# Performance of INLA analysing bivariate meta-regression and age-period-cohort models

Andrea Riebler

Biostatistics Unit, Institute of Social and Preventive Medicine
University of Zurich

INLA workshop, May 2009

Joint work with Lucas Bachmann, Leonhard Held, Michaela Paul and Håvard Rue

# Outline

# 1. Introduction

## Bivariate meta-analysis

Comparison of the performance of `inla` and the performance obtained by the maximum likelihood procedure `SAS PROC NLMIXED` (Paul et al., 2009).

## Age-period-cohort models

Comparison of the performance of `inla` and an MCMC algorithm implemented in C using the `GMRFLib` library (Rue and Held, 2005, Appendix).

All analyses were run under Kubuntu 8.04 on a laptop with Intel(R) Core(TM) 2 Duo T7200 processor with 2.00 GHz.


University of Zurich

# Bivariate meta-analysis

Meta-analyses are used to summarise the results of separately performed studies, here diagnostic studies.

Diagnostic studies often report two-by-two tables

$$\Rightarrow \text{Sensitivity Se} = \frac{\text{TP}}{\text{TP}+\text{FN}} \text{ and specificity Sp} = \frac{\text{TN}}{\text{TN}+\text{FP}}.$$

## Bivariate meta-analysis:

Models the relationship between sensitivity and specificity (after logit transformation), including random effects for both and allowing for correlation between them.

Focus: Estimation of the expected sensitivity and specificity

TP = true positives, FP = false positives, TN = true negatives, FN = false negatives.


University of Zurich

# Model formulation

## 1. Level

$$\mathsf{TP}_i \,|\, \mathsf{Se}_i \sim \mathsf{Binomial}(\mathsf{TP}_i + \mathsf{FN}_i, \mathsf{Se}_i)$$
$$\mathsf{TN}_i \,|\, \mathsf{Sp}_i \sim \mathsf{Binomial}(\mathsf{TN}_i + \mathsf{FP}_i, \mathsf{Sp}_i)$$

## 2. Level

$$\begin{aligned} \mathsf{logit}(\mathsf{Se}_i) &= \mu + \mathbf{U}_i\boldsymbol{\alpha} + \phi_i, \\ \mathsf{logit}(\mathsf{Sp}_i) &= \nu + \mathbf{V}_i\boldsymbol{\beta} + \psi_i, \end{aligned} \quad \text{with} \quad \begin{pmatrix} \phi_i \\ \psi_i \end{pmatrix} \sim \mathcal{N}\left[ \begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1/\tau_\phi & \rho/\sqrt{\tau_\phi\tau_\psi} \\ \rho/\sqrt{\tau_\phi\tau_\psi} & 1/\tau_\psi \end{pmatrix} \right],$$

where $i = 1, \ldots, I$ is the study index (Chu and Cole, 2006).

# Inference

## Likelihood approaches

- Numerical maximisation might fail in complex problems.
- Construction of confidence intervals is problematic.

## Bayesian approaches

- Markov chain Monte Carlo (MCMC) is very time-consuming.
- Credible intervals are obtained as the quantiles of the samples.

Comparison of `inla` and `SAS PROC NLMIXED` using an extensive simulation study.

# Simulation study

72 different scenarios where each scenario contains 1000 meta-analyses sampled from the model.

We varied

- the number of studies per meta-analysis.
- the overall sensitivity and specificity.
- the between-studies precisions.
- the correlation between logit sensitivity and logit specificity.

The number of participants is sampled for each study separately.

# Settings

In a Bayesian context all parameters are treated as random and prior distributions are assigned (determined by a sensitivity analysis):

- For $\tau_\phi, \tau_\psi$: Gamma(shape=0.25, rate=0.025).
- For Fisher's z-transformed correlation $\tilde{\rho}$:

$$\tilde{\rho} \sim \mathcal{N}(0, 0.2^{-1})$$

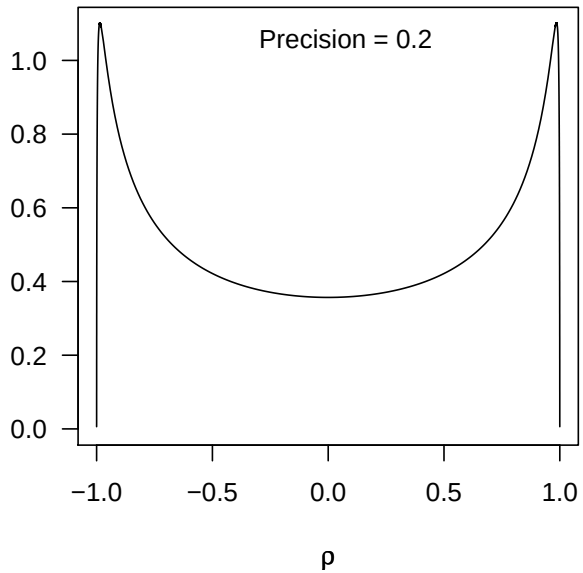


University of Zurich

# Results

Comparison using bias, SD, MSE and coverage probabilities:

- Bias and MSE of `inla` and `NLMIXED` are almost the same.
- Bias and MSE depend on choice of sensitivity and specificity.
- The estimates are more precise for more studies.
- Precision of estimates and MSE are hardly influenced by the value of $\rho$.
- In general `inla` produces better coverage.

# Performance and running time

## Performance:

Out of 72 000 analyses

- `inla` failed 2 times,
- `NLMIXED` failed 7 482 times (10.4%).

## Running time:

For one scenario of 1000 meta-analyses

- `inla` took on average 6.0 minutes (min: 4.7, max: 7.8),
- `NLMIXED` took on average 38.1 minutes (min: 20.5, max: 89.3).

# Radiological evaluation of lymph node metastases

Three types of diagnostic imaging are compared for detecting lymph node metastases in patients with cervical cancer (Scheidler et al., 1997).

The meta-analysis consists of a total of 46 studies:

- 17 studies for lymphangiography (LAG)
- 19 studies for computed tomography (CT)
- 10 studies for magnetic resonance (MR)

with each containing at least 20 patients.

# INLA call using the R-Interface

```
> library(INLA)

> data(BivMetaAnalysis)
> head(BivMetaAnalysis)

  N  Y diid lag.tp lag.tn ct.tp ct.tn mr.tp mr.tn
1 29 19    1      1      0     0     0     0     0
2 82 81    2      0      1     0     0     0     0
3 10  8    3      1      0     0     0     0     0
4 22 13    4      0      1     0     0     0     0
5 53 41    5      1      0     0     0     0     0
6 50 49    6      0      1     0     0     0     0

> formula <- Y ~ f(diid, model = "2diid",
+   param = c(0.25, 0.025, 0.25, 0.025, 0, 0.2)) +
+   lag.tp + lag.tn + ct.tp + ct.tn + mr.tp + mr.tn - 1

> model <- inla(formula, family = "binomial", Ntrials = N,
+   data = BivMetaAnalysis, quantiles = c(0.025, 0.5, 0.975))
```

The analysis took about $\sim$ 0.6 seconds.


University of Zurich

# Summary estimates

| Imaging | Sensitivity | | |
|---|---|---|---|
| | Median | 2.5%-quantile | 97.5%-quantile |
| LAG | 0.69 | 0.57 | 0.80 |
| CT | 0.49 | 0.36 | 0.62 |
| MR | 0.55 | 0.37 | 0.71 |

| Imaging | Specificity | | |
|---|---|---|---|
| | Median | 2.5%-quantile | 97.5%-quantile |
| LAG | 0.83 | 0.76 | 0.89 |
| CT | 0.93 | 0.89 | 0.96 |
| MR | 0.95 | 0.91 | 0.98 |

The correlation $\rho$ was estimated to $-0.48$ $(-0.76, -0.04)$.

University of Zurich

# Discussion

Similar performance of `inla` and `NLMIXED` regarding bias and MSE.

## Advantage of `inla`

- Better coverage
- More stable and faster

Since sensitivity and specificity are jointly analysed, a joint confidence ellipse for these measures might be of interest.

Comparison of `NLMIXED` and `inla` using
an empirical Bayes approach?

# 3. Age-period-cohort model

Data on cancer often consist of yearly counts for different age groups and gender in pre-defined geographical areas.

Our goal lies in:

- Detecting temporal patterns.
- Providing predictions for subsequent periods.

## Age-period-cohort (APC) model

to describe incidence or mortality rates using three time scales.

- **A**ge: age at diagnosis.
- **P**eriod: date of diagnosis.
- **C**ohort: date of birth.

