



## *bamlss*: A Lego toolbox for flexible regression models

Nikolaus Umlauf

<http://eeecon.uibk.ac.at/~umlauf/>

# Overview

Joint work with Achim Zeileis, Nadja Klein, Meike Köhler, Thorsten Simon and Stanislaus Stadlmann.

- ① Introduction
- ② Model specification
- ③ Model fitting
- ④ R package *bamlss*
- ⑤ Application


# Introduction

- **Computational power** has tremendously increased.
- **Complicated inferential problems**, e.g., with MCMC simulation, possible on virtually any modern computer.
- To embed many **different approaches** suggested in literature and software, a **unified modeling architecture** for flexible regression models is particularly helpful.
- With the *bamlss* framework, implementing (new) algorithms, integration of already existing software, is relatively straightforward.
- The original idea came from flexible **Bayesian distributional regression** models.

# Introduction

Linz daily maximum temperature data (2000/01-2018/09)


$$T \sim N(\mu, \sigma^2).$$



# Introduction

Linz daily maximum temperature data (2000/01-2018/09)


$$T \sim N(\mu = f(T_{t-1}), \log(\sigma^2) = \beta_0).$$



# Introduction

Linz daily maximum temperature data (2000/01-2018/09)


$$T \sim N(\mu = f(T_{t-1}), \log(\sigma^2) = f(T_{t-1})).$$



# Introduction

Linz daily maximum temperature data (2000/01-2018/09)

$$T \sim N(\mu = f(T_{t-1}), \log(\sigma^2) = f(T_{t-1})).$$



# Model specification

Any parameter of a population distribution  $\mathcal{D}$  may be modeled by explanatory variables

$$y \sim \mathcal{D} (h_1(\theta_1) = \eta_1, h_2(\theta_2) = \eta_2, \dots, h_K(\theta_K) = \eta_K),$$

Each parameter is linked to a structured additive predictor

$$h_k(\theta_k) = \eta_k = \eta_k(\mathbf{x}; \boldsymbol{\beta}_k) = f_{1k}(\mathbf{x}; \boldsymbol{\beta}_{1k}) + \dots + f_{J_k k}(\mathbf{x}; \boldsymbol{\beta}_{J_k k}),$$

$j = 1, \dots, J_k$  and  $k = 1, \dots, K$  and  $h_k(\cdot)$  are known monotonic link functions.


Vector of function evaluations  $\mathbf{f}_{jk} = (f_{jk}(\mathbf{x}_1; \boldsymbol{\beta}_{jk}), \dots, f_{jk}(\mathbf{x}_n; \boldsymbol{\beta}_{jk}))^\top$

$$\mathbf{f}_{jk} = \begin{pmatrix} f_{jk}(\mathbf{x}_1; \boldsymbol{\beta}_{jk}) \\ \vdots \\ f_{jk}(\mathbf{x}_n; \boldsymbol{\beta}_{jk}) \end{pmatrix} = f_{jk}(\mathbf{X}_{jk}; \boldsymbol{\beta}_{jk}).$$




# Model specification


**Nonlinear effects of continuous covariates**




**Two-dimensional surfaces**



**Spatially correlated effects  $f(x) = f(s)$**




**Random intercepts  $f(x) = f(id)$**





# Model specification

Model terms  $f_{jk}(\mathbf{x}; \beta_{jk})$  with LASSO-type penalties  $J_c(\beta_{jk})$ .



# Model specification

Model terms  $f_{jk}(\mathbf{x}; \beta_{jk})$  with LASSO-type penalties  $J_f(\beta_{jk})$ .



# Model fitting

The main building block of regression model algorithms is the probability density function  $d_y(\mathbf{y}|\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_K)$ .

Estimation typically requires to evaluate

$$\ell(\boldsymbol{\beta}; \mathbf{y}, \mathbf{X}) = \sum_{i=1}^n \log d_y(y_i; \theta_{i1} = h_1^{-1}(\eta_{i1}(\mathbf{x}_i, \boldsymbol{\beta}_1)), \dots, \dots, \theta_{iK} = h_K^{-1}(\eta_{iK}(\mathbf{x}_i, \boldsymbol{\beta}_K))),$$

with  $\boldsymbol{\beta} = (\boldsymbol{\beta}_1^\top, \dots, \boldsymbol{\beta}_K^\top)^\top$  and  $\mathbf{X} = (\mathbf{X}_1, \dots, \mathbf{X}_K)$ .

