
RALFit Documentation

Release 0.0.1

STFC Numerical Analysis Group

Dec 20, 2018

Contents

| | | |
|----------|--|----------|
| 1 | Contents | 3 |
| 1.1 | Installation | 3 |
| 1.2 | How to use the package | 3 |
| 1.3 | Description of the method used | 14 |
| 1.4 | Examples | 20 |
| 1.5 | Indices and tables | 23 |

RALFit computes a solution \mathbf{x} to the non-linear least-squares problem

$$\min_{\mathbf{x}} F(\mathbf{x}) := \frac{1}{2} \|\mathbf{r}(\mathbf{x})\|_{\mathbf{W}}^2 + \frac{\sigma}{p} \|\mathbf{x}\|_2^p,$$

where $\mathbf{W} \in \mathbb{R}^{m \times m}$ is a diagonal, non-negative, weighting matrix, and $\mathbf{r}(\mathbf{x}) = (r_1(\mathbf{x}), r_2(\mathbf{x}), \dots, r_m(\mathbf{x}))^T$ is a non-linear function.

A typical use may be to fit a function $f(\mathbf{x})$ to the data y_i, t_i , weighted by the uncertainty of the data, σ_i , so that

$$r_i(\mathbf{x}) := y_i - f(\mathbf{x}; t_i),$$

and \mathbf{W} is the diagonal matrix such that $\mathbf{W}_{ii} = (1/\sqrt{\sigma_i})$. For this reason we refer to the function \mathbf{r} as the *residual* function.

Various algorithms for solving this problem are implemented – see [Description of the method used](#).

1.1 Installation

1.1.1 Obtaining the code

The latest version of the code can be downloaded from GitHub by issuing the command

```
git clone https://github.com/ralna/RALFit.git
```

1.1.2 Building the library

From the `RALFit/libRALFit/` directory, issue the commands:

```
mkdir build
cd build
cmake ..
make
```

1.2 How to use the package

1.2.1 Overview

Calling sequences

Access to the package requires a `USE` statement

```
use ral_nlls_double
```

The user can then call one of the procedures:

`nlls_solve()` solves the non-linear least squares problem.

`nlls_iterate()` performs one iteration of the non-linear least squares solver.

The calling sequences of these subroutines are outlined in *Argument lists and calling sequences*.

The derived data types

For each problem, the user must employ the derived types defined by the module to declare scalars of the types `nlls_inform` and `nlls_options`. If `nlls_iterate()` is to be used, then a scalar of the type `nlls_workspace` must also be defined. The following pseudocode illustrates this.

```
use nlls_module
!...
type (NLLS_inform) :: inform
type (NLLS_options) :: options
type (NLLS_workspace) :: work ! needed if nlls_iterate to be called
!...
```

The components of `nlls_options` and `nlls_inform` are explained below in *Data types*.

1.2.2 Argument lists and calling sequences

We use square brackets to indicate arguments. In each call, optional arguments follow the argument inform. Since we reserve the right to add additional optional arguments in future releases of the code, **we strongly recommend that all optional arguments be called by keyword, not by position**.

The term **package type** is used to mean double precision.

To solve the non-linear least squares problem

subroutine nlls_solve (*n*, *m*, *X*, *eval_r*, *eval_J*, *eval_Hf*, *params*, *options*, *inform*[, *weights*, *eval_HP*]
])

Solves the non-linear least squares problem.

Parameters

- **n** [*integer,in*] :: holds the number *n* of variables to be fitted; i.e., *n* is the length of the unknown vector *x*. **Restriction:** *n* > 0.
- **m** [*integer,in*] :: holds the number *m* of data points available; i.e., *m* is the number of residuals *r_i*. **Restriction:** *m* ≥ 0.
- **X** (*n*) [*real,inout*] :: on entry, it must hold the initial guess for *x*, and on successful exit it holds the solution to the non-linear least squares problem.
- **eval_r** [*procedure*] :: given a point *x_k*, returns the vector *r*(*x_k*). Further details of the format required are given in `eval_r()` in *User-supplied function evaluation routines*.
- **eval_J** [*procedure*] :: given a point *x_k*, returns the *m* × *n* Jacobian matrix, *J_k*, of *r* evaluated at *x_k*. Further details of the format required are given in `eval_J()` in *User-supplied function evaluation routines*.
- **eval_Hf** [*procedure*] :: given vectors *x* ∈ ℝ^{*n*} and *r* ∈ ℝ^{*m*}, returns the quantity $\sum_{i=1}^m (\mathbf{r})_i \nabla^2 \mathbf{r}_i(\mathbf{x})$. Further details of the format required are given in `eval_Hf()`

in *User-supplied function evaluation routines*. If `exact_second_derivative = .false.` in `nlls_options`, then this is not referenced.

- **params** [`params_base_type`, `in`] :: holds parameters to be passed to the user-defined routines `eval_r()`, `eval_J()`, and `eval_Hf()`. Further details of its use are given in *User-supplied function evaluation routines*.
- **options** [`nlls_options`, `in`] :: controls execution of algorithm.
- **inform** [`nlls_inform`, `out`] :: components provide information about the execution of the subroutine.

Options

- **weights** (`n`) [`real`] :: If present, this holds the square-roots of the diagonal entries of the weighting matrix, \mathbf{W} . If absent, then the norm in the least squares problem is taken to be the 2-norm, that is, $\mathbf{W} = \mathbf{I}$.
- **eval_HP** [`procedure`] :: If present, this is a routine that, given vectors $\mathbf{x}, \mathbf{y} \in \mathbb{R}^m$, returns the matrix $P(\mathbf{x}, \mathbf{y}) := (H_1(\mathbf{x})\mathbf{y} \dots H_m(\mathbf{x})\mathbf{y})$. Further details of the format required are given in `eval_HP()` in *User-supplied function evaluation routines*. This is only referenced if `model = 4` in `nlls_options`.

To iterate once

subroutine nlls_iterate (`n`, `m`, `X`, `eval_r`, `eval_J`, `eval_Hf`, `params`, `options`, `inform` [, `weights`])

A call of this form allows the user to step through the solution process one iteration at a time.

`n`, `m`, `eval_F`, `eval_J`, `eval_HF`, `params`, `info`, `options` and `weights` are as in the description of `nlls_solve()`.

Parameters

- **X** (`n`) [`real`, `inout`] :: on the first call it must hold the initial guess for \mathbf{x} . On return it holds the value of \mathbf{x} at the current iterate, and must be passed unaltered to any subsequent call to `nlls_iterate()`.
- **w** [`nlls_workspace`, `inout`] :: is used to store the current state of the iteration and should not be altered by the user.

The user may use the components `convergence_normf` and `convergence_normg` and `convergence_norms` in `nlls_inform` to determine whether the iteration has converged.

1.2.3 User-supplied function evaluation routines

The user must supply routines to evaluate the residual, Jacobian and Hessian at a point. **RALFit** will call these routines internally.

In order to pass user-defined data into the evaluation calls, `params_base_type` is extended to a `user_type`, as follows:

```
type, extends( params_base_type ) :: user_type
! code declaring components of user_type
end type user_type
```

We recommend this type is wrapped in a module with the user-defined routines for evaluating the function, Jacobian, and Hessian.

