



Norwegian University of
Science and Technology

Department of Mathematical Sciences

Examination paper for
**ST2304 Statistical modelling
for biologists and biotechnologists**

Academic contact during examination: Bob O'Hara

Phone: +47 915 54 416

Examination date: 4th June 2019

Examination time (from–to): 09:00–13:00

Permitted examination support material: C: One yellow A4 sheet with your own handwritten notes (stamped by the Department of Mathematical Sciences), Tabeller og formler i statistikk (Tapir forlag, Fagbokforlaget), Matematiske formelsamling (K. Rottmann), specified calculator.

Other information:

All answers must be justified, and relevant calculations provided. Help pages for some R functions you may need are enclosed.

Language: English

Number of pages: 9

Number of pages enclosed: 3

Checked by:

Informasjon om trykking av eksamensoppgave

Originalen er:

1-sidig 2-sidig

sort/hvit farger

skal ha flervalgskjema

Date

Signature

Problem 1 Beetles

One problem in toxicology is to work out how toxic a compound is to an organism. This is estimated by calculating the LD_{50} ("lethal dose"). Animals are treated with different concentrations of a toxin, and from this the concentration that would kill half of the animals is estimated: this is the LD_{50} . Here we want to estimate the LD_{50} for beetles exposed to carbon disulphide at 8 different concentrations. The data are plotted in Figure 1.

All log values are the natural log.




Figure 1: Proportion of beetles killed by exposure to different concentrations of carbon disulphide.

The LD_{50} can be estimated by fitting a GLM, assuming a binomial distribution for each concentration, and modelling how the proportion of beetles that are killed (which we call p) varies with the $\log(\text{concentration})$. For the analysis here we assume a logit link function, and a linear effect of $\log(\text{concentration})$.

- a) What are the parameters of the binomial distribution?
- b) What assumptions do we make when using a GLM with a binomial distribution?
- c) What other link functions could be used in this analysis?

The model was fitted, with the summary shown in Figure 2.

- d) Write down the model for the effect of $\log(\text{concentration})$ on $\text{logit}(p)$, in mathematical notation.

```
beetle <- glm(cbind(NKilled, NSurv) ~ log(concentration),
+           family=binomial("logit"))
> summary(beetle)

Call:
glm(formula = cbind(NKilled, NSurv) ~ log(concentration),
    family = binomial("logit"))

Deviance Residuals:
    Min       1Q   Median       3Q      Max
-1.5213  -0.6270   0.8705   1.2575   1.6486

Coefficients:
                Estimate Std. Error z value Pr(>|z|)
(Intercept)      -59.869     5.108  -11.72  <2e-16 ***
log(concentration)  33.784     2.870   11.77  <2e-16 ***
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

(Dispersion parameter for binomial family taken to be 1)

    Null deviance: 280.866  on 7  degrees of freedom
Residual deviance:  11.474  on 6  degrees of freedom
AIC: 41.803

Number of Fisher Scoring iterations: 4
```

Figure 2: R code and results from fitting a GLM for the proportion of beetles killed as a response to log(concentration)

- e) Does the analysis support the assumption that there is an effect of $\log(\text{concentration})$ on the proportion of beetles killed? Explain your answer.

The LD_{50} can be calculated from the model as $-\frac{\text{intercept}}{\text{slope}}$.

- f) What is the estimate of LD_{50} in this experiment?

Because the estimate is a ratio, the confidence interval for it is not trivial to work out. Therefore replicate data was simulated from the model, and the LD_{50} was estimated from the simulations. This gave the density in Figure 3.