The log-posterior

$$\log \pi(\boldsymbol{\beta}, \boldsymbol{\tau}; \mathbf{y}, \mathbf{X}, \boldsymbol{\alpha}) \propto \ell(\boldsymbol{\beta}; \mathbf{y}, \mathbf{X}) + \sum_{k=1}^K \sum_{j=1}^{J_k} [\log p_{jk}(\boldsymbol{\beta}_{jk}; \boldsymbol{\tau}_{jk}, \boldsymbol{\alpha}_{jk})],$$

where  $\boldsymbol{\tau} = (\boldsymbol{\tau}_1^\top, \dots, \boldsymbol{\tau}_K^\top)^\top = (\boldsymbol{\tau}_{11}^\top, \dots, \boldsymbol{\tau}_{J_1 1}^\top, \dots, \boldsymbol{\tau}_{1K}^\top, \dots, \boldsymbol{\tau}_{J_K K}^\top)^\top$   
(frequentist, penalized log-likelihood).

# Model fitting


A simple generic algorithm for flexible regression models:

```
1  while(eps > ε & t < maxit) {
2    for(k in 1:K) {
3      for(j in 1:J[k]) {
4        Compute  $\tilde{\eta} = \eta_k - \mathbf{f}_{jk}$ .
5        Obtain new  $(\beta_{jk}^*, \tau_{jk}^*)^\top = U_{jk}(\mathbf{X}_{jk}, \mathbf{y}, \tilde{\eta}, \beta_{jk}^{[t]}, \tau_{jk}^{[t]}, \alpha_{jk})$ .
6        Update  $\eta_k = \tilde{\eta} + \mathbf{f}_{jk}^*$ .
7      }
8    }
9    t = t + 1
10   Compute new eps.
11 }
```

Functions  $U_{jk}(\cdot)$  could either return updates from an optimizing algorithm or proposals from a MCMC sampler.

# Model fitting

For complicated models use combination of algorithms, e.g., gradient boosting for finding starting values for MCMC.



# R package *bamlss*

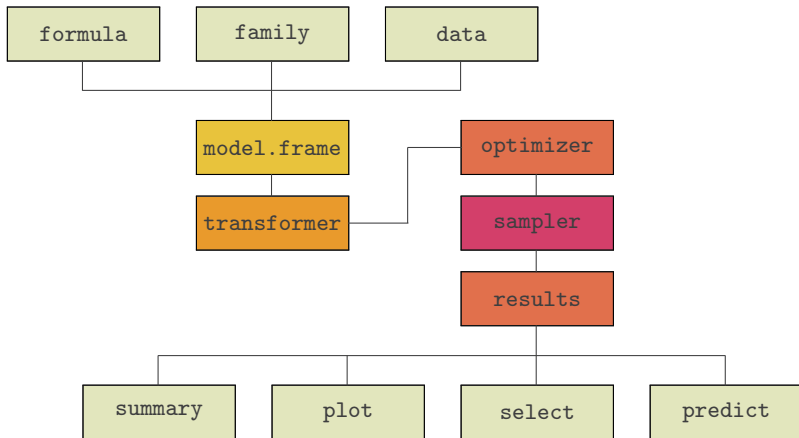
The package is available at

<https://CRAN.R-project.org/package=bamlss>

Development version, in R simply type

```
R> install.packages("bamlss",  
+   repos = "http://R-Forge.R-project.org")
```

# R package *bamlss*



In principle, the setup does not restrict to any specific type of engine (Bayesian or frequentist).



## R package *bamlss*

Type	Function
Parser	<code>bamlss.frame()</code>
Transformer	<code>bamlss.engine.setup()</code> , <code>randomize()</code>
Optimizer	<code>bfit()</code> , <code>opt()</code> , <code>cox.mode()</code> , <code>jm.mode()</code> <code>boost()</code> , <code>stabsel()</code> , <code>bboost()</code> , <code>lasso()</code>
Sampler	<code>GMCMC()</code> , <code>JAGS()</code> , <code>STAN()</code> , <code>BayesX()</code> , <code>cox.mcmc()</code> , <code>jm.mcmc()</code>
Results	<code>results.bamlss.default()</code>

To implement new engines, only the building block functions have to be exchanged.

# R package *bamlss*

The package makes heavy uses of **mgcv infrastructures** using `smooth.construct()`, however, optimizers and samplers may use special model terms, e.g., the LASSO constructor `la()`.

```
R> f <- list(
+   num ~ s(x1) + s(x2) + la(id),
+   sigma ~ s(x1) + s(x2) + la(id)
+ )
R> bf <- bamlss.frame(f, data = d, family = "gaussian")
R> names(bf)

[1] "call"          "model.frame"  "y"            "formula"
[5] "terms"        "family"       "x"

R> names(bf$x$mu)

[1] "formula"          "fake.formula"  "terms"
[4] "model.matrix"    "smooth.construct"

R> names(bf$x$mu$smooth.construct)

[1] "s(x1)" "s(x2)" "la(id)"
```

# R package *bamlss*

Work in progress ...

