

Package ‘nlmixr2’

February 20, 2025

Title Nonlinear Mixed Effects Models in Population PK/PD

Version 3.0.2

Description Fit and compare nonlinear mixed-effects models in differential equations with flexible dosing information commonly seen in pharmacokinetics and pharmacodynamics (Almquist, Leander, and Jirstrand 2015 <[doi:10.1007/s10928-015-9409-1](https://doi.org/10.1007/s10928-015-9409-1)>). Differential equation solving is by compiled C code provided in the 'rxode2' package (Wang, Hallow, and James 2015 <[doi:10.1002/psp4.12052](https://doi.org/10.1002/psp4.12052)>).

License GPL (>= 3)

Encoding UTF-8

RoxygenNote 7.3.2

Imports nlmixr2est (>= 2.2.2), nlmixr2extra (>= 2.0.10), rxode2 (>= 2.1.3), lotri (>= 0.4.3), nlmixr2plot (>= 2.0.8), magrittr, crayon, cli

Depends nlmixr2data

Suggests rmarkdown, knitr, devtools, ggplot2, testthat, n1qn1, withr

BugReports <https://github.com/nlmixr2/nlmixr2/issues/>

URL <https://nlmixr2.org/>, <https://github.com/nlmixr2/nlmixr2/>

NeedsCompilation no

Author Matthew Fidler [aut, cre] (<<https://orcid.org/0000-0001-8538-6691>>),
Anne Keunecke [ctb],
Rik Schoemaker [ctb] (<<https://orcid.org/0000-0002-7538-3005>>),
Justin Wilkins [ctb] (<<https://orcid.org/0000-0002-7099-9396>>),
Wenping Wang [ctb],
Mirjam Trame [ctb],
John Harrold [ctb],
Bill Denney [ctb] (<<https://orcid.org/0000-0002-5759-428X>>),
Theodoros Papathanasiou [ctb],
Richard Hooijmaijers [ctb]

Maintainer Matthew Fidler <matthew.fidler@gmail.com>

Repository CRAN

Date/Publication 2025-02-20 03:40:02 UTC

Contents

addCwres	2
addNpde	4
addTable	5
bobyqaControl	7
bootplot	14
bootstrapFit	14
foceiControl	16
lbfgsb3cControl	30
n1qn1Control	37
newuoaControl	43
nlmControl	50
nlmeControl	58
nlminbControl	62
nlmixr2	69
nlmixr2CheckInstall	81
nlsControl	81
optimControl	89
preconditionFit	97
profileFixed	98
profileFixedSingle	99
profileLlp	99
saemControl	100
setOfv	104
tableControl	105
traceplot	107
uobyqaControl	108
vpcCens	115
vpcCensTad	116
vpcPlot	117
vpcPlotTad	120
vpcSim	121
Index	123

addCwres

Add CWRES

Description

This returns a new fit object with CWRES attached

Usage

```
addCwres(fit, focei = TRUE, updateObject = TRUE, envir = parent.frame(1))
```

Arguments

fit	nlmixr2 fit without WRES/CWRES
focei	Boolean indicating if the focei objective function is added. If not the foce objective function is added.
updateObject	Boolean indicating if the original fit object should be updated. By default this is true.
envir	Environment that should be checked for object to update. By default this is the global environment.

Value

fit with CWRES

Author(s)

Matthew L. Fidler

Examples

```

one.cmt <- function() {
  ini({
    ## You may label each parameter with a comment
    tka <- 0.45 # Log Ka
    tcl <- log(c(0, 2.7, 100)) # Log Cl
    ## This works with interactive models
    ## You may also label the preceding line with label("label text")
    tv <- 3.45; label("log V")
    ## the label("Label name") works with all models
    eta.ka ~ 0.6
    eta.cl ~ 0.3
    eta.v ~ 0.1
    add.sd <- 0.7
  })
  model({
    ka <- exp(tka + eta.ka)
    cl <- exp(tcl + eta.cl)
    v <- exp(tv + eta.v)
    linCmt() ~ add(add.sd)
  })
}

f <- try(nlmixr2(one.cmt, theo_sd, "saem"))

print(f)

# even though you may have forgotten to add the cwres, you can add it to the data.frame:

if (!inherits(f, "try-error")) {
  f <- try(addCwres(f))
}

```

```
    print(f)
  }

# Note this also adds the FOCEi objective function
```

addNpde	<i>NPDE calculation for nlmixr2</i>
---------	-------------------------------------

Description

NPDE calculation for nlmixr2

Usage

```
addNpde(
  object,
  updateObject = TRUE,
  table = tableControl(),
  ...,
  envir = parent.frame(1)
)
```

Arguments

object	nlmixr2 fit object
updateObject	Boolean indicating if original object should be updated. By default this is TRUE.
table	'tableControl()' list of options
...	Additional arguments passed to <code>nlmixr2est::addNpde()</code> .
envir	Environment that should be checked for object to update. By default this is the global environment.

Value

New nlmixr2 fit object

Author(s)

Matthew L. Fidler

Examples

```

one.cmt <- function() {
  ini({
    ## You may label each parameter with a comment
    tka <- 0.45 # Log Ka
    tcl <- log(c(0, 2.7, 100)) # Log Cl
    ## This works with interactive models
    ## You may also label the preceding line with label("label text")
    tv <- 3.45; label("log V")
    ## the label("Label name") works with all models
    eta.ka ~ 0.6
    eta.cl ~ 0.3
    eta.v ~ 0.1
    add.sd <- 0.7
  })
  model({
    ka <- exp(tka + eta.ka)
    cl <- exp(tcl + eta.cl)
    v <- exp(tv + eta.v)
    linCmt() ~ add(add.sd)
  })
}

f <- nlmixr2(one.cmt, theo_sd, "saem")

# even though you may have forgotten to add the NPDE, you can add it to the data.frame:

f <- addNpde(f)

```

addTable

Add table information to nlmixr2 fit object without tables

Description

Add table information to nlmixr2 fit object without tables

Usage

```

addTable(
  object,
  updateObject = FALSE,
  data = object$dataSav,
  thetaEtaParameters = object$foceiThetaEtaParameters,
  table = tableControl(),
  keep = NULL,
  drop = NULL,

```

```

    envir = parent.frame(1)
  )

```

Arguments

object	nlmixr2 family of objects
updateObject	Update the object (default FALSE)
data	Saved data from
thetaEtaParameters	Internal theta/eta parameters
table	a 'tableControl()' list of options
keep	Character Vector of items to keep
drop	Character Vector of items to drop or NULL
envir	Environment to search for updating

Value

Fit with table information attached

Author(s)

Matthew Fidler

Examples

```

one.cmt <- function() {
  ini({
    ## You may label each parameter with a comment
    tka <- 0.45 # Log Ka
    tcl <- log(c(0, 2.7, 100)) # Log Cl
    ## This works with interactive models
    ## You may also label the preceding line with label("label text")
    tv <- 3.45; label("log V")
    ## the label("Label name") works with all models
    eta.ka ~ 0.6
    eta.cl ~ 0.3
    eta.v ~ 0.1
    add.sd <- 0.7
  })
  model({
    ka <- exp(tka + eta.ka)
    cl <- exp(tcl + eta.cl)
    v <- exp(tv + eta.v)
    linCmt() ~ add(add.sd)
  })
}

# run without tables step
f <- nlmixr2(one.cmt, theo_sd, "saem", control=list(calcTables=FALSE))

```

```

print(f)

# Now add the tables

f <- addTable(f)

print(f)

```

bobyqaControl

Control for bobyqa estimation method in nlmixr2

Description

Control for bobyqa estimation method in nlmixr2

Usage

```

bobyqaControl(
  npt = NULL,
  rhobeg = NULL,
  rhoend = NULL,
  iprint = 0L,
  maxfun = 100000L,
  returnBobyqa = FALSE,
  stickyRecalcN = 4,
  maxOdeRecalc = 5,
  odeRecalcFactor = 10^(0.5),
  useColor = crayon::has_color(),
  printNcol = floor((getOption("width") - 23)/12),
  print = 1L,
  normType = c("rescale2", "mean", "rescale", "std", "len", "constant"),
  scaleType = c("nlmixr2", "norm", "mult", "multAdd"),
  scaleCmax = 1e+05,
  scaleCmin = 1e-05,
  scaleC = NULL,
  scaleTo = 1,
  rxControl = NULL,
  optExpression = TRUE,
  sumProd = FALSE,
  literalFix = TRUE,
  addProp = c("combined2", "combined1"),
  calcTables = TRUE,
  compress = TRUE,
  covMethod = c("r", ""),
  adjObf = TRUE,

```

```

    ci = 0.95,
    sigdig = 4,
    sigdigTable = NULL,
    ...
)

```

Arguments

npt	The number of points used to approximate the objective function via a quadratic approximation. The value of npt must be in the interval $[n+2, (n+1)(n+2)/2]$ where n is the number of parameters in 'par'. Choices that exceed $2*n+1$ are not recommended. If not defined, it will be set to $\min(n * 2, n+2)$.
rhobeg	'rhobeg' and 'rhoend' must be set to the initial and final values of a trust region radius, so both must be positive with $0 < \text{rhoend} < \text{rhobeg}$. Typically 'rhobeg' should be about one tenth of the greatest expected change to a variable. If the user does not provide a value, this will be set to $\min(0.95, 0.2 * \max(\text{abs}(\text{par})))$. Note also that smallest difference $\text{abs}(\text{upper-lower})$ should be greater than or equal to rhobeg^2 . If this is not the case then 'rhobeg' will be adjusted.
rhoend	The smallest value of the trust region radius that is allowed. If not defined, then $1e-6$ times the value set for 'rhobeg' will be used.
iprint	The value of 'iprint' should be set to an integer value in '0, 1, 2, 3, ...', which controls the amount of printing. Specifically, there is no output if 'iprint=0' and there is output only at the start and the return if 'iprint=1'. Otherwise, each new value of 'rho' is printed, with the best vector of variables so far and the corresponding value of the objective function. Further, each new value of the objective function with its variables are output if 'iprint=3'. If 'iprint > 3', the objective function value and corresponding variables are output every 'iprint' evaluations. Default value is '0'.
maxfun	The maximum allowed number of function evaluations. If this is exceeded, the method will terminate.
returnBobyqa	return the bobyqa output instead of the nlmixr2 fit
stickyRecalcN	The number of bad ODE solves before reducing the atol/rtol for the rest of the problem.
maxOdeRecalc	Maximum number of times to reduce the ODE tolerances and try to resolve the system if there was a bad ODE solve.
odeRecalcFactor	The ODE recalculation factor when ODE solving goes bad, this is the factor the rtol/atol is reduced
useColor	Boolean indicating if foci can use ASCII color codes
printNcol	Number of columns to printout before wrapping parameter estimates/gradient
print	Integer representing when the outer step is printed. When this is 0 or do not print the iterations. 1 is print every function evaluation (default), 5 is print every 5 evaluations.
normType	This is the type of parameter normalization/scaling used to get the scaled initial values for nlmixr2. These are used with scaleType of.

With the exception of rescale2, these come from **Feature Scaling**. The rescale2 The rescaling is the same type described in the **OptdesX** software manual. In general, all all scaling formula can be described by:

$$= \left(\frac{v_{scaled}}{v_{unscaled} - C_1} \right) / C_2$$

Where

The other data normalization approaches follow the following formula

$$= \left(\frac{v_{scaled}}{v_{unscaled} - C_1} \right) / C_2$$

- rescale2 This scales all parameters from (-1 to 1). The relative differences between the parameters are preserved with this approach and the constants are:

$$C_1 = (\max(\text{all unscaled values}) + \min(\text{all unscaled values})) / 2$$

$$C_2 = (\max(\text{all unscaled values}) - \min(\text{all unscaled values})) / 2$$

- rescale or min-max normalization. This rescales all parameters from (0 to 1). As in the rescale2 the relative differences are preserved. In this approach:

$$C_1 = \min(\text{all unscaled values})$$

$$C_2 = \max(\text{all unscaled values}) - \min(\text{all unscaled values})$$

- mean or mean normalization. This rescales to center the parameters around the mean but the parameters are from 0 to 1. In this approach:

$$C_1 = \text{mean}(\text{all unscaled values})$$

$$C_2 = \max(\text{all unscaled values}) - \min(\text{all unscaled values})$$

- `std` or standardization. This standardizes by the mean and standard deviation. In this approach:

$$C_1$$

= mean(all unscaled values)

$$C_2$$

= sd(all unscaled values)

- `len` or unit length scaling. This scales the parameters to the unit length. For this approach we use the Euclidean length, that is:

$$C_1$$

= 0

$$C_2$$

=

$$\sqrt{(v_1^2 + v_2^2 + \dots + v_n^2)}$$

- `constant` which does not perform data normalization. That is

$$C_1$$

= 0

$$C_2$$

= 1

scaleType

The scaling scheme for `nlmixr2`. The supported types are:

- `nlmixr2` In this approach the scaling is performed by the following equation:

$$v_{scaled} = (v_{current} - v_{init}) * scaleC[i] + scaleTo$$

The `scaleTo` parameter is specified by the `normType`, and the scales are specified by `scaleC`.

- `norm` This approach uses the simple scaling provided by the `normType` argument.
- `mult` This approach does not use the data normalization provided by `normType`, but rather uses multiplicative scaling to a constant provided by the `scaleTo` argument.

In this case:

$$v_{scaled}$$

$$= \frac{v_{current}}{v_{init}} *scaleTo$$

- `multAdd` This approach changes the scaling based on the parameter being specified. If a parameter is defined in an exponential block (ie $\exp(\theta)$), then it is scaled on a linearly, that is:

$$= (\frac{v_{scaled}}{v_{current} - v_{init}}) + scaleTo$$

Otherwise the parameter is scaled multiplicatively.

$$= \frac{v_{scaled}}{v_{current}} / v_{init} *scaleTo$$

`scaleCmax`

Maximum value of the `scaleC` to prevent overflow.

`scaleCmin`

Minimum value of the `scaleC` to prevent underflow.

`scaleC`

The scaling constant used with `scaleType=nlmixr2`. When not specified, it is based on the type of parameter that is estimated. The idea is to keep the derivatives similar on a log scale to have similar gradient sizes. Hence parameters like $\log(\exp(\theta))$ would have a scaling factor of 1 and $\log(\theta)$ would have a scaling factor of `ini_value` (to scale by $1/value$; ie $d/dt(\log(ini_value)) = 1/ini_value$ or `scaleC=ini_value`)

- For parameters in an exponential (ie $\exp(\theta)$) or parameters specifying powers, `boxCox` or `yeoJohnson` transformations, this is 1.
- For additive, proportional, lognormal error structures, these are given by $0.5 * \text{abs}(\text{initial_estimate})$
- Factorials are scaled by $\text{abs}(1/\text{digamma}(\text{initial_estimate}+1))$
- parameters in a log scale (ie $\log(\theta)$) are transformed by $\log(\text{abs}(\text{initial_estimate})) * \text{abs}(\text{initial_estimate})$

These parameter scaling coefficients are chosen to try to keep similar slopes among parameters. That is they all follow the slopes approximately on a log-scale.

While these are chosen in a logical manner, they may not always apply. You can specify each parameters scaling factor by this parameter if you wish.

`scaleTo`

Scale the initial parameter estimate to this value. By default this is 1. When zero or below, no scaling is performed.

rxControl	'rxode2' ODE solving options during fitting, created with 'rxControl()'
optExpression	Optimize the rxode2 expression to speed up calculation. By default this is turned on.
sumProd	Is a boolean indicating if the model should change multiplication to high precision multiplication and sums to high precision sums using the PreciseSums package. By default this is FALSE.
literalFix	boolean, substitute fixed population values as literals and re-adjust ui and parameter estimates after optimization; Default is 'TRUE'.
addProp	specifies the type of additive plus proportional errors, the one where standard deviations add (combined1) or the type where the variances add (combined2). The combined1 error type can be described by the following equation:

$$y = f + (a + b \times f^c) \times \varepsilon$$

The combined2 error model can be described by the following equation:

$$y = f + \sqrt{a^2 + b^2 \times f^{2 \times c}} \times \varepsilon$$

Where:

- y represents the observed value
- f represents the predicted value
- a is the additive standard deviation
- b is the proportional/power standard deviation
- c is the power exponent (in the proportional case c=1)

calcTables	This boolean is to determine if the foceiFit will calculate tables. By default this is TRUE
compress	Should the object have compressed items
covMethod	Method for calculating covariance. In this discussion, R is the Hessian matrix of the objective function. The S matrix is the sum of individual gradient cross-product (evaluated at the individual empirical Bayes estimates). <ul style="list-style-type: none"> • "r, s" Uses the sandwich matrix to calculate the covariance, that is: solve(R) %*% S %*% solve(R) • "r" Uses the Hessian matrix to calculate the covariance as 2 %*% solve(R) • "s" Uses the cross-product matrix to calculate the covariance as 4 %*% solve(S) • "" Does not calculate the covariance step.
adjObf	is a boolean to indicate if the objective function should be adjusted to be closer to NONMEM's default objective function. By default this is TRUE
ci	Confidence level for some tables. By default this is 0.95 or 95% confidence.
sigdig	Optimization significant digits. This controls: <ul style="list-style-type: none"> • The tolerance of the inner and outer optimization is $10^{-\text{sigdig}}$ • The tolerance of the ODE solvers is $0.5 \times 10^{-(\text{sigdig}-2)}$; For the sensitivity equations and steady-state solutions the default is $0.5 \times 10^{-(\text{sigdig}-1.5)}$ (sensitivity changes only applicable for liblsoda)

- The tolerance of the boundary check is $5 * 10^{(-sigdig + 1)}$
- sigdigTable Significant digits in the final output table. If not specified, then it matches the significant digits in the 'sigdig' optimization algorithm. If 'sigdig' is NULL, use 3.
- ... Additional arguments passed to `nlmixr2est::bobyqaControl()`.

Value

bobqya control structure

Author(s)

Matthew L. Fidler

Examples

```
# A logit regression example with emax model

dsn <- data.frame(i=1:1000)
dsn$time <- exp(rnorm(1000))
dsn$DV=rbinom(1000,1,exp(-1+dsn$time)/(1+exp(-1+dsn$time)))

mod <- function() {
  ini({
    E0 <- 0.5
    Em <- 0.5
    E50 <- 2
    g <- fix(2)
  })
  model({
    v <- E0+Em*time^g/(E50^g+time^g)
    ll(bin) ~ DV * v - log(1 + exp(v))
  })
}

fit2 <- nlmixr(mod, dsn, est="bobyqa")

print(fit2)

# you can also get the nlm output with

fit2$bobyqa

# The nlm control has been modified slightly to include
# extra components and name the parameters
```

bootplot	<i>Produce delta objective function for bootstrap</i>
----------	---

Description

Produce delta objective function for bootstrap

Usage

```
bootplot(x, ...)
```

Arguments

x	fit object
...	Additional arguments passed to <code>nlmixr2extra::bootplot()</code> .

Value

Fit traceplot or nothing.

Author(s)

Vipul Mann, Matthew L. Fidler

References

R Niebecker, MO Karlsson. (2013) *Are datasets for NLME models large enough for a bootstrap to provide reliable parameter uncertainty distributions?* PAGE 2013. <https://www.page-meeting.org/?abstract=2899>

bootstrapFit	<i>Bootstrap nlmixr2 fit</i>
--------------	------------------------------

Description

Bootstrap input dataset and rerun the model to get confidence bounds and aggregated parameters

Usage

```
bootstrapFit(
  fit,
  nboot = 200,
  nSampIndiv,
  stratVar,
  stdErrType = c("perc", "sd", "se"),
  ci = 0.95,
  pvalues = NULL,
  restart = FALSE,
  plotHist = FALSE,
  fitName = as.character(substitute(fit))
)
```

Arguments

<code>fit</code>	the <code>nlmixr2</code> fit object
<code>nboot</code>	an integer giving the number of bootstrapped models to be fit; default value is 200
<code>nSampIndiv</code>	an integer specifying the number of samples in each bootstrapped sample; default is the number of unique subjects in the original dataset
<code>stratVar</code>	Variable in the original dataset to stratify on; This is useful to distinguish between sparse and full sampling and other features you may wish to keep distinct in your bootstrap
<code>stdErrType</code>	This gives the standard error type for the updated standard errors; The current possibilities are: "perc" which gives the standard errors by percentiles (default), "sd" which gives the standard errors by the using the normal approximation of the mean with standard deviation, or "se" which uses the normal approximation with standard errors calculated with <code>nSampIndiv</code>
<code>ci</code>	Confidence interval level to calculate. Default is 0.95 for a 95 percent confidence interval
<code>pvalues</code>	a vector of pvalues indicating the probability of each subject to get selected; default value is NULL implying that probability of each subject is the same
<code>restart</code>	A boolean to try to restart an interrupted or incomplete bootstrap. By default this is FALSE
<code>plotHist</code>	A boolean indicating if a histogram plot to assess how well the bootstrap is doing. By default this is turned off (FALSE)
<code>fitName</code>	is the fit name that is used for the name of the bootstrap files. By default it is the fit provided though it could be something else.

Value

Nothing, called for the side effects; The original fit is updated with the bootstrap confidence bands

Author(s)

Vipul Mann, Matthew Fidler

Examples

```

## Not run:
one.cmt <- function() {
  ini({
    tka <- 0.45; label("Ka")
    tcl <- 1; label("Cl")
    tv <- 3.45; label("V")
    eta.ka ~ 0.6
    eta.cl ~ 0.3
    eta.v ~ 0.1
    add.sd <- 0.7
  })
  model({
    ka <- exp(tka + eta.ka)
    cl <- exp(tcl + eta.cl)
    v <- exp(tv + eta.v)
    linCmt() ~ add(add.sd)
  })
}

fit <- nlmixr2(one.cmt, nlmixr2data::theo_sd, est = "saem", control = list(print = 0))

withr::with_tempdir({ # Run example in temp dir

  bootstrapFit(fit, nboot = 5, restart = TRUE) # overwrites any of the existing data or model files
  bootstrapFit(fit, nboot = 7) # resumes fitting using the stored data and model files

  # Note this resumes because the total number of bootstrap samples is not 10

  bootstrapFit(fit, nboot=10)

  # Note the bootstrap standard error and variance/covariance matrix is retained.
  # If you wish to switch back you can change the covariance matrix by

  nlmixr2est::setCov(fit, "linFim")

  # And change it back again

  nlmixr2est::setCov(fit, "boot10")

  # This change will affect any simulations with uncertainty in their parameters

  # You may also do a chi-square diagnostic plot check for the bootstrap with
  bootplot(fit)
})

## End(Not run)

```


Description

Control Options for FOCEi

Usage

```
foceiControl(
  sigdig = 3,
  ...,
  epsilon = NULL,
  maxInnerIterations = 1000,
  maxOuterIterations = 5000,
  n1qn1nsim = NULL,
  print = 1L,
  printNcol = floor((getOption("width") - 23)/12),
  scaleTo = 1,
  scaleObjective = 0,
  normType = c("rescale2", "mean", "rescale", "std", "len", "constant"),
  scaleType = c("nlmixr2", "norm", "mult", "multAdd"),
  scaleCmax = 1e+05,
  scaleCmin = 1e-05,
  scaleC = NULL,
  scaleC0 = 1e+05,
  derivEps = rep(20 * sqrt(.Machine$double.eps), 2),
  derivMethod = c("switch", "forward", "central"),
  derivSwitchTol = NULL,
  covDerivMethod = c("central", "forward"),
  covMethod = c("r,s", "r", "s", ""),
  hessEps = (.Machine$double.eps)^(1/3),
  hessEpsLlik = (.Machine$double.eps)^(1/3),
  optimHessType = c("central", "forward"),
  optimHessCovType = c("central", "forward"),
  eventType = c("central", "forward"),
  centralDerivEps = rep(20 * sqrt(.Machine$double.eps), 2),
  lbfgsLmm = 7L,
  lbfgsPgtol = 0,
  lbfgsFactr = NULL,
  eigen = TRUE,
  addPosthoc = TRUE,
  diagXform = c("sqrt", "log", "identity"),
  sumProd = FALSE,
  optExpression = TRUE,
  literalFix = TRUE,
  ci = 0.95,
  useColor = crayon::has_color(),
  boundTol = NULL,
  calcTables = TRUE,
  noAbort = TRUE,
  interaction = TRUE,
```

```
cholSEtol = (.Machine$double.eps)^(1/3),
cholAccept = 0.001,
resetEtaP = 0.15,
resetThetaP = 0.05,
resetThetaFinalP = 0.15,
diagOmegaBoundUpper = 5,
diagOmegaBoundLower = 100,
cholSEOpt = FALSE,
cholSECov = FALSE,
fo = FALSE,
covTryHarder = FALSE,
outerOpt = c("nlsminb", "bobyqa", "lbfgsb3c", "L-BFGS-B", "mma", "lbfgsbLG", "slsqp",
  "Rvmin"),
innerOpt = c("n1qn1", "BFGS"),
rhobeg = 0.2,
rhoend = NULL,
npt = NULL,
rel.tol = NULL,
x.tol = NULL,
eval.max = 4000,
iter.max = 2000,
abstol = NULL,
reltol = NULL,
resetHessianAndEta = FALSE,
stateTrim = Inf,
shi21maxOuter = 0L,
shi21maxInner = 20L,
shi21maxInnerCov = 20L,
shi21maxFD = 20L,
gillK = 10L,
gillStep = 4,
gillFtol = 0,
gillRtol = sqrt(.Machine$double.eps),
gillKcov = 10L,
gillKcovLlik = 10L,
gillStepCovLlik = 4.5,
gillStepCov = 2,
gillFtolCov = 0,
gillFtolCovLlik = 0,
rmatNorm = TRUE,
rmatNormLlik = TRUE,
smatNorm = TRUE,
smatNormLlik = TRUE,
covGillF = TRUE,
optGillF = TRUE,
covSmall = 1e-05,
adjLlik = TRUE,
gradTrim = Inf,
```

```

maxOdeRecalc = 5,
odeRecalcFactor = 10^(0.5),
gradCalcCentralSmall = 1e-04,
gradCalcCentralLarge = 10000,
etaNudge = qnorm(1 - 0.05/2)/sqrt(3),
etaNudge2 = qnorm(1 - 0.05/2) * sqrt(3/5),
nRetries = 3,
seed = 42,
resetThetaCheckPer = 0.1,
etaMat = NULL,
repeatGillMax = 1,
stickyRecalcN = 4,
gradProgressOfvTime = 10,
addProp = c("combined2", "combined1"),
badSolveObjfAdj = 100,
compress = TRUE,
rxControl = NULL,
sigdigTable = NULL,
fallbackFD = FALSE,
smatPer = 0.6,
sdLowerFact = 0.001,
zeroGradFirstReset = TRUE,
zeroGradRunReset = TRUE,
zeroGradBobyqa = TRUE
)

```

Arguments

sigdig	Optimization significant digits. This controls: <ul style="list-style-type: none"> • The tolerance of the inner and outer optimization is $10^{-\text{sigdig}}$ • The tolerance of the ODE solvers is $0.5 \times 10^{-(\text{sigdig}-2)}$; For the sensitivity equations and steady-state solutions the default is $0.5 \times 10^{-(\text{sigdig}-1.5)}$ (sensitivity changes only applicable for liblsoda) • The tolerance of the boundary check is $5 \times 10^{-(\text{sigdig} + 1)}$
...	Additional arguments passed to <code>nlmixr2est::foceiControl()</code> .
epsilon	Precision of estimate for n1qn1 optimization.
maxInnerIterations	Number of iterations for n1qn1 optimization.
maxOuterIterations	Maximum number of L-BFGS-B optimization for outer problem.
n1qn1nsim	Number of function evaluations for n1qn1 optimization.
print	Integer representing when the outer step is printed. When this is 0 or do not print the iterations. 1 is print every function evaluation (default), 5 is print every 5 evaluations.
printNcol	Number of columns to printout before wrapping parameter estimates/gradient

scaleTo	Scale the initial parameter estimate to this value. By default this is 1. When zero or below, no scaling is performed.
scaleObjective	Scale the initial objective function to this value. By default this is 0 (meaning do not scale)
normType	This is the type of parameter normalization/scaling used to get the scaled initial values for nlmixr2. These are used with scaleType of. With the exception of rescale2, these come from Feature Scaling . The rescale2 The rescaling is the same type described in the OptdesX software manual. In general, all all scaling formula can be described by:

$$= \left(\frac{v_{scaled}}{v_{unscaled} - C_1} \right) / C_2$$

Where

The other data normalization approaches follow the following formula

$$= \left(\frac{v_{scaled}}{v_{unscaled} - C_1} \right) / C_2$$

- **rescale2** This scales all parameters from (-1 to 1). The relative differences between the parameters are preserved with this approach and the constants are:

$$C_1 = (\max(\text{all unscaled values}) + \min(\text{all unscaled values})) / 2$$

$$C_2 = (\max(\text{all unscaled values}) - \min(\text{all unscaled values})) / 2$$

- **rescale** or min-max normalization. This rescales all parameters from (0 to 1). As in the rescale2 the relative differences are preserved. In this approach:

$$C_1 = \min(\text{all unscaled values})$$

$$C_2 = \max(\text{all unscaled values}) - \min(\text{all unscaled values})$$

- mean or mean normalization. This rescales to center the parameters around the mean but the parameters are from 0 to 1. In this approach:

$$C_1$$

= mean(all unscaled values)

$$C_2$$

= max(all unscaled values) - min(all unscaled values)

- std or standardization. This standardizes by the mean and standard deviation. In this approach:

$$C_1$$

= mean(all unscaled values)

$$C_2$$

= sd(all unscaled values)

- len or unit length scaling. This scales the parameters to the unit length. For this approach we use the Euclidean length, that is:

$$C_1$$

= 0

$$C_2$$

=

$$\sqrt{(v_1^2 + v_2^2 + \dots + v_n^2)}$$

- constant which does not perform data normalization. That is

$$C_1$$

= 0

$$C_2$$

= 1

scaleType

The scaling scheme for nlmixr2. The supported types are:

- nlmixr2 In this approach the scaling is performed by the following equation:

$$v_{scaled}$$

= (

$$v_{current} - v_{init}$$

)*scaleC[i] + scaleTo

The scaleTo parameter is specified by the normType, and the scales are specified by scaleC.

