int NN_ind[N - 1, M];
  matrix[N - 1, M] NN_dist;
  matrix[N - 1, (M * (M - 1) / 2)] NN_distM;
}

```

Here the design matrix X contains an initial column of 1s, P is the number of regression coefficients, and M is the number of nearest neighbors (maximum number of elements in each row of $n \times n$ sparse matrix A). Notice that we provide three matrices `NN_ind`, `NN_dist` and `NN_distM`:

`NN_ind` is a two-dimensional array of indices whose $i - 1$ th row shows at most M closest points to s_i among the locations indexed less than i .

`NN_dist` is a matrix whose $i - 1$ th row contains the distance of i th location to its selected neighbors.

`NN_distM` is a matrix whose $i - 1$ th row contains the strictly lower triangular part of the distance matrix of the selected neighbors of i th location.

These three matrices are required for constructing the sparse lower triangular matrix A , and the diagonal matrix D . Since they are fixed across the MCMC updates, we recommend user to provide them in the data segment. Next, I will show how to efficiently generate the matrices listed above.

2.2.1 Build neighbor index

File “NNmatrix.R” provides a wrapper function in R that uses package `spNNGP`, which has a fast algorithm of building neighbor index, to generate the required matrices in the Stan data segment. Here `n.omp.threads` indicates the number of threads to use for parallel processing.

```

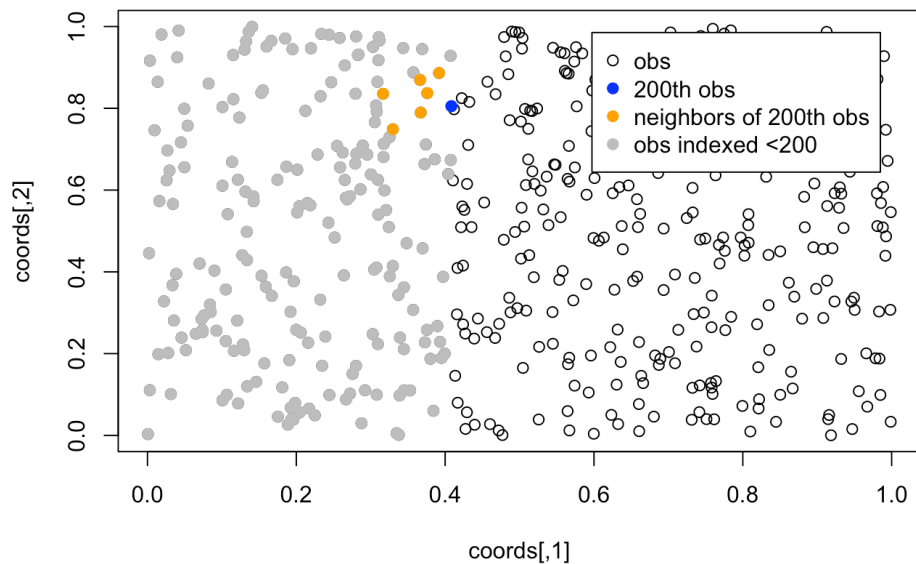
source("NNmatrix.R")
M = 6 # Number of Nearest Neighbors
NN.matrix <- NNMatrix(coords = coords, n.neighbors = M, n
.omp.threads = 2)
str(NN.matrix)

```

2.2.2 Check Neighbors (for fun)

We can use function `Check_Neighbors` in “NNmatrix.R” for checking the nearest neighbor index. It is important to point out that `NNmatrix` sorts coordinates on the first column before building the neighbor index. Thus we should use the sorted response and design matrix instead of the raw data in the data block.

```
Check_Neighbors(NN.matrix$coords.ord, n.neighbors = M, NN
.matrix, ind = 200)
```



2.3 Parameter and model block in Stan

We assign Gaussian priors for σ and τ , an inverse gamma prior for ϕ and a Gaussian prior for $\beta = \{\beta_0, \beta_1\}$. The following is the parameter and model block for a random-effects NNGP model.