The components of the extended type are accessed through a `select type` construct:

```
select type(params)
type is(user_type)
! code that accesses components of params that were defined within user_type
end select
```

For evaluating the function $r(x)$

A subroutine must be supplied to calculate $r(x)$ for a given vector x . It must implement the following interface:

```
abstract interface
  subroutine eval_r(status, n, m, x, r, params)
    integer, intent(inout) :: status
    integer, intent(in) :: n
    integer, intent(in) :: m
    double precision, dimension(n), intent(in) :: x
    double precision, dimension(m), intent(out) :: r
    class(params_base_type), intent(in) :: params
  end subroutine eval_r
end interface
```

subroutine eval_r (*status, n, m, x, r, params*)

Parameters

- **status** [*integer,inout*] :: is initialised to 0 before the routine is called. If it is set to a non-zero value by the routine, then `nlls_solve()` / `nlls_iterate()` will exit with error.
- **n** [*integer,in*] :: is passed unchanged as provided in the call to `nlls_solve()`/`nlls_iterate()`.
- **m** [*integer,in*] :: is passed unchanged as provided in the call to `nlls_solve()`/`nlls_iterate()`.
- **X** (*n*) [*real,in*] :: holds the current point x_k at which to evaluate $r(x_k)$.
- **r** (*m*) [*real,out*] :: must be set by the routine to hold the residual function evaluated at the current point x_k , $r(x_k)$.
- **params** [*params_base_type,in*] :: is passed unchanged as provided in the call to `nlls_solve()`/`nlls_iterate()`.

For evaluating the function $J = \nabla r(x)$

A subroutine must be supplied to calculate $J = \nabla r(x)$ for a given vector x . It must implement the following interface:

```
abstract interface
  subroutine eval_J(status, n, m, x, J, params)
    integer, intent(inout) :: status
    integer, intent(in) :: n
    integer, intent(in) :: m
    double precision, dimension(n), intent(in) :: x
    double precision, dimension(n*m), intent(out) :: J
    class(params_base_type), intent(in) :: params
  end subroutine eval_J
end interface
```

subroutine eval_J (*status, n, m, x, J, params*)

Parameters

- **status** [*integer,inout*] :: is initialised to 0 before the routine is called. If it is set to a non-zero value by the routine, then `nlls_solve()` / `nlls_iterate()` will exit with error.
- **n** [*integer,in*] :: is passed unchanged as provided in the call to `nlls_solve()`/`nlls_iterate()`.
- **m** [*integer,in*] :: is passed unchanged as provided in the call to `nlls_solve()`/`nlls_iterate()`.
- **X** (n) [*real,in*] :: holds the current point \mathbf{x}_k at which to evaluate $\mathbf{J}(\mathbf{x}_k)$.
- **J** (m*n) [*real,out*] :: must be set by the routine to hold the Jacobian of the residual function evaluated at the current point \mathbf{x}_k , $\mathbf{r}(\mathbf{x}_k)$. $J(i*m+j)$ must be set to hold $\nabla_{x_j} r_i(\mathbf{x}_k)$.
- **params** [*params_base_type,in*] :: is passed unchanged as provided in the call to `nlls_solve()`/`nlls_iterate()`.

For evaluating the function $Hf = \sum_{i=1}^m r_i(\mathbf{x}) \mathbf{W} \nabla^2 r_i(\mathbf{x})$

A subroutine must be supplied to calculate $Hf = \sum_{i=1}^m (\mathbf{r})_i \nabla^2 r_i(\mathbf{x})$ for given vectors $\mathbf{x} \in \mathbb{R}^n$ and $\mathbf{r} \in \mathbb{R}^m$; here $(\mathbf{r})_i$ denotes the i th component of the vector \mathbf{r} . The subroutine must implement the following interface:

```
abstract interface
  subroutine eval_Hf_type(status, n, m, x, r, Hf, params)
    integer, intent(inout) :: status
    integer, intent(in) :: n
    integer, intent(in) :: m
    double precision, dimension(n), intent(in) :: x
    double precision, dimension(m), intent(in) :: r
    double precision, dimension(n*n), intent(out) :: Hf
    class(params_base_type), intent(in) :: params
  end subroutine eval_Hf_type
end interface
:language: fortran
```

subroutine eval_Hf (*status, n, m, x, r, Hf, params*)

Parameters

- **status** [*integer,inout*] :: is initialised to 0 before the routine is called. If it is set to a non-zero value by the routine, then `nlls_solve()` / `nlls_iterate()` will exit with error.
- **n** [*integer,in*] :: is passed unchanged as provided in the call to `nlls_solve()`/`nlls_iterate()`.
- **m** [*integer,in*] :: is passed unchanged as provided in the call to `nlls_solve()`/`nlls_iterate()`.
- **X** (n) [*real,in*] :: holds the current point \mathbf{x}_k at which to evaluate $\sum_{i=1}^m (\mathbf{r})_i \nabla^2 r_i(\mathbf{x})$.
- **r** (m) [*real,in*] :: holds $\mathbf{W} \mathbf{r}(\mathbf{x})$, the (weighted) residual, as computed from vector returned by the last call to `eval_r()`.
- **Hf** (n*n) [*real,out*] :: must be set by the routine to hold the matrix $\sum_{i=1}^m (\mathbf{r})_i \nabla^2 r_i(\mathbf{x}_k)$, held by columns as a vector, where $(\mathbf{r})_i$ denotes the i th component of \mathbf{r} , the vector passed to the routine.
- **params** [*params_base_type,in*] :: is passed unchanged as provided in the call to `nlls_solve()`/`nlls_iterate()`.

For evaluating the function $P(\mathbf{x}, \mathbf{y}) := (H_1(\mathbf{x})\mathbf{y} \dots H_m(\mathbf{x})\mathbf{y})$

A subroutine may be supplied to calculate $P(\mathbf{x}, \mathbf{y}) := (H_1(\mathbf{x})\mathbf{y} \dots H_m(\mathbf{x})\mathbf{y})$ for given vectors $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$. The subroutine must implement the following interface:

```
abstract interface
  subroutine eval_HP_type(status, n, m, x, y, HP, params)
    integer, intent(inout) :: status
    integer, intent(in) :: n
    integer, intent(in) :: m
    double precision, dimension(n), intent(in) :: x
    double precision, dimension(n), intent(in) :: y
    double precision, dimension(n*m), intent(out) :: HP
    class(params_base_type), intent(in) :: params
  end subroutine eval_HP_type
end interface
:language: fortran
```

subroutine eval_HP (*status, n, m, x, y, HP, params*)

Parameters

- **status** [*integer,inout*] :: is initialised to 0 before the routine is called. If it is set to a non-zero value by the routine, then `nlls_solve()` / `nlls_iterate()` will exit with error.
- **n** [*integer,in*] :: is passed unchanged as provided in the call to `nlls_solve()` / `nlls_iterate()`.
- **m** [*integer,in*] :: is passed unchanged as provided in the call to `nlls_solve()` / `nlls_iterate()`.
- **x** (n) [*real,in*] :: holds the current point \mathbf{x}_k at which to evaluate the Hessians $\nabla^2 r_i(\mathbf{x}_k)$.
- **y** (n) [*real,in*] :: holds \mathbf{y} , the vector which multiplies each Hessian.
- **HP** (n*m) [*real,out*] :: must be set by the routine to hold the matrix $P(\mathbf{x}, \mathbf{y}) := (H_1(\mathbf{x})\mathbf{y} \dots H_m(\mathbf{x})\mathbf{y})$, held by columns as a vector.
- **params** [*params_base_type,in*] :: is passed unchanged as provided in the call to `nlls_solve()` / `nlls_iterate()`.

1.2.4 Data types

The derived data type for holding options

type nlls_options

This is used to hold controlling data. The components are automatically given default values in the definition of the type.

