

Package ‘dde’

October 13, 2022

Title Solve Delay Differential Equations

Version 1.0.1

Description Solves ordinary and delay differential equations, where the objective function is written in either R or C. Suitable only for non-stiff equations, the solver uses a 'Dormand-Prince' method that allows interpolation of the solution at any point. This approach is as described by Hairer, Norsett and Wanner (1993) <ISBN:3540604529>. Support is also included for iterating difference equations.

License MIT + file LICENSE

URL <https://github.com/mrc-ide/dde>

BugReports <https://github.com/mrc-ide/dde/issues>

Depends R (>= 3.1.0)

LinkingTo ring (>= 1.0.0)

Suggests deSolve, knitr, microbenchmark, rmarkdown, testthat

RoxygenNote 6.1.1

VignetteBuilder knitr

Encoding UTF-8

Language en-GB

NeedsCompilation yes

Author Rich FitzJohn [aut, cre],
Wes Hinsley [aut],
Imperial College of Science, Technology and Medicine [cph]

Maintainer Rich FitzJohn <rich.fitzjohn@gmail.com>

Repository CRAN

Date/Publication 2020-01-16 23:20:05 UTC

R topics documented:

difeq	2
difeq_replicate	4
dopri	5
dopri_interpolate	10

Index	13
--------------	-----------

difeq *Solve difference equation*

Description

Solve a difference (or recurrence) equation by iterating it a number of times.

Usage

```
difeq(y, steps, target, parms, ..., n_out = 0L, n_history = 0L,
      grow_history = FALSE, return_history = n_history > 0, dllname = "",
      parms_are_real = TRUE, ynames = names(y), outnames = NULL,
      return_by_column = TRUE, return_initial = TRUE, return_step = TRUE,
      return_output_with_y = TRUE, restartable = FALSE,
      return_minimal = FALSE)
```

```
difeq_continue(obj, steps, y = NULL, ..., copy = FALSE, parms = NULL,
               return_history = NULL, return_by_column = NULL,
               return_initial = NULL, return_step = NULL,
               return_output_with_y = NULL, restartable = NULL)
```

```
yprev(step, i = NULL)
```

Arguments

y	The initial state of the system. Must be a numeric vector (and will be passed through as <code>.numeric</code> by this function).
steps	A vector of steps to return the system at. The <i>first</i> step is taken as step zero, and the solution will be recorded at every other step in the vector. So to step a system from time zero to times 1, 2, 3, ..., n use 0:n. Must be integer values and will be passed through as <code>.integer</code> (which may truncate or otherwise butcher non-integer values).
target	The target function to advance. This can either be an R function taking arguments <code>n</code> , <code>i</code> , <code>t</code> , <code>y</code> , <code>parms</code> or be a scalar character with the name of a compiled function with arguments <code>size_t n</code> , <code>size_t step</code> , <code>double time</code> , <code>const double *y</code> , <code>double *dydt</code> , <code>size_t n_out</code> , <code>double *output</code> void <code>*data</code> .
parms	Parameters to pass through to the difference function

