

# 1 `logit.gam`: Generalized Additive Model for Dichotomous Dependent Variables

This function runs a nonparametric Generalized Additive Model (GAM) for dichotomous dependent variables.

## 1.1 Syntax

```
> z.out <- zelig(y ~ x1 + s(x2), model = "logit.gam", data = mydata)
> x.out <- setx(z.out)
> s.out <- sim(z.out, x = x.out)
```

Where `s()` indicates a variable to be estimated via nonparametric smooth. All variables for which `s()` is not specified, are estimated via standard parametric methods.

## 1.2 Additional Inputs

In addition to the standard inputs, `zelig()` takes the following additional options for GAM models.

- **method**: Controls the fitting method to be used. Fitting methods are selected via a list environment within `method=gam.method()`. See `gam.method()` for details.
- **scale**: Generalized Cross Validation (GCV) is used if `scale = 0` (see the “Model” section for details) except for Logit models where a Un-Biased Risk Estimator (UBRE) (also see the “Model” section for details) is used with a scale parameter assumed to be 1. If `scale` is greater than 1, it is assumed to be the scale parameter/variance and UBRE is used. If `scale` is negative GCV is used.
- **knots**: An optional list of knot values to be used for the construction of basis functions.
- **H**: A user supplied fixed quadratic penalty on the parameters of the GAM can be supplied with this as its coefficient matrix. For example, ridge penalties can be added to the parameters of the GAM to aid in identification on the scale of the linear predictor.
- **sp**: A vector of smoothing parameters for each term.
- **...**: additional options passed to the `logit.gam` model. See the `mgcv` library for details.

## 1.3 Examples

### 1. Basic Example

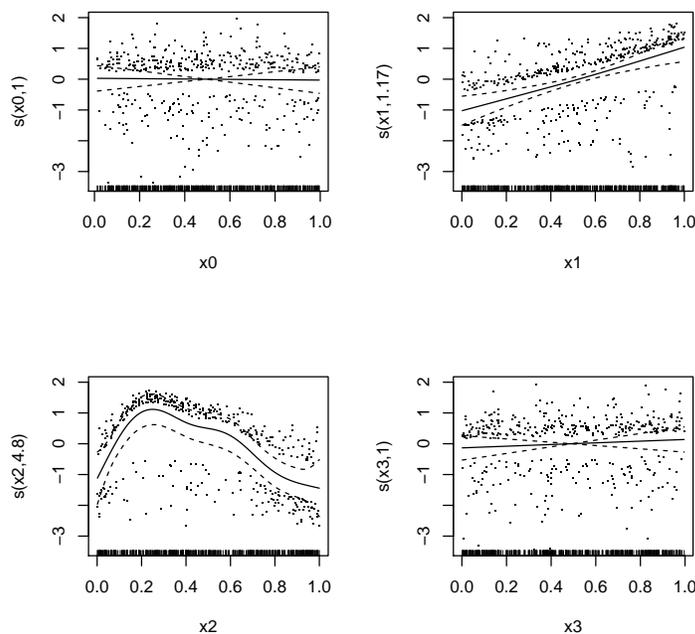
Create some count data:

```
> set.seed(0); n <- 400; sig <- 2;
> x0 <- runif(n, 0, 1); x1 <- runif(n, 0, 1)
> x2 <- runif(n, 0, 1); x3 <- runif(n, 0, 1)
> g <- (f-5)/3
> g <- binomial()$linkinv(g)
> y <- rbinom(g,1,g)
> my.data <- as.data.frame(cbind(y, x0, x1, x2, x3))
```

Estimate the model, summarize the results, and plot nonlinearities:

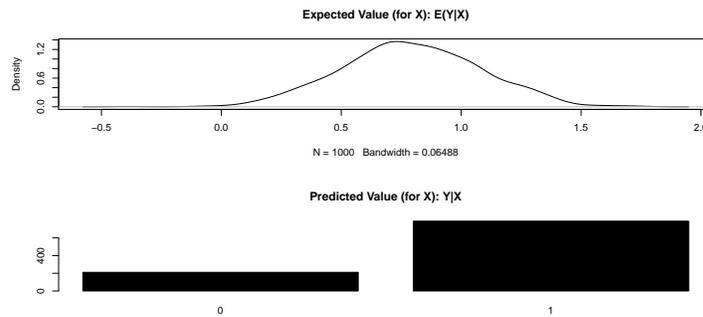
```
> z.out <- zelig(y~s(x0)+s(x1)+s(x2)+s(x3), model="logit.gam", data=my.data)
> summary(z.out)
> plot(z.out$result,pages=1,residuals=TRUE)
```

Note that the `plot()` function can be used after model estimation and before simulation to view the nonlinear relationships in the independent variables:



Set values for the explanatory variables to their default (mean/mode) values, then simulate, summarize and plot quantities of interest:

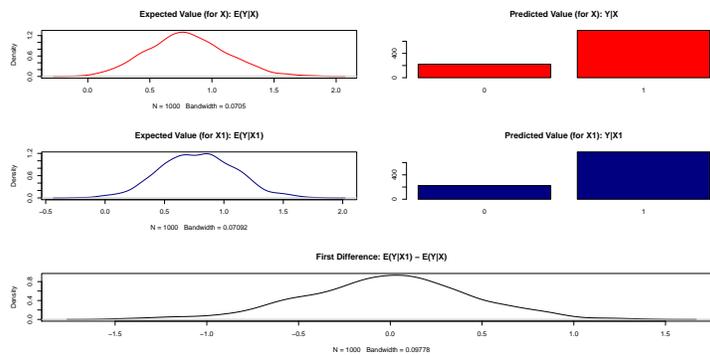
```
> x.out <- setx(z.out)
> s.out <- sim(z.out, x = x.out)
> summary(s.out)
> plot(s.out)
```



## 2. Simulating First Differences

Estimating the risk difference (and risk ratio) between low values (20th percentile) and high values (80th percentile) of the explanatory variable  $x_3$  while all the other variables are held at their default (mean/mode) values.

```
> x.high <- setx(z.out, x3= quantile(my.data$x3, 0.8))
> x.low <- setx(z.out, x3 = quantile(my.data$x3, 0.2))
> s.out <- sim(z.out, x=x.high, x1=x.low)
> summary(s.out)
> plot(s.out)
>
```



3. Variations in GAM model specification. Note that `setx` and `sim` work as shown in the above examples for any GAM model. As such, in the interest of parsimony, I will not re-specify the simulations of quantities of interest.