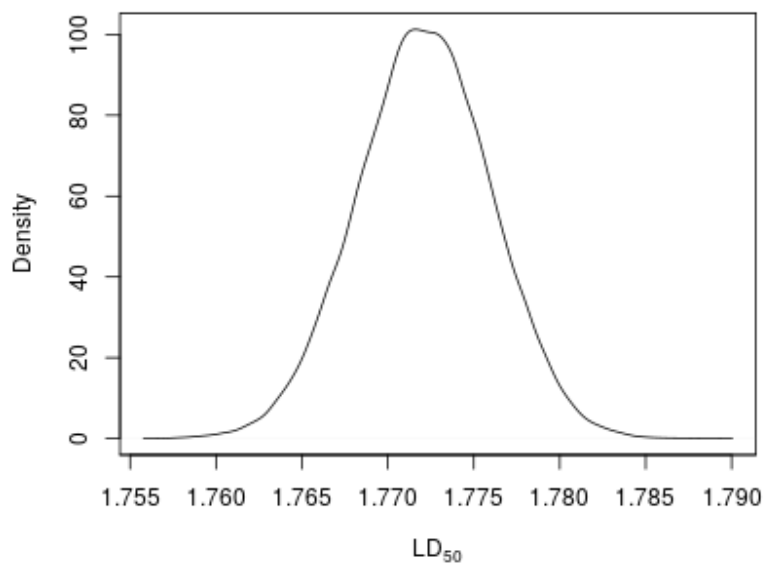


Figure 3: Density curve for LD_{50} for carbon disulphide in beetles.

- g) Using just the plot, suggest an approximate 95% confidence interval. (Your estimate will be rough, because from the plot you can only give a guess)

The residuals from the model in Figure 2 are plotted in Figure 4.

- h) Does the residual plot (Figure 4) suggest that the model assumptions seem reasonable? If not, what could you do to improve it?




Figure 4: Residuals of the model in Figure 2

Problem 2 Seed Retention Time (SRT)

A new piece of research recently looked at seed retention time (SRT) in different animal species, i.e. the time from when an animal eats a seed to when it excretes the seed in its faeces. Data on this, along with the body size of the species, has been collected. The study asked how long it would take a seed to go through any animal: this can be linked to movement to predict how far an animal can disperse seeds. The study looked at two taxa, birds and mammals. As an example, we can ask how long it would take a seed to go through a *Tyrannosaurus rex* (which we can consider to be a bird). The data are plotted in Figure 5.

All log values are the natural log.




Figure 5: Seed retention time and body mass for mammals and birds, both plotted on the log scale.

We can fit a regression model and use that to make this prediction. The model and its summary are in Figure 6.

- a) Write down the model for the effect of $\log(\text{body mass})$ on $\log(\text{SRT})$ in mathematical notation.
- b) What assumptions are we making when using a linear regression model?

```

> model1 <- lm(log(SRT) ~ log(Body.mass), data=Data)
> summary(model1)

Call:
lm(formula = log(SRT) ~ log(Body.mass), data = Data)

Residuals:
    Min       1Q   Median       3Q      Max
-2.52500 -0.70189 -0.04724  0.75092  2.72079

Coefficients:
              Estimate Std. Error t value Pr(>|t|)
(Intercept)    1.69101    0.12377   13.66  <2e-16 ***
log(Body.mass)  0.52801    0.01454   36.32  <2e-16 ***
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Residual standard error: 0.9425 on 384 degrees of freedom
(4 observations deleted due to missingness)
Multiple R-squared:  0.7745,    Adjusted R-squared:  0.7739
F-statistic: 1319 on 1 and 384 DF, p-value: < 2.2e-16

```

Figure 6: R code and summary from a linear regression model to predict log(seed retention time (SRT)) by log(body mass) of a species.

- c) Which unknown parameters does the model contain?
- d) What is the effect of body mass on SRT?

The body mass of *T. rex* has been estimated as about 7700 kg, and we can assume it behaved like a bird.

- e) Predict the SRT for a *T. rex* using the output of model 1 (Figure 6).
($\log(7700) = 8.95$)

The model fit of model 1 was checked graphically: see Figure 7.

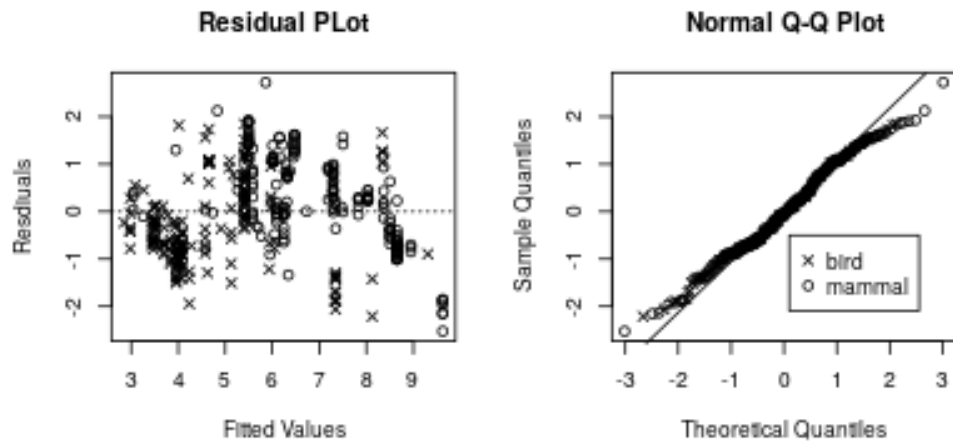


Figure 7: Residuals for a model of the effect of body mass for mammals and birds on seed retention time.

