

Approximate Bayesian Inference for Spatial Econometrics Models

V. Gómez-Rubio

Department of Mathematics
U. of Castilla-La Mancha, Albacete, Spain

Innsbruck, 7 November 2012

joint work with Roger S. Bivand and Havard Rue

Talk Outline

- 1 Spatial models
- 2 Regression models
- 3 Spatial Econometrics
- 4 Bayesian Inference
- 5 Approximate Inference

Spatial models



Modelling spatial dependence is important when events close in space are thought to have a similar behaviour

Spatial statistics

- Point patterns
- Geostatistics
- Lattice data


Point patterns

- Point patterns record the locations of a series of events in a study area
- The aim is to determine the distribution of points in the study area, i.e., their probability distribution in space
- Also, it is of interest whether the events appear independently (Complete Spatial Randomness), clustered or repelled




Geostatistics

- Some environmental processes show a strong spatial pattern (e.g., temperature, wind speed, concentration of heavy metals, etc.)
- If a survey is conducted in the study area there is a strong spatial dependence
- Samples close in space will have closer values than values that are further apart



Lattice data

- Lattice data involves data measured at different areas, e.g., neighbourhoods, cities, provinces, states, etc.
- Spatial dependence appears because neighbour areas will show similar values of the variable of interest



Approaches to spatial modelling

- In point patterns, we will estimate the density of the point pattern
- In geostatistics, we will model how the variable changes with distance
- In the analysis of lattice data we will model the correlation between the outcome at the different areas
- **This talk focuses on the analysis of lattice data using different models**

Multivariate Normal

- The vector of observed values \mathbf{y} can be supposed to have a multivariate Normal distribution:

$$\mathbf{y} \sim MVN(\boldsymbol{\mu}, \boldsymbol{\Sigma})$$

- Note that each observation y_i has been measured at a different location
- The mean can be modelled to depend on relevant covariates
- The variance covariance matrix can be defined so that observations close in space have higher covariance/correlation
- In the case of lattice data, a common approach is that if two regions are neighbours they will have a higher correlation
- There are many ways of defining $\boldsymbol{\Sigma}$ to account for spatial dependence

Continuous processes

- In geostatistics we have a continuous process $Z(x)$
- Observations are taken at some survey sites $\{x_i\}_{i=1}^n$
- Many times we have multivariate data. For example, concentration of heavy metals (zinc, copper, etc.)
- A common approach is to measure how the process change, i.e., $Z(x) - Z(y)$
- We can compute the (sem-i)variogram to measure spatial dependence

$$\gamma(x, y) = \frac{1}{2} E[(Z(x) - Z(y))^2]$$

- Krigging uses the semivariogram to produce estimates such as

$$\hat{Z}(x) = \sum_{i=1}^n \lambda_i(x) Z(x_i)$$


where $\lambda_i(x)$ are a number of weights

- $\lambda_i(x)$ is higher for points x_i closer to x

Models for lattice data

- We have observations $y = \{y_i\}_{i=1}^n$ from the n areas
- y is assigned a multivariate distribution that *accounts for spatial dependence*
- A common way of describing spatial proximity in lattice data is by means of an *adjacency matrix* W
- $W[i, j]$ is non-zero if areas i and j are neighbours
- Usually, two areas are neighbours if they share a common boundary
- There are other definitions of neighbourhood

Adjacency matrix



Regression models

- It is often the case that, in addition to y_i , we have a number of covariates x_i
- Hence, we may want to regress y_i on x_i
- In addition, to the covariates we may want to account for the spatial structure of the data
- Different types of regression models can be used to model lattice data:
 - Generalized Linear Models (with spatial random effects)
 - Spatial econometrics models