- `norm` This approach uses the simple scaling provided by the `normType` argument.
- `mult` This approach does not use the data normalization provided by `normType`, but rather uses multiplicative scaling to a constant provided by the `scaleTo` argument.

In this case:

$$v_{scaled} = \frac{v_{current}}{v_{init}}$$

*`scaleTo`

- `multAdd` This approach changes the scaling based on the parameter being specified. If a parameter is defined in an exponential block (ie `exp(theta)`), then it is scaled on a linearly, that is:

$$v_{scaled} = (v_{current} - v_{init}) + scaleTo$$

Otherwise the parameter is scaled multiplicatively.

$$v_{scaled} = \frac{v_{current}}{v_{init}}$$

*`scaleTo`

`scaleCmax` Maximum value of the `scaleC` to prevent overflow.

`scaleCmin` Minimum value of the `scaleC` to prevent underflow.

`scaleC` The scaling constant used with `scaleType=nlmixr2`. When not specified, it is based on the type of parameter that is estimated. The idea is to keep the derivatives similar on a log scale to have similar gradient sizes. Hence parameters like `log(exp(theta))` would have a scaling factor of 1 and `log(theta)` would have a scaling factor of `ini_value` (to scale by `1/value`; ie `d/dt(log(ini_value)) = 1/ini_value` or `scaleC=ini_value`)

- For parameters in an exponential (ie `exp(theta)`) or parameters specifying powers, `boxCox` or `yeoJohnson` transformations, this is 1.
- For additive, proportional, lognormal error structures, these are given by `0.5*abs(initial_estimate)`
- Factorials are scaled by `abs(1/digamma(initial_estimate+1))`
- parameters in a log scale (ie `log(theta)`) are transformed by `log(abs(initial_estimate))*abs(initial_estimate)`

	<p>These parameter scaling coefficients are chose to try to keep similar slopes among parameters. That is they all follow the slopes approximately on a log-scale.</p> <p>While these are chosen in a logical manner, they may not always apply. You can specify each parameters scaling factor by this parameter if you wish.</p>
scaleC0	Number to adjust the scaling factor by if the initial gradient is zero.
derivEps	<p>Forward difference tolerances, which is a vector of relative difference and absolute difference. The central/forward difference step size h is calculated as:</p> $h = \text{abs}(x) * \text{derivEps}[1] + \text{derivEps}[2]$
derivMethod	<p>indicates the method for calculating derivatives of the outer problem. Currently supports "switch", "central" and "forward" difference methods. Switch starts with forward differences. This will switch to central differences when $\text{abs}(\text{delta}(\text{OFV})) \leq \text{derivSwitchTol}$ and switch back to forward differences when $\text{abs}(\text{delta}(\text{OFV})) > \text{derivSwitchTol}$.</p>
derivSwitchTol	The tolerance to switch forward to central differences.
covDerivMethod	indicates the method for calculating the derivatives while calculating the covariance components (Hessian and S).
covMethod	<p>Method for calculating covariance. In this discussion, R is the Hessian matrix of the objective function. The S matrix is the sum of individual gradient cross-product (evaluated at the individual empirical Bayes estimates).</p> <ul style="list-style-type: none"> • "r, s" Uses the sandwich matrix to calculate the covariance, that is: $\text{solve}(R) \%*\% S \%*\% \text{solve}(R)$ • "r" Uses the Hessian matrix to calculate the covariance as $2 \%*\% \text{solve}(R)$ • "s" Uses the cross-product matrix to calculate the covariance as $4 \%*\% \text{solve}(S)$ • "" Does not calculate the covariance step.
hessEps	is a double value representing the epsilon for the Hessian calculation. This is used for the R matrix calculation.
hessEpsLlik	is a double value representing the epsilon for the Hessian calculation when doing focei generalized log-likelihood estimation. This is used for the R matrix calculation.
optimHessType	The hessian type for when calculating the individual hessian by numeric differences (in generalized log-likelihood estimation). The options are "central", and "forward". The central differences is what R's 'optimHess()' uses and is the default for this method. (Though the "forward" is faster and still reasonable for most cases). The Shi21 cannot be changed for the Gill83 algorithm with the optimHess in a generalized likelihood problem.
optimHessCovType	The hessian type for when calculating the individual hessian by numeric differences (in generalized log-likelihood estimation). The options are "central", and "forward". The central differences is what R's 'optimHess()' uses. While this takes longer in optimization, it is more accurate, so for calculating the covariance and final likelihood, the central differences are used. This also uses the modified Shi21 method
eventType	Event gradient type for dosing events; Can be "central" or "forward"

centralDerivEps	Central difference tolerances. This is a numeric vector of relative difference and absolute difference. The central/forward difference step size h is calculated as: $h = \text{abs}(x) * \text{derivEps}[1] + \text{derivEps}[2]$
lbfgsLmm	An integer giving the number of BFGS updates retained in the "L-BFGS-B" method, It defaults to 7.
lbfgsPgtol	is a double precision variable. On entry $\text{pgtol} \geq 0$ is specified by the user. The iteration will stop when: $\max(\ \text{proj } g_i \ \mid i = 1, \dots, n) \leq \text{lbfgsPgtol}$ where pg_i is the i th component of the projected gradient. On exit pgtol is unchanged. This defaults to zero, when the check is suppressed.
lbfgsFactr	Controls the convergence of the "L-BFGS-B" method. Convergence occurs when the reduction in the objective is within this factor of the machine tolerance. Default is $1e10$, which gives a tolerance of about $2e-6$, approximately 4 sigdigs. You can check your exact tolerance by multiplying this value by <code>.Machine\$double.eps</code>
eigen	A boolean indicating if eigenvectors are calculated to include a condition number calculation.
addPosthoc	Boolean indicating if posthoc parameters are added to the table output.
diagXform	This is the transformation used on the diagonal of the <code>chol(solve(omega))</code> . This matrix and values are the parameters estimated in FOCEi. The possibilities are: <ul style="list-style-type: none"> • <code>sqrt</code> Estimates the sqrt of the diagonal elements of <code>chol(solve(omega))</code>. This is the default method. • <code>log</code> Estimates the log of the diagonal elements of <code>chol(solve(omega))</code> • <code>identity</code> Estimates the diagonal elements without any transformations
sumProd	Is a boolean indicating if the model should change multiplication to high precision multiplication and sums to high precision sums using the <code>PreciseSums</code> package. By default this is <code>FALSE</code> .
optExpression	Optimize the <code>rxode2</code> expression to speed up calculation. By default this is turned on.
literalFix	boolean, substitute fixed population values as literals and re-adjust <code>ui</code> and parameter estimates after optimization; Default is <code>'TRUE'</code> .
ci	Confidence level for some tables. By default this is 0.95 or 95% confidence.
useColor	Boolean indicating if focei can use ASCII color codes
boundTol	Tolerance for boundary issues.
calcTables	This boolean is to determine if the foceiFit will calculate tables. By default this is <code>TRUE</code>
noAbort	Boolean to indicate if you should abort the FOCEi evaluation if it runs into troubles. (default <code>TRUE</code>)
interaction	Boolean indicate FOCEi should be used (<code>TRUE</code>) instead of FOCE (<code>FALSE</code>)
cholSEtol	tolerance for Generalized Cholesky Decomposition. Defaults to suggested $(\text{.Machine\$double.eps})^{(1/3)}$

cholAccept	Tolerance to accept a Generalized Cholesky Decomposition for a R or S matrix.
resetEtaP	represents the p-value for resetting the individual ETA to 0 during optimization (instead of the saved value). The two test statistics used in the z-test are either $\text{chol}(\omega^{-1})$ $\%*\% \text{ eta}$ or $\text{eta}/\text{sd}(\text{allEtas})$. A p-value of 0 indicates the ETAs never reset. A p-value of 1 indicates the ETAs always reset.
resetThetaP	represents the p-value for resetting the population mu-referenced THETA parameters based on ETA drift during optimization, and resetting the optimization. A p-value of 0 indicates the THETAs never reset. A p-value of 1 indicates the THETAs always reset and is not allowed. The theta reset is checked at the beginning and when nearing a local minima. The percent change in objective function where a theta reset check is initiated is controlled in resetThetaCheckPer.
resetThetaFinalP	represents the p-value for resetting the population mu-referenced THETA parameters based on ETA drift during optimization, and resetting the optimization one final time.
diagOmegaBoundUpper	This represents the upper bound of the diagonal omega matrix. The upper bound is given by $\text{diag}(\omega) * \text{diagOmegaBoundUpper}$. If diagOmegaBoundUpper is 1, there is no upper bound on Omega.
diagOmegaBoundLower	This represents the lower bound of the diagonal omega matrix. The lower bound is given by $\text{diag}(\omega) / \text{diagOmegaBoundUpper}$. If diagOmegaBoundLower is 1, there is no lower bound on Omega.
cholSEOpt	Boolean indicating if the generalized Cholesky should be used while optimizing.
cholSECov	Boolean indicating if the generalized Cholesky should be used while calculating the Covariance Matrix.
fo	is a boolean indicating if this is a FO approximation routine.
covTryHarder	If the R matrix is non-positive definite and cannot be corrected to be non-positive definite try estimating the Hessian on the unscaled parameter space.
outerOpt	optimization method for the outer problem
innerOpt	optimization method for the inner problem (not implemented yet.)
rhobeg	Beginning change in parameters for bobyqa algorithm (trust region). By default this is 0.2 or 20 parameters when the parameters are scaled to 1. rhobeg and rhoend must be set to the initial and final values of a trust region radius, so both must be positive with $0 < \text{rhoend} < \text{rhobeg}$. Typically rhobeg should be about one tenth of the greatest expected change to a variable. Note also that smallest difference $\text{abs}(\text{upper}-\text{lower})$ should be greater than or equal to $\text{rhobeg} * 2$. If this is not the case then rhobeg will be adjusted. (bobyqa)
rhoend	The smallest value of the trust region radius that is allowed. If not defined, then $10^{-(\text{sigdig}-1)}$ will be used. (bobyqa)
npt	The number of points used to approximate the objective function via a quadratic approximation for bobyqa. The value of npt must be in the interval $[\text{n}+2, (\text{n}+1)(\text{n}+2)/2]$ where n is the number of parameters in par. Choices that exceed $2 * \text{n} + 1$ are not recommended. If not defined, it will be set to $2 * \text{n} + 1$. (bobyqa)
rel.tol	Relative tolerance before nlmimb stops (nlmimb).

x.tol	X tolerance for nlmixr2 optimizer
eval.max	Number of maximum evaluations of the objective function (nlmimb)
iter.max	Maximum number of iterations allowed (nlmimb)
abstol	Absolute tolerance for nlmixr2 optimizer (BFGS)
reltol	tolerance for nlmixr2 (BFGS)
resetHessianAndEta	is a boolean representing if the individual Hessian is reset when ETAs are reset using the option resetEtaP.
stateTrim	Trim state amounts/concentrations to this value.
shi21maxOuter	The maximum number of steps for the optimization of the forward-difference step size. When not zero, use this instead of Gill differences.
shi21maxInner	The maximum number of steps for the optimization of the individual Hessian matrices in the generalized likelihood problem. When 0, un-optimized finite differences are used.
shi21maxInnerCov	The maximum number of steps for the optimization of the individual Hessian matrices in the generalized likelihood problem for the covariance step. When 0, un-optimized finite differences are used.
shi21maxFD	The maximum number of steps for the optimization of the forward difference step size when using dosing events (lag time, modeled duration/rate and bioavailability)
gillK	The total number of possible steps to determine the optimal forward/central difference step size per parameter (by the Gill 1983 method). If 0, no optimal step size is determined. Otherwise this is the optimal step size determined.
gillStep	When looking for the optimal forward difference step size, this is This is the step size to increase the initial estimate by. So each iteration the new step size = (prior step size)*gillStep
gillFtol	The gillFtol is the gradient error tolerance that is acceptable before issuing a warning/error about the gradient estimates.
gillRtol	The relative tolerance used for Gill 1983 determination of optimal step size.
gillKcov	The total number of possible steps to determine the optimal forward/central difference step size per parameter (by the Gill 1983 method) during the covariance step. If 0, no optimal step size is determined. Otherwise this is the optimal step size determined.
gillKcovLlik	The total number of possible steps to determine the optimal forward/central difference step per parameter when using the generalized focei log-likelihood method (by the Gill 1986 method). If 0, no optimal step size is determined. Otherwise this is the optimal step size is determined
gillStepCovLlik	Same as above but during generalized focei log-likelihood
gillStepCov	When looking for the optimal forward difference step size, this is This is the step size to increase the initial estimate by. So each iteration during the covariance step is equal to the new step size = (prior step size)*gillStepCov

gillFtolCov	The gillFtol is the gradient error tolerance that is acceptable before issuing a warning/error about the gradient estimates during the covariance step.
gillFtolCovLlik	Same as above but applied during generalized log-likelihood estimation.
rmatNorm	A parameter to normalize gradient step size by the parameter value during the calculation of the R matrix
rmatNormLlik	A parameter to normalize gradient step size by the parameter value during the calculation of the R matrix if you are using generalized log-likelihood Hessian matrix.
smatNorm	A parameter to normalize gradient step size by the parameter value during the calculation of the S matrix
smatNormLlik	A parameter to normalize gradient step size by the parameter value during the calculation of the S matrix if you are using the generalized log-likelihood.
covGillF	Use the Gill calculated optimal Forward difference step size for the instead of the central difference step size during the central difference gradient calculation.
optGillF	Use the Gill calculated optimal Forward difference step size for the instead of the central difference step size during the central differences for optimization.
covSmall	The covSmall is the small number to compare covariance numbers before rejecting an estimate of the covariance as the final estimate (when comparing sandwich vs R/S matrix estimates of the covariance). This number controls how small the variance is before the covariance matrix is rejected.
adjLlik	In nlmixr2, the objective function matches NONMEM's objective function, which removes a 2π constant from the likelihood calculation. If this is TRUE, the likelihood function is adjusted by this 2π factor. When adjusted this number more closely matches the likelihood approximations of nlme, and SAS approximations. Regardless of if this is turned on or off the objective function matches NONMEM's objective function.
gradTrim	The parameter to adjust the gradient to if the $ \text{grad} $ is very large.
maxOdeRecalc	Maximum number of times to reduce the ODE tolerances and try to resolve the system if there was a bad ODE solve.
odeRecalcFactor	The ODE recalculation factor when ODE solving goes bad, this is the factor the rtol/atol is reduced
gradCalcCentralSmall	A small number that represents the value where $ \text{grad} < \text{gradCalcCentralSmall}$ where forward differences switch to central differences.
gradCalcCentralLarge	A large number that represents the value where $ \text{grad} > \text{gradCalcCentralLarge}$ where forward differences switch to central differences.
etaNudge	By default initial ETA estimates start at zero; Sometimes this doesn't optimize appropriately. If this value is non-zero, when the nlqn1 optimization didn't perform appropriately, reset the Hessian, and nudge the ETA up by this value; If the ETA still doesn't move, nudge the ETA down by this value. By default this value is $\text{qnorm}(1-0.05/2)*1/\text{sqrt}(3)$, the first of the Gauss Quadrature numbers times by the 0.95% normal region. If this is not successful try the second eta

	nudge number (below). If +-etaNudge2 is not successful, then assign to zero and do not optimize any longer
etaNudge2	This is the second eta nudge. By default it is $qnorm(1-0.05/2)*sqrt(3/5)$, which is the n=3 quadrature point (excluding zero) times by the 0.95% normal region
nRetries	If FOCEi doesn't fit with the current parameter estimates, randomly sample new parameter estimates and restart the problem. This is similar to 'PsN' resampling.
seed	an object specifying if and how the random number generator should be initialized
resetThetaCheckPer	represents objective function % percentage below which resetThetaP is checked.
etaMat	Eta matrix for initial estimates or final estimates of the ETAs.
repeatGillMax	If the tolerances were reduced when calculating the initial Gill differences, the Gill difference is repeated up to a maximum number of times defined by this parameter.
stickyRecalcN	The number of bad ODE solves before reducing the atol/rtol for the rest of the problem.
gradProgressOfvTime	This is the time for a single objective function evaluation (in seconds) to start progress bars on gradient evaluations
addProp	specifies the type of additive plus proportional errors, the one where standard deviations add (combined1) or the type where the variances add (combined2). The combined1 error type can be described by the following equation: $y = f + (a + b \times f^c) \times \varepsilon$ <p>The combined2 error model can be described by the following equation:</p> $y = f + \sqrt{a^2 + b^2 \times f^{2 \times c}} \times \varepsilon$ <p>Where:</p> <ul style="list-style-type: none"> - y represents the observed value - f represents the predicted value - a is the additive standard deviation - b is the proportional/power standard deviation - c is the power exponent (in the proportional case c=1)
badSolveObjfAdj	The objective function adjustment when the ODE system cannot be solved. It is based on each individual bad solve.
compress	Should the object have compressed items
rxControl	'rxode2' ODE solving options during fitting, created with 'rxControl()'
sigdigTable	Significant digits in the final output table. If not specified, then it matches the significant digits in the 'sigdig' optimization algorithm. If 'sigdig' is NULL, use 3.
fallbackFD	Fallback to the finite differences if the sensitivity equations do not solve.

smatPer	A percentage representing the number of failed parameter gradients for each individual (which are replaced with the overall gradient for the parameter) out of the total number of gradients parameters (ie 'ntheta*nsub') before the S matrix is considered to be a bad matrix.
sdLowerFact	A factor for multiplying the estimate by when the lower estimate is zero and the error is known to represent a standard deviation of a parameter (like add.sd, prop.sd, pow.sd, lnorm.sd, etc). When zero, no factor is applied. If your initial estimate is 0.15 and your lower bound is zero, then the lower bound would be assumed to be 0.00015.
zeroGradFirstReset	boolean, when 'TRUE' if the first gradient is zero, reset the zero gradient to 'sqrt(.Machine\$double.eps)' to get past the bad initial estimate, otherwise error (and possibly reset), when 'FALSE' error when the first gradient is zero. When 'NA' on the last reset, have the zero gradient ignored, otherwise error and look for another value. Default is 'TRUE'
zeroGradRunReset	boolean, when 'TRUE' if a gradient is zero, reset the zero gradient to 'sqrt(.Machine\$double.eps)' to get past the bad estimate while running. Otherwise error (and possibly reset). Default is 'TRUE'
zeroGradBobyqa	boolean, when 'TRUE' if a gradient is zero, the reset will change the method to the gradient free bobyqa method. When 'NA', the zero gradient will change to bobyqa only when the first gradient is zero. Default is 'TRUE'

Details

Note this uses the R's L-BFGS-B in `optim` for the outer problem and the BFGS `n1qn1` with that allows restoring the prior individual Hessian (for faster optimization speed).

However the inner problem is not scaled. Since most eta estimates start near zero, scaling for these parameters do not make sense.

This process of scaling can fix some ill conditioning for the unscaled problem. The covariance step is performed on the unscaled problem, so the condition number of that matrix may not be reflective of the scaled problem's condition-number.

Value

The control object that changes the options for the FOCEi family of estimation methods

Author(s)

Matthew L. Fidler

References

Gill, P.E., Murray, W., Saunders, M.A., & Wright, M.H. (1983). Computing Forward-Difference Intervals for Numerical Optimization. *Siam Journal on Scientific and Statistical Computing*, 4, 310-321.

Shi, H.M., Xie, Y., Xuan, M.Q., & Nocedal, J. (2021). Adaptive Finite-Difference Interval Estimation for Noisy Derivative-Free Optimization.

See Also[optim](#)[n1qn1](#)[rxSolve](#)Other Estimation control: [nlmixr2NlmeControl\(\)](#), [saemControl\(\)](#)

`lbfgsb3cControl`*Control for lbfgsb3c estimation method in nlmixr2*

Description

Control for lbfgsb3c estimation method in nlmixr2

Usage

```

lbfgsb3cControl(
  trace = 0,
  factr = 1e+07,
  pgtol = 0,
  abstol = 0,
  reltol = 0,
  lmm = 5L,
  maxit = 10000L,
  returnLbfgsb3c = FALSE,
  stickyRecalcN = 4,
  maxOdeRecalc = 5,
  odeRecalcFactor = 10^(0.5),
  useColor = crayon::has_color(),
  printNcol = floor((getOption("width") - 23)/12),
  print = 1L,
  normType = c("rescale2", "mean", "rescale", "std", "len", "constant"),
  scaleType = c("nlmixr2", "norm", "mult", "multAdd"),
  scaleCmax = 1e+05,
  scaleCmin = 1e-05,
  scaleC = NULL,
  scaleTo = 1,
  gradTo = 1,
  rxControl = NULL,
  optExpression = TRUE,
  sumProd = FALSE,
  literalFix = TRUE,
  addProp = c("combined2", "combined1"),
  calcTables = TRUE,
  compress = TRUE,
  covMethod = c("r", ""),

```

```

    adjObf = TRUE,
    ci = 0.95,
    sigdig = 4,
    sigdigTable = NULL,
    ...
)

```

Arguments

trace	If positive, tracing information on the progress of the optimization is produced. Higher values may produce more tracing information: for method "L-BFGS-B" there are six levels of tracing. (To understand exactly what these do see the source code: higher levels give more detail.)
factr	controls the convergence of the "L-BFGS-B" method. Convergence occurs when the reduction in the objective is within this factor of the machine tolerance. Default is 1e7, that is a tolerance of about 1e-8.
pgtol	helps control the convergence of the "L-BFGS-B" method. It is a tolerance on the projected gradient in the current search direction. This defaults to zero, when the check is suppressed.
abstol	helps control the convergence of the "L-BFGS-B" method. It is an absolute tolerance difference in x values. This defaults to zero, when the check is suppressed.
reltol	helps control the convergence of the "L-BFGS-B" method. It is a relative tolerance difference in x values. This defaults to zero, when the check is suppressed.
lmm	is an integer giving the number of BFGS updates retained in the "L-BFGS-B" method, It defaults to 5.
maxit	maximum number of iterations.
returnLbfgsb3c	return the lbfgsb3c output instead of the nlmixr2 fit
stickyRecalcN	The number of bad ODE solves before reducing the atol/rtol for the rest of the problem.
maxOdeRecalc	Maximum number of times to reduce the ODE tolerances and try to resolve the system if there was a bad ODE solve.
odeRecalcFactor	The ODE recalculation factor when ODE solving goes bad, this is the factor the rtol/atol is reduced
useColor	Boolean indicating if focei can use ASCII color codes
printNcol	Number of columns to printout before wrapping parameter estimates/gradient
print	Integer representing when the outer step is printed. When this is 0 or do not print the iterations. 1 is print every function evaluation (default), 5 is print every 5 evaluations.
normType	This is the type of parameter normalization/scaling used to get the scaled initial values for nlmixr2. These are used with scaleType of. With the exception of rescale2, these come from Feature Scaling . The rescale2 The rescaling is the same type described in the OptdesX software manual.

