

HMSC-R 3.0: Getting started with Hmsc-R: spatial models

Gleb Tikhonov Øystein H. Opedal Nerea Abrego Aleksi Lehikoinen
Melinda M. J. de Jonge Jari Oksanen Otso Ovaskainen

26 August 2025

Introduction

The Hierarchical Modelling of Species Communities (HMSC) framework is a statistical framework for analysis of multivariate data, typically from species communities. We assume that the reader has already gone through the vignette “Hmsc-R 3.0: Getting started with Hmsc-R: univariate models” and “Hmsc-R 3.0: Getting started with Hmsc-R: low-dimensional multivariate models”. In the first vignette we shortly discussed how to fit spatially explicit models to the univariate case. Here, we continue to demonstrate how to use HMSC to make spatially explicit models for the multivariate case and for large datasets.

To get Hmsc-R in use, you need to load it and other needed packages.

```
library(Hmsc)
library(MASS)
set.seed(6)
```

We also set the random number seed to make the results presented here reproducible.

Generating simulated data

To illustrate how to use spatial models in Hmsc, we generate data for 5 species (**ns**) on 100 sampling units (**n**). We include only one environmental predictor (**x1**) and give the true intercept (**alpha**) and slope (**beta1**) parameters to construct the matrix of the linear predictors. So far, this is similar to what we did in the low-dimensional multivariate case, but note that we now call this linear predictor **Lf** instead of **L**. This indicates that this is the fixed effects part of the linear predictor, i.e. the effect that can be explained by the environmental covariates. In addition to the effect from the environmental covariates, we now add spatially structured residuals which follow the same spatial structure for all species, indicated by **Lr**. To generate these spatial residuals, we first simulate some random x and y coordinates for all sampling units (**xycoords**). We then generate a spatially structured latent predictor (**eta1**) using an exponentially decreasing spatial covariance function where we have set the spatial scale parameter (**alpha**) to 0.35. Next, we set the true slope parameters (**lambda1**) for the species responses to the latent predictor and compute the spatial residuals for each species in **Lr**.

```
n = 100
ns = 5
beta1 = c(-2,-1,0,1,2)
alpha = rep(0,ns)
beta = cbind(alpha,beta1)
x = cbind(rep(1,n),rnorm(n))
```

```

Lf = x%*%t(beta)

xycoords = matrix(runif(2*n),ncol=2)
colnames(xycoords) = c("x-coordinate","y-coordinate")
rownames(xycoords) = 1:n

sigma.spatial = c(2)
alpha.spatial = c(0.35)
Sigma = sigma.spatial^2*exp(-as.matrix(dist(xycoords))/alpha.spatial)
eta1 = mvrnorm(mu=rep(0,n), Sigma=Sigma)
lambda1 = c(1,2,-2,-1,0)
Lr = eta1%*%t(lambda1)
L = Lf + Lr
y = as.matrix(L + matrix(rnorm(n*ns),ncol=ns))
yprob = 1*((L +matrix(rnorm(n*ns),ncol=ns))>0)
XData = data.frame(x1=x[,2])

```

We can now visualize the species response matrix y as function of the x and y coordinates (Fig. 1). This shows that indeed nearby sampling units have similar responses for species that have a non-zero loading (λ) to the spatially structured latent variable.

```

rbPal = colorRampPalette(c('cyan','red'))
par(mfrow=c(2,3))
Col = rbPal(10)[as.numeric(cut(x[,2],breaks = 10))]
plot(xycoords[,2],xycoords[,1],pch = 20,col = Col,main=paste('x'), asp=1)
for(s in 1:ns){
  Col = rbPal(10)[as.numeric(cut(y[,s],breaks = 10))]
  plot(xycoords[,2],xycoords[,1],pch = 20,col = Col,main=paste('Species',s), asp=1)
}

