



NTNU  
Norwegian University of  
Science and Technology

## **INLA - Introduction**

Elias T. Krainski

# Outline

Tokyo example

Hierarchical model

On the Tokyo model


Heart model example

Bayesian inference

INLA overview

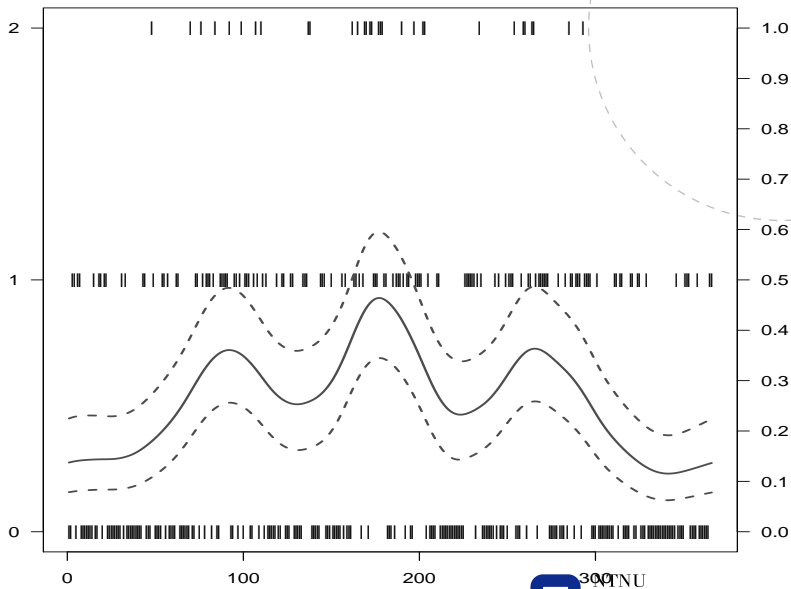


# Number of raining days in Tokyo, for each yearly day in two years



NTNU  
Norwegian University of  
Science and Technology

# Number of raining days in Tokyo, for each yearly day in two years



Norwegian University of Science and Technology

# A model for Tokyo data

Observation model

$$y_i \sim \text{Binomial}(n_i, p_i),$$

for  $i = 1, 2, \dots, 366$

$$n_i = \begin{cases} 1, & \text{for 29 February} \\ 2, & \text{other days} \end{cases} \quad (1)$$

$$y_i \in \begin{cases} \{0, 1\}, & \text{for 29 February} \\ \{0, 1, 2\}, & \text{other days} \end{cases} \quad (2)$$

$$p_i = \frac{1}{1 + \exp(-x_i)}$$

probability on day  $i$  depends on  $x_i$



# Smoothing $x$

— Let

$$x_i | \mathbf{x}_{-i} \sim N(\bar{x}_i, \frac{\sigma^2}{2})$$

where

$$\bar{x}_i = \begin{cases} \frac{x_2 + x_{366}}{2} & \text{if } i = 1 \\ \frac{x_{i-1} + x_{i+1}}{2} & \text{if } 1 < i < 366 \\ \frac{x_{365} + x_1}{2} & \text{if } i = 366 \end{cases} . \quad (3)$$

and  $\theta = 1/\sigma^2$

—  $\theta$  is controls the variation of  $\mathbf{x}$

- so, related to variation of  $p_i$

— as  $\theta > 0$ : people usually use  $\pi(\theta) \sim \text{Gamma}(a, b)$



# Outline

Tokyo example

Hierarchical model

On the Tokyo model

Heart model example

Bayesian inference

INLA overview



NTNU  
Norwegian University of  
Science and Technology

# Bayesian hierarchical model

- $\mathbf{y}$ : observed response data
- $\theta_1$ : likelihood parameter(s)

$$\mathbf{y}|\mathbf{x}, \theta_1 \sim \pi(\mathbf{y}|\mathbf{x}, \theta_1) = \prod_{i=1}^n \pi(y_i|\mathbf{x}, \theta_1) \quad (\text{ind. cond.})$$





# Bayesian hierarchical model

- $\mathbf{y}$ : observed response data
- $\theta_1$ : likelihood parameter(s)

$$\mathbf{y}|\mathbf{x}, \theta_1 \sim \pi(\mathbf{y}|\mathbf{x}, \theta_1) = \prod_{i=1}^n \pi(y_i|\mathbf{x}, \theta_1) \quad (\text{ind. cond.})$$

- $\mathbf{x}$ : latent/unobserved field
  - Gaussian  $\rightarrow$  to use INLA
- $\theta_2$ : latent field parameter(s)

$$\mathbf{x}|\theta_2 \sim \pi(\mathbf{x}|\theta_2) = N(\mathbf{0}, \mathbf{Q}(\theta_2)^{-1})$$