Linear Mixed Models

- A common approach (for Gaussian data) is to use a linear regression with random effects

$$Y = X\beta + Zu + \varepsilon$$

- The vector random effects u is modelled as a MVN:

$$u \sim N(0, \sigma_u^2 \Sigma)$$

- Σ is defined such as it induces higher correlation in adjacent areas
- Z is a design matrix for the random effects
- $\varepsilon_i \sim N(0, \sigma^2), i = 1, \dots, n$: error term
- The full model can be written as:

$$Y \sim N(X\beta, \sigma^2 I + \sigma_u^2 Z\Sigma Z^T)$$

Generalized Linear Models

- Y_i random variable of the Exponential Family

$$f_Y(y; \theta, \phi) = \exp\{(y\theta - b(\theta))/a(\phi) + c(y, \phi)\}$$

ϕ is a *scale* parameter and θ is the *canonical* parameter

- Linear predictor

$$\eta = \beta_1 x_1 + \dots + \beta_k x_k$$

- *Link* function

$$g(\mu) = \eta = \beta_1 x_1 + \dots + \beta_k x_k; \quad \mu \equiv E[Y]$$

Generalized Linear Models

Parameters of some common distributions:

Distribution	Range y	$\mu(\theta)$	Canon. link	ϕ
$N(\mu, \sigma^2)$	$(-\infty, \infty)$	θ	identity	σ^2
$P(\mu)$	$0, 1, \dots, \infty$	$\exp(\theta)$	log	1
$B(n, p)/n$	$(0, 1, \dots, n)/n$	$\frac{\exp\theta}{1+\exp\theta}$	logit	$1/n$

Some link functions:

$$\eta = \log(\mu)$$

$$\eta = \text{logit}(\mu) = \log\left(\frac{\mu}{1-\mu}\right)$$

$$\eta = \text{probit}(\mu) = \Phi^{-1}(\mu)$$

$$\eta = \text{cloglog}(\mu) = \log(-\log(1-\mu))$$

GLMs with random effects

- Y_i random variable of the exponential family

$$f_Y(y; \theta, \phi) = \exp\{(y\theta - b(\theta))/a(\phi) + c(y, \phi)\}$$

- Linear predictor

$$\eta = \beta_1 x_1 + \dots + \beta_k x_k + Zu$$

- *Link* function

$$g(\mu) = \eta = \beta_1 x_1 + \dots + \beta_k x_k + Zu; \quad \mu \equiv E[Y]$$

- u is a vector of random effects, which are MVN-distributed

$$u \sim N(0, \sigma^2 \Sigma)$$

- Z is a design matrix for the random effects

Structure of spatial random effects

There are **many** different ways of including spatial dependence in Σ :

- Simultaneous autoregressive (SAR):

$$\Sigma = [(I - \rho W)'(I - \rho W)]^{-1}$$

- Conditional autoregressive (CAR):

$$\Sigma = (I - \rho W)^{-1}$$

- $\Sigma_{i,j}$ depends on a function of $d(i,j)$. For example:

$$\Sigma_{i,j} = \exp\{-d(i,j)/\phi\}$$

SAR

- The variance-covariance matrix has the following structure:

$$\Sigma = [(I - \rho W)'(I - \rho W)]^{-1}$$

- W can have many different forms.
- Any of them will produce a symmetric variace-covariance matrix (but some may be singular!)
- W can be the usual 0/1 matrix
- W is often taken as a row-standardised matrix
- This implies that in many cases $\rho \in (-1, 1)$
- See (Haining, 2003) for details

CAR

- The variance-covariance matrix has the following structure:

$$\Sigma = (I - \rho W)^{-1}$$

- In order to have a valid variance covariance matrix W must be symmetric
- W is often taken as a binary matrix
- If $\rho = 1$ we have an intrinsic CAR model but then Σ is singular
- Fitting this model will require the use of a generalized inverse
- Seldom a problem if you are Bayesian!!

Spatial Econometrics Models

- Slightly different approach to spatial modelling
- Instead of using latent effects, spatial dependence is modelled explicitly
- Autoregressive models are used to make the response variable to depend on the values at its neighbours

Spatial autoregressive models

- Autoregression on the response term

$$y = \rho W y + \alpha + e; e \sim N(0, \sigma^2 I)$$

- Autoregression on the error term

$$y = \alpha + u; u = \rho W u + e; e \sim N(0, \sigma^2 I)$$

- ρ measures spatial dependence
- If $\rho = 0$ there is no spatial dependence

Regression models on the response

$$y = \alpha + \rho W y + e; e \sim N(0, \sigma^2 I)$$

$$y = (I - \rho W)^{-1}(\alpha + e); e \sim N(0, \sigma^2 I)$$

$$y = (I - \rho W)^{-1}\alpha + \varepsilon; \varepsilon \sim N(0, \sigma^2 [(I - \rho W')(I - \rho W)]^{-1})$$