In general, all all scaling formula can be described by:

$$v_{scaled} = (v_{unscaled} - C_1) / C_2$$

Where

The other data normalization approaches follow the following formula

$$v_{scaled} = (v_{unscaled} - C_1) / C_2$$

- **rescale2** This scales all parameters from (-1 to 1). The relative differences between the parameters are preserved with this approach and the constants are:

$$C_1 = (\max(\text{all unscaled values}) + \min(\text{all unscaled values})) / 2$$

$$C_2 = (\max(\text{all unscaled values}) - \min(\text{all unscaled values})) / 2$$

- **rescale** or **min-max normalization**. This rescales all parameters from (0 to 1). As in the **rescale2** the relative differences are preserved. In this approach:

$$C_1 = \min(\text{all unscaled values})$$

$$C_2 = \max(\text{all unscaled values}) - \min(\text{all unscaled values})$$

- **mean** or **mean normalization**. This rescales to center the parameters around the mean but the parameters are from 0 to 1. In this approach:

$$C_1 = \text{mean}(\text{all unscaled values})$$

$$C_2 = \max(\text{all unscaled values}) - \min(\text{all unscaled values})$$

- `std` or standardization. This standardizes by the mean and standard deviation. In this approach:

$$C_1$$

= mean(all unscaled values)

$$C_2$$

= sd(all unscaled values)

- `len` or unit length scaling. This scales the parameters to the unit length. For this approach we use the Euclidean length, that is:

$$C_1$$

= 0

$$C_2$$

=

$$\sqrt{(v_1^2 + v_2^2 + \dots + v_n^2)}$$

- `constant` which does not perform data normalization. That is

$$C_1$$

= 0

$$C_2$$

= 1

`scaleType` The scaling scheme for `nlmixr2`. The supported types are:

- `nlmixr2` In this approach the scaling is performed by the following equation:

$$v_{scaled} = (v_{current} - v_{init}) * scaleC[i] + scaleTo$$

The `scaleTo` parameter is specified by the `normType`, and the scales are specified by `scaleC`.

- `norm` This approach uses the simple scaling provided by the `normType` argument.
- `mult` This approach does not use the data normalization provided by `normType`, but rather uses multiplicative scaling to a constant provided by the `scaleTo` argument.

In this case:

$$v_{scaled}$$

$$= \frac{v_{current}}{v_{init}} *scaleTo$$

- `multAdd` This approach changes the scaling based on the parameter being specified. If a parameter is defined in an exponential block (ie $\exp(\theta)$), then it is scaled on a linearly, that is:

$$= \left(\frac{v_{scaled}}{v_{current} - v_{init}} \right) + scaleTo$$

Otherwise the parameter is scaled multiplicatively.

$$= \frac{v_{scaled}}{v_{current}} / v_{init} *scaleTo$$

`scaleCmax` Maximum value of the `scaleC` to prevent overflow.

`scaleCmin` Minimum value of the `scaleC` to prevent underflow.

`scaleC` The scaling constant used with `scaleType=nlmixr2`. When not specified, it is based on the type of parameter that is estimated. The idea is to keep the derivatives similar on a log scale to have similar gradient sizes. Hence parameters like $\log(\exp(\theta))$ would have a scaling factor of 1 and $\log(\theta)$ would have a scaling factor of `ini_value` (to scale by $1/\text{value}$; ie $d/dt(\log(\text{ini_value})) = 1/\text{ini_value}$ or `scaleC=ini_value`)

- For parameters in an exponential (ie $\exp(\theta)$) or parameters specifying powers, `boxCox` or `yeoJohnson` transformations, this is 1.
- For additive, proportional, lognormal error structures, these are given by $0.5 * \text{abs}(\text{initial_estimate})$
- Factorials are scaled by $\text{abs}(1/\text{digamma}(\text{initial_estimate}+1))$
- parameters in a log scale (ie $\log(\theta)$) are transformed by $\log(\text{abs}(\text{initial_estimate})) * \text{abs}(\text{initial_estimate})$

These parameter scaling coefficients are chosen to try to keep similar slopes among parameters. That is they all follow the slopes approximately on a log-scale.

While these are chosen in a logical manner, they may not always apply. You can specify each parameters scaling factor by this parameter if you wish.

`scaleTo` Scale the initial parameter estimate to this value. By default this is 1. When zero or below, no scaling is performed.

gradTo	this is the factor that the gradient is scaled to before optimizing. This only works with scaleType="nlmixr2".
rxControl	'rxode2' ODE solving options during fitting, created with 'rxControl()'
optExpression	Optimize the rxode2 expression to speed up calculation. By default this is turned on.
sumProd	Is a boolean indicating if the model should change multiplication to high precision multiplication and sums to high precision sums using the PreciseSums package. By default this is FALSE.
literalFix	boolean, substitute fixed population values as literals and re-adjust ui and parameter estimates after optimization; Default is 'TRUE'.
addProp	specifies the type of additive plus proportional errors, the one where standard deviations add (combined1) or the type where the variances add (combined2). The combined1 error type can be described by the following equation:

$$y = f + (a + b \times f^c) \times \varepsilon$$

The combined2 error model can be described by the following equation:

$$y = f + \sqrt{a^2 + b^2 \times f^{2 \times c}} \times \varepsilon$$

Where:

- y represents the observed value
- f represents the predicted value
- a is the additive standard deviation
- b is the proportional/power standard deviation
- c is the power exponent (in the proportional case c=1)

calcTables	This boolean is to determine if the foceiFit will calculate tables. By default this is TRUE
compress	Should the object have compressed items
covMethod	Method for calculating covariance. In this discussion, R is the Hessian matrix of the objective function. The S matrix is the sum of individual gradient cross-product (evaluated at the individual empirical Bayes estimates). <ul style="list-style-type: none"> • "r, s" Uses the sandwich matrix to calculate the covariance, that is: solve(R) %*% S %*% solve(R) • "r" Uses the Hessian matrix to calculate the covariance as 2 %*% solve(R) • "s" Uses the cross-product matrix to calculate the covariance as 4 %*% solve(S) • "" Does not calculate the covariance step.
adjObf	is a boolean to indicate if the objective function should be adjusted to be closer to NONMEM's default objective function. By default this is TRUE
ci	Confidence level for some tables. By default this is 0.95 or 95% confidence.
sigdig	Optimization significant digits. This controls: <ul style="list-style-type: none"> • The tolerance of the inner and outer optimization is $10^{-\text{sigdig}}$

- The tolerance of the ODE solvers is $0.5 \times 10^{-(\text{sigdig}-2)}$; For the sensitivity equations and steady-state solutions the default is $0.5 \times 10^{-(\text{sigdig}-1.5)}$ (sensitivity changes only applicable for liblsoda)
 - The tolerance of the boundary check is $5 \times 10^{-(\text{sigdig}+1)}$
- sigdigTable Significant digits in the final output table. If not specified, then it matches the significant digits in the 'sigdig' optimization algorithm. If 'sigdig' is NULL, use 3.
- ... Additional arguments passed to `nlmixr2est::lbfgsb3cControl()`.

Value

bobqya control structure

Author(s)

Matthew L. Fidler

Examples

```
# A logit regression example with emax model

dsn <- data.frame(i=1:1000)
dsn$time <- exp(rnorm(1000))
dsn$DV=rbinom(1000,1,exp(-1+dsn$time)/(1+exp(-1+dsn$time)))

mod <- function() {
  ini({
    E0 <- 0.5
    Em <- 0.5
    E50 <- 2
    g <- fix(2)
  })
  model({
    v <- E0+Em*time^g/(E50^g+time^g)
    ll(bin) ~ DV * v - log(1 + exp(v))
  })
}

fit2 <- nlmixr(mod, dsn, est="lbfgsb3c")

print(fit2)

# you can also get the nlm output with fit2$lbfgsb3c

fit2$lbfgsb3c

# The nlm control has been modified slightly to include
# extra components and name the parameters
```

n1qn1Control	<i>Control for n1qn1 estimation method in nlmixr2</i>
--------------	---

Description

Control for n1qn1 estimation method in nlmixr2

Usage

```
n1qn1Control(
  epsilon = (.Machine$double.eps)^0.25,
  max_iterations = 10000,
  nsim = 10000,
  imp = 0,
  print.functions = FALSE,
  returnN1qn1 = FALSE,
  stickyRecalcN = 4,
  maxOdeRecalc = 5,
  odeRecalcFactor = 10^(0.5),
  useColor = crayon::has_color(),
  printNcol = floor((getOption("width") - 23)/12),
  print = 1L,
  normType = c("rescale2", "mean", "rescale", "std", "len", "constant"),
  scaleType = c("nlmixr2", "norm", "mult", "multAdd"),
  scaleCmax = 1e+05,
  scaleCmin = 1e-05,
  scaleC = NULL,
  scaleTo = 1,
  gradTo = 1,
  rxControl = NULL,
  optExpression = TRUE,
  sumProd = FALSE,
  literalFix = TRUE,
  addProp = c("combined2", "combined1"),
  calcTables = TRUE,
  compress = TRUE,
  covMethod = c("r", "n1qn1", ""),
  adjObf = TRUE,
  ci = 0.95,
  sigdig = 4,
  sigdigTable = NULL,
  ...
)
```

Arguments

epsilon Precision of estimate for n1qn1 optimization.

max_iterations	Number of iterations
nsim	Number of function evaluations
imp	Verbosity of messages.
print.functions	Boolean to control if the function value and parameter estimates are echoed every time a function is called.
returnN1qn1	return the nlqn1 output instead of the nlmixr2 fit
stickyRecalcN	The number of bad ODE solves before reducing the atol/rtol for the rest of the problem.
maxOdeRecalc	Maximum number of times to reduce the ODE tolerances and try to resolve the system if there was a bad ODE solve.
odeRecalcFactor	The ODE recalculation factor when ODE solving goes bad, this is the factor the rtol/atol is reduced
useColor	Boolean indicating if focei can use ASCII color codes
printNcol	Number of columns to printout before wrapping parameter estimates/gradient
print	Integer representing when the outer step is printed. When this is 0 or do not print the iterations. 1 is print every function evaluation (default), 5 is print every 5 evaluations.
normType	This is the type of parameter normalization/scaling used to get the scaled initial values for nlmixr2. These are used with scaleType of. With the exception of rescale2, these come from Feature Scaling . The rescale2 The rescaling is the same type described in the OptdesX software manual. In general, all all scaling formula can be described by:

$$v_{scaled} = \left(\frac{v_{unscaled} - C_1}{C_2} \right)$$

Where

The other data normalization approaches follow the following formula

$$v_{scaled} = \left(\frac{v_{unscaled} - C_1}{C_2} \right)$$

- `rescale2` This scales all parameters from (-1 to 1). The relative differences between the parameters are preserved with this approach and the constants are:

$$C_1$$

$$= (\max(\text{all unscaled values}) + \min(\text{all unscaled values})) / 2$$

$$C_2$$

$$= (\max(\text{all unscaled values}) - \min(\text{all unscaled values})) / 2$$

- `rescale` or min-max normalization. This rescales all parameters from (0 to 1). As in the `rescale2` the relative differences are preserved. In this approach:

$$C_1$$

$$= \min(\text{all unscaled values})$$

$$C_2$$

$$= \max(\text{all unscaled values}) - \min(\text{all unscaled values})$$

- `mean` or mean normalization. This rescales to center the parameters around the mean but the parameters are from 0 to 1. In this approach:

$$C_1$$

$$= \text{mean}(\text{all unscaled values})$$

$$C_2$$

$$= \max(\text{all unscaled values}) - \min(\text{all unscaled values})$$

- `std` or standardization. This standardizes by the mean and standard deviation. In this approach:

$$C_1$$

$$= \text{mean}(\text{all unscaled values})$$

$$C_2$$

$$= \text{sd}(\text{all unscaled values})$$

- `len` or unit length scaling. This scales the parameters to the unit length. For this approach we use the Euclidean length, that is:

$$C_1$$

$$= 0$$

$$C_2$$

$$= \sqrt{v_1^2 + v_2^2 + \dots + v_n^2}$$

- constant which does not perform data normalization. That is

$$= 0 \quad C_1$$

$$= 1 \quad C_2$$

scaleType

The scaling scheme for nlmixr2. The supported types are:

- nlmixr2 In this approach the scaling is performed by the following equation:

$$= (\quad v_{scaled} \\ \quad v_{current} - v_{init} \\) * scaleC[i] + scaleTo$$

The scaleTo parameter is specified by the normType, and the scales are specified by scaleC.

- norm This approach uses the simple scaling provided by the normType argument.
- mult This approach does not use the data normalization provided by normType, but rather uses multiplicative scaling to a constant provided by the scaleTo argument.

In this case:

$$= \quad v_{scaled} \\ \quad v_{current} \\ / \\ \quad v_{init}$$

*scaleTo

- multAdd This approach changes the scaling based on the parameter being specified. If a parameter is defined in an exponential block (ie exp(theta)), then it is scaled on a linearly, that is:

$$= (\quad v_{scaled} \\ \quad v_{current} - v_{init} \\) + scaleTo$$

Otherwise the parameter is scaled multiplicatively.

$$= \quad v_{scaled} \\ \quad v_{current}$$

	/	v_{init}
		*scaleTo
scaleCmax	Maximum value of the scaleC to prevent overflow.	
scaleCmin	Minimum value of the scaleC to prevent underflow.	
scaleC	<p>The scaling constant used with scaleType=nlmixr2. When not specified, it is based on the type of parameter that is estimated. The idea is to keep the derivatives similar on a log scale to have similar gradient sizes. Hence parameters like log(exp(theta)) would have a scaling factor of 1 and log(theta) would have a scaling factor of ini_value (to scale by 1/value; ie d/dt(log(ini_value)) = 1/ini_value or scaleC=ini_value)</p> <ul style="list-style-type: none"> • For parameters in an exponential (ie exp(theta)) or parameters specifying powers, boxCox or yeoJohnson transformations , this is 1. • For additive, proportional, lognormal error structures, these are given by 0.5*abs(initial_estimate) • Factorials are scaled by abs(1/digamma(initial_estimate+1)) • parameters in a log scale (ie log(theta)) are transformed by log(abs(initial_estimate))*abs(initial_estim <p>These parameter scaling coefficients are chose to try to keep similar slopes among parameters. That is they all follow the slopes approximately on a log-scale.</p> <p>While these are chosen in a logical manner, they may not always apply. You can specify each parameters scaling factor by this parameter if you wish.</p>	
scaleTo	Scale the initial parameter estimate to this value. By default this is 1. When zero or below, no scaling is performed.	
gradTo	this is the factor that the gradient is scaled to before optimizing. This only works with scaleType="nlmixr2".	
rxControl	'rxode2' ODE solving options during fitting, created with 'rxControl()'	
optExpression	Optimize the rxode2 expression to speed up calculation. By default this is turned on.	
sumProd	Is a boolean indicating if the model should change multiplication to high precision multiplication and sums to high precision sums using the PreciseSums package. By default this is FALSE.	
literalFix	boolean, substitute fixed population values as literals and re-adjust ui and parameter estimates after optimization; Default is 'TRUE'.	
addProp	<p>specifies the type of additive plus proportional errors, the one where standard deviations add (combined1) or the type where the variances add (combined2). The combined1 error type can be described by the following equation:</p>	

$$y = f + (a + b \times f^c) \times \varepsilon$$

The combined2 error model can be described by the following equation:

$$y = f + \sqrt{a^2 + b^2 \times f^{2 \times c}} \times \varepsilon$$

	Where:
	- y represents the observed value
	- f represents the predicted value
	- a is the additive standard deviation
	- b is the proportional/power standard deviation
	- c is the power exponent (in the proportional case c=1)
calcTables	This boolean is to determine if the foceiFit will calculate tables. By default this is TRUE
compress	Should the object have compressed items
covMethod	Method for calculating covariance. In this discussion, R is the Hessian matrix of the objective function. The S matrix is the sum of individual gradient cross-product (evaluated at the individual empirical Bayes estimates). <ul style="list-style-type: none"> • "r, s" Uses the sandwich matrix to calculate the covariance, that is: <code>solve(R) %*% S %*% solve(R)</code> • "r" Uses the Hessian matrix to calculate the covariance as <code>2 %*% solve(R)</code> • "s" Uses the cross-product matrix to calculate the covariance as <code>4 %*% solve(S)</code> • "" Does not calculate the covariance step.
adj0bf	is a boolean to indicate if the objective function should be adjusted to be closer to NONMEM's default objective function. By default this is TRUE
ci	Confidence level for some tables. By default this is 0.95 or 95% confidence.
sigdig	Optimization significant digits. This controls: <ul style="list-style-type: none"> • The tolerance of the inner and outer optimization is $10^{-\text{sigdig}}$ • The tolerance of the ODE solvers is $0.5 \times 10^{-(\text{sigdig}-2)}$; For the sensitivity equations and steady-state solutions the default is $0.5 \times 10^{-(\text{sigdig}-1.5)}$ (sensitivity changes only applicable for liblsoda) • The tolerance of the boundary check is $5 \times 10^{-(\text{sigdig} + 1)}$
sigdigTable	Significant digits in the final output table. If not specified, then it matches the significant digits in the 'sigdig' optimization algorithm. If 'sigdig' is NULL, use 3.
...	Additional arguments passed to <code>nlmixr2est::n1qn1Control()</code> .

Value

bobqya control structure

Author(s)

Matthew L. Fidler

Examples

```
# A logit regression example with emax model

dsn <- data.frame(i=1:1000)
dsn$time <- exp(rnorm(1000))
```

```

dsn$DV=rbinom(1000,1,exp(-1+dsn$time)/(1+exp(-1+dsn$time)))

mod <- function() {
  ini({
    E0 <- 0.5
    Em <- 0.5
    E50 <- 2
    g <- fix(2)
  })
  model({
    v <- E0+Em*time^g/(E50^g+time^g)
    ll(bin) ~ DV * v - log(1 + exp(v))
  })
}

fit2 <- nlmixr(mod, dsn, est="n1qn1")

print(fit2)

# you can also get the nlm output with fit2$n1qn1

fit2$n1qn1

# The nlm control has been modified slightly to include
# extra components and name the parameters

```

newuoaControl

Control for newuoa estimation method in nlmixr2

Description

Control for newuoa estimation method in nlmixr2

Usage

```

newuoaControl(
  npt = NULL,
  rhobeg = NULL,
  rhoend = NULL,
  iprint = 0L,
  maxfun = 100000L,
  returnNewuoa = FALSE,
  stickyRecalcN = 4,
  maxOdeRecalc = 5,
  odeRecalcFactor = 10^(0.5),
  useColor = crayon::has_color(),
  printNcol = floor((getOption("width") - 23)/12),
  print = 1L,

```

```

normType = c("rescale2", "mean", "rescale", "std", "len", "constant"),
scaleType = c("nlmixr2", "norm", "mult", "multAdd"),
scaleCmax = 1e+05,
scaleCmin = 1e-05,
scaleC = NULL,
scaleTo = 1,
rxControl = NULL,
optExpression = TRUE,
sumProd = FALSE,
literalFix = TRUE,
addProp = c("combined2", "combined1"),
calcTables = TRUE,
compress = TRUE,
covMethod = c("r", ""),
adjObf = TRUE,
ci = 0.95,
sigdig = 4,
sigdigTable = NULL,
...
)

```

Arguments

npt	The number of points used to approximate the objective function via a quadratic approximation for bobyqa. The value of npt must be in the interval $[n+2, (n+1)(n+2)/2]$ where n is the number of parameters in par. Choices that exceed $2*n+1$ are not recommended. If not defined, it will be set to $2*n + 1$. (bobyqa)
rhobeg	Beginning change in parameters for bobyqa algorithm (trust region). By default this is 0.2 or 20 parameters when the parameters are scaled to 1. rhobeg and rhoend must be set to the initial and final values of a trust region radius, so both must be positive with $0 < \text{rhoend} < \text{rhobeg}$. Typically rhobeg should be about one tenth of the greatest expected change to a variable. Note also that smallest difference $\text{abs}(\text{upper}-\text{lower})$ should be greater than or equal to rhobeg^2 . If this is not the case then rhobeg will be adjusted. (bobyqa)
rhoend	The smallest value of the trust region radius that is allowed. If not defined, then $10^{-(\text{sigdig}-1)}$ will be used. (bobyqa)
iprint	The value of 'iprint' should be set to an integer value in '0, 1, 2, 3, ...', which controls the amount of printing. Specifically, there is no output if 'iprint=0' and there is output only at the start and the return if 'iprint=1'. Otherwise, each new value of 'rho' is printed, with the best vector of variables so far and the corresponding value of the objective function. Further, each new value of the objective function with its variables are output if 'iprint=3'. If 'iprint > 3', the objective function value and corresponding variables are output every 'iprint' evaluations. Default value is '0'.
maxfun	The maximum allowed number of function evaluations. If this is exceeded, the method will terminate.
returnNewuoa	return the newuoa output instead of the nlmixr2 fit

stickyRecalcN	The number of bad ODE solves before reducing the atol/rtol for the rest of the problem.
maxOdeRecalc	Maximum number of times to reduce the ODE tolerances and try to resolve the system if there was a bad ODE solve.
odeRecalcFactor	The ODE recalculation factor when ODE solving goes bad, this is the factor the rtol/atol is reduced
useColor	Boolean indicating if focei can use ASCII color codes
printNcol	Number of columns to printout before wrapping parameter estimates/gradient
print	Integer representing when the outer step is printed. When this is 0 or do not print the iterations. 1 is print every function evaluation (default), 5 is print every 5 evaluations.
normType	This is the type of parameter normalization/scaling used to get the scaled initial values for nlmixr2. These are used with scaleType of. With the exception of rescale2, these come from Feature Scaling . The rescale2 The rescaling is the same type described in the OptdesX software manual. In general, all all scaling formula can be described by:

$$= \left(\frac{v_{scaled}}{v_{unscaled} - C_1} \right) / C_2$$

Where

The other data normalization approaches follow the following formula

$$= \left(\frac{v_{scaled}}{v_{unscaled} - C_1} \right) / C_2$$

- rescale2 This scales all parameters from (-1 to 1). The relative differences between the parameters are preserved with this approach and the constants are:

$$= \frac{C_1}{(\max(\text{all unscaled values}) + \min(\text{all unscaled values})) / 2}$$

$$= \frac{C_2}{(\max(\text{all unscaled values}) - \min(\text{all unscaled values})) / 2}$$

- `rescale` or min-max normalization. This rescales all parameters from (0 to 1). As in the `rescale2` the relative differences are preserved. In this approach:

$$C_1$$

= min(all unscaled values)

$$C_2$$

= max(all unscaled values) - min(all unscaled values)

- `mean` or mean normalization. This rescales to center the parameters around the mean but the parameters are from 0 to 1. In this approach:

$$C_1$$

= mean(all unscaled values)

$$C_2$$

= max(all unscaled values) - min(all unscaled values)

- `std` or standardization. This standardizes by the mean and standard deviation. In this approach:

$$C_1$$

= mean(all unscaled values)

$$C_2$$

= sd(all unscaled values)

- `len` or unit length scaling. This scales the parameters to the unit length. For this approach we use the Euclidean length, that is:

$$C_1$$

= 0

$$C_2$$

=

$$\sqrt{(v_1^2 + v_2^2 + \dots + v_n^2)}$$

- `constant` which does not perform data normalization. That is

$$C_1$$

= 0

$$C_2$$

= 1

`scaleType`

The scaling scheme for `nlmixr2`. The supported types are:

- `nlmixr2` In this approach the scaling is performed by the following equation:

$$v_{scaled} = (v_{current} - v_{init}) * scaleC[i] + scaleTo$$

The `scaleTo` parameter is specified by the `normType`, and the scales are specified by `scaleC`.

- `norm` This approach uses the simple scaling provided by the `normType` argument.
- `mult` This approach does not use the data normalization provided by `normType`, but rather uses multiplicative scaling to a constant provided by the `scaleTo` argument.

In this case:

$$v_{scaled} = v_{current} / v_{init} * scaleTo$$

- `multAdd` This approach changes the scaling based on the parameter being specified. If a parameter is defined in an exponential block (ie `exp(theta)`), then it is scaled on a linearly, that is:

$$v_{scaled} = (v_{current} - v_{init}) + scaleTo$$

Otherwise the parameter is scaled multiplicatively.

$$v_{scaled} = v_{current} / v_{init} * scaleTo$$

`scaleCmax` Maximum value of the `scaleC` to prevent overflow.
`scaleCmin` Minimum value of the `scaleC` to prevent underflow.

scaleC	<p>The scaling constant used with scaleType=nlmixr2. When not specified, it is based on the type of parameter that is estimated. The idea is to keep the derivatives similar on a log scale to have similar gradient sizes. Hence parameters like $\log(\exp(\theta))$ would have a scaling factor of 1 and $\log(\theta)$ would have a scaling factor of ini_value (to scale by 1/value; ie $d/dt(\log(\text{ini_value})) = 1/\text{ini_value}$ or $\text{scaleC}=\text{ini_value}$)</p> <ul style="list-style-type: none"> • For parameters in an exponential (ie $\exp(\theta)$) or parameters specifying powers, boxCox or yeoJohnson transformations , this is 1. • For additive, proportional, lognormal error structures, these are given by $0.5*\text{abs}(\text{initial_estimate})$ • Factorials are scaled by $\text{abs}(1/\text{digamma}(\text{initial_estimate}+1))$ • parameters in a log scale (ie $\log(\theta)$) are transformed by $\log(\text{abs}(\text{initial_estimate}))*\text{abs}(\text{initial_estimate})$ <p>These parameter scaling coefficients are chose to try to keep similar slopes among parameters. That is they all follow the slopes approximately on a log-scale.</p> <p>While these are chosen in a logical manner, they may not always apply. You can specify each parameters scaling factor by this parameter if you wish.</p>
scaleTo	Scale the initial parameter estimate to this value. By default this is 1. When zero or below, no scaling is performed.
rxControl	'rxode2' ODE solving options during fitting, created with 'rxControl()'
optExpression	Optimize the rxode2 expression to speed up calculation. By default this is turned on.
sumProd	Is a boolean indicating if the model should change multiplication to high precision multiplication and sums to high precision sums using the PreciseSums package. By default this is FALSE.
literalFix	boolean, substitute fixed population values as literals and re-adjust ui and parameter estimates after optimization; Default is 'TRUE'.
addProp	<p>specifies the type of additive plus proportional errors, the one where standard deviations add (combined1) or the type where the variances add (combined2). The combined1 error type can be described by the following equation:</p> $y = f + (a + b \times f^c) \times \varepsilon$ <p>The combined2 error model can be described by the following equation:</p> $y = f + \sqrt{a^2 + b^2 \times f^{2 \times c}} \times \varepsilon$
	<p>Where:</p> <ul style="list-style-type: none"> - y represents the observed value - f represents the predicted value - a is the additive standard deviation - b is the proportional/power standard deviation - c is the power exponent (in the proportional case c=1)
calcTables	This boolean is to determine if the fociFit will calculate tables. By default this is TRUE

compress	Should the object have compressed items
covMethod	Method for calculating covariance. In this discussion, R is the Hessian matrix of the objective function. The S matrix is the sum of individual gradient cross-product (evaluated at the individual empirical Bayes estimates). <ul style="list-style-type: none"> • "r, s" Uses the sandwich matrix to calculate the covariance, that is: <code>solve(R) %*% S %*% solve(R)</code> • "r" Uses the Hessian matrix to calculate the covariance as <code>2 %*% solve(R)</code> • "s" Uses the cross-product matrix to calculate the covariance as <code>4 %*% solve(S)</code> • "" Does not calculate the covariance step.
adj0bf	is a boolean to indicate if the objective function should be adjusted to be closer to NONMEM's default objective function. By default this is TRUE
ci	Confidence level for some tables. By default this is 0.95 or 95% confidence.
sigdig	Optimization significant digits. This controls: <ul style="list-style-type: none"> • The tolerance of the inner and outer optimization is $10^{-\text{sigdig}}$ • The tolerance of the ODE solvers is $0.5 \times 10^{-(\text{sigdig}-2)}$; For the sensitivity equations and steady-state solutions the default is $0.5 \times 10^{-(\text{sigdig}-1.5)}$ (sensitivity changes only applicable for liblsoda) • The tolerance of the boundary check is $5 \times 10^{-(\text{sigdig} + 1)}$
sigdigTable	Significant digits in the final output table. If not specified, then it matches the significant digits in the 'sigdig' optimization algorithm. If 'sigdig' is NULL, use 3.
...	Additional arguments passed to <code>nlmixr2est::newuoaControl()</code> .