...	Dummy arguments - nothing is allowed here, but this means that all further arguments <i>must</i> be specified by name (not order) so I can easily reorder them later on.
n_out	The number of output variables (in addition to the difference equation variables). If given, then an R function must return an <i>attribute</i> output with the output variables.
n_history	The number of iterations of history to save during the simulation. By default, no history is saved.
grow_history	Logical indicating if history should be grown during the simulation. If FALSE (the default) then when history is used it is overwritten as needed (so only the most recent n_history elements are saved. This may require some tuning so that you have enough history to run your simulation (i.e. to the longest delay) or an error will be thrown when it underflows. The required history length will vary with your delay sizes and with the timestep for dopri. If TRUE, then history will grow as the buffer is exhausted. The growth is geometric, so every time it reaches the end of the buffer it will increase by a factor of about 1.6 (see the ring documentation). This may consume more memory than necessary, but may be useful where you don't want to care about picking the history length carefully.
return_history	Logical indicating if history is to be returned. By default, history is returned if n_history is nonzero.
dllname	Name of the shared library (without extension) to find the function func in the case where func refers to compiled function.
parms_are_real	Logical, indicating if parms should be treated as vector of doubles by func (when it is a compiled function). If TRUE (the default), then REAL(parms), which is double* is passed through. If FALSE then if parms is an externalptr type (EXTPTRSXP) we pass through the result of R_ExternalPtrAddr, otherwise we pass parms through unmodified as a SEXP. In the last case, in your target function you will need to include <Rinternals.h>, cast to SEXP and then pull it apart using the R API (or Rcpp).
y_names	Logical, indicating if the output should be named following the names of the input vector y. Alternatively, if y_names is a character vector of the same length as y, these will be used as the output names.
out_names	An optional character vector, used when n_out is greater than 0, to name the model output matrix.
return_by_column	Logical, indicating if the output should be returned organised by column (rather than row). This incurs a slight cost for transposing the matrices. If you can work with matrices that are transposed relative to deSolve, then set this to FALSE.
return_initial	Logical, indicating if the output should include the initial conditions. Specifying FALSE avoids binding this onto the output.
return_step	Logical, indicating if a row (or column if return_by_column is TRUE) representing step is included.
return_output_with_y	Logical, indicating if the output should be bound together with the returned matrix y (as it is with deSolve). If FALSE, then output will be returned as the attribute output.

restartable	Logical, indicating if the problem should be restartable. If TRUE, then the return value of a simulation can be passed to <code>difeq_restart</code> to continue the simulation after arbitrary changes to the state or the parameters. Note that this is really only useful for delay difference equations where you want to keep the history but make changes to the parameters or to the state vector while keeping the history of the problem so far.
return_minimal	Shorthand option - if set to TRUE then it sets all of <code>return_by_column</code> , <code>return_initial</code> , <code>return_time</code> , <code>return_output_with_y</code> to FALSE
obj	An object to continue from; this must be the results of running a simulation with the option <code>restartable = TRUE</code> . Note that continuing a problem moves the pointer along in time (unless <code>copy = TRUE</code> , and that the incoming time (<code>times[[1]]</code>) must equal the previous time <i>exactly</i> .
copy	Logical, indicating if the pointer should be copied before continuing. If TRUE, this is non-destructive with respect to the data in the original pointer so the problem can be restarted multiple times. By default this is FALSE because there is a (potentially very small) cost to this operation.
step	The step to access (not that this is not an offset, but the actual step; within your target function you'd write things like <code>yprev(step - 1)</code> to get the previous step.
i	index within the state vector <code>y</code> to return. The index here is R-style base-1 indexing, so pass 1 in to access the first element. This can be left NULL to return all the elements or a vector longer than one.

Examples

```
# Here is a really simple equation that just increases by 'p' each
# time (p is the parameter vector and could be any R structure).
rhs <- function(i, y, p) y + p

y0 <- 1
t <- 0:10
p <- 5
dde::difeq(y0, t, rhs, p)
```

difeq_replicate *Solve difference equations repeatedly*

Description

Solve a replicate set of difference (or recurrence) equation by iterating it a number of times. This is a wrapper around `difeq` that does not (yet) do anything clever to avoid many allocations.

Usage

```
difeq_replicate(n, y, ..., as_array = TRUE)
```

Arguments

n	Number of replicates. It is an error to request zero replicates.
y	The initial state of the system. Must be either a numeric vector or a list of numeric vectors. If the latter, it must have length n.
...	Additional arguments passed through to <code>difeq</code> .
as_array	Logical, indicating if the output should be converted into an array. If TRUE then <code>res[, , i]</code> will contain the <i>i</i> 'th replicate, if FALSE then <code>res[[i]]</code> does instead. If both <code>as_array</code> and <code>restartable</code> are TRUE, then the attributes <code>ptr</code> and <code>restart_data</code> will be present as a list of restarting information for <code>difeq_continue</code> , though using these is not yet supported.

Details

It is not currently possible to replicate over a set of parameters at once yet; the same parameter set will be used for all replications.

The details of how replication is done here are all considered implementation details and are up for change in the future - in particular if the models are run in turn or simultaneously (and the effect that has on the random number stream). Logic around naming output may change in future too; note that varying names in the `y` here will have some unexpected behaviours.