- f) How well does model 1 fit? Are there any deviations from the assumptions of the model?
- g) How could you improve the model in the light of the conclusions in qu 2f? (Other than carrying out the analyses described below)?

To follow up this initial model, two models of increasing complexity were fitted: see Figure 8 for the code.

```

> model2 <- lm(log(SRT) ~ log(Body.mass) + Taxon, data=Data)
> model3 <- lm(log(SRT) ~ log(Body.mass) * Taxon, data=Data)
    
```

Figure 8: R code to fit two multiple regression models to predict log(seed retention time (SRT)) against log(body mass) of a species, and its taxon (whether it is a bird or a mammal).

- h) Describe the differences between how model 1 and model 2 assume body mass and taxon (bird or mammal) affect SRT.
- i) Describe the differences between how model 2 and model 3 assume body mass and taxon (bird or mammal) affect SRT.

We have a summary of model 3 in Figure 9.

```

> summary(model3)

Call:
lm(formula = log(SRT) ~ log(Body.mass) * Taxon, data = Data)

Residuals:
    Min       1Q   Median       3Q      Max
-2.13565 -0.53808 -0.00752  0.49367  2.12382

Coefficients:
              Estimate Std. Error t value Pr(>|t|)
(Intercept)      1.32220    0.16185   8.169 4.59e-15 ***
log(Body.mass)    0.54183    0.02806  19.310 < 2e-16 ***
Taxonmammal      3.11416    0.29379  10.600 < 2e-16 ***
log(Body.mass):Taxonmammal -0.25509    0.03652 -6.986 1.26e-11 ***
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Residual standard error: 0.7951 on 382 degrees of freedom
(4 observations deleted due to missingness)
Multiple R-squared:  0.8404,    Adjusted R-squared:  0.8391
F-statistic: 670.3 on 3 and 382 DF,  p-value: < 2.2e-16

```

Figure 9: R code and summary from a multiple regression model to predict $\log(\text{seed retention time (SRT)})$ by $\log(\text{body mass})$ of a species, and its taxon (bird or mammal).

- j)** Draw, by hand, approximate fitted lines for the two taxa (birds and mammals), from model 3 (Figure 9). You may want to exaggerate the differences to illustrate the patterns in the fitted model.
- k)** Predict the SRT for a *T. rex* from model 3 (assuming it is a bird).

The three models were compared in Figure 10.

- l)** On the basis of the comparison in Figure 10, and also the summaries of model 1 and model 3, which model do you prefer? What statistics did you take into consideration?

```

> anova(model1, model2, model3)

Analysis of Variance Table

Model 1: log(SRT) ~ log(Body.mass)
Model 2: log(SRT) ~ log(Body.mass) + Taxon
Model 3: log(SRT) ~ log(Body.mass) * Taxon
  Res.Df  RSS Df Sum of Sq    F        Pr(>F)
1     384 341.08
2     383 272.32  1    68.765 108.785    < 2.2e-16 ***
3     382 241.47  1    30.849  48.803 0.00000000001265 ***
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

```

Figure 10: R code and output from a comparison of models 1 to 3, to predict $\log(\text{seed retention time (SRT)})$ by $\log(\text{body mass})$ of a species, and its taxon.

- m)** Compare the predictions for SRT in *T. rex* from model 1 and model 3 (the standard error for the prediction of $\log(\text{SRT})$ is around 0.9 in both models): what are the differences between the predictions, and how important are they?

lm	<i>Fitting Linear Models</i>
----	------------------------------

Description

lm is used to fit linear models. It can be used to carry out regression, single stratum analysis of variance and analysis of covariance (although [aov](#) may provide a more convenient interface for these).

