

Package ‘smof’

March 4, 2024

Type Package

Title Scoring Methodology for Ordered Factors

Version 1.1.0

Date 2024-03-04

Maintainer Adelchi Azzalini <adelchi.azzalini@unipd.it>

Depends R (>= 4.0.0)

Imports stats, methods

Suggests ggplot2, survival

ByteCompile yes

NeedsCompilation no

Description Starting from a given object representing a fitted model (within a certain set of model classes) whose linear predictor includes some ordered factor(s) among the explanatory variables, a new model is constructed and fitted where each named factor is replaced by a single numeric score, suitably chosen so that the new variable produces a fit comparable with the standard methodology based on a set of polynomial contrasts. Reference: Azzalini (2023) <doi:10.1002/sta4.624>.

License GPL-2 | GPL-3

Encoding UTF-8

Author Adelchi Azzalini [aut, cre] (<<https://orcid.org/0000-0002-7583-1269>>)

R topics documented:

smof-package	2
smof	3
smof-methods	7

Index	9
--------------	----------

smof-package

*Scoring Methodology for Ordered Factors***Description**

Starting from a given object representing a fitted model (within a certain set of model classes) whose linear predictor includes some ordered factor(s) among the explanatory variables, a new model is constructed and fitted where each named factor is replaced by a single numeric score, suitably chosen so that the new variable produces a fit comparable with the standard methodology based on a set of polynomial contrasts.

Details

The DESCRIPTION file:

```
Package:          smof
Type:             Package
Title:            Scoring Methodology for Ordered Factors
Version:          1.1.0
Date:             2024-03-04
Authors@R:        person(given = "Adelchi", family = "Azzalini", email = "adelchi.azzalini@unipd.it", role = c("aut",
Maintainer:       Adelchi Azzalini <adelchi.azzalini@unipd.it>
Depends:          R (>= 4.0.0)
Imports:          stats, methods
Suggests:         ggplot2, survival
ByteCompile:     yes
NeedsCompilation: no
Description:      Starting from a given object representing a fitted model (within a certain set of model classes) whos
License:          GPL-2 | GPL-3
Encoding:         UTF-8
Author:           Adelchi Azzalini [aut, cre] (<https://orcid.org/0000-0002-7583-1269>)
```

Index of help topics:

```
print.smof      Methods for 'smof' objects
smof            Scoring Methodology for Ordered Factors
smof-package    Scoring Methodology for Ordered Factors
```

Author(s)

Author: Adelchi Azzalini [aut, cre] (<https://orcid.org/0000-0002-7583-1269>) Maintainer: Adelchi Azzalini <adelchi.azzalini@unipd.it>

References

Azzalini, A. (2023). On the use of ordered factors as explanatory variables. *Stat* **12**, e624. doi:10.1002/sta4.624

Examples

```
library(datasets)
data(esoph)
contrasts(esoph$agegp, 2) <- contr.poly(6) # optional
contrasts(esoph$tobgp, 1) <- contr.poly(4) # optional
obj1 <- glm(cbind(ncases, ncontrols) ~ agegp + tobgp + alcgp, family=binomial(), data=esoph)
out0 <- smof(obj1, esoph, "alcgp", distr.type="gh")
print(summary(out0$object))
```

 smof

Scoring Methodology for Ordered Factors

Description

Starting from an object representing a fitted model whose linear predictor includes some ordered factor(s) among the explanatory variables, a new model is constructed where each named factor is replaced by a single numeric score, suitably chosen so that the new variable produces a data fit comparable with the standard methodology based on a set of polynomial contrasts.

Usage

```
smof(object, data, factors, distr.type = "gh", fast.fit = FALSE, trace = FALSE)
```

Arguments

object	an object produced by a fitting function; see ‘Details’ below for specification of the admissible classes of objects.
data	the data frame used for producing object.
factors	a character vector with the names of the ordered factors of data which must be converted to numeric scores.
distr.type	a character string with the name of the parametric family of distributions used to construct the numeric scores. See ‘Details’ for the set of admissible choices; default value: "gh".
fast.fit	a logical value (default value: FALSE) indicating whether a fast-fitting procedure must be used. This option is available only under certain circumstances specified in the ‘Details’ below.
trace	a logical value (default value: FALSE) indicating whether details of the iterative procedure for parameter estimation must be printed at each iteration.

Details

Function `smof` implements the methodology proposed by Azzalini (2023), briefly summarized in the ‘Background’ section. It is recommended to read at least that section in case the referenced paper is not examined. The published paper has open access.