Value

newuoa control structure

Author(s)

Matthew L. Fidler

Examples

```
# A logit regression example with emax model

dsn <- data.frame(i=1:1000)
dsn$time <- exp(rnorm(1000))
dsn$DV=rbinom(1000,1,exp(-1+dsn$time)/(1+exp(-1+dsn$time)))

mod <- function() {
  ini({
    E0 <- 0.5
    Em <- 0.5
    E50 <- 2
    g <- fix(2)
  })
  model({
```

```

      v <- E0+Em*time^g/(E50^g+time^g)
      ll(bin) ~ DV * v - log(1 + exp(v))
    })
  }

fit2 <- nlmixr(mod, dsn, est="newuoa")

print(fit2)

# you can also get the nlm output with

fit2$newuoa

# The nlm control has been modified slightly to include
# extra components and name the parameters

```

nlmControl

nlmixr2 defaults controls for nlm

Description

nlmixr2 defaults controls for nlm

Usage

```

nlmControl(
  tysize = NULL,
  fscale = 1,
  print.level = 0,
  ndigit = NULL,
  gradtol = 1e-06,
  stepmax = NULL,
  steptol = 1e-06,
  iterlim = 10000,
  check.analyticals = FALSE,
  returnNlm = FALSE,
  solveType = c("hessian", "grad", "fun"),
  stickyRecalcN = 4,
  maxOdeRecalc = 5,
  odeRecalcFactor = 10^(0.5),
  eventType = c("central", "forward"),
  shiErr = (.Machine$double.eps)^(1/3),
  shi21maxFD = 20L,
  optimHessType = c("central", "forward"),
  hessErr = (.Machine$double.eps)^(1/3),
  shi21maxHess = 20L,
  useColor = crayon::has_color(),

```

```

printNcol = floor((getOption("width") - 23)/12),
print = 1L,
normType = c("rescale2", "mean", "rescale", "std", "len", "constant"),
scaleType = c("nlmixr2", "norm", "mult", "multAdd"),
scaleCmax = 1e+05,
scaleCmin = 1e-05,
scaleC = NULL,
scaleTo = 1,
gradTo = 1,
rxControl = NULL,
optExpression = TRUE,
sumProd = FALSE,
literalFix = TRUE,
addProp = c("combined2", "combined1"),
calcTables = TRUE,
compress = TRUE,
covMethod = c("r", "nlm", ""),
adjObf = TRUE,
ci = 0.95,
sigdig = 4,
sigdigTable = NULL,
...
)

```

Arguments

<code>typsize</code>	an estimate of the size of each parameter at the minimum.
<code>fscale</code>	an estimate of the size of f at the minimum.
<code>print.level</code>	this argument determines the level of printing which is done during the minimization process. The default value of 0 means that no printing occurs, a value of 1 means that initial and final details are printed and a value of 2 means that full tracing information is printed.
<code>ndigit</code>	the number of significant digits in the function f .
<code>gradtol</code>	a positive scalar giving the tolerance at which the scaled gradient is considered close enough to zero to terminate the algorithm. The scaled gradient is a measure of the relative change in f in each direction $p[i]$ divided by the relative change in $p[i]$.
<code>stepmax</code>	a positive scalar which gives the maximum allowable scaled step length. <code>stepmax</code> is used to prevent steps which would cause the optimization function to overflow, to prevent the algorithm from leaving the area of interest in parameter space, or to detect divergence in the algorithm. <code>stepmax</code> would be chosen small enough to prevent the first two of these occurrences, but should be larger than any anticipated reasonable step.
<code>steptol</code>	A positive scalar providing the minimum allowable relative step length.
<code>iterlim</code>	a positive integer specifying the maximum number of iterations to be performed before the program is terminated.

check.analyticals	a logical scalar specifying whether the analytic gradients and Hessians, if they are supplied, should be checked against numerical derivatives at the initial parameter values. This can help detect incorrectly formulated gradients or Hessians.
returnNlm	is a logical that allows a return of the 'nlm' object
solveType	tells if 'nlm' will use nlmixr2's analytical gradients when available (finite differences will be used for event-related parameters like parameters controlling lag time, duration/rate of infusion, and modeled bioavailability). This can be: <ul style="list-style-type: none"> - "hessian" which will use the analytical gradients to create a Hessian with finite differences. - "gradient" which will use the gradient and let 'nlm' calculate the finite difference hessian - "fun" where nlm will calculate both the finite difference gradient and the finite difference Hessian When using nlmixr2's finite differences, the "ideal" step size for either central or forward differences are optimized for with the Shi2021 method which may give more accurate derivatives
stickyRecalcN	The number of bad ODE solves before reducing the atol/rtol for the rest of the problem.
maxOdeRecalc	Maximum number of times to reduce the ODE tolerances and try to resolve the system if there was a bad ODE solve.
odeRecalcFactor	The ODE recalculation factor when ODE solving goes bad, this is the factor the rtol/atol is reduced
eventType	Event gradient type for dosing events; Can be "central" or "forward"
shiErr	This represents the epsilon when optimizing the ideal step size for numeric differentiation using the Shi2021 method
shi21maxFD	The maximum number of steps for the optimization of the forward difference step size when using dosing events (lag time, modeled duration/rate and bioavailability)
optimHessType	The hessian type for when calculating the individual hessian by numeric differences (in generalized log-likelihood estimation). The options are "central", and "forward". The central differences is what R's 'optimHess()' uses and is the default for this method. (Though the "forward" is faster and still reasonable for most cases). The Shi21 cannot be changed for the Gill83 algorithm with the optimHess in a generalized likelihood problem.
hessErr	This represents the epsilon when optimizing the Hessian step size using the Shi2021 method.
shi21maxHess	Maximum number of times to optimize the best step size for the hessian calculation
useColor	Boolean indicating if focei can use ASCII color codes
printNcol	Number of columns to printout before wrapping parameter estimates/gradient

print	Integer representing when the outer step is printed. When this is 0 or do not print the iterations. 1 is print every function evaluation (default), 5 is print every 5 evaluations.
normType	This is the type of parameter normalization/scaling used to get the scaled initial values for nlmixr2. These are used with scaleType of. With the exception of rescale2, these come from Feature Scaling . The rescale2 The rescaling is the same type described in the OptdesX software manual. In general, all all scaling formula can be described by:

$$v_{scaled} = (v_{unscaled} - C_1) / C_2$$

Where

The other data normalization approaches follow the following formula

$$v_{scaled} = (v_{unscaled} - C_1) / C_2$$

- rescale2 This scales all parameters from (-1 to 1). The relative differences between the parameters are preserved with this approach and the constants are:

$$C_1 = (\max(\text{all unscaled values}) + \min(\text{all unscaled values})) / 2$$

$$C_2 = (\max(\text{all unscaled values}) - \min(\text{all unscaled values})) / 2$$

- rescale or min-max normalization. This rescales all parameters from (0 to 1). As in the rescale2 the relative differences are preserved. In this approach:

$$C_1 = \min(\text{all unscaled values})$$

$$C_2 = \max(\text{all unscaled values}) - \min(\text{all unscaled values})$$

- mean or mean normalization. This rescales to center the parameters around the mean but the parameters are from 0 to 1. In this approach:

$$C_1$$

= mean(all unscaled values)

$$C_2$$

= max(all unscaled values) - min(all unscaled values)

- std or standardization. This standardizes by the mean and standard deviation. In this approach:

$$C_1$$

= mean(all unscaled values)

$$C_2$$

= sd(all unscaled values)

- len or unit length scaling. This scales the parameters to the unit length. For this approach we use the Euclidean length, that is:

$$C_1$$

= 0

$$C_2$$

=

$$\sqrt{(v_1^2 + v_2^2 + \dots + v_n^2)}$$

- constant which does not perform data normalization. That is

$$C_1$$

= 0

$$C_2$$

= 1

scaleType

The scaling scheme for nlmixr2. The supported types are:

- nlmixr2 In this approach the scaling is performed by the following equation:

$$v_{scaled}$$

= (

$$v_{current} - v_{init}$$

)*scaleC[i] + scaleTo

The scaleTo parameter is specified by the normType, and the scales are specified by scaleC.

- `norm` This approach uses the simple scaling provided by the `normType` argument.
- `mult` This approach does not use the data normalization provided by `normType`, but rather uses multiplicative scaling to a constant provided by the `scaleTo` argument.

In this case:

$$v_{scaled} = \frac{v_{current}}{v_{init}}$$

*`scaleTo`

- `multAdd` This approach changes the scaling based on the parameter being specified. If a parameter is defined in an exponential block (ie `exp(theta)`), then it is scaled on a linearly, that is:

$$v_{scaled} = (v_{current} - v_{init}) + scaleTo$$

Otherwise the parameter is scaled multiplicatively.

$$v_{scaled} = \frac{v_{current}}{v_{init}}$$

*`scaleTo`

`scaleCmax`

Maximum value of the `scaleC` to prevent overflow.

`scaleCmin`

Minimum value of the `scaleC` to prevent underflow.

`scaleC`

The scaling constant used with `scaleType=nlmixr2`. When not specified, it is based on the type of parameter that is estimated. The idea is to keep the derivatives similar on a log scale to have similar gradient sizes. Hence parameters like `log(exp(theta))` would have a scaling factor of 1 and `log(theta)` would have a scaling factor of `ini_value` (to scale by `1/value`; ie `d/dt(log(ini_value)) = 1/ini_value` or `scaleC=ini_value`)

- For parameters in an exponential (ie `exp(theta)`) or parameters specifying powers, `boxCox` or `yeoJohnson` transformations, this is 1.
- For additive, proportional, lognormal error structures, these are given by `0.5*abs(initial_estimate)`
- Factorials are scaled by `abs(1/digamma(initial_estimate+1))`
- parameters in a log scale (ie `log(theta)`) are transformed by `log(abs(initial_estimate))*abs(initial_estim`

These parameter scaling coefficients are chosen to try to keep similar slopes among parameters. That is they all follow the slopes approximately on a log-scale.

While these are chosen in a logical manner, they may not always apply. You can specify each parameter's scaling factor by this parameter if you wish.

scaleTo	Scale the initial parameter estimate to this value. By default this is 1. When zero or below, no scaling is performed.
gradTo	this is the factor that the gradient is scaled to before optimizing. This only works with scaleType="nlmixr2".
rxControl	'rxode2' ODE solving options during fitting, created with 'rxControl()'
optExpression	Optimize the rxode2 expression to speed up calculation. By default this is turned on.
sumProd	Is a boolean indicating if the model should change multiplication to high precision multiplication and sums to high precision sums using the PreciseSums package. By default this is FALSE.
literalFix	boolean, substitute fixed population values as literals and re-adjust ui and parameter estimates after optimization; Default is 'TRUE'.
addProp	specifies the type of additive plus proportional errors, the one where standard deviations add (combined1) or the type where the variances add (combined2). The combined1 error type can be described by the following equation:

$$y = f + (a + b \times f^c) \times \varepsilon$$

The combined2 error model can be described by the following equation:

$$y = f + \sqrt{a^2 + b^2 \times f^{2 \times c}} \times \varepsilon$$

Where:

- y represents the observed value
- f represents the predicted value
- a is the additive standard deviation
- b is the proportional/power standard deviation
- c is the power exponent (in the proportional case c=1)

calcTables	This boolean is to determine if the fociFit will calculate tables. By default this is TRUE
compress	Should the object have compressed items
covMethod	allows selection of "r", which uses nlmixr2's 'nlmixr2Hess()' for the hessian calculation or "nlm" which uses the hessian from 'stats::nlm(..., hessian=TRUE)'. When using 'nlmixr2's' hessian for optimization or 'nlmixr2's' gradient for solving this defaults to "nlm" since 'stats::optimHess()' assumes an accurate gradient and is faster than 'nlmixr2Hess'
adjObf	is a boolean to indicate if the objective function should be adjusted to be closer to NONMEM's default objective function. By default this is TRUE
ci	Confidence level for some tables. By default this is 0.95 or 95% confidence.

sigdig	Optimization significant digits. This controls: <ul style="list-style-type: none"> • The tolerance of the inner and outer optimization is $10^{-\text{sigdig}}$ • The tolerance of the ODE solvers is $0.5 \times 10^{-(\text{sigdig}-2)}$; For the sensitivity equations and steady-state solutions the default is $0.5 \times 10^{-(\text{sigdig}-1.5)}$ (sensitivity changes only applicable for liblsoda) • The tolerance of the boundary check is $5 \times 10^{-(\text{sigdig} + 1)}$
sigdigTable	Significant digits in the final output table. If not specified, then it matches the significant digits in the 'sigdig' optimization algorithm. If 'sigdig' is NULL, use 3.
...	Additional arguments passed to <code>nlmixr2est::nlmControl()</code> .

Value

nlm control object

Author(s)

Matthew L. Fidler

Examples

```
# A logit regression example with emax model

dsn <- data.frame(i=1:1000)
dsn$time <- exp(rnorm(1000))
dsn$DV=rbinom(1000,1,exp(-1+dsn$time)/(1+exp(-1+dsn$time)))

mod <- function() {
  ini({
    E0 <- 0.5
    Em <- 0.5
    E50 <- 2
    g <- fix(2)
  })
  model({
    v <- E0+Em*time^g/(E50^g+time^g)
    ll(bin) ~ DV * v - log(1 + exp(v))
  })
}

fit2 <- nlmixr(mod, dsn, est="nlm")

print(fit2)

# you can also get the nlm output with fit2$nlm

fit2$nlm

# The nlm control has been modified slightly to include
# extra components and name the parameters
```

nlmeControl

Control Values for nlme Fit with extra options for nlmixr

Description

The values supplied in the function call replace the defaults and a list with all possible arguments is returned. The returned list is used as the 'control' argument to the 'nlme' function.

Usage

```
nlmeControl(
  maxIter = 100,
  pnlsMaxIter = 100,
  msMaxIter = 100,
  minScale = 0.001,
  tolerance = 1e-05,
  niterEM = 25,
  pnlsTol = 0.001,
  msTol = 1e-06,
  returnObject = FALSE,
  msVerbose = FALSE,
  msWarnNoConv = TRUE,
  gradHess = TRUE,
  apVar = TRUE,
  .relStep = .Machine$double.eps^(1/3),
  minAbsParApVar = 0.05,
  opt = c("nlminb", "nlm"),
  natural = TRUE,
  sigma = NULL,
  optExpression = TRUE,
  literalFix = TRUE,
  sumProd = FALSE,
  rxControl = NULL,
  method = c("ML", "REML"),
  random = NULL,
  fixed = NULL,
  weights = NULL,
  verbose = TRUE,
  returnNlme = FALSE,
  addProp = c("combined2", "combined1"),
  calcTables = TRUE,
  compress = TRUE,
  adjObf = TRUE,
  ci = 0.95,
  sigdig = 4,
  sigdigTable = NULL,
  muRefCovAlg = TRUE,
```

```
    ...
  )
```

Arguments

maxIter	maximum number of iterations for the nlme optimization algorithm. Default is 50.
pnlsMaxIter	maximum number of iterations for the PNLs optimization step inside the nlme optimization. Default is 7.
msMaxIter	maximum number of iterations for <code>nlminb</code> (<code>iter.max</code>) or the <code>nlm</code> (<code>iterlim</code> , from the 10-th step) optimization step inside the nlme optimization. Default is 50 (which may be too small for e.g. for overparametrized cases).
minScale	minimum factor by which to shrink the default step size in an attempt to decrease the sum of squares in the PNLs step. Default <code>0.001</code> .
tolerance	tolerance for the convergence criterion in the nlme algorithm. Default is <code>1e-6</code> .
niterEM	number of iterations for the EM algorithm used to refine the initial estimates of the random effects variance-covariance coefficients. Default is 25.
pnlsTol	tolerance for the convergence criterion in PNLs step. Default is <code>1e-3</code> .
msTol	tolerance for the convergence criterion in <code>nlm</code> , passed as the <code>gradtol</code> argument to the function (see documentation on <code>nlm</code>). Default is <code>1e-7</code> .
returnObject	a logical value indicating whether the fitted object should be returned when the maximum number of iterations is reached without convergence of the algorithm. Default is <code>FALSE</code> .
msVerbose	a logical value passed as the <code>trace</code> to <code>nlminb(..., control=list(trace=*, ...))</code> or as argument <code>print.level</code> to <code>nlm()</code> . Default is <code>FALSE</code> .
msWarnNoConv	logical indicating if a <code>warning</code> should be signalled whenever the minimization (by <code>opt</code>) in the LME step does not converge; defaults to <code>TRUE</code> .
gradHess	a logical value indicating whether numerical gradient vectors and Hessian matrices of the log-likelihood function should be used in the <code>nlm</code> optimization. This option is only available when the correlation structure (<code>corStruct</code>) and the variance function structure (<code>varFunc</code>) have no "varying" parameters and the <code>pdMat</code> classes used in the random effects structure are <code>pdSymm</code> (general positive-definite), <code>pdDiag</code> (diagonal), <code>pdIdent</code> (multiple of the identity), or <code>pdCompSymm</code> (compound symmetry). Default is <code>TRUE</code> .
apVar	a logical value indicating whether the approximate covariance matrix of the variance-covariance parameters should be calculated. Default is <code>TRUE</code> .
.relStep	relative step for numerical derivatives calculations. Default is <code>.Machine\$double.eps^(1/3)</code> .
minAbsParApVar	numeric value - minimum absolute parameter value in the approximate variance calculation. The default is <code>0.05</code> .
opt	the optimizer to be used, either <code>"nlminb"</code> (the default) or <code>"nlm"</code> .
natural	a logical value indicating whether the <code>pdNatural</code> parametrization should be used for general positive-definite matrices (<code>pdSymm</code>) in <code>reStruct</code> , when the approximate covariance matrix of the estimators is calculated. Default is <code>TRUE</code> .

sigma	optionally a positive number to fix the residual error at. If NULL, as by default, or 0, sigma is estimated.
optExpression	Optimize the rxode2 expression to speed up calculation. By default this is turned on.
literalFix	boolean, substitute fixed population values as literals and re-adjust ui and parameter estimates after optimization; Default is 'TRUE'.
sumProd	Is a boolean indicating if the model should change multiplication to high precision multiplication and sums to high precision sums using the PreciseSums package. By default this is FALSE.
rxControl	'rxode2' ODE solving options during fitting, created with 'rxControl()'
method	a character string. If "REML" the model is fit by maximizing the restricted log-likelihood. If "ML" the log-likelihood is maximized. Defaults to "ML".
random	optionally, any of the following: (i) a two-sided formula of the form $r_1 + \dots + r_n \sim x_1 + \dots + x_m$ $g_1 / \dots / g_Q$, with r_1, \dots, r_n naming parameters included on the right hand side of model, $x_1 + \dots + x_m$ specifying the random-effects model for these parameters and $g_1 / \dots / g_Q$ the grouping structure (Q may be equal to 1, in which case no / is required). The random effects formula will be repeated for all levels of grouping, in the case of multiple levels of grouping; (ii) a two-sided formula of the form $r_1 + \dots + r_n \sim x_1 + \dots + x_m$, a list of two-sided formulas of the form $r_1 \sim x_1 + \dots + x_m$, with possibly different random-effects models for different parameters, a pdMat object with a two-sided formula, or list of two-sided formulas (i.e. a non-NULL value for formula(random)), or a list of pdMat objects with two-sided formulas, or lists of two-sided formulas. In this case, the grouping structure formula will be given in groups, or derived from the data used to fit the nonlinear mixed-effects model, which should inherit from class groupedData; (iii) a named list of formulas, lists of formulas, or pdMat objects as in (ii), with the grouping factors as names. The order of nesting will be assumed the same as the order of the elements in the list; (iv) an reStruct object. See the documentation on pdClasses for a description of the available pdMat classes. Defaults to fixed, resulting in all fixed effects having also random effects.
fixed	a two-sided linear formula of the form $f_1 + \dots + f_n \sim x_1 + \dots + x_m$, or a list of two-sided formulas of the form $f_1 \sim x_1 + \dots + x_m$, with possibly different models for different parameters. The f_1, \dots, f_n are the names of parameters included on the right hand side of model and the $x_1 + \dots + x_m$ expressions define linear models for these parameters (when the left hand side of the formula contains several parameters, they all are assumed to follow the same linear model, described by the right hand side expression). A 1 on the right hand side of the formula(s) indicates a single fixed effects for the corresponding parameter(s).
weights	an optional varFunc object or one-sided formula describing the within-group heteroscedasticity structure. If given as a formula, it is used as the argument to varFixed, corresponding to fixed variance weights. See the documentation on varClasses for a description of the available varFunc classes. Defaults to NULL, corresponding to homoscedastic within-group errors.
verbose	an optional logical value. If TRUE information on the evolution of the iterative algorithm is printed. Default is FALSE.

returnNlme	Returns the nlme object instead of the nlmixr object (by default FALSE). If any of the nlme specific options of 'random', 'fixed', 'sens', the nlme object is returned
addProp	specifies the type of additive plus proportional errors, the one where standard deviations add (combined1) or the type where the variances add (combined2). The combined1 error type can be described by the following equation:

$$y = f + (a + b \times f^c) \times \varepsilon$$

The combined2 error model can be described by the following equation:

$$y = f + \sqrt{a^2 + b^2 \times f^{2 \times c}} \times \varepsilon$$

Where:

- y represents the observed value
- f represents the predicted value
- a is the additive standard deviation
- b is the proportional/power standard deviation
- c is the power exponent (in the proportional case c=1)

calcTables	This boolean is to determine if the foceiFit will calculate tables. By default this is TRUE
compress	Should the object have compressed items
adjObf	is a boolean to indicate if the objective function should be adjusted to be closer to NONMEM's default objective function. By default this is TRUE
ci	Confidence level for some tables. By default this is 0.95 or 95% confidence.
sigdig	Optimization significant digits. This controls: <ul style="list-style-type: none"> • The tolerance of the inner and outer optimization is $10^{-\text{sigdig}}$ • The tolerance of the ODE solvers is $0.5 \times 10^{-(\text{sigdig}-2)}$; For the sensitivity equations and steady-state solutions the default is $0.5 \times 10^{-(\text{sigdig}-1.5)}$ (sensitivity changes only applicable for liblsoda) • The tolerance of the boundary check is $5 \times 10^{-(\text{sigdig} + 1)}$
sigdigTable	Significant digits in the final output table. If not specified, then it matches the significant digits in the 'sigdig' optimization algorithm. If 'sigdig' is NULL, use 3.
muRefCovAlg	This controls if algebraic expressions that can be mu-referenced are treated as mu-referenced covariates by: <ol style="list-style-type: none"> 1. Creating a internal data-variable 'nlmixrMuDerCov#' for each algebraic mu-referenced expression 2. Change the algebraic expression to 'nlmixrMuDerCov# * mu_cov_theta' 3. Use the internal mu-referenced covariate for saem 4. After optimization is completed, replace 'model()' with old 'model()' expression 5. Remove 'nlmixrMuDerCov#' from nlmix2 output <p>In general, these covariates should be more accurate since it changes the system to a linear compartment model. Therefore, by default this is 'TRUE'.</p>
...	Additional arguments passed to <code>nlmixr2est::nlmeControl()</code> .