Usage

```
lm(formula, data, subset, weights, na.action,
   method = "qr", model = TRUE, x = FALSE, y = FALSE, qr = TRUE,
   singular.ok = TRUE, contrasts = NULL, offset, ...)
```

Arguments

formula	an object of class <code>"formula"</code> (or one that can be coerced to that class): a symbolic description of the model to be fitted. The details of model specification are given under 'Details'.
data	an optional data frame, list or environment (or object coercible by <code>as.data.frame</code> to a data frame) containing the variables in the model. If not found in data, the variables are taken from environment(<code>formula</code>), typically the environment from which <code>lm</code> is called.
subset	an optional vector specifying a subset of observations to be used in the fitting process.
weights	an optional vector of weights to be used in the fitting process. Should be <code>NULL</code> or a numeric vector. If non- <code>NULL</code> , weighted least squares is used with weights (that is, minimizing $\sum(w_i e_i^2)$); otherwise ordinary least squares is used. See also 'Details'.
na.action	a function which indicates what should happen when the data contain NAs. The default is set by the <code>na.action</code> setting of <code>options</code> , and is <code>na.fail</code> if that is unset. The 'factory-fresh' default is <code>na.omit</code> . Another possible value is <code>NULL</code> , no action. Value <code>na.exclude</code> can be useful.
method	the method to be used; for fitting, currently only <code>method = "qr"</code> is supported; <code>method = "model.frame"</code> returns the model frame (the same as with <code>model = TRUE</code> , see below).
model, x, y, qr	logicals. If <code>TRUE</code> the corresponding components of the fit (the model frame, the model matrix, the response, the QR decomposition) are returned.
singular.ok	logical. If <code>FALSE</code> (the default in S but not in R) a singular fit is an error.
contrasts	an optional list. See the <code>contrasts.arg</code> of <code>model.matrix.default</code> .
offset	this can be used to specify an <i>a priori</i> known component to be included in the linear predictor during fitting. This should be <code>NULL</code> or a numeric vector of length equal to the number of cases. One or more <code>offset</code> terms can be included in the formula instead or as well, and if more than one are specified their sum is used. See <code>model.offset</code> .
...	additional arguments to be passed to the low level regression fitting functions (see below).

Details

Models for `lm` are specified symbolically. A typical model has the form `response ~ terms` where `response` is the (numeric) response vector and `terms` is a series of terms which specifies a linear predictor for response. A terms specification of the form `first + second` indicates all the terms in `first` together with all the terms in `second` with duplicates removed. A specification of the form `first:second` indicates the set of terms obtained by taking the interactions of all terms in `first` with all terms in `second`. The specification `first*second` indicates the *cross* of `first` and `second`. This is the same as `first + second + first:second`.

If the formula includes an `offset`, this is evaluated and subtracted from the response.

If response is a matrix a linear model is fitted separately by least-squares to each column of the matrix.

See `model.matrix` for some further details. The terms in the formula will be re-ordered so that main effects come first, followed by the interactions, all second-order, all third-order and so on: to avoid this pass a terms object as the formula (see `aov` and `demo(glm.vr)` for an example).

A formula has an implied intercept term. To remove this use either `y ~ x - 1` or `y ~ 0 + x`. See `formula` for more details of allowed formulae.

Non-`NULL` weights can be used to indicate that different observations have different variances (with the values in `weights` being inversely proportional to the variances); or equivalently, when the elements of `weights` are positive integers w_i , that each response y_i is the mean of w_i unit-weight observations (including the case that there are w_i observations equal to y_i and the data have been summarized).

`lm` calls the lower level functions `lm.fit`, etc, see below, for the actual numerical computations. For programming only, you may consider doing likewise.

All of `weights`, `subset` and `offset` are evaluated in the same way as variables in `formula`, that is first in `data` and then in the environment of `formula`.

Value

`lm` returns an object of class `"lm"` or for multiple responses of class `c("mlm", "lm")`.

The functions `summary` and `anova` are used to obtain and print a summary and analysis of variance table of the results. The generic accessor functions `coefficients`, `effects`, `fitted.values` and `residuals` extract various useful features of the value returned by `lm`.

