Extending the Integrated Laplace Approximation (using MCMC)

Virgilio Gómez-Rubio[1] and Håvard Rue[2]
1. University of Castilla-La Mancha, 02071 Albacete, Spain
2. Norwegian University for Science and Technology, 7491 Trondheim, Norway

NTNU – Trondheim
Norwegian University of
Science and Technology

The Integrated Nested Laplace Approximation provides a feasible framework for approximate Bayesian inference. However, the type of models that can be fitted in practice is constrained to those implemented in the **R-INLA** package. We describe here a novel approach to fit more general models by combining MCMC and INLA.

Extending INLA (using MCMC)

- Bivand et al. (2014, 2015) provide a way to extend the models that can be fitted with **R-INLA**:
  - Fix one or several hyperparameters so that the model can be fitted with **R-INLA**
  - Fit these conditional models for different values of the fixed hyperparameters
  - Combine the resulting models using Bayesian Model Averaging (BMA)
- Using a regular grid of values of the hyperparameters may not be feasible:
  - Unbounded parameters
  - Large number of parameters
- Split the vector of parameters $\theta = (\theta_c, \theta_{-c})$
  - Models can be fitted with **R-INLA** conditioning on θ_c
- General idea:
  - Use MCMC to estimate $\pi(\theta_c|y)$
  - Use **R-INLA** to estimate $\pi(\theta_{-c}|\theta_c, y)$
  - Use BMA to obtain $\pi(\theta_{-c}|y)$

Metropolis-Hastings with INLA

- Step (n) in M-H requires:
  - Sampling a new proposal $\theta_c^{(n)}$ from $q(\cdot|\theta_c^{(n-1)})$
  - Fitting conditional model: $\pi(\theta_{-c}|y, \theta_c^{(n)})$ and $\pi(y|\theta_c^{(n)})$
  - Accept/reject $\theta_c^{(n)}$: requires $\pi(y|\theta_c^{(n-1)})$ and $\pi(y|\theta_c^{(n)})$

$$\frac{\pi(\theta_c^{(n)}|y)q(\theta_c^{(n-1)}|\theta_c^{(n)})}{\pi(\theta_c^{(n-1)}|y)q(\theta_c^{(n)}|\theta_c^{(n-1)})} = \frac{\pi(y|\theta_c^{(n)})\pi(\theta_c^{(n)})q(\theta_c^{(n-1)}|\theta_c^{(n)})}{\pi(y|\theta_c^{(n-1)})\pi(\theta_c^{(n-1)})q(\theta_c^{(n)}|\theta_c^{(n-1)})}$$

- Output is:
  - Sample from $\theta_c|y$
  - List of fitted models (with conditional marginals for the other parameters)

$$\left\{\pi(\theta_{-c}|y, \theta_c^{(i)})\right\}_{i=1}^{k}$$

- Marginals are obtained as follows:
  - $\pi(\theta_c|y)$ from the MCMC output
  - $\pi(\theta_{-c}|y)$ by BMA'ing $\left\{\pi(\theta_{-c}|y, \theta_c^{(i)})\right\}_{i=1}^{k}$

Toy Example: Univariate Case

- Simple linear regression with one covariate:

$$y_i = \alpha + \beta x_i + \varepsilon_i;\ \ i = 1, \ldots, 100$$

Here, ε_i is a Gaussian error term with zero mean and precision τ.

- β is sampled using M-H
- $\pi(\beta|y)$ is computed from sampled values
- α, τ are estimated with **R-INLA**
- $\pi(\alpha|y)$, $\pi(\tau|y)$ are computed by BMA'ing the fitted models
- We have computed the posterior marginals of the parameters in 3 different ways:
  - R-INLA
  - MCMC
  - INLA+MCMC