# Bayesian hierarchical model

- $\mathbf{y}$ : observed response data
- $\theta_1$ : likelihood parameter(s)

$$\mathbf{y}|\mathbf{x}, \theta_1 \sim \pi(\mathbf{y}|\mathbf{x}, \theta_1) = \prod_{i=1}^n \pi(y_i|\mathbf{x}, \theta_1) \quad (\text{ind. cond.})$$

- $\mathbf{x}$ : latent/unobserved field
  - Gaussian  $\rightarrow$  to use INLA
- $\theta_2$ : latent field parameter(s)

$$\mathbf{x}|\theta_2 \sim \pi(\mathbf{x}|\theta_2) = N(\mathbf{0}, \mathbf{Q}(\theta_2)^{-1})$$

- in short:  $\theta = \{\theta_1, \theta_2\}$  (hyperparameter)

$$\theta \sim \pi(\theta) \rightarrow \text{to be Bayesian}$$



# $\pi(\mathbf{y}|\mathbf{x}, \theta)$ : likelihood

Depends on

- which kind of data values we have
  - binary (yes/no response, binary image)
  - counts (people infected with a disease in each area)
  - continuous - or + (stock return, temperature)
  - continuous + (rainfall amount, fish weight)
  - survival (recovery time, time to death)



# $\pi(\mathbf{y}|\mathbf{x}, \theta)$ : likelihood

Depends on

- which kind of data values we have
  - binary (yes/no response, binary image)
  - counts (people infected with a disease in each area)
  - continuous - or + (stock return, temperature)
  - continuous + (rainfall amount, fish weight)
  - survival (recovery time, time to death)
- the way it is collected
  - **usual**: each observational unit gives one value
  - each observational unit gives more than one value
  - point process: locations (time, spatial) of events



# $\pi(\mathbf{x}|\mathbf{Q}(\theta))$ : The latent field prior

— It is

- unobserved
- called Gaussian latent random field
- the **most important** ingredient in INLA



# $\pi(\mathbf{x}|\mathbf{Q}(\theta))$ : The latent field prior

- It is
  - unobserved
  - called Gaussian latent random field
  - the **most important** ingredient in INLA
- it represents
  - covariate coefficients
  - unobserved effects



# $\pi(\mathbf{x}|\mathbf{Q}(\theta))$ : The latent field prior

- It is
  - unobserved
  - called Gaussian latent random field
  - the **most important** ingredient in INLA
- it represents
  - covariate coefficients
  - unobserved effects
- it can be
  - unstructured (Tokyo:  $p_i$  doesn't depend on  $p_j$ )
  - structured (Tokyo:  $p_i$  depends on neighbour days)
  - more than one (structured(s) + unstructured(s) + covariate(s))



# $\pi(\theta)$ : The $\theta$ prior

- parameters from the likelihood and  $\mathbf{x}$  distribution
- examples (likelihood):
  - precision parameter of the Gaussian or Gamma
  - dispersion parameter in Beta, negative binomial
  - zero-inflation probability





# $\pi(\theta)$ : The $\theta$ prior

- parameters from the likelihood and  $\mathbf{x}$  distribution
- examples (likelihood):
  - precision parameter of the Gaussian or Gamma
  - dispersion parameter in Beta, negative binomial
  - zero-inflation probability
- examples (latent field)
  - precision parameter in, usually, all of those
  - correlation parameter (in some for time series modelling)
  - range parameter (in some for spatial data modelling)



# Outline

Tokyo example

Hierarchical model

On the Tokyo model

Heart model example

Bayesian inference

INLA overview



# A model for Tokyo data

Observation model

$$y_i \sim \text{Binomial}(n_i, p_i)$$

$$p_i = \frac{1}{1 + \exp(-x_i)}$$

the likelihood has no  $\theta$

$$\pi(\mathbf{y}|\mathbf{x}) = \prod_{i=1}^{366} \pi(y_i|x_i)$$



# Latent model

$$\pi(\mathbf{x}|\theta) \propto \exp \left\{ -\frac{\theta}{2} \left[ (x_1 - x_{366})^2 + \sum_{i=2}^{366} (x_i - x_{i-1})^2 \right] \right\} \quad (4)$$

$$= \exp \left\{ -\frac{\theta}{2} \mathbf{x}^T \mathbf{R} \mathbf{x} \right\} \quad (5)$$





# Latent model **warning**

$$\exp \left\{ -\frac{\theta}{2} \left[ (x_1 - x_{366})^2 + \sum_{i=2}^{366} (x_i - x_{i-1})^2 \right] \right\} \quad (6)$$