Printing Controls

Type fields

- **% error** [*integer,default=6*] :: the Fortran unit number for error messages. If it is negative, these messages will be suppressed.
- **% out** [*integer,default=6*] :: the Fortran unit number for general messages. If it is negative, these messages will be suppressed.

- % **print_level** [*integer,default=0*] :: controls the level of output required. Options are:

| | |
|-----|---|
| < 1 | No informational output will occur. |
| 1 | Gives a one-line summary for each iteration. |
| 2 | As 1, plus gives a summary of the inner iteration for each iteration. |
| 3 | As 2, plus gives more verbose (debugging) output. |

Choice of Algorithm

Type fields

- % **model** [*integer,default=3*] :: specifies the model, $m_k(\cdot)$, used. Possible values are:

| | |
|---|---|
| 1 | Gauss-Newton (no Hessian). |
| 2 | (Quasi-)Newton (uses exact Hessian if <code>exact_second_derivatives</code> is <code>true</code> , otherwise builds an approximation to the Hessian). |
| 3 | Hybrid method (mixture of Gauss-Newton/(Quasi-)Newton, as determined by the package). |
| 4 | Newton-tensor method. |

See [The models](#) for details.

- % **type_of_method** [*integer,default=1*] :: specifies the type of globalization method used. Possible values are:

| | |
|---|----------------------|
| 1 | Trust-region method. |
| 2 | Regularization. |

See [Subproblem solves](#) for details.

- % **nlls_method** [*integer,default=4*] :: specifies the method used to solve (or approximate the solution to) the trust-region sub problem. Possible values are:

| | |
|---|--|
| 1 | Powell's dogleg method (approximates the solution). |
| 2 | The Adachi-Iwata-Nakatsukasa-Takeda (AINT) method. |
| 3 | The More-Sorensen method. |
| 4 | Galahad's DTRS method if <code>type_of_method=1</code> , or Galahad's DRQS method if <code>type_of_method=2</code> . |

See [Subproblem solves](#) for details.

- % **exact_second_derivatives** [*logical,default=false*] :: if `true`, signifies that the exact second derivatives are available (and, if `false`, approximates them using a secant method).

Solving a regularized problem

Type fields

- % **regularization** [*integer,default=0*] :: specifies if a regularized non-linear least squares problem needs to be solved, and if so, what method is used to solve it. Possible values are:

| | |
|---|--|
| 0 | $\sigma = 0$, and an unregularized problem is solved |
| 1 | a non-linear least-squares problem of size $n + m$ is solved implicitly. Can only be used if $p = 2$. |
| 2 | a non-linear least-squares problem of size $n + 1$ is solved implicitly. |

See *Incorporating the regularization term* for details.

- % **regularization_term** [*real,default=0.0*] :: specifies the regularization weight, σ , used in the least-squares problem.
- % **regularization_power** [*real,default=0.0*] :: specifies the regularization index, p , used in the least-squares problem.

Stopping rules

Type fields

- % **maxit** [*integer,default=100*] :: gives the number of iterations the algorithm is allowed to take before being stopped. This is not accessed if `nlls_iterate()` is used.
- % **stop_g_absolute** [*real,default=1e-5*] :: specifies the absolute tolerance used in the convergence test on $\|\mathbf{J}_k^T \mathbf{r}(\mathbf{x}_k)\|/\|\mathbf{r}(\mathbf{x}_k)\|$.
- % **stop_g_relative** [*real,default=1e-8*] :: specifies the relative tolerance used in the convergence test on $\|\mathbf{J}_k^T \mathbf{r}(\mathbf{x}_k)\|/\|\mathbf{r}(\mathbf{x}_k)\|$.
- % **stop_f_absolute** [*real,default=1e-5*] :: specifies the absolute tolerance used in the convergence test on $\|\mathbf{r}(\mathbf{x}_k)\|$.
- % **stop_f_relative** [*real,default=1e-8*] :: specifies the relative tolerance used in the convergence test on $\|\mathbf{r}(\mathbf{x}_k)\|$.
- % **stop_s** [*real,default=eps*] :: specifies the tolerance used in the convergence test on $\|\mathbf{s}_k\|$.

Trust region radius/regularization behaviour

Type fields

- % **relative_tr_radius** [*integer,default=0*] :: specifies whether the initial trust region radius should be scaled.
- % **initial_radius_scale** [*real,default=1.0*] :: specifies the scaling parameter for the initial trust region radius, which is only used if `relative_tr_radius = 1`.
- % **initial_radius** [*real,default=100.0*] :: specifies the initial trust-region radius, Δ .
- % **maximum_radius** [*real,default=1e8*] :: specifies the maximum size permitted for the trust-region radius.
- % **eta_successful** [*real,default=1e-8*] :: specifies the smallest value of ρ such that the step is accepted. .. `success_but_reduce` is also available, but not documented
- % **eta_very_successful** [*real,default=0.9*] :: specifies the value of ρ after which the trust-region radius is increased.
- % **eta_too_successful** [*real,default=2.0*] :: specifies that value of ρ after which the step is accepted, but keep the trust-region radius unchanged.
- % **radius_increase** [*real,default=2.0*] :: specifies the factor to increase the trust-region radius by.

- % **radius_reduce** [*real,default=0.5*] :: specifies the factor to decrease the trust-region radius by.
- % **tr_update_strategy** [*integer,default=1*] :: specifies the strategy used to update Δ_k . Possible values are:

| | |
|---|------------------------------|
| 1 | use the usual step function. |
| 2 | use a the continuous method. |

Scaling options

Type fields

- % **scale** [*integer,default=1*] :: specifies how, if at all, we scale the Jacobian. We calculate a diagonal scaling matrix, D , as follows:
- % **scale_trim_max** [*logical,default=true*] :: specifies whether or not to trim large values of the scaling matrix, D . If `true`, $D_{i,i} \leftarrow \min(D_{i,i}, \text{scale_max})$.
- % **scale_max** [*real,default=1e11*] :: specifies the maximum value allowed if `scale_trim_max = true`.
- % **scale_trim_min** [*logical,default=true*] :: specifies whether or not to trim small values of the scaling matrix, D . If `true`, $D_{i,i} \leftarrow \max(D_{i,i}, \text{scale_min})$.
- % **scale_min** [*real,default=1e-11*] :: specifies the minimum value allowed if `scale_trim_max = true`.
- % **scale_require_increase** [*logical,default=false*] :: specifies whether or not to require $D_{i,i}$ to increase before updating it.

Hybrid method options These options are used if `model=3`

Type fields

- % **hybrid_switch** [*real,default=0.1*] :: specifies the value, if `model=3`, at which second derivatives are used.
- % **hybrid_tol** [*real,default=2.0*] :: if `model=3`, specifies the value such that if $\|J_k^T W r(x_k)\|_2 < \text{hybrid_tol} * 0.5 \|r(x_k)\|_W^2$ the method switches to a quasi-Newton method.
- % **hybrid_switch_its** [*integer,default=1*] :: if `model=3`, sets how many iterates in a row must the condition in the definition of `hybrid_tol` hold before a switch.