```

parameters{
  vector[P + 1] beta;
  real<lower = 0> sigma;
  real<lower = 0> tau;
  real<lower = 0> phi;
  vector[N] w;
}

transformed parameters {
  real sigmasq = sigma^2;
  real tausq = tau^2;
}

model{
  beta ~ multi_normal_cholesky(uB, L_VB);
  phi ~ inv_gamma(ap, bp);
  sigma ~ normal(0, ss);
  tau ~ normal(0, st);
  w ~ nngp_w(sigmasq, phi, NN_dist, NN_distM, NN_ind,
N, M);
  Y ~ normal(X * beta + w, tau);
}

```

A small modification will make the code work for a response NNGP model:

```

parameters{
  vector[P + 1] beta;
  real<lower = 0> sigma;
  real<lower = 0> tau;
  real<lower = 0> phi;
}

transformed parameters {
  real sigmasq = sigma^2;
  real tausq = tau^2;
}

model{
  beta ~ multi_normal_cholesky(uB, L_VB);
  phi ~ inv_gamma(ap, bp);
  sigma ~ normal(0, ss);
  tau ~ normal(0, st);
  Y ~ nngp(X * beta, sigmasq, tausq, phi, NN_dist, NN
_distM, NN_ind, N, M);
}

```

Here, the user-defined functions `nngp_w` and `nngp` will be given in the simulation study section.

2.4 User-defined likelihood function for NNGP models

The hardest part in coding NNGP in Stan is the user-defined likelihood, specifically, the function `nngp_w` and `nngp` in the last subsection. Here we use `nngp` to illustrate the main idea of coding NNGP likelihood.

The log-likelihood of $y(S)$ in (5) is given by:

$$-\frac{1}{2} \sum_{i=1}^N \log D_{ii} - \frac{1}{2} (y(S) - X(S)^\top \beta)^\top (I - A)^\top D^{-1} (I - A) (y(S) - X(S)^\top \beta)$$

In the code below, vector u saves the results of $(I - A)(y(S) - X(S)^\top \beta)$, and vector v saves all the diagonal elements of Matrix D scaled by σ^2 . With this notation, the log-likelihood can be simplified as

$$-\frac{1}{2} \left\{ \sum_{i=1}^N \log V_i + N \log(\sigma^2) + \frac{1}{\sigma^2} U^\top (U \oslash V) \right\} \quad (10)$$

where all the elements in the likelihood are vectors.

In the calculation of vector $U = (I - A)(y(S) - X(S)^\top \beta)$, since we know that matrix A has at most M nonzero elements and the index of nonzero elements is given in `NN_ind`, there is no need for saving the $n \times n$ matrix $I - A$. Instead, we use a for loop to calculate U . Within each iteration, we first use `NN_dist` and `NN_distM` along with the updated parameter to obtain $A(i, N(s_i))$ by (7) and $D(i, i)$ by (6), then use `NN_ind` and $y(S) - X(S)^\top \beta$ to calculate the i th element of U . The flops required in each iteration is in the order of M^3 .

```

functions{
  real nngp_lpdf(vector Y, vector X_beta, real sigmasq,
    real tausq,
    real phi, matrix NN_dist, matrix NN_
    distM, int[, ] NN_ind,
    int N, int M){

    vector[N] V;
    vector[N] Yxb = Y - X_beta;
    vector[N] U = Yxb;
    real kappa_p_1 = tausq / sigmasq + 1;
    int dim;
    int h;
  }

```

```

        for (i in 2:N) {
            matrix[ i < (M + 1) ? (i - 1) : M, i < (M +
1) ? (i - 1): M]
                iNNdistM;
            matrix[ i < (M + 1) ? (i - 1) : M, i < (M +
1) ? (i - 1): M]
                iNNCholL;
            vector[ i < (M + 1) ? (i - 1) : M] iNNcorr;
            vector[ i < (M + 1) ? (i - 1) : M] v;
            row_vector[i < (M + 1) ? (i - 1) : M] v2;
            dim = (i < (M + 1)) ? (i - 1) : M;

            if(dim == 1){iNNdistM[1, 1] = kappa_p_1;}
            else{
                h = 0;
                for (j in 1:(dim - 1)){
                    for (k in (j + 1):dim){
                        h = h + 1;
                        iNNdistM[j, k] = exp(- phi * NN
_distM[(i - 1), h]);
                        iNNdistM[k, j] = iNNdistM[j, k]
;
                    }
                }
                for(j in 1:dim){
                    iNNdistM[j, j] = kappa_p_1;
                }
            }

            iNNCholL = cholesky_decompose(iNNdistM);
            for (j in 1: dim){
                iNNcorr[j] = exp(- phi * NN_dist[(i - 1
), j]);
            }

            v = mdivide_left_tri_low(iNNCholL, iNNcorr);