Examples

```
# Here is a really simple equation that does a random walk with
# steps that are normally distributed:
rhs <- function(i, y, p) y + runif(1)
y0 <- 1
t <- 0:10
p <- 5
dde::difeq_replicate(10, y0, t, rhs, p)
```

dopri

Integrate ODE/DDE with dopri

Description

Integrate an ODE or DDE with dopri.

Usage

```
dopri(y, times, func, parms, ..., n_out = 0L, output = NULL,
      rtol = 1e-06, atol = 1e-06, step_size_min = 0,
      step_size_max = Inf, step_size_initial = 0, step_max_n = 100000L,
      tcrit = NULL, event_time = NULL, event_function = NULL,
      method = "dopri5", stiff_check = 0, verbose = FALSE,
      callback = NULL, n_history = 0, grow_history = FALSE,
```

```

return_history = n_history > 0, dllname = "",
parms_are_real = TRUE, ynames = names(y), outnames = NULL,
return_by_column = TRUE, return_initial = TRUE, return_time = TRUE,
return_output_with_y = TRUE, return_statistics = FALSE,
restartable = FALSE, return_minimal = FALSE)

dopri5(y, times, func, parms, ...)

dopri853(y, times, func, parms, ...)

dopri_continue(obj, times, y = NULL, ..., copy = FALSE, parms = NULL,
  tcrit = NULL, return_history = NULL, return_by_column = NULL,
  return_initial = NULL, return_statistics = NULL,
  return_time = NULL, return_output_with_y = NULL,
  restartable = NULL)

ylag(t, i = NULL)

```

Arguments

<code>y</code>	Initial conditions for the integration
<code>times</code>	Times where output is needed. Unlike <code>deSolve</code> we won't actually stop at these times, but instead interpolate back to get the result.
<code>func</code>	Function to integrate. Can be an R function of arguments <code>t</code> , <code>y</code> , <code>parms</code> , returning a numeric vector, or it can be the name or address of a C function with arguments <code>size_t n</code> , <code>double t</code> , <code>const double *y</code> , <code>double *dydt</code> , <code>void *data</code> .
<code>parms</code>	Parameters to pass through to the derivatives.
<code>...</code>	Dummy arguments - nothing is allowed here, but this means that all further arguments <i>must</i> be specified by name (not order) so I can easily reorder them later on.
<code>n_out</code>	Number of "output" variables (not differential equation variables) to compute via the routine output.
<code>output</code>	The output routine; either an R function taking arguments <code>t</code> , <code>y</code> , <code>parms</code> or the name/address of a C function taking arguments <code>size_t n</code> , <code>double t</code> , <code>const double *y</code> , <code>size_t n_out</code> , <code>double *out</code> , <code>void *data</code> .
<code>rtol</code>	The per-step relative tolerance. The total accuracy will be less than this.
<code>atol</code>	The per-step absolute tolerance.
<code>step_size_min</code>	The minimum step size. The actual minimum used will be the largest of the absolute value of this <code>step_size_min</code> or <code>.Machine\$double.eps</code> . If the integration attempts to make a step smaller than this, it will throw an error, stopping the integration (note that this differs from the treatment of <code>hmin</code> in <code>deSolve::lsoda</code>).
<code>step_size_max</code>	The largest step size. By default there is no maximum step size (<code>Inf</code>) so the solver can take as large a step as it wants to. If you have short-lived fluctuations in your rhs that the solver may skip over by accident, then specify a smaller maximum step size here (or use <code>tcrit</code> below).