An extra ridge penalty (useful with convergence problems):

```
> z.out <- zelig(y~s(x0)+s(x1)+s(x2)+s(x3), H=diag(0.5,37),
+   model="logit.gam", data=my.data)
> summary(z.out)
> plot(z.out$result,pages=1,residuals=TRUE)
>
```

Set the smoothing parameter for the first term, estimate the rest:

```
> z.out <- zelig(y~s(x0)+s(x1)+s(x2)+s(x3),sp=c(0.01,-1,-1,-1),
+   model="logit.gam", data=my.data)
> summary(z.out)
> plot(z.out$result,pages=1)
>
```

Set lower bounds on smoothing parameters:

```
> z.out <- zelig(y~s(x0)+s(x1)+s(x2)+s(x3),min.sp=c(0.001,0.01,0,10),
+   model="logit.gam", data=my.data)
> summary(z.out)
> plot(z.out$result, pages=1)
>
```

A GAM with 3df regression spline term & 2 penalized terms:

```
> z.out <-zelig(y~s(x0,k=4,fx=TRUE,bs="tp")+s(x1,k=12)+s(x2,k=15),
+   model="logit.gam", data=my.data)
> summary(z.out)
> plot(z.out$result,pages=1)
>
```

## 1.4 Model

GAM models use families the same way GLM models do: they specify the distribution and link function to use in model fitting. In the case of `logit.gam` a logistic link function is used. Specifically, let  $Y_i$  be the binary dependent variable for observation  $i$  which takes the value of either 0 or 1.

- The logistic distribution has *stochastic component*

$$\begin{aligned} Y_i &\sim \text{Bernoulli}(y_i|\pi_i) \\ &= \pi_i^{y_i} (1 - \pi_i)^{1-y_i} \end{aligned}$$

where  $\pi_i = \text{Pr}(Y_i = 1)$ .

- The *systematic component* is given by:

$$\pi_i = \frac{1}{1 + \exp\left(-x_i\beta + \sum_{j=1}^J f_j(Z_j)\right)},$$

where  $x_i$  is the vector of covariates,  $\beta$  is the vector of coefficients and  $f_j(Z_j)$  for  $j = 1, \dots, J$  is the set of smooth terms..

Generalized additive models (GAMs) are similar in many respects to generalized linear models (GLMs). Specifically, GAMs are generally fit by penalized maximum likelihood estimation and GAMs have (or can have) a parametric component identical to that of a GLM. The difference is that GAMs also include in their linear predictors a specified sum of smooth functions.

In this GAM implementation, smooth functions are represented using penalized regression splines. Two techniques may be used to estimate smoothing parameters: Generalized Cross Validation (GCV),

$$n \frac{D}{(n - DF)^2}, \tag{1}$$

or an Un-Biased Risk Estimator (UBRE) (which is effectively just a rescaled AIC),

$$\frac{D}{n} + 2s \frac{DF}{n - s}, \tag{2}$$

where  $D$  is the deviance,  $n$  is the number of observations,  $s$  is the scale parameter, and  $DF$  is the effective degrees of freedom of the model. The use of GCV or UBRE can be set by the user with the `scale` command described in the “Additional Inputs” section and in either case, smoothing parameters are chosen to minimize the GCV or UBRE score for the model.

Estimation for GAM models proceeds as follows: first, basis functions and a set (one or more) of quadratic penalty coefficient matrices are constructed for each smooth term. Second, a model matrix is obtained for the parametric component of the GAM. These matrices are combined to produce a complete model matrix and a set of penalty matrices for the smooth terms. Iteratively Reweighted Least Squares (IRLS) is then used to estimate the model; at each iteration of the IRLS, a penalized weighted least squares model is run and the smoothing parameters of that model are estimated by GCV or UBRE. This process is repeated until convergence is achieved.

Further details of the GAM fitting process are given in Wood (2000, 2004, 2006).

## 1.5 Quantities of Interest

The quantities of interest for the `logit.gam` model are the same as those for the standard logistic regression.

- The expected value (`qi$ev`) for the `logit.gam` model is the mean of simulations from the stochastic component,

$$\pi_i = \frac{1}{1 + \exp\left(-x_i\beta + \sum_{j=1}^J f_j(Z_j)\right)},$$

- The predicted values (`qi$pr`) are draws from the Binomial distribution with mean equal to the simulated expected value  $\pi_i$ .
- The first difference (`qi$fd`) for the `logit.gam` model is defined as

$$FD = \Pr(Y|w_1) - \Pr(Y|w)$$

for  $w = \{X, Z\}$ .

## 1.6 Output Values

The output of each Zelig command contains useful information which you may view. For example, if you run `z.out <- zelig(y ~ x, model = "logit.gam", data)`, then you may examine the available information in `z.out` by using `names(z.out)`, see the coefficients by using `coefficients(z.out)`, and a default summary of information through `summary(z.out)`. Other elements available through the `$` operator are listed below.

- From the `zelig()` output stored in `z.out`, you may extract:
  - `coefficients`: parameter estimates for the explanatory variables.
  - `fitted.values`: the vector of fitted values for the explanatory variables.
  - `residuals`: the working residuals in the final iteration of the IRLS fit.
  - `linear.predictors`: the vector of  $x_i\beta$ .
  - `aic`: Akaike's Information Criterion (minus twice the maximized log-likelihood plus twice the number of coefficients).
  - `method`: the fitting method used.
  - `converged`: logical indicating weather the model converged or not.
  - `smooth`: information about the smoothed parameters.
  - `df.residual`: the residual degrees of freedom.
  - `df.null`: the residual degrees of freedom for the null model.
  - `data`: the input data frame.
  - `model`: the model matrix used.
- From `summary(z.out)` (as well as from `zelig()`), you may extract:

- `p.coeff`: the coefficients of the parametric components of the model.
  - `se`: the standard errors of the entire model.
  - `p.table`: the coefficients, standard errors, and associated  $t$  statistics for the parametric portion of the model.
  - `s.table`: the table of estimated degrees of freedom, estimated rank,  $F$  statistics, and  $p$ -values for the nonparametric portion of the model.
  - `cov.scaled`: a  $k \times k$  matrix of scaled covariances.
  - `cov.unscaled`: a  $k \times k$  matrix of unscaled covariances.
- From the `sim()` output stored in `s.out`, you may extract:
    - `qi$ev`: the simulated expected probabilities for the specified values of `x`.
    - `qi$pr`: the simulated predicted values for the specified values of `x`.
    - `qi$fd`: the simulated first differences in the expected probabilities simulated from `x` and `x1`.

## How to Cite the Logistic General Additive Model

### How to Cite the Zelig Software Package

To cite Zelig as a whole, please reference these two sources:

Kosuke Imai, Gary King, and Olivia Lau. 2007. “Zelig: Everyone’s Statistical Software,” <http://GKing.harvard.edu/zelig>.

Imai, Kosuke, Gary King, and Olivia Lau. (2008). “Toward A Common Framework for Statistical Analysis and Development.” *Journal of Computational and Graphical Statistics*, Vol. 17, No. 4 (December), pp. 892-913.

### See also

The `logit.gam` model is adapted from the `mgcv` package by Simon N. Wood [7]. Advanced users may wish to refer to `help(gam)`, [6], [5], and other documentation accompanying the `mgcv` package. All examples are reproduced and extended from `mgcv`’s `gam()` help pages.

## 2 `normal.gam`: Generalized Additive Model for Continuous Dependent Variables

This function runs a nonparametric Generalized Additive Model (GAM) for continuous dependent variables.

## 2.1 Syntax

```
> z.out <- zelig(y ~ x1 + s(x2), model = "normal.gam", data = mydata)
> x.out <- setx(z.out)
> s.out <- sim(z.out, x = x.out)
```

Where `s()` indicates a variable to be estimated via nonparametric smooth. All variables for which `s()` is not specified, are estimated via standard parametric methods.