# Univariate age-period-cohort model

$y_{ij}$: number of deaths or disease cases in age group $i$ at period $j$
$n_{ij}$: number of persons at risk in age group $i$ at period $j$

$$y_{ij} \sim \text{Poisson}(n_{ij} \exp(\xi_{ij})) \qquad \xi_{ij} = \mu + \alpha_i + \beta_j + \gamma_k + z_{ij}$$

with age effect $\alpha_i$, period effect $\beta_j$, cohort effect $\gamma_k$ and additional random effect $z_{ij} \sim \mathcal{N}(0, \delta^{-1})$ to adjust for overdispersion.

To assure identifiability of the intercept $\mu$, we set

$$\sum_{i=1}^{I} \alpha_i = \sum_{j=1}^{J} \beta_j = \sum_{k=1}^{K} \gamma_k = 0.$$

Note: Because of the linear relationship $k = I - i + j$, the age, period and cohort effects are still not identifiable


University of Zurich

# Bayesian age-period-cohort model

Non-parametric smoothing priors are used for the main effects with gamma hyperpriors for the associated smoothing parameters.

## Second-order random walk (RW2)

$$\alpha_i \sim \mathcal{N}(2\alpha_{i-1} - \alpha_{i-2}, \kappa^{-1}) \qquad i = 3, \ldots, I$$

RW2 penalises deviations from a linear trend $\alpha_i = 2\alpha_{i-1} - \alpha_{i-2}$.

## Note:

Non-identifiability of the latent parameters remains, but does not require further constraints.

# Case study: Lung cancer mortality in West Germany

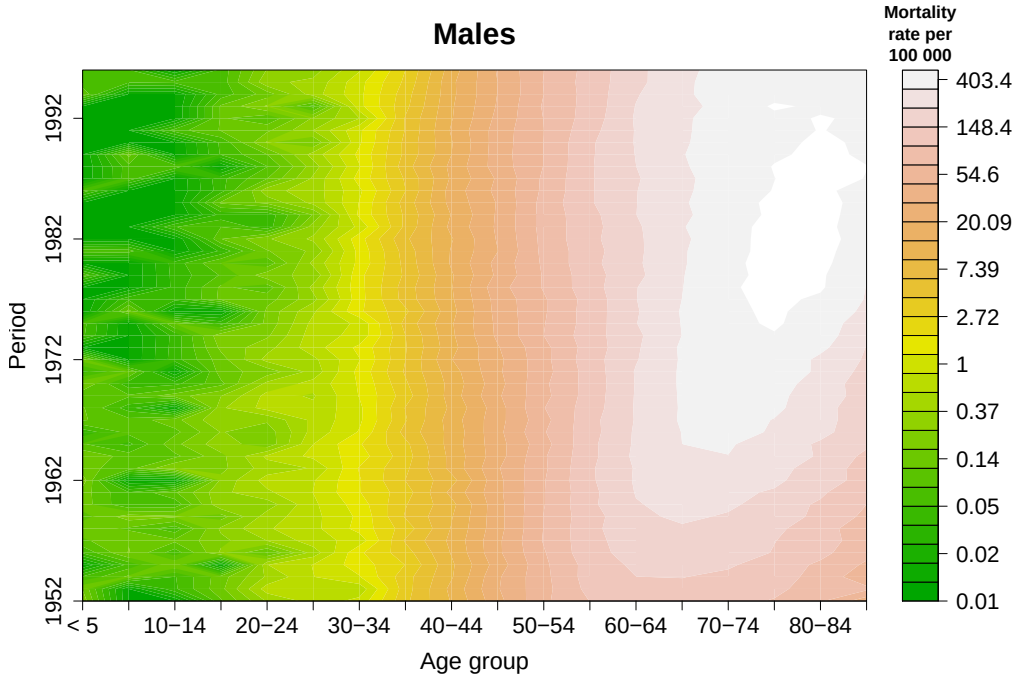


- 18 age groups: $< 5$, 5-9, 10-14,. . . , 80-84, $\geq 85$.
- 45 periods: 1952 - 1996.
- 130 cohorts: 1862-1867, 1863-1868, . . . , 1991-1996.