```

A spatially explicit model in Hmsc

To fit a spatially explicit model with Hmsc, we construct the random effect using the `sData` input argument where we give the coordinates of the sampling units.

```

studyDesign = data.frame(sample = as.factor(1:n))
rL.spatial = HmscRandomLevel(sData = xycoords)
rL.spatial = setPriors(rL.spatial,nfMin=1,nfMax=1) #We limit the model to one latent variables for visu
m.spatial = Hmsc(Y=yprob, XData=XData, XFormula=~x1,
studyDesign=studyDesign, ranLevels=list("sample"=rL.spatial),distr="probit")

```

Model fitting and evaluation of explanatory and predictive power can be done as before. We first set the MCMC sampling parameters.

```

nChains = 2
test.run = FALSE
if (test.run){
  # with this option, the vignette runs fast but results are not reliable
  thin = 1
  samples = 10
}

```

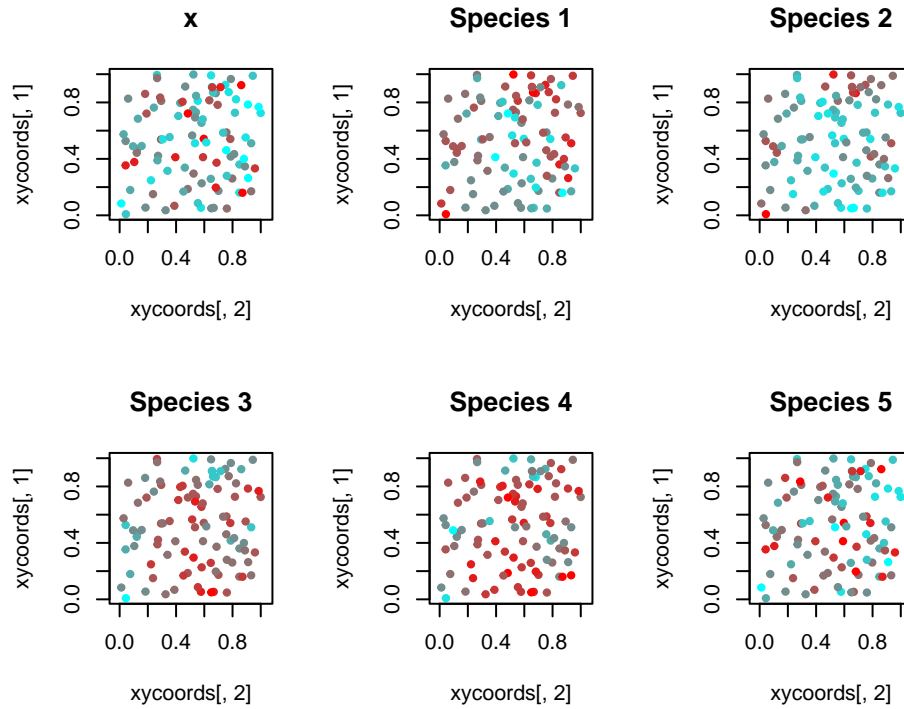


Figure 1: Plots of simulated spatially structured data.

```

    transient = 5
    verbose = 0
  } else {
    # with this option, the vignette evaluates slow but it reproduces the results of
    # the .pdf version
    thin = 10
    samples = 1000
    transient = 1000
    verbose = 0
  }

m.spatial = sampleMcmc(m.spatial, thin = thin, samples = samples, transient = transient,
                      nChains = nChains, nParallel = nChains, verbose = verbose,
                      updater=list(GammaEta=FALSE))

```

The explanatory and predictive power of the model can now be calculated in the same way as before.

```

#Explanatory power
preds.spatial = computePredictedValues(m.spatial)
MF.spatial = evaluateModelFit(hM=m.spatial, predY=preds.spatial)
MF.spatial

## $RMSE
## [1] 0.2723915 0.1787932 0.2355696 0.2512485 0.3014344
##
## $AUC
## [1] 0.9642857 0.9952575 0.9728752 0.9706381 0.9504831