## 2.2 Additional Inputs

In addition to the standard inputs, `zelig()` takes the following additional options for GAM models.

- **method:** Controls the fitting method to be used. Fitting methods are selected via a list environment within `method=gam.method()`. See `gam.method()` for details.
- **scale:** Generalized Cross Validation (GCV) is used if `scale = 0` (see the “Model” section for details) except for Normal models where a Un-Biased Risk Estimator (UBRE) (also see the “Model” section for details) is used with a scale parameter assumed to be 1. If `scale` is greater than 1, it is assumed to be the scale parameter/variance and UBRE is used. If `scale` is negative GCV is used.
- **knots:** An optional list of knot values to be used for the construction of basis functions.
- **H:** A user supplied fixed quadratic penalty on the parameters of the GAM can be supplied with this as its coefficient matrix. For example, ridge penalties can be added to the parameters of the GAM to aid in identification on the scale of the linear predictor.
- **sp:** A vector of smoothing parameters for each term.
- **...:** additional options passed to the `normal.gam` model. See the `mgcv` library for details.

## 2.3 Examples

1. Basic Example:

Create some data:

```
> set.seed(0); n <- 400; sig <- 2;
> x0 <- runif(n, 0, 1); x1 <- runif(n, 0, 1)
> x2 <- runif(n, 0, 1); x3 <- runif(n, 0, 1)
> f0 <- function(x) 2 * sin(pi * x)
> f1 <- function(x) exp(2 * x)
```

```

> f2 <- function(x) 0.2 * x^11 * (10 * (1 - x))^6 + 10 * (10 *
+   x)^3 * (1 - x)^10
> f3 <- function(x) 0 * x
> f <- f0(x0) + f1(x1) + f2(x2)
> e <- rnorm(n, 0, sig); y <- f + e
> my.data <- as.data.frame(cbind(y, x0, x1, x2, x3))

```

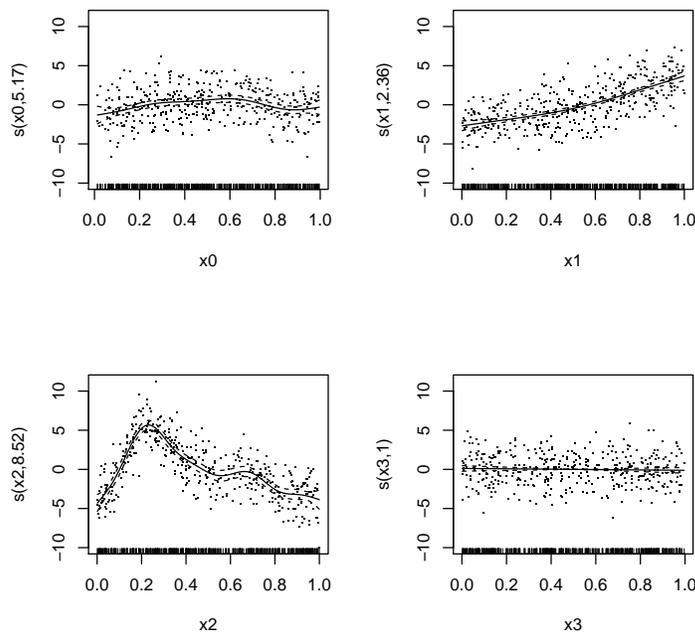
Estimate the model, summarize the results, and plot nonlinearities:

```

> z.out <- zelig(y~s(x0)+s(x1)+s(x2)+s(x3), model="normal.gam", data=my.data)
> summary(z.out)
> plot(z.out$result,pages=1,residuals=TRUE)

```

Note that the `plot()` function can be used after model estimation and before simulation to view the nonlinear relationships in the independent variables:



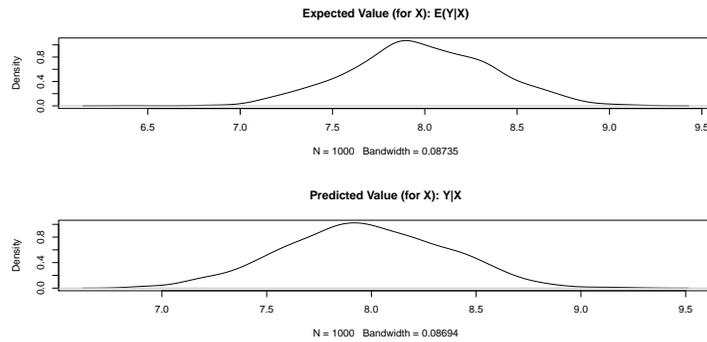
Set values for the explanatory variables to their default (mean/mode) values, then simulate, summarize and plot quantities of interest:

```

> x.out <- setx(z.out)
> s.out <- sim(z.out, x = x.out)

```

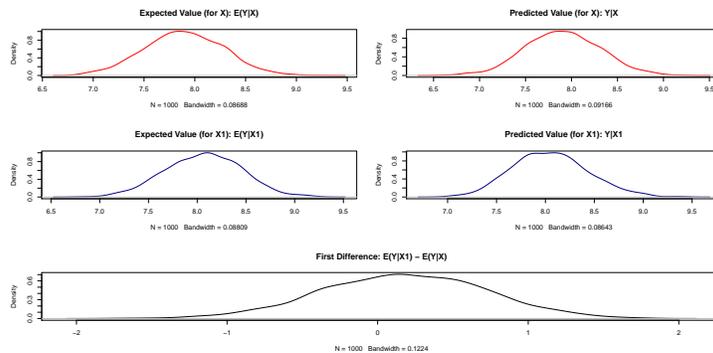
```
> summary(s.out)
> plot(s.out)
```



## 2. Simulating First Differences

Estimating the risk difference (and risk ratio) between low values (20th percentile) and high values (80th percentile) of the explanatory variable  $x_3$  while all the other variables are held at their default (mean/mode) values.

```
> x.high <- setx(z.out, x3= quantile(my.data$x3, 0.8))
> x.low <- setx(z.out, x3 = quantile(my.data$x3, 0.2))
> s.out <- sim(z.out, x=x.high, x1=x.low)
> summary(s.out)
> plot(s.out)
>
```



- ## 3. Variations in GAM model specification. Note that `setx` and `sim` work as shown in the above examples for any GAM model. As such, in the interest of parsimony, I will not re-specify the simulations of quantities of interest. An extra ridge penalty (useful with convergence problems):

```

> z.out <- zelig(y~s(x0)+s(x1)+s(x2)+s(x3), H=diag(0.5,37),
+   model="normal.gam", data=my.data)
> summary(z.out)
> plot(z.out$result,pages=1,residuals=TRUE)
>

```

Set the smoothing parameter for the first term, estimate the rest:

```

> z.out <- zelig(y~s(x0)+s(x1)+s(x2)+s(x3),sp=c(0.01,-1,-1,-1),
+   model="normal.gam", data=my.data)
> summary(z.out)
> plot(z.out$result,pages=1)
>

```

Set lower bounds on smoothing parameters:

```

> z.out <- zelig(y~s(x0)+s(x1)+s(x2)+s(x3),min.sp=c(0.001,0.01,0,10),
+   model="normal.gam", data=my.data)
> summary(z.out)
> plot(z.out$result, pages=1)
>

```

A GAM with 3df regression spline term & 2 penalized terms:

```

> z.out <-zelig(y~s(x0,k=4,fx=TRUE,bs="tp")+s(x1,k=12)+s(x2,k=15),
+   model="normal.gam", data=my.data)
> summary(z.out)
> plot(z.out$result,pages=1)
>

```

## 2.4 Model

GAM models use families the same way GLM models do: they specify the distribution and link function to use in model fitting. In the case of `normal.gam` a normal link function is used. Specifically, let  $Y_i$  be the continuous dependent variable for observation  $i$ .