(7)

intrinsic/improper

$$\begin{array}{llll} x_i = 20, & x_{i-1} = 10 & \rightarrow & x_i - x_{i-1} = 10 \\ x_i = 10020, & x_{i-1} = 10010 & \rightarrow & x_i - x_{i-1} = 10 \end{array}$$

constraint or take the intercept out



# $\pi(\theta)$ problem

- Tokyo example:  $Q(\theta) = \theta \mathbf{R}$ 
  - bigger  $\theta$  less variation of  $\mathbf{x}$ 
    - related to the variation of  $p_i$
- $\theta > 0$ : people usually use  $\theta \sim \text{Gamma}(a, b)$
- improper distribution:  $\theta$  values depends on  $\mathbf{R}$ 
  - hard to interpret  $\theta$  ( $a=?????$ ,  $b=?????$ )



## $\pi(\mathbf{x}|\theta = 1)$ and $n$

The marginal variance and  $n$  relation


```
rw.var <- function(n, order) {
  R <- as.matrix(INLA:::inla.rw(n, order=order))
  mean(diag(INLA:::inla.ginv(R, rankdef=order)))
}
n <- c(10, 100, 366, 1000); names(n) <- n
rbind(rw1=sapply(n, rw.var, order=1),
      rw2=sapply(n, rw.var, order=2))
```

##	10	100	366	1000
## rw1	1.65	16.665	60.99954	166.6665
## rw2	2.40	2381.190	116733.95702	2380955.1304





$\pi(\mathbf{x}|\theta = 1)$ : one realization



We need to control the marginal variance!



NTNU  
Norwegian University of  
Science and Technology

# $\pi(\theta)$ solution

1. scale the model  $\rightarrow$  easy to interpret  $\theta$ 
  - Tutorial on `scale.option` at [www.r-inla.org/](http://www.r-inla.org/)



# $\pi(\theta)$ solution

1. scale the model  $\rightarrow$  easy to interpret  $\theta$ 
  - Tutorial on `scale.option` at [www.r-inla.org/](http://www.r-inla.org/)
2. AND (new idea) Penalized complexity prior
  - P0: basic model:  $p_i = p_0$
  - P1: complex model:  $p_i$  varies
  - Kullback-Leibler divergence (KLD)
    - a distance from P1 model to P0,  $\text{KLD}(P0/P0) = 0$
  - allow variation on  $p_i$
  - AND supports the basic model
    - **Gamma(a, b) always overfits**



# Outline

Tokyo example

Hierarchical model

On the Tokyo model







Heart model example

Bayesian inference

INLA overview



NTNU  
Norwegian University of  
Science and Technology



Heart data, from `catdata` package



NTNU  
Norwegian University of  
Science and Technology

# Model?

- Generalized Linear Model?
  - linear predictor
  - nonlinear link (logit, probit, and others)
- Nonlinear effect from covariate?
  - parametric nonlinear function?
  - non parametric nonlinear function?
- Bayesian?

```
l <- glm(y ~ age, family=binomial, data=heart)
b <- inla(y ~ age, family='binomial', data=heart)
round(cbind(l=coef(summary(l))[,1:2], b=b$summary.fix[,1:2]))
```

##	l.Estimate	l.Std. Error	b.mean	b.sd
## (Intercept)	-3.5217	0.4160	-3.5257	0.4160
## age	0.0641	0.0085	0.0643	0.0085



# GxxMs: Different names for the same thing

GLMM/GAM/GAMM/+++

- Perhaps the most important class of statistical models
- Many “different” models belong to this class
- No good (enough) MCMC solution around
- Even frequentist approaches does not scale well computationally



# Back to linear models

Consider the linear model

$$y_i = \beta_0 \mathbf{F}_{i1} + \beta_1 \mathbf{F}_{i2} + \beta_2 \mathbf{F}_{i3} + u_i + \epsilon_i$$

where  $\mathbf{F}$  is the design matrix (with ones at first column)

- $y_i$  is an observation
- $\mu$  is the intercept
- $\beta_0, \beta_1$  and  $\beta_2$  are the regression coefficients
- $\mathbf{u}$  is a random effect
- $\epsilon_j$  is i.i.d. normal observation noise.

How it works in a Bayesian framework?