Newton-Tensor options These options are used if `model=4`

Type fields

- % **reg_order** [*real,default=0.0*] :: if `nlls_method = 4`, the order of the regularization used (p in **TODO** (eq.: `reg_subproblem`)). If `reg_order = 0.0`, then the algorithm chooses an appropriate value of p .
- % **inner_method** [*integer,default=2*] :: if `nlls_method = 4`, specifies the method used to pass in the regularization parameter to the inner non-linear least squares solver. Possible values are:

| | |
|---|--|
| 1 | The current regularization parameter is passed in as a base regularization parameter. |
| 2 | A larger non-linear least squares problem is explicitly formed to be solved as the subproblem. |
| 3 | The regularization is handled implicitly by calling RALFit recursively. |

More-Sorensen options These options are used if `nlls_method=3`

Type fields

- % **more_sorensen_maxits** [*integer,default=3*] :: if `nlls_method = 3`, specifies the maximum number of iterations allowed in the More-Sorensen method.
- % **more_sorensen_maxits** :: if `nlls_method = 3`, specifies the maximum number of iterations allowed in the More-Sorensen method.
- % **more_sorensen_shift** [*real,default=1e-13*] :: if `nlls_method = 3`, specifies the shift to be used in the More-Sorensen method.
- % **more_sorensen_tiny** [*real,default=10.0*eps*] :: if `nlls_method = 3`, specifies the value below which numbers are considered to be essentially zero.
- % **more_sorensen_tol** [*real,default=1e-3*] :: if `nlls_method = 3`, specifies the tolerance to be used in the More-Sorensen method.

Other options

Type fields

- % **calculate_svd_J** [*logical,default=false*] :: specifies whether or not to calculate the singular value decomposition of `J` at each iteration.
- % **output_progress_vectors** [*logical,default=false*] :: if true, outputs the progress vectors `nlls_inform%resvec` and `nlls_inform%gradvec` at the end of the routine.

The derived data type for holding information

`type nlls_inform`

This is used to hold information about the progress of the algorithm.

Type fields

- % **status** [*integer*] :: gives the exit status of the subroutine. See [Warning and error messages](#) for details.
- % **error_message** (80) [*character*] :: holds the error message corresponding to the exit status.
- % **alloc_status** [*integer*] :: gives the status of the last attempted allocation/deallocation.
- % **bad_alloc** (80) [*character*] :: holds the name of the array that was being allocated when an error was flagged.
- % **iter** [*integer*] :: gives the total number of iterations performed.
- % **f_eval** [*integer*] :: gives the total number of evaluations of the objective function.
- % **g_eval** [*integer*] :: gives the total number of evaluations of the gradient of the objective function.
- % **h_eval** [*integer*] :: gives the total number of evaluations of the Hessian of the objective function.
- % **convergence_normf** [*integer*] :: tells us if the test on the size of `r` is satisfied.
- % **convergence_normf** :: that tells us if the test on the size of the gradient is satisfied.
- % **convergence_normf** :: that tells us if the test on the step length is satisfied.

- % **resvec** (iter+1) [real] :: if `output_progress_vectors=true` in `nlls_options`, holds the vector of residuals.
- % **resvec** :: if `output_progress_vectors=true` in `nlls_options`, holds the vector of gradients.
- % **obj** [real] :: holds the value of the objective function at the best estimate of the solution determined by the algorithm.
- % **norm_g** [real] :: holds the gradient of the objective function at the best estimate of the solution determined by the package.
- % **scaled_g** [real] :: holds the gradient of the objective function at the best estimate of the solution determined by the package.
- % **external_return** [integer] :: gives the error code that was returned by a call to an external routine.
- % **external_name** (80) [character] :: holds the name of the external code that flagged an error.

The workspace derived data type

type nlls_workspace

This is used to save the state of the algorithm in between calls to `nlls_iterate()`, and must be used if that subroutine is required. It's components are not designed to be accessed by the user.

1.2.5 Warning and error messages

A successful return from a subroutine in the package is indicated by `status` in `nlls_inform` having the value zero. A non-zero value is associated with an error message, which will be output on `error` in `nlls_inform`.

Possible values are:

| | |
|------|---|
| -1 | Maximum number of iterations reached without convergence. |
| -2 | Error from evaluating a function/Jacobian/Hessian. |
| -3 | Unsupported choice of model. |
| -4 | Error return from an external routine. |
| -5 | Unsupported choice of method. |
| -6 | Allocation error. |
| -7 | Maximum number of reductions of the trust radius reached. |
| -8 | No progress being made in the solution. |
| -9 | n > m. |
| -10 | Unsupported trust region update strategy. |
| -11 | Unable to valid step when solving trust region subproblem. |
| -12 | Unsupported scaling method. |
| -13 | Error accessing pre-allocated workspace. |
| -14 | Unsupported value in <code>type_of_method</code> |
| -101 | Unsupported model in <code>dogleg</code> (<code>nlls_method = 1</code>). |
| -201 | All eigenvalues are imaginary (<code>nlls_method=2</code>). |
| -202 | Matrix with odd number of columns sent to <code>max_eig</code> subroutine (<code>nlls_method=2</code>). |
| -301 | <code>more_sorensen_max_its</code> is exceeded in <code>more_sorensen</code> subroutine (<code>nlls_method=3</code>). |
| -302 | Too many shifts taken in <code>more_sorensen</code> subroutine (<code>nlls_method=3</code>). |
| -303 | No progress being made in <code>more_sorensen</code> subroutine (<code>nlls_method=3</code>). |
| -401 | <code>model = 4</code> selected, but <code>exact_second_derivatives</code> is set to <code>false</code> . |

1.3 Description of the method used

RALFit computes a solution \mathbf{x} to the non-linear least-squares problem

$$\min_{\mathbf{x}} F(\mathbf{x}) := \frac{1}{2} \|\mathbf{r}(\mathbf{x})\|_{\mathbf{W}}^2 + \frac{\sigma}{p} \|\mathbf{x}\|_2^p, \quad (1.1)$$

Here we describe the method used to solve (1.1). **RALFit** implements an iterative method that, at each iteration, calculates and returns a step \mathbf{s} that reduces the model by an acceptable amount by solving (or approximating a solution to) a subproblem, as detailed in [Subproblem solves](#).

The algorithm is iterative. At each point, \mathbf{x}_k , the algorithm builds a model of the function at the next step, $F(\mathbf{x}_k + \mathbf{s}_k)$, which we refer to as $m_k(\cdot)$. We allow either a Gauss-Newton model, a (quasi-)Newton model, or a Newton-tensor model; see [The models](#) for more details.

Once the model has been formed we find a candidate for the next step by solving a suitable subproblem. The quantity

$$\rho = \frac{F(\mathbf{x}_k) - F(\mathbf{x}_k + \mathbf{s}_k)}{m_k(\mathbf{x}_k) - m_k(\mathbf{x}_k + \mathbf{s}_k)} \quad (1.2)$$

is then calculated. If this is sufficiently large we accept the step, and \mathbf{x}_{k+1} is set to $\mathbf{x}_k + \mathbf{s}_k$; if not, the parameter Δ_k is reduced and the resulting new trust-region sub-problem is solved. If the step is very successful – in that ρ is close to one – Δ_k is increased. Details are explained in [Accepting the step and updating the parameter](#).

This process continues until either the residual, $\|\mathbf{r}(\mathbf{x}_k)\|_{\mathbf{W}}$, or a measure of the gradient, $\|\mathbf{J}_k^T \mathbf{W} \mathbf{r}(\mathbf{x}_k)\|_2 / \|\mathbf{r}(\mathbf{x}_k)\|_{\mathbf{W}}$, is sufficiently small.