Value

a nlmixr-nlme list

See Also

Other Estimation control: [foceiControl\(\)](#), [saemControl\(\)](#)

Examples

```
nlmeControl()
nlmixr2NlmeControl()
```

nlminbControl	<i>nlmixr2 nlminb defaults</i>
---------------	--------------------------------

Description

nlmixr2 nlminb defaults

Usage

```
nlminbControl(
  eval.max = 200,
  iter.max = 150,
  trace = 0,
  abs.tol = 0,
  rel.tol = 1e-10,
  x.tol = 1.5e-08,
  xf.tol = 2.2e-14,
  step.min = 1,
  step.max = 1,
  sing.tol = rel.tol,
  scale = 1,
  scale.init = NULL,
  diff.g = NULL,
  rxControl = NULL,
  optExpression = TRUE,
  sumProd = FALSE,
  literalFix = TRUE,
  returnNlminb = FALSE,
  solveType = c("hessian", "grad", "fun"),
  stickyRecalcN = 4,
  maxOdeRecalc = 5,
  odeRecalcFactor = 10^(0.5),
  eventType = c("central", "forward"),
  shiErr = (.Machine$double.eps)^(1/3),
  shi21maxFD = 20L,
```

```

optimHessType = c("central", "forward"),
hessErr = (.Machine$double.eps)^(1/3),
shi21maxHess = 20L,
useColor = crayon::has_color(),
printNcol = floor((getOption("width") - 23)/12),
print = 1L,
normType = c("rescale2", "mean", "rescale", "std", "len", "constant"),
scaleType = c("nlmixr2", "norm", "mult", "multAdd"),
scaleCmax = 1e+05,
scaleCmin = 1e-05,
scaleC = NULL,
scaleTo = 1,
gradTo = 1,
addProp = c("combined2", "combined1"),
calcTables = TRUE,
compress = TRUE,
covMethod = c("r", "nlminb", ""),
adjObf = TRUE,
ci = 0.95,
sigdig = 4,
sigdigTable = NULL,
...
)

```

Arguments

eval.max	Maximum number of evaluations of the objective function allowed. Defaults to 200.
iter.max	Maximum number of iterations allowed. Defaults to 150.
trace	The value of the objective function and the parameters is printed every trace'th iteration. When 0 no trace information is to be printed
abs.tol	Absolute tolerance. Defaults to 0 so the absolute convergence test is not used. If the objective function is known to be non-negative, the previous default of '1e-20' would be more appropriate
rel.tol	Relative tolerance. Defaults to '1e-10'.
x.tol	X tolerance. Defaults to '1.5e-8'.
xf.tol	false convergence tolerance. Defaults to '2.2e-14'.
step.min	Minimum step size. Default to '1.'
step.max	Maximum step size. Default to '1.'
sing.tol	singular convergence tolerance; defaults to 'rel.tol;
scale	See PORT documentation (or leave alone).
scale.init	... probably need to check PORT documentation
diff.g	an estimated bound on the relative error in the objective function value
rxControl	'rxode2' ODE solving options during fitting, created with 'rxControl()'

optExpression	Optimize the rxode2 expression to speed up calculation. By default this is turned on.
sumProd	Is a boolean indicating if the model should change multiplication to high precision multiplication and sums to high precision sums using the PreciseSums package. By default this is FALSE.
literalFix	boolean, substitute fixed population values as literals and re-adjust ui and parameter estimates after optimization; Default is 'TRUE'.
returnNlminb	logical; when TRUE this will return the nlminb result instead of the nlmixr2 fit object
solveType	tells if 'nlm' will use nlmixr2's analytical gradients when available (finite differences will be used for event-related parameters like parameters controlling lag time, duration/rate of infusion, and modeled bioavailability). This can be: <ul style="list-style-type: none"> - "hessian" which will use the analytical gradients to create a Hessian with finite differences. - "gradient" which will use the gradient and let 'nlm' calculate the finite difference hessian - "fun" where nlm will calculate both the finite difference gradient and the finite difference Hessian When using nlmixr2's finite differences, the "ideal" step size for either central or forward differences are optimized for with the Shi2021 method which may give more accurate derivatives
stickyRecalcN	The number of bad ODE solves before reducing the atol/rtol for the rest of the problem.
maxOdeRecalc	Maximum number of times to reduce the ODE tolerances and try to resolve the system if there was a bad ODE solve.
odeRecalcFactor	The ODE recalculation factor when ODE solving goes bad, this is the factor the rtol/atol is reduced
eventType	Event gradient type for dosing events; Can be "central" or "forward"
shiErr	This represents the epsilon when optimizing the ideal step size for numeric differentiation using the Shi2021 method
shi21maxFD	The maximum number of steps for the optimization of the forward difference step size when using dosing events (lag time, modeled duration/rate and bioavailability)
optimHessType	The hessian type for when calculating the individual hessian by numeric differences (in generalized log-likelihood estimation). The options are "central", and "forward". The central differences is what R's 'optimHess()' uses and is the default for this method. (Though the "forward" is faster and still reasonable for most cases). The Shi21 cannot be changed for the Gill83 algorithm with the optimHess in a generalized likelihood problem.
hessErr	This represents the epsilon when optimizing the Hessian step size using the Shi2021 method.
shi21maxHess	Maximum number of times to optimize the best step size for the hessian calculation

useColor	Boolean indicating if focei can use ASCII color codes
printNcol	Number of columns to printout before wrapping parameter estimates/gradient
print	Integer representing when the outer step is printed. When this is 0 or do not print the iterations. 1 is print every function evaluation (default), 5 is print every 5 evaluations.
normType	This is the type of parameter normalization/scaling used to get the scaled initial values for nlmixr2. These are used with scaleType of. With the exception of rescale2, these come from Feature Scaling . The rescale2 The rescaling is the same type described in the OptdesX software manual. In general, all all scaling formula can be described by:

$$v_{scaled} = (v_{unscaled} - C_1) / C_2$$

Where

The other data normalization approaches follow the following formula

$$v_{scaled} = (v_{unscaled} - C_1) / C_2$$

- rescale2 This scales all parameters from (-1 to 1). The relative differences between the parameters are preserved with this approach and the constants are:

$$C_1 = (\max(\text{all unscaled values}) + \min(\text{all unscaled values})) / 2$$

$$C_2 = (\max(\text{all unscaled values}) - \min(\text{all unscaled values})) / 2$$

- rescale or min-max normalization. This rescales all parameters from (0 to 1). As in the rescale2 the relative differences are preserved. In this approach:

$$C_1 = \min(\text{all unscaled values})$$

$$C_2 = \max(\text{all unscaled values}) - \min(\text{all unscaled values})$$

- mean or mean normalization. This rescales to center the parameters around the mean but the parameters are from 0 to 1. In this approach:

$$C_1$$

= mean(all unscaled values)

$$C_2$$

= max(all unscaled values) - min(all unscaled values)

- std or standardization. This standardizes by the mean and standard deviation. In this approach:

$$C_1$$

= mean(all unscaled values)

$$C_2$$

= sd(all unscaled values)

- len or unit length scaling. This scales the parameters to the unit length. For this approach we use the Euclidean length, that is:

$$C_1$$

= 0

$$C_2$$

=

$$\sqrt{(v_1^2 + v_2^2 + \dots + v_n^2)}$$

- constant which does not perform data normalization. That is

$$C_1$$

= 0

$$C_2$$

= 1

scaleType

The scaling scheme for nlmixr2. The supported types are:

- nlmixr2 In this approach the scaling is performed by the following equation:

$$v_{scaled}$$

= (

$$v_{current} - v_{init}$$

)*scaleC[i] + scaleTo

The scaleTo parameter is specified by the normType, and the scales are specified by scaleC.

- `norm` This approach uses the simple scaling provided by the `normType` argument.
- `mult` This approach does not use the data normalization provided by `normType`, but rather uses multiplicative scaling to a constant provided by the `scaleTo` argument.

In this case:

$$v_{scaled} = \frac{v_{current}}{v_{init}}$$

*`scaleTo`

- `multAdd` This approach changes the scaling based on the parameter being specified. If a parameter is defined in an exponential block (ie `exp(theta)`), then it is scaled on a linearly, that is:

$$v_{scaled} = (v_{current} - v_{init}) + scaleTo$$

Otherwise the parameter is scaled multiplicatively.

$$v_{scaled} = \frac{v_{current}}{v_{init}}$$

*`scaleTo`

`scaleCmax`

Maximum value of the `scaleC` to prevent overflow.

`scaleCmin`

Minimum value of the `scaleC` to prevent underflow.

`scaleC`

The scaling constant used with `scaleType=nlmixr2`. When not specified, it is based on the type of parameter that is estimated. The idea is to keep the derivatives similar on a log scale to have similar gradient sizes. Hence parameters like `log(exp(theta))` would have a scaling factor of 1 and `log(theta)` would have a scaling factor of `ini_value` (to scale by `1/value`; ie `d/dt(log(ini_value)) = 1/ini_value` or `scaleC=ini_value`)

- For parameters in an exponential (ie `exp(theta)`) or parameters specifying powers, `boxCox` or `yeoJohnson` transformations, this is 1.
- For additive, proportional, lognormal error structures, these are given by `0.5*abs(initial_estimate)`
- Factorials are scaled by `abs(1/digamma(initial_estimate+1))`
- parameters in a log scale (ie `log(theta)`) are transformed by `log(abs(initial_estimate))*abs(initial_esti`

These parameter scaling coefficients are chosen to try to keep similar slopes among parameters. That is they all follow the slopes approximately on a log-scale.

While these are chosen in a logical manner, they may not always apply. You can specify each parameter's scaling factor by this parameter if you wish.

scaleTo	Scale the initial parameter estimate to this value. By default this is 1. When zero or below, no scaling is performed.
gradTo	this is the factor that the gradient is scaled to before optimizing. This only works with scaleType="nlmixr2".
addProp	specifies the type of additive plus proportional errors, the one where standard deviations add (combined1) or the type where the variances add (combined2). The combined1 error type can be described by the following equation:

$$y = f + (a + b \times f^c) \times \varepsilon$$

The combined2 error model can be described by the following equation:

$$y = f + \sqrt{a^2 + b^2 \times f^{2 \times c}} \times \varepsilon$$

Where:

- y represents the observed value
- f represents the predicted value
- a is the additive standard deviation
- b is the proportional/power standard deviation
- c is the power exponent (in the proportional case c=1)

calcTables	This boolean is to determine if the foceiFit will calculate tables. By default this is TRUE
compress	Should the object have compressed items
covMethod	Method for calculating covariance. In this discussion, R is the Hessian matrix of the objective function. The S matrix is the sum of individual gradient cross-product (evaluated at the individual empirical Bayes estimates). <ul style="list-style-type: none"> • "r, s" Uses the sandwich matrix to calculate the covariance, that is: solve(R) %*% S %*% solve(R) • "r" Uses the Hessian matrix to calculate the covariance as 2 %*% solve(R) • "s" Uses the cross-product matrix to calculate the covariance as 4 %*% solve(S) • "" Does not calculate the covariance step.
adj0bf	is a boolean to indicate if the objective function should be adjusted to be closer to NONMEM's default objective function. By default this is TRUE
ci	Confidence level for some tables. By default this is 0.95 or 95% confidence.
sigdig	Optimization significant digits. This controls: <ul style="list-style-type: none"> • The tolerance of the inner and outer optimization is $10^{-\text{sigdig}}$ • The tolerance of the ODE solvers is $0.5 \times 10^{-(\text{sigdig}-2)}$; For the sensitivity equations and steady-state solutions the default is $0.5 \times 10^{-(\text{sigdig}-1.5)}$ (sensitivity changes only applicable for liblsoda)

- The tolerance of the boundary check is $5 * 10^{(-sigdig + 1)}$
- sigdigTable Significant digits in the final output table. If not specified, then it matches the significant digits in the 'sigdig' optimization algorithm. If 'sigdig' is NULL, use 3.
- ... Additional arguments passed to `nlmixr2est::nlminbControl()`.

Author(s)

Matthew L. Fidler

Examples

```
# A logit regression example with emax model

dsn <- data.frame(i=1:1000)
dsn$time <- exp(rnorm(1000))
dsn$DV=rbinom(1000,1,exp(-1+dsn$time)/(1+exp(-1+dsn$time)))

mod <- function() {
  ini({
    E0 <- 0.5
    Em <- 0.5
    E50 <- 2
    g <- fix(2)
  })
  model({
    v <- E0+Em*time^g/(E50^g+time^g)
    ll(bin) ~ DV * v - log(1 + exp(v))
  })
}

fit2 <- nlmixr(mod, dsn, est="nlminb")

print(fit2)

# you can also get the nlm output with fit2$nlminb

fit2$nlminb
```

nlmixr2

nlmixr2 fits population PK and PKPD non-linear mixed effects models.

Description

nlmixr2 is an R package for fitting population pharmacokinetic (PK) and pharmacokinetic-pharmacodynamic (PKPD) models.

Usage

```
nlmixr2(
  object,
  data,
  est = NULL,
  control = list(),
  table = tableControl(),
  ...,
  save = NULL,
  envir = parent.frame()
)
```

Arguments

object	Fitted object or function specifying the model.
data	nlmixr data
est	estimation method (all methods are shown by ‘nlmixr2AllEst()’). Methods can be added for other tools
control	The estimation control object. These are expected to be different for each type of estimation method
table	The output table control object (like ‘tableControl()’)
...	Additional arguments passed to <code>nlmixr2est::nlmixr2()</code> .
save	Boolean to save a nlmixr2 object in a rds file in the working directory. If NULL, uses option "nlmixr2.save"
envir	Environment where the nlmixr object/function is evaluated before running the estimation routine.

Details

The nlmixr2 generalized function allows common access to the nlmixr2 estimation routines.

The nlmixr object has the following fields:

Field	Description
conditionNumber	Condition number, that is the highest divided by the lowest eigenvalue in the population covariance matrix
cor	Correlation matrix
phiR	correlation matrix of each individual’s eta (if present)
objDF	Data frame containing objective function information (AIC, BIC, etc.)
time	Duration of different parts of the analysis (e.g. setup, optimization, calculation of covariance, etc.)
theta	Estimates for eta for each individual
etaObf	Estimates for eta for each individual, This also includes the objective function for each individual
fixef	Estimates of fixed effects
foceiControl	Estimation options if focei was used
ui	Final estimates for the model
dataMergeFull	Full data merge with the fit output and the original dataset; Also includes nlmixrLlikObs which includes t
censInfo	Gives the censornrg information abot the fit (the type of censoring that was send and handled in the dataset)
dataLloq	Gives the lloq from the dataset (average) when cesoring has occurred; Requires the fit to have a table step

dataUloq	Gives the uloq from the dataset (average) when censoring has occurred; requires the fit to have a table step
eta	IIV values for each individual
dataMergeInner	Inner data merge with the fit output and the original dataset; Also includes nlmixrLlikObs which includes t
rxControl	Integration options used to control rxode2
dataMergeLeft	Left data merge with the fit output and the original dataset; Also includes nlmixrLlikObs which includes t
omega	Matrix containing the estimates of the multivariate normal covariance matrix for between subject variability
covMethod	Method used to calculate covariance of the fixed effects
modelName	Name of the R object containing the model
origData	Original dataset
phiRSE	Relative standard error of each individual's eta
dataMergeRight	Right data merge with the fit output and the original dataset; Also includes nlmixrLlikObs which includes t
ipredModel	rxode2 estimation model for fit (internal will likely be removed from visibility)
phiSE	Standard error of each individual's eta
parFixed	Table of parameter estimates (rounded and pretty looking)
parFixedDF	Table of parameter estimates as a data frame
omegaR	The correlation matrix of omega with standard deviations for the diagonal pieces
iniUi	The initial model used to start the estimation
finalUi	The model with the estimates replaced as values
scaleInfo	The scaling factors used for nlmixr2 estimation in focei; The can be changed by foceiControl(scaleC=...
table	These are the table options that were used when generating the table output (were CWRES included, etc
shrink	This is a table of shrinkages for all the individual ETAs as well as the variance shrinkage as well as summ
env	This is the environment where all the information for the fit is stored outside of the data-frame. It is an R
seed	This is the initial seed used for saem
simInfo	This returns a list of all the fit information used for a traditional rxode2 simulation, which you can tweak
runInfo	This returns a list of all the warnings or fit information
parHistStacked	Value of objective function and parameters at each iteration (tall format)
parHist	Value of objective function and parameters at each iteration (wide format)
cov	Variance-covariance matrix

Value

Either a nlmixr2 model or a nlmixr2 fit object

nlmixr modeling mini-language

Rationale

nlmixr estimation routines each have their own way of specifying models. Often the models are specified in ways that are most intuitive for one estimation routine, but do not make sense for another estimation routine. Sometimes, legacy estimation routines like [nlme](#) have their own syntax that is outside of the control of the nlmixr package.

The unique syntax of each routine makes the routines themselves easier to maintain and expand, and allows interfacing with existing packages that are outside of nlmixr (like [nlme](#)). However, a model definition language that is common between estimation methods, and an output object that is uniform, will make it easier to switch between estimation routines and will facilitate interfacing output with external packages like Xpose.

The nlmixr mini-modeling language, attempts to address this issue by incorporating a common language. This language is inspired by both R and NONMEM, since these languages are familiar

to many pharmacometricians.

Initial Estimates and boundaries for population parameters

nlmixr models are contained in a R function with two blocks: ini and model. This R function can be named anything, but is not meant to be called directly from R. In fact if you try you will likely get an error such as Error: could not find function "ini".

The ini model block is meant to hold the initial estimates for the model, and the boundaries of the parameters for estimation routines that support boundaries (note nlmixr's saem and nlme do not currently support parameter boundaries).

To explain how these initial estimates are specified we will start with an annotated example:

```
f <- function(){ ## Note the arguments to the function are currently
                  ## ignored by nlmixr
  ini({
    ## Initial conditions for population parameters (sometimes
    ## called theta parameters) are defined by either `<-` or `=`
    lCl <- 1.6      #log Cl (L/hr)
    ## Note that simple expressions that evaluate to a number are
    ## OK for defining initial conditions (like in R)
    lVc = log(90)  #log V (L)
    ## Also a comment on a parameter is captured as a parameter label
    lKa <- 1 #log Ka (1/hr)
    ## Bounds may be specified by c(lower, est, upper), like NONMEM:
    ## Residuals errors are assumed to be population parameters
    prop.err <- c(0, 0.2, 1)
  })
  ## The model block will be discussed later
  model({})
}
```

As shown in the above examples:

- Simple parameter values are specified as a R-compatible assignment
- Boundaries may be specified by c(lower, est, upper).
- Like NONMEM, c(lower, est) is equivalent to c(lower, est, Inf)
- Also like NONMEM, c(est) does not specify a lower bound, and is equivalent to specifying the parameter without R's 'c' function.
- The initial estimates are specified on the variance scale, and in analogy with NONMEM, the square roots of the diagonal elements correspond to coefficients of variation when used in the exponential IIV implementation

These parameters can be named almost any R compatible name. Please note that:

- Residual error estimates should be coded as population estimates (i.e. using an '=' or '<-' statement, not a '~').
- Naming variables that start with "_" are not supported. Note that R does not allow variable starting with "_" to be assigned without quoting them.

- Naming variables that start with "rx_" or "nlmixr_" is not supported since [rxode2](#) and nlmixr2 use these prefixes internally for certain estimation routines and calculating residuals.
- Variable names are case sensitive, just like they are in R. "CL" is not the same as "Cl".

Initial Estimates for between subject error distribution (NONMEM's \$OMEGA)

In mixture models, multivariate normal individual deviations from the population parameters are estimated (in NONMEM these are called eta parameters). Additionally the variance/covariance matrix of these deviations is also estimated (in NONMEM this is the OMEGA matrix). These also have initial estimates. In nlmixr these are specified by the '~' operator that is typically used in R for "modeled by", and was chosen to distinguish these estimates from the population and residual error parameters.

Continuing the prior example, we can annotate the estimates for the between subject error distribution

```
f <- function(){
  ini({
    lCl <- 1.6      #log Cl (L/hr)
    lVc = log(90)  #log V (L)
    lKa <- 1 #log Ka (1/hr)
    prop.err <- c(0, 0.2, 1)
    ## Initial estimate for ka IIV variance
    ## Labels work for single parameters
    eta.ka ~ 0.1 # BSV Ka

    ## For correlated parameters, you specify the names of each
    ## correlated parameter separated by a addition operator `+`
    ## and the left handed side specifies the lower triangular
    ## matrix initial of the covariance matrix.
    eta.cl + eta.vc ~ c(0.1,
                       0.005, 0.1)
    ## Note that labels do not currently work for correlated
    ## parameters. Also do not put comments inside the lower
    ## triangular matrix as this will currently break the model.
  })
  ## The model block will be discussed later
  model({})
}
```

As shown in the above examples:

- Simple variances are specified by the variable name and the estimate separated by '~'.
- Correlated parameters are specified by the sum of the variable labels and then the lower triangular matrix of the covariance is specified on the left handed side of the equation. This is also separated by '~'.

Currently the model syntax does not allow comments inside the lower triangular matrix.

Model Syntax for ODE based models (NONMEM's \$PK, \$PRED, \$DES and \$ERROR)

Once the initialization block has been defined, you can define a model in terms of the defined variables in the ini block. You can also mix in RxODE blocks into the model.

The current method of defining a nlmixr model is to specify the parameters, and then possibly the RxODE lines:

Continuing describing the syntax with an annotated example:

```
f <- function(){
  ini({
    lCl <- 1.6      #log Cl (L/hr)
    lVc <- log(90)  #log Vc (L)
    lKA <- 0.1     #log Ka (1/hr)
    prop.err <- c(0, 0.2, 1)
    eta.Cl ~ 0.1 ## BSV Cl
    eta.Vc ~ 0.1 ## BSV Vc
    eta.KA ~ 0.1 ## BSV Ka
  })
  model({
    ## First parameters are defined in terms of the initial estimates
    ## parameter names.
    Cl <- exp(lCl + eta.Cl)
    Vc = exp(lVc + eta.Vc)
    KA <- exp(lKA + eta.KA)
    ## After the differential equations are defined
    kel <- Cl / Vc;
    d/dt(depot) = -KA*depot;
    d/dt(centr) = KA*depot-kel*centr;
    ## And the concentration is then calculated
    cp = centr / Vc;
    ## Last, nlmixr is told that the plasma concentration follows
    ## a proportional error (estimated by the parameter prop.err)
    cp ~ prop(prop.err)
  })
}
```

A few points to note:

- Parameters are often defined before the differential equations.
- The differential equations, parameters and error terms are in a single block, instead of multiple sections.
- State names, calculated variables cannot start with either "rx_" or "nlmixr_" since these are used internally in some estimation routines.
- Errors are specified using the '~'. Currently you can use either add(parameter) for additive error, prop(parameter) for proportional error or add(parameter1) + prop(parameter2) for additive plus proportional error. You can also specify norm(parameter) for the additive error, since it follows a normal distribution.
- Some routines, like saem require parameters in terms of Pop.Parameter + Individual.Deviation.Parameter + Covariate*Covariate.Parameter. The order of these parameters do not matter. This is similar to NONMEM's mu-referencing, though not quite so restrictive.

- The type of parameter in the model is determined by the initial block; Covariates used in the model are missing in the ini block. These variables need to be present in the modeling dataset for the model to run.

Model Syntax for solved PK systems

Solved PK systems are also currently supported by nlmixr with the ‘linCmt()’ pseudo-function. An annotated example of a solved system is below:

```
##
f <- function(){
  ini({
    lCl <- 1.6      #log Cl (L/hr)
    lVc <- log(90)  #log Vc (L)
    lKA <- 0.1     #log Ka (1/hr)
    prop.err <- c(0, 0.2, 1)
    eta.Cl ~ 0.1 ## BSV Cl
    eta.Vc ~ 0.1 ## BSV Vc
    eta.KA ~ 0.1 ## BSV Ka
  })
  model({
    Cl <- exp(lCl + eta.Cl)
    Vc = exp(lVc + eta.Vc)
    KA <- exp(lKA + eta.KA)
    ## Instead of specifying the ODEs, you can use
    ## the linCmt() function to use the solved system.
    ##
    ## This function determines the type of PK solved system
    ## to use by the parameters that are defined. In this case
    ## it knows that this is a one-compartment model with first-order
    ## absorption.
    linCmt() ~ prop(prop.err)
  })
}
```

A few things to keep in mind:

- While RxODE allows mixing of solved systems and ODEs, this has not been implemented in nlmixr yet.
- The solved systems implemented are the one, two and three compartment models with or without first-order absorption. Each of the models support a lag time with a tlag parameter.
- In general the linear compartment model figures out the model by the parameter names. nlmixr currently knows about numbered volumes, Vc/Vp, Clearances in terms of both Cl and Q/CLD. Additionally nlmixr knows about elimination micro-constants (ie K12). Mixing of these parameters for these models is currently not supported.

Checking model syntax

After specifying the model syntax you can check that nlmixr is interpreting it correctly by using the nlmixr function on it.

Using the above function we can get:

```

> nlmixr(f)
## 1-compartment model with first-order absorption in terms of Cl
## Initialization:
#####
Fixed Effects ($theta):
      lCl      lVc      lKA
1.60000 4.49981 0.10000

Omega ($omega):
      [,1] [,2] [,3]
[1,] 0.1 0.0 0.0
[2,] 0.0 0.1 0.0
[3,] 0.0 0.0 0.1

## Model:
#####
Cl <- exp(lCl + eta.Cl)
Vc = exp(lVc + eta.Vc)
KA <- exp(lKA + eta.KA)
## Instead of specifying the ODEs, you can use
## the linCmt() function to use the solved system.
##
## This function determines the type of PK solved system
## to use by the parameters that are defined. In this case
## it knows that this is a one-compartment model with first-order
## absorption.
linCmt() ~ prop(prop.err)

```

In general this gives you information about the model (what type of solved system/RxODE), initial estimates as well as the code for the model block.

Using the model syntax for estimating a model

Once the model function has been created, you can use it and a dataset to estimate the parameters for a model given a dataset.

This dataset has to have RxODE compatible events IDs. Both Monolix and NONMEM use a very similar standard to what nlmixr can support.

Once the data has been converted to the appropriate format, you can use the nlmixr function to run the appropriate code.