# Bayesian linear models

Linear model:

$$y_i = \beta_0 \mathbf{F}_{i1} + \beta_1 \mathbf{F}_{i2} + \beta_2 \mathbf{F}_{i1} + \mathbf{u}_i + \epsilon_i$$

Bayesian model: chose priors. Usual choices:

- $\boldsymbol{\beta} = (\beta_0, \beta_1, \beta_2)^T \sim N(\mathbf{0}, \tau_{\text{fix}}^{-1} \mathbf{I})$ , where  $\tau_{\text{fix}}$  is a small number
- $\mathbf{u} \sim N(\mathbf{0}, \mathbf{Q}_u^{-1})$  where the *precision matrix*  $\mathbf{Q}$  is known
- $\epsilon \sim N(\mathbf{0}, \tau_n^{-1} \mathbf{I})$



# What does this look like? (Horror slide!)

$(\mathbf{y}, \mathbf{u}, \beta)$  are jointly Gaussian!

$$\begin{aligned} \pi(\mathbf{y}|\mathbf{u}, \beta) &\propto \exp\left(-\frac{\tau_n}{2}(\mathbf{y} - \mathbf{u} - \mathbf{F}^T\beta)^T(\mathbf{y} - \mathbf{u} - \mathbf{F}^T\beta)\right) \\ &= \exp\left(-\frac{\tau_n}{2}(\mathbf{y}^T \quad \mathbf{u}^T \quad \beta^T) \begin{pmatrix} \mathbf{I} & -\mathbf{I} & -\mathbf{F}^T \\ -\mathbf{I} & \mathbf{I} & \mathbf{F}^T \\ -\mathbf{F} & \mathbf{F} & \mathbf{F}^T\mathbf{F} \end{pmatrix} \begin{pmatrix} \mathbf{y} \\ \mathbf{u} \\ \beta \end{pmatrix}\right) \end{aligned}$$

It follows that

$$\begin{aligned} \pi(\mathbf{y}, \mathbf{u}, \beta) &= \pi(\mathbf{y}|\mathbf{u}, \beta)\pi(\mathbf{u})\pi(\beta) \\ &\propto \exp\left(-\frac{\tau_n}{2}(\mathbf{y}^T \quad \mathbf{u}^T \quad \beta^T) \begin{pmatrix} \mathbf{I} & -\mathbf{I} & -\mathbf{F}^T \\ -\mathbf{I} & \mathbf{I} + \tau_n^{-1}\mathbf{Q}_u & \mathbf{F}^T \\ -\mathbf{F} & \mathbf{F} & \mathbf{F}^T\mathbf{F} + \frac{\tau_{\text{fix}}}{\tau_n}\mathbf{I} \end{pmatrix} \begin{pmatrix} \mathbf{y} \\ \mathbf{u} \\ \beta \end{pmatrix}\right) \end{aligned}$$



# How can we use this?

From multivariate Gaussian distributions:

If

$$\mathbf{x} \equiv \begin{pmatrix} \mathbf{x}_A \\ \mathbf{x}_B \end{pmatrix} \sim N \left( \begin{pmatrix} \boldsymbol{\mu}_A \\ \boldsymbol{\mu}_B \end{pmatrix}, \begin{pmatrix} \mathbf{Q}_{AA} & \mathbf{Q}_{AB} \\ \mathbf{Q}_{BA} & \mathbf{Q}_{BB} \end{pmatrix}^{-1} \right),$$

then the conditional distribution is given by

$$\mathbf{x}_A | \mathbf{x}_B \sim N \left( \boldsymbol{\mu}_A - \mathbf{Q}_{AA}^{-1} \mathbf{Q}_{AB} (\mathbf{x}_B - \boldsymbol{\mu}_B), \mathbf{Q}_{AA}^{-1} \right).$$

We can easily compute the marginal distributions for  $u_i | \mathbf{y}$  and  $\beta_i | \mathbf{y}$ .



# Non Gaussian likelihood?

General framework:

- include the linear predictor in  $\mathbf{x}$ 
  - with a small fixed variance

—

$$\pi(\mathbf{y}|\mathbf{x}) = \prod_{i=1}^{\text{\#data}} \pi(y_i|x_i)$$

- Gaussian approximation does the rest



# Further Examples

- Dynamic linear models
- Stochastic volatility models (famously difficult with MCMC)
- Generalised linear (mixed) models
- Generalised additive (mixed) models
- Spline smoothing
- Semiparametric regression
- Space-varying (semiparametric) regression models
- Disease mapping
- Log-Gaussian Cox-processes
- Model-based geostatistics (\*)
- Spatio-temporal models
- Survival analysis
- +++



# Outline

Tokyo example

Hierarchical model

On the Tokyo model

Heart model example

**Bayesian inference**

INLA overview



NTNU  
Norwegian University of  
Science and Technology

# On our Bayesian hierarchical model

- Inference on (what we know about)  $\theta$  and  $\mathbf{x}$  given  $\mathbf{y}$ 
  - in maths:  $\pi(\mathbf{x}|\mathbf{y})$  and  $\pi(\theta|\mathbf{y})$
- considering  $\pi(\mathbf{y}|\mathbf{x}, \theta)$ ,  $\pi(\mathbf{x}|\theta)$  and  $\pi(\theta)$