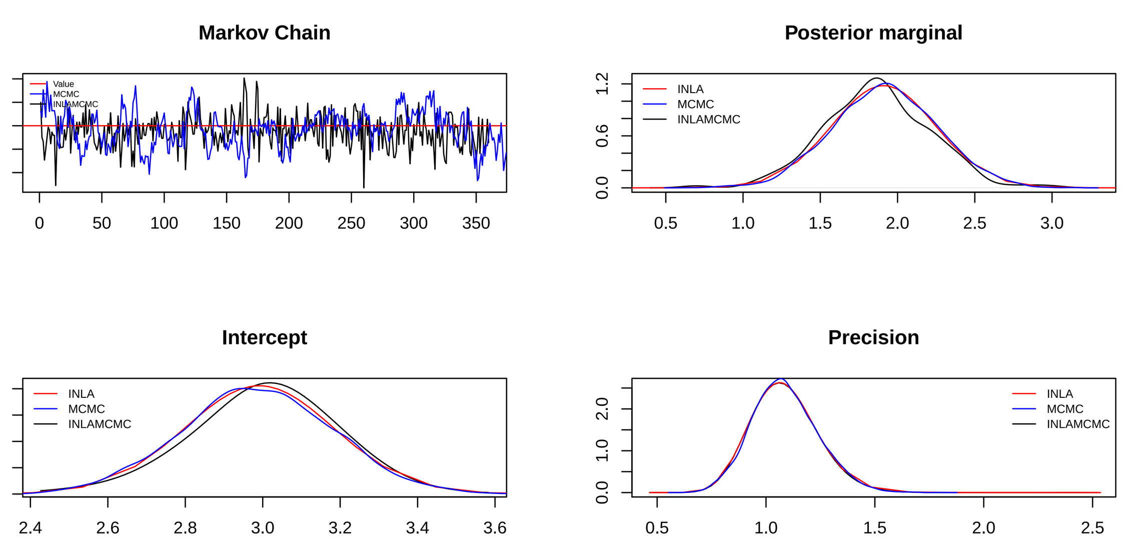

Toy Example: Bivariate Case

- Simple linear regression with two covariate:

$$y_i = \alpha + \beta_1 x_{1i} + \beta_2 x_{2i} + \varepsilon_i;\ \ i = 1, \ldots, 100$$

Here, ε_i is a Gaussian error term with zero mean and precision τ.

- β_1, β_2 are sampled using M-H
- $\pi(\beta_i|y)$; $i = 1, 2$ are computed from sampled values
- $\pi(\beta_1, \beta_2|y)$ can be estimated from sampled values
- α, τ are estimated with **R-INLA**
- $\pi(\alpha|y)$, $\pi(\tau|y)$ are computed by BMA'ing the fitted models
- We have computed the posterior marginals of the parameters in 3 different ways:
  - R-INLA
  - MCMC
  - INLA+MCMC

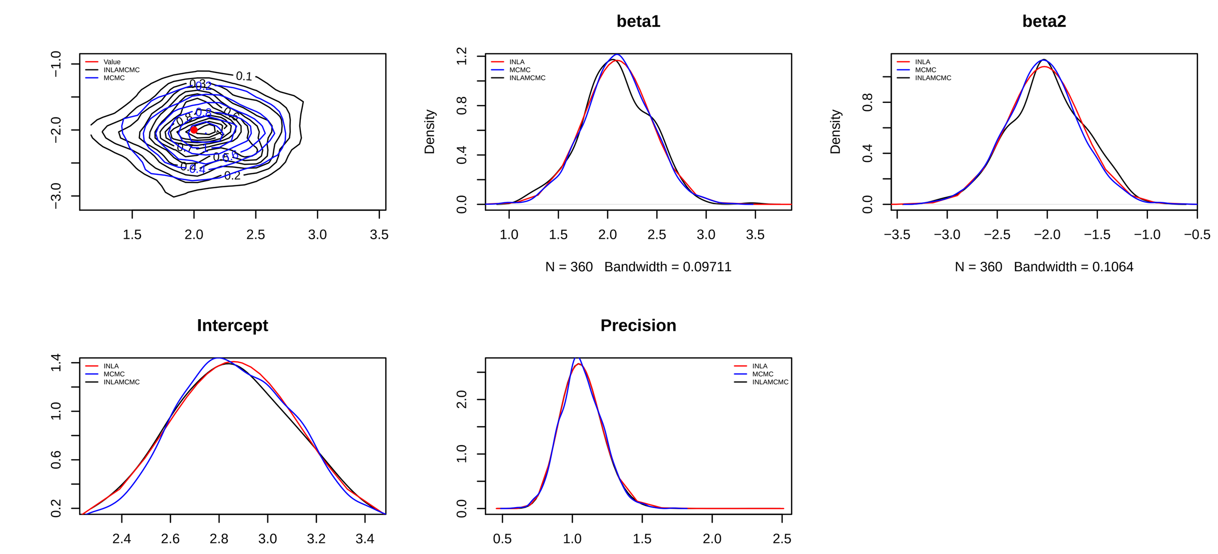

Toy Example: Imputation of Missing Covariates

- Simple linear regression with one covariate
- Missing values in the covariates
- In principle, **R-INLA** cannot handle this...
- We treat missing values as model parameters
- Missing covariates x_m are sampled using block updating in M-H
- $\pi(x_m|y)$ is computed from sampled values
- β_1, τ are estimated with **R-INLA**
- $\pi(\beta_1|y)$, $\pi(\tau|y)$ are computed by BMA'ing the fitted models
- We have computed the posterior marginals of the parameters in 2 different ways:
  - MCMC
  - INLA+MCMC