Function	Distribution
<code>beta_bamlss()</code>	Beta distribution
<code>binomial_bamlss()</code>	Binomial distribution
<code>cnorm_bamlss()</code>	Censored normal distribution
<code>cox_bamlss()</code>	Continuous time Cox-model
<code>gaussian_bamlss()</code>	Gaussian distribution
<code>gamma_bamlss()</code>	Gamma distribution
<code>gpareto_bamlss()</code>	Generalized Pareto distribution
<code>jm_bamlss()</code>	Continuous time joint-model
<code>multinomial_bamlss()</code>	Multinomial distribution
<code>mvn_bamlss()</code>	Multivariate normal distribution
<code>poisson_bamlss()</code>	Poisson distribution
...	

New families only require density, distribution, random number generator, quantile, score and hess functions. Wrapper for R package *gamlss* families.

# R package *bamlss*

Wrapper function:

```
R> f <- list(y ~ la(id,fuse=2), sigma ~ la(id,fuse=1))
R> b <- bamlss(f, family = "gaussian", sampler = FALSE,
+   optimizer = lasso, criterion = "BIC", multiple = TRUE)
```

Standard extractor and plotting functions:

```
summary(), plot(), fitted(), residuals(), predict(),
coef(), logLik(), DIC(), samples(), ...
```

# R package *bamlss*

**Example:** model fitting functions.

```
bfit(x, y, family, start = NULL, weights = NULL, offset = NULL,  
     update = "iwls", criterion = c("AICc", "BIC", "AIC"), ...)
```

```
boost(x, y, family, weights = NULL, offset = NULL,  
      nu = 0.1, df = 4, maxit = 400, ...)
```

```
GMCMC(x, y, family, start = NULL, weights = NULL, offset = NULL,  
      n.iter = 1200, burnin = 200, thin = 1, ...)
```

# R package *bamlss*

**Example:** updating functions.

```
bfit_iwls(x, family, y, eta, id, weights, criterion, ...)
```

```
boost_fit(x, y, nu, hatmatrix = TRUE, weights = NULL, ...)
```

```
GMCMC_iwls(family, theta, id, eta, y, data,  
  weights = NULL, offset = NULL, ...)
```

```
GMCMC_slice(family, theta, id, eta, y, data, ...)
```

# Forecasting Lightning

**Objective:** Develop a NWP postprocessing for probabilistic forecasting of lightning counts in the Eastern Alps.


## Requirements:

- Handle **nonlinear** relationships between the response and covariates.
- **Select** objectively important exploratory variables.
- Provide **inference** of scores and predictions.

# Data

## ALDIS lightning counts:

- Summer: May–August.
- Afternoons: 12–18 UTC.
- Gridded on  $18 \times 18 \text{ km}^2$ .
- 2010–2017.






# Data

## ALDIS lightning counts:


- Summer: May–August.
- Afternoons: 12–18 UTC.
- Gridded on  $18 \times 18 \text{ km}^2$ .
- 2010–2017.



# Data

## ALDIS lightning counts:


- Summer: May–August.
- Afternoons: 12–18 UTC.
- Gridded on  $18 \times 18 \text{ km}^2$ .
- 2010–2017.



## ECMWF ENS:

- Forecast horizons: 1–5 days.
- 2016–2017.
- Convective precipitation, CAPE, temperature, relative humidity, vertical velocity, radiation, heat fluxes, . . .
- Median and interquartile range.

# Lightning Counts



# Methods

- Excess zeros and overdispersion in count data:  
**Hurdle approach.**
- Nonlinear relationships:  
**Additive predictors.**
- Selection of terms:  
**Gradient boosting with stability selection.**
- Inference:  
**Markov chain Monte Carlo simulation.**

# Hurdle approach

Form an **intensity model** of lightning events by combining

- **Binary hurdle part:**

Logit binomial with parameter  $\pi$ .

Serves as isolated model for the **occurrence** of lightning events.

- **Truncated count part:**

Zero-truncated negative binomial with location parameter  $\mu > 0$  and dispersion parameter  $\theta > 0$ .