Regression models on the error term

$$y = \alpha + u; u = Wu + \varepsilon; \varepsilon \sim N(0, \sigma^2 I)$$

$$y = \alpha + u; u = (I - \rho W)^{-1} \varepsilon; \varepsilon \sim N(0, \sigma^2 I)$$

$$y = \alpha + u; u \sim N(0, \sigma^2 [(I - \rho W')(I - \rho W)]^{-1})$$

Simultaneous Autoregressive Model (SEM)

- This model includes covariates
- Autoregressive on the error term

$$y = X\beta + u; u = \rho Wu + e; e \sim N(0, \sigma^2)$$

$$y = X\beta + \varepsilon; \varepsilon \sim N(0, \sigma^2(I - \rho W)^{-1}(I - \rho W')^{-1})$$

Spatial Lag Model (SLM)

- This model includes covariates
- Autoregressive on the response

$$y = \rho W y + X\beta + e; e \sim N(0, \sigma^2)$$

$$y = (I - \rho W)^{-1} X\beta + \varepsilon; \varepsilon \sim N(0, \sigma^2 (I - \rho W)^{-1} (I - \rho W')^{-1})$$

Spatial Durbin Model (SDM)

- This model includes covariates
- Autoregressive on the response
- In addition, we include the lagged-covariates WX as another extra term in the regression

$$y = \rho Wy + X\beta + WX\gamma + e = [X, WX][\beta, \gamma] + e; \quad e \sim N(0, \sigma^2)$$

$$y = \rho Wy + XWX\mathcal{B} + e; \quad XWX = [X, WX]; \quad \mathcal{B} = [\beta, \gamma]$$

$$y = (I - \rho W)^{-1} XWX \mathcal{B} + \varepsilon; \quad \varepsilon \sim N(0, \sigma^2 (I - \rho W)^{-1} (I - \rho W')^{-1})$$

Probit models

- LeSage et al. (2011) use a non-Gaussian model for the probability of reopening a business in the aftermath of hurricane Katrina in New Orleans.
- A non-linear link function is used between the response y ($=0/1$) and a latent variable y^* which represents (unobserved) net profits.
- y^* is modelled with a SLM
- Other models (SEM, SDM) could be used for the latent variable y^*
- Note also that a GLM can be fit to this data

Software

- The Spatial Econometrics Toolbox (<http://www.spatial-econometrics.com/>) provides Matlab code to fit a wealth of Spatial Econometrics models
- The R software provides number of functions to fit LMs, LMMs, GLMs and GLMMs
- The spdep package implements some functions for spatial econometrics. In particular, it includes functions to fit SEM, SLM and SDM
- Other generic software packages can be used to fit some of the models presented so far

Bayes Inference for Spatial Models

- Bayesian inference is based on Bayes' rule to compute the probability of the parameters in the model (θ) given the observed data (y):

$$\pi(\theta|y) = \frac{\pi(y|\theta)\pi(\theta)}{\pi(y)}$$

- $\pi(y|\theta)$ is the likelihood of the model
- $\pi(\theta)$ is the prior distribution of the parameters in the model
- $\pi(y)$ is a normalising constant that is often ignored
- In spatial statistics, the prior distribution of the random effects can be used to encode spatial dependence
- Vague priors are often used for most parameters in the model

Overview of Bayesian inference

- The aim is computing the (multivariate) posterior probability of θ
- Given that $\pi(\theta|y)$ is a probability distribution, all statements are made in terms of probabilities
- Bayesian inference is 'exact'
- Obtaining $\pi(\theta|y)$ is usually hard
- However, 'recent' computational approaches have made Bayesian inference easier

Model fitting and computational issues

- Fitting a Bayesian model means computing $\pi(\theta|y)$
- θ contains all parameters in the model and, possibly, other derived quantities
- For example, we could compute posterior probabilities of linear predictors, random effects, sums of random effects, etc.
- Depending on the likelihood and the prior distribution computing $\pi(\theta|y)$ can be very difficult
- In the last 20-30 years some computational approaches have been proposed to estimate $\pi(\theta|y)$ with Monte Carlo methods

