

Package ‘cheapr’

May 14, 2025

Title Simple Functions to Save Time and Memory
Version 1.2.0
Maintainer Nick Christofides <nick.christofides.r@gmail.com>
Description Fast and memory-efficient (or 'cheap') tools to facilitate efficient programming, saving time and memory. It aims to provide 'cheaper' alternatives to common base R functions, as well as some additional functions.
License MIT + file LICENSE
BugReports <https://github.com/NicChr/cheapr/issues>
Depends R (>= 4.0.0)
Imports collapse (>= 2.0.0)
Suggests bench, data.table, testthat (>= 3.0.0)
LinkingTo cpp11
Config/testthat/edition 3
Encoding UTF-8
RoxygenNote 7.3.2
NeedsCompilation yes
Author Nick Christofides [aut, cre] (ORCID:
 <<https://orcid.org/0000-0002-9743-7342>>)
Repository CRAN
Date/Publication 2025-05-14 09:20:02 UTC

Contents

cheapr-package	2
address	3
as_discrete	3
attrs	6
bin	7
case	8

cheapr_c	9
cheapr_if_else	10
cheapr_table	10
copy	11
factor_	12
gcd	15
get_breaks	17
int_sign	19
is_na	19
lag_	22
list_lengths	26
named_list	28
new_df	28
overview	29
reconstruct	31
recycle	32
rep	33
sequence_	33
setdiff_	35
set_abs	38
sset	40
sset_df	42
val_count	43
which_	44
Index	46

cheapr-package	<i>cheapr: Simple Functions to Save Time and Memory</i>
----------------	---

Description

In this package, 'cheap' means fast and efficient.

cheapr aims to provide a set of functions for programmers to write cheaper code, saving time and memory.

Author(s)

Maintainer: Nick Christofides <nick.christofides.r@gmail.com> ([ORCID](#))

See Also

- Useful links:
- Report bugs at <https://github.com/NicChr/cheapr/issues>

address	<i>Memory address of R object</i>
---------	-----------------------------------

Description

Memory address of R object

Usage

```
address(x)
```

Arguments

x An R object.

Value

Memory address of R object.

as_discrete	<i>Turn continuous data into discrete bins</i>
-------------	--

Description

This is a cheapr version of `cut.numeric()` which is more efficient and prioritises pretty-looking breaks by default through the use of `get_breaks()`. Out-of-bounds values can be included naturally through the `include_oob` argument. Left-closed (right-open) intervals are returned by default in contrast to `cut`'s default right-closed intervals. Furthermore there is flexibility in formatting the interval bins, allowing the user to specify formatting functions and symbols for the interval close and open symbols.

Usage

```
as_discrete(x, ...)
```

```
## S3 method for class 'numeric'
```

```
as_discrete(
  x,
  breaks = if (left_closed) get_breaks(x) else cheapr_rev(-get_breaks(-x)),
  left_closed = TRUE,
  include_endpoint = FALSE,
  include_oob = FALSE,
  ordered = FALSE,
  intv_start_fun = prettyNum,
  intv_end_fun = prettyNum,
  intv_closers = c("[", "]"),
```

```

    intv_openers = c("(", ")"),
    intv_sep = ", ",
    inf_label = NULL,
    ...
)

## S3 method for class 'integer64'
as_discrete(x, ...)

```

Arguments

<code>x</code>	A numeric vector.
<code>...</code>	Extra arguments passed onto methods.
<code>breaks</code>	Break-points. The default option creates pretty looking breaks. Unlike <code>cut()</code> , the <code>breaks</code> arg cannot be a number denoting the number of breaks you want. To generate breakpoints this way use <code>get_breaks()</code> .
<code>left_closed</code>	Left-closed intervals or right-closed intervals?
<code>include_endpoint</code>	Include endpoint? Default is FALSE.
<code>include_oob</code>	Include out-of-bounds values? Default is FALSE. This is equivalent to <code>breaks = c(breaks, Inf)</code> or <code>breaks = c(-Inf, breaks)</code> when <code>left_closed = FALSE</code> . If <code>include_endpoint = TRUE</code> , the endpoint interval is prioritised before the out-of-bounds interval. This behaviour cannot be replicated easily with <code>cut()</code> . For example, these 2 expressions are not equivalent: <pre>cut(10, c(9, 10, Inf), right = F, include.lowest = T) != as_discrete(10, c(9, 10), include_endpoint = T, include_oob = T)</pre>
<code>ordered</code>	Should result be an ordered factor? Default is FALSE.
<code>intv_start_fun</code>	Function used to format interval start points.
<code>intv_end_fun</code>	Function used to format interval end points.
<code>intv_closers</code>	A length 2 character vector denoting the symbol to use for closing either left or right closed intervals.
<code>intv_openers</code>	A length 2 character vector denoting the symbol to use for opening either left or right closed intervals.
<code>intv_sep</code>	A length 1 character vector used to separate the start and end points.
<code>inf_label</code>	Label to use for intervals that include infinity. If left NULL the Unicode infinity symbol is used.

Value

A factor of discrete bins (intervals of start/end pairs).

See Also

[bin get_breaks](#)

Examples

```
library(cheapr)

# `as_discrete()` is very similar to `cut()`
# but more flexible as it allows you to supply
# formatting functions and symbols for the discrete bins

# Here is an example of how to use the formatting functions to
# categorise age groups nicely

ages <- 1:100

age_group <- function(x, breaks){
  age_groups <- as_discrete(
    x,
    breaks = breaks,
    intv_sep = "-",
    intv_end_fun = function(x) x - 1,
    intv