- The *stochastic component* is described by a univariate normal model with a vector of means  $\mu_i$  and scalar variance  $\sigma^2$ :

$$Y_i \sim \text{Normal}(\mu_i, \sigma^2).$$

- The *systematic component* is given by:

$$\mu_i = x_i\beta + \sum_{j=1}^J f_j(Z_j).$$

where  $x_i$  is the vector of  $k$  explanatory variables,  $\beta$  is the vector of coefficients and  $f_j(Z_j)$  for  $j = 1, \dots, J$  is the set of smooth terms.

Generalized additive models (GAMs) are similar in many respects to generalized linear models (GLMs). Specifically, GAMs are generally fit by penalized maximum likelihood estimation and GAMs have (or can have) a parametric component identical to that of a GLM. The difference is that GAMs also include in their linear predictors a specified sum of smooth functions.

In this GAM implementation, smooth functions are represented using penalized regression splines. Two techniques may be used to estimate smoothing parameters: Generalized Cross Validation (GCV),

$$n \frac{D}{(n - DF)^2}, \quad (3)$$

or an Un-Biased Risk Estimator (UBRE) (which is effectively just a rescaled AIC),

$$\frac{D}{n} + 2s \frac{DF}{n - s}, \quad (4)$$

where  $D$  is the deviance,  $n$  is the number of observations,  $s$  is the scale parameter, and  $DF$  is the effective degrees of freedom of the model. The use of GCV or UBRE can be set by the user with the `scale` command described in the “Additional Inputs” section and in either case, smoothing parameters are chosen to minimize the GCV or UBRE score for the model.

Estimation for GAM models proceeds as follows: first, basis functions and a set (one or more) of quadratic penalty coefficient matrices are constructed for each smooth term. Second, a model matrix is obtained for the parametric component of the GAM. These matrices are combined to produce a complete model matrix and a set of penalty matrices for the smooth terms. Iteratively Reweighted Least Squares (IRLS) is then used to estimate the model; at each iteration of the IRLS, a penalized weighted least squares model is run and the smoothing parameters of that model are estimated by GCV or UBRE. This process is repeated until convergence is achieved.

Further details of the GAM fitting process are given in Wood (2000, 2004, 2006).

## 2.5 Quantities of Interest

The quantities of interest for the `normal.gam` model are the same as those for the standard Normal regression.

- The expected value (`qi$ev`) for the `normal.gam` model is the mean of simulations from the stochastic component,

$$E(Y) = \mu_i = x_i \beta + \sum_{j=1}^J f_j(Z_j).$$

- The predicted value (`qi$pr`) is a draw from the Normal distribution defined by the set of parameters  $(\mu_i, \sigma^2)$ .

- The first difference (`qi$fd`) for the `normal.gam` model is defined as

$$FD = \Pr(Y|w_1) - \Pr(Y|w)$$

for  $w = \{X, Z\}$ .

## 2.6 Output Values

The output of each Zelig command contains useful information which you may view. For example, if you run `z.out <- zelig(y ~ x, model = "normal.gam", data)`, then you may examine the available information in `z.out` by using `names(z.out)`, see the coefficients by using `coefficients(z.out)`, and a default summary of information through `summary(z.out)`. Other elements available through the `$` operator are listed below.

- From the `zelig()` output stored in `z.out`, you may extract:
  - `coefficients`: parameter estimates for the explanatory variables.
  - `fitted.values`: the vector of fitted values for the explanatory variables.
  - `residuals`: the working residuals in the final iteration of the IRLS fit.
  - `linear.predictors`: the vector of  $x_i\beta$ .
  - `aic`: Akaike's Information Criterion (minus twice the maximized log-likelihood plus twice the number of coefficients).
  - `method`: the fitting method used.
  - `converged`: logical indicating whether the model converged or not.
  - `smooth`: information about the smoothed parameters.
  - `df.residual`: the residual degrees of freedom.
  - `df.null`: the residual degrees of freedom for the null model.
  - `data`: the input data frame.
  - `model`: the model matrix used.
- From `summary(z.out)` (as well as from `zelig()`), you may extract:
  - `p.coeff`: the coefficients of the parametric components of the model.
  - `se`: the standard errors of the entire model.
  - `p.table`: the coefficients, standard errors, and associated  $t$  statistics for the parametric portion of the model.
  - `s.table`: the table of estimated degrees of freedom, estimated rank,  $F$  statistics, and  $p$ -values for the nonparametric portion of the model.
  - `cov.scaled`: a  $k \times k$  matrix of scaled covariances.
  - `cov.unscaled`: a  $k \times k$  matrix of unscaled covariances.

- From the `sim()` output stored in `s.out`, you may extract:
  - `qi$ev`: the simulated expected probabilities for the specified values of `x`.
  - `qi$pr`: the simulated predicted values for the specified values of `x`.
  - `qi$fd`: the simulated first differences in the expected probabilities simulated from `x` and `x1`.

## How to Cite the Normal General Additive Model

## How to Cite the Zelig Software Package

To cite Zelig as a whole, please reference these two sources:

Kosuke Imai, Gary King, and Olivia Lau. 2007. “Zelig: Everyone’s Statistical Software,” <http://GKing.harvard.edu/zelig>.

Imai, Kosuke, Gary King, and Olivia Lau. (2008). “Toward A Common Framework for Statistical Analysis and Development.” *Journal of Computational and Graphical Statistics*, Vol. 17, No. 4 (December), pp. 892-913.

## See also

The `gam.logit` model is adapted from the `mgcv` package by Simon N. Wood [7]. Advanced users may wish to refer to `help(gam)`, [6], [5], and other documentation accompanying the `mgcv` package. All examples are reproduced and extended from `mgcv`’s `gam()` help pages.

## 3 poisson.gam: Generalized Additive Model for Count Dependent Variables

This function runs a nonparametric Generalized Additive Model (GAM) for count dependent variables.

### 3.1 Syntax

```
> z.out <- zelig(y ~ x1 + s(x2), model = "poisson.gam", data = mydata)
> x.out <- setx(z.out)
> s.out <- sim(z.out, x = x.out)
```

Where `s()` indicates a variable to be estimated via nonparametric smooth. All variables for which `s()` is not specified, are estimated via standard parametric methods.