Markov Chain Monte Carlo

- MCMC is a family of algorithms to obtain draws from the posterior distribution
- In all cases, a starting point is chosen to start the simulation
- At every iteration k we draw a sample of the model parameters $\hat{\theta}_i^{(k)}$ using a particular rule
- After a number of iterations (burn-in period) the algorithm is in fact sampling from $\pi(\theta|y)$
- The iterations generated during the burn-in period are discarded
- A large number of simulations is generated and the posterior distribution is estimated from these samples
- Summary statistics for the model parameters can be easily computed from the simulations

Metropolis-Hasting Sampling

- Generic algorithm to sample from any probability density $f(y)$
- A candidate-generating probability density $q(v|u)$ is required for every parameter in the model
- This will give us the probabilities of sampling v given that we are at u
- We draw a value from this density
- This new value is only accepted with a certain probability, which is

$$\min\left\{1, \frac{\pi(v|y)q(u|v)}{\pi(u|y)q(v|u)}\right\}$$

- Note that

$$\frac{\pi(v|y)q(u|v)}{\pi(u|y)q(v|u)} = \frac{\pi(y|v)\pi(v)q(u|v)}{\pi(y|u)\pi(u)q(v|u)}$$

and that the probability can be computed

Gibbs Sampling

- Particular case of Metropolis-Hastings algorithm
- The proposal distribution is the conditional distribution given the :

$$\pi(\theta_i^{(k+1)} | \theta_1^{(k+1)}, \dots, \theta_{i-1}^{(k+1)}, \theta_{i+1}^{(k)}, \dots, \theta_N^{(k)})$$

- This ensures that the acceptance probability is always 1
- This means that we always accept a new candidate point, i.e., "we always move to a new point"
- Sampling from the conditional probability distribution is usually very easy

Inference with MCMC

- MCMC provides simulations from the ensemble of model parameters, i.e., a multivariate distribution
- This will allow us to estimate the joint posterior distribution
- However, we may be interested in a single parameter or a subset of the parameters
- Inference for this subset of parameters can be done by simply ignoring the samples for the other parameters
- Using the samples it is possible to compute the posterior distribution of any function on the model parameters
- MCMC may require lots of simulations to make valid inference
- Also, we must check that the burn-in period has ended, i.e., we have reached the posterior distribution

Integrated Nested Laplace Approximation

- Sometimes we only need marginal inference on some parameters, i.e., we need $\pi(\theta_i|y)$
- Rue et al. (2009) propose a way of approximating the marginal distributions
- Now we are dealing with dealing with (many) univariate distributions
- This is computationally faster because numerical integration techniques are used instead of Monte Carlo sampling

Integrated Nested Laplace Approximation

- We assume that observations \mathbf{y} are independent given \mathbf{x} (latent effects) and $\theta = (\theta_1, \theta_2)$ (two sets of hyperparameters)
- The model likelihood can be written down as

$$\pi(\mathbf{y}|\mathbf{x}, \theta) = \prod_{i \in \mathcal{I}} \pi(y_i|x_i, \theta)$$

- x_i is the latent linear predictor η_i and other latent effects
- \mathcal{I} represents the indices of the observations (missing observations are not include here, for example)
- $\theta = (\theta_1, \theta_2)$ is a vector of hyperparameters for the likelihood and the distribution of the latent effects

Integrated Nested Laplace Approximation

- \mathbf{x} is assumed to be distributed as a Gaussian Markov Random Field with precision matrix $\mathbf{Q}(\theta_2)$
- The posterior distribution of the model parameters and hyperparameters is:

$$\pi(\mathbf{x}, \theta | \mathbf{y}) \propto \pi(\theta) \pi(\mathbf{x} | \theta) \prod_{i \in \mathcal{I}} \pi(y_i | x_i, \theta) \propto$$

$$\pi(\theta) |\mathbf{Q}(\theta)|^{n/2} \exp\left\{-\frac{1}{2} \mathbf{x}^T \mathbf{Q}(\theta) \mathbf{x} + \sum_{i \in \mathcal{I}} \log(\pi(y_i | x_i, \theta))\right\}$$