# On our Bayesian hierarchical model

- Inference on (what we know about)  $\theta$  and  $\mathbf{x}$  given  $\mathbf{y}$ 
  - in maths:  $\pi(\mathbf{x}|\mathbf{y})$  and  $\pi(\theta|\mathbf{y})$
- considering  $\pi(\mathbf{y}|\mathbf{x}, \theta)$ ,  $\pi(\mathbf{x}|\theta)$  and  $\pi(\theta)$
- using the Bayes theorem,

$$\pi(\mathbf{x}|\mathbf{y}) = \int \pi(\mathbf{y}|\mathbf{x}, \theta)\pi(\mathbf{x}|\theta)\pi(\theta)d\theta$$

$$\pi(\theta|\mathbf{y}) = \int \pi(\mathbf{y}|\mathbf{x}, \theta)\pi(\mathbf{x}|\theta)\pi(\theta)d\mathbf{x}$$





# On our Bayesian hierarchical model

- Inference on (what we know about)  $\theta$  and  $\mathbf{x}$  given  $\mathbf{y}$ 
  - in maths:  $\pi(\mathbf{x}|\mathbf{y})$  and  $\pi(\theta|\mathbf{y})$
- considering  $\pi(\mathbf{y}|\mathbf{x}, \theta)$ ,  $\pi(\mathbf{x}|\theta)$  and  $\pi(\theta)$
- using the Bayes theorem,

$$\pi(\mathbf{x}|\mathbf{y}) = \int \pi(\mathbf{y}|\mathbf{x}, \theta)\pi(\mathbf{x}|\theta)\pi(\theta)d\theta$$

$$\pi(\theta|\mathbf{y}) = \int \pi(\mathbf{y}|\mathbf{x}, \theta)\pi(\mathbf{x}|\theta)\pi(\theta)d\mathbf{x}$$

- even more...
  - $\pi(\theta_j|\mathbf{y})$ ,  $j = 1, \dots, \dim(\theta)$
  - $\pi(x_i|\mathbf{y})$ ,  $i = 1, \dots, \dim(\mathbf{x})$



# The inference problem

— we have to compute

$$\pi(\mathbf{x}_i | \mathbf{y}) \propto \int_{\mathbf{x}_{\{-i\}}} \int_{\boldsymbol{\theta}} \pi(\mathbf{y} | \mathbf{x}, \boldsymbol{\theta}) \pi(\mathbf{x} | \boldsymbol{\theta}) \pi(\boldsymbol{\theta}) d\boldsymbol{\theta} d\mathbf{x}_{\{-i\}}$$



# The inference problem

— we have to compute

$$\pi(\mathbf{x}_j | \mathbf{y}) \propto \int_{\mathbf{x}_{\{-j\}}} \int_{\boldsymbol{\theta}} \pi(\mathbf{y} | \mathbf{x}, \boldsymbol{\theta}) \pi(\mathbf{x} | \boldsymbol{\theta}) \pi(\boldsymbol{\theta}) d\boldsymbol{\theta} d\mathbf{x}_{\{-j\}}$$

and

$$\pi(\boldsymbol{\theta}_j | \mathbf{y}) \propto \int_{\mathbf{x}} \int_{\boldsymbol{\theta}_{\{-j\}}} \pi(\mathbf{y} | \mathbf{x}, \boldsymbol{\theta}) \pi(\mathbf{x} | \boldsymbol{\theta}) \pi(\boldsymbol{\theta}) d\boldsymbol{\theta}_{\{-j\}} d\mathbf{x}$$



# The inference problem

— we have to compute

$$\pi(\mathbf{x}_j | \mathbf{y}) \propto \int_{\mathbf{x}_{\{-j\}}} \int_{\boldsymbol{\theta}} \pi(\mathbf{y} | \mathbf{x}, \boldsymbol{\theta}) \pi(\mathbf{x} | \boldsymbol{\theta}) \pi(\boldsymbol{\theta}) d\boldsymbol{\theta} d\mathbf{x}_{\{-j\}}$$

and

$$\pi(\boldsymbol{\theta}_j | \mathbf{y}) \propto \int_{\mathbf{x}} \int_{\boldsymbol{\theta}_{\{-j\}}} \pi(\mathbf{y} | \mathbf{x}, \boldsymbol{\theta}) \pi(\mathbf{x} | \boldsymbol{\theta}) \pi(\boldsymbol{\theta}) d\boldsymbol{\theta}_{\{-j\}} d\mathbf{x}$$

— remember

- $\dim(\boldsymbol{\theta})$  is small
- $\dim(\mathbf{x})$  is not small
- we have to compute very high dimensional integrals