Start from an object obtained as the outcome from some fitting procedure, whose linear predictor includes one or more ordered factor(s) among the explanatory variables. For each ordered factor whose name is included in the vector `factors`, a suitable vector of numeric scores is constructed. The selection process examines the quantiles of the members of a specified parametric class of distributions and selects the member with optimizes (i.e. minimizes) a pertaining target criterion.

The admissible parametric families whose quantiles are used to construct the scores of the factors are all obtained by monotonic transformations of a standard normal variate. Specifically, the admissible families and corresponding strings to be specified in `distr.type` are as follows:

Johnson's S_U	"SU"
Tukey's g -and- h	"gh", "g-and-h"
Jones and Pewsey's \sinh - \arcsinh	"sinh-arcsinh", "SAS"

where either string name can be used when two of them are indicated. All these families involve two parameters for shape regulation; location and scale parameters are not considered, because irrelevant for our purposes. In each case, the adopted parameterization is the 'standard' one, but explicit specifications are provided in the reference below. The same `distr.type` is employed for all the components of factors.

The admissible classes for `object` are currently as follows, listed along the corresponding target criteria:

<code>class</code>	fitting function (package)	target criterion
<code>lm</code>	<code>lm (stats)</code>	sum of squared residuals
<code>mlm</code>	<code>lm (stats)</code>	[see below]
<code>glm</code>	<code>glm (stats)</code>	deviance
<code>survreg</code>	<code>survreg (survival)</code>	-loglikelihood
<code>coxph</code>	<code>coxph (survival)</code>	-loglikelihood
<code>coxph.penal</code>	<code>coxph (survival)</code>	-loglikelihood

For an object of class `mlm`, the target function is formed by summing terms where the contribution from the j -th response variable is $(1 - R_j^2)$, where R_j^2 is the r -squared statistic for that component of the fitted model. Note that, in the case of a single response variable, its $(1 - R^2)$ value is equivalent, up to an algebraic transformation, to the sum of squared residuals used for `lm` objects; hence the chosen target criterion for `mlm` models is a direct extension of the one for `lm`'s. The above list of classes may be expanded in the future, depending on feedback.

The rest of this section is slightly of more technical nature, and it may be not of interest to the casual user, especially if the option `fast.fit=TRUE` is not selected. Operationally, estimation of the `distr.type` parameters is performed via optimization of the pertaining target criterion, as indicated by the table above. For each candidate set of parameters, each factor included in `factors` is replaced by values determined by the quantiles of `distr.type` and the current parameters. The name of the new constructed variable is formed by adding `.score` to the original name. For instance, an ordered factor called `ordfac` is replaced by the numeric variable `ordfac.score` both in the linear predictor of `object` and in the data frame. A call to `update` using the modified linear predictor and data delivers a new fitting, with attached a value of the target criterion. An iterative optimization process the target criterion leads to the estimated parameters of `distr.type` with a corresponding fitted model.

There are in fact two variants of the procedure. What has been just described refers to the more 'general' variant form. However, in the prominent cases of an object of class `lm` or `glm`, the procedure can be speeded-up by setting `fast.fit=TRUE`, provided the fitted model is of a basic form, that is, a model specification via a formula, and a family in the `glm` case, without non-basic arguments such as `offset`, `subset` and alike. If these non-basic arguments are included in the object call, they are ignored for estimation of the `distr.type` parameter. However, they are included for producing the final object returned by the function. With this option, the sequence of calls to `lm` and `glm` involved by the iterative search procedure is replaced by faster calls to `lm.fit` and `glm.fit`. Correspondingly, the internal target function (`target.fit`) is slightly different from

the one used on the more general case (`target.gen`). Since the selection of the parameters involves an iterative process with dimensionality equal to twice the length of factors and each iteration involves a new data fitting process, the saving in execution time can be substantial in some cases.

Value

A list with the following components:

<code>call</code>	the calling statement
<code>new.object</code>	an updated version of the original object, with the components of factors in the model replaced by new variables; this object is itself a list, whose structure depends on its class.
<code>new.data</code>	a new data frame where the ordered factors are replaced by numeric variables representing scores.
<code>distr</code>	a list with two components: [1] type, the name of of the distribution type; [2] param, a matrix of fitted parameters with as many rows as <code>length(factors)</code> .
<code>factor.scores</code>	a list of numeric vectors with the scores assigned to the levels of each factor.
<code>original.factors</code>	a list with the names and the levels of the original factors.

