

```
logLik.bamlss <- function(object, ..., optimizer = FALSE, samples = FALSE)
{
  Call <- match.call()
  Call <- Call[!(names(Call) %in% c("optimizer", "samples"))]
  mn <- as.character(Call)[-1L]
  object <- list(object, ...)
  mstop <- object$mstop
  if(any(names(object) != "") {
    i <- names(object) == ""
    object <- object[i]
    mn <- mn[i]
  }
  object <- object[mn != "mstop"]
}
```

Distributional Modelling in R

Model Checking and Predictive Evaluation

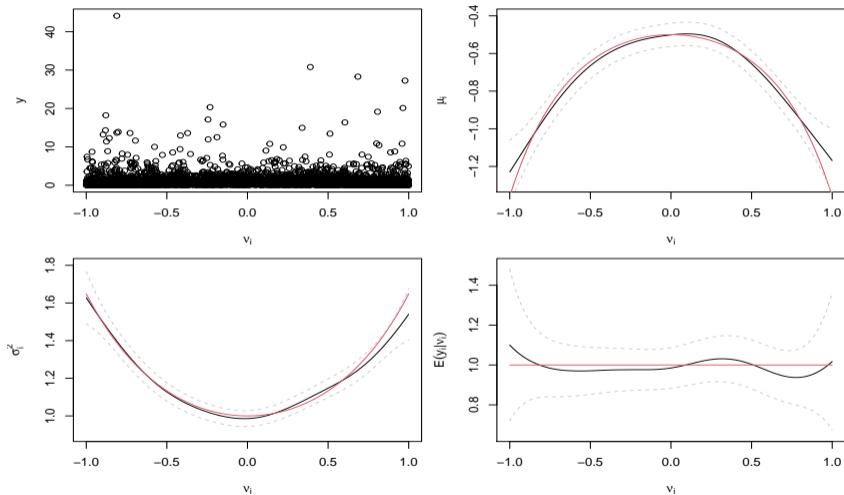
Thomas Kneib, Nikolaus Umlauf

<https://nikum.org/dmr.html>

Challenges in Distributional Regression

- Conceptually, distributional regression is very appealing and intuitive, but it comes with a number of specific challenges:
 - Interpretation of the estimated effects more difficult due to link functions and multi-parameter setup.
 - Model choice and checking to avoid model miss-specification.

Illustration for Simulated Log-Normal Data



Goals of Model Checking and Predictive Evaluation

- Finding a suitable response distribution.
- Determining the predictor (covariates and appropriate modelling variant).
- Checking model adequacy in general.

Quantile Residuals

- For a continuous random variable Y with cumulative distribution function F , the probability integral transform yields

$$F(Y) \sim U(0, 1)$$

or

$$\Phi^{-1}(F(Y)) \sim N(0, 1).$$

- For a correctly specified distributional regression model, we should therefore have

$$u_i = \Phi^{-1}(F(y_i | \hat{\theta}(\mathbf{x}_i))) \stackrel{a}{\sim} N(0, 1)$$

and the quantile residuals u_i can, e.g., be visualized in a quantile-quantile plot.

- For discrete or multivariate data, appropriate generalisations are needed.

Information Criteria

- Information criteria such as AIC or BIC can be used in the distributional regression context, e.g.

$$\text{AIC} = -2l(\hat{\gamma}) + 2 \text{df}(\hat{\gamma})$$

- The model fit is evaluated based on the (negative) log-likelihood $-2l(\hat{\gamma})$.
- The degrees of freedom $\text{df}(\hat{\gamma})$ of the model have to take the impact of the regularisation penalty into account.

Proper Scoring Rules

- In a distributional setting, typical predictive measures such as the mean squared error of prediction or the mean absolute error of prediction are usually not the most adequate choice.
- Proper scoring rules provide a framework for evaluating predictive distributions rather than point predictions.
- Underlying theory ensures that proper scores encourage the analyst to honestly report their uncertainty in terms of the predictive distribution.
- The cross-validated log-likelihood is the most commonly used proper score.

Scoring Rules for Real-Valued Outcomes

- Evaluate a predictive density $f(y)$ based on an observed outcome y_0 .
- Spherical score

$$\text{SPS}(f(y), y_0) = -\frac{f(y_0)}{(\int f^2(t)dt)^{1/2}}.$$

- Logarithmic score

$$\text{LS}(f(y), y_0) = -\log(f(y_0)).$$

- Continuously ranked probability score

$$\text{CRPS}(f(y), y_0) = \int [F(t) - \mathbf{1}_{[y_0, \infty)}(t)]^2 dt.$$

- Note: All scores are negatively oriented, i.e. smaller values indicate a better agreement between the predictive distribution and the observed values.