The method to estimate the model is:

```
fit <- nlmixr(model.function, dataset, est="est", control=estControl(options))
```

Currently nlme and saem are implemented. For example, to run the above model with saem, we could have the following:

```

> f <- function(){
  ini({
    lCl <- 1.6      #log Cl (L/hr)

```

```

    lVc <- log(90)    #log Vc (L)
    lKA <- 0.1        #log Ka (1/hr)
    prop.err <- c(0, 0.2, 1)
    eta.Cl ~ 0.1 ## BSV Cl
    eta.Vc ~ 0.1 ## BSV Vc
    eta.KA ~ 0.1 ## BSV Ka
  })
  model({
    ## First parameters are defined in terms of the initial estimates
    ## parameter names.
    Cl <- exp(lCl + eta.Cl)
    Vc = exp(lVc + eta.Vc)
    KA <- exp(lKA + eta.KA)
    ## After the differential equations are defined
    kel <- Cl / Vc;
    d/dt(depot) = -KA*depot;
    d/dt(centr) = KA*depot-kel*centr;
    ## And the concentration is then calculated
    cp = centr / Vc;
    ## Last, nlmixr is told that the plasma concentration follows
    ## a proportional error (estimated by the parameter prop.err)
    cp ~ prop(prop.err)
  })
}
> fit.s <- nlmixr(f,d,est="saem",control=saemControl(n.burn=50,n.em=100,print=50));
Compiling RxODE differential equations...done.
c:/Rtools/mingw_64/bin/g++ -I"c:/R/R-34~1.1/include" -DNDEBUG -I"d:/Compiler/gcc-4.9.3/local330/i
In file included from c:/R/R-34~1.1/library/RCPAR~1/include/armadillo:52:0,
      from c:/R/R-34~1.1/library/RCPAR~1/include/RcppArmadilloForward.h:46,
      from c:/R/R-34~1.1/library/RCPAR~1/include/RcppArmadillo.h:31,
      from saem3090757b4bd1x64.cpp:1:
c:/R/R-34~1.1/library/RCPAR~1/include/armadillo_bits/compiler_setup.hpp:474:96: note: #pragma messa
      #pragma message ("WARNING: use of OpenMP disabled; this compiler doesn't support OpenMP 3.0+")
      ^
c:/Rtools/mingw_64/bin/g++ -shared -s -static-libgcc -o saem3090757b4bd1x64.dll tmp.def saem3090757b4
done.
1:    1.8174    4.6328    0.0553    0.0950    0.0950    0.0950    0.6357
50:    1.3900    4.2039    0.0001    0.0679    0.0784    0.1082    0.1992
100:   1.3894    4.2054    0.0107    0.0686    0.0777    0.1111    0.1981
150:   1.3885    4.2041    0.0089    0.0683    0.0778    0.1117    0.1980
Using sympy via SnakeCharmR
## Calculate ETA-based prediction and error derivatives:
Calculate Jacobian.....done.
Calculate sensitivities.....
done.
## Calculate d(f)/d(eta)
## ...
## done

```

```
## ...
## done
The model-based sensitivities have been calculated
Calculating Table Variables...
done
```

The options for saem are controlled by `saemControl`. You may wish to make sure the minimization is complete in the case of saem. You can do that with `traceplot` which shows the iteration history with the divided by burn-in and EM phases. In this case, the burn in seems reasonable; you may wish to increase the number of iterations in the EM phase of the estimation. Overall it is probably a semi-reasonable solution.

nlmixr output objects

In addition to unifying the modeling language sent to each of the estimation routines, the outputs currently have a unified structure.

You can see the fit object by typing the object name:

```
> fit.s
-- nlmixr SAEM fit (ODE); OBJF calculated from FOCEi approximation -----
      OBJF      AIC      BIC Log-likelihood Condition Number
62337.09 62351.09 62399.01      -31168.55      82.6086

-- Time (sec; fit.s$time): -----
      saem setup Likelihood Calculation covariance table
elapsed 430.25 31.64      1.19      0 3.44

-- Parameters (fit.s$par.fixed): -----
      Parameter Estimate      SE
1Cl    log Cl (L/hr)      1.39 0.0240 1.73      4.01 (3.83, 4.20) 26.6
1Vc    log Vc (L)      4.20 0.0256 0.608      67.0 (63.7, 70.4) 28.5
1KA    log Ka (1/hr) 0.00924 0.0323 349.      1.01 (0.947, 1.08) 34.3
prop.err    prop.err    0.198      19.8
      Shrink(SD)
1Cl    0.248
1Vc    1.09
1KA    4.19
prop.err    1.81

      No correlations in between subject variability (BSV) matrix
      Full BSV covariance (fit.s$omega) or correlation (fit.s$omega.R; diagonals=SDs)
      Distribution stats (mean/skewness/kurtosis/p-value) available in fit.s$shrink

-- Fit Data (object fit.s is a modified data.frame): -----
# A tibble: 6,947 x 22
  ID    TIME    DV  PRED    RES    WRES IPRED  IRES  IWRES CPRED  CRES
* <fct> <dbl> <dbl> <dbl> <dbl> <dbl> <dbl> <dbl> <dbl> <dbl> <dbl>
1 1      0.25  205.  198.   6.60  0.0741  189.  16.2  0.434  198.   6.78
2 1      0.5    311.  349. -38.7 -0.261  330. -19.0 -0.291  349. -38.3
3 1      0.75  389.  464. -74.5 -0.398  434. -45.2 -0.526  463. -73.9
```

```
# ... with 6,944 more rows, and 11 more variables: CWRES <dbl>, eta.Cl <dbl>,
# eta.Vc <dbl>, eta.KA <dbl>, depot <dbl>, centr <dbl>, Cl <dbl>, Vc <dbl>,
# KA <dbl>, kel <dbl>, cp <dbl>
```

This example shows what is typical printout of a nlmixr fit object. The elements of the fit are:

- The type of fit (`nlme`, `saem`, etc)
- Metrics of goodness of fit (`AIC`, `BIC`, and `logLik`).
 - To align the comparison between methods, the FOCEi likelihood objective is calculated regardless of the method used and used for goodness of fit metrics.
 - This FOCEi likelihood has been compared to NONMEM's objective function and gives the same values (based on the data in Wang 2007)
 - Also note that `saem` does not calculate an objective function, and the FOCEi is used as the only objective function for the fit.
 - Even though the objective functions are calculated in the same manner, caution should be used when comparing fits from various estimation routines.
- The next item is the timing of each of the steps of the fit.
 - These can be also accessed by (`fit.s$time`).
 - As a mnemonic, the access for this item is shown in the printout. This is true for almost all of the other items in the printout.
- After the timing of the fit, the parameter estimates are displayed (can be accessed by `fit.s$par.fixed`)
 - While the items are rounded for R printing, each estimate without rounding is still accessible by the '\$' syntax. For example, the '\$Untransformed' gives the untransformed parameter values.
 - The Untransformed parameter takes log-space parameters and back-transforms them to normal parameters. Not the CIs are listed on the back-transformed parameter space.
 - Proportional Errors are converted to
- Omega block (accessed by `fit.s$omega`)
- The table of fit data. Please note:
 - A nlmixr fit object is actually a data frame. Saving it as a Rdata object and then loading it without nlmixr will just show the data by itself. Don't worry; the fit information has not vanished, you can bring it back by simply loading nlmixr, and then accessing the data.
 - Special access to fit information (like the `$omega`) needs nlmixr to extract the information.
 - If you use the \$ to access information, the order of precedence is:
 - * Fit data from the overall data.frame
 - * Information about the parsed nlmixr model (via `$uif`)
 - * Parameter history if available (via `$par.hist` and `$par.hist.stacked`)
 - * Fixed effects table (via `$par.fixed`)
 - * Individual differences from the typical population parameters (via `$eta`)
 - * Fit information from the list of information generated during the post-hoc residual calculation.
 - * Fit information from the environment where the post-hoc residual were calculated
 - * Fit information about how the data and options interacted with the specified model (such as estimation options or if the solved system is for an infusion or an IV bolus).

- While the printout may displays the data as a `data.table` object or `tbl` object, the data is NOT any of these objects, but rather a derived data frame.
- Since the object *is* a `data.frame`, you can treat it like one.

In addition to the above properties of the fit object, there are a few additional that may be helpful for the modeler:

- `$theta` gives the fixed effects parameter estimates (in NONMEM the `thetas`). This can also be accessed in `fixed.effects` function. Note that the residual variability is treated as a fixed effect parameter and is included in this list.
- `$eta` gives the random effects parameter estimates, or in NONMEM the `etas`. This can also be accessed in using the `random.effects` function.

Author(s)

Matthew L. Fidler

Examples

```
one.cmt <- function() {
  ini({
    ## You may label each parameter with a comment
    tka <- 0.45 # Ka
    tcl <- log(c(0, 2.7, 100)) # Log Cl
    ## This works with interactive models
    ## You may also label the preceding line with label("label text")
    tv <- 3.45; label("log V")
    ## the label("Label name") works with all models
    eta.ka ~ 0.6
    eta.cl ~ 0.3
    eta.v ~ 0.1
    add.sd <- 0.7
    prop.sd <- 0.01
  })
  model({
    ka <- exp(tka + eta.ka)
    cl <- exp(tcl + eta.cl)
    v <- exp(tv + eta.v)
    linCmt() ~ add(add.sd) + prop(prop.sd)
  })
}

# fitF <- nlmixr(one.cmt, theo_sd, "focei")

fitS <- nlmixr(one.cmt, theo_sd, "saem")
```

nlmixr2CheckInstall *Check your nlmixr2 installation for potential issues*

Description

Check your nlmixr2 installation for potential issues

Usage

```
nlmixr2CheckInstall()
```

Examples

```
nlmixr2CheckInstall()
```

nlsControl *nlmixr2 defaults controls for nls*

Description

nlmixr2 defaults controls for nls

Usage

```
nlsControl(  
  maxiter = 10000,  
  tol = 1e-05,  
  minFactor = 1/1024,  
  printEval = FALSE,  
  warnOnly = FALSE,  
  scaleOffset = 0,  
  nDcentral = FALSE,  
  algorithm = c("LM", "default", "plinear", "port"),  
  ftol = sqrt(.Machine$double.eps),  
  ptol = sqrt(.Machine$double.eps),  
  gtol = 0,  
  diag = list(),  
  epsfcn = 0,  
  factor = 100,  
  maxfev = integer(),  
  nprint = 0,  
  solveType = c("grad", "fun"),  
  stickyRecalcN = 4,  
  maxOdeRecalc = 5,  
  odeRecalcFactor = 10^(0.5),
```

```

eventType = c("central", "forward"),
shiErr = (.Machine$double.eps)^(1/3),
shi21maxFD = 20L,
useColor = crayon::has_color(),
printNcol = floor((getOption("width") - 23)/12),
print = 1L,
normType = c("rescale2", "mean", "rescale", "std", "len", "constant"),
scaleType = c("nlmixr2", "norm", "mult", "multAdd"),
scaleCmax = 1e+05,
scaleCmin = 1e-05,
scaleC = NULL,
scaleTo = 1,
gradTo = 1,
trace = FALSE,
rxControl = NULL,
optExpression = TRUE,
sumProd = FALSE,
literalFix = TRUE,
returnNls = FALSE,
addProp = c("combined2", "combined1"),
calcTables = TRUE,
compress = TRUE,
adjObf = TRUE,
ci = 0.95,
sigdig = 4,
sigdigTable = NULL,
...
)

```

Arguments

maxiter	A positive integer specifying the maximum number of iterations allowed.
tol	A positive numeric value specifying the tolerance level for the relative offset convergence criterion.
minFactor	A positive numeric value specifying the minimum step-size factor allowed on any step in the iteration. The increment is calculated with a Gauss-Newton algorithm and successively halved until the residual sum of squares has been decreased or until the step-size factor has been reduced below this limit.
printEval	a logical specifying whether the number of evaluations (steps in the gradient direction taken each iteration) is printed.
warnOnly	a logical specifying whether <code>nls()</code> should return instead of signalling an error in the case of termination before convergence. Termination before convergence happens upon completion of <code>maxiter</code> iterations, in the case of a singular gradient, and in the case that the step-size factor is reduced below <code>minFactor</code> .
scaleOffset	a constant to be added to the denominator of the relative offset convergence criterion calculation to avoid a zero divide in the case where the fit of a model to data is very close. The default value of 0 keeps the legacy behaviour of <code>nls()</code> .

	A value such as 1 seems to work for problems of reasonable scale with very small residuals.
nDcentral	only when <i>numerical</i> derivatives are used: <code>logical</code> indicating if <i>central</i> differences should be employed, i.e., <code>numericDeriv(*, central=TRUE)</code> be used.
algorithm	character string specifying the algorithm to use. The default algorithm is a Gauss-Newton algorithm. Other possible values are "plinear" for the Golub-Pereyra algorithm for partially linear least-squares models and "port" for the 'nl2sol' algorithm from the Port library – see the references. Can be abbreviated.
ftol	non-negative numeric. Termination occurs when both the actual and predicted relative reductions in the sum of squares are at most ftol. Therefore, ftol measures the relative error desired in the sum of squares.
ptol	non-negative numeric. Termination occurs when the relative error between two consecutive iterates is at most ptol. Therefore, ptol measures the relative error desired in the approximate solution.
gtol	non-negative numeric. Termination occurs when the cosine of the angle between result of fn evaluation <i>fvec</i> and any column of the Jacobian is at most gtol in absolute value. Therefore, gtol measures the orthogonality desired between the function vector and the columns of the Jacobian.
diag	a list or numeric vector containing positive entries that serve as multiplicative scale factors for the parameters. Length of diag should be equal to that of par. If not, user-provided diag is ignored and diag is internally set.
epsfcn	(used if jac is not provided) is a numeric used in determining a suitable step for the forward-difference approximation. This approximation assumes that the relative errors in the functions are of the order of epsfcn. If epsfcn is less than the machine precision, it is assumed that the relative errors in the functions are of the order of the machine precision.
factor	positive numeric, used in determining the initial step bound. This bound is set to the product of factor and the $ \text{diag} * \text{par} $ if nonzero, or else to factor itself. In most cases factor should lie in the interval (0.1,100). 100 is a generally recommended value.
maxfev	integer; termination occurs when the number of calls to fn has reached maxfev. Note that <code>nls.lm</code> sets the value of maxfev to $100 * (\text{length}(\text{par}) + 1)$ if <code>maxfev = integer()</code> , where par is the list or vector of parameters to be optimized.
nprint	is an integer; set nprint to be positive to enable printing of iterates
solveType	tells if 'nlm' will use <code>nlmixr2</code> 's analytical gradients when available (finite differences will be used for event-related parameters like parameters controlling lag time, duration/rate of infusion, and modeled bioavailability). This can be: <ul style="list-style-type: none"> - "hessian" which will use the analytical gradients to create a Hessian with finite differences. - "gradient" which will use the gradient and let 'nlm' calculate the finite difference hessian - "fun" where nlm will calculate both the finite difference gradient and the finite difference Hessian <p>When using <code>nlmixr2</code>'s finite differences, the "ideal" step size for either central or forward differences are optimized for with the Shi2021 method which may give more accurate derivatives</p>

stickyRecalcN	The number of bad ODE solves before reducing the atol/rtol for the rest of the problem.
maxOdeRecalc	Maximum number of times to reduce the ODE tolerances and try to resolve the system if there was a bad ODE solve.
odeRecalcFactor	The ODE recalculation factor when ODE solving goes bad, this is the factor the rtol/atol is reduced
eventType	Event gradient type for dosing events; Can be "central" or "forward"
shiErr	This represents the epsilon when optimizing the ideal step size for numeric differentiation using the Shi2021 method
shi21maxFD	The maximum number of steps for the optimization of the forward difference step size when using dosing events (lag time, modeled duration/rate and bioavailability)
useColor	Boolean indicating if focei can use ASCII color codes
printNcol	Number of columns to printout before wrapping parameter estimates/gradient
print	Integer representing when the outer step is printed. When this is 0 or do not print the iterations. 1 is print every function evaluation (default), 5 is print every 5 evaluations.
normType	This is the type of parameter normalization/scaling used to get the scaled initial values for nlmixr2. These are used with scaleType of. With the exception of rescale2, these come from Feature Scaling . The rescale2 The rescaling is the same type described in the OptdesX software manual. In general, all all scaling formula can be described by:

$$= \left(\frac{v_{scaled}}{v_{unscaled} - C_1} \right) / C_2$$

Where

The other data normalization approaches follow the following formula

$$= \left(\frac{v_{scaled}}{v_{unscaled} - C_1} \right) / C_2$$

- rescale2 This scales all parameters from (-1 to 1). The relative differences between the parameters are preserved with this approach and the constants are:

$$C_1$$

$$= (\max(\text{all unscaled values}) + \min(\text{all unscaled values})) / 2$$

$$C_2$$

$$= (\max(\text{all unscaled values}) - \min(\text{all unscaled values})) / 2$$

- `rescale` or min-max normalization. This rescales all parameters from (0 to 1). As in the `rescale2` the relative differences are preserved. In this approach:

$$C_1$$

$$= \min(\text{all unscaled values})$$

$$C_2$$

$$= \max(\text{all unscaled values}) - \min(\text{all unscaled values})$$

- `mean` or mean normalization. This rescales to center the parameters around the mean but the parameters are from 0 to 1. In this approach:

$$C_1$$

$$= \text{mean}(\text{all unscaled values})$$

$$C_2$$

$$= \max(\text{all unscaled values}) - \min(\text{all unscaled values})$$

- `std` or standardization. This standardizes by the mean and standard deviation. In this approach:

$$C_1$$

$$= \text{mean}(\text{all unscaled values})$$

$$C_2$$

$$= \text{sd}(\text{all unscaled values})$$

- `len` or unit length scaling. This scales the parameters to the unit length. For this approach we use the Euclidean length, that is:

$$C_1$$

$$= 0$$

$$C_2$$

$$=$$

$$\sqrt{(v_1^2 + v_2^2 + \dots + v_n^2)}$$

- `constant` which does not perform data normalization. That is

$$C_1$$

$$= 0$$

C_2

= 1

scaleType The scaling scheme for nlmixr2. The supported types are:

- nlmixr2 In this approach the scaling is performed by the following equation:

$$v_{scaled} = (v_{current} - v_{init}) * scaleC[i] + scaleTo$$

The scaleTo parameter is specified by the normType, and the scales are specified by scaleC.

- norm This approach uses the simple scaling provided by the normType argument.
- mult This approach does not use the data normalization provided by normType, but rather uses multiplicative scaling to a constant provided by the scaleTo argument. In this case:

$$v_{scaled} = v_{current} / v_{init} * scaleTo$$

- multAdd This approach changes the scaling based on the parameter being specified. If a parameter is defined in an exponential block (ie exp(theta)), then it is scaled on a linearly, that is:

$$v_{scaled} = (v_{current} - v_{init}) + scaleTo$$

Otherwise the parameter is scaled multiplicatively.

$$v_{scaled} = v_{current} / v_{init} * scaleTo$$

scaleCmax	Maximum value of the scaleC to prevent overflow.
scaleCmin	Minimum value of the scaleC to prevent underflow.
scaleC	<p>The scaling constant used with <code>scaleType=nlmixr2</code>. When not specified, it is based on the type of parameter that is estimated. The idea is to keep the derivatives similar on a log scale to have similar gradient sizes. Hence parameters like $\log(\exp(\theta))$ would have a scaling factor of 1 and $\log(\theta)$ would have a scaling factor of <code>ini_value</code> (to scale by $1/\text{value}$; ie $d/dt(\log(\text{ini_value})) = 1/\text{ini_value}$ or <code>scaleC=ini_value</code>)</p> <ul style="list-style-type: none"> • For parameters in an exponential (ie $\exp(\theta)$) or parameters specifying powers, <code>boxCox</code> or <code>yeoJohnson</code> transformations, this is 1. • For additive, proportional, lognormal error structures, these are given by $0.5 * \text{abs}(\text{initial_estimate})$ • Factorials are scaled by $\text{abs}(1/\text{digamma}(\text{initial_estimate}+1))$ • parameters in a log scale (ie $\log(\theta)$) are transformed by $\log(\text{abs}(\text{initial_estimate})) * \text{abs}(\text{initial_estimate})$ <p>These parameter scaling coefficients are chosen to try to keep similar slopes among parameters. That is they all follow the slopes approximately on a log-scale.</p> <p>While these are chosen in a logical manner, they may not always apply. You can specify each parameter's scaling factor by this parameter if you wish.</p>
scaleTo	Scale the initial parameter estimate to this value. By default this is 1. When zero or below, no scaling is performed.
gradTo	this is the factor that the gradient is scaled to before optimizing. This only works with <code>scaleType="nlmixr2"</code> .
trace	logical value indicating if a trace of the iteration progress should be printed. Default is FALSE. If TRUE the residual (weighted) sum-of-squares, the convergence criterion and the parameter values are printed at the conclusion of each iteration. Note that <code>format()</code> is used, so these mostly depend on <code>getOption("digits")</code> . When the "plinear" algorithm is used, the conditional estimates of the linear parameters are printed after the nonlinear parameters. When the "port" algorithm is used the objective function value printed is half the residual (weighted) sum-of-squares.
rxControl	'rxode2' ODE solving options during fitting, created with <code>'rxControl()'</code>
optExpression	Optimize the rxode2 expression to speed up calculation. By default this is turned on.
sumProd	Is a boolean indicating if the model should change multiplication to high precision multiplication and sums to high precision sums using the <code>PreciseSums</code> package. By default this is FALSE.
literalFix	boolean, substitute fixed population values as literals and re-adjust ui and parameter estimates after optimization; Default is 'TRUE'.
returnNls	logical; when TRUE, will return the nls object instead of the nlmixr object
addProp	specifies the type of additive plus proportional errors, the one where standard deviations add (<code>combined1</code>) or the type where the variances add (<code>combined2</code>). The <code>combined1</code> error type can be described by the following equation:

$$y = f + (a + b \times f^c) \times \varepsilon$$

The combined2 error model can be described by the following equation:

$$y = f + \sqrt{a^2 + b^2 \times f^{2 \times c}} \times \varepsilon$$

Where:

- y represents the observed value
- f represents the predicted value
- a is the additive standard deviation
- b is the proportional/power standard deviation
- c is the power exponent (in the proportional case c=1)

calcTables	This boolean is to determine if the foceiFit will calculate tables. By default this is TRUE
compress	Should the object have compressed items
adj0bf	is a boolean to indicate if the objective function should be adjusted to be closer to NONMEM's default objective function. By default this is TRUE
ci	Confidence level for some tables. By default this is 0.95 or 95% confidence.
sigdig	Optimization significant digits. This controls: <ul style="list-style-type: none"> • The tolerance of the inner and outer optimization is $10^{-\text{sigdig}}$ • The tolerance of the ODE solvers is $0.5 \times 10^{-(\text{sigdig}-2)}$; For the sensitivity equations and steady-state solutions the default is $0.5 \times 10^{-(\text{sigdig}-1.5)}$ (sensitivity changes only applicable for liblsoda) • The tolerance of the boundary check is $5 \times 10^{-(\text{sigdig} + 1)}$
sigdigTable	Significant digits in the final output table. If not specified, then it matches the significant digits in the 'sigdig' optimization algorithm. If 'sigdig' is NULL, use 3.
...	Additional arguments passed to <code>nlmixr2est::nlsControl()</code> .

Value

nls control object

Author(s)

Matthew L. Fidler

Examples

```
if (rxode2:::linCmtSensB()) {
  one.cmt <- function() {
    ini({
      tka <- 0.45
      tc1 <- log(c(0, 2.7, 100))
    })
  }
}
```



```

    tv <- 3.45
    add.sd <- 0.7
  })
  model({
    ka <- exp(tka)
    cl <- exp(tcl)
    v <- exp(tv)
    linCmt() ~ add(add.sd)
  })
}

# Uses nlsLM from minpack.lm if available

fit1 <- nlmixr(one.cmt, nlmixr2data::theo_sd, est="nls", nlsControl(algorithm="LM"))

# Uses port and respect parameter boundaries
fit2 <- nlmixr(one.cmt, nlmixr2data::theo_sd, est="nls", nlsControl(algorithm="port"))

# You can access the underlying nls object with `fit2$nls`
fit2$nls
}

```

optimControl

nlmixr2 optim defaults

Description

nlmixr2 optim defaults

Usage

```

optimControl(
  method = c("Nelder-Mead", "BFGS", "CG", "L-BFGS-B", "SANN", "Brent"),
  trace = 0,
  fnscale = 1,
  parscale = 1,
  ndeps = 0.001,
  maxit = 10000,
  abstol = 1e-08,
  reltol = 1e-08,
  alpha = 1,
  beta = 0.5,
  gamma = 2,
  REPORT = NULL,
  warn.1d.NelderMead = TRUE,
  type = NULL,
  lmm = 5,
  factr = 1e+07,

```

```

pgtol = 0,
temp = 10,
tmax = 10,
stickyRecalcN = 4,
maxOdeRecalc = 5,
odeRecalcFactor = 10^(0.5),
eventType = c("central", "forward"),
shiErr = (.Machine$double.eps)^(1/3),
shi21maxFD = 20L,
solveType = c("grad", "fun"),
useColor = crayon::has_color(),
printNcol = floor((getOption("width") - 23)/12),
print = 1L,
normType = c("rescale2", "mean", "rescale", "std", "len", "constant"),
scaleType = c("nlmixr2", "norm", "mult", "multAdd"),
scaleCmax = 1e+05,
scaleCmin = 1e-05,
scaleC = NULL,
scaleTo = 1,
gradTo = 1,
rxControl = NULL,
optExpression = TRUE,
sumProd = FALSE,
literalFix = TRUE,
returnOptim = FALSE,
addProp = c("combined2", "combined1"),
calcTables = TRUE,
compress = TRUE,
covMethod = c("r", "optim", ""),
adjObf = TRUE,
ci = 0.95,
sigdig = 4,
sigdigTable = NULL,
...
)