<code>step_size_initial</code>	The initial step size. By default the integrator will guess the step size automatically, but one can be given here instead.
<code>step_max_n</code>	The maximum number of steps allowed. If the solver takes more steps than this it will throw an error. Note the number of evaluations of <code>func</code> will be about 6 times the number of steps (or 11 times if using <code>method = "dopri853"</code>).
<code>tcrit</code>	An optional vector of critical times that the solver must stop at (rather than interpolating over). This can include an end time that we can't go past, or points within the integration that must be stopped at exactly (for example cases where the derivatives change abruptly). Note that this differs from the interpretation of this parameter in <code>deSolve</code> ; there <code>tcrit</code> is a single time that integration may not go past – with <code>dde</code> we never go past the final time, and this is just for times that fall <i>within</i> the range of times in <code>times</code> .
<code>event_time</code>	Vector of times to fire events listed in <code>event_function</code> at
<code>event_function</code>	Function to fire at events. For R models (<code>func</code> is an R function and <code>dllname</code> is empty), this must be either a single R function (same function for all events) or a list of R functions. For C models, this must be a single C function (same requirements as <code>func</code> or <code>output</code> or a list/vector of these as appropriate).
<code>method</code>	The integration method to use, as a string. The supported methods are <code>"dopri5"</code> (5th order method with 4th order dense output) and <code>"dopri853"</code> (8th order method with 7th order output and embedded 5th and 3rd order schemes). Alternatively, use the functions <code>dopri5</code> or <code>dopri853</code> which simply sets this argument.
<code>stiff_check</code>	How often to check that the problem has become stiff. If zero, then the problem is never checked, and if positive then the problem is checked every <code>stiff_check</code> accepted steps. The actual check is based off the algorithm in Hairer's implementation of the solvers and may be overly strict, especially for delay equations with the 853 method (in my limited experience with it).
<code>verbose</code>	Be verbose, and print information about each step. This may be useful for learning about models that misbehave. Valid values are <code>TRUE</code> (enable debugging) or <code>FALSE</code> (disable debugging) or use one of <code>dopri:::VERBOSE_QUIET</code> , <code>dopri:::VERBOSE_STEP</code> or <code>VERBOSE:::VERBOSE_EVAL</code> . If an R function is provided as the argument <code>callback</code> then this function will also be called at each step or evaluation (see below for details).
<code>callback</code>	Callback function that can be used to make verbose output more useful. This can be used to return more information about the evaluation as it proceeds, generally as information printed to the screen. The function must accept arguments <code>t</code> , <code>y</code> and <code>dydt</code> . See <code>Details</code> for further information.
<code>n_history</code>	Number of history points to retain. This needs to be greater than zero for delay differential equations to work. Alternatively, this may be greater than zero to return model outputs that can be inspected later.
<code>grow_history</code>	Logical indicating if history should be grown during the simulation. If <code>FALSE</code> (the default) then when history is used it is overwritten as needed (so only the most recent <code>n_history</code> elements are saved. This may require some tuning so that you have enough history to run your simulation (i.e. to the longest delay) or an error will be thrown when it underflows. The required history length will vary with your delay sizes and with the timestep for <code>dopri</code> . If <code>TRUE</code> , then history

will grow as the buffer is exhausted. The growth is geometric, so every time it reaches the end of the buffer it will increase by a factor of about 1.6 (see the `ring` documentation). This may consume more memory than necessary, but may be useful where you don't want to care about picking the history length carefully.