```

```
##
## $TjurR2
## [1] 0.5611244 0.6638560 0.5500152 0.5509740 0.6186115
```

```
#Predictive power
partition = createPartition(m.spatial, nfolds = 2, column = "sample")
cvpreds.spatial = computePredictedValues(m.spatial, partition=partition,
  nParallel = nChains, updater=list(GammaEta=FALSE))
```

```
## Cross-validation, fold 1 out of 2
## Cross-validation, fold 2 out of 2
```

```
cvMF.spatial = evaluateModelFit(hM=m.spatial, predY=cvpreds.spatial)
cvMF.spatial
```

```
## $RMSE
## [1] 0.3512856 0.3329677 0.3582368 0.3581119 0.3147683
##
## $AUC
## [1] 0.8980952 0.8570461 0.8125377 0.8599661 0.9432367
##
## $TjurR2
## [1] 0.3643824 0.2485915 0.2329053 0.2492040 0.5774692
```

As the model includes a spatially structured random effect, its predictive power is based on both the fixed and the random effects. Concerning the latter, the model can utilize observed data from nearby sampling units included in model fitting when predicting the response for a focal sampling unit that is not included in model fitting.

The estimated spatial scale of the random effect is given by the parameter `Alpha[[1]]`. Let's have a look at the MCMC trace plot for this parameter (Fig. 3).

```
mpost.spatial = convertToCodaObject(m.spatial)
plot(mpost.spatial$Alpha[[1]])
```

```
summary(mpost.spatial$Alpha[[1]])
```

```
##
## Iterations = 1010:11000
## Thinning interval = 10
## Number of chains = 2
## Sample size per chain = 1000
##
## 1. Empirical mean and standard deviation for each variable,
##    plus standard error of the mean:
##
##           Mean           SD      Naive SE Time-series SE
##      0.291687      0.102924      0.002301      0.004570
##
## 2. Quantiles for each variable:
##
##      2.5%    25%    50%    75%   97.5%
## 0.1538 0.2237 0.2657 0.3356 0.5453
```