An object of class `"lm"` is a list containing at least the following components:

<code>coefficients</code>	a named vector of coefficients
<code>residuals</code>	the residuals, that is response minus fitted values.
<code>fitted.values</code>	the fitted mean values.
<code>rank</code>	the numeric rank of the fitted linear model.
<code>weights</code>	(only for weighted fits) the specified weights.
<code>df.residual</code>	the residual degrees of freedom.
<code>call</code>	the matched call.
<code>terms</code>	the <code>terms</code> object used.
<code>contrasts</code>	(only where relevant) the contrasts used.
<code>xlevels</code>	(only where relevant) a record of the levels of the factors used in fitting.
<code>offset</code>	the offset used (missing if none were used).
<code>y</code>	if requested, the response used.

<code>x</code>	if requested, the model matrix used.
<code>model</code>	if requested (the default), the model frame used.
<code>na.action</code>	(where relevant) information returned by <code>model.frame</code> on the special handling of NAs.

In addition, non-null fits will have components `assign`, `effects` and (unless not requested) `qr` relating to the linear fit, for use by extractor functions such as `summary` and `effects`.

Using time series

Considerable care is needed when using `lm` with time series.

Unless `na.action = NULL`, the time series attributes are stripped from the variables before the regression is done. (This is necessary as omitting NAs would invalidate the time series attributes, and if NAs are omitted in the middle of the series the result would no longer be a regular time series.)

Even if the time series attributes are retained, they are not used to line up series, so that the time shift of a lagged or differenced regressor would be ignored. It is good practice to prepare a data argument by `ts.intersect(..., dframe = TRUE)`, then apply a suitable `na.action` to that data frame and call `lm` with `na.action = NULL` so that residuals and fitted values are time series.

Note

Offsets specified by `offset` will not be included in predictions by `predict.lm`, whereas those specified by an offset term in the formula will be.

Author(s)

The design was inspired by the S function of the same name described in Chambers (1992). The implementation of model formula by Ross Ihaka was based on Wilkinson & Rogers (1973).

References

Chambers, J. M. (1992) *Linear models*. Chapter 4 of *Statistical Models in S* eds J. M. Chambers and T. J. Hastie, Wadsworth & Brooks/Cole.

Wilkinson, G. N. and Rogers, C. E. (1973) Symbolic descriptions of factorial models for analysis of variance. *Applied Statistics*, **22**, 392–9.

See Also

`summary.lm` for summaries and `anova.lm` for the ANOVA table; `aov` for a different interface.

The generic functions `coef`, `effects`, `residuals`, `fitted`, `vcov`.

`predict.lm` (via `predict`) for prediction, including confidence and prediction intervals; `confint` for confidence intervals of *parameters*.

`lm.influence` for regression diagnostics, and `glm` for **generalized** linear models.

The underlying low level functions, `lm.fit` for plain, and `lm.wfit` for weighted regression fitting. More `lm()` examples are available e.g., in `anscombe`, `attitude`, `freeny`, `LifeCycleSavings`, `longley`, `stackloss`, `swiss`.

`biglm` in package `biglm` for an alternative way to fit linear models to large datasets (especially those with many cases).

Examples

```
require(graphics)

## Annette Dobson (1990) "An Introduction to Generalized Linear Models".
## Page 9: Plant Weight Data.
ctl <- c(4.17, 5.58, 5.18, 6.11, 4.50, 4.61, 5.17, 4.53, 5.33, 5.14)
trt <- c(4.81, 4.17, 4.41, 3.59, 5.87, 3.83, 6.03, 4.89, 4.32, 4.69)
group <- gl(2, 10, 20, labels = c("ctl", "trt"))
weight <- c(ctl, trt)
lm.D9 <- lm(weight ~ group)
lm.D90 <- lm(weight ~ group - 1) # omitting intercept

anova(lm.D9)
summary(lm.D90)

opar <- par(mfrow = c(2,2), oma = c(0, 0, 1.1, 0))
plot(lm.D9, las = 1) # Residuals, Fitted, ...
par(opar)

### less simple examples in "See Also" above
```