1.3.1 The models

A vital component of the algorithm is the choice of model employed. There are four choices available, controlled by the parameter `model` of `nlls_options`.

model = 1 this implements the **Gauss-Newton** model. Here we replace $\mathbf{r}(\mathbf{x}_k + \mathbf{s})$ by its first-order Taylor approximation, $\mathbf{r}(\mathbf{x}_k) + \mathbf{J}_k \mathbf{s}$. The model is therefore given by

$$m_k^{GN}(\mathbf{s}) = \frac{1}{2} \|\mathbf{r}(\mathbf{x}_k) + \mathbf{J}_k \mathbf{s}\|_{\mathbf{W}}^2. \quad (1.3)$$

model = 2 this implements the **Newton** model. Here, instead of approximating the residual, $\mathbf{r}(\cdot)$, we take as our model the second-order Taylor approximation of the function, $F(\mathbf{x}_{k+1})$. Namely, we use

where $\mathbf{g}_k = \mathbf{J}_k^T \mathbf{W} \mathbf{r}(\mathbf{x}_k)$ and $\mathbf{H}_k = \sum_{i=1}^m r_i(\mathbf{x}_k) \mathbf{W} \nabla^2 r_i(\mathbf{x}_k)$. Note that $m_k^N(\mathbf{s}) = m_k^{GN}(\mathbf{s}) + \frac{1}{2} \mathbf{s}^T \mathbf{H}_k \mathbf{s}$.

If the second derivatives of $\mathbf{r}(\cdot)$ are not available (i.e., the option `exact_second_derivatives` is set to `false`), then the method approximates the matrix \mathbf{H}_k ; see [Approximating the Hessian](#).

model = 3 This implements a **hybrid** model. In practice the Gauss-Newton model tends to work well far away from the solution, whereas Newton performs better once we are near to the minimum (particularly if the residual is large at the solution). This option will try to switch between these two models, picking the model that is most appropriate for the step. In particular, we start using $m_k^{GN}(\cdot)$, and switch to $m_k^N(\cdot)$ if $\|\mathbf{g}_k\|_2 \leq \text{hybrid_tol} \frac{1}{2} \|\mathbf{r}(\mathbf{x}_k)\|_{\mathbf{W}}$ for more than `hybrid_switch_its` iterations in a row. If, in subsequent iterations, we fail to get a decrease in the function value, then the algorithm interprets this as being not sufficiently close to the solution, and thus switches back to using the Gauss-Newton model.

The exact method used is described below:

```

if use_second_derivatives    // previous step used Newton model
    if ||gk+1|| > ||gk||
        use_second_derivatives = false    // Switch back to Gauss-Newton
        Hftemp = Hfk, Hfk = 0    // Copy Hessian back to temp array
    endif
else
    if ||gk+1||/normFk+1 < hybrid_tol
        hybrid_count = hybrid_count + 1    // Update the no of successive failures
        if hybrid_count = hybrid_count_switch_its
            use_second_derivatives = true
            hybrid_count = 0
            Hftemp = Hfk // Copy approximate Hessian back
        end if
    end if
end if
end if

```

model = 4 this implements a **Newton-tensor** model. This uses a second order Taylor approximation to the residual, namely

$$r_i(\mathbf{x}_k + \mathbf{s}) \approx (\mathbf{t}_k(\mathbf{s}))_i := r_i(\mathbf{x}_k) + (\mathbf{J}_k)_i \mathbf{s} + \frac{1}{2} \mathbf{s}^T B_{ik} \mathbf{s},$$

where $(\mathbf{J}_k)_i$ is the i th row of \mathbf{J}_k , and B_{ik} is $\nabla^2 r_i(\mathbf{x}_k)$. We use this to define our model

$$m_k^{NT}(\mathbf{s}) = \frac{1}{2} \|\mathbf{t}_k(\mathbf{s})\|_{\mathbf{W}}^2. \quad (1.4)$$

1.3.2 Approximating the Hessian

If the exact Hessian is not available, we approximate it using the method of Dennis, Gay, and Welsch⁴. The method used is given as follows:

```

function Hfk+1 = rank_one_update(d, gk, gk+1, rk+1, Jk, Hfk)
    y = gk - gk+1
    ŷ = JkT rk+1 - gk+1
    Hfk = min(1, |dT ŷ| / |dT Hfk d|) Hfk
    Hfk+1 = Hfk + ((ŷk+1 - Hfk d)T d) / yT d

```

It is sometimes the case that this approximation becomes corrupted, and the algorithm may not recover from this. To guard against this, if **model = 3** in *nlls_options* and we are using this approximation to the Hessian in our (quasi-Newton) model, we test against the Gauss-Newton model if the first step is unsuccessful. If the Gauss-Newton step would have been successful, we discard the approximate Hessian information, and recompute the step using Gauss-Newton.

In the case where **model=3**, the approximation to the Hessian is updated at each step whether or not it is needed for the current calculation.

⁴ Nocedal, J., & Wright, S. (2006). Numerical optimization. Springer Science & Business Media.

1.3.3 Subproblem solves

The main algorithm calls a number of subroutines. The most vital is the subroutine `calculate_step`, which finds a step that minimizes the model chosen, subject to a globalization strategy. The algorithm supports the use of two such strategies: using a trust-region, and regularization. If Gauss-Newton, (quasi-)Newton, or a hybrid method is used (`model = 1, 2, 3` in `nlls_options`), then the model function is quadratic, and the methods available to solve the subproblem are described in *The trust region method* and *Regularization*. If the Newton-Tensor model is selected (`model = 4` in `nlls_options`), then this model is not quadratic, and the methods available are described in *Newton-Tensor subproblem*.

Note that, when calculating the step, if the initial regularization parameter σ in (1.1) is non-zero, then we must modify $J_k^T J_k$ to take into account the Jacobian of the modified least squares problem being solved. Practically, this amounts to making the change

$$J_k^T J_k = J_k^T J_k + \begin{cases} \sigma I & \text{if } p = 2 \\ \frac{\sigma p}{2} \|x_k\|^{p-4} x_k x_k^T & \text{otherwise} \end{cases}.$$

The trust region method

If `model = 1, 2, or 3`, and `type_of_method=1`, then we solve the subproblem

$$s_k = \arg \min_s m_k(s) \quad \text{s.t.} \quad \|s\|_B \leq \Delta_k, \quad (1.5)$$

and we take as our next step the minimum of the model within some radius of the current point. The method used to solve this is dependent on the control parameter `optionsnlls_method`. The algorithms called for each of the options are listed below:

nlls_method = 1 approximates the solution to (1.5) by using Powell's dogleg method. This takes as the step a linear combination of the Gauss-Newton step and the steepest descent step, and the method used is described here:

```
function dogleg(J, r, Hf, g, Δ)
α = ||g||2 / ||J * g||2
dsd = α g
solve dgn = arg minx ||Jx - r||2
if ||dgn|| ≤ Δ then
    d = dgn
else if ||α dsd|| ≥ Δ
    d = (Δ / ||dsd||) dsd
else
    d = α dsd + β (dgn - α dsd), where β is chosen such that ||d|| = Δ
end if
```

nlls_method = 2 solves the trust region subproblem using the trust region solver of Adachi, Iwata, Nakatsukasa, and Takeda. This reformulates the problem (1.5) as a generalized eigenvalue problem, and solves that. See¹ for more details.

¹ Adachi, Satoru and Iwata, Satoru and Nakatsukasa, Yuji and Takeda, Akiko (2015). Solving the trust region subproblem by a generalized eigenvalue problem. Technical report, Mathematical Engineering, The University of Tokyo.

nlls_method = 3 this solves (1.5) using a variant of the More-Sorensen method. In particular, we implement Algorithm 7.3.6 in Trust Region Methods by Conn, Gould and Toint².

nlls_method = 4 this solves (1.5) by first converting the problem into the form

$$\min_p \mathbf{w}^T \mathbf{p} + \frac{1}{2} \mathbf{p}^T \mathbf{D} \mathbf{p} \quad \text{s.t.} \quad \|\mathbf{p}\| \leq \Delta,$$

where \mathbf{D} is a diagonal matrix. We do this by performing an eigen-decomposition of the Hessian in the model. Then, we call the Galahad routine DTRS; see the Galahad³ documentation for further details.