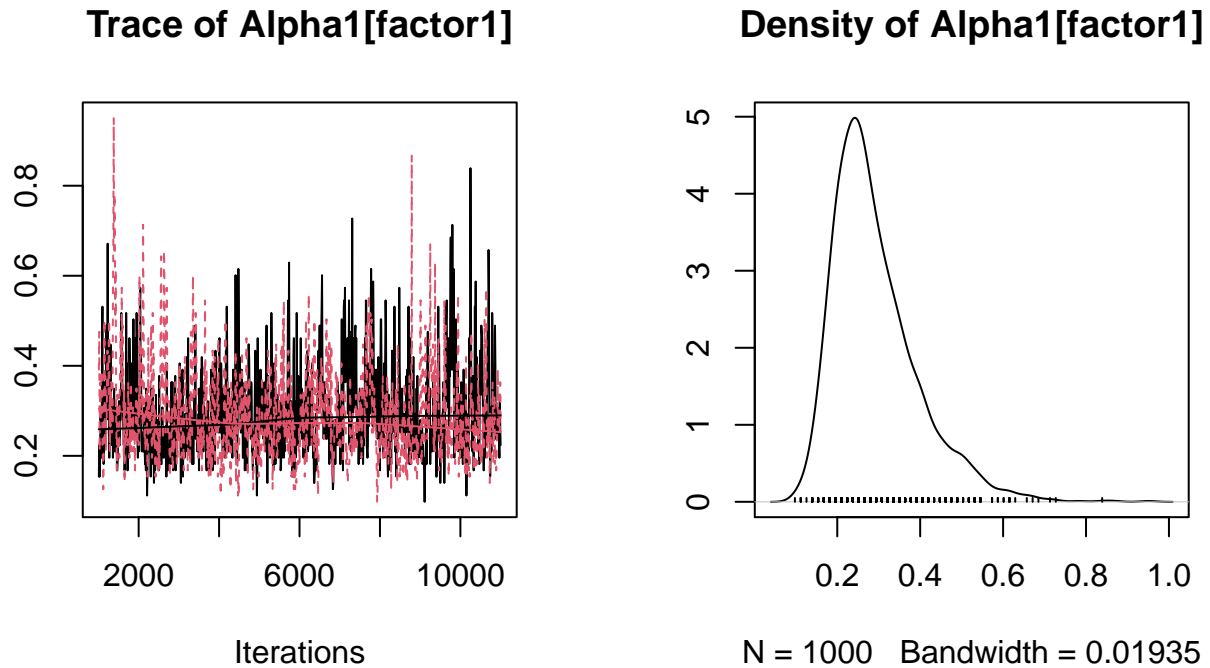


Figure 2: Posterior trace plot of the spatial scale parameter of the spatial model.

For comparison, let us fit a non-spatial model to the same data.

```
m = Hmsc(Y=yprob, XData=XData, XFormula=~x1, studyDesign = studyDesign, distr="probit")
m = sampleMcmc(m, thin = thin, samples = samples, transient = transient,
               nChains = nChains, nParallel = nChains, verbose = verbose)
preds = computePredictedValues(m)
MF = evaluateModelFit(hM=m, predY=preds)
MF
```

```
## $RMSE
## [1] 0.3946252 0.3797230 0.4071993 0.4099658 0.3018431
##
## $AUC
## [1] 0.8133333 0.6029810 0.5346594 0.6357990 0.9504831
##
## $TjurR2
## [1] 0.2607040533 0.0270019929 0.0007665274 0.0491697035 0.6147594362
```

```
partition = createPartition(m, nfolds = 2, column = "sample")
preds = computePredictedValues(m, partition=partition, nParallel = nChains)
```

```
## Cross-validation, fold 1 out of 2
```

```
## setting updater$GammaEta=FALSE due to absence of random effects included to the model
```

```
## Cross-validation, fold 2 out of 2
```

```
## setting updater$GammaEta=FALSE due to absence of random effects included to the model
```

```
MF = evaluateModelFit(hM=m, predY=preds)
MF
```

```
## $RMSE
## [1] 0.3991611 0.3815012 0.4131798 0.4110814 0.3408304
##
## $AUC
## [1] 0.8047619 0.5901084 0.4442435 0.6363636 0.9239130
##
## $TjurR2
## [1] 0.238276454 0.025757870 -0.007983475 0.048322622 0.563647450
```

We observe that both the explanatory and the predictive power is lower than for the model with a spatial random effect for those species that had spatially structured responses.

Spatial models for big datasets

For large datasets, i.e. those with > 1000 sampling units, the standard spatial models as described above may become computationally infeasible. For such datasets, we implemented two alternative approaches to account for the spatial structure in the data: the ‘Nearest Neighbour Gaussian Process (NNGP)’ and the ‘Gaussian Predictive Process (GPP)’. The details of this method are described in Tikhonov et al. (2020).

NNGP models

If we want to fit a NNGP model we have to set `sMethod` to ‘NNGP’ when constructing the random level. Additionally, we can specify how many neighbours we want to use by setting the `nNeighbours` argument. When we do not explicitly set `nNeighbours` when constructing the random level, this parameter is set to 10 as a standard.

```
rL.nngp = HmscRandomLevel(sData = xycoords, sMethod = 'NNGP', nNeighbours = 10)
rL.nngp = setPriors(rL.nngp, nfMin=1, nfMax=1)
```