(Knorr-Held and Rainer, 2001)

University of Zurich

# INLA call using the R-Interface

| y | n | i | j | k | z |
|---|---|---|---|---|---|
| 3 | 250 | 1 | 1 | 2 | 1 |
| 20 | 260 | 2 | 1 | 1 | 2 |
| 9 | 230 | 1 | 2 | 3 | 3 |
| 12 | 270 | 2 | 2 | 2 | 4 |
| 7 | 260 | 1 | 3 | 4 | 5 |
| 10 | 290 | 2 | 3 | 3 | 6 |
| ⋮ | | | | | |

For predictions, set $y_{ij} = \text{NA}$.

```
> library(INLA)
> lungm <- read.table("data/lungm4inla.txt", header=T)
> formula <- y ~ f(i, model="rw2", param=c(1,0.00005)) +
+           f(j, model="rw2", param=c(1,0.00005)) +
+           f(k, model="rw2", param=c(1,0.00005)) +
+           f(z, model="iid", param=c(1,0.005))
> model <- inla(formula, family="poisson", data=lungm, E=lungm$n,
+            quantiles=c(0.1, 0.5, 0.9), control.compute=list(cpo=TRUE),
+            control.predictor=list(compute=TRUE))
> hyper <- inla.hyperpar(model)
```

# Results for complete dataset

- MCMC needed for 120 000 iterations about 10 minutes.
- INLA needed for the model estimation about 17 seconds and for the improved hyperparameter estimation about 2 minutes.

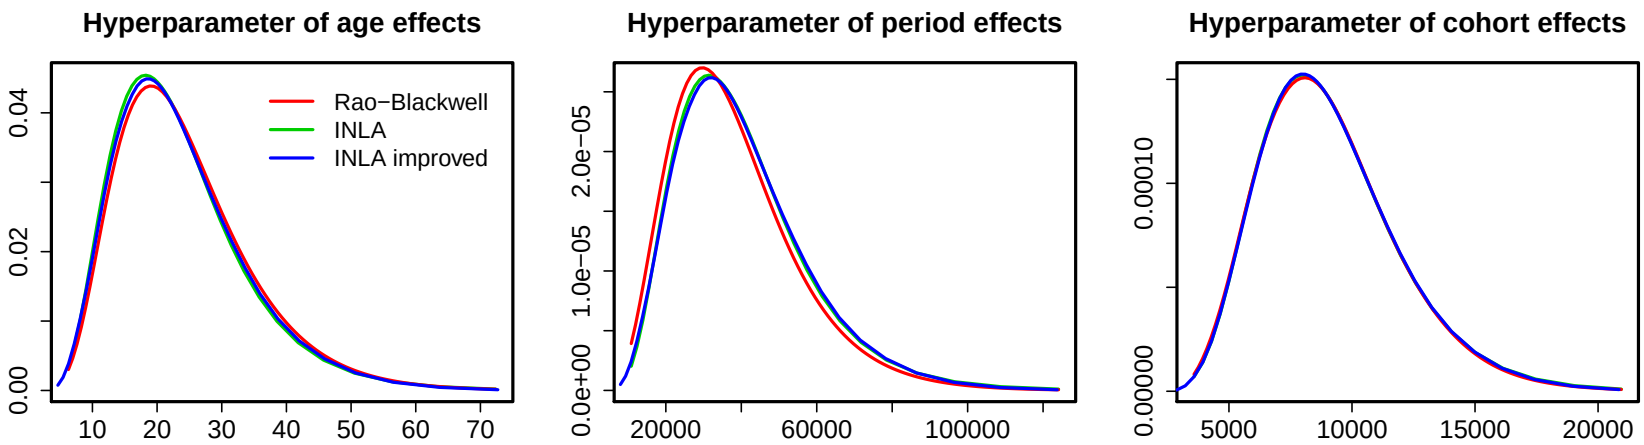


The inspection of identifiable measures gave similar results.