# The inference problem

— we have to compute

$$\pi(\mathbf{x}_j | \mathbf{y}) \propto \int_{\mathbf{x}_{\{-j\}}} \int_{\boldsymbol{\theta}} \pi(\mathbf{y} | \mathbf{x}, \boldsymbol{\theta}) \pi(\mathbf{x} | \boldsymbol{\theta}) \pi(\boldsymbol{\theta}) d\boldsymbol{\theta} d\mathbf{x}_{\{-j\}}$$

and

$$\pi(\boldsymbol{\theta}_j | \mathbf{y}) \propto \int_{\mathbf{x}} \int_{\boldsymbol{\theta}_{\{-j\}}} \pi(\mathbf{y} | \mathbf{x}, \boldsymbol{\theta}) \pi(\mathbf{x} | \boldsymbol{\theta}) \pi(\boldsymbol{\theta}) d\boldsymbol{\theta}_{\{-j\}} d\mathbf{x}$$

— remember

- $\dim(\boldsymbol{\theta})$  is small
- $\dim(\mathbf{x})$  is not small
- we have to compute very high dimensional integrals

— typically they are not analytically tractable

- $\rightarrow$  we have to approach



# using MCMC

- single-site: compute (the expressions) for
  - $p(\theta_j | \boldsymbol{\theta}_{-j}, \mathbf{x}, \mathbf{y})$
  - $p(x_i | \mathbf{x}_{-i}, \boldsymbol{\theta}, \mathbf{y})$



# using MCMC

- single-site: compute (the expressions) for
  - $p(\theta_j | \theta_{-j}, \mathbf{x}, \mathbf{y})$
  - $p(x_i | \mathbf{x}_{-i}, \theta, \mathbf{y})$
- draw samples from such conditionals
  - WinBUGS, OpenBUGS, JAGS, and others
- use these samples to summarize  $p(\mathbf{x})$  and  $p(\theta)$



# using MCMC

- single-site: compute (the expressions) for
  - $p(\theta_j | \theta_{-j}, \mathbf{x}, \mathbf{y})$
  - $p(x_i | \mathbf{x}_{-i}, \theta, \mathbf{y})$
- draw samples from such conditionals
  - WinBUGS, OpenBUGS, JAGS, and others
- use these samples to summarize  $p(\mathbf{x})$  and  $p(\theta)$
- **warning**
  - sampling from  $x_i | \mathbf{x}_{-i}, \theta, \mathbf{y}$ 
    - slow convergence when strong dependence
    - **does not work for our example...**
  - better: draw joint sample from  $\mathbf{x} | \theta, \mathbf{y}$
  - best: use INLA





# Outline

Tokyo example

Hierarchical model

On the Tokyo model

Heart model example

Bayesian inference

**INLA overview**



# What INLA does

- INLA does:
  - compute marginals of  $\pi(x_j|\mathbf{y})$  and  $\pi(\theta_j|\mathbf{y})$
- how?
  - approach  $\pi(\mathbf{x}|\boldsymbol{\theta}, \mathbf{y})$  to approach  $\pi(\boldsymbol{\theta}|\mathbf{y})$
  - explore  $\pi(\boldsymbol{\theta}|\mathbf{y})$ 
    - approach  $\pi(\theta_j|\mathbf{y})$
  - approach  $\pi(x_j|\mathbf{x}_{-j})$



# Important ingredient

The GMRF-approximation

$$\pi(\mathbf{x} \mid \boldsymbol{\theta}, \mathbf{y}) \propto \exp\left(-\frac{1}{2}\mathbf{x}^T \mathbf{Q} \mathbf{x} + \sum_i \log \pi(y_i | x_i)\right)$$



# Important ingredient

The GMRF-approximation

$$\begin{aligned} \pi(\mathbf{x} \mid \boldsymbol{\theta}, \mathbf{y}) &\propto \exp\left(-\frac{1}{2}\mathbf{x}^T \mathbf{Q} \mathbf{x} + \sum_i \log \pi(y_i \mid x_i)\right) \\ &\approx \exp\left(-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu})^T (\mathbf{Q} + \text{diag}(\mathbf{c})) (\mathbf{x} - \boldsymbol{\mu})\right) \\ &= \pi_G(\mathbf{x} \mid \boldsymbol{\theta}, \mathbf{y}) \end{aligned}$$

$$c_i = -\frac{d^2 l_i}{dx_i^2} \text{ where } l_i = \log(\pi(y_i \mid x_i)), i = 1, \dots, \# \text{ data}$$