Running the model and checking the fit is done in the same way as earlier.

```
m.nngp = Hmsc(Y=yprob, XData=XData, XFormula=~x1,
              studyDesign=studyDesign, ranLevels=list("sample"=rL.nngp), distr="probit")
m.nngp = sampleMcmc(m.nngp, thin = thin, samples = samples, transient = transient,
                  nChains = nChains, nParallel = nChains, verbose = verbose,
                  updater=list(GammaEta=FALSE))
preds.nngp = computePredictedValues(m.nngp, updater=list(GammaEta=FALSE))
MF.nngp = evaluateModelFit(hM=m.nngp, predY=preds.nngp)
MF.nngp
```

```
## $RMSE
## [1] 0.2720166 0.1707469 0.2417962 0.2564901 0.3013351
##
## $AUC
## [1] 0.9633333 0.9959350 0.9680530 0.9666855 0.9504831
##
## $TjurR2
## [1] 0.5642137 0.6910263 0.5416389 0.5466876 0.6199860
```

```
partition = createPartition(m.nngp, nfolds = 2, column = "sample")
cvpreds.nngp = computePredictedValues(m.nngp, partition=partition,
  nParallel = nChains, updater=list(GammaEta=FALSE))
```

```
## Cross-validation, fold 1 out of 2
## Cross-validation, fold 2 out of 2
```

```
cvMF.nngp = evaluateModelFit(hM=m.nngp, predY=cvpreds.nngp)
cvMF.nngp
```

```
## $RMSE
## [1] 0.3492980 0.3055479 0.3352653 0.3434148 0.3076929
##
## $AUC
## [1] 0.8914286 0.9112466 0.8704039 0.8938453 0.9460548
##
## $TjurR2
## [1] 0.3764224 0.2954774 0.2484652 0.2741600 0.5864402
```

Let's have a look at the MCMC trace plots for the spatial scale parameter `Alpha[[1]]`.

```
mpost.nngp = convertToCodaObject(m.nngp)
plot(mpost.nngp$Alpha[[1]])
```

