

Package ‘ucminf’

April 19, 2009

Version 1.0-5

Date 2009-03-11

Title General-purpose unconstrained non-linear optimization

Author Hans Bruun Nielsen and Stig Bousgaard Mortensen

Maintainer Stig Bousgaard Mortensen <sbm@imm.dtu.dk>

Description An algorithm for general-purpose unconstrained non-linear optimization. The algorithm is of quasi-Newton type with BFGS updating of the inverse Hessian and soft line search with a trust region type monitoring of the input to the line search algorithm. The interface of ‘ucminf’ is designed for easy interchange with ‘optim’.

Depends numDeriv, R (>= 1.9.0)

License GPL (>= 2)

Repository CRAN

Date/Publication 2009-03-11 09:56:11

R topics documented:

ucminf	1
Index	6

ucminf

*General-purpose unconstrained non-linear optimization***Description**

An algorithm for general-purpose unconstrained non-linear optimization. The algorithm is of quasi-Newton type with BFGS updating of the inverse Hessian and soft line search with a trust region type monitoring of the input to the line search algorithm. The interface of ‘ucminf’ is designed for easy interchange with ‘optim’.

Usage

```
ucminf(par, fn, gr = NULL, ..., control = list(), hessian=0)
```

Arguments

<code>par</code>	Initial estimate of minimum for <code>fn</code> .
<code>fn</code>	Objective function to be minimized.
<code>gr</code>	Gradient of objective function. If <code>NULL</code> a finite difference approximation is used.
<code>...</code>	Optional arguments passed to the objective and gradient functions.
<code>control</code>	A list of control parameters. See ‘Details’.
<code>hessian</code>	Integer value: <ul style="list-style-type: none"> 0 No hessian approximation is returned. 1 Returns a numerical approximation of the Hessian using ‘hessian’ in the package ‘numDeriv’. 2 Returns final approximation of the inverse Hessian based on the series of BFGS updates during optimization. 3 Same at 2, but will also return the Hessian (the inverse of 2). If a <code>TRUE</code> or <code>FALSE</code> value is given it will switch between option 1 or 0.

Details

The algorithm is documented in (Nielsen, 2000) (see References below) together with a comparison to the Fortran subroutine ‘MINF’ and the Matlab function ‘fminunc’. The implementation of ‘ucminf’ in R uses the original Fortran version of the algorithm.

The interface in R is designed so that it is very easy to switch between using ‘ucminf’ and ‘[optim](#)’. The arguments `par`, `fn`, `gr`, and `hessian` are all the same (with a few extra options for `hessian` in ‘ucminf’). The difference is that there is no `method` argument in ‘ucminf’ and that some of the components in the `control` argument are different due to differences in the algorithms.

The algorithm can be given an initial estimate of the Hessian for the optimization and it is possible to get the final approximation of the Hessian based on the series of BFGS updates. This extra functionality may be useful for optimization in a series of related problems.

The functions `fn` and `gr` can return `Inf` or `NaN` if the functions cannot be evaluated at the supplied value, but the functions must be computable at the initial value. The functions are not allowed to return `NA`. Any names given to `par` will be copied to the vectors passed to `fn` and `gr`. No other attributes of `par` are copied over.

The `control` argument is a list that can supply any of the following components:

trace If `trace` is positive then detailed tracing information is printed for each iteration.

grtol The algorithm stops when $\|F'(x)\|_\infty \leq \text{grtol}$, that is when the largest absolute value of the gradient is less than `grtol`. Default value is `grtol = 1e-6`.

xtol The algorithm stops when $\|x - x_p\|_2 \leq \text{xtol} \cdot (\text{xtol} + \|x\|_2)$, where x_p and x are the previous and current estimate of the minimizer. Thus the algorithm stops when the last relative step length is sufficiently small. Default value is `xtol = 1e-12`.

stepmax Initial maximal allowed step length (radius of trust-region). The value is updated during the optimization. Default value is `stepmax = 1`.

maxeval The maximum number of function evaluations. A function evaluation is counted as one evaluation of the objective function and of the gradient function. Default value is `maxeval = 500`.

grad Either 'forward' or 'central'. Controls the type of finite difference approximation to be used for the gradient if no gradient function is given in the input argument 'gr'. Default value is `grad = 'forward'`.

gradstep Vector of length 2. The step length in finite difference approximation for the gradient. Step length is $|x_i| \cdot \text{gradstep}[1] + \text{gradstep}[2]$. Default value is `gradstep = c(1e-6, 1e-8)`.

invhessian.lt A vector with an initial approximation to the lower triangle of the inverse Hessian. If not given, the inverse Hessian is initialized as the identity matrix. If `H0` is the initial hessian matrix then the lower triangle of the inverse of `H0` can be found as `invhessian.lt = solve(H0)[lower.tri(H0, diag=TRUE)]`.