University of Zurich

# Predictions for 1987 - 1996

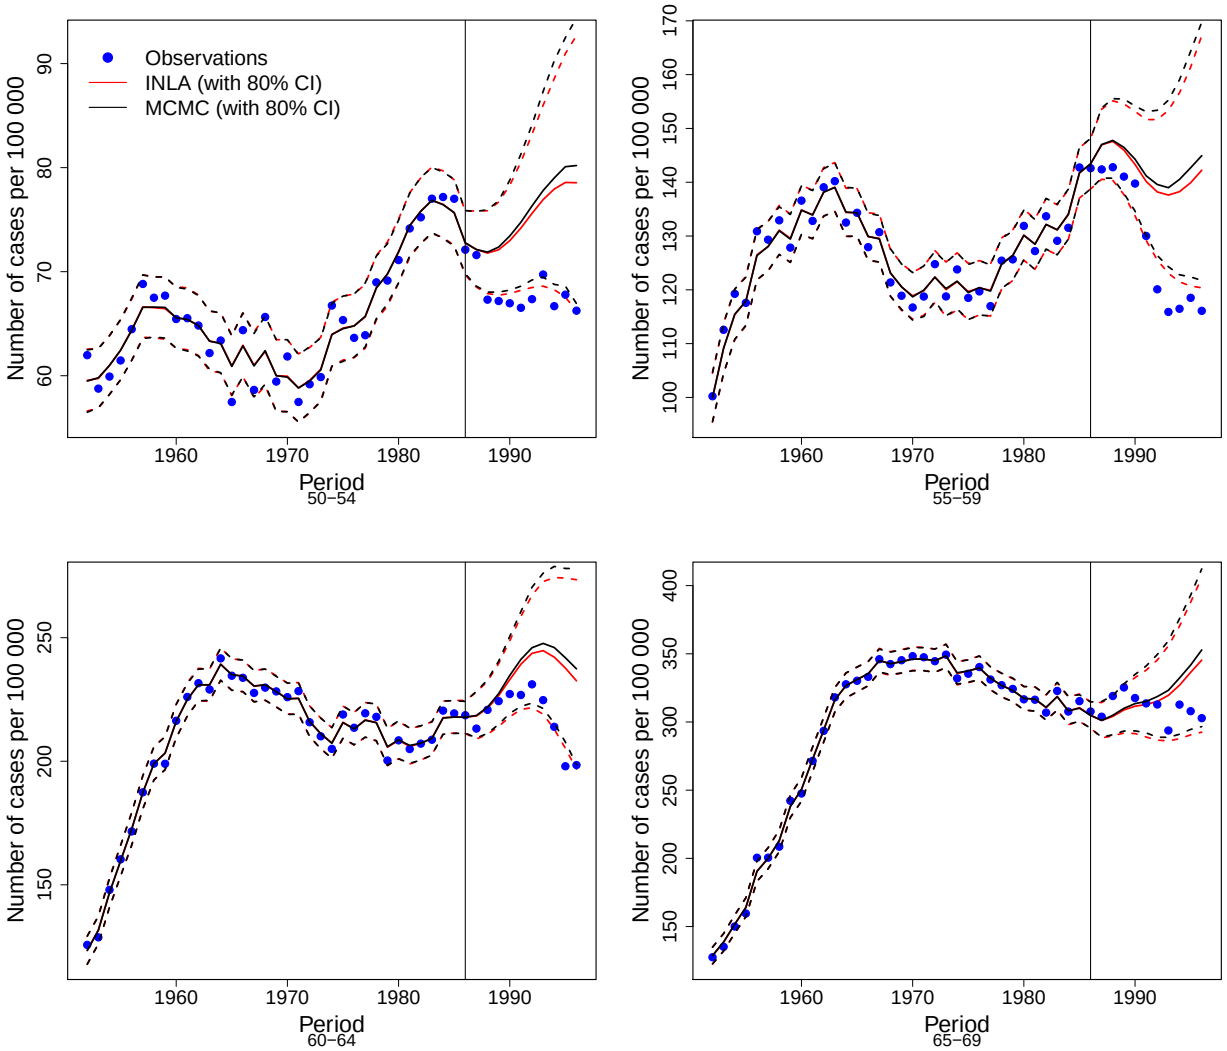


University of Zurich

# Inclusion of smoking data in the APC model

The inclusion of appropriate covariate information in the APC model could improve the predictions.

## Model formulation:

Assuming a time-constant effect $\beta$:

$$\xi_{ij} = \mu + \alpha_i + \beta \cdot x_{j-L} + \gamma_k + z_{ij}.$$

Assuming a time-varying effect $\beta_j$:

$$\xi_{ij} = \mu + \alpha_i + \beta_j \cdot x_{j-L} + \gamma_k + z_{ij},$$

assigning a RW2 smoothing prior to $\beta_j$.

- $x_j$: number of cigarettes sold per 1/1000 capita in 1955-1994.
- $L = 20$ years: latency period.