            V[i] = kappa_p_1 - dot_self(v);

            v2 = mdivide_right_tri_low(v', iNNCholL);

            for (j in 1:dim){
                U[i] = U[i] - v2[j] * YXb[NN_ind[(i - 1
), j]];
            }
        }
        V[1] = kappa_p_1;

```



```

        return - 0.5 * ( 1 / sigmasq * dot_product(U, (
U ./ V)) +
                                sum(log(V)) + N * log(sigmasq))
;
    }
}

```

3 Simulation study

Now let's run the NNGP based models for the simulation data in the last section. First set parameters of priors:

```

P = 1                # number of regression coefficient
s
uB = rep(0, P + 1)  # mean vector in the Gaussian prior
of beta
VB = diag(P + 1)*1000 # covariance matrix in the Gaussian
prior of beta
ss = 3 * sqrt(2)    # scale parameter in the normal prior
of sigma
st = 3 * sqrt(0.1)  # scale parameter in the normal prior
of tau
ap = 3; bp = 0.5    # shape and scale parameter in the
inv-gamma prior of phi

```

3.1 Response NNGP models in Stan

The following chunk is the R code for running response NNGP models. We use the response and design matrix sorted by the order from `NNMatrix` instead of the raw `Y` and `x` in the data block.

```

library(rstan)
options(mc.cores = parallel::detectCores())
data <- list(N = N, M = M, P = P,
             Y = Y[NN.matrix$ord], X = X[NN.matrix$ord, ]
,           # sorted Y and X
             NN_ind = NN.matrix$NN_ind,
             NN_dist = NN.matrix$NN_dist,
             NN_distM = NN.matrix$NN_distM,
             uB = uB, VB = VB, ss = ss, st = st, ap = ap,
bp = bp)

myinits <-list(list(beta = c(1, 5), sigma = 1, tau = 0.5,
phi = 12),
               list(beta = c(5, 5), sigma = 1.5, tau = 0.
2, phi = 5),
               list(beta = c(0, 0), sigma = 2.5, tau = 0.
1, phi = 9))

parameters <- c("beta", "sigmasq", "tausq", "phi")
samples <- stan(
  file = "nngp_response.stan",
  data = data,
  init = myinits,
  pars = parameters,
  iter = 600,
  chains = 3,
  thin = 1,
  seed = 123
)

```

The full Stan program for response NNGP model is in the file “nngp_response.stan”.

```
writeLines(readLines('nngp_response.stan'))
```

```

/* Response NNGP model */

functions{
  real nngp_lpdf(vector Y, vector X_beta, real sigmas
q, real tausq,
                real phi, matrix NN_dist, matrix NN_
distM, int[, ] NN_ind,
                int N, int M){

    vector[N] V;
    vector[N] YXb = Y - X_beta;

```

```

    vector[N] U = YXb;
    real kappa_p_1 = tausq / sigmasq + 1;
    int dim;
    int h;

    for (i in 2:N) {
        matrix[ i < (M + 1) ? (i - 1) : M, i < (M +
1) ? (i - 1): M]
            iNNdistM;
        matrix[ i < (M + 1) ? (i - 1) : M, i < (M +
1) ? (i - 1): M]
            iNNCholL;
        vector[ i < (M + 1) ? (i - 1) : M] iNNcorr;
        vector[ i < (M + 1) ? (i - 1) : M] v;
        row_vector[i < (M + 1) ? (i - 1) : M] v2;
        dim = (i < (M + 1)) ? (i - 1) : M;

        if(dim == 1){iNNdistM[1, 1] = kappa_p_1;}
        else{
            h = 0;
            for (j in 1:(dim - 1)){
                for (k in (j + 1):dim){
                    h = h + 1;
                    iNNdistM[j, k] = exp(- phi * NN
_distM[(i - 1), h]);
                    iNNdistM[k, j] = iNNdistM[j, k]
;
                }
            }
            for(j in 1:dim){
                iNNdistM[j, j] = kappa_p_1;
            }
        }

        iNNCholL = cholesky_decompose(iNNdistM);
        iNNcorr = to_vector(exp(- phi * NN_dist[(i
- 1), 1: dim]));

        v = mdivide_left_tri_low(iNNCholL, iNNcorr);