## 3.2 Additional Inputs

In addition to the standard inputs, `zelig()` takes the following additional options for GAM models.

- **method**: Controls the fitting method to be used. Fitting methods are selected via a list environment within `method=gam.method()`. See `gam.method()` for details.
- **scale**: Generalized Cross Validation (GCV) is used if `scale = 0` (see the “Model” section for details) except for Poisson models where a Un-Biased Risk Estimator (UBRE) (also see the “Model” section for details) is used with a scale parameter assumed to be 1. If `scale` is greater than 1, it is assumed to be the scale parameter/variance and UBRE is used. If `scale` is negative GCV is used.
- **knots**: An optional list of knot values to be used for the construction of basis functions.
- **H**: A user supplied fixed quadratic penalty on the parameters of the GAM can be supplied with this as its coefficient matrix. For example, ridge penalties can be added to the parameters of the GAM to aid in identification on the scale of the linear predictor.
- **sp**: A vector of smoothing parameters for each term.
- **...**: additional options passed to the `poisson.gam` model. See the `mgcv` library for details.

## 3.3 Examples

### 1. Basic Example

Create some count data:

```
> set.seed(0); n <- 400; sig <- 2;
> x0 <- runif(n, 0, 1); x1 <- runif(n, 0, 1)
> x2 <- runif(n, 0, 1); x3 <- runif(n, 0, 1)
> f0 <- function(x) 2 * sin(pi * x)
> f1 <- function(x) exp(2 * x)
> f2 <- function(x) 0.2 * x^11 * (10 * (1 - x))^6 + 10 * (10 *
+   x)^3 * (1 - x)^10
> f3 <- function(x) 0 * x
> f <- f0(x0) + f1(x1) + f2(x2)
> g <- exp(f/4); y <- rpois(rep(1, n), g)
> my.data <- as.data.frame(cbind(y, x0, x1, x2, x3))
```

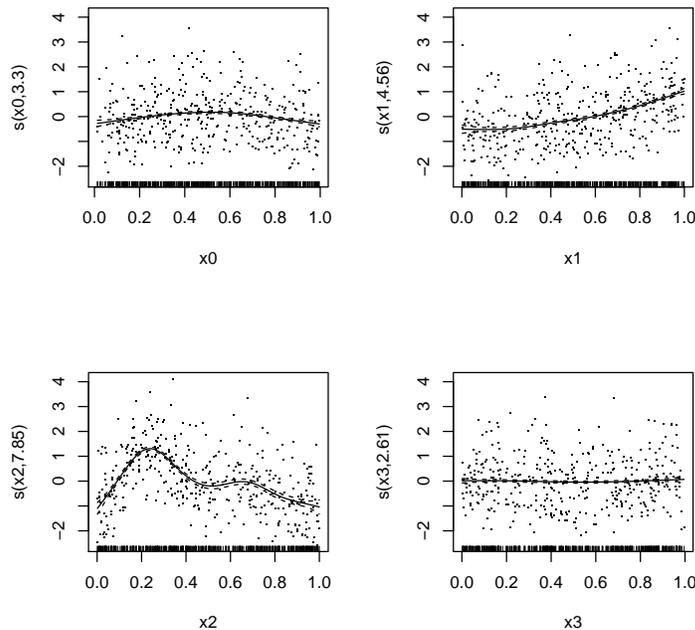
Estimate the model, summarize the results, and plot nonlinearities:

```

> z.out <- zelig(y~s(x0)+s(x1)+s(x2)+s(x3), model="poisson.gam", data=my.data)
> summary(z.out)
> plot(z.out$result,pages=1,residuals=TRUE)

```

Note that the `plot()` function can be used after model estimation and before simulation to view the nonlinear relationships in the independent variables:



Set values for the explanatory variables to their default (mean/mode) values, then simulate, summarize and plot quantities of interest:

```

> x.out <- setx(z.out)
> s.out <- sim(z.out, x = x.out)
> summary(s.out)
> plot(s.out)

```

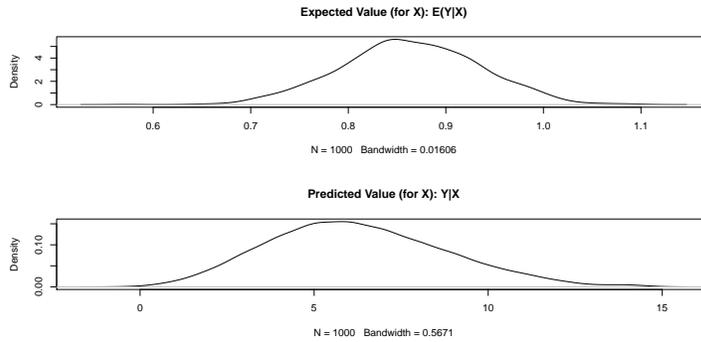
## 2. Simulating First Differences

Estimating the risk difference (and risk ratio) between low values (20th percentile) and high values (80th percentile) of the explanatory variable `x3` while all the other variables are held at their default (mean/mode) values.

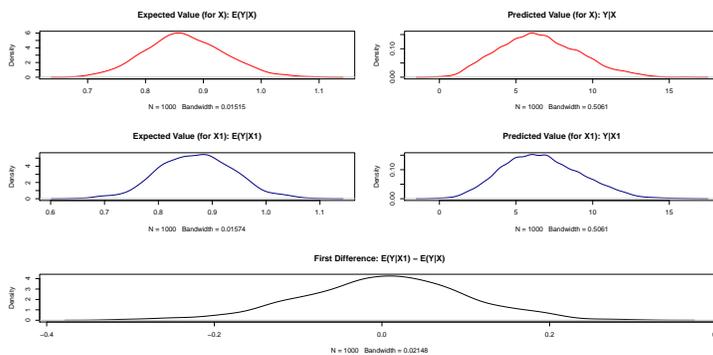
```

> x.high <- setx(z.out, x3= quantile(my.data$x3, 0.8))
> x.low <- setx(z.out, x3 = quantile(my.data$x3, 0.2))

```



```
> s.out <- sim(z.out, x=x.high, x1=x.low)
> summary(s.out)
> plot(s.out)
>
```



3. Variations in GAM model specification. Note that `setx` and `sim` work as shown in the above examples for any GAM model. As such, in the interest of parsimony, I will not re-specify the simulations of quantities of interest.

An extra ridge penalty (useful with convergence problems):

```
> z.out <- zelig(y~s(x0)+s(x1)+s(x2)+s(x3), H=diag(0.5,37),
+   model="poisson.gam", data=my.data)
> summary(z.out)
> plot(z.out$result,pages=1,residuals=TRUE)
>
```

Set the smoothing parameter for the first term, estimate the rest:

```
> z.out <- zelig(y~s(x0)+s(x1)+s(x2)+s(x3), sp=c(0.01, -1, -1, -1),
+   model="poisson.gam", data=my.data)
```

```
> summary(z.out)
> plot(z.out$result,pages=1)
>
```

Set lower bounds on smoothing parameters:

```
> z.out <- zelig(y~s(x0)+s(x1)+s(x2)+s(x3),min.sp=c(0.001,0.01,0,10),
+   model="poisson.gam", data=my.data)
> summary(z.out)
> plot(z.out$result, pages=1)
>
```

A GAM with 3df regression spline term & 2 penalized terms:

```
> z.out <-zelig(y~s(x0,k=4,fx=TRUE,bs="tp")+s(x1,k=12)+s(x2,k=15),
+   model="poisson.gam", data=my.data)
> summary(z.out)
> plot(z.out$result,pages=1)
>
```

### 3.4 Model

GAM models use families the same way GLM models do: they specify the distribution and link function to use in model fitting. In the case of `poisson.gam` a Poisson link function is used. Specifically, let  $Y_i$  be the dependent variable for observation  $i$ .  $Y_i$  is thus the number of independent events that occur during a fixed time period. This variable can take any non-negative integer.