# Inclusion of covariates in R-inla

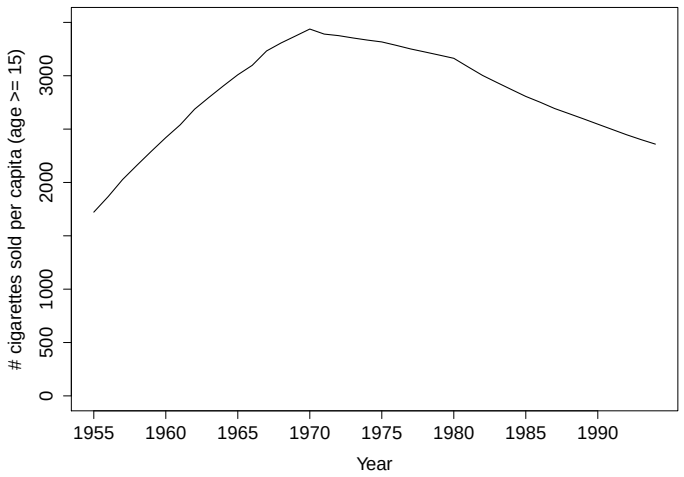


Goal: Prediction until 2010.

Note: Because of $L = 20$ years, only data from 1975 onwards can enter.

- Assuming a time-constant effect $\beta$:

```
formula_const <- y ~ f(i, model="rw2", param=c(1,0.00005), constr=1) +
  f(k, model="rw2", param=c(1,0.00005), constr=1) +
  f(z, model="iid", param=c(1,0.005) ) + cig_cov
```

- Assuming a time-varying effect $\beta_j$:

```
formula_vary <- y ~ f(i, model="rw2", param=c(1,0.00005), constr=1) +
  f(j, model="rw2", param=c(1,0.00005), constr=0, weights=cig_cov) +
  f(k, model="rw2", param=c(1,0.00005), constr=1) +
  f(z, model="iid", param=c(1,0.005) )
```

University of Zurich

# Covariate effects

Time constant effect $\exp(\beta)$:

| 10%-quantile | Median | 90%-quantile |
|---|---|---|
| 1.11 | 1.13 | 1.15 |

Time-varying effect $\exp(\beta_j)$:

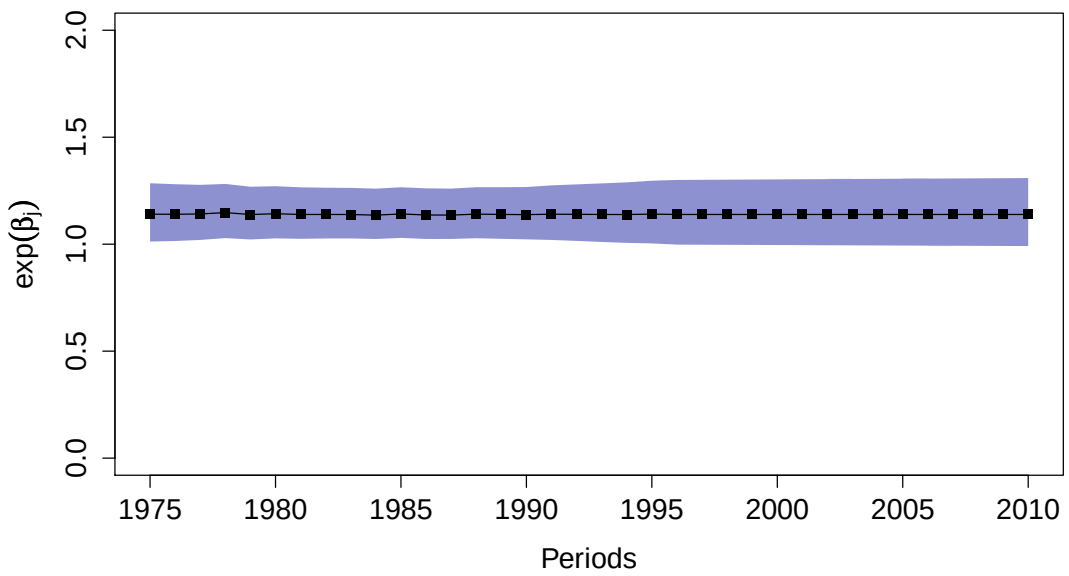


University of Zurich

# Model assessment

- PIT histogram for count data (Czado et al. 2009):