        V[i] = kappa_p_1 - dot_self(v);

        v2 = mdivide_right_tri_low(v', iNNCholL);

        U[i] = U[i] - v2 * YXb[NN_ind[(i - 1), 1:dim
]];
    }

```

```

        V[1] = kappa_p_1;
        return - 0.5 * ( 1 / sigmasq * dot_product(U, (
U ./ V)) +
                                sum(log(V)) + N * log(sigmasq))
;
    }
}

data {
    int<lower=1> N;
    int<lower=1> M;
    int<lower=1> P;
    vector[N] Y;
    matrix[N, P + 1] X;
    int NN_ind[N - 1, M];
    matrix[N - 1, M] NN_dist;
    matrix[N - 1, (M * (M - 1) / 2)] NN_distM;
    vector[P + 1] uB;
    matrix[P + 1, P + 1] VB;
    real ss;
    real st;
    real ap;
    real bp;
}

transformed data {
    cholesky_factor_cov[P + 1] L_VB;
    L_VB = cholesky_decompose(VB);
}

parameters{
    vector[P + 1] beta;
    real<lower = 0> sigma;
    real<lower = 0> tau;
    real<lower = 0> phi;
}

transformed parameters {
    real sigmasq = sigma^2;
    real tausq = tau^2;
}

model{
    beta ~ multi_normal_cholesky(uB, L_VB);
    phi ~ gamma(ap, bp);
    sigma ~ normal(0, ss);
    tau ~ normal(0, st);
}

```

```

      Y ~ nngp(X * beta, sigmasq, tausq, phi, NN_dist, NN
_distM, NN_ind, N, M);
    }

```

3.2 Random-effects NNGP models for simulation study

The following chunk is the R code for running Random-effects NNGP models:

```

options(mc.cores = parallel::detectCores())
data <- list(N = N, M = M, P = P,
             Y = Y[NN.matrix$ord], X = X[NN.matrix$ord, ]
,           # sorted Y and X
             NN_ind = NN.matrix$NN_ind,
             NN_dist = NN.matrix$NN_dist,
             NN_distM = NN.matrix$NN_distM,
             uB = uB, VB = VB, ss = ss, st = st, ap = ap,
             bp = bp)

myinits <-list(list(beta = c(1, 5), sigma = 1, tau = 0.5,
                    phi = 12,
                    w_b1 = rep(0, N)),
               list(beta = c(5, 5), sigma = 1.5, tau = 0.
2, phi = 5,
                    w_b1 = rep(0.1, N)),
               list(beta = c(0, 0), sigma = 2.5, tau = 0.
1, phi = 9 ,
                    w_b1 = rep(0, N)))

parameters <- c("beta", "sigmasq", "tausq", "phi", "w")
samples_w <- stan(
  file = "nngp_random.stan",
  data = data,
  init = myinits,
  pars = parameters,
  iter = 600,
  chains = 3,
  thin = 1,
  seed = 123
)

```

The full Stan program for random-effects NNGP model is in the file “nngp_random.stan”.

```

writeLines(readLines('nngp_random.stan'))

```

```

/* Random-effects NNGP models*/

functions{
  real nngp_w_lpdf(vector w, real sigmasq, real phi,
matrix NN_dist,
                    matrix NN_distM, int[, ] NN_ind, in
t N, int M){

  vector[N] V;
  vector[N] I_Aw = w;
  int dim;
  int h;

  for (i in 2:N) {

    matrix[ i < (M + 1)? (i - 1) : M, i < (M +
1)? (i - 1): M]
      iNNdistM;
    matrix[ i < (M + 1)? (i - 1) : M, i < (M +
1)? (i - 1): M]
      iNNCholL;
    vector[ i < (M + 1)? (i - 1) : M] iNNcorr;
    vector[ i < (M + 1)? (i - 1) : M] v;
    row_vector[i < (M + 1)? (i - 1) : M] v2;