Regularization

If `model = 1, 2, or 3`, and `type_of_method=2`, then the next step is taken to be the minimum of the model with a regularization term added:

$$\mathbf{s}_k = \arg \min_s m_k(\mathbf{s}) + \frac{1}{\Delta_k} \cdot \frac{1}{p} \|\mathbf{s}\|_B^p, \quad (1.6)$$

At present, only one method of solving this subproblem is supported:

nlls_method = 4: this solves (1.6) by first converting the problem into the form

$$\min_p \mathbf{w}^T \mathbf{p} + \frac{1}{2} \mathbf{p}^T \mathbf{D} \mathbf{p} + \frac{1}{p} \|\mathbf{p}\|_2^p,$$

where \mathbf{D} is a diagonal matrix. We do this by performing an eigen-decomposition of the Hessian in the model. Then, we call the Galahad routine DRQS; see the Galahad³ documentation for further details.

Newton-Tensor subproblem

If `model=4`, then the non-quadratic Newton-Tensor model is used. As such, none of the established subproblem solvers described in *The trust region method* or *Regularization* can be used.

If we use regularization (with $p = 2$), then the subproblem we need to solve is of the form

$$\min_s \frac{1}{2} \sum_{i=1}^m W_{ii}(\mathbf{t}_k(\mathbf{s}))^2 + \frac{1}{2\Delta_k} \|\mathbf{s}\|_2^2 \quad (1.7)$$

Note that (1.7) is a sum-of-squares, and as such can be solved by calling `nlls_solve()` recursively. We support two options:

inner_method = 1 if this option is selected, then `nlls_solve()` is called to solve (1.4) directly. The current regularization parameter of the ‘outer’ method is used as a base regularization in the ‘inner’ method, so that the (quadratic) subproblem being solved in the ‘inner’ call is of the form

$$\min_s m_k(\mathbf{s}) + \frac{1}{2} \left(\frac{1}{\Delta_k} + \frac{1}{\delta_k} \right) \|\mathbf{s}\|_B^2,$$

where $m_k(\mathbf{s})$ is a quadratic model of (1.4), Δ_k is the (fixed) regularization parameter of the outer iteration, and δ_k the regularization parameter of the inner iteration, which is free to be updated as required by the method.

² Conn, A. R., Gould, N. I., & Toint, P. L. (2000). Trust region methods. SIAM.

³ Gould, N. I., Orban, D., & Toint, P. L. (2003). GALAHAD, a library of thread-safe Fortran 90 packages for large-scale nonlinear optimization. ACM Transactions on Mathematical Software (TOMS), 29(4), 353-372.

`inner_method = 2` in this case we use `nlls_solve()` to solve the regularized model (1.7) directly. The number of parameters for this subproblem is $n + m$. Specifically, we have a problem of the form

$$\min_s \frac{1}{2} \|\hat{\mathbf{r}}(\mathbf{s})\|_{\mathbf{W}}^2, \quad \text{where } (\hat{\mathbf{r}}(\mathbf{s}))_i = \begin{cases} (\mathbf{t}_k(\mathbf{s}))_i & 1 \leq i \leq m \\ \frac{1}{\sqrt{\Delta_k}} s_i & m+1 \leq i \leq n+m \end{cases}.$$

This subproblem can then be solved using any of the methods described in *The trust region method* or *Regularization*.

`inner_method = 3`

In this case, `nlls_solve()` is called recursively with the inbuilt feature of solving a regularized problem, as described in *Incorporating the regularization term*

1.3.4 Accepting the step and updating the parameter

Once a step has been suggested, we must decide whether or not to accept the step, and whether the trust region radius or regularization parameter, as appropriate, should grow, shrink, or remain the same.

These decisions are made with reference to the parameter, ρ (1.2), which measures the ratio of the actual reduction in the model to the predicted reduction in the model. If this is larger than `eta_successful` in `nlls_options`, then the step is accepted.

The value of Δ_k then needs to be updated, if appropriate. The package supports two options:

`tr_update_strategy = 1` a step-function is used to decide whether or not to increase or decrease Δ_k , as described here:

```

if  $\rho \leq \text{eta\_success\_but\_reduce}$  then
     $\Delta = \text{radius\_reduce} * \Delta$     // reduce  $\Delta$ 
else if  $\rho \leq \text{eta\_very\_successful}$ 
     $\Delta = \Delta$     //  $\Delta$  stays unchanged
else if  $\rho \leq \text{eta\_too\_successful}$ 
     $\Delta = \text{radius\_increase} * \Delta$     // increase  $\Delta$ 
else if  $\rho > \text{eta\_too\_successful}$ 
     $\Delta = \Delta$     // too successful: accept step, but don't change  $\Delta$ 
end if

```

`tr_update_strategy = 2` a continuous function is used to make the decision⁵, as described below. On the first call, the parameter ν is set to 2.0.

```

if  $\rho \geq \text{eta\_too\_successful}$ 
     $\Delta = \Delta$     //  $\Delta$  stays unchanged
else if  $\rho > \text{eta\_successful}$ 
     $\Delta = \Delta * \min(\text{radius\_increase}, 1 - ((\text{radius\_increase} - 1) * ((1 - 2 * \rho)^3)))$ 
     $\nu = \text{radius\_reduce}$ 
else if  $\rho \leq \text{eta\_successful}$ 
     $\Delta = \nu * \Delta$ 
     $\nu = 0.5 * \nu$ 
end if

```

⁵ Nielsen, Hans Bruun (1999). Damping parameter in Marquadt's MMethod. Technical report TR IMM-REP-1999-05, Department of Mathematical Modelling, Technical University of Denmark (http://www2.imm.dtu.dk/documents/ftp/tr99/tr05_99.pdf)

1.3.5 Incorporating the regularization term

If a non-zero regularization term is required in (1.1), then this is handled by transforming the problem internally into a new non-linear least-squares problem. The formulation used will depend on the value of `regularization` in `nlls_options`.

regularization = 1 This is only supported if $p = 2$. We solve a least squares problem with n additional degrees of freedom. The new function, $\hat{\mathbf{r}} : \mathbb{R}^n \rightarrow \mathbb{R}^{m+n}$, is defined as

$$\hat{\mathbf{r}}_i(\mathbf{x}) = \begin{cases} \mathbf{r}_i(\mathbf{x}) & \text{for } i = 1, \dots, m \\ \sqrt{\sigma}[\mathbf{x}]_j & \text{for } i = m + j, j = 1, \dots, n \end{cases}$$

where $[\mathbf{x}]_j$ denotes the j th component of \mathbf{x} .