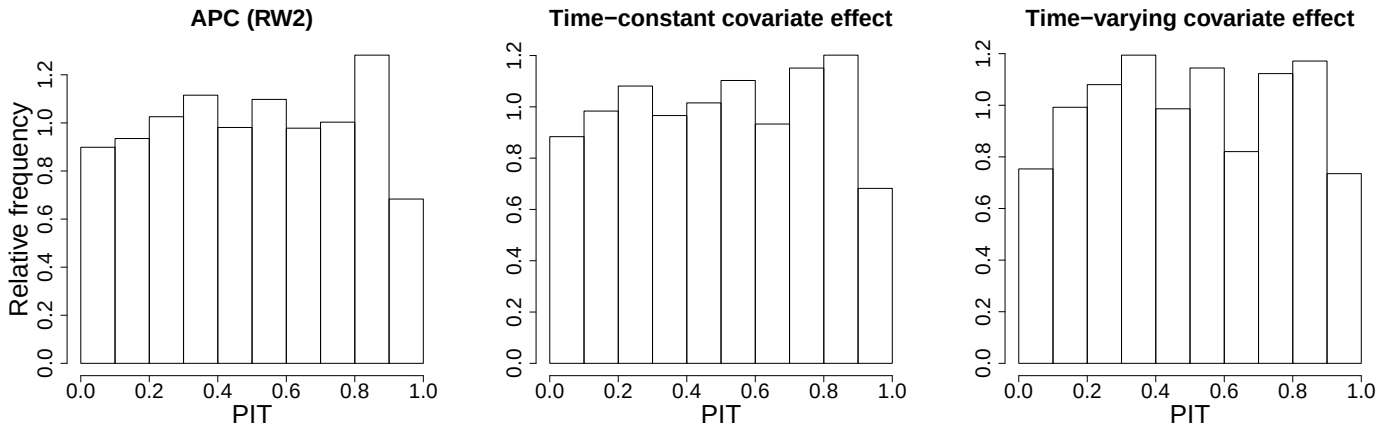


- Log-score:

| | APC | constant | RW2 |
|---|---|---|---|
| $-\overline{\log(\text{CPO})}$ | 3.895* | 3.887* | 3.905* |

*Two CPO values were removed as they were classified as unreliable.