glm	Fitting Generalized Linear Models
Description	glm is used to fit generalized linear models, specified by giving a symbolic description of the linear predictor and a description of the error distribution.
Usage	<pre>glm(formula, family = gaussian, data, weights, subset, na.action, start = NULL, etastart, mustart, offset, control = list(...), model = TRUE, method = "glm.fit", x = FALSE, y = TRUE, contrasts = NULL, ...) glm.fit(x, y, weights = rep(1, nobs), start = NULL, etastart = NULL, mustart = NULL, offset = rep(0, nobs), family = gaussian(), control = list(), intercept = TRUE) ## S3 method for class 'glm' weights(object, type = c("prior", "working"), ...)</pre>
Arguments	<p>formula an object of class "formula" (or one that can be coerced to that class): a symbolic description of the model to be fitted. The details of model specification are given under 'Details'.</p> <p>family a description of the error distribution and link function to be used in the model. For <code>glm</code> this can be a character string naming a family function, a family function or the result of a call to a family function. For <code>glm.fit</code> only the third option is supported. (See family for details of family functions).</p> <p>data an optional data frame, list or environment (or object coercible by <code>as.data.frame</code> to a data frame) containing the variables in the model. If not found in data, the variables are taken from environment (<code>formula</code>), typically the environment from which <code>glm</code> is called.</p> <p>weights an optional vector of 'prior weights' to be used in the fitting process. Should be NULL or a numeric vector.</p> <p>subset an optional vector specifying a subset of observations to be used in the fitting process.</p> <p>na.action a function which indicates what should happen when the data contain NAs. The default is set by the <code>na.action</code> setting of <code>options</code>, and is <code>na.fail</code> if that is unset. The 'factory-fresh' default is <code>na.omit</code>. Another possible value is NULL, no action. Value <code>na.exclude</code> can be useful.</p> <p>start starting values for the parameters in the linear predictor.</p> <p>etastart starting values for the linear predictor.</p> <p>mustart starting values for the vector of means.</p> <p>offset this can be used to specify an <i>a priori</i> known component to be included in the linear predictor during fitting. This should be NULL or a numeric vector of length equal to the number of cases. One or more <code>offset</code> terms can be included in the formula instead or as well, and if more than one is specified their sum is used. See <code>model.offset</code>.</p> <p>control a list of parameters for controlling the fitting process. For <code>glm.fit</code> this is passed to <code>glm.control</code>.</p> <p>model a logical value indicating whether <i>model frame</i> should be included as a component of the returned value.</p> <p>method the method to be used in fitting the model. The default method <code>"glm.fit"</code> uses iteratively reweighted least squares (IWLS); the alternative <code>"model.frame"</code> returns the model frame and does no fitting. User-supplied fitting functions can be supplied either as a function or a character string naming a function, with a function which takes the same arguments as <code>glm.fit</code>. If specified as a character string it is looked up from within the <code>stats</code> namespace.</p> <p>x, y For <code>glm</code>: logical values indicating whether the response vector and model matrix used in the fitting process should be returned as components of the returned value. For <code>glm.fit</code>: <code>x</code> is a design matrix of dimension $n \times p$, and <code>y</code> is a vector of observations of length <code>n</code>.</p> <p>contrasts an optional list. See the <code>contrasts.arg</code> of <code>model.matrix.default</code>.</p> <p>intercept logical. Should an intercept be included in the <i>null</i> model?</p> <p>object an object inheriting from class <code>"glm"</code>.</p> <p>type character, partial matching allowed. Type of weights to extract from the fitted model object. Can be abbreviated.</p> <p>... For <code>glm</code>: arguments to be used to form the default <code>control</code> argument if it is not supplied directly. For <code>weights</code>: further arguments passed to or from other methods.</p>
Details	<p>A typical predictor has the form <code>response ~ terms</code> where <code>response</code> is the (numeric) response vector and <code>terms</code> is a series of terms which specifies a linear predictor for response. For binomial and quasibinomial families the response can also be specified as a factor (when the first level denotes failure and all others success) or as a two-column matrix with the columns giving the numbers of successes and failures. A terms specification of the form <code>first + second</code> indicates all the terms in <code>first</code> together with all the terms in <code>second</code> with any duplicates removed.</p> <p>A specification of the form <code>first:second</code> indicates the set of terms obtained by taking the interactions of all terms in <code>first</code> with all terms in <code>second</code>. The specification <code>first*second</code> indicates the <i>cross</i> of <code>first</code> and <code>second</code>. This is the same as <code>first + second + first:second</code>.</p> <p>The terms in the formula will be re-ordered so that main effects come first, followed by the interactions, all second-order, all third-order and so on: to avoid this pass a <code>terms</code> object as the formula.</p> <p>Non-NULL weights can be used to indicate that different observations have different dispersions (with the values in <code>weights</code> being inversely proportional to the dispersions); or equivalently, when the elements of <code>weights</code> are positive integers w_i, that each response y_i is the mean of w_i unit-weight observations. For a binomial GLM prior weights are used to give the number of trials when the response is the proportion of successes: they would rarely be used for a Poisson GLM.</p>