- The Poisson distribution has *stochastic component*

$$Y_i \sim \text{Poisson}(\lambda_i),$$

where  $\lambda_i$  is the mean and variance parameter.

- The *systematic component* is given by:

$$\lambda_i = \exp \left( x_i \beta + \sum_{j=1}^J f_j(Z_j) \right).$$

where  $x_i$  is the vector of explanatory variables,  $\beta$  is the vector of coefficients and  $f_j(Z_j)$  for  $j = 1, \dots, J$  is the set of smooth terms.

Generalized additive models (GAMs) are similar in many respects to generalized linear models (GLMs). Specifically, GAMs are generally fit by penalized maximum likelihood estimation and GAMs have (or can have) a parametric component identical to that of a GLM. The difference is that GAMs also include in their linear predictors a specified sum of smooth functions.

In this GAM implementation, smooth functions are represented using penalized regression splines. Two techniques may be used to estimate smoothing parameters: Generalized Cross Validation (GCV),

$$n \frac{D}{(n - DF)^2}, \quad (5)$$

or an Un-Biased Risk Estimator (UBRE) (which is effectively just a rescaled AIC),

$$\frac{D}{n} + 2s \frac{DF}{n - s}, \quad (6)$$

where  $D$  is the deviance,  $n$  is the number of observations,  $s$  is the scale parameter, and  $DF$  is the effective degrees of freedom of the model. The use of GCV or UBRE can be set by the user with the `scale` command described in the “Additional Inputs” section and in either case, smoothing parameters are chosen to minimize the GCV or UBRE score for the model.

Estimation for GAM models proceeds as follows: first, basis functions and a set (one or more) of quadratic penalty coefficient matrices are constructed for each smooth term. Second, a model matrix is obtained for the parametric component of the GAM. These matrices are combined to produce a complete model matrix and a set of penalty matrices for the smooth terms. Iteratively Reweighted Least Squares (IRLS) is then used to estimate the model; at each iteration of the IRLS, a penalized weighted least squares model is run and the smoothing parameters of that model are estimated by GCV or UBRE. This process is repeated until convergence is achieved.

Further details of the GAM fitting process are given in Wood (2000, 2004, 2006).

### 3.5 Quantities of Interest

The quantities of interest for the `poisson.gam` model are the same as those for the standard Poisson regression.

- The expected value (`qi$ev`) for the `poisson.gam` model is the mean of simulations from the stochastic component,

$$E(Y) = \lambda_i = \exp \left( x_i \beta \sum_{j=1}^J f_j(Z_j) \right).$$

- The predicted value (`qi$pr`) is a random draw from the Poisson distribution defined by mean  $\lambda_i$ .
- The first difference (`qi$fd`) for the `poisson.gam` model is defined as

$$FD = \Pr(Y|w_1) - \Pr(Y|w)$$

for  $w = \{X, Z\}$ .

### 3.6 Output Values

The output of each Zelig command contains useful information which you may view. For example, if you run `z.out <- zelig(y ~ x, model = "poisson.gam", data)`, then you may examine the available information in `z.out` by using `names(z.out)`, see the coefficients by using `coefficients(z.out)`, and a default summary of information through `summary(z.out)`. Other elements available through the `$` operator are listed below.

- From the `zelig()` output stored in `z.out`, you may extract:
  - `coefficients`: parameter estimates for the explanatory variables.
  - `fitted.values`: the vector of fitted values for the explanatory variables.
  - `residuals`: the working residuals in the final iteration of the IRLS fit.
  - `linear.predictors`: the vector of  $x_i\beta$ .
  - `aic`: Akaike's Information Criterion (minus twice the maximized log-likelihood plus twice the number of coefficients).
  - `method`: the fitting method used.
  - `converged`: logical indicating weather the model converged or not.
  - `smooth`: information about the smoothed parameters.
  - `df.residual`: the residual degrees of freedom.
  - `df.null`: the residual degrees of freedom for the null model.
  - `data`: the input data frame.
  - `model`: the model matrix used.
- From `summary(z.out)` (as well as from `zelig()`), you may extract:
  - `p.coef`: the coefficients of the parametric components of the model.
  - `se`: the standard errors of the entire model.
  - `p.table`: the coefficients, standard errors, and associated  $t$  statistics for the parametric portion of the model.
  - `s.table`: the table of estimated degrees of freedom, estimated rank,  $F$  statistics, and  $p$ -values for the nonparametric portion of the model.
  - `cov.scaled`: a  $k \times k$  matrix of scaled covariances.
  - `cov.unscaled`: a  $k \times k$  matrix of unscaled covariances.
- From the `sim()` output stored in `s.out`, you may extract:
  - `qi$ev`: the simulated expected probabilities for the specified values of `x`.
  - `qi$pr`: the simulated predicted values for the specified values of `x`.
  - `qi$fd`: the simulated first differences in the expected probabilities simulated from `x` and `x1`.

## How to Cite the Poisson General Additive Model

## How to Cite the Zelig Software Package

To cite Zelig as a whole, please reference these two sources:

Kosuke Imai, Gary King, and Olivia Lau. 2007. “Zelig: Everyone’s Statistical Software,” <http://GKing.harvard.edu/zelig>.

Imai, Kosuke, Gary King, and Olivia Lau. (2008). “Toward A Common Framework for Statistical Analysis and Development.” *Journal of Computational and Graphical Statistics*, Vol. 17, No. 4 (December), pp. 892-913.

### See also

The `gam.logit` model is adapted from the `mgcv` package by Simon N. Wood [7]. Advanced users may wish to refer to `help(gam)`, [6], [5], and other documentation accompanying the `mgcv` package. All examples are reproduced and extended from `mgcv`’s `gam()` help pages.

## 4 `probit.gam`: Generalized Additive Model for Dichotomous Dependent Variables

This function runs a nonparametric Generalized Additive Model (GAM) for dichotomous dependent variables.

### 4.1 Syntax

```
> z.out <- zelig(y ~ x1 + s(x2), model = "probit.gam", data = mydata)
> x.out <- setx(z.out)
> s.out <- sim(z.out, x = x.out)
```

Where `s()` indicates a variable to be estimated via nonparametric smooth. All variables for which `s()` is not specified, are estimated via standard parametric methods.

### 4.2 Additional Inputs

In addition to the standard inputs, `zelig()` takes the following additional options for GAM models.