<code>return_history</code>	Logical indicating if history should be returned alongside the output or discarded. By default, history is retained if <code>n_history</code> is greater than 0, but that might change (and may not be desirable unless you plan on actually using it).
<code>dllname</code>	Name of the shared library (without extension) to find the function <code>func</code> (and output if given) in the case where <code>func</code> refers to compiled function.
<code>parms_are_real</code>	Logical, indicating if <code>parms</code> should be treated as vector of doubles by <code>func</code> (when it is a compiled function). If <code>TRUE</code> (the default), then <code>REAL(parms)</code> , which is <code>double*</code> is passed through. If <code>FALSE</code> then if <code>parms</code> is an externalptr type (<code>EXTPTRSXP</code>) we pass through the result of <code>R_ExternalPtrAddr</code> , otherwise we pass <code>parms</code> through unmodified as a <code>SEXP</code> . In the last case, in your target function you will need to include <code><Rinternals.h></code> , cast to <code>SEXP</code> and then pull it apart using the R API (or <code>Rcpp</code>).
<code>ynames</code>	Logical, indicating if the output should be named following the names of the input vector <code>y</code> . Alternatively, if <code>ynames</code> is a character vector of the same length as <code>y</code> , these will be used as the output names.
<code>outnames</code>	An optional character vector, used when <code>n_out</code> is greater than 0, to name the model output matrix.
<code>return_by_column</code>	Logical, indicating if the output should be returned organised by column (rather than row). This incurs a slight cost for transposing the matrices. If you can work with matrices that are transposed relative to <code>deSolve</code> , then set this to <code>FALSE</code> .
<code>return_initial</code>	Logical, indicating if the output should include the initial conditions. Specifying <code>FALSE</code> avoids binding this onto the output.
<code>return_time</code>	Logical, indicating if a row (or column if <code>return_by_column</code> is <code>TRUE</code>) representing time is included. If <code>FALSE</code> , this is not added.
<code>return_output_with_y</code>	Logical, indicating if the output should be bound together with the returned matrix <code>y</code> (as it is with <code>deSolve</code>). If <code>FALSE</code> , then output will be returned as the attribute output.
<code>return_statistics</code>	Logical, indicating if statistics about the run should be included. If <code>TRUE</code> , then an integer vector containing the number of target evaluations, steps, accepted steps and rejected steps is returned (the vector is named).
<code>restartable</code>	Logical, indicating if the problem should be restartable. If <code>TRUE</code> , then the return value of an integration can be passed to <code>dopri_restart</code> to continue the integration after arbitrary changes to the state or the parameters. Note that when using delay models, the integrator is fairly naive about how abrupt changes in the state space are dealt with, and may perform very badly with <code>method = "dopri853"</code> which assumes a fairly smooth problem. Note that this is really only useful for delay differential equations where you want to keep the history but make changes to the parameters or to the state vector while keeping the history of the problem so far.

<code>return_minimal</code>	Shorthand option - if set to TRUE then it sets all of <code>return_by_column</code> , <code>return_initial</code> , <code>return_time</code> , <code>return_output_with_y</code> to FALSE
<code>obj</code>	An object to continue from; this must be the results of running an integration with the option <code>restartable = TRUE</code> . Note that continuing a problem moves the pointer along in time (unless <code>copy = TRUE</code> , and that the incoming time (<code>times[[1]]</code>) must equal the previous time <i>exactly</i> .
<code>copy</code>	Logical, indicating if the pointer should be copied before continuing. If TRUE, this is non-destructive with respect to the data in the original pointer so the problem can be restarted multiple times. By default this is FALSE because there is a (potentially very small) cost to this operation.
<code>t</code>	The time to access (not that this is not an offset, but the actual time; within your target function you'd write things like <code>tlag(t - 1)</code> to get 1 time unit ago.
<code>i</code>	index within the state vector <code>y</code> to return. The index here is R-style base-1 indexing, so pass 1 in to access the first element. This can be left NULL to return all the elements or a vector longer than one.

Details

Like `deSolve::lsoda`, this function has *many* arguments. This is far from ideal, and I would welcome any approach for simplifying it a bit.

The options `return_by_column`, `return_initial`, `return_time`, `return_output_with_y` exist because these options all carry out modifications of the data at the end of solving the ODE and this can incur a small but measurable cost. When solving an ODE repeatedly (e.g., in the context of an MCMC or optimisation) it may be useful to do as little as possible. For simple problems this can save around 5-10% of the total computational time (especially the transpose). The shorthand option `return_minimal` will set all to FALSE when used.

Value

At present the return value is transposed relative to `deSolve`. This might change in future.

Verbose output and callbacks

Debugging a failed integration can be difficult, but `dopri` provides a couple of tools to get more information about where a failure might have occurred. Most simply, one can pass `verbose = TRUE` which will print information about the time and the step size at each point just before the step is stated. Passing in `verbose = dde::VERBOSE_EVAL` will print information just before every evaluation of the target function (there are several evaluations per step).

However, this does not provide information about the state just before failure. To get that, one must provide a callback function - this is an R function that will be called just before a step or evaluation (based on the value of the `verbose` argument) in place of the default print. Define a callback function with arguments `t`, `h` and `y` where `t` is the time (beginning of a step or location of an evaluation), `h` is the step size (or NA for an evaluation) and `y` is the state at the point of the step or evaluation. Your callback function can do anything - you can print to the screen (using `cat` or `message`), you can store results using a closure and `<<-` or you could conditionally use a `browser()` call to debug interactively. However, it is not possible for the callback to affect the solution short of throwing an error and interrupting it. See the Examples for an example of use.