    dim = (i < (M + 1))? (i - 1) : M;

    if(dim == 1){iNNdistM[1, 1] = 1;}
    else{
      h = 0;
      for (j in 1:(dim - 1)){
        for (k in (j + 1):dim){
          h = h + 1;
          iNNdistM[j, k] = exp(- phi * NN
_distM[(i - 1), h]);
          iNNdistM[k, j] = iNNdistM[j, k]
;
        }
      }
      for(j in 1:dim){
        iNNdistM[j, j] = 1;
      }
    }

    iNNCholL = cholesky_decompose(iNNdistM);
    iNNcorr = to_vector(exp(- phi * NN_dist[(i

```

```

- 1), 1: dim]));

        v = mdivide_left_tri_low(iNNCholL, iNNcorr)
;

        V[i] = 1 - dot_self(v);

        v2 = mdivide_right_tri_low(v', iNNCholL);

        I_Aw[i] = I_Aw[i] - v2 * w[NN_ind[(i - 1),
1:dim]];

    }
    V[1] = 1;
    return - 0.5 * ( 1 / sigmasq * dot_product(I_Aw
, (I_Aw ./ V)) +
                                sum(log(V)) + N * log(sigmasq))
;
}
}

data {
  int<lower=1> N;
  int<lower=1> M;
  int<lower=1> P;
  vector[N] Y;
  matrix[N, P + 1] X;
  int NN_ind[N - 1, M];
  matrix[N - 1, M] NN_dist;
  matrix[N - 1, (M * (M - 1) / 2)] NN_distM;
  vector[P + 1] uB;
  matrix[P + 1, P + 1] VB;
  real ss;
  real st;
  real ap;
  real bp;
}

transformed data {
  cholesky_factor_cov[P + 1] L_VB;
  L_VB = cholesky_decompose(VB);
}

parameters{
  vector[P + 1] beta;
  real<lower = 0> sigma;

```

```

    real<lower = 0> tau;
    real<lower = 0> phi;
    vector[N] w;
  }

  transformed parameters {
    real sigmasq = sigma^2;
    real tausq = tau^2;
  }

  model{
    beta ~ multi_normal_cholesky(uB, L_VB);
    phi ~ gamma(ap, bp);
    sigma ~ normal(0, ss);
    tau ~ normal(0, st);
    w ~ nngp_w(sigmasq, phi, NN_dist, NN_distM, NN_ind,
N, M);
    Y ~ normal(X * beta + w, tau);
  }

```

4 Results and Discussion

In this section, we will show the results of the simulation study, compare response and random-effects NNGP models, and provide suggestions on how to use NNGP based models.

4.1 Response NNGP model for simulation study

Response NNGP model is faster and easier to sample from posterior distribution than the random-effects NNGP models. The following shows the summary table and trace plot of the posterior samples from response NNGP model.