Background

The methodology proposed in the reference below deals with the presence of ordered factors used as explanatory variables, hence included in the linear predictor of some model under consideration. For any given ordered factor with K levels, say, a set of K numeric scores is introduced, with a certain value assigned to each factor level. In the end, the original factor is effectively replaced by a numeric variable. This scheme represents a refinement of the elementary scoring system based on the basic sequence $1, \dots, K$, which constitutes a simple time-honoured option to deal with ordered factors, but it is not always appropriate.

The actual construction of numeric scores proceeds by selecting K quantiles of a distribution belonging to some parametric family. The adoption of a sufficiently flexible parametric family helps to find a scoring system best suited for the data under consideration, hence improving upon the basic sequence $1, \dots, K$. A concomitant product of this scheme is the identification of numeric values which indicate how the K levels are “really” spaced. Combining these two features, the key feature of the proposal is interpretability of the construction.

The proposed method represents an alternative to the use of polynomial contrasts, which is the default action taken by R for ordered factors; see the documentation of `contr.poly`.

In the proposed logic, the constructed scores are intended to be used, and interpreted, without further manipulation. Hence, for instance, building a polynomial form using one such variable would diverge somewhat from the proposed logic, although still conceivable. With a single numeric variable to represent a given factor, one cannot expect to achieve the same numerical fit to the data as obtained the polynomial contrasts built for the original factor, when these constrasts involve high degrees polynomials, and correspondingly several parameters. However, a range of numerical explorations has indicated that in many cases the resulting fit is equal or similar to the one achieved via polynomial constrasts, with non-negligible simplification in the model specification, and easier interpretation,

In a nutshell, the aim of the approach is to achieve a satisfactory data fit while improving an model parsimony, with simple interpretability of the score system.

For a more comprehensive exposition and discussion, see the reference below.

Note

For subsequent computations on the object returned by `smof`, difficulties may arise if the call to the fitting function does not set `model=TRUE`. This is not a problem with `lm` and `glm`, if their default setting `model=TRUE` has not been modified. The default setting of `coxph` is instead `model=FALSE`. This implies, for instance, that issuing the survival command `survfit(smof4$new.object)`, right after running the code of Example 4 below, would cause an error. There exist various ways to overcome this snag; the simplest one is to write

```
new.data <- smof4$new.data
s <- survfit(smof4$new.object)
```

This indication is temporary and it may be superseded by a different design in future versions of the package.

Author(s)

Adelchi Azzalini

References

Azzalini, A. (2023). On the use of ordered factors as explanatory variables. *Stat* **12**, e624. doi:10.1002/sta4.624

See Also

[contr.poly](#), [update](#), [lm](#), [lm.fit](#), [glm](#), [glm.fit](#)

Examples

```
# Example 1, reconstructs Table 2 (first part) of the reference
message("--- Example 1: esoph data ---")
library(datasets)
data(esoph)
contrasts(esoph$agegp, 2) <- contr.poly(6) # optional
contrasts(esoph$tobgp, 1) <- contr.poly(4) # optional
fit1 <- glm(cbind(ncases, ncontrols) ~ agegp + tobgp + alcgp, family=binomial(), data=esoph)
message("original fit:")
print(summary(fit1))
smof1 <- smof(fit1, esoph, "alcgp", distr.type="SU")
print(smof1, type="b", pch=20, col=4)
print(summary(smof1))
plot(smof1)
#
# Example 2 , reconstructs Table 4 (first part) of the reference
if(require(ggplot2, quietly=TRUE)) {
message("--- Example 2: diamonds data ---")
data(diamonds, package="ggplot2")
dmd <- data.frame(diamonds[seq(1, 53940, by=100),]) # use a subset of the data
dmd <- dmd[-c(518, 519, 523),] # remove three outliers
contrasts(dmd$cut, 1) <- contr.poly(5)
fit2 <- lm(sqrt(price) ~ carat + clarity + color + cut, data=dmd)
smof2 <- smof(fit2, dmd, c("color", "clarity"), distr.type="gh")
message("smof fit:")
print(smof2)
print(summary(smof2))
```

```

plot(smof2, which="clarity")
} # end diamonds example
#
# Example 3
if(require(survival, quietly=TRUE)) {
message("--- Example 3: lung data ---")
lung0 <- lung
lung0$ph.karno <- ordered(lung0$ph.karno)
contrasts(lung0$ph.karno, 3) <- contr.poly(6)
fit3 <- survreg(Surv(time, status) ~ ph.karno, data=lung0)
smof3 <- smof(fit3, lung0, "ph.karno")
print(summary(smof3))
plot(smof3) # Karnofsky scores do not seem to be linearly spaced
#
message("--- Example 4: PBC data ---")
data(pbc, package="survival")
pbc$stage <- ordered(pbc$stage)
fit4 <- coxph(Surv(time) ~ strata(status) + stage, data=pbc)
smof4 <- smof(fit4, data=pbc, factors="stage")
print(summary(smof4))
plot(smof4)
} # end of survival examples