# Additive predictors

$$\text{logit}(\pi) = \underbrace{\beta_0 + f_1(\text{doy}) + f_2(\text{lon}, \text{lat})}_{\text{climatology}} + f_3(\mathbf{x}_3) + \dots + f_p(\mathbf{x}_p).$$

$$\log(\mu) = \dots$$

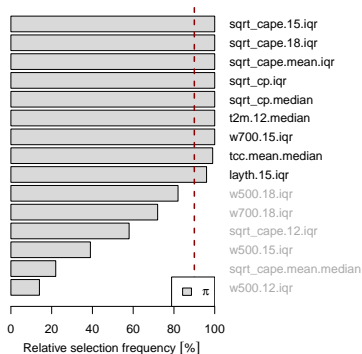
$$\log(\theta) = \dots$$

- $\pi, \mu, \theta$ : Distributional parameters.
- $f_i(\cdot)$ : Potentially non-linear functions.
- $\text{doy}$ : Day of the year.
- $\text{lon}, \text{lat}$ : Longitude, latitude.
- $\mathbf{x}_i$ : ECMWF-ENS covariate.
- $p$ : Number of potential terms (here:  $p \approx 200$ ).

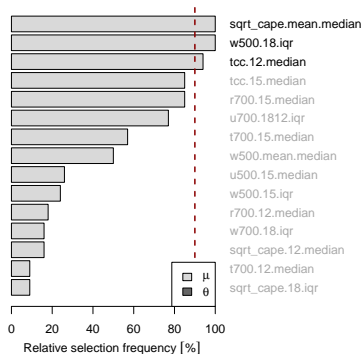
# Stability Selection

Function `stabse1()` in *bamlss*.


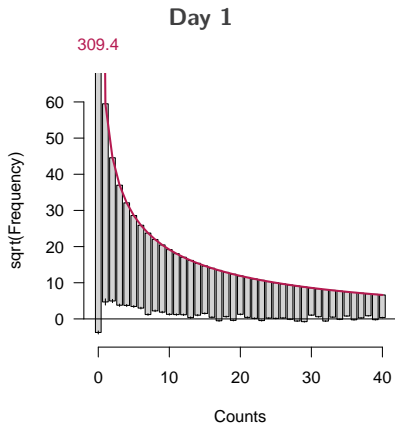
## Binary hurdle part



## Truncated count part

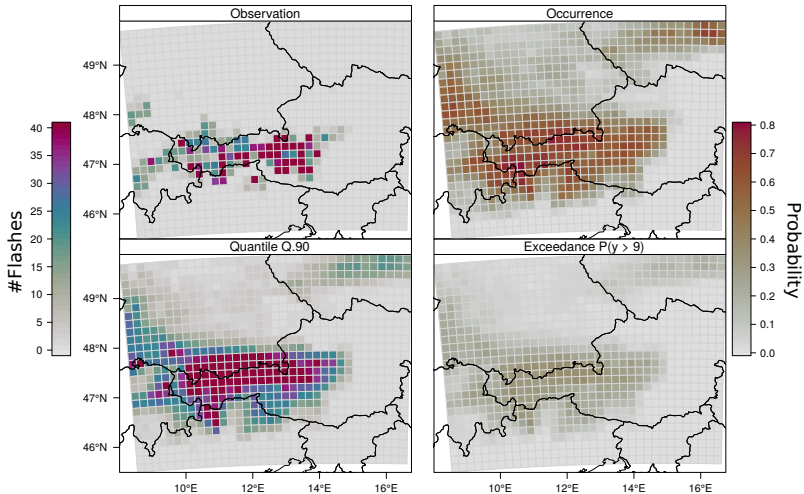


# Predictive Performance





# Predictions



# References

- ▶ A. Groll, et al. *LASSO-Type Penalization in the Framework of Generalized Additive Models for Location Scale and Shape*. Working Papers, Faculty of Economics and Statistics, University of Innsbruck, 2018.  
<https://EconPapers.repec.org/RePEc:inn:wpaper:2018-16>
- ▶ M. Köhler, et al. *Nonlinear Association Structures in Flexible Bayesian Additive Joint Models*. STAT MED, 2018. doi:10.1002/sim.7967.
- ▶ T. Simon, et al. *Lightning Prediction Using Model Output Statistics*. Working Papers, Faculty of Economics and Statistics, University of Innsbruck, 2018.  
<https://EconPapers.repec.org/RePEc:inn:wpaper:2018-14>.
- ▶ N. Umlauf, et al. *bamlss: Bayesian Additive Models for Location Scale and Shape (and Beyond)*. R package version 1.0-0, 2018.  
<http://CRAN.R-project.org/package=bamlss>.
- ▶ N. Umlauf, et al. *BAMLSS: Bayesian Additive Models for Location, Scale and Shape (and Beyond)*. J COMPUT GRAPH STAT, 2017.  
doi:10.1080/10618600.2017.1407325.
- ▶ A. Zeileis, et al. *Visualizing Count Data Regressions Using Rootograms*. AM STAT, 2016. doi:10.1080/00031305.2016.1173590.



Thank you for your attention!

Nikolaus Umlauf

<http://eeecon.uibk.ac.at/~umlauf/>