Description	<p><code>glm.fit</code> is the workhorse function: it is not normally called directly but can be more efficient where the response vector, design matrix and family have already been calculated.</p> <p>If more than one of <code>etastart</code>, <code>start</code> and <code>mustart</code> is specified, the first in the list will be used. It is often advisable to supply starting values for a quasi family, and also for families with unusual links such as <code>gaussian("log")</code>.</p> <p>All of <code>weights</code>, <code>subset</code>, <code>offset</code>, <code>etastart</code> and <code>mustart</code> are evaluated in the same way as variables in <code>formula</code>, that is first in <code>data</code> and then in the environment of <code>formula</code>.</p> <p>For the background to warning messages about 'fitted probabilities numerically 0 or 1 occurred' for binomial GLMs, see Venables & Ripley (2002, pp. 197–8).</p>
Value	<p><code>glm</code> returns an object of class inheriting from <code>"glm"</code> which inherits from the class <code>"lm"</code>. See later in this section. If a non-standard method is used, the object will also inherit from the class (if any) returned by that function.</p> <p>The function <code>summary</code> (i.e., <code>summary.glm</code>) can be used to obtain or print a summary of the results and the function <code>anova</code> (i.e., <code>anova.glm</code>) to produce an analysis of variance table.</p> <p>The generic accessor functions <code>coefficients</code>, <code>effects</code>, <code>fitted.values</code> and <code>residuals</code> can be used to extract various useful features of the value returned by <code>glm</code>.</p> <p><code>weights</code> extracts a vector of weights, one for each case in the fit (after subsetting and <code>na.action</code>).</p> <p>An object of class <code>"glm"</code> is a list containing at least the following components:</p> <ul style="list-style-type: none"> <code>coefficients</code> a named vector of coefficients <code>residuals</code> the <i>working</i> residuals, that is the residuals in the final iteration of the IWLS fit. Since cases with zero weights are omitted, their working residuals are NA. <code>fitted.values</code> the fitted mean values, obtained by transforming the linear predictors by the inverse of the link function. <code>rank</code> the numeric rank of the fitted linear model. <code>family</code> the family object used. <code>linear.predictors</code> the linear fit on link scale. <code>deviance</code> up to a constant, minus twice the maximized log-likelihood. Where sensible, the constant is chosen so that a saturated model has deviance zero. <code>aic</code> A version of Akaike's <i>An Information Criterion</i>, minus twice the maximized log-likelihood plus twice the number of parameters, computed by the <code>aic</code> component of the family. For binomial and Poisson families the dispersion is fixed at one and the number of parameters is the number of coefficients. For gaussian, Gamma and inverse gaussian families the dispersion is estimated from the residual deviance, and the number of parameters is the number of coefficients plus one. For a gaussian family the MLE of the dispersion is used so this is a valid value of AIC, but for Gamma and inverse gaussian families it is not. For families fitted by quasi-likelihood the value is NA. <code>null.deviance</code> The deviance for the null model, comparable with deviance. The null model will include the offset, and an intercept if there is one in the model. Note that this will be incorrect if the link function depends on the data other than through the fitted mean: specify a zero offset to force a correct calculation. <code>iter</code> the number of iterations of IWLS used. <code>weights</code> the <i>working</i> weights, that is the weights in the final iteration of the IWLS fit. <code>prior.weights</code> the weights initially supplied, a vector of 1s if none were. <code>df.residual</code> the residual degrees of freedom. <code>df.null</code> the residual degrees of freedom for the null model. <code>y</code> if requested (the default) the <code>y</code> vector used. (It is a vector even for a binomial model.) <code>x</code> if requested, the model matrix. <code>model</code> if requested (the default), the model frame. <code>converged</code> logical. Was the IWLS algorithm judged to have converged? <code>boundary</code> logical. Is the fitted value on the boundary of the attainable values? <code>call</code> the matched call. <code>formula</code> the formula supplied. <code>terms</code> the terms object used. <code>data</code> the <code>data</code> argument. <code>offset</code> the offset vector used. <code>control</code> the value of the <code>control</code> argument used. <code>method</code> the name of the fitter function used, currently always <code>"glm.fit"</code>. <code>contrasts</code> (where relevant) the contrasts used. <code>xlevels</code> (where relevant) a record of the levels of the factors used in fitting. <code>na.action</code> (where relevant) information returned by <code>model.frame</code> on the special handling of NAs. <p>In addition, non-empty fits will have components <code>qr</code>, <code>R</code> and <code>effects</code> relating to the final weighted linear fit.</p> <p>Objects of class <code>"glm"</code> are normally of class <code>c("glm", "lm")</code>, that is inherit from class <code>"lm"</code>, and well-designed methods for class <code>"lm"</code> will be applied to the weighted linear model at the final iteration of IWLS. However, care is needed, as extractor functions for class <code>"glm"</code> such as <code>residuals</code> and <code>weights</code> do not just pick out the component of the fit with the same name.</p> <p>If a binomial <code>glm</code> model was specified by giving a two-column response, the <code>weights</code> returned by <code>prior.weights</code> are the total numbers of cases (factored by the supplied case weights) and the component <code>y</code> of the result is the proportion of successes.</p>
Fitting functions	<p>The argument <code>method</code> serves two purposes. One is to allow the model frame to be recreated with no fitting. The other is to allow the default fitting function <code>glm.fit</code> to be replaced by a function which takes the same arguments and uses a different fitting algorithm. If <code>glm.fit</code> is supplied as a character string it is used to search for a function of that name, starting in the <code>stats</code> namespace.</p> <p>The class of the object return by the fitter (if any) will be prepended to the class returned by <code>glm</code>.</p>
Author(s)	<p>The original R implementation of <code>glm</code> was written by Simon Davies working for Ross Ihaka at the University of Auckland, but has since been extensively re-written by members of the R Core team. The design was inspired by the S function of the same name described in Hastie & Pregibon (1992).</p>