# Important ingredient

The GMRF-approximation

$$\begin{aligned}\pi(\mathbf{x} \mid \boldsymbol{\theta}, \mathbf{y}) &\propto \exp\left(-\frac{1}{2}\mathbf{x}^T \mathbf{Q} \mathbf{x} + \sum_i \log \pi(y_i | x_i)\right) \\ &\approx \exp\left(-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu})^T (\mathbf{Q} + \text{diag}(\mathbf{c})) (\mathbf{x} - \boldsymbol{\mu})\right) \\ &= \pi_G(\mathbf{x} | \boldsymbol{\theta}, \mathbf{y})\end{aligned}$$

$c_i = -\frac{d^2 l_i}{dx_i^2}$  where  $l_i = \log(\pi(y_i | x_i))$ ,  $i = 1, \dots, \# \text{ data}$

— Markov and computational properties (on  $\mathbf{Q}$ ) are preserved



# Important ingredient

The GMRF-approximation

$$\begin{aligned}\pi(\mathbf{x} \mid \boldsymbol{\theta}, \mathbf{y}) &\propto \exp\left(-\frac{1}{2}\mathbf{x}^T \mathbf{Q} \mathbf{x} + \sum_i \log \pi(y_i \mid x_i)\right) \\ &\approx \exp\left(-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu})^T (\mathbf{Q} + \text{diag}(\mathbf{c})) (\mathbf{x} - \boldsymbol{\mu})\right) \\ &= \pi_G(\mathbf{x} \mid \boldsymbol{\theta}, \mathbf{y})\end{aligned}$$

$c_i = -\frac{dl_i^2}{dx_i^2}$  where  $l_i = \log(\pi(y_i \mid x_i))$ ,  $i = 1, \dots, \# \text{ data}$

- Markov and computational properties (on  $\mathbf{Q}$ ) are preserved
- $\tilde{\pi}(\mathbf{x} \mid \boldsymbol{\theta}, \mathbf{y})$  costs
  - temporal:  $O(n)$
  - spatial:  $O(n \log(n))$

If  $\mathbf{y} \mid \mathbf{x}, \boldsymbol{\theta}$  is *Gaussian*, the “approximation” is exact.



# INLA, $\pi(\theta_j|\mathbf{y})$

— Considering

$$\pi(\boldsymbol{\theta}|\mathbf{y}) = \frac{\pi(\boldsymbol{\theta}, \mathbf{x}|\mathbf{y})}{\pi(\mathbf{x}|\boldsymbol{\theta}, \mathbf{y})}$$



# INLA, $\pi(\theta_j|\mathbf{y})$

— Considering

$$\begin{aligned}\pi(\boldsymbol{\theta}|\mathbf{y}) &= \frac{\pi(\boldsymbol{\theta}, \mathbf{x}|\mathbf{y})}{\pi(\mathbf{x}|\boldsymbol{\theta}, \mathbf{y})} \\ &\propto \frac{\pi(\boldsymbol{\theta})\pi(\mathbf{x}|\boldsymbol{\theta})\pi(\mathbf{y}|\mathbf{x}, \boldsymbol{\theta})}{\pi(\mathbf{x}|\boldsymbol{\theta}, \mathbf{y})}\end{aligned}$$





# INLA, $\pi(\theta_j | \mathbf{y})$

— Considering

$$\begin{aligned}\pi(\boldsymbol{\theta} | \mathbf{y}) &= \frac{\pi(\boldsymbol{\theta}, \mathbf{x} | \mathbf{y})}{\pi(\mathbf{x} | \boldsymbol{\theta}, \mathbf{y})} \\ &\propto \frac{\pi(\boldsymbol{\theta})\pi(\mathbf{x} | \boldsymbol{\theta})\pi(\mathbf{y} | \mathbf{x}, \boldsymbol{\theta})}{\pi(\mathbf{x} | \boldsymbol{\theta}, \mathbf{y})}\end{aligned}$$

— Gaussian approximation to denominator

$$\pi(\boldsymbol{\theta} | \mathbf{y}) \approx \frac{\pi(\boldsymbol{\theta})\pi(\mathbf{x} | \boldsymbol{\theta})\pi(\mathbf{y} | \mathbf{x}, \boldsymbol{\theta})}{\pi_G(\mathbf{x} | \boldsymbol{\theta}, \mathbf{y})} \Big|_{\mathbf{x}=\mathbf{x}^*(\boldsymbol{\theta})}$$



# INLA, $\pi(\theta_j|\mathbf{y})$

— Considering

$$\begin{aligned}\pi(\boldsymbol{\theta}|\mathbf{y}) &= \frac{\pi(\boldsymbol{\theta}, \mathbf{x}|\mathbf{y})}{\pi(\mathbf{x}|\boldsymbol{\theta}, \mathbf{y})} \\ &\propto \frac{\pi(\boldsymbol{\theta})\pi(\mathbf{x}|\boldsymbol{\theta})\pi(\mathbf{y}|\mathbf{x}, \boldsymbol{\theta})}{\pi(\mathbf{x}|\boldsymbol{\theta}, \mathbf{y})}\end{aligned}$$