See Also

[dopri_interpolate](#) which can be used to efficiently sample from output of `dopri`, and the package vignette which shows in more detail how to solve delay differential equations and to use compiled objective functions.

Examples

```
# The lorenz attractor:
lorenz <- function(t, y, p) {
  sigma <- p[[1L]]
  R <- p[[2L]]
  b <- p[[3L]]
  c(sigma * (y[[2L]] - y[[1L]]),
    R * y[[1L]] - y[[2L]] - y[[1L]] * y[[3L]],
    -b * y[[3L]] + y[[1L]] * y[[2L]])
}

p <- c(10, 28, 8 / 3)
y0 <- c(10, 1, 1)

tt <- seq(0, 100, length.out = 40000)
y <- dde::dopri(y0, tt, lorenz, p, return_time = FALSE)
plot(y[, c(1, 3)], type = "l", lwd = 0.5, col = "#00000066")

# If we want to print progress as the integration progresses we can
# use the verbose argument:
y <- dde::dopri(y0, c(0, 0.1), lorenz, p, verbose = TRUE)

# Or print the y values too using a callback:
callback <- function(t, h, y) {
  message(sprintf("t: %f, h: %e, y: [%s]", t, h,
    paste(format(y, 5), collapse = ", ")))
}
y <- dde::dopri(y0, c(0, 0.1), lorenz, p, verbose = TRUE,
  callback = callback)
```

dopri_interpolate

Interpolate Dormand-Prince output

Description

Interpolate the Dormand-Prince output after an integration. This only interpolates the core integration variables and not any additional output variables.

Usage

```
dopri_interpolate(h, t)
```

Arguments

- h** The interpolation history. This can be the output running dopri with `return_history = TRUE`, or the history attribute of this object (retrievable with `attr(res, "history")`).
- t** The times at which interpolated output is required. These times must fall within the included history (i.e., the times that the original simulation was run) or an error will be thrown.

Details

This decouples the integration of the equations and the generation of output; it is not necessary for use of the package, but may come in useful where you need to do (for example) root finding on the time course of a problem, or generate minimal output in some cases and interrogate the solution more deeply in others. See the examples and the package vignette for a full worked example.

Author(s)

Rich FitzJohn

Examples

```
# Here is the Lorenz attractor implemented as an R function
lorenz <- function(t, y, p) {
  sigma <- p[[1L]]
  R <- p[[2L]]
  b <- p[[3L]]
  c(sigma * (y[[2L]] - y[[1L]]),
    R * y[[1L]] - y[[2L]] - y[[1L]] * y[[3L]],
    -b * y[[3L]] + y[[1L]] * y[[2L]])
}

# Standard parameters and a reasonable starting point:
p <- c(10, 28, 8 / 3)
y0 <- c(10, 1, 1)

# Run the integration for times [0, 50] and return minimal output,
# but *do* record and return history.
y <- dopri(y0, c(0, 50), lorenz, p,
  n_history = 5000, return_history = TRUE,
  return_time = FALSE, return_initial = FALSE,
  return_by_column = FALSE)

# Very little output is returned (just 3 numbers being the final
# state of the system), but the "history" attribute is fairly
# large matrix of history information. It is not printed though
# as its contents should not be relied on. What does matter is
# the range of supported times printed (i.e., [0, 50]) and the
# number of entries (~2000).
y

# Generate an interpolated set of variables using this; first for
# 1000 steps over the full range:
```

```
tt <- seq(0, 50, length.out = 1000)
yy <- dopri_interpolate(y, tt)
plot(yy[, c(1, 3)], type = "l")

# Then for 50000
tt <- seq(0, 50, length.out = 50000)
yy <- dopri_interpolate(y, tt)
plot(yy[, c(1, 3)], type = "l")
```

Index

difeq, [2](#), [4](#), [5](#)
difeq_continue (difeq), [2](#)
difeq_replicate, [4](#)
dopri, [5](#)
dopri5 (dopri), [5](#)
dopri853 (dopri), [5](#)
dopri_continue (dopri), [5](#)
dopri_interpolate, [10](#), [10](#)

ylag (dopri), [5](#)
yprev (difeq), [2](#)