References

- Dobson, A. J. (1990) *An Introduction to Generalized Linear Models*. London: Chapman and Hall.
- Hastie, T. J. and Pregibon, D. (1992) *Generalized linear models*. Chapter 6 of *Statistical Models in S* eds J. M. Chambers and T. J. Hastie, Wadsworth & Brooks/Cole.
- McCullagh P. and Nelder, J. A. (1989) *Generalized Linear Models*. London: Chapman and Hall.
- Venables, W. N. and Ripley, B. D. (2002) *Modern Applied Statistics with S*. New York: Springer.

See Also

[anova.glm](#), [summary.glm](#), etc. for glm methods, and the generic functions [anova](#), [summary](#), [effects.fitted.values](#), and [residuals](#).

[lm](#) for non-generalized *linear* models (which SAS calls GLMs, for 'general' linear models).

[loglin](#) and [loglm](#) (package **MASS**) for fitting log-linear models (which binomial and Poisson GLMs are) to contingency tables.

[bigglm](#) in package **biglm** for an alternative way to fit GLMs to large datasets (especially those with many cases).

[esoph](#), [infert](#) and [predict.glm](#) have examples of fitting binomial glms.

Examples

```
## Dobson (1990) Page 93: Randomized Controlled Trial :
counts <- c(18,17,15,20,10,20,25,13,12)
outcome <- gl(3,1,9)
treatment <- gl(3,3)
print(d.AD <- data.frame(treatment, outcome, counts))
glm.D93 <- glm(counts ~ outcome + treatment, family = poisson())
anova(glm.D93)
summary(glm.D93)
```

```
## an example with offsets from Venables & Ripley (2002, p.189)
utils::data(anorexia, package = "MASS")
```

```
anorex.1 <- glm(Postwt ~ Prewt + Treat + offset(Prewt),
               family = gaussian, data = anorexia)
summary(anorex.1)
```

```
# A Gamma example, from McCullagh & Nelder (1989, pp. 300-2)
clotting <- data.frame(
  u = c(5,10,15,20,30,40,60,80,100),
  lot1 = c(118,58,42,35,27,25,21,19,18),
  lot2 = c(69,35,26,21,18,16,13,12,12))
summary(glm(lot1 ~ log(u), data = clotting, family = Gamma))
summary(glm(lot2 ~ log(u), data = clotting, family = Gamma))
```

```
## Not run:
## for an example of the use of a terms object as a formula
demo(glm.vr)
```

```
## End(Not run)
```