— Gaussian approximation to denominator

$$\pi(\boldsymbol{\theta}|\mathbf{y}) \approx \frac{\pi(\boldsymbol{\theta})\pi(\mathbf{x}|\boldsymbol{\theta})\pi(\mathbf{y}|\mathbf{x}, \boldsymbol{\theta})}{\pi_G(\mathbf{x}|\boldsymbol{\theta}, \mathbf{y})} \Big|_{\mathbf{x}=\mathbf{x}^*(\boldsymbol{\theta})}$$

— mode of  $\tilde{\pi}(\boldsymbol{\theta}|\mathbf{y})$  (optimization)

- explore  $\tilde{\pi}(\boldsymbol{\theta}|\mathbf{y})$ 
  - approach  $\pi(\theta_j|\mathbf{y})$  (numerical integration)



# INLA, $\pi(x_i | \mathbf{y}, \boldsymbol{\theta})$

Approaching  $\pi(x_i | \mathbf{y}, \boldsymbol{\theta})$

— Problem

- $\dim(\mathbf{x})=n$  is not small
- $n$  marginals to compute

— Laplace approximation

$$\tilde{\pi}(x_i | \mathbf{y}, \boldsymbol{\theta}) \approx \frac{\pi(\mathbf{x}, \boldsymbol{\theta} | \mathbf{y})}{\tilde{\pi}_{GG}(\mathbf{x}_{-i} | x_i, \mathbf{y}, \boldsymbol{\theta})} \Big|_{\mathbf{x}_{-i} = \mathbf{x}_{-i}^*(x_i, \boldsymbol{\theta})}$$



# INLA, $\pi(x_i | \mathbf{y}, \boldsymbol{\theta})$

Approaching  $\pi(x_i | \mathbf{y}, \boldsymbol{\theta})$

— Problem

- $\dim(\mathbf{x})=n$  is not small
- $n$  marginals to compute

— Laplace approximation

$$\tilde{\pi}(x_i | \mathbf{y}, \boldsymbol{\theta}) \approx \frac{\pi(\mathbf{x}, \boldsymbol{\theta} | \mathbf{y})}{\tilde{\pi}_{GG}(\mathbf{x}_{-i} | x_i, \mathbf{y}, \boldsymbol{\theta})} \Big|_{\mathbf{x}_{-i} = \mathbf{x}_{-i}^*(x_i, \boldsymbol{\theta})}$$

— simpler/crunder (fast) approximation (from  $\pi_G(\mathbf{x} | \mathbf{y}, \boldsymbol{\theta})$ )

$$\hat{\pi}(x_i | \mathbf{y}, \boldsymbol{\theta}) = N(x_i; \mu_i(\boldsymbol{\theta}), \sigma_i^2(\boldsymbol{\theta}))$$



# INLA, $\pi(x_j | \mathbf{y})$

Approaching  $\pi(x_j | \mathbf{y}, \theta)$

- integrate  $\theta$  out from  $\tilde{\pi}(x_j | \mathbf{y}, \theta)$
- select values for  $\theta$
- use weighted sum

$$\tilde{\pi}(x_j | \mathbf{y}) \propto \sum_j \tilde{\pi}(x_j | \mathbf{y}, \theta_j) \times \tilde{\pi}(\theta_j | \mathbf{y})$$



# Remarks

1. Expect  $\tilde{\pi}(\boldsymbol{\theta}|\mathbf{y})$  to be accurate, since
  - $\mathbf{x}|\boldsymbol{\theta}$  is *a priori* Gaussian
  - Likelihood models are 'well-behaved' so

$$\pi(\mathbf{x} | \boldsymbol{\theta}, \mathbf{y})$$

is *almost* Gaussian.

2. There are no distributional assumptions on  $\boldsymbol{\theta}|\mathbf{y}$
3. Similar remarks are valid to

$$\tilde{\pi}(x_i | \boldsymbol{\theta}, \mathbf{y})$$



# How can we assess the error in the approximations?

**Tool 1:** Compare a sequence of improved approximations

1. Gaussian approximation
2. Simplified Laplace
3. Laplace

No big differences  $\rightarrow$  good approximation



# How can we assess the error in the approximations?

**Tool 2:** Estimate the “effective” number of parameters as defined in the Deviance Information Criteria:

$$\rho_D(\theta) = \bar{D}(\mathbf{x}; \theta) - D(\bar{\mathbf{x}}; \theta)$$

and compare this with the number of observations  
Low ratio is good.

This criteria has theoretical justification.