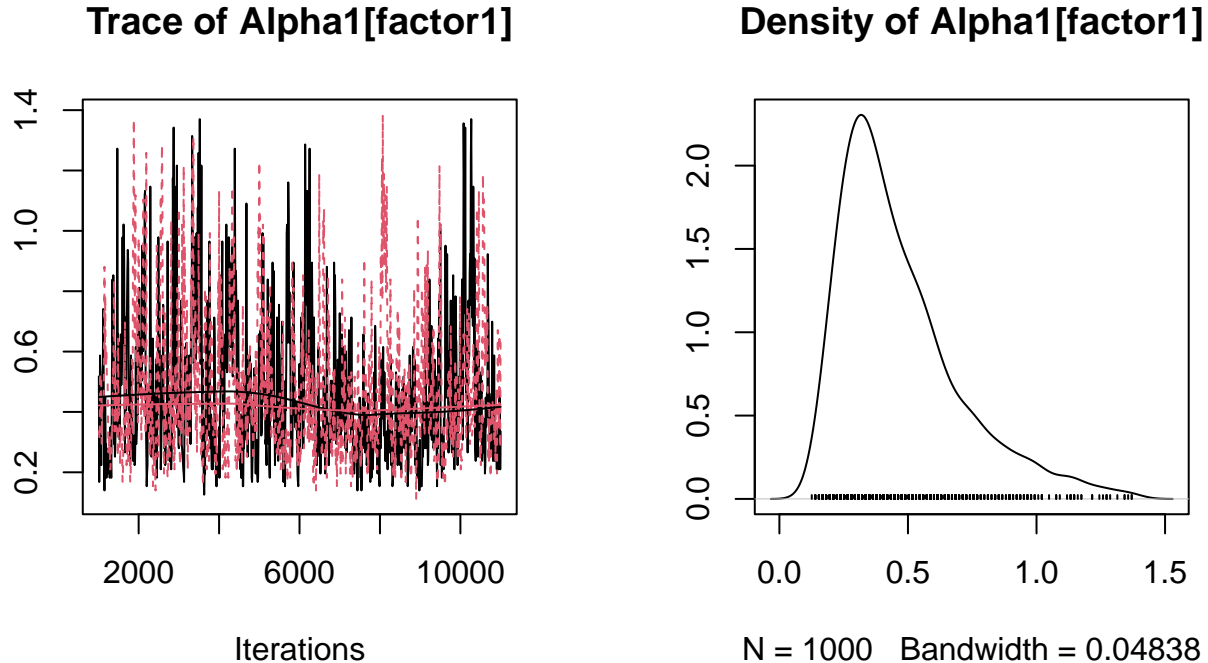


Figure 3: Posterior trace plot of the spatial scale parameter of the nngp spatial model.

```
summary(mpost.nngp$Alpha[[1]])
```

```
##
## Iterations = 1010:11000
## Thinning interval = 10
## Number of chains = 2
## Sample size per chain = 1000
##
## 1. Empirical mean and standard deviation for each variable,
##    plus standard error of the mean:
##
##           Mean           SD      Naive SE Time-series SE
##      0.467748      0.234661      0.005247      0.012961
##
## 2. Quantiles for each variable:
##
##   2.5%   25%   50%   75%  97.5%
## 0.1818 0.2936 0.4055 0.5733 1.1047
```

GPP models

The Gaussian Predictive Process (GPP) assumes that all information on the spatial structure of the data can be summarized at a small number of so called ‘knot’ locations. The locations of these knots have to be specified by the user. To help the user with this, we implemented a function to construct a uniform grid of knots based on the locations of the dataset. This function either need the wanted number of knots along the shortest spatial axis `nKnots` or the wanted distance between knots `KnotDist`. Additionally, the user can specify the maximum distance of a knot to the nearest data point, this ensures that the created grid does contain knots in locations with no datapoints.

```
# Setting the knots
Knots = constructKnots(xycoords, knotDist = 0.2, minKnotDist = 0.4)

plot(xycoords[,1],xycoords[,2],pch=18, asp=1)
points(Knots[,1],Knots[,2],col='red',pch=18)
```

We can now construct the random level by by setting `sMethod` to ‘GPP’ and supplying the knot locations to `sKnot`.

```
rL.gpp = HmscRandomLevel(sData = xycoords, sMethod = 'GPP', sKnot = Knots)
rL.gpp = setPriors(rL.gpp,nfMin=1,nfMax=1)
m.gpp = Hmsc(Y=yprob, XData=XData, XFormula=~x1,
             studyDesign=studyDesign, ranLevels=list("sample"=rL.gpp),distr="probit")
m.gpp = sampleMcmc(m.gpp, thin = thin, samples = samples, transient = transient,
                  nChains = nChains, nParallel = nChains, verbose = verbose,
                  updater=list(GammaEta=FALSE))
preds.gpp = computePredictedValues(m.gpp, updater=list(GammaEta=FALSE))
MF.gpp = evaluateModelFit(hM=m.gpp, predY=preds.gpp)
MF.gpp

## $RMSE
## [1] 0.2824912 0.1548371 0.2435126 0.2349641 0.3011816
##
## $AUC
## [1] 0.9595238 0.9986450 0.9740808 0.9785432 0.9504831
```





Figure 4: Locations of the created knots in red together with the locations of the plots in black.

```
##
## $TjurR2
## [1] 0.5373093 0.7082397 0.5254423 0.5853250 0.6189267
```

```
cvpreds.gpp = computePredictedValues(m.gpp, partition=partition,
  nParallel = nChains, updater=list(GammaEta=FALSE))
```

```
## Cross-validation, fold 1 out of 2
## Cross-validation, fold 2 out of 2
```

```
cvMF.gpp = evaluateModelFit(hM=m.gpp, predY=cvpreds.gpp)
cvMF.gpp
```

```
## $RMSE
## [1] 0.3634370 0.3285627 0.3639366 0.3579537 0.3081450
##
## $AUC
## [1] 0.8780952 0.8882114 0.8089210 0.8769057 0.9468599
##
## $TjurR2
## [1] 0.3322393 0.2088455 0.1595653 0.2173909 0.5843595
```

```
mpost.gpp = convertToCodaObject(m.gpp)
plot(mpost.gpp$Alpha[[1]])
```