```

Arguments

method	The method to be used. See ‘Details’. Can be abbreviated.
trace	Non-negative integer. If positive, tracing information on the progress of the optimization is produced. Higher values may produce more tracing information: for method “L-BFGS-B”, there are six levels of tracing. See ‘optim()’ for more information
fnscale	An overall scaling to be applied to the value of ‘fn’ and ‘gr’ during optimization. If negative, turns the problem into a maximization problem. Optimization is performed on ‘fn(par)/fnscale’
parscale	A vector of scaling values for the parameters. Optimization is performed on ‘par/parscale’ and these should be comparable in the sense that a unit change

	in any element produces about a unit change in the scaled value. Not used (nor needed) for <code>method = "Brent"</code>
<code>ndeps</code>	A vector of step sizes for the finite-difference approximation to the gradient, on <code>'par/parscale'</code> scale. Defaults to <code>'1e-3'</code>
<code>maxit</code>	The maximum number of iterations. Defaults to <code>'100'</code> for the derivative-based methods, and <code>'500'</code> for <code>"Nelder-Mead"</code> .
<code>abstol</code>	The absolute convergence tolerance. Only useful for non-negative functions, as a tolerance for reaching zero.
<code>reltol</code>	Relative convergence tolerance. The algorithm stops if it is unable to reduce the value by a factor of <code>'reltol * (abs(val) + reltol)'</code> at a step
<code>alpha</code>	Reflection factor for the <code>"Nelder-Mead"</code> method.
<code>beta</code>	Contraction factor for the <code>"Nelder-Mead"</code> method
<code>gamma</code>	Expansion factor for the <code>"Nelder-Mead"</code> method
<code>REPORT</code>	The frequency of reports for the <code>"BFGS"</code> , <code>"L-BFGS-B"</code> and <code>"SANN"</code> methods if <code>'control\$trace'</code> is positive. Defaults to every 10 iterations for <code>"BFGS"</code> and <code>"L-BFGS-B"</code> , or every 100 temperatures for <code>"SANN"</code>
<code>warn.1d.NelderMead</code>	a logical indicating if the (default) <code>"Nelder-Mead"</code> method should signal a warning when used for one-dimensional minimization. As the warning is sometimes inappropriate, you can suppress it by setting this option to <code>'FALSE'</code>
<code>type</code>	for the conjugate-gradients method. Takes value <code>'1'</code> for the Fletcher-Reeves update, <code>'2'</code> for Polak-Ribiere and <code>'3'</code> for Beale-Sorenson.
<code>lmm</code>	is an integer giving the number of BFGS updates retained in the <code>"L-BFGS-B"</code> method, It defaults to <code>'5'</code>
<code>factr</code>	controls the convergence of the <code>"L-BFGS-B"</code> method. Convergence occurs when the reduction in the objective is within this factor of the machine tolerance. Default is <code>'1e7'</code> , that is a tolerance of about <code>'1e-8'</code> .
<code>pgtol</code>	helps control the convergence of the <code>"L-BFGS-B"</code> method. It is a tolerance on the projected gradient in the current search direction. This defaults to zero, when the check is suppressed
<code>temp</code>	controls the <code>"SANN"</code> method. It is the starting temperature for the cooling schedule. Defaults to <code>'10'</code> .
<code>tmax</code>	is the number of function evaluations at each temperature for the <code>"SANN"</code> method. Defaults to <code>'10'</code> .
<code>stickyRecalcN</code>	The number of bad ODE solves before reducing the <code>atol/rtol</code> for the rest of the problem.
<code>maxOdeRecalc</code>	Maximum number of times to reduce the ODE tolerances and try to resolve the system if there was a bad ODE solve.
<code>odeRecalcFactor</code>	The ODE recalculation factor when ODE solving goes bad, this is the factor the <code>rtol/atol</code> is reduced
<code>eventType</code>	Event gradient type for dosing events; Can be <code>"central"</code> or <code>"forward"</code>

shiErr	This represents the epsilon when optimizing the ideal step size for numeric differentiation using the Shi2021 method
shi21maxFD	The maximum number of steps for the optimization of the forward difference step size when using dosing events (lag time, modeled duration/rate and bioavailability)
solveType	<p>tells if 'optim' will use nlmixr2's analytical gradients when available (finite differences will be used for event-related parameters like parameters controlling lag time, duration/rate of infusion, and modeled bioavailability). This can be:</p> <ul style="list-style-type: none"> - "gradient" which will use the gradient and let 'optim' calculate the finite difference hessian - "fun" where optim will calculate both the finite difference gradient and the finite difference Hessian <p>When using nlmixr2's finite differences, the "ideal" step size for either central or forward differences are optimized for with the Shi2021 method which may give more accurate derivatives</p> <p>These are only applied in the gradient based methods: "BFGS", "CG", "L-BFGS-B"</p>
useColor	Boolean indicating if focei can use ASCII color codes
printNcol	Number of columns to printout before wrapping parameter estimates/gradient
print	Integer representing when the outer step is printed. When this is 0 or do not print the iterations. 1 is print every function evaluation (default), 5 is print every 5 evaluations.
normType	<p>This is the type of parameter normalization/scaling used to get the scaled initial values for nlmixr2. These are used with scaleType of.</p> <p>With the exception of rescale2, these come from Feature Scaling. The rescale2 The rescaling is the same type described in the OptdesX software manual.</p> <p>In general, all all scaling formula can be described by:</p>

$$v_{scaled} = \frac{v_{unscaled} - C_1}{C_2}$$

Where

The other data normalization approaches follow the following formula

$$v_{scaled} = \frac{v_{unscaled} - C_1}{C_2}$$

- `rescale2` This scales all parameters from (-1 to 1). The relative differences between the parameters are preserved with this approach and the constants are:

$$C_1$$

$$= (\max(\text{all unscaled values}) + \min(\text{all unscaled values})) / 2$$

$$C_2$$

$$= (\max(\text{all unscaled values}) - \min(\text{all unscaled values})) / 2$$

- `rescale` or min-max normalization. This rescales all parameters from (0 to 1). As in the `rescale2` the relative differences are preserved. In this approach:

$$C_1$$

$$= \min(\text{all unscaled values})$$

$$C_2$$

$$= \max(\text{all unscaled values}) - \min(\text{all unscaled values})$$

- `mean` or mean normalization. This rescales to center the parameters around the mean but the parameters are from 0 to 1. In this approach:

$$C_1$$

$$= \text{mean}(\text{all unscaled values})$$

$$C_2$$

$$= \max(\text{all unscaled values}) - \min(\text{all unscaled values})$$

- `std` or standardization. This standardizes by the mean and standard deviation. In this approach:

$$C_1$$

$$= \text{mean}(\text{all unscaled values})$$

$$C_2$$

$$= \text{sd}(\text{all unscaled values})$$

- `len` or unit length scaling. This scales the parameters to the unit length. For this approach we use the Euclidean length, that is:

$$C_1$$

$$= 0$$

$$C_2$$

$$= \sqrt{v_1^2 + v_2^2 + \dots + v_n^2}$$

- constant which does not perform data normalization. That is

$$= 0 \quad C_1$$

$$= 1 \quad C_2$$

scaleType

The scaling scheme for nlmixr2. The supported types are:

- nlmixr2 In this approach the scaling is performed by the following equation:

$$= (\quad v_{scaled}$$

$$\quad v_{current} - v_{init}$$

$$) * scaleC[i] + scaleTo$$

The scaleTo parameter is specified by the normType, and the scales are specified by scaleC.

- norm This approach uses the simple scaling provided by the normType argument.
- mult This approach does not use the data normalization provided by normType, but rather uses multiplicative scaling to a constant provided by the scaleTo argument.

In this case:

$$= \quad v_{scaled}$$

$$\quad v_{current}$$

$$/ \quad v_{init}$$

$$* scaleTo$$

- multAdd This approach changes the scaling based on the parameter being specified. If a parameter is defined in an exponential block (ie exp(theta)), then it is scaled on a linearly, that is:

$$= (\quad v_{scaled}$$

$$\quad v_{current} - v_{init}$$

$$) + scaleTo$$

Otherwise the parameter is scaled multiplicatively.

$$= \quad v_{scaled}$$

$$\quad v_{current}$$

	$/$
	v_{init}
	*scaleTo
scaleCmax	Maximum value of the scaleC to prevent overflow.
scaleCmin	Minimum value of the scaleC to prevent underflow.
scaleC	<p>The scaling constant used with scaleType=nlmixr2. When not specified, it is based on the type of parameter that is estimated. The idea is to keep the derivatives similar on a log scale to have similar gradient sizes. Hence parameters like log(exp(theta)) would have a scaling factor of 1 and log(theta) would have a scaling factor of ini_value (to scale by 1/value; ie d/dt(log(ini_value)) = 1/ini_value or scaleC=ini_value)</p> <ul style="list-style-type: none"> • For parameters in an exponential (ie exp(theta)) or parameters specifying powers, boxCox or yeoJohnson transformations, this is 1. • For additive, proportional, lognormal error structures, these are given by 0.5*abs(initial_estimate) • Factorials are scaled by abs(1/digamma(initial_estimate+1)) • parameters in a log scale (ie log(theta)) are transformed by log(abs(initial_estimate))*abs(initial_estim <p>These parameter scaling coefficients are chose to try to keep similar slopes among parameters. That is they all follow the slopes approximately on a log-scale.</p> <p>While these are chosen in a logical manner, they may not always apply. You can specify each parameters scaling factor by this parameter if you wish.</p>
scaleTo	Scale the initial parameter estimate to this value. By default this is 1. When zero or below, no scaling is performed.
gradTo	this is the factor that the gradient is scaled to before optimizing. This only works with scaleType="nlmixr2".
rxControl	'rxode2' ODE solving options during fitting, created with 'rxControl()'
optExpression	Optimize the rxode2 expression to speed up calculation. By default this is turned on.
sumProd	Is a boolean indicating if the model should change multiplication to high precision multiplication and sums to high precision sums using the PreciseSums package. By default this is FALSE.
literalFix	boolean, substitute fixed population values as literals and re-adjust ui and parameter estimates after optimization; Default is 'TRUE'.
returnOptim	logical; when TRUE this will return the optim list instead of the nlmixr2 fit object
addProp	<p>specifies the type of additive plus proportional errors, the one where standard deviations add (combined1) or the type where the variances add (combined2). The combined1 error type can be described by the following equation:</p>

$$y = f + (a + b \times f^c) \times \varepsilon$$

The combined2 error model can be described by the following equation:

$$y = f + \sqrt{a^2 + b^2 \times f^{2 \times c}} \times \varepsilon$$

Where:

- y represents the observed value
- f represents the predicted value
- a is the additive standard deviation
- b is the proportional/power standard deviation
- c is the power exponent (in the proportional case c=1)

calcTables	This boolean is to determine if the foceiFit will calculate tables. By default this is TRUE
compress	Should the object have compressed items
covMethod	allows selection of "r", which uses nlmixr2's 'nlmixr2Hess()' for the hessian calculation or "optim" which uses the hessian from 'stats::optim(..., hessian=TRUE)'
adj0bf	is a boolean to indicate if the objective function should be adjusted to be closer to NONMEM's default objective function. By default this is TRUE
ci	Confidence level for some tables. By default this is 0.95 or 95% confidence.
sigdig	Optimization significant digits. This controls: <ul style="list-style-type: none"> • The tolerance of the inner and outer optimization is $10^{-\text{sigdig}}$ • The tolerance of the ODE solvers is $0.5 \times 10^{-(\text{sigdig}-2)}$; For the sensitivity equations and steady-state solutions the default is $0.5 \times 10^{-(\text{sigdig}-1.5)}$ (sensitivity changes only applicable for liblsoda) • The tolerance of the boundary check is $5 \times 10^{-(\text{sigdig} + 1)}$
sigdigTable	Significant digits in the final output table. If not specified, then it matches the significant digits in the 'sigdig' optimization algorithm. If 'sigdig' is NULL, use 3.
...	Additional arguments passed to <code>nlmixr2est::optimControl()</code> .

Value

optimControl object for nlmixr2

Author(s)

Matthew L. Fidler

Examples

```
# A logit regression example with emax model

dsn <- data.frame(i=1:1000)
dsn$time <- exp(rnorm(1000))
dsn$DV=rbinom(1000,1,exp(-1+dsn$time)/(1+exp(-1+dsn$time)))

mod <- function() {
  ini({
```



```

    E0 <- 0.5
    Em <- 0.5
    E50 <- 2
    g <- fix(2)
  })
  model({
    v <- E0+Em*time^g/(E50^g+time^g)
    ll(bin) ~ DV * v - log(1 + exp(v))
  })
}

fit2 <- nlmixr(mod, dsn, est="optim", optimControl(method="BFGS"))
fit2

```

preconditionFit	<i>Linearly re-parameterize the model to be less sensitive to rounding errors</i>
-----------------	---

Description

Linearly re-parameterize the model to be less sensitive to rounding errors

Usage

```
preconditionFit(fit, estType = c("full", "posthoc", "none"), ntry = 10L)
```

Arguments

fit	A nlmixr2 fit to be preconditioned
estType	Once the fit has been linearly reparameterized, should a "full" estimation, "posthoc" estimation or simply a estimation of the covariance matrix "none" before the fit is updated
ntry	number of tries before giving up on a pre-conditioned covariance estimate

Value

A nlmixr2 fit object that was preconditioned to stabilize the variance/covariance calculation

References

Aoki Y, Nordgren R, Hooker AC. Preconditioning of Nonlinear Mixed Effects Models for Stabilisation of Variance-Covariance Matrix Computations. AAPS J. 2016;18(2):505-518. doi:10.1208/s12248-016-9866-5

profileFixed	<i>Estimate the objective function values for a model while fixing defined parameter values</i>
--------------	---

Description

Estimate the objective function values for a model while fixing defined parameter values

Usage

```
profileFixed(fitted, which, control = list())
```

Arguments

fitted	The fit model
which	A data.frame with column names of parameters to fix and values of the fitted value to fix (one row only).
control	A list passed to fixedControl() (currently unused)

Value

which with a column named OFV added with the objective function value of the fitted estimate fixing the parameters in the other columns

Functions

- profileFixedSingle(): Estimate the objective function value for a model while fixing a single set of defined parameter values (for use in parameter profiling)

See Also

Other Profiling: [fixedControl\(\)](#), [llpControl\(\)](#), [profile.nlmixr2FitCore\(\)](#), [profileLlp\(\)](#), [profileNlmixr2FitCoreRet\(\)](#)

Other Profiling: [fixedControl\(\)](#), [llpControl\(\)](#), [profile.nlmixr2FitCore\(\)](#), [profileLlp\(\)](#), [profileNlmixr2FitCoreRet\(\)](#)

profileFixedSingle	<i>Estimate the objective function values for a model while fixing defined parameter values</i>
--------------------	---

Description

Estimate the objective function values for a model while fixing defined parameter values

Usage

```
profileFixedSingle(fitted, which)
```

Arguments

fitted	The fit model
which	A data.frame with column names of parameters to fix and values of the fitted value to fix (one row only).

Value

which with a column named OFV added with the objective function value of the fitted estimate fixing the parameters in the other columns

Functions

- `profileFixedSingle()`: Estimate the objective function value for a model while fixing a single set of defined parameter values (for use in parameter profiling)

See Also

Other Profiling: [fixedControl\(\)](#), [llpControl\(\)](#), [profile.nlmixr2FitCore\(\)](#), [profileLlp\(\)](#), [profileNlmixr2FitCoreRet\(\)](#)

Other Profiling: [fixedControl\(\)](#), [llpControl\(\)](#), [profile.nlmixr2FitCore\(\)](#), [profileLlp\(\)](#), [profileNlmixr2FitCoreRet\(\)](#)

profileLlp	<i>Profile confidence intervals with log-likelihood profiling</i>
------------	---

Description

Profile confidence intervals with log-likelihood profiling

Usage

```
profileLlp(fitted, which, control)
```

Arguments

fitted	The fit model
which	Either NULL to profile all parameters or a character vector of parameters to estimate
control	A list passed to llpControl()

Value

A data.frame with columns named "Parameter" (the parameter name(s) that were fixed), OFV (the objective function value), and the current estimate for each of the parameters. In addition, if any boundary is found, the OFV increase will be indicated by the absolute value of the "profileBound" column and if that boundary is the upper or lower boundary will be indicated by the "profileBound" column being positive or negative, respectively.

See Also

Other Profiling: [fixedControl\(\)](#), [llpControl\(\)](#), [profile.nlmixr2FitCore\(\)](#), [profileFixed\(\)](#), [profileNlmixr2FitCoreRet\(\)](#)

 saemControl

Control Options for SAEM

Description

Control Options for SAEM

Usage

```
saemControl(
  seed = 99,
  nBurn = 200,
  nEm = 300,
  nmc = 3,
  nu = c(2, 2, 2),
  print = 1,
  trace = 0,
  covMethod = c("linFim", "fim", "r,s", "r", "s", ""),
  calcTables = TRUE,
  logLik = FALSE,
  nnodesGq = 3,
  nsdGq = 1.6,
  optExpression = TRUE,
  literalFix = FALSE,
  adjObf = TRUE,
  sumProd = FALSE,
  addProp = c("combined2", "combined1"),
```

```

    tol = 1e-06,
    itmax = 30,
    type = c("nelder-mead", "newuoa"),
    powRange = 10,
    lambdaRange = 3,
    odeRecalcFactor = 10^(0.5),
    maxOdeRecalc = 5L,
    perSa = 0.75,
    perNoCor = 0.75,
    perFixOmega = 0.1,
    perFixResid = 0.1,
    compress = TRUE,
    rxControl = NULL,
    sigdig = NULL,
    sigdigTable = NULL,
    ci = 0.95,
    muRefCov = TRUE,
    muRefCovAlg = TRUE,
    handleUninformativeEtas = TRUE,
    ...
)

```

Arguments

seed	Random Seed for SAEM step. (Needs to be set for reproducibility.) By default this is 99.
nBurn	Number of iterations in the first phase, ie the MCMC/Stochastic Approximation steps. This is equivalent to Monolix's K_0 or K_b .
nEm	Number of iterations in the Expectation-Maximization (EM) Step. This is equivalent to Monolix's K_1 .
nmc	Number of Markov Chains. By default this is 3. When you increase the number of chains the numerical integration by MC method will be more accurate at the cost of more computation. In Monolix this is equivalent to L.
nu	<p>This is a vector of 3 integers. They represent the numbers of transitions of the three different kernels used in the Hasting-Metropolis algorithm. The default value is $c(2, 2, 2)$, representing 40 for each transition initially (each value is multiplied by 20).</p> <p>The first value represents the initial number of multi-variate Gibbs samples are taken from a normal distribution.</p> <p>The second value represents the number of uni-variate, or multi- dimensional random walk Gibbs samples are taken.</p> <p>The third value represents the number of bootstrap/reshuffling or uni-dimensional random samples are taken.</p>
print	The number it iterations that are completed before anything is printed to the console. By default, this is 1.
trace	An integer indicating if you want to trace(1) the SAEM algorithm process. Useful for debugging, but not for typical fitting.

covMethod	<p>Method for calculating covariance. In this discussion, R is the Hessian matrix of the objective function. The S matrix is the sum of each individual's gradient cross-product (evaluated at the individual empirical Bayes estimates).</p> <p>"linFim" Use the Linearized Fisher Information Matrix to calculate the covariance.</p> <p>"fim" Use the SAEM-calculated Fisher Information Matrix to calculate the covariance.</p> <p>"r, s" Uses the sandwich matrix to calculate the covariance, that is: $R^{-1} \times S \times R^{-1}$</p> <p>"r" Uses the Hessian matrix to calculate the covariance as $2 \times R^{-1}$</p> <p>"s" Uses the crossproduct matrix to calculate the covariance as $4 \times S^{-1}$</p> <p>"" Does not calculate the covariance step.</p>
calcTables	This boolean is to determine if the foceiFit will calculate tables. By default this is TRUE
logLik	boolean indicating that log-likelihood should be calculate by Gaussian quadrature.
nnodesGq	number of nodes to use for the Gaussian quadrature when computing the likelihood with this method (defaults to 1, equivalent to the Laplacian likelihood)
nsdGq	span (in SD) over which to integrate when computing the likelihood by Gaussian quadrature. Defaults to 3 (eg 3 times the SD)
optExpression	Optimize the rxode2 expression to speed up calculation. By default this is turned on.
literalFix	boolean, substitute fixed population values as literals and re-adjust ui and parameter estimates after optimization; Default is 'TRUE'.
adj0bf	is a boolean to indicate if the objective function should be adjusted to be closer to NONMEM's default objective function. By default this is TRUE
sumProd	Is a boolean indicating if the model should change multiplication to high precision multiplication and sums to high precision sums using the PreciseSums package. By default this is FALSE.
addProp	specifies the type of additive plus proportional errors, the one where standard deviations add (combined1) or the type where the variances add (combined2). The combined1 error type can be described by the following equation:

$$y = f + (a + b \times f^c) \times \varepsilon$$

The combined2 error model can be described by the following equation:

$$y = f + \sqrt{a^2 + b^2 \times f^{2 \times c}} \times \varepsilon$$

Where:

- y represents the observed value
- f represents the predicted value
- a is the additive standard deviation
- b is the proportional/power standard deviation
- c is the power exponent (in the proportional case c=1)

tol	This is the tolerance for the regression models used for complex residual errors (ie add+prop etc)
itmax	This is the maximum number of iterations for the regression models used for complex residual errors. The number of iterations is itmax*number of parameters
type	indicates the type of optimization for the residuals; Can be one of c("nelder-mead", "newuoa")
powRange	This indicates the range that powers can take for residual errors; By default this is 10 indicating the range is c(-10, 10)
lambdaRange	This indicates the range that Box-Cox and Yeo-Johnson parameters are constrained to be; The default is 3 indicating the range c(-3,3)
odeRecalcFactor	The ODE recalculation factor when ODE solving goes bad, this is the factor the rtol/atol is reduced
maxOdeRecalc	Maximum number of times to reduce the ODE tolerances and try to resolve the system if there was a bad ODE solve.
perSa	This is the percent of the time the 'nBurn' iterations in phase runs runs a simulated annealing.
perNoCor	This is the percentage of the MCMC phase of the SAEM algorithm where the variance/covariance matrix has no correlations. By default this is 0.75 or 75 Monte-carlo iteration.
perFixOmega	This is the percentage of the 'nBurn' phase where the omega values are unfixed to allow better exploration of the likelihood surface. After this time, the omegas are fixed during optimization.
perFixResid	This is the percentage of the 'nBurn' phase where the residual components are unfixed to allow better exploration of the likelihood surface.
compress	Should the object have compressed items
rxControl	'rxode2' ODE solving options during fitting, created with 'rxControl()'
sigdig	Specifies the "significant digits" that the ode solving requests. When specified this controls the relative and absolute tolerances of the ODE solvers. By default the tolerance is $0.5 \times 10^{-(\text{sigdig}-2)}$ for regular ODEs. For the sensitivity equations the default is $0.5 \times 10^{-(\text{sigdig}-1.5)}$ (sensitivity changes only applicable for liblsoda). This also controls the atol/rtol of the steady state solutions. The ssAtol/ssRtol is $0.5 \times 10^{-(\text{sigdig})}$ and for the sensitivities $0.5 \times 10^{-(\text{sigdig}+0.625)}$. By default this is unspecified (NULL) and uses the standard atol/rtol.
sigdigTable	Significant digits in the final output table. If not specified, then it matches the significant digits in the 'sigdig' optimization algorithm. If 'sigdig' is NULL, use 3.
ci	Confidence level for some tables. By default this is 0.95 or 95% confidence.
muRefCov	This controls if mu-referenced covariates in 'saem' are handled differently than non mu-referenced covariates. When 'TRUE', mu-referenced covariates have special handling. When 'FALSE' mu-referenced covariates are treated the same as any other input parameter.

muRefCovAlg	<p>This controls if algebraic expressions that can be mu-referenced are treated as mu-referenced covariates by:</p> <ol style="list-style-type: none"> 1. Creating a internal data-variable ‘nlmixrMuDerCov#’ for each algebraic mu-referenced expression 2. Change the algebraic expression to ‘nlmixrMuDerCov# * mu_cov_theta’ 3. Use the internal mu-referenced covariate for saem 4. After optimization is completed, replace ‘model()’ with old ‘model()’ expression 5. Remove ‘nlmixrMuDerCov#’ from nlmix2 output <p>In general, these covariates should be more accurate since it changes the system to a linear compartment model. Therefore, by default this is ‘TRUE’.</p>
handleUninformativeEtas	boolean that tells nlmixr2’s saem to calculate uninformative etas and handle them specially (default is ‘TRUE’).
...	Additional arguments passed to <code>nlmixr2est::saemControl()</code> .

Value

List of options to be used in `nlmixr2` fit for SAEM.

Author(s)

Wenping Wang & Matthew L. Fidler

See Also

Other Estimation control: `foceiControl()`, `nlmixr2NlmeControl()`

setOfv

Set/get Objective function type for a nlmixr2 object

Description

Set/get Objective function type for a nlmixr2 object

Usage

```
setOfv(x, type)
```

Arguments

x	nlmixr2 fit object
type	Type of objective function to use for AIC, BIC, and \$objective

Value

Nothing

Author(s)

Matthew L. Fidler

tableControl	<i>Output table/data.frame options</i>
--------------	--

Description

Output table/data.frame options

Usage

```
tableControl(
  npde = NULL,
  cwres = NULL,
  nsim = 300,
  ties = TRUE,
  censMethod = c("truncated-normal", "cdf", "ipred", "pred", "epred", "omit"),
  seed = 1009,
  cholSEtol = (.Machine$double.eps)^(1/3),
  state = TRUE,
  lhs = TRUE,
  eta = TRUE,
  covariates = TRUE,
  addDosing = FALSE,
  subsetNonmem = TRUE,
  cores = NULL,
  keep = NULL,
  drop = NULL
)
```

Arguments

npde	When TRUE, request npde regardless of the algorithm used.
cwres	When TRUE, request CWRES and FOCEi likelihood regardless of the algorithm used.
nsim	represents the number of simulations. For rxode2, if you supply single subject event tables (created with [eventTable()])
ties	When 'TRUE' jitter prediction-discrepancy points to discourage ties in cdf.
censMethod	Handle censoring method: <ul style="list-style-type: none"> - "truncated-normal" Simulates from a truncated normal distribution under the assumption of the model and censoring. - "cdf" Use the cdf-method for censoring with npde and use this for any other residuals ('cwres' etc) - "omit" omit the residuals for censoring

seed	an object specifying if and how the random number generator should be initialized
cholSEtol	The tolerance for the 'rxode2::choleSE' function
state	is a Boolean indicating if 'state' values will be included (default 'TRUE')
lhs	is a Boolean indicating if remaining 'lhs' values will be included (default 'TRUE')
eta	is a Boolean indicating if 'eta' values will be included (default 'TRUE')
covariates	is a Boolean indicating if covariates will be included (default 'TRUE')
addDosing	<p>Boolean indicating if the solve should add rxode2 EVID and related columns. This will also include dosing information and estimates at the doses. Be default, rxode2 only includes estimates at the observations. (default FALSE). When addDosing is NULL, only include EVID=0 on solve and exclude any model-times or EVID=2. If addDosing is NA the classic rxode2 EVID events are returned. When addDosing is TRUE add the event information in NONMEM-style format; If subsetNonmem=FALSE rxode2 will also include extra event types (EVID) for ending infusion and modeled times:</p> <ul style="list-style-type: none"> • EVID=-1 when the modeled rate infusions are turned off (matches rate=-1) • EVID=-2 When the modeled duration infusions are turned off (matches rate=-2) • EVID=-10 When the specified rate infusions are turned off (matches rate>0) • EVID=-20 When the specified dur infusions are turned off (matches dur>0) • EVID=101, 102, 103, . . . Modeled time where 101 is the first model time, 102 is the second etc.
subsetNonmem	subset to NONMEM compatible EVIDs only. By default TRUE.
cores	Number of cores used in parallel ODE solving. This is equivalent to calling setRxThreads()
keep	is the keep sent to the table
drop	is the dropped variables sent to the table

Details

If you ever want to add CWRES/FOCEi objective function you can use the [addCwres](#)

If you ever want to add NPDE/EPRED columns you can use the [addNpde](#)

Value

A list of table options for nlmixr2

Author(s)

Matthew L. Fidler

traceplot	<i>Produce trace-plot for fit if applicable</i>
-----------	---

Description

Produce trace-plot for fit if applicable

Usage

```
traceplot(x, ...)
```

Arguments

x	fit object
...	Additional arguments passed to <code>nlmixr2plot::traceplot()</code> .