This problem is now in the format of a standard non-linear least-squares problem. In addition to the function values, the we also need a Jacobian and some more information about the Hessian. For our modified function, the Jacobian is

$$\hat{\mathbf{J}}(\mathbf{x}) = \begin{bmatrix} \mathbf{J}(\mathbf{x}) \\ \sqrt{\sigma} \mathbf{I} \end{bmatrix},$$

and the other function that needs to be supplied is given by

$$\hat{\mathbf{H}}_k(\mathbf{x}) = \sum_{i=1}^{n+m} \hat{\mathbf{r}}_i(\mathbf{x}) \nabla^2 \hat{\mathbf{r}}_i(\mathbf{x}) = \sum_{i=1}^n \mathbf{r}_i(\mathbf{x}) \nabla^2 \mathbf{r}_i(\mathbf{x}) = \mathbf{H}_k(\mathbf{x}).$$

We solve these problems implicitly by modifying the code so that the user does not need do any additional work. We can simply note that

$$\|\hat{\mathbf{r}}(\mathbf{x})\|^2 = \|\mathbf{r}(\mathbf{x})\|^2 + \sigma \|\mathbf{x}\|^2,$$

$$\hat{\mathbf{J}}^T \hat{\mathbf{r}} = \mathbf{J}^T \mathbf{r} + \sigma \mathbf{x},$$

and that

$$\hat{\mathbf{J}}^T \hat{\mathbf{J}} = \mathbf{J}^T \mathbf{J} + \sigma \mathbf{I}.$$

We also need to update the value of the model. Since the Hessian vanishes, we only need to be concerned with the Gauss-Newton model. We have that

$$\begin{aligned} \hat{m}_k^{GN}(\mathbf{s}) &= \frac{1}{2} \|\hat{\mathbf{r}}(\mathbf{x}_k) + \hat{\mathbf{J}}_k \mathbf{s}\|^2 \\ &= \frac{1}{2} \left(\hat{\mathbf{r}}^T \hat{\mathbf{r}} + 2 \mathbf{s}^T \hat{\mathbf{J}}_k^T \hat{\mathbf{r}} + \mathbf{s}^T \hat{\mathbf{J}}_k^T \hat{\mathbf{J}}_k \mathbf{s} \right) \\ &= \frac{1}{2} \left(\mathbf{r}^T \mathbf{r} + \sigma \mathbf{x}^T \mathbf{x} + 2(\mathbf{s}^T \mathbf{J}_k^T \mathbf{r} + \sigma \mathbf{s}^T \mathbf{x}) + \mathbf{s}^T \mathbf{J}_k^T \mathbf{J}_k \mathbf{s} + \sigma \mathbf{s}^T \mathbf{s} \right) \\ &= m_k^{GN}(\mathbf{s}) + \frac{1}{2} \sigma (\mathbf{x}^T \mathbf{x} + 2 \mathbf{s}^T \mathbf{x} + \mathbf{s}^T \mathbf{s}) \\ &= m_k^{GN}(\mathbf{s}) + \frac{\sigma}{2} \|\mathbf{x} + \mathbf{s}\|^2 \end{aligned}$$

`regularization=2`

We solve a non-linear least-squares problem with one additional degree of freedom.

Since the term $\frac{\sigma}{p} \|\mathbf{x}\|_2^p$ is non-negative, we can write

$$F_\sigma(\mathbf{x}) = \frac{1}{2} \left(\|\mathbf{r}(\mathbf{x})\|^2 + \left(\left(\frac{2\sigma}{p} \right)^{1/2} \|\mathbf{x}\|^{p/2} \right)^2 \right),$$

thereby defining a new non-linear least squares problem involving the function $\mathbf{r} : \mathbb{R}^n \rightarrow \mathbb{R}^{m+1}$ such that

$$\bar{r}_i(\mathbf{x}) = \begin{cases} r_i(\mathbf{x}) & 1 \leq i \leq m \\ \frac{2\sigma}{p} \|\mathbf{x}\|^{p/2} & i = m+1 \end{cases}.$$

The Jacobian for this new function is given by

$$\bar{\mathbf{J}}(\mathbf{x}) = \begin{bmatrix} \mathbf{J}(\mathbf{x}) \\ \left(\frac{\sigma p}{2}\right)^{1/2} \|\mathbf{x}\|^{(p-4)/2} \mathbf{x}^T \end{bmatrix},$$

and we get that

$$\nabla^2 \bar{r}_{m+1} = \left(\frac{\sigma p}{2}\right)^{1/2} \|\mathbf{x}\|^{(p-4)/2} \left(I + \frac{\mathbf{x}\mathbf{x}^T}{\|\mathbf{x}\|^2} \right).$$

As for the case where `regularization=1`, we simply need to update quantities in our non-linear least squares code to solve this problem, and the changes needed in this case are

$$\|\bar{\mathbf{r}}(\mathbf{x})\|^2 = \|\mathbf{r}(\mathbf{x})\|^2 + \frac{2\sigma}{p} \|\mathbf{x}\|^p,$$

$$\bar{\mathbf{J}}^T \bar{\mathbf{r}} = \mathbf{J}^T \mathbf{r} + \sigma \|\mathbf{x}\|^{p-2} \mathbf{x},$$

$$\bar{\mathbf{J}}^T \bar{\mathbf{J}} = \mathbf{J}^T \mathbf{J} + \frac{\sigma p}{2} \|\mathbf{x}\|^{p-4} \mathbf{x}\mathbf{x}^T,$$

$$\sum_{i=1}^{m+1} \bar{r}_i(\mathbf{x}) \bar{\mathbf{H}}_i(\mathbf{x}) = \sigma \|\mathbf{x}\|^{p-4} (\|\mathbf{x}\|^2 I + \mathbf{x}\mathbf{x}^T) + \sum_{i=1}^m r_i(\mathbf{x}) \mathbf{H}_i(\mathbf{x})$$

We also need to update the model. Here we must consider the Gauss-Newton and Newton models separately.

$$\begin{aligned} \bar{m}_k^{GN}(\mathbf{s}) &= \frac{1}{2} \|\bar{\mathbf{r}}(\mathbf{x}_k) + \bar{\mathbf{J}}_k \mathbf{s}\|^2 \\ &= \frac{1}{2} (\bar{\mathbf{r}}^T \bar{\mathbf{r}} + 2\mathbf{s}^T \bar{\mathbf{J}}_k^T \bar{\mathbf{r}} + \mathbf{s}^T \bar{\mathbf{J}}_k^T \bar{\mathbf{J}}_k \mathbf{s}) \\ &= \frac{1}{2} \left(\mathbf{r}^T \mathbf{r} + \frac{2\sigma}{p} \|\mathbf{x}\|^p + 2(\mathbf{s}^T \mathbf{J}_k^T \mathbf{r} + \sigma \|\mathbf{x}\|^{p-2} \mathbf{s}^T \mathbf{x}) + \mathbf{s}^T \mathbf{J}_k^T \mathbf{J}_k \mathbf{s} + \frac{\sigma p}{2} \|\mathbf{x}\|^{p-4} (\mathbf{s}^T \mathbf{x})^2 \right) \\ &= m_k^{GN}(\mathbf{s}) + \sigma \left(\frac{1}{p} \|\mathbf{x}\|^p + \|\mathbf{x}\|^{p-2} \mathbf{s}^T \mathbf{x} + \frac{p}{4} \|\mathbf{x}\|^{p-4} (\mathbf{s}^T \mathbf{x})^2 \right). \end{aligned}$$

If we use a Newton model then

$$\begin{aligned} \bar{m}_k^N(\mathbf{s}) &= \bar{m}_k^{GN}(\mathbf{s}) + \frac{1}{2} \mathbf{s}^T \bar{\mathbf{H}}_k \mathbf{s} \\ &= \bar{m}_k^{GN}(\mathbf{s}) + \frac{1}{2} \mathbf{s}^T \left(\mathbf{H}_k + \sigma \|\mathbf{x}\|^{p-2} \left(I + \frac{\mathbf{x}\mathbf{x}^T}{\|\mathbf{x}\|^2} \right) \right) \mathbf{s} \\ &= \bar{m}_k^{GN}(\mathbf{s}) + \frac{1}{2} \mathbf{s}^T \mathbf{H}_k \mathbf{s} + \frac{\sigma}{2} \|\mathbf{x}\|^{p-4} \mathbf{s}^T (\mathbf{x}^T \mathbf{x} I + \mathbf{x}\mathbf{x}^T) \mathbf{s} \\ &= \bar{m}_k^{GN}(\mathbf{s}) + \frac{1}{2} \mathbf{s}^T \mathbf{H}_k \mathbf{s} + \frac{\sigma}{2} \|\mathbf{x}\|^{p-4} ((\mathbf{x}^T \mathbf{x})(\mathbf{s}^T \mathbf{s}) + (\mathbf{x}^T \mathbf{s})^2) \end{aligned}$$

1.4 Examples

Consider fitting the function $y(t) = x_1 e^{x_2 t}$ to data (\mathbf{t}, \mathbf{y}) using a non-linear least squares fit.