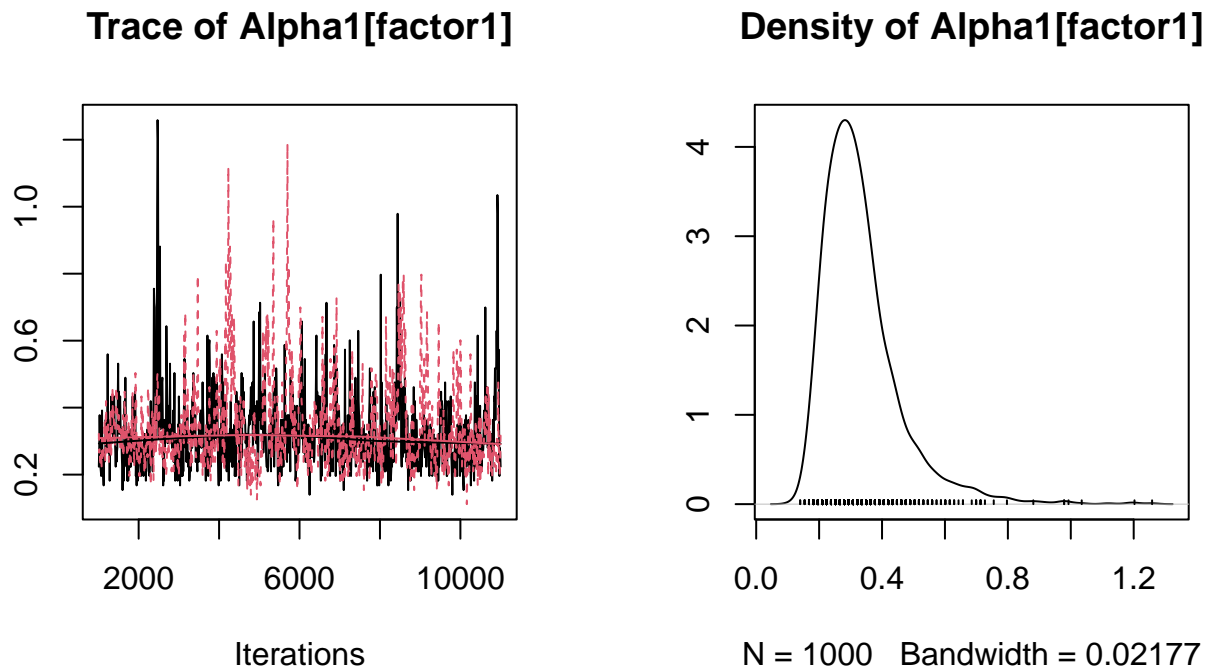


Figure 5: Posterior trace plot of the spatial scale parameter of the ggp spatial model.

```
summary(mpost.gpp$Alpha[[1]])
```

```
##
## Iterations = 1010:11000
## Thinning interval = 10
## Number of chains = 2
## Sample size per chain = 1000
##
## 1. Empirical mean and standard deviation for each variable,
##    plus standard error of the mean:
##
##           Mean           SD      Naive SE Time-series SE
##      0.333196      0.126428      0.002827      0.007967
##
## 2. Quantiles for each variable:
##
##   2.5%   25%   50%   75%  97.5%
## 0.1818 0.2517 0.3076 0.3775 0.6712
```

For this small dataset of 200 sampling units, you will notice that the sampling times for the three spatial models are very similar. However, for larger datasets, the sampling times of the NNGP and GPP method are much shorter than for the standard spatial model.

As a last remark, you may have noticed that we set the `updater` parameter in `sampleMCMC` when running the NNGP and GPP method. With this parameter we specify which MCMC updaters we want to include during the MCMC sampling. The standard setting is to use all available updaters. However, the GammaEta updater is currently not available for these methods. This is not a problem however because this is an optional updater which helps reach convergence in less iterations in some situations.

References

Tikhonov, G., L. Duan, N. Abrego, G. Newell, M. White, D. Dunson, and O. Ovaskainen. 2020. "Computationally Efficient Joint Species Distribution Modeling of Big Spatial Data." *Ecology* 101: e02929. doi:10.1002/ecy.2929.