Value

<code>par</code>	Computed minimizer.
<code>value</code>	Objective function value at computed minimizer.
<code>convergence</code>	Flag for reason of termination: <ol style="list-style-type: none"> 1 Stopped by small gradient (<code>grtol</code>). 2 Stopped by small step (<code>xtol</code>). 3 Stopped by function evaluation limit (<code>maxeval</code>). 4 Stopped by zero step from line search -2 Computation did not start: <code>length(par) = 0</code>. -4 Computation did not start: <code>stepmax</code> is too small. -5 Computation did not start: <code>grtol</code> or <code>xtol</code> ≤ 0. -6 Computation did not start: <code>maxeval</code> ≤ 0. -7 Computation did not start: given Hessian not pos. definite.
<code>message</code>	String with reason of termination.

hessian, invhessian Estimate of (inv.) Hessian at computed minimizer. The type of estimate is given by the input argument 'hessian'.

invhessian.lt The lower triangle of the final approximation to the inverse Hessian based on the series of BFGS updates during optimization.

info Information about the search:
maxgradient $\|F'(x)\|_\infty$, the largest element in the absolute value of the gradient at the computed minimizer.
laststep Length of last step.
stepmax Final maximal allowed step length.
neval Number of calls to both objective and gradient function.

Author(s)

'UCMINF' algorithm design and Fortran code by Hans Bruun Nielsen.
 Implementation in R by Stig B. Mortensen, <sbm@imm.dtu.dk>.

References

Nielsen, H. B. (2000) 'UCMINF - An Algorithm For Unconstrained, Nonlinear Optimization', Report IMM-REP-2000-18, Department of Mathematical Modelling, Technical University of Denmark. <http://www2.imm.dtu.dk/~hbn/publ/TR0019.ps> or <http://orbit.dtu.dk/record/200975>.

The original Fortran source code can be found at <http://www2.imm.dtu.dk/~hbn/Software/ucminf.f>. The code has been slightly modified in this package to be suitable for use with R.

The general structure of the implementation in R is based on the package 'FortranCallsR' by Diethelm Wuertz.

See Also

[optim](#), [nlminb](#), [nlm](#).

Examples

```
## Rosenbrock Banana function
fR = function(x) { (1-x[1])^2+100*(x[2]-x[1]^2)^2 }
gR = function(x) { c( -400*x[1]*(x[2]-x[1]*x[1]) - 2*(1-x[1]),
                    200*(x[2]-x[1]*x[1])) }

# Find minimum
ucminf(par = c(2,.5), fn = fR, gr=gR)

# Compare hessian approximations
ucminf(par = c(2,.5), fn = fR, gr=gR, hessian=1)$hessian
ucminf(par = c(2,.5), fn = fR, gr=gR, hessian=3)$hessian

# Compare run times with optim's BFGS method
# (chosen convergence criteria result in similar accuracy)
system.time( for(i in 1:500)
```

```
        ucminf(par = c(2,0.5), fn = fR, gr=gR)
      )
system.time( for(i in 1:500)
              optim(par = c(2,0.5), fn = fR, gr=gR,method='BFGS')
            )

## Quadratic function
fQ = function(x) { sum((4*x-1)^2) }
gQ = function(x) { 32*x -8}

# Find minimum with too small stepmax and print trace
ucminf(par = c(20.5,20.0), fn = fQ, gr = gQ,
       control=list(stepmax=1,trace=TRUE))

# The same again with a larger stepmax
ucminf(par = c(20.5,20.0), fn = fQ, gr = gQ,
       control=list(stepmax=100,trace=TRUE))
```

Index

*Topic **nonlinear**

ucminf, 1

*Topic **optimize**

ucminf, 1

nlm, 4

nlminb, 4

optim, 2, 4

ucminf, 1