```

smof-methods

*Methods for smof objects***Description**

The list of methods that apply to smof objects

Usage

```

## S3 method for class 'smof'
print(x, ...)
## S3 method for class 'smof'
plot(x, which, ...)
## S3 method for class 'smof'
summary(object, ...)
## S3 method for class 'summary.smof'
print(x, ...)
## S3 method for class 'smof'
predict(object, newdata, ...)

```

Arguments

object	an object returned by smof.
x	an object returned by smof, except for print.summary.smof where x is the outcome of summary.smof.
which	either a vector or a two-elements list; see ‘Details’ for full specification.
newdata	a data frame which includes ordered factors with the same names and levels as those in the data frame used to produce object; see ‘Details’ for additional information.
...	arguments passed through to other methods.

Details

There are two main parts in the outcome of `summary.smof`. One is the outcome of the selection of the factor(s) transformation(s); the other part, denoted `Final fitting call`, represents the newly fitted model with the original ordered factors replaced by numeric scores. The corresponding `data.frame`, denoted `new.data`, is one of the components of the object returned by `smof`.

If `which` is a character vector, its components are interpreted as names of the factors in the calling statement of the object to be plotted, producing a set of graphs where the numeric scores of each named factor are plotted versus the equally spaced ticks associated to the original level names. The same effect is obtained when `which` is a numeric vector of integers, which then select the corresponding components of the factors sequence.

If `which` is a list, its first element is assumed to be a vector having the meaning just described. After the pertaining set of graphs has been completed, the second element of the list is passed to the plotting method for the object produced by the model fitting procedure. Currently this option operates only for objects which inherits from class `lm`; specifically, it works for objects originated by a call to `lm` or to `glm`.

With `predict.smof`, the outcome of a `smof` fit is applied to new data frame which includes ordered factors analogous to those used to compute the `smof` transformation. Only the factors of `newdata` with the same name as those processed by `smof` to produce `object` are examined. The levels of these factors must coincide with or be a subset of those of the original data frame.

There is a difference between the role played by `predict.smof` and the one of `predict` methods for other classes. Usually the returned values pertain to the response variable, or to some related entity, while here the outcome refers to explanatory variables.

Value

For `summary.smof`, a list. For `predict.smof`, a data frame.

Author(s)

Adelchi Azzalini

See Also

[smof](#), [lm](#), [glm](#)

Examples

```
library(datasets)
data(esoph)
contrasts(esoph$agegp, 2) <- contr.poly(6)
contrasts(esoph$tobgp, 1) <- contr.poly(4)
fit1 <- glm(cbind(ncases, ncontrols) ~ agegp + tobgp + alcgp, family=binomial(), data=esoph)
smof1 <- smof(fit1, esoph, "alcgp")
print(smof1)
print(summary(smof1))
plot(smof1, type="b", pch=19, col="blue")
plot(smof1, which=list(1, 1:4))
predict(smof1, newdata=esoph[seq(1, 88, by=8), ])
```


Index

*Topic **categorical data**

smof, [3](#)

smof-methods, [7](#)

*Topic **category**

smof, [3](#)

smof-methods, [7](#)

*Topic **classes**

smof, [3](#)

smof-methods, [7](#)

*Topic **models**

smof, [3](#)

smof-methods, [7](#)

*Topic **ordered factor**

smof, [3](#)

smof-methods, [7](#)

*Topic **package**

smof-package, [2](#)

contr.poly, [6](#)

glm, [6](#), [8](#)

glm.fit, [6](#)

lm, [6](#), [8](#)

lm.fit, [6](#)

plot.smof (smof-methods), [7](#)

predict.smof (smof-methods), [7](#)

print.smof (smof-methods), [7](#)

print.summary.smof (smof-methods), [7](#)

smof, [3](#), [8](#)

smof-methods, [7](#)

smof-package, [2](#)

summary.smof (smof-methods), [7](#)

update, [6](#)