```
print(samples)
```



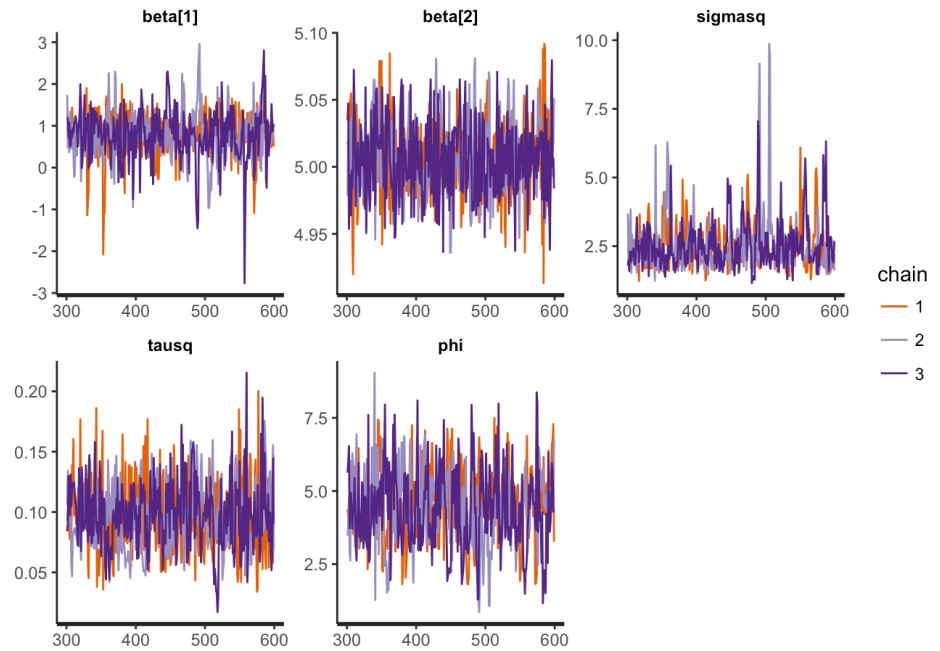
```
Inference for Stan model: nngp_response.
3 chains, each with iter=600; warmup=300; thin=1;
post-warmup draws per chain=300, total post-warmup draws=
900.
```

	mean	se_mean	sd	2.5%	25%	50%	75%
97.5% n_eff							
beta[1]	0.77	0.03	0.58	-0.44	0.47	0.80	1.12
1.91	358						
beta[2]	5.01	0.00	0.03	4.95	4.99	5.01	5.02
5.06	756						
sigmasq	2.49	0.05	0.92	1.49	1.94	2.28	2.79
4.91	282						
tausq	0.10	0.00	0.03	0.05	0.08	0.10	0.12
0.16	488						
phi	4.47	0.07	1.30	1.91	3.57	4.47	5.31
6.98	327						
lp__	-99.27	0.11	1.79	-103.93	-100.13	-98.86	-98.02
-97.03	282						
Rhat							
beta[1]	1.00						
beta[2]	1.00						
sigmasq	1.00						
tausq	1.01						
phi	1.01						
lp__	1.01						

Samples were drawn using NUTS(diag_e) at Tue Nov 28 17:11:10 2017.

For each parameter, n_eff is a crude measure of effective sample size, and Rhat is the potential scale reduction factor on split chains (at convergence, Rhat=1).

```
stan_trace(samples)
```



4.2 Random-effects NNGP models for simulation study