Value

Fit traceplot or nothing.

Author(s)

Rik Schoemaker, Wenping Wang & Matthew L. Fidler

Examples

```
library(nlmixr2est)
## The basic model consists of an ini block that has initial estimates
one.compartment <- function() {
  ini({
    tka <- 0.45 # Log Ka
    tc1 <- 1 # Log Cl
    tv <- 3.45 # Log V
    eta.ka ~ 0.6
    eta.cl ~ 0.3
    eta.v ~ 0.1
    add.sd <- 0.7
  })
  # and a model block with the error specification and model specification
  model({
    ka <- exp(tka + eta.ka)
    cl <- exp(tc1 + eta.cl)
    v <- exp(tv + eta.v)
    d/dt(depot) = -ka * depot
    d/dt(center) = ka * depot - cl / v * center
    cp = center / v
    cp ~ add(add.sd)
  })
}
```

```

}

## The fit is performed by the function nlmixr/nlmix2 specifying the model, data and estimate
fit <- nlmixr2(one.compartment, theo_sd, est="saem", saemControl(print=0))

# This shows the traceplot of the fit (useful for saem)
traceplot(fit)

```

uobyqaControl

Control for uobyqa estimation method in nlmixr2

Description

Control for uobyqa estimation method in nlmixr2

Usage

```

uobyqaControl(
  npt = NULL,
  rhobeg = NULL,
  rhoend = NULL,
  iprint = 0L,
  maxfun = 100000L,
  returnUobyqa = FALSE,
  stickyRecalcN = 4,
  maxOdeRecalc = 5,
  odeRecalcFactor = 10^(0.5),
  useColor = crayon::has_color(),
  printNcol = floor((getOption("width") - 23)/12),
  print = 1L,
  normType = c("rescale2", "mean", "rescale", "std", "len", "constant"),
  scaleType = c("nlmixr2", "norm", "mult", "multAdd"),
  scaleCmax = 1e+05,
  scaleCmin = 1e-05,
  scaleC = NULL,
  scaleTo = 1,
  rxControl = NULL,
  optExpression = TRUE,
  sumProd = FALSE,
  literalFix = TRUE,
  addProp = c("combined2", "combined1"),
  calcTables = TRUE,
  compress = TRUE,
  covMethod = c("r", ""),
  adjObf = TRUE,
  ci = 0.95,

```

```

    sigdig = 4,
    sigdigTable = NULL,
    ...
)

```

Arguments

npt	The number of points used to approximate the objective function via a quadratic approximation for bobyqa. The value of npt must be in the interval $[n+2, (n+1)(n+2)/2]$ where n is the number of parameters in par. Choices that exceed $2*n+1$ are not recommended. If not defined, it will be set to $2*n + 1$. (bobyqa)
rhobeg	Beginning change in parameters for bobyqa algorithm (trust region). By default this is 0.2 or 20 parameters when the parameters are scaled to 1. rhobeg and rhoend must be set to the initial and final values of a trust region radius, so both must be positive with $0 < \text{rhoend} < \text{rhobeg}$. Typically rhobeg should be about one tenth of the greatest expected change to a variable. Note also that smallest difference $\text{abs}(\text{upper-lower})$ should be greater than or equal to rhobeg^2 . If this is not the case then rhobeg will be adjusted. (bobyqa)
rhoend	The smallest value of the trust region radius that is allowed. If not defined, then $10^{-(\text{sigdig}-1)}$ will be used. (bobyqa)
iprint	The value of 'iprint' should be set to an integer value in '0, 1, 2, 3, ...', which controls the amount of printing. Specifically, there is no output if 'iprint=0' and there is output only at the start and the return if 'iprint=1'. Otherwise, each new value of 'rho' is printed, with the best vector of variables so far and the corresponding value of the objective function. Further, each new value of the objective function with its variables are output if 'iprint=3'. If 'iprint > 3', the objective function value and corresponding variables are output every 'iprint' evaluations. Default value is '0'.
maxfun	The maximum allowed number of function evaluations. If this is exceeded, the method will terminate.
returnUobyqa	return the uobyqa output instead of the nlmixr2 fit
stickyRecalcN	The number of bad ODE solves before reducing the atol/rtol for the rest of the problem.
maxOdeRecalc	Maximum number of times to reduce the ODE tolerances and try to resolve the system if there was a bad ODE solve.
odeRecalcFactor	The ODE recalculation factor when ODE solving goes bad, this is the factor the rtol/atol is reduced
useColor	Boolean indicating if foci can use ASCII color codes
printNcol	Number of columns to printout before wrapping parameter estimates/gradient
print	Integer representing when the outer step is printed. When this is 0 or do not print the iterations. 1 is print every function evaluation (default), 5 is print every 5 evaluations.
normType	This is the type of parameter normalization/scaling used to get the scaled initial values for nlmixr2. These are used with scaleType of.

With the exception of rescale2, these come from **Feature Scaling**. The rescale2 The rescaling is the same type described in the **OptdesX** software manual. In general, all all scaling formula can be described by:

$$= \left(\frac{v_{scaled}}{v_{unscaled} - C_1} \right) / C_2$$

Where

The other data normalization approaches follow the following formula

$$= \left(\frac{v_{scaled}}{v_{unscaled} - C_1} \right) / C_2$$

- rescale2 This scales all parameters from (-1 to 1). The relative differences between the parameters are preserved with this approach and the constants are:

$$C_1 = (\max(\text{all unscaled values}) + \min(\text{all unscaled values})) / 2$$

$$C_2 = (\max(\text{all unscaled values}) - \min(\text{all unscaled values})) / 2$$

- rescale or min-max normalization. This rescales all parameters from (0 to 1). As in the rescale2 the relative differences are preserved. In this approach:

$$C_1 = \min(\text{all unscaled values})$$

$$C_2 = \max(\text{all unscaled values}) - \min(\text{all unscaled values})$$

- mean or mean normalization. This rescales to center the parameters around the mean but the parameters are from 0 to 1. In this approach:

$$C_1 = \text{mean}(\text{all unscaled values})$$

$$C_2 = \max(\text{all unscaled values}) - \min(\text{all unscaled values})$$

- **std** or standardization. This standardizes by the mean and standard deviation. In this approach:

$$C_1 = \text{mean}(\text{all unscaled values})$$

$$C_2 = \text{sd}(\text{all unscaled values})$$

- **len** or unit length scaling. This scales the parameters to the unit length. For this approach we use the Euclidean length, that is:

$$C_1 = 0$$

$$C_2 = \sqrt{(v_1^2 + v_2^2 + \dots + v_n^2)}$$

- **constant** which does not perform data normalization. That is

$$C_1 = 0$$

$$C_2 = 1$$

scaleType

The scaling scheme for nlmixr2. The supported types are:

- **nlmixr2** In this approach the scaling is performed by the following equation:

$$v_{scaled} = (v_{current} - v_{init}) * \text{scaleC}[i] + \text{scaleTo}$$

The `scaleTo` parameter is specified by the `normType`, and the scales are specified by `scaleC`.

- **norm** This approach uses the simple scaling provided by the `normType` argument.
- **mult** This approach does not use the data normalization provided by `normType`, but rather uses multiplicative scaling to a constant provided by the `scaleTo` argument.

In this case:

$$v_{scaled}$$

$$= \frac{v_{current}}{v_{init}} *scaleTo$$

- `multAdd` This approach changes the scaling based on the parameter being specified. If a parameter is defined in an exponential block (ie $\exp(\theta)$), then it is scaled on a linearly, that is:

$$= (\frac{v_{scaled}}{v_{current} - v_{init}}) + scaleTo$$

Otherwise the parameter is scaled multiplicatively.

$$= \frac{v_{scaled}}{v_{current}} / v_{init} *scaleTo$$

`scaleCmax`
`scaleCmin`
`scaleC`

Maximum value of the `scaleC` to prevent overflow.

Minimum value of the `scaleC` to prevent underflow.

The scaling constant used with `scaleType=nlmixr2`. When not specified, it is based on the type of parameter that is estimated. The idea is to keep the derivatives similar on a log scale to have similar gradient sizes. Hence parameters like $\log(\exp(\theta))$ would have a scaling factor of 1 and $\log(\theta)$ would have a scaling factor of `ini_value` (to scale by $1/value$; ie $d/dt(\log(ini_value)) = 1/ini_value$ or `scaleC=ini_value`)

- For parameters in an exponential (ie $\exp(\theta)$) or parameters specifying powers, `boxCox` or `yeoJohnson` transformations, this is 1.
- For additive, proportional, lognormal error structures, these are given by $0.5 * \text{abs}(\text{initial_estimate})$
- Factorials are scaled by $\text{abs}(1/\text{digamma}(\text{initial_estimate}+1))$
- parameters in a log scale (ie $\log(\theta)$) are transformed by $\log(\text{abs}(\text{initial_estimate})) * \text{abs}(\text{initial_estimate})$

These parameter scaling coefficients are chosen to try to keep similar slopes among parameters. That is they all follow the slopes approximately on a log-scale.

While these are chosen in a logical manner, they may not always apply. You can specify each parameters scaling factor by this parameter if you wish.

`scaleTo`

Scale the initial parameter estimate to this value. By default this is 1. When zero or below, no scaling is performed.

rxControl	'rxode2' ODE solving options during fitting, created with 'rxControl()'
optExpression	Optimize the rxode2 expression to speed up calculation. By default this is turned on.
sumProd	Is a boolean indicating if the model should change multiplication to high precision multiplication and sums to high precision sums using the PreciseSums package. By default this is FALSE.
literalFix	boolean, substitute fixed population values as literals and re-adjust ui and parameter estimates after optimization; Default is 'TRUE'.
addProp	specifies the type of additive plus proportional errors, the one where standard deviations add (combined1) or the type where the variances add (combined2). The combined1 error type can be described by the following equation:

$$y = f + (a + b \times f^c) \times \varepsilon$$

The combined2 error model can be described by the following equation:

$$y = f + \sqrt{a^2 + b^2 \times f^{2 \times c}} \times \varepsilon$$

Where:

- y represents the observed value
- f represents the predicted value
- a is the additive standard deviation
- b is the proportional/power standard deviation
- c is the power exponent (in the proportional case c=1)

calcTables	This boolean is to determine if the foceiFit will calculate tables. By default this is TRUE
compress	Should the object have compressed items
covMethod	Method for calculating covariance. In this discussion, R is the Hessian matrix of the objective function. The S matrix is the sum of individual gradient cross-product (evaluated at the individual empirical Bayes estimates). <ul style="list-style-type: none"> • "r, s" Uses the sandwich matrix to calculate the covariance, that is: solve(R) %*% S %*% solve(R) • "r" Uses the Hessian matrix to calculate the covariance as 2 %*% solve(R) • "s" Uses the cross-product matrix to calculate the covariance as 4 %*% solve(S) • "" Does not calculate the covariance step.
adjObf	is a boolean to indicate if the objective function should be adjusted to be closer to NONMEM's default objective function. By default this is TRUE
ci	Confidence level for some tables. By default this is 0.95 or 95% confidence.
sigdig	Optimization significant digits. This controls: <ul style="list-style-type: none"> • The tolerance of the inner and outer optimization is $10^{-\text{sigdig}}$ • The tolerance of the ODE solvers is $0.5 \times 10^{-(\text{sigdig}-2)}$; For the sensitivity equations and steady-state solutions the default is $0.5 \times 10^{-(\text{sigdig}-1.5)}$ (sensitivity changes only applicable for liblsoda)

- The tolerance of the boundary check is $5 * 10^{(-sigdig + 1)}$
- sigdigTable Significant digits in the final output table. If not specified, then it matches the significant digits in the 'sigdig' optimization algorithm. If 'sigdig' is NULL, use 3.
- ... Additional arguments passed to `nlmixr2est::uobyqaControl()`.

Value

uobyqa control structure

Author(s)

Matthew L. Fidler

Examples

```
# A logit regression example with emax model

dsn <- data.frame(i=1:1000)
dsn$time <- exp(rnorm(1000))
dsn$DV=rbinom(1000,1,exp(-1+dsn$time)/(1+exp(-1+dsn$time)))

mod <- function() {
  ini({
    E0 <- 0.5
    Em <- 0.5
    E50 <- 2
    g <- fix(2)
  })
  model({
    v <- E0+Em*time^g/(E50^g+time^g)
    ll(bin) ~ DV * v - log(1 + exp(v))
  })
}

fit2 <- nlmixr(mod, dsn, est="uobyqa")

print(fit2)

# you can also get the nlm output with fit2$nlm

fit2$uobyqa

# The nlm control has been modified slightly to include
# extra components and name the parameters
```

vpcCens	<i>VPC based on ui model</i>
---------	------------------------------

Description

VPC based on ui model

Usage

```
vpcCens(..., cens = TRUE, idv = "time")
```

Arguments

...	Additional arguments passed to <code>nlmixr2plot::vpcCens()</code> .
cens	is a boolean to show if this is a censoring plot or not. When <code>cens=TRUE</code> this is actually a censoring vpc plot (with <code>vpcCens()</code> and <code>vpcCensTad()</code>). When <code>cens=FALSE</code> this is traditional VPC plot (<code>vpcPlot()</code> and <code>vpcPlotTad()</code>).
idv	Name of independent variable. For <code>vpcPlot()</code> and <code>vpcCens()</code> the default is "time" for <code>vpcPlotTad()</code> and <code>vpcCensTad()</code> this is "tad"

Value

Simulated dataset (invisibly)

Author(s)

Matthew L. Fidler

Examples

```
one.cmt <- function() {  
  ini({  
    tka <- 0.45; label("Ka")  
    tcl <- log(c(0, 2.7, 100)); label("Cl")  
    tv <- 3.45; label("V")  
    eta.ka ~ 0.6  
    eta.cl ~ 0.3  
    eta.v ~ 0.1  
    add.sd <- 0.7; label("Additive residual error")  
  })  
  model({  
    ka <- exp(tka + eta.ka)  
    cl <- exp(tcl + eta.cl)  
    v <- exp(tv + eta.v)  
    linCmt() ~ add(add.sd)  
  })  
}
```

```

fit <-
  nlmixr2est::nlmixr(
    one.cmt,
    data = nlmixr2data::theo_sd,
    est = "saem",
    control = list(print = 0)
  )

vpcPlot(fit)

```

vpcCensTad

VPC based on ui model

Description

VPC based on ui model

Usage

```
vpcCensTad(..., cens = TRUE, idv = "tad")
```

Arguments

...	Additional arguments passed to <code>nlmixr2plot::vpcCensTad()</code> .
cens	is a boolean to show if this is a censoring plot or not. When <code>cens=TRUE</code> this is actually a censoring vpc plot (with <code>vpcCens()</code> and <code>vpcCensTad()</code>). When <code>cens=FALSE</code> this is traditional VPC plot (<code>vpcPlot()</code> and <code>vpcPlotTad()</code>).
idv	Name of independent variable. For <code>vpcPlot()</code> and <code>vpcCens()</code> the default is "time" for <code>vpcPlotTad()</code> and <code>vpcCensTad()</code> this is "tad"

Value

Simulated dataset (invisibly)

Author(s)

Matthew L. Fidler

Examples

```

one.cmt <- function() {
  ini({
    tka <- 0.45; label("Ka")
    tc1 <- log(c(0, 2.7, 100)); label("Cl")
    tv <- 3.45; label("V")
    eta.ka ~ 0.6
    eta.cl ~ 0.3
    eta.v ~ 0.1
  })
}

```

```
    add.sd <- 0.7; label("Additive residual error")
  })
  model({
    ka <- exp(tka + eta.ka)
    cl <- exp(tcl + eta.cl)
    v <- exp(tv + eta.v)
    linCmt() ~ add(add.sd)
  })
}

fit <-
  nlmixr2est::nlmixr(
    one.cmt,
    data = nlmixr2data::theo_sd,
    est = "saem",
    control = list(print = 0)
  )

vpcPlot(fit)
```

vpcPlot

VPC based on ui model

Description

VPC based on ui model

Usage

```
vpcPlot(
  fit,
  data = NULL,
  n = 300,
  bins = "jenks",
  n_bins = "auto",
  bin_mid = "mean",
  show = NULL,
  stratify = NULL,
  pred_corr = FALSE,
  pred_corr_lower_bnd = 0,
  pi = c(0.05, 0.95),
  ci = c(0.05, 0.95),
  uloq = fit$dataUloq,
  lloq = fit$dataLloq,
  log_y = FALSE,
  log_y_min = 0.001,
  xlab = NULL,
  ylab = NULL,
```

```

title = NULL,
smooth = TRUE,
vpc_theme = NULL,
facet = "wrap",
scales = "fixed",
labeller = NULL,
vpcdb = FALSE,
verbose = FALSE,
...,
seed = 1009,
idv = "time",
cens = FALSE
)

```

Arguments

<code>fit</code>	nlmixr2 fit object
<code>data</code>	this is the data to use to augment the VPC fit. By default is the fitted data, (can be retrieved by getData), but it can be changed by specifying this argument.
<code>n</code>	Number of VPC simulations
<code>bins</code>	either "density", "time", or "data", "none", or one of the approaches available in <code>classInterval()</code> such as "jenks" (default) or "pretty", or a numeric vector specifying the bin separators.
<code>n_bins</code>	when using the "auto" binning method, what number of bins to aim for
<code>bin_mid</code>	either "mean" for the mean of all timepoints (default) or "middle" to use the average of the bin boundaries.
<code>show</code>	what to show in VPC (<code>obs_dv</code> , <code>obs_ci</code> , <code>pi</code> , <code>pi_as_area</code> , <code>pi_ci</code> , <code>obs_median</code> , <code>sim_median</code> , <code>sim_median_ci</code>)
<code>stratify</code>	character vector of stratification variables. Only 1 or 2 stratification variables can be supplied.
<code>pred_corr</code>	perform prediction-correction?
<code>pred_corr_lower_bnd</code>	lower bound for the prediction-correction
<code>pi</code>	simulated prediction interval to plot. Default is <code>c(0.05, 0.95)</code> ,
<code>ci</code>	confidence interval to plot. Default is <code>(0.05, 0.95)</code>
<code>uloq</code>	Number or NULL indicating upper limit of quantification. Default is NULL.
<code>lloq</code>	Number or NULL indicating lower limit of quantification. Default is NULL.
<code>log_y</code>	Boolean indicating whether y-axis should be shown as logarithmic. Default is FALSE.
<code>log_y_min</code>	minimal value when using <code>log_y</code> argument. Default is <code>1e-3</code> .
<code>xlab</code>	label for x axis
<code>ylab</code>	label for y axis
<code>title</code>	title

smooth	"smooth" the VPC (connect bin midpoints) or show bins as rectangular boxes. Default is TRUE.
vpc_theme	theme to be used in VPC. Expects list of class vpc_theme created with function vpc_theme()
facet	either "wrap", "columns", or "rows"
scales	either "fixed" (default), "free_y", "free_x" or "free"
labeller	ggplot2 labeller function to be passed to underlying ggplot object
vpcdb	Boolean whether to return the underlying vpcdb rather than the plot
verbose	show debugging information (TRUE or FALSE)
...	Additional arguments passed to <code>nlmixr2plot::vpcPlot()</code> .
seed	an object specifying if and how the random number generator should be initialized
idv	Name of independent variable. For <code>vpcPlot()</code> and <code>vpcCens()</code> the default is "time" for <code>vpcPlotTad()</code> and <code>vpcCensTad()</code> this is "tad"
cens	is a boolean to show if this is a censoring plot or not. When <code>cens=TRUE</code> this is actually a censoring vpc plot (with <code>vpcCens()</code> and <code>vpcCensTad()</code>). When <code>cens=FALSE</code> this is traditional VPC plot (<code>vpcPlot()</code> and <code>vpcPlotTad()</code>).

Value

Simulated dataset (invisibly)

Author(s)

Matthew L. Fidler

Examples

```
one.cmt <- function() {
  ini({
    tka <- 0.45; label("Ka")
    tcl <- log(c(0, 2.7, 100)); label("Cl")
    tv <- 3.45; label("V")
    eta.ka ~ 0.6
    eta.cl ~ 0.3
    eta.v ~ 0.1
    add.sd <- 0.7; label("Additive residual error")
  })
  model({
    ka <- exp(tka + eta.ka)
    cl <- exp(tcl + eta.cl)
    v <- exp(tv + eta.v)
    linCmt() ~ add(add.sd)
  })
}

fit <-
  nlmixr2est::nlmixr(
```

```

    one.cmt,
    data = nlmixr2data::theo_sd,
    est = "saem",
    control = list(print = 0)
  )

vpcPlot(fit)

```

vpcPlotTad

VPC based on ui model

Description

VPC based on ui model

Usage

```
vpcPlotTad(..., idv = "tad")
```

Arguments

... Additional arguments passed to `nlmixr2plot::vpcPlotTad()`.

idv Name of independent variable. For `vpcPlot()` and `vpcCens()` the default is "time" for `vpcPlotTad()` and `vpcCensTad()` this is "tad"

Value

Simulated dataset (invisibly)

Author(s)

Matthew L. Fidler

Examples

```

one.cmt <- function() {
  ini({
    tka <- 0.45; label("Ka")
    tcl <- log(c(0, 2.7, 100)); label("Cl")
    tv <- 3.45; label("V")
    eta.ka ~ 0.6
    eta.cl ~ 0.3
    eta.v ~ 0.1
    add.sd <- 0.7; label("Additive residual error")
  })
  model({
    ka <- exp(tka + eta.ka)
    cl <- exp(tcl + eta.cl)

```



```

      v <- exp(tv + eta.v)
      linCmt() ~ add(add.sd)
    })
  }

  fit <-
  nlmixr2est::nlmixr(
    one.cmt,
    data = nlmixr2data::theo_sd,
    est = "saem",
    control = list(print = 0)
  )

  vpcPlot(fit)

```

vpcSim

VPC simulation

Description

VPC simulation

Usage

```

vpcSim(
  object,
  ...,
  keep = NULL,
  n = 300,
  pred = FALSE,
  seed = 1009,
  nretry = 50,
  minN = 10,
  normRelated = TRUE
)

```

Arguments

object	This is the nlmixr2 fit object
...	Additional arguments passed to nlmixr2est::vpcSim() .
keep	Column names to keep in the output simulated dataset
n	Number of simulations
pred	Should predictions be added to the simulation
seed	Seed to set for the VPC simulation
nretry	Number of times to retry the simulation if there is NA values in the simulation
minN	With retries, the minimum number of studies to restimulate (by default 10)
normRelated	should the VPC style simulation be for normal related variables only

Value

data frame of the VPC simulation

Author(s)

Matthew L. Fidler

Examples

```
if (rxode2:::linCmtSensB()) {  
  
  one.cmt <- function() {  
    ini({  
      ## You may label each parameter with a comment  
      tka <- 0.45 # Log Ka  
      tcl <- log(c(0, 2.7, 100)) # Log Cl  
      ## This works with interactive models  
      ## You may also label the preceding line with label("label text")  
      tv <- 3.45; label("log V")  
      ## the label("Label name") works with all models  
      eta.ka ~ 0.6  
      eta.cl ~ 0.3  
      eta.v ~ 0.1  
      add.sd <- 0.7  
    })  
    model({  
      ka <- exp(tka + eta.ka)  
      cl <- exp(tcl + eta.cl)  
      v <- exp(tv + eta.v)  
      linCmt() ~ add(add.sd)  
    })  
  }  
  
  fit <- nlmixr(one.cmt, theo_sd, est="focei")  
  
  head(vpcSim(fit, pred=TRUE))  
  
}
```

Index

addCwres, [2](#), [106](#)
addNpde, [4](#), [106](#)
addTable, [5](#)
AIC, [79](#)

BIC, [79](#)
bobyqaControl, [7](#)
bootplot, [14](#)
bootstrapFit, [14](#)

fixed.effects, [80](#)
fixedControl, [98–100](#)
foceiControl, [16](#), [62](#), [104](#)
format, [87](#)

getData, [118](#)
getOption, [87](#)

lbfgsb3cControl, [30](#)
llpControl, [98–100](#)
logical, [83](#)
logLik, [79](#)

n1qn1, [29](#), [30](#)
n1qn1Control, [37](#)
newuoaControl, [43](#)
nlm, [59](#)
nlmControl, [50](#)
nlme, [71](#), [79](#)
nlmeControl, [58](#)
nlminb, [59](#)
nlminbControl, [62](#)
nlmixr2, [69](#), [104](#)
nlmixr2CheckInstall, [81](#)
nlmixr2est::addNpde(), [4](#)
nlmixr2est::bobyqaControl(), [13](#)
nlmixr2est::foceiControl(), [19](#)
nlmixr2est::lbfgsb3cControl(), [36](#)
nlmixr2est::n1qn1Control(), [42](#)
nlmixr2est::newuoaControl(), [49](#)
nlmixr2est::nlmControl(), [57](#)

nlmixr2est::nlmeControl(), [61](#)
nlmixr2est::nlminbControl(), [69](#)
nlmixr2est::nlmixr2(), [70](#)
nlmixr2est::nlsControl(), [88](#)
nlmixr2est::optimControl(), [96](#)
nlmixr2est::saemControl(), [104](#)
nlmixr2est::uobyqaControl(), [114](#)
nlmixr2est::vpcSim(), [121](#)
nlmixr2extra::bootplot(), [14](#)
nlmixr2NlmeControl, [30](#), [104](#)
nlmixr2plot::traceplot(), [107](#)
nlmixr2plot::vpcCens(), [115](#)
nlmixr2plot::vpcCensTad(), [116](#)
nlmixr2plot::vpcPlot(), [119](#)
nlmixr2plot::vpcPlotTad(), [120](#)
nls, [82](#)
nlsControl, [81](#)
numericDeriv, [83](#)

optim, [29](#), [30](#)
optimControl, [89](#)

pdClasses, [60](#)
preconditionFit, [97](#)
profile.nlmixr2FitCore, [98–100](#)
profileFixed, [98](#), [100](#)
profileFixedSingle, [99](#)
profileLlp, [98](#), [99](#), [99](#)
profileNlmixr2FitCoreRet, [98–100](#)

random.effects, [80](#)
rxode2, [73](#)
rxSolve, [30](#)

saemControl, [30](#), [62](#), [78](#), [100](#)
setOfv, [104](#)
setRxThreads(), [106](#)

tableControl, [105](#)
traceplot, [107](#)

uobyqaControl, [108](#)

varClasses, [60](#)

vpcCens, [115](#)

vpcCensTad, [116](#)

vpcPlot, [117](#)

vpcPlotTad, [120](#)

vpcSim, [121](#)

warning, [59](#)