Integrated Nested Laplace Approximation

The marginal distributions for the latent effects and hyper-parameters can be written as

$$\pi(x_i|\mathbf{y}) \propto \int \pi(x_i|\theta, \mathbf{y})\pi(\theta|\mathbf{y})d\theta$$

and

$$\pi(\theta_j|\mathbf{y}) \propto \int \pi(\theta|\mathbf{y})d\theta_{-j}$$

Integrated Nested Laplace Approximation

Rue et al. (2009) provide a simple approximation to $\pi(\theta|\mathbf{y})$, denoted by $\tilde{\pi}(\theta|\mathbf{y})$, which is then used to compute the approximate marginal distribution of a latent parameter x_i :

$$\tilde{\pi}(x_i|\mathbf{y}) = \sum_k \tilde{\pi}(x_i|\theta_k, \mathbf{y}) \times \tilde{\pi}(\theta_k|\mathbf{y}) \times \Delta_k$$

Δ_k are the weights of a particular vector of values θ_k in a grid for the ensemble of hyperparameters .

INLA & Spatial econometrics models

- In principle, INLA can handle a large number of models
- The R-INLA package for the R software implements a number of likelihoods and latent effects
- This includes GLMs
- This includes several latent effects for lattice data
- The CAR model is implemented
- SEM, SLM and SDM are not implemented
- The SAR specification is not implemented as a random effects
- Linear predictors are multiplied by $(I - \rho W)^{-1}$, and this is not implemented either
- What to do then?

A possible approach

- Conditioning on ρ , SEM, SLM and SDM become models that R-INLA can fit
- We can fit different models conditioning on different values of ρ . This will provide

$$\pi(\theta|y, \rho = \rho_k), \quad k = 1, 2, \dots$$

- The values of ρ can be chosen equally spaced in $(-1,1)$
- For each fitted model, we can compute the marginal likelihood, i.e., the likelihood of that model: $\pi(y|\rho = \rho_k)$
- Our inference can be based on the model with the largest likelihood
- However, we cannot obtain a marginal distribution for ρ and cannot compute summary statistics

Bayesian Model Averaging

- A better approach is to combine the different fitted models in some way
- Bayesian Model Averaging provides a way of combining all these models
- For each fitted model (conditioned on a value of ρ) we have $\pi(\theta_i|y, \rho = \rho_k)$ and $\pi(y|\rho = \rho_k)$
- We can choose a prior distribution for ρ : $\pi(\rho)$
- It should be noted that the marginal distribution of ρ is

$$\pi(\rho|y) = \frac{\pi(y|\rho)\pi(\rho)}{\pi(y)} \propto \pi(y|\rho)\pi(\rho)$$

- The marginal distribution for ρ can be computed by fitting a curve to the values

$$[\rho_k, \pi(y|\rho = \rho_k)\pi(\rho = \rho_k)]$$

Marginal distribution of ρ

- The marginal distribution of a parameter can be written as

$$\pi(\theta_i|y) = \int \pi(\theta_i, \rho|y) d\rho = \int \pi(\theta_i|y, \rho) \pi(\rho|y) d\rho$$

- The previous integral can be approximated as follows:

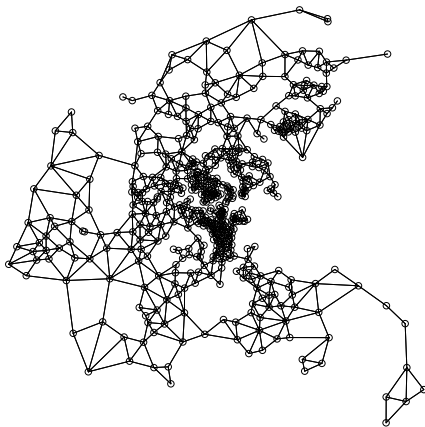
$$\sum_k \pi(\theta_i|y, \rho = \rho_k) \frac{\pi(y|\rho = \rho_k) \pi(\rho_k)}{\sum_k \pi(y|\rho = \rho_k) \pi(\rho_k)} = \sum_k w_k \pi(\theta_i|y, \rho = \rho_k)$$

- Finally, a spline can be fitted to the resulting function so that it can be used to compute other quantities, such as the mean, mode, quantiles, etc.