# Prediction until 2010

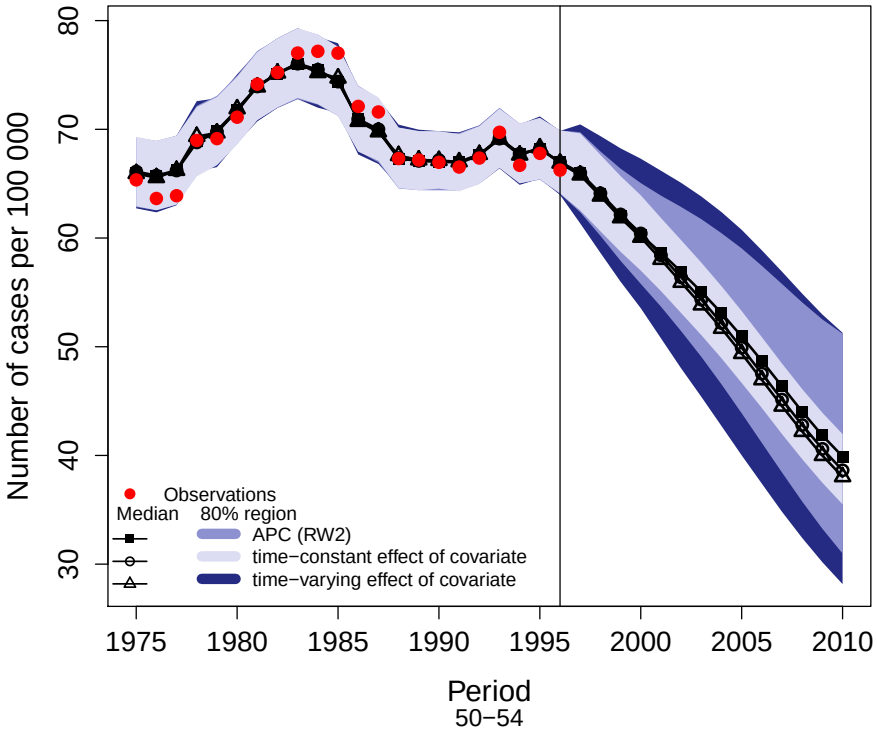


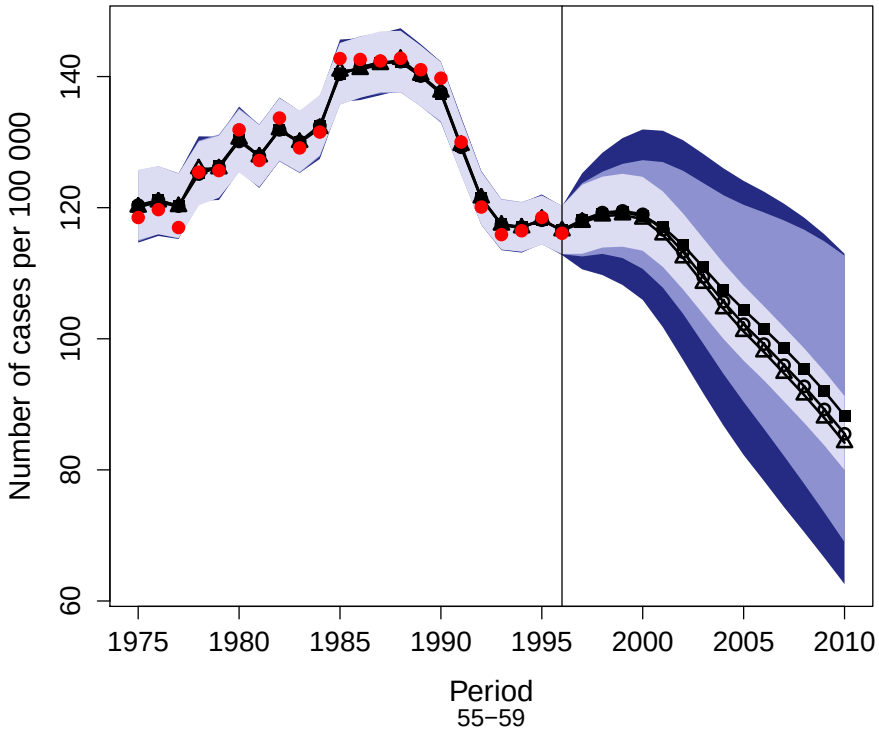


University of Zurich

# Discussion

- INLA facilitates the analysis of Bayesian APC models.
- Prediction is straightforward.
- Covariate information can be easily incorporated.
- Model diagnostics available, but not completely robust.

# 4. Summary

For both applications presented, INLA is an alternative to the standard used inference approaches (ML, MCMC). It is:

- User-friendly and easy to apply
- Fast
- Flexible

Issues for future work might be:

- Improved model diagnostics,
- Calculation of joint credibility intervals,
- Calculation of predictive distribution for response.

# Thank you for your attention

Chu, H. and Cole, S.R. (2006).
Bivariate meta-analysis of sensitivity and specificity with sparse data: a generalised linear mixed model approach. *Journal of Clinical Epidemiology*, 59, 1331–1333.

Czado, C., Gneiting, T. and Held, L. (2009).
Predictive Model Assessment for Count Data. *Biometrics*, to appear.

Knorr-Held, L. and Rainer, E. (2001).
Projections of lung cancer mortality in West Germany: a case study in Bayesian prediction. *Biostatistics*, 2, 109–129.

Paul, M., Riebler, A., Bachmann, L., Rue, H. and Held, L. (2009). Bivariate meta-analysis with INLA: an approximate Bayesian inference. *Statistics in Medicine*, submitted.

Rue, H. and Held, L. (2005). Gaussian Markov Random Fields: Theory and Applications. Volume 104 of *Monographs on Statistics and Applied Probability*, Chapmann & Hall/CRC.

Rue, H., Martino, S. and Chopin, N. (2009). Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations (with discussion). *Journal of the Royal Statistical Society: Series B*, 71, 319–392.

Scheidler, J., Hricak, H., Yu, K. K., Subak, L. Segal, M. R. (1997). Radiological evaluation of lymph node metastases in patients with cervical cancer. *Journal of the American Medical Association*,278, 1096–1101.


University of Zurich