- **method**: Controls the fitting method to be used. Fitting methods are selected via a list environment within `method=gam.method()`. See `gam.method()` for details.

- **scale**: Generalized Cross Validation (GCV) is used if `scale = 0` (see the “Model” section for details) except for Logit models where a Un-Biased Risk Estimator (UBRE) (also see the “Model” section for details) is used with a scale parameter assumed to be 1. If `scale` is greater than 1, it is assumed to be the scale parameter/variance and UBRE is used. If `scale` is negative GCV is used.
- **knots**: An optional list of knot values to be used for the construction of basis functions.
- **H**: A user supplied fixed quadratic penalty on the parameters of the GAM can be supplied with this as its coefficient matrix. For example, ridge penalties can be added to the parameters of the GAM to aid in identification on the scale of the linear predictor.
- **sp**: A vector of smoothing parameters for each term.
- **...**: additional options passed to the `probit.gam` model. See the `mgcv` library for details.

### 4.3 Examples

#### 1. Basic Example

Create some count data:

```
> set.seed(0); n <- 400; sig <- 2;
> x0 <- runif(n, 0, 1); x1 <- runif(n, 0, 1)
> x2 <- runif(n, 0, 1); x3 <- runif(n, 0, 1)
> g <- (f-5)/3
> g <- binomial()$linkinv(g)
> y <- rbinom(g,1,g)
> my.data <- as.data.frame(cbind(y, x0, x1, x2, x3))
```

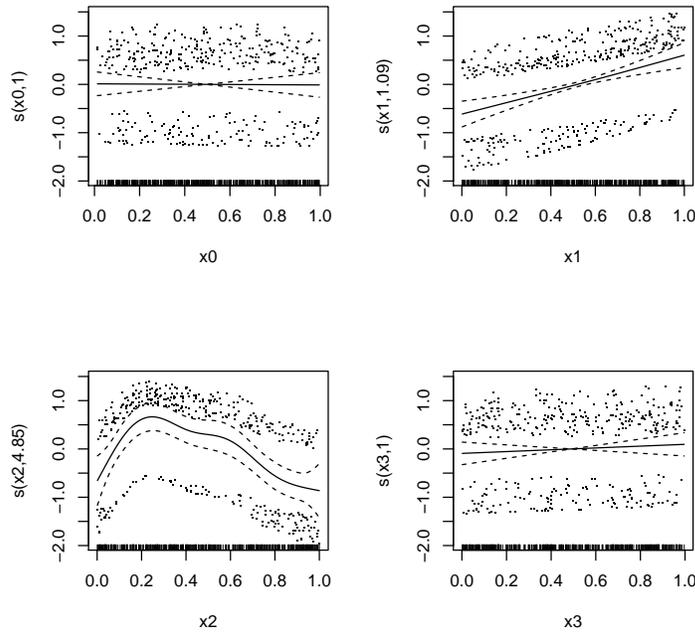
Estimate the model, summarize the results, and plot nonlinearities:

```
> z.out <- zelig(y~s(x0)+s(x1)+s(x2)+s(x3), model="probit.gam", data=my.data)
> summary(z.out)
> plot(z.out$result,pages=1,residuals=TRUE)
```

Note that the `plot()` function can be used after model estimation and before simulation to view the nonlinear relationships in the independent variables:

Set values for the explanatory variables to their default (mean/mode) values, then simulate, summarize and plot quantities of interest:

```
> x.out <- setx(z.out)
> s.out <- sim(z.out, x = x.out)
> summary(s.out)
> plot(s.out)
```



## 2. Simulating First Differences

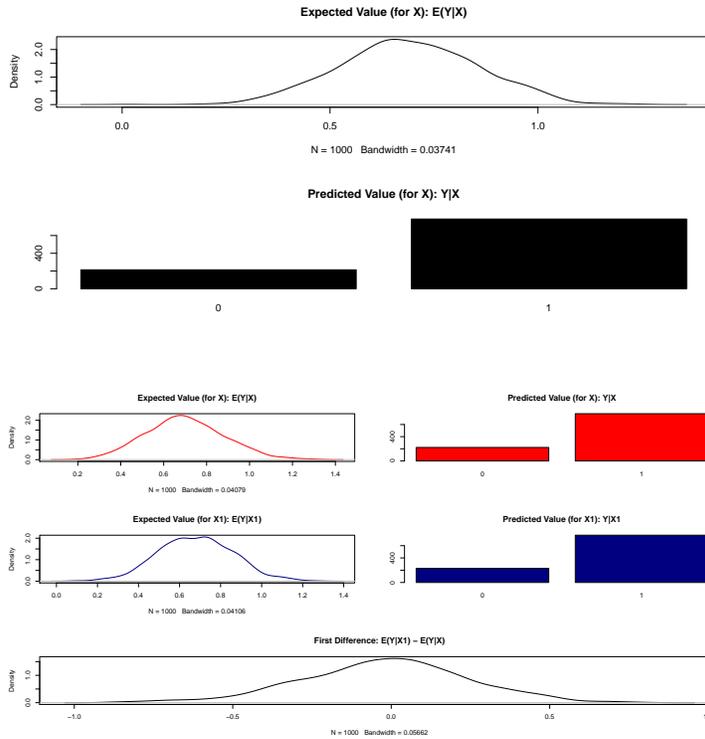
Estimating the risk difference (and risk ratio) between low values (20th percentile) and high values (80th percentile) of the explanatory variable  $x_3$  while all the other variables are held at their default (mean/mode) values.

```
> x.high <- setx(z.out, x3= quantile(my.data$x3, 0.8))
> x.low <- setx(z.out, x3 = quantile(my.data$x3, 0.2))
> s.out <- sim(z.out, x=x.high, x1=x.low)
> summary(s.out)
> plot(s.out)
>
```

## 3. Variations in GAM model specification. Note that `setx` and `sim` work as shown in the above examples for any GAM model. As such, in the interest of parsimony, I will not re-specify the simulations of quantities of interest.

An extra ridge penalty (useful with convergence problems):

```
> z.out <- zelig(y~s(x0)+s(x1)+s(x2)+s(x3), H=diag(0.5,37),
+   model="probit.gam", data=my.data)
> summary(z.out)
```



```
> plot(z.out$result, pages=1, residuals=TRUE)
>
```

Set the smoothing parameter for the first term, estimate the rest:

```
> z.out <- zelig(y~s(x0)+s(x1)+s(x2)+s(x3), sp=c(0.01, -1, -1, -1),
+   model="probit.gam", data=my.data)
> summary(z.out)
> plot(z.out$result, pages=1)
>
```

Set lower bounds on smoothing parameters:

```
> z.out <- zelig(y~s(x0)+s(x1)+s(x2)+s(x3), min.sp=c(0.001, 0.01, 0, 10),
+   model="probit.gam", data=my.data)
> summary(z.out)
> plot(z.out$result, pages=1)
>
```

A GAM with 3df regression spline term & 2 penalized terms:

```
> z.out <- zelig(y~s(x0, k=4, fx=TRUE, bs="tp")+s(x1, k=12)+s(x2, k=15),
+   model="probit.gam", data=my.data)
```

```

> summary(z.out)
> plot(z.out$result,pages=1)
>

```

## 4.4 Model

GAM models use families the same way GLM models do: they specify the distribution and link function to use in model fitting. In the case of `probit.gam` a normal link function is used. Specifically, let  $Y_i$  be the binary dependent variable for observation  $i$  which takes the value of either 0 or 1.