Model Checking in *gamlss2*

Example:

Load the Munich rent data.

```
R> library("gamlss2")
R> data("rent", package = "gamlss.data")
```

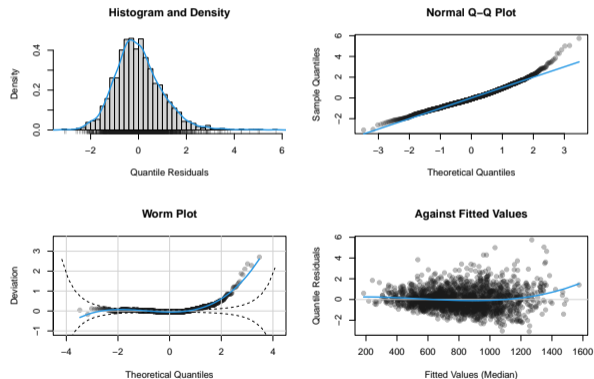
Estimate models.

```
R> f <- R ~ s(F1) + s(A) + loc + H
R> b1 <- gamlss2(f, data = rent, family = NO)
GAMLSS-RS iteration 1: Global Deviance = 28158.9399 eps = 0.027957
GAMLSS-RS iteration 2: Global Deviance = 28063.8112 eps = 0.003378
GAMLSS-RS iteration 3: Global Deviance = 28063.535 eps = 0.000009
R> b2 <- gamlss2(f, data = rent, family = GA)
GAMLSS-RS iteration 1: Global Deviance = 27694.9306 eps = 0.086618
GAMLSS-RS iteration 2: Global Deviance = 27686.4006 eps = 0.000308
GAMLSS-RS iteration 3: Global Deviance = 27686.3861 eps = 0.000000
```

Model Checking in *gamlss2*

Diagnostic plots.

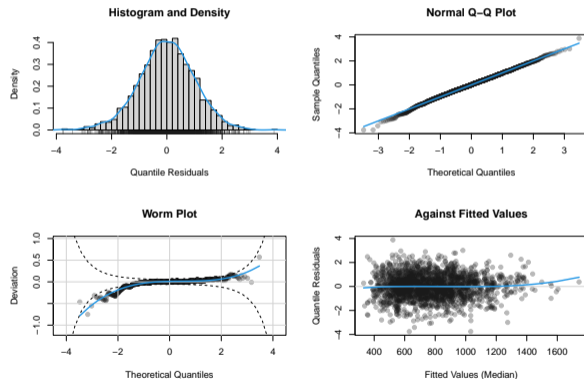
```
R> plot(b1, which = "resid")
```



Model Checking in *gamlss2*

Diagnostic plots.

```
R> plot(b2, which = "resid")
```



Model Checking in *gamlss2*

Log-likelihood and AIC.

```
R> logLik(b1)
```

```
'log Lik.' -14031.77 (df=11.5149)
```

```
R> logLik(b2)
```

```
'log Lik.' -13843.19 (df=10.08221)
```

```
R> AIC(b1, b2)
```

	AIC	df
b1	28086.56	11.51490
b2	27706.55	10.08221

```
R> GAIC(b1, b2, k = log(nrow(rent)))
```

	AIC	df
b1	28150.88	11.51490
b2	27762.86	10.08221

Model Checking in *gamlss2*

Scoring rules.

```
R> library("scoringRules")
R> p1 <- predict(b1, type = "parameter")
R> p2 <- predict(b2, type = "parameter")
R> mean(crps_norm(rent$R, mean = p1$mu, sd = p1$sigma))
[1] 165.4144
```

Using *bamlss* infrastructures.

```
R> s1 <- bamlss:::.CRPS(rent$R, p1, family(b1))
R> s2 <- bamlss:::.CRPS(rent$R, p2, family(b2))
R> mean(s1)
[1] 165.4146
R> mean(s2)
[1] 161.2407
```

Model Checking in *gamlss2*

Quantile residuals by hand. First predict parameters.

```
R> par <- predict(b1, newdata = rent, type = "parameter")
```

Calculate probabilities using the `$p()` function.

```
R> u <- family(b1)$p(rent$R, par)
```

Compute standard normal quantiles, aka quantile residuals.

```
R> e <- qnorm(u)
```

Model Checking in *gamlss2*

Q-Q plot.

```
R> qqnorm(e); qqline(e)
```