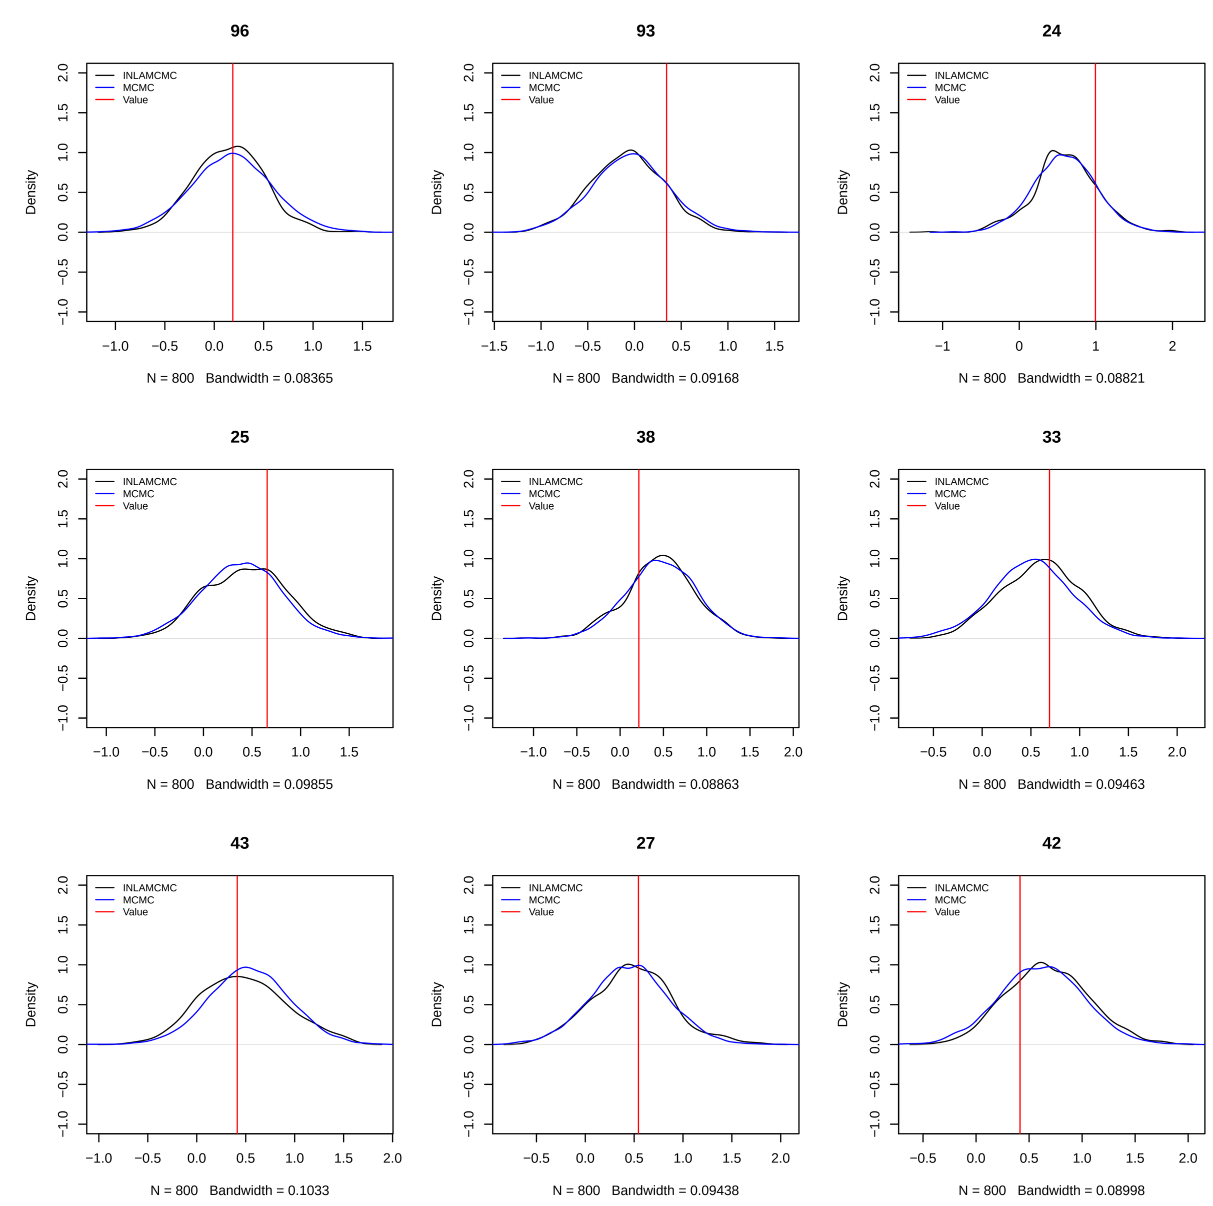

Work in progress. We hope to apply this new approach to:

- Models with complex structures of the random effects that involve several hyperparameters
- Models with a hierarchical structure on the hyperparameters (for example, the precision)
- Spatial econometrics models that include several spatial autocorrelation parameters
- Imputation of missing covariates in real datasets

To download code and references, please visit:
`http://www.uclm.es/profesorado/vgomez/INLAMCMC`