Example: Boston housing data

In our first example we will re-analyse the Boston housing data (Harrison and Rubinfeld, 1978). Here the interest is in estimating the median of owner-occupied houses using relevant covariates and the spatial structure of the data (Pace and Gilley, 1997). We have fitted the 3 spatial econometrics models described in this paper plus a spatial model with a CAR error term. In addition, we have fitted the spatial econometrics models using maximum likelihood to compare the estimates of the model parameters

Boston housing data: Adjacency matrix



Boston housing data: Fixed effects (I)

	MLCAR	INLACAR	MLSEM	INLASEM	MLSLM	INLASLM
(Intercept)	4.11e+00	4.116596	3.84e+00	3.923540	2.279626	2.308686
CRIM	-8.14e-03	-0.008197	-5.29e-03	-0.005886	-0.007105	-0.007267
ZN	1.56e-04	0.000150	4.73e-04	0.000372	0.000380	0.000365
INDUS	-4.87e-05	-0.000019	-2.52e-05	0.000158	0.001257	0.001376
CHAS1	3.18e-02	0.033327	-3.88e-02	-0.008519	0.007368	0.012609
I(NOX^2)	-4.61e-01	-0.467009	-2.23e-01	-0.339600	-0.268916	-0.286864
I(RM^2)	8.08e-03	0.008042	7.96e-03	0.008052	0.006724	0.006763
AGE	-4.88e-04	-0.000478	-1.05e-03	-0.000787	-0.000277	-0.000226
log(DIS)	-1.64e-01	-0.164862	-1.18e-01	-0.136482	-0.158301	-0.157636
log(RAD)	7.83e-02	0.078420	6.55e-02	0.074014	0.070689	0.073167
TAX	-3.97e-04	-0.000399	-5.00e-04	-0.000475	-0.000366	-0.000364
PTRATIO	-2.10e-02	-0.021258	-1.77e-02	-0.020105	-0.012011	-0.012743
B	5.07e-04	0.000507	5.94e-04	0.000585	0.000284	0.000289
log(LSTAT)	-3.21e-01	-0.321995	-2.66e-01	-0.276658	-0.232161	-0.232201


Boston housing data: Fixed effects

	MLSDM	INLASDM	MLSDM (LAG)	INLASDM (LAG)
(Intercept)	1.898178	1.996513		
CRIM	-0.005710	-0.005832	-0.004642	-0.004722
ZN	0.000691	0.000630	-0.000379	-0.000257
INDUS	-0.001113	-0.001495	0.000251	0.001275
CHAS1	-0.041632	-0.028907	0.125183	0.100803
I(NOX ²)	-0.010349	-0.006888	-0.386407	-0.418605
I(RM ²)	0.007950	0.007825	-0.004513	-0.004207
AGE	-0.001288	-0.001214	0.001497	0.001523
log(DIS)	-0.124041	-0.122187	-0.004539	-0.011010
log(RAD)	0.058635	0.055938	-0.009407	0.003412
TAX	-0.000491	-0.000474	0.000411	0.000350
PTRATIO	-0.013199	-0.013531	0.000603	-0.000588
B	0.000564	0.000540	-0.000508	-0.000429
log(LSTAT)	-0.247245	-0.249881	0.098467	0.091499

Boston housing data: Estimates of ρ

	rho	rhosd	LLrho	ULrho
MLCAR	0.186	NA	NA	NA
INLACAR	0.970	0.0127	0.938	0.987
MLSEM	0.715	NA	NA	NA
INLASEM	0.687	0.0318	0.623	0.749
MLSLM	0.485	0.0294	0.428	0.543
INLASLM	0.477	0.0274	0.423	0.529
MLSDM	0.596	0.0384	0.520	0.671
INLASDM	0.571	0.0378	0.496	0.643

Example: Boston housing data




Example: New Orleans business data

In this example we will look at the data analysed in LeSage et al. (2011) regarding the probability of re-opening a business in the aftermath of hurricane Katrina. In this case we have a non-Gaussian model because we are modelling a probability and the response variable can take either 1 (the business re-opened) or 0 (the business didn't re-open). Similarly as in the previous example, we have fitted four models. However, now we have used a GLM with a Binomial family and a probit link.