The residual function is given by

$$r_i(\mathbf{x}) = x_1 e^{x_2 t_i} - y_i.$$

We can calculate the Jacobian and Hessian of the residual as

$$\nabla r_i(\mathbf{x}) = \begin{pmatrix} e^{x_2 t_i} & t_i x_1 e^{x_2 t_i} \end{pmatrix},$$

$$\nabla^2 r_i(\mathbf{x}) = \begin{pmatrix} 0 & t_i e^{x_2 t_i} \\ t_i e^{x_2 t_i} & x_1 t_i^2 e^{x_2 t_i} \end{pmatrix}.$$

For some data points, $y_i, t_i, (i = 1, \dots, m)$ the user must return

$$\mathbf{r}(\mathbf{x}) = \begin{bmatrix} r_1(\mathbf{x}) \\ \vdots \\ r_m(\mathbf{x}) \end{bmatrix}, \quad \mathbf{J}(\mathbf{x}) = \begin{bmatrix} \nabla r_1(\mathbf{x}) \\ \vdots \\ \nabla r_m(\mathbf{x}) \end{bmatrix}, \quad Hf(\mathbf{x}) = \sum_{i=1}^m (\mathbf{r})_i \nabla^2 r_i(\mathbf{x}),$$

where, in the case of the Hessian, $(\mathbf{r})_i$ is the i th component of a residual vector passed to the user.

| | | | | | |
|-------|---|---|---|----|----|
| i | 1 | 2 | 3 | 4 | 5 |
| t_i | 1 | 2 | 4 | 5 | 8 |
| y_i | 3 | 4 | 6 | 11 | 20 |

and initial guess $\mathbf{x} = (2.5, 0.25)$, the following code performs the fit (with no weightings, i.e., $\mathbf{W} = \mathbf{I}$).

```
! examples/Fortran/nlls_example.f90
!
! Attempts to fit the model  $y_i = x_1 e^{x_2 t_i}$ 
! For parameters  $x_1$  and  $x_2$ , and input data  $(t_i, y_i)$ 
module fndef_example
  use ral_nlls_double, only : params_base_type
  implicit none

  integer, parameter :: wp = kind(0d0)

  type, extends(params_base_type) :: params_type
    real(wp), dimension(:), allocatable :: t ! The m data points  $t_i$ 
    real(wp), dimension(:), allocatable :: y ! The m data points  $y_i$ 
  end type

contains
  ! Calculate  $r_i(x; t_i, y_i) = x_1 e^{x_2 * t_i} - y_i$ 
  subroutine eval_r(status, n, m, x, r, params)
    integer, intent(out) :: status
    integer, intent(in) :: n
    integer, intent(in) :: m
    real(wp), dimension(*), intent(in) :: x
    real(wp), dimension(*), intent(out) :: r
    class(params_base_type), intent(inout) :: params

    real(wp) :: x1, x2

    x1 = x(1)
    x2 = x(2)
    select type(params)
    type is(params_type)
      r(1:m) = x1 * exp(x2*params%t(:)) - params%y(:)
```

```

    end select

    status = 0 ! Success
end subroutine eval_r
! Calculate:
! J_i1 = e^(x_2 * t_i)
! J_i2 = t_i x_1 e^(x_2 * t_i)
subroutine eval_J(status, n, m, x, J, params)
    integer, intent(out) :: status
    integer, intent(in) :: n
    integer, intent(in) :: m
    real(wp), dimension(*), intent(in) :: x
    real(wp), dimension(*), intent(out) :: J
    class(params_base_type), intent(inout) :: params

    real(wp) :: x1, x2

    x1 = x(1)
    x2 = x(2)
    select type(params)
    type is(params_type)
        J( 1: m) = exp(x2*params%t(1:m)) ! J_i1
        J(m+1:2*m) = params%t(1:m) * x1 * exp(x2*params%t(1:m)) ! J_i2
    end select

    status = 0 ! Success
end subroutine eval_J
! Calculate
! HF = sum_i r_i H_i
! Where H_i = [ 0 t_i e^(x_2 t_i) ]
! [ t_i e^(x_2 t_i) t_i^2 x_1 e^(x_2 t_i) ]
subroutine eval_HF(status, n, m, x, r, HF, params)
    integer, intent(out) :: status
    integer, intent(in) :: n
    integer, intent(in) :: m
    real(wp), dimension(*), intent(in) :: x
    real(wp), dimension(*), intent(in) :: r
    real(wp), dimension(*), intent(out) :: HF
    class(params_base_type), intent(inout) :: params

    real(wp) :: x1, x2

    x1 = x(1)
    x2 = x(2)
    select type(params)
    type is(params_type)
        HF( 1) = sum(r(1:m) * 0) ! H_11
        HF( 2) = sum(r(1:m) * params%t(1:m) * exp(x2*params%t(1:m))) ! H_21
        HF(1*n+1) = HF(2) ! H_12
        HF(1*n+2) = sum(r(1:m) * (params%t(1:m)**2) * x1 * exp(x2*params%t(1:m))) ! H_
↪22
    end select

    status = 0 ! Success
end subroutine eval_HF
end module fndef_example

program nlls_example

```

```

use ral_nlls_double
use indef_example
implicit none

type(nlls_options) :: options
type(nlls_inform) :: inform

integer :: m,n
real(wp), allocatable :: x(:)
type(params_type) :: params

! Data to be fitted
m = 5
allocate(params%t(m), params%y(m))
params%t(:) = (/ 1.0, 2.0, 4.0, 5.0, 8.0 /)
params%y(:) = (/ 3.0, 4.0, 6.0, 11.0, 20.0 /)

! Call fitting routine
n = 2
allocate(x(n))
x = (/ 2.5, 0.25 /) ! Initial guess
call nlls_solve(n, m, x, eval_r, eval_J, eval_HF, params, options, inform)
if(inform%status.ne.0) then
    print *, "ral_nlls() returned with error flag ", inform%status
    stop
endif

! Print result
print *, "Found a local optimum at x = ", x
print *, "Took ", inform%iter, " iterations"
print *, "      ", inform%f_eval, " function evaluations"
print *, "      ", inform%g_eval, " gradient evaluations"
print *, "      ", inform%h_eval, " hessian evaluations"
end program nlls_example

```

This returns the following output:

1.5 Indices and tables

- `genindex`
- `search`

E

eval_Hf() (fortran subroutine), [7](#)
eval_HP() (fortran subroutine), [8](#)
eval_J() (fortran subroutine), [6](#)
eval_r() (fortran subroutine), [6](#)

N

nlls_inform (fortran type), [12](#)
nlls_iterate() (fortran subroutine), [5](#)
nlls_options (fortran type), [8](#)
nlls_solve() (fortran subroutine), [4](#)
nlls_workspace (fortran type), [13](#)