- The normal distribution has *stochastic component*

$$Y_i \sim \text{Bernoulli}(\pi_i)$$

where  $\pi_i = \Pr(Y_i = 1)$ .

- The *systematic component* is given by:

$$\pi_i = \Phi \left( x_i \beta + \sum_{j=1}^J f_j(Z_j) \right),$$

where  $\Phi(\mu)$  is the cumulative distribution function of the Normal distribution with mean 0 and unit variance and  $f_j(Z_j)$  for  $j = 1, \dots, J$  is the set of smooth terms.

Generalized additive models (GAMs) are similar in many respects to generalized linear models (GLMs). Specifically, GAMs are generally fit by penalized maximum likelihood estimation and GAMs have (or can have) a parametric component identical to that of a GLM. The difference is that GAMs also include in their linear predictors a specified sum of smooth functions.

In this GAM implementation, smooth functions are represented using penalized regression splines. Two techniques may be used to estimate smoothing parameters: Generalized Cross Validation (GCV),

$$n \frac{D}{(n - DF)^2}, \tag{7}$$

or an Un-Biased Risk Estimator (UBRE) (which is effectively just a rescaled AIC),

$$\frac{D}{n} + 2s \frac{DF}{n - s}, \tag{8}$$

where  $D$  is the deviance,  $n$  is the number of observations,  $s$  is the scale parameter, and  $DF$  is the effective degrees of freedom of the model. The use of GCV or UBRE can be set by the user with the `scale` command described in the “Additional Inputs” section and in either case, smoothing parameters are chosen to minimize the GCV or UBRE score for the model.

Estimation for GAM models proceeds as follows: first, basis functions and a set (one or more) of quadratic penalty coefficient matrices are constructed for each smooth term. Second, a model matrix is obtained for the parametric component of the GAM. These matrices are combined to produce a complete model matrix and a set of penalty matrices for the smooth terms. Iteratively Reweighted Least Squares (IRLS) is then used to estimate the model; at each iteration of the IRLS, a penalized weighted least squares model is run and the smoothing parameters of that model are estimated by GCV or UBRE. This process is repeated until convergence is achieved.

Further details of the GAM fitting process are given in Wood (2000, 2004, 2006).

## 4.5 Quantities of Interest

The quantities of interest for the `probit.gam` model are the same as those for the standard normal regression.

- The expected value (`qi$ev`) for the `probit.gam` model is the mean of simulations from the stochastic component,

$$\pi_i = \Phi \left( x_i \beta + \sum_{j=1}^J f_j(Z_j) \right).$$

- The predicted values (`qi$pr`) are draws from the Binomial distribution with mean equal to the simulated expected value  $\pi_i$ .
- The first difference (`qi$fd`) for the `probit.gam` model is defined as

$$FD = \Pr(Y|w_1) - \Pr(Y|w)$$

for  $w = \{X, Z\}$ .

## 4.6 Output Values

The output of each Zelig command contains useful information which you may view. For example, if you run `z.out <- zelig(y ~ x, model = "probit.gam", data)`, then you may examine the available information in `z.out` by using `names(z.out)`, see the coefficients by using `coefficients(z.out)`, and a default summary of information through `summary(z.out)`. Other elements available through the `$` operator are listed below.

- From the `zelig()` output stored in `z.out`, you may extract:
  - `coefficients`: parameter estimates for the explanatory variables.
  - `fitted.values`: the vector of fitted values for the explanatory variables.

- **residuals**: the working residuals in the final iteration of the IRLS fit.
  - **linear.predictors**: the vector of  $x_i\beta$ .
  - **aic**: Akaike's Information Criterion (minus twice the maximized log-likelihood plus twice the number of coefficients).
  - **method**: the fitting method used.
  - **converged**: logical indicating weather the model converged or not.
  - **smooth**: information about the smoothed parameters.
  - **df.residual**: the residual degrees of freedom.
  - **df.null**: the residual degrees of freedom for the null model.
  - **data**: the input data frame.
  - **model**: the model matrix used.
- From `summary(z.out)` (as well as from `zelig()`), you may extract:
    - **p.coeff**: the coefficients of the parametric components of the model.
    - **se**: the standard errors of the entire model.
    - **p.table**: the coefficients, standard errors, and associated  $t$  statistics for the parametric portion of the model.
    - **s.table**: the table of estimated degrees of freedom, estimated rank,  $F$  statistics, and  $p$ -values for the nonparametric portion of the model.
    - **cov.scaled**: a  $k \times k$  matrix of scaled covariances.
    - **cov.unscaled**: a  $k \times k$  matrix of unscaled covariances.
  - From the `sim()` output stored in `s.out`, you may extract:
    - **qi\$ev**: the simulated expected probabilities for the specified values of `x`.
    - **qi\$pr**: the simulated predicted values for the specified values of `x`.
    - **qi\$fd**: the simulated first differences in the expected probabilities simulated from `x` and `x1`.

## How to Cite the Probit General Additive Model

## How to Cite the Zelig Software Package

To cite Zelig as a whole, please reference these two sources:

Kosuke Imai, Gary King, and Olivia Lau. 2007. "Zelig: Everyone's Statistical Software," <http://GKing.harvard.edu/zelig>.

Imai, Kosuke, Gary King, and Olivia Lau. (2008). "Toward A Common Framework for Statistical Analysis and Development." *Journal of Computational and Graphical Statistics*, Vol. 17, No. 4 (December), pp. 892-913.

## See also

The `gam.logit` model is adapted from the `mgcv` package by Simon N. Wood [7]. Advanced users may wish to refer to `help(gam)`, [6], [5], and other documentation accompanying the `mgcv` package. All examples are reproduced and extended from `mgcv`'s `gam()` help pages.

## References

- [1] Matt Owen and Skyler Cranmer. *Generalized Additive Model for Logistic Regression of Dichotomous Dependent Variables*, 2011.
- [2] Matt Owen and Skyler Cranmer. *Generalized Additive Model for Normal Regression of Continuous Dependent Variables*, 2011.
- [3] Matt Owen and Skyler Cranmer. *Generalized Additive Model for Poisson Regression of Count Dependent Variables*, 2011.
- [4] Matt Owen and Skyler Cranmer. *Generalized Additive Model for Probit Regression of Dichotomous Dependent Variables*, 2011.
- [5] Simon N. Wood. Modeling and smoothing parameter estimation with multiple quadratic penalties. *Journal of the Royal Statistical Society*, 62(2):413–428, 2000.
- [6] Simon N. Wood. Stable and efficient multiple smoothing parameter estimation for generalized additive models. *Journal of the American Statistical Association*, 99:673–686, 2004.
- [7] Simon N. Wood. *Generalized Additive Models: An Introduction with R*. CRC Press, London, 2006.