LeSage et al. (2011) split the data into four periods according to different time frames. In our analysis we will focus on the first period, i.e., the business re-opened during the first 3 months (90 days). The model used therein is the one that we have termed Spatial Lag Model in this paper.

New Orleans business data: Adjacency matrix



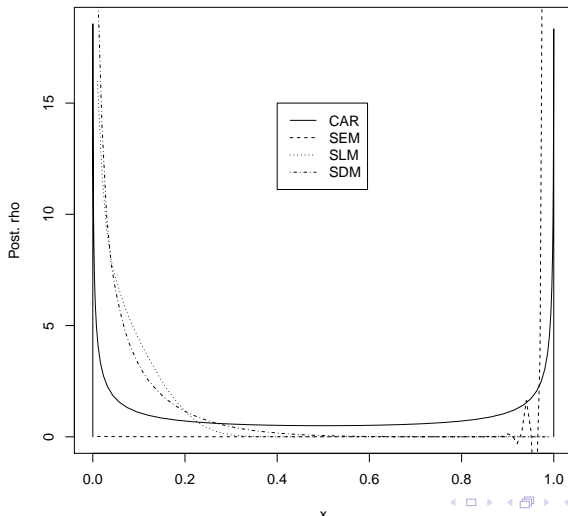
New Orleans business data: fixed effects

	INLACAR	INLASEM	INLASLM	INLASDM
(Intercept)	-8.5136	-8.6076	-11.3811	-7.8558
flood	0.2691	0.3255	0.2862	0.2892
log(medinc)	0.9004	0.8356	1.1083	1.2858
sizeempsmall	-0.1318	-0.2718	-0.3082	-0.2881
sizeemplarge	-0.4292	-0.2828	-0.2933	-0.2982
sesstatuslow	-0.5072	-0.2534	-0.4243	-0.0949
sesstatushigh	0.0657	0.0699	0.0658	0.0807
owntypesole	0.3656	0.6063	0.6284	0.6382
owntypenational	0.1972	0.2098	0.1301	0.1155

New Orleans business data: Estimates of ρ

	rho	rhosd	LLrho	ULrho
INLACAR	0.5002	0.3738	0.002004	0.998
INLASEM	1.6094	0.7975	0.974572	0.990
INLASLM	0.0801	0.0636	0.001155	0.127
INLASDM	0.0918	0.0963	0.000846	0.117

New Orleans business data



Other issues and current work

- Impacts are difficult to compute as their are based on bivariate inference:

$$\frac{\partial y_i}{\partial x'_{v,j}}$$

- Exploit INLA for bivariate inference
- Computational issues with INLA with the Probit model
- Simulations to assess differences between GLMs and Spatial Econometrics models
- Develop R code similar to the Spatial Econometrics Toolbox (GSoc project) to compare with our results with INLA

Concluding remarks

- INLA provides a novel approach to Bayesian inference based on approximating the marginal distribution of the model parameters
- The R-INLA software can fit many different types of spatial models
- We have proposed a new way of extending R-INLA to cover some widely used spatial econometrics models
- In our approach, we fit many conditioned models and combine them using Bayesian Model Averaging
- For Gaussian models, our implementation and results match those of ML estimation
- For the Probit models our approach seems feasible as well
- INLA may struggle when providing good estimates of some derived quantities, such as the impacts

References

- Haining, R. (2003). *Spatial Data Analysis: Theory and Practice*. Cambridge University Press.
- Harrison, D. and D. L. Rubinfeld (1978). Hedonic housing prices and the demand for clean air. *Journal of Environmental Economics and Management* 5, 81–102.
- LeSage, J. P., R. Kelley Pace, N. Lam, R. Campanella, and X. Liu (2011). New orleans business recovery in the aftermath of hurricane katrina. *Journal of the Royal Statistical Society: Series A (Statistics in Society)* 174(4), 1007–1027.
- Pace, R. K. and O. Gilley (1997). Using the spatial configuration of the data to improve estimation. *Journal of the Real Estate Finance and Economics* 14, 333–340.
- Rue, H., S. Martino, and N. Chopin (2009). Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations. *JOURNAL OF THE ROYAL STATISTICAL SOCIETY SERIES B-STATISTICAL METHODOLOGY* 71(Part 2), 319–392.