

# Tutorial for the `dcglm` package

Péter Sólymos  
University of Alberta

---

## Abstract

This tutorial package demonstrates the capabilities of data cloning algorithm via the infrastructure provided by the `dclone` package. The functions developed here reproduce main features of the `glm` base function in R by using data cloning.

*Keywords:* Bayesian statistics, data cloning, maximum likelihood inference, generalized linear models, R.

---

## 1. Introduction

Data cloning is a statistical computing method introduced by [Lele \*et al.\* \(2007\)](#). It exploits the computational simplicity of the Markov chain Monte Carlo (MCMC) algorithms used in the Bayesian statistical framework, but it provides valid frequentist inferences such as the maximum likelihood estimates and their standard errors for complex hierarchical models. The use of the data cloning algorithm is especially straightforward for complex models, where the number of unknowns increases with sample size (i.e. mixed models), because inference and prediction procedures are often hard to implement in such situations.

The `dclone` R package ([Sólymos 2009](#)) aims to provide low level functionality to easily implement more specific higher level procedures based on data cloning for users familiar with the Bayesian methodology. This tutorial, we develop high level functions to reproduce the `glm` base function of R by using the data cloning algorithm building on the infrastructure of the `dclone` package.

## 2. Data generation

We generate random data for Poisson and Binomial GLMs. First we define the number of locations ( $n$ ) and the independent covariate ( $x$ ).  $X$  represents the design matrix:

```
R> library(dclone)
R> set.seed(1234)
R> n <- 20
R> x <- runif(n, -1, 1)
R> X <- model.matrix(~x)
```

Parameters ( $\beta_1$ ), linear predictor ( $\mu_1$ ) and random response ( $Y_1$ ) for the Poisson case (log link function):

```
R> beta1 <- c(2, -1)
```

```
R> mu1 <- X %*% beta1
R> Y1 <- rpois(n, exp(mu1))
```

Parameters (`beta2`), linear predictor (`mu2`) and random response (`Y2`) for the Binomial (Bernoulli) case (logistic link function):

```
R> beta2 <- c(0, -1)
R> mu2 <- X %*% beta2
R> Y2 <- rbinom(n, 1, exp(mu2) / (1 + exp(mu2)))
```

### 3. GLM based on the glm function

Now we fit the Poisson and Binomial GLM by using the `glm` base function and inspect their summaries:

```
R> m1 <- glm(Y1 ~ x, family=poisson)
R> summary(m1)
```

Call:

```
glm(formula = Y1 ~ x, family = poisson)
```

Deviance Residuals:

Min	1Q	Median	3Q	Max
-1.6336	-0.8138	-0.2134	0.5154	2.0309

Coefficients:

	Estimate	Std. Error	z value	Pr(> z )
(Intercept)	1.94162	0.09482	20.478	< 2e-16 ***
x	-1.27648	0.16017	-7.969	1.60e-15 ***

---

Signif. codes: 0 '\*\*\*' 0.001 '\*\*' 0.01 '\*' 0.05 '.' 0.1 ' ' 1

(Dispersion parameter for poisson family taken to be 1)

Null deviance: 90.897 on 19 degrees of freedom  
 Residual deviance: 19.060 on 18 degrees of freedom  
 AIC: 99.543

Number of Fisher Scoring iterations: 4

```
R> m2 <- glm(Y2 ~ x, family=binomial)
R> summary(m2)
```

Call:

```
glm(formula = Y2 ~ x, family = binomial)
```

Deviance Residuals:

Min	1Q	Median	3Q	Max
-1.7850	-0.8787	-0.5794	1.0077	1.5424

Coefficients:

	Estimate	Std. Error	z value	Pr(> z )
--	----------	------------	---------	----------

```

(Intercept) -0.3388      0.5021  -0.675   0.500
x           -1.7377      1.0186  -1.706   0.088 .
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

(Dispersion parameter for binomial family taken to be 1)

Null deviance: 27.526  on 19  degrees of freedom
Residual deviance: 24.097  on 18  degrees of freedom
AIC: 28.097

Number of Fisher Scoring iterations: 4

```

## 4. The full Bayesian model for GLM and data cloning

Here is the JAGS model for the Poisson case, with flat Normal priors for the regression coefficients:

```

R> glm.pois <- function() {
+   for (i in 1:n) {
+     Y[i] ~ dpois(lambda[i])
+     log(lambda[i]) <- inprod(X[i,], beta[1,])
+   }
+   for (j in 1:np) {
+     beta[1,j] ~ dnorm(0, 0.001)
+   }
+ }

```

The data will be represented as a list (if we want to use data cloning consistently, we have to use the list format instead of the global environment):

```

R> dat1 <- list(n = length(Y1), Y = Y1, X = X, np = ncol(X))
R> str(dat1)

```

```

List of 4
 $ n : int 20
 $ Y : num [1:20] 14 3 2 4 5 6 25 22 3 9 ...
 $ X : num [1:20, 1:2] 1 1 1 1 1 1 1 1 1 1 ...
 .. attr(*, "dimnames")=List of 2
 .. ..$ : chr [1:20] "1" "2" "3" "4" ...
 .. ..$ : chr [1:2] "(Intercept)" "x"
 ..- attr(*, "assign")= int [1:2] 0 1
 $ np: int 2