The following gives the summary table posterior samples and trace plots of the MCMC Chains from random-effects NNGP model:

```
print(samples_w, pars = c("beta", "sigmasq", "tausq", "phi", "w[1]", "w[2]", "w[3]", "w[4]"))
```

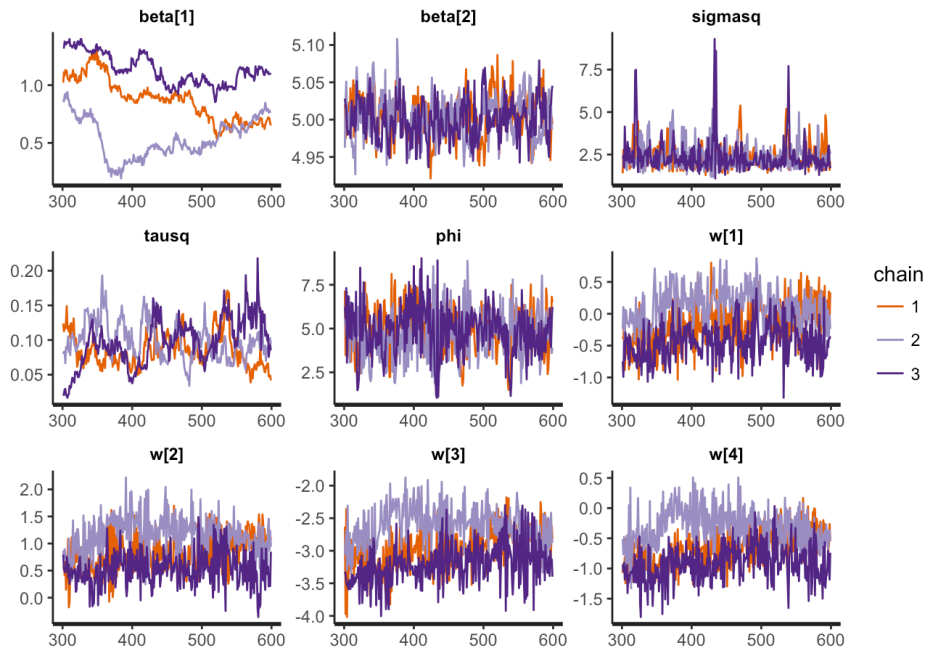
```
Inference for Stan model: nngp_random.
3 chains, each with iter=600; warmup=300; thin=1;
post-warmup draws per chain=300, total post-warmup draws=
900.
```

	mean	se_mean	sd	2.5%	25%	50%	75%	97.5%
n_eff								
Rhat								
beta[1]	0.85	0.20	0.29	0.31	0.64	0.89	1.09	1.34
2	2.44							
beta[2]	5.00	0.00	0.03	4.95	4.98	5.00	5.02	5.06
272	1.00							
sigmasq	2.37	0.04	0.82	1.49	1.89	2.17	2.64	4.29
357	1.01							
tausq	0.09	0.01	0.03	0.04	0.07	0.09	0.11	0.15
33	1.13							
phi	4.79	0.08	1.34	2.15	3.90	4.79	5.66	7.40
274	1.01							
w[1]	-0.17	0.20	0.38	-0.90	-0.43	-0.18	0.12	0.59
4	1.32							
w[2]	0.84	0.23	0.42	0.06	0.52	0.83	1.16	1.63
3	1.34							
w[3]	-2.94	0.21	0.39	-3.69	-3.23	-2.94	-2.65	-2.20
3	1.36							
w[4]	-0.62	0.22	0.40	-1.33	-0.91	-0.64	-0.35	0.17
3	1.42							

```
Samples were drawn using NUTS(diag_e) at Tue Nov 28 17:18
:17 2017.
```

```
For each parameter, n_eff is a crude measure of effective
sample size,
and Rhat is the potential scale reduction factor on split
chains (at
convergence, Rhat=1).
```

```
stan_trace(samples_w, pars = c("beta", "sigmasq", "tausq"
, "phi", "w[1]",
" w[2]", "w[3]", "w[4]"))
```



It is not surprising to see a slower convergence rate and worse mixing of the MCMC Chains. Response NNGP model marginalizes out the spatial effects w , yields a lower-dimensional parameter space, hence drastically improves the posterior geometry. While the parameters to be estimated in a random-effects NNGP model $\{\theta, w\}$ are highly correlated, and the number of parameters is on the scale of the number of observations. Thus the convergence rate of MCMC chains from a random-effects NNGP is slow because of the high correlation and dimension of the parameter space.

Notice the trace plot of random-effects w are highly correlated with the intercept, we modified the code and make the random-effects $w(s)$ centered at the intercept. The code of modified random-effects NNGP model can be found in “nngp.R” and “nngp_random_b1.stan”. Here we suppress the details of the code and show the results directly:

```
print(samples_wb1, pars = c("beta", "sigmasq", "tausq", "
  phi", "w_b1[1]",
                                "w_b1[2]", "w_b1[3]", "w_b1[4]
  ]"))
```

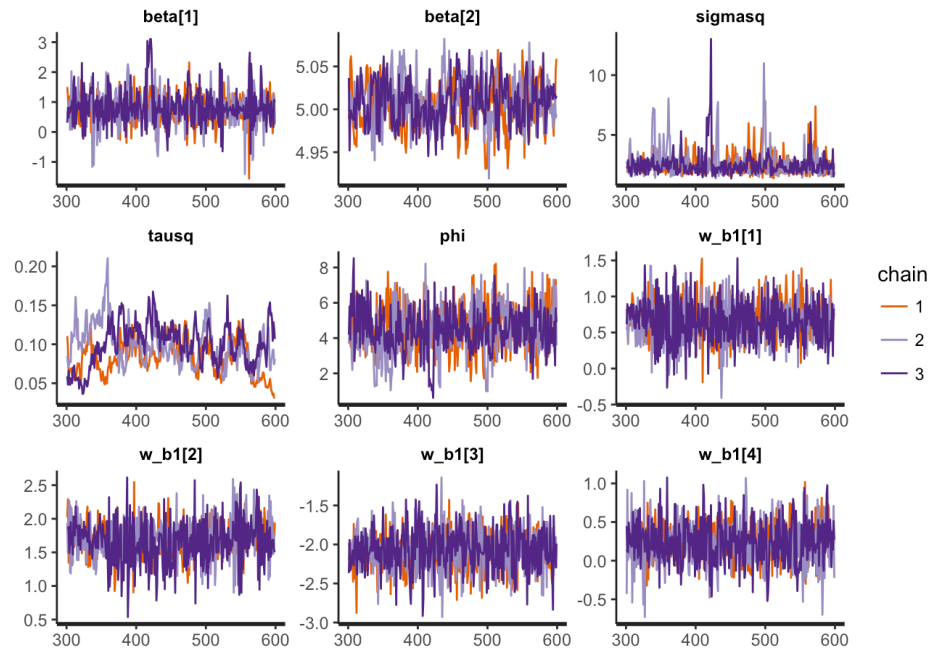
```
Inference for Stan model: nngp_random_b1.
3 chains, each with iter=600; warmup=300; thin=1;
post-warmup draws per chain=300, total post-warmup draws=
900.
```

	mean	se_mean	sd	2.5%	25%	50%	75%	97.5%
n_eff								
Rhat								
beta[1]	0.79	0.02	0.55	-0.28	0.46	0.76	1.08	2.01
900	1.00							
beta[2]	5.01	0.00	0.03	4.95	4.99	5.01	5.03	5.06
333	1.01							
sigmasq	2.56	0.06	1.05	1.54	1.94	2.31	2.82	5.50
323	1.01							
tausq	0.09	0.01	0.03	0.05	0.07	0.09	0.11	0.15
24	1.09							
phi	4.50	0.07	1.35	1.73	3.61	4.44	5.45	7.14
407	1.01							
w_b1[1]	0.67	0.01	0.29	0.11	0.47	0.66	0.86	1.24
900	1.01							
w_b1[2]	1.69	0.01	0.30	1.04	1.51	1.69	1.89	2.27
900	1.00							
w_b1[3]	-2.08	0.01	0.27	-2.61	-2.25	-2.07	-1.90	-1.57
900	1.00							
w_b1[4]	0.23	0.01	0.27	-0.27	0.05	0.24	0.40	0.78
900	1.00							

```
Samples were drawn using NUTS(diag_e) at Tue Nov 28 17:22
:49 2017.
```

```
For each parameter, n_eff is a crude measure of effective
sample size,
and Rhat is the potential scale reduction factor on split
chains (at
convergence, Rhat=1).
```

```
stan_trace(samples_wb1, pars = c("beta", "sigmasq", "tausq",
"phi", "w_b1[1]",
"w_b1[2]", "w_b1[3]", "w_b1[4]"))
```



4.3 Discussion

We recommend a response NNGP model for a large scale data analysis when the study focuses on the inference of parameter set θ . On the other hand, random-effects NNGP models are preferred when the study needs the recovery of latent process $w(s)$. However, the convergence rate of the MCMC Chains from random-effects model could be prohibitively slow when the sample size is large, so we only recommend coding random-effects NNGP model in Stan when the sample size is small. For recovering latent process when the sample size is large, Package spNNGP provides an algorithm for random-effects NNGP model, which implements a sequential Gibbs sampler for updating the latent process. Conjugate NNGP models are also good options for recovering latent process $w(s)$. More details of NNGP based models can be found in Finley et al. (2017)

References

- Datta, Abhirup, Sudipto Banerjee, Andrew O Finley, and Alan E Gelfand. 2016. “Hierarchical Nearest-Neighbor Gaussian Process Models for Large Geostatistical Datasets” 111. Taylor & Francis: 800–812.
- Finley, Andrew O, Abhirup Datta, Bruce C Cook, Douglas C Morton, Hans E Andersen, and Sudipto Banerjee. 2017. “Applying Nearest Neighbor Gaussian Processes to Massive Spatial Data Sets Forest Canopy Height Prediction Across Tanana Valley Alaska.”
- STAN DEVELOPMENT TEAM. 2017. “Stan Modeling Language User’s Guide and Reference Manual.”

Vecchia, A. V. 1988. "Estimation and Model Identification for Continuous Spatial Processes."