```

Now let's clone the data set (note that `n` should be multiplied, while `np` must remain unchanged):

```

R> n.clones <- 5
R> dcdat1 <- dclone(dat1, n.clones, multiply = "n", unchanged = "np")
R> str(dcdat1)

```

```

List of 4
 $ n : atomic [1:1] 100

```

```

..- attr(*, "n.clones")= atomic [1:1] 5
.. ..- attr(*, "method")= chr "multi"
$ Y : atomic [1:100] 14 3 2 4 5 6 25 22 3 9 ...
..- attr(*, "n.clones")= atomic [1:1] 5
.. ..- attr(*, "method")= chr "rep"
$ X : num [1:100, 1:2] 1 1 1 1 1 1 1 1 1 1 ...
..- attr(*, "dimnames")=List of 2
.. ..$ : chr [1:100] "1_1" "2_1" "3_1" "4_1" ...
.. ..$ : chr [1:2] "(Intercept)" "x"
..- attr(*, "n.clones")= atomic [1:1] 5
.. ..- attr(*, "method")= chr "rep"
$ np: int 2

```

To fit the model to the data with data cloning is as easy as this:

```
R> mod1 <- jags.fit(dcdat1, "beta", glm.pois)
```

```
R> summary(mod1)
```

```

Iterations = 1001:2000
Thinning interval = 1
Number of chains = 3
Sample size per chain = 1000
Number of clones = 5

```

1. Empirical mean and standard deviation for each variable, plus standard error of the mean:

	Mean	SD	DC SD	Naive SE	Time-series SE	R hat
beta[1]	1.940	0.04385	0.09804	0.0008005	0.001800	1.005
beta[2]	-1.276	0.07336	0.16403	0.0013393	0.002854	1.005

2. Quantiles for each variable:

	2.5%	25%	50%	75%	97.5%
beta[1]	1.852	1.912	1.942	1.970	2.023
beta[2]	-1.423	-1.325	-1.274	-1.228	-1.133

Let's compare this `mcmc.list` (more accurately an `mcmc.list.dc`) object with the `glm` results:

```

R> cbind(true.values=beta1,
+        glm.estimates=coef(m1), glm.se=summary(m1)$coefficients[,2],
+        dc.estimates=coef(mod1), dc.se=dcsd(mod1))

```

	true.values	glm.estimates	glm.se	dc.estimates	dc.se
(Intercept)	2	1.941624	0.09481704	1.940468	0.09804054
x	-1	-1.276479	0.16017431	-1.275982	0.16403481

Here is the JAGS model for the Binomial (Bernoulli, because only one trial) case:

```

R> glm.bin <- function() {
+   for (i in 1:n) {
+     Y[i] ~ dbin(p[i], k)
+     logit(p[i]) <- inprod(X[i,], beta[1,])
+   }
+ }

```

```
+     }
+     for (j in 1:np) {
+         beta[1,j] ~ dnorm(0, 0.001)
+     }
+ }
```

Putting together the data set is similar to the Poisson case:

```
R> dat2 <- list(n = length(Y2), Y = Y2, k = 1, X = X, np = ncol(X))
R> str(dat2)
```

```
List of 5
 $ n : int 20
 $ Y : num [1:20] 1 0 0 1 0 1 0 1 0 0 ...
 $ k : num 1
 $ X : num [1:20, 1:2] 1 1 1 1 1 1 1 1 1 1 ...
 .. attr(*, "dimnames")=List of 2
 .. ..$ : chr [1:20] "1" "2" "3" "4" ...
 .. ..$ : chr [1:2] "(Intercept)" "x"
 ..- attr(*, "assign")= int [1:2] 0 1
 $ np: int 2
```

but data cloning setup is a bit different. We don't have to repeat the data vectors or columns in the design matrix `n.clones` times, because there is an easier way. We multiply `Y` and `k` with `n.clones` and leave the other elements unchanged:

```
R> dcdat2 <- dclone(dat2, n.clones, multiply = c("Y", "k"), unchanged = c("n", "np", "X"))
R> str(dcdat2)
```

```
List of 5
 $ n : int 20
 $ Y : atomic [1:20] 5 0 0 5 0 5 0 5 0 0 ...
 ..- attr(*, "n.clones")= atomic [1:1] 5
 .. ..- attr(*, "method")= chr "multi"
 $ k : atomic [1:1] 5
 ..- attr(*, "n.clones")= atomic [1:1] 5
 .. ..- attr(*, "method")= chr "multi"
 $ X : num [1:20, 1:2] 1 1 1 1 1 1 1 1 1 1 ...
 ..- attr(*, "dimnames")=List of 2
 .. ..$ : chr [1:20] "1" "2" "3" "4" ...
 .. ..$ : chr [1:2] "(Intercept)" "x"
 ..- attr(*, "assign")= int [1:2] 0 1
 $ np: int 2
```

Now fit the model to the data with data cloning:

```
R> mod2 <- jags.fit(dcdat2, "beta", glm.bin)
R> summary(mod2)
```

```
Iterations = 1001:2000
Thinning interval = 1
Number of chains = 3
Sample size per chain = 1000
```

Number of clones = 5

1. Empirical mean and standard deviation for each variable, plus standard error of the mean:

	Mean	SD	DC SD	Naive SE	Time-series SE	R hat
beta[1]	-0.3528	0.2287	0.5114	0.004176	0.005744	1.0004
beta[2]	-1.7878	0.4631	1.0355	0.008455	0.011879	0.9998

2. Quantiles for each variable:

	2.5%	25%	50%	75%	97.5%
beta[1]	-0.8183	-0.5061	-0.3491	-0.1949	0.08165
beta[2]	-2.7756	-2.0775	-1.7726	-1.4760	-0.92991

and compare results with `glm` results:

```
R> cbind(true.values=beta2,
+        glm.estimates=coef(m2), glm.se=summary(m2)$coefficients[,2],
+        dc.estimates=coef(mod2), dc.se=dcSD(mod2))
```

	true.values	glm.estimates	glm.se	dc.estimates	dc.se
(Intercept)	0	-0.338286	0.5020926	-0.3527595	0.5114491
x	-1	-1.7377202	1.0185540	-1.7878365	1.0355356

## 5. The 'custommodel' function

The `custommodel` function enables us to reuse the same JAGS model with minor modifications. For example we combine the above Poisson and Binomial model into one.

```
R> glm.model <- function() {
+   for (i in 1:n) {
+     Y[i] ~ dpois(lambda[i])
+     Y[i] ~ dbin(p[i], k)
+     log(lambda[i]) <- inprod(X[i,], beta[1,])
+     logit(p[i]) <- inprod(X[i,], beta[1,])
+   }
+   for (j in 1:np) {
+     beta[1,j] ~ dnorm(0, 0.001)
+   }
+ }
```

If we want to use this in the `jags.fit` function, it would give the error message about the attempt to define the nodes more than once. To avoid this, we tell the function which line(s) should be excluded:

```
R> custommodel(glm.model, c(4,6))
```

```
c("model {", "   for (i in 1:n) {", "       Y[i] ~ dpois(lambda[i])",
"       log(lambda[i]) <- inprod(X[i,], beta[1,])", "   }",
"   for (j in 1:np) {", "       beta[1,j] ~ dnorm(0, 0.001)",
"   }", "}")
```

```
R> custommodel(glm.model, c(3,5))
```

```
c("model {", "    for (i in 1:n) {", "        Y[i] ~ dbin(p[i], k)",
"        logit(p[i]) <- inprod(X[i,], beta[1,])", "    }", "    for (j in 1:np) {",
"        beta[1,j] ~ dnorm(0, 0.001)", "    }", "}")
```

so eventually we get back our original models. But these are not functions, but character vectors of the class ‘`custommodel`’. `jags.fit` will recognize this.

Why do we want to complicate our lives with the `custommodel`? Because `dpois`, `dbin`, and `inprod` are not recognised as valid R objects or functions. This is, however, a requisite of building a valid R package that that passes R CMD check. A way to overcome this is to define fake objects as e.g. `inprod <- function() NULL`, but this option should be regarded as ugly and inefficient (unnecessary) as compared to a clean `custommodel` approach that will be presented in the next section.

## 6. The main function `dcglm`

Here is our main function for the data cloning based estimating procedure for the Poisson and Binomial GLMs:

```
R> dcglm <-
+ function(formula, data = parent.frame(),
+ family=c("poisson", "binomial"), n.clones=5, ...)
+ {
+     glm.model <- c("model {",
+                   "    for (i in 1:n) {",
+                   "        Y[i] ~ dpois(lambda[i])",
+                   "        Y[i] ~ dbin(p[i], k)",
+                   "        log(lambda[i]) <- inprod(X[i,], beta[1,])",
+                   "        logit(p[i]) <- inprod(X[i,], beta[1,])",
+                   "    }",
+                   "    for (j in 1:np) {",
+                   "        beta[1,j] ~ dnorm(0, 0.001)",
+                   "    }",
+                   "}")
+     family <- match.arg(family)
+     lhs <- formula[[2]]
+     formula.orig <- formula
+     Y <- eval(lhs, data)
+     formula[[2]] <- NULL
+     rhs <- model.frame(formula, data)
+     X <- model.matrix(attr(rhs, "terms"), rhs)
+     if (family == "poisson") {
+         dat <- list(n = length(Y), Y = Y, X = X, np = ncol(X))
+         dcdat <- dclone(dat, n.clones, multiply = "n", unchanged = "np")
+         model <- dclone:::custommodel(glm.model, c(4,6))
+     } else {
+         dat <- list(n = length(Y), Y = Y, X = X, np = ncol(X), k = 1)
+         dcdat <- dclone(dat, n.clones, multiply = c("Y", "k"), unchanged = c("n", "np", "X"))
+         model <- dclone:::custommodel(glm.model, c(3,5))
+     }
+ }
```

```

+   mod <- jags.fit(dcdat, "beta", model, ...)
+   COEF <- coef(mod)
+   SE <- dcsd(mod)
+   names(COEF) <- names(SE) <- colnames(X)
+   mu <- X %*% COEF
+   if (family == "poisson") {
+     fitval <- drop(exp(mu))
+     ll <- sum(dpois(Y, fitval, log=TRUE))
+   } else {
+     fitval <- drop(exp(mu) / (1 + exp(mu)))
+     ll <- sum(dbinom(Y, 1, fitval, log=TRUE))
+   }
+   rval <- list(call=match.call(),
+               mcmc = mod,
+               y = Y,
+               x = rhs,
+               model = X,
+               fitted.values = fitval,
+               linear.predictors = mu,
+               formula = formula.orig,
+               coefficients = COEF,
+               std.error = SE,
+               loglik = ll,
+               family = family,
+               df.residual = length(Y) - length(COEF),
+               df.null = length(Y) - 1)
+   class(rval) <- c("dcglm")
+   rval
+ }

```

Let's go through this function step-by-step as pseudo-code:

1. `glm.model` is the `custommodel` version of the BUGS model, unifying the Poisson and Binomial cases, as we have seen before.
2. The `family` argument is recognized, and as a result, it can be given not only in full (e.g. `family = "p"` is equivalent of `family = "poisson"`).
3. `lhs` is the left-hand-side of the formula, `Y` is the value as a result of evaluating `lhs` in `data` (that is the parent frame, which is usually the global environment if not called from inside of a function).
4. `formula[[2]] <- NULL` removes the left-hand-side from the formula (that's why we keep a copy of it named as `formula.orig`).
5. `rhs` is the right-hand-side, that is a model frame with variables defined in `data`.
6. The design matrix `X` is a result of using the `"terms"` attribute of `rhs` and evaluated in `rhs`.
7. The the Bayesian (`dat`) and data cloned (`dcdat`) data representation, and `model` depends on the `family` argument.



8. `mod` is the fitted `mcmc.list` object. Dots (...) represents all the additional arguments that can be passed, including `n.update`, `n.iter`, and `n.chains`.
9. `COEF` is the `coef` method evaluated on the `mcmc.list` object `mod`. `SD` is the data cloned standard error (scaled by  $\sqrt{k}$ ). Names of `COEF` and `SD` follow column names of `X`.
10. `mu` is the linear predictor (on log/logit scale), while `fitval` is the fitted value (response scale after using the appropriate inverse link function) and `ll` is the log-likelihood calculated from the probability mass function.
11. `rval` is the return value, that is a list with elements commonly applied in objects representing model fit (cf. for example element names with `names(m1)`):

**call** the function call,  
**mcmc** the fitted `mcmc.list` object,  
**y** the response,  
**x** the model frame (right-hand-side),  
**model** the design matrix,  
**fitted.values** fitted values,  
**linear.predictors** linear predictors,  
**formula** the formula argument of the call,  
**coefficients** means of the joint posterior distribution (maximum likelihood estimates),  
**std.error** standard errors of the MLE,  
**loglik** log-likelihood,  
**family** family argument of the call,  
**df.residual** residual degrees of freedom,  
**df.null** degrees of freedom in the null model.

12. Finally, we attach the class attribute and return `rval`.

Fun, isn't it? See if it is actually working:

```
R> dcm1 <- dcglm(Y1 ~ x)
R> dcm2 <- dcglm(Y2 ~ x, family = "binomial")
```

If we are about to inspect these objects, well, it is a mess without some additional helper functions. The most basic such functions (called methods in R jargon) are `print` and `summary`. For our convenience, we also define some other methods, too. These are based on the so called S3 method dispatch system. That is, if a generic function is defined, we can add class specific methods to it.

In our case, the most simple methods are the `coef` and `fitted`, because these only extract an element from the objects<sup>1</sup>:

```
R> coef.dcglm <- function(object, ...) object$coefficients
R> fitted.dcglm <- function(object, ...) object$fitted.values
```

---

<sup>1</sup>Note that these are identical to the `coef.default` and `fitted.default` functions, but shown here for didactic purposes.

Compare with the `glm` results:

```
R> rbind(glm=coef(m1), dcglm=coef(dcm1))
```

```
      (Intercept)      x
glm      1.941624 -1.276479
dcglm    1.939610 -1.278539
```

```
R> rbind(glm=coef(m2), dcglm=coef(dcm2))
```

```
      (Intercept)      x
glm    -0.3388286 -1.737720
dcglm  -0.3485690 -1.801005
```

```
R> rbind(glm=fitted(m1), dcglm=fitted(dcm1))
```

```
      1      2      3      4      5      6      7
glm  18.68691 5.100808 5.273269 5.086763 2.773811 4.871574 24.38239
dcglm 18.67900 5.087980 5.260289 5.073948 2.764118 4.858962 24.38253
      8      9     10     11     12     13     14
glm  13.79645 4.561353 6.721032 4.252018 6.214003 12.13746 2.364617
dcglm 13.78387 4.549061 6.707113 4.240080 6.200350 12.12388 2.355747
      15     16     17     18     19     20
glm  11.84414 2.946219 12.02980 12.64069 15.50880 13.80789
dcglm 11.83042 2.936208 12.01617 12.62738 15.49758 13.79531
```

```
R> rbind(glm=fitted(m2), dcglm=fitted(dcm2))
```

```
      1      2      3      4      5      6      7
glm  0.7317898 0.3178060 0.3276999 0.3169928 0.1689386 0.3043920 0.7967153
dcglm 0.7394043 0.3123635 0.3225281 0.3115285 0.1612997 0.2986009 0.8050611
      8      9     10     11     12     13     14
glm  0.6435203 0.2857649 0.4041127 0.2666583 0.3786844 0.6025899 0.1405842
dcglm 0.6490319 0.2795280 0.4013359 0.2600158 0.3750602 0.6068348 0.1331049
      15     16     17     18     19     20
glm  0.5945881 0.1807789 0.5996819 0.6157557 0.6791673 0.6437790
dcglm 0.5985702 0.1731434 0.6038318 0.6204232 0.6856496 0.6492982
```

For the `logLik` method, it is necessary to follow the standard rules, because AIC calculations depend on this method (this means, that we don't have to define a method for AIC if the `logLik` method exists for a class):

```
R> logLik.dcglm <- function (object, ...)
+   structure(object$loglik,
+             df = object$df.null + 1 - object$df.residual,
+             nobs = object$df.null + 1,
+             class = "logLik")
```

Compare with the `glm` results:

```
R> logLik(m1)
```

```
'log Lik.' -47.77174 (df=2)
```

```
R> logLik(dcm1)
'log Lik.' -47.77196 (df=2)
R> logLik(m2)
'log Lik.' -12.04875 (df=2)
R> logLik(dcm2)
'log Lik.' -12.05068 (df=2)
R> AIC(m1, dcm1, m2, dcm2)
```

	df	AIC
m1	2	99.54347
dcm1	2	99.54392
m2	2	28.09749
dcm2	2	28.10137

Now it is possible to write the `print` method:

```
R> print.dcgglm <- function(x, digits = max(3, getOption("digits") - 3), ...) {
+   cat("\nCall: ", deparse(x$call), "\n\n")
+   cat("Coefficients:\n")
+   print.default(format(x$coefficients, digits = digits), print.gap = 2, quote = FALSE)
+   cat("\nDegrees of Freedom:", x$df.null, "Total (i.e. Null); ", x$df.residual, "Residual\n")
+   cat("Log Likelihood:\t  ", format(signif(x$loglik, digits)), "\n")
+   invisible(x)
+ }
```

Let's have a look at the resulting objects of our `dcgglm` function:

```
R> dcm1
Call: dcgglm(formula = Y1 ~ x, n.iter = 1000)
Coefficients:
(Intercept)          x
      1.940         -1.279
Degrees of Freedom: 19 Total (i.e. Null);  18 Residual
Log Likelihood:          -47.77
R> dcm2
Call: dcgglm(formula = Y2 ~ x, family = "binomial", n.iter = 1000)
Coefficients:
(Intercept)          x
     -0.3486        -1.8010
Degrees of Freedom: 19 Total (i.e. Null);  18 Residual
Log Likelihood:          -12.05
```

Well done so far!

## 7. Methods for inference

The `summary` method returns the ML estimates, data cloning standard errors, and Wald-type  $z$  statistics and  $p$ -values:

```
R> summary.dcglm <- function(object, ...){
+   COEF <- coef(object)
+   SE <- object$std.error
+   z <- COEF / SE
+   p <- 2 * pnorm(-abs(z))
+   stab <- cbind("Estimate" = COEF, "Std. Error" = SE,
+     "z value" = z, "Pr(>|z|)" = p)
+   rval <- list(call = object$call,
+     coefficients = stab,
+     loglik = object$loglik,
+     df.residual = object$df.residual,
+     df.null = object$df.null)
+   class(rval) <- "summary.dcglm"
+   rval
+ }
```

The return value here is also a list, repeating some of the elements of the fitted `object`. To appropriately format the summary, we use the `print` method for the object class `'summary.dcglm'`:

```
R> print.summary.dcglm <-
+   function (x, digits = max(3, getOption("digits") - 3),
+     signif.stars = getOption("show.signif.stars"), ...)
+   {
+     cat("\nCall:\n")
+     cat(paste(deparse(x$call), sep = "\n", collapse = "\n"), "\n", sep = "")
+     cat("\nCoefficients:\n")
+     printCoefmat(x$coefficients, digits = digits, signif.stars = signif.stars, na.print = "NA",
+       cat("\nDegrees of Freedom:", x$df.null, "Total (i.e. Null); ", x$df.residual, "Residual\n"),
+       cat("Log Likelihood:\t ", format(signif(x$loglik, digits)), "\n")
+     invisible(x)
+   }
```

Summaries of the `glm()` results and our models:

```
R> summary(m1)
```

Call:

```
glm(formula = Y1 ~ x, family = poisson)
```

Deviance Residuals:

Min	1Q	Median	3Q	Max
-1.6336	-0.8138	-0.2134	0.5154	2.0309

Coefficients:

	Estimate	Std. Error	z value	Pr(> z )
(Intercept)	1.94162	0.09482	20.478	< 2e-16 ***

```
x          -1.27648    0.16017   -7.969 1.60e-15 ***
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

(Dispersion parameter for poisson family taken to be 1)

```
Null deviance: 90.897  on 19  degrees of freedom
Residual deviance: 19.060  on 18  degrees of freedom
AIC: 99.543
```

Number of Fisher Scoring iterations: 4

```
R> summary(dcm1)
```

```
Call:
dcglm(formula = Y1 ~ x, n.iter = 1000)
```

Coefficients:

```
          Estimate Std. Error z value Pr(>|z|)
(Intercept)  1.93961    0.09578  20.251 < 2e-16 ***
x            -1.27854    0.16201  -7.892 2.98e-15 ***
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

```
Degrees of Freedom: 19 Total (i.e. Null);  18 Residual
Log Likelihood:          -47.77
```

```
R> summary(m2)
```

```
Call:
glm(formula = Y2 ~ x, family = binomial)
```

Deviance Residuals:

```
      Min       1Q   Median       3Q      Max
-1.7850 -0.8787 -0.5794  1.0077  1.5424
```

Coefficients:

```
          Estimate Std. Error z value Pr(>|z|)
(Intercept) -0.3388    0.5021  -0.675  0.500
x            -1.7377    1.0186  -1.706  0.088 .
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

(Dispersion parameter for binomial family taken to be 1)

```
Null deviance: 27.526  on 19  degrees of freedom
Residual deviance: 24.097  on 18  degrees of freedom
AIC: 28.097
```

Number of Fisher Scoring iterations: 4

```
R> summary(dcm2)
```

Call:

```
dcglm(formula = Y2 ~ x, family = "binomial", n.iter = 1000)
```

Coefficients:

	Estimate	Std. Error	z value	Pr(> z )
(Intercept)	-0.3486	0.5160	-0.675	0.4994
x	-1.8010	1.0362	-1.738	0.0822

---

Signif. codes: 0 '\*\*\*' 0.001 '\*\*' 0.01 '\*' 0.05 '.' 0.1 ' ' 1

Degrees of Freedom: 19 Total (i.e. Null); 18 Residual

Log Likelihood: -12.05

Piece of cake!

For the `confint`, we use the asymptotic normality result of the data cloning theory ([Lele et al. 2007](#)), and the `confint` method defined for the data cloned `mcmc.list` part of the fitted model object:

```
R> confint.dcglm <- function(object, parm, level = 0.95, ...) {
+   rval <- confint(object$mcmc, parm, level, ...)
+   rownames(rval) <- names(coef(object))
+   rval
+ }
```

The 95% confidence intervals for the model estimates are:

```
R> confint(m1)
```

	2.5 %	97.5 %
(Intercept)	1.748410	2.120651
x	-1.596540	-0.968007

```
R> confint(dcm1)
```

	2.5 %	97.5 %
(Intercept)	1.751886	2.1273333
x	-1.596064	-0.9610133

```
R> confint(m2)
```

	2.5 %	97.5 %
(Intercept)	-1.404579	0.62478977
x	-4.014429	0.09589234

```
R> confint(dcm2)
```

	2.5 %	97.5 %
(Intercept)	-1.359989	0.6628509
x	-3.831866	0.2298568

Differences are due to the fact, that `confint` for `glm` uses profile likelihood, while `dcglm` confidence intervals are based on the asymptotic normality assumption. Profile likelihood can

be computed based on data cloning (Ponciano *et al.* 2009) but that procedure is not covered here.

## 8. Prediction based on the joint posterior distribution

In the prediction, we use MCMC. The likelihood part of the BUGS model for the prediction is the same as for the estimation. The only difference is in the prior specification:

```
R> glm.pred <- function() {
+   for (i in 1:n) {
+     Y[i] ~ dpois(z[i])
+     Y[i] ~ dbin(z[i], k)
+     log(z[i]) <- mu[i]
+     logit(z[i]) <- mu[i]
+     mu[i] <- inprod(X[i,], beta[1,])
+   }
+   beta[1,1:np] <- mvn[1:np]
+   mvn[1:np] ~ dnorm(coefs[], prec[,])
+ }
```

Note that we denote  $\lambda$  or  $p$  as  $z$ , this will make life easier later. We use again the `custommodel` approach to differentiate between the Poisson and Binomial cases:

```
R> custommodel(glm.pred, c(4,6))
```

```
c("model {", "   for (i in 1:n) {", "       Y[i] ~ dpois(z[i])",
"       log(z[i]) <- mu[i]", "       mu[i] <- inprod(X[i,], beta[1,])",
"   }", "   beta[1,1:np] <- mvn[1:np]", "   mvn[1:np] ~ dnorm(coefs[], prec[,])",
"}")
```

```
R> custommodel(glm.pred, c(3,5))
```

```
c("model {", "   for (i in 1:n) {", "       Y[i] ~ dbin(z[i], k)",
"       logit(z[i]) <- mu[i]", "       mu[i] <- inprod(X[i,], beta[1,])",
"   }", "   beta[1,1:np] <- mvn[1:np]", "   mvn[1:np] ~ dnorm(coefs[], prec[,])",
"}")
```

Let's consider the Poisson case only (the Binomial differs from it only by the specification of the `model` argument based on the `custommodel` approach, and the fitted model used). The prediction can be done by `jags.fit`, only the data specification is somewhat different. We will define the model parameters based on the MLE (`coefs`) and the variance-covariance matrix. We define a Multivariate Normal node for all the model parameters, by using the inverse of the variance-covariance matrix as a precision matrix (`prec`). Be careful, the check for symmetry in JAGS is stricter than the usual numerical precision in R, consequently we ensure that this condition is met by using the `make.symmetric` function. The data specification will look like (note, we are using the observed data in `X`, but algorithmically, this doesn't make any difference):

```
R> prec <- make.symmetric(solve(vcov(mod1)))
R> coefs <- coef(mod1)
R> prdat <- list(n = nrow(X), X = X,
+   np = ncol(X), k = 1, coefs = coefs, prec = prec)
```

We use the `jags.fit` function. One chain is usually enough:

```
R> prval <- jags.fit(prdat, "z", custommodel(glm.pred, c(4,6)), n.chains = 1)
```

The resulting `mcmc.list` object contains the conditional posterior distribution for our Poisson GLM based prediction with prediction intervals.

## 9. Methods for prediction

For our convenience, we can write a `vcov` method. We simply use the `vcov` method defined for the `mcmc.list` part of the fitted model object and do some cosmetics on the names:

```
R> vcov.dcglm <- function(object, ...) {
+   rval <- vcov(object$mcmc, ...)
+   rownames(rval) <- colnames(rval) <- names(coef(object))
+   rval
+ }
```

Comparison of the `glm` and `dcglm` approaches:

```
R> vcov(m1)
```

```
              (Intercept)          x
(Intercept) 0.008990272 0.009590252
x            0.009590252 0.025655809
```

```
R> vcov(dcm1)
```

```
              (Intercept)          x
(Intercept) 0.009173621 0.009808561
x            0.009808561 0.026245824
```

```
R> vcov(m2)
```

```
              (Intercept)          x
(Intercept)  0.2520969 0.1053385
x            0.1053385 1.0374523
```

```
R> vcov(dcm2)
```

```
              (Intercept)          x
(Intercept)  0.2662973 0.1083411
x            0.1083411 1.0736542
```

Quite similar as we expected.

The `predict` function will look like:

```
R> predict.dcglm <-
+   function(object, newdata = NULL,
+   type = c("link", "response"), se = FALSE, ...)
+   {
+     glm.pred <- c("model {",
+                   "   for (i in 1:n) {",
```



```

+           "      Y[i] ~ dpois(z[i])",
+           "      Y[i] ~ dbin(z[i], k)",
+           "      log(z[i]) <- mu[i]",
+           "      logit(z[i]) <- mu[i]",
+           "      mu[i] <- inprod(X[i,], beta[1,])",
+           "    }",
+           "      beta[1,1:np] <- mvn[1:np]",
+           "      mvn[1:np] ~ dmnorm(coefs[], prec[,])",
+           "    }")
+   prec <- make.symmetric(solve(vcov(object)))
+   coefs <- coef(object)
+   if (is.null(newdata)) {
+     X <- object$model
+   } else {
+     formul <- object$formula
+     formul[[2]] <- NULL
+     rhs <- model.frame(formul, newdata)
+     X <- model.matrix(attr(rhs, "terms"), rhs)
+   }
+   type <- match.arg(type)
+   params <- switch(type,
+     "link" = "mu",
+     "response" = "z")
+   model <- switch(object$family,
+     "poisson" = custommodel(glm.pred, c(4,6)),
+     "binomial" = custommodel(glm.pred, c(3,5)))
+   prdat <- list(n = nrow(X), X = X,
+     np = ncol(X), coefs = coefs, prec = prec)
+   if (object$family == "binomial")
+     prdat[["k"]] <- 1
+   prval <- jags.fit(prdat, params, model, ...)
+   if (!se) {
+     rval <- coef(prval)
+   } else {
+     rval <- list(fit = coef(prval),
+       se.fit = mcmcapply(prval, sd))
+   }
+   rval
+ }

```

The pseudo-code for `predict` is:

1. `glm.predict` is the familiar `custommodel` specification.
2. `prec` and `coefs` are needed for the data specification.
3. If `newdata` is `NULL`, we use the extracted design matrix of our fitted model (`object`). Else, we create the design matrix corresponding to our model from `newdata` (a data frame, containing the same covariates, but possibly with different values). For this extraction, we use the `formula` of the fitted model `object`.
4. Based on the `type` argument, we will monitor (sample) the nodes `mu` (if `type = "link"`) or `z` (if `type = "response"`). `mu` corresponds to the values on the scale of the linear

predictors, while  $\mathbf{z}$  corresponds to the values on the response scale.

5. `model` is determined by the `family` of the fitted model object.
6. `prdat` is the data, `prmod` is the fitted MCMC object (we use `k` only in case of the Bionimal family, otherwise `rjags` warns us that `k` is unused for the Poisson case).
7. If the `se` argument is `FALSE`, the return value will be the point estimate vector of the prediction. If the `se` argument is `TRUE`, the return value will be a list including point estimates (`fit`) and standard errors (`se.fit`). Then, return the value.

Now let's do the prediction for a range of  $x$  values from  $-1$  to  $1$  (call it `px`):

```
R> px <- data.frame(x=seq(-1, 1, len = 10))
R> px
```

```
      x
1 -1.0000000
2 -0.7777778
3 -0.5555556
4 -0.3333333
5 -0.1111111
6  0.1111111
7  0.3333333
8  0.5555556
9  0.7777778
10 1.0000000
```

The `glm` based predictions are:

```
R> pm1link <- predict(m1, newdata=px, type="link", se=TRUE)
R> pm1resp <- predict(m1, newdata=px, type="response", se=TRUE)
R> pm2link <- predict(m2, newdata=px, type="link", se=TRUE)
R> pm2resp <- predict(m2, newdata=px, type="response", se=TRUE)
```

The `dcglm` based predictions are:

```
R> pdcm1link <- predict(dcm1, newdata=px, type="link", se=TRUE)
R> pdcm1resp <- predict(dcm1, newdata=px, type="response", se=TRUE)
R> pdcm2link <- predict(dcm2, newdata=px, type="link", se=TRUE)
R> pdcm2resp <- predict(dcm2, newdata=px, type="response", se=TRUE)
```

Fig. 1 shows prediction results.

## 10. Making the *dcglm* package

The easiest part now comes:

```
R> package.skeleton("dcglm", c("coef.dcglm", "confint.dcglm", "dcglm",
+   "fitted.dcglm", "logLik.dcglm", "predict.dcglm",
+   "print.dcglm", "print.summary.dcglm", "summary.dcglm", "vcov.dcglm"))
```

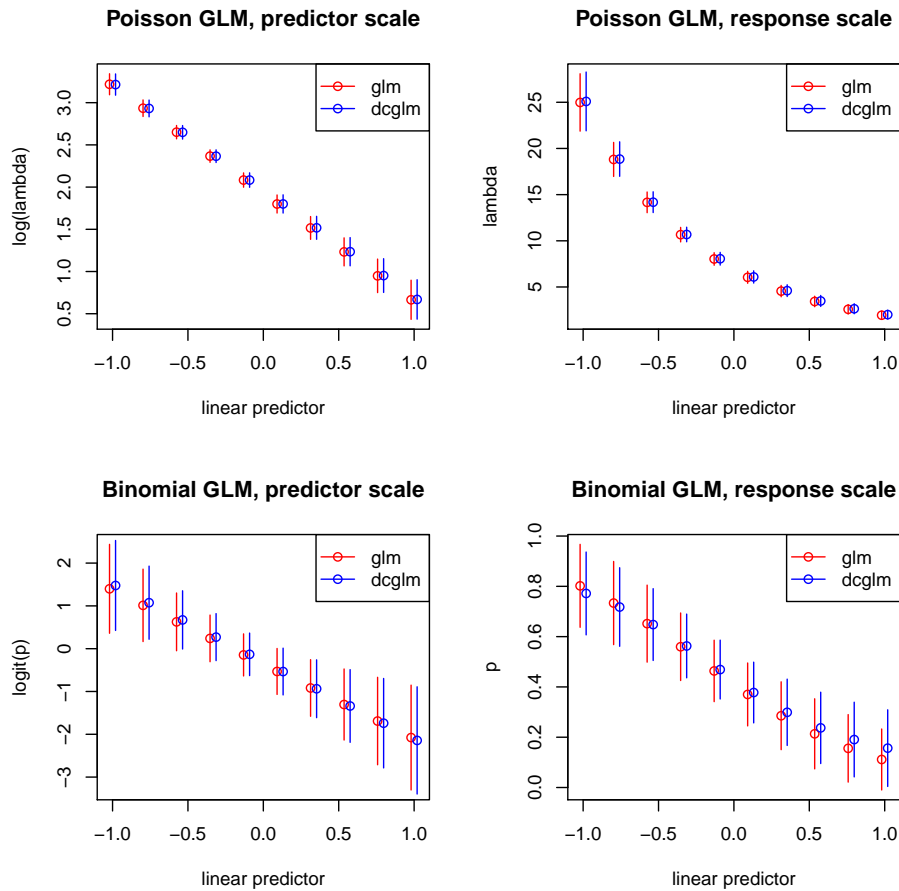


Figure 1: Prediction results based on the `glm` and `dcglm` approaches for the Poisson GLM. Points are prediction estimates, whiskers are prediction standard errors.

Follow this workflow for your own model and estimating procedure, then edit the files (read the *Writing R Extensions* manual) in the package directory, run `R CMD check` and `R CMD build`, distribute your package, and go out for a walk!

## References

- Lele SR, Dennis B, Lutscher F (2007). “Data cloning: easy maximum likelihood estimation for complex ecological models using Bayesian Markov chain Monte Carlo methods.” *Ecology Letters*, **10**, 551–563.
- Ponciano JM, Taper ML, Dennis B, Lele SR (2009). “Hierarchical models in ecology: confidence intervals, hypothesis testing, and model selection using data cloning.” *Ecology*, **90**, 356–362.

Sólymos P (2009). *dclone: Data Cloning and MCMC Tools for Maximum Likelihood Methods*.  
R package version 1.0-0, URL <http://cran.r-project.org/packages=dclone>.

**Affiliation:**

Péter Sólymos  
Alberta Biodiversity Monitoring Institute  
Department of Biological Sciences  
CW 405, Biological Sciences Bldg  
University of Alberta  
Edmonton, Alberta, T6G 2E9, Canada  
E-mail: [solyms@ualberta.ca](mailto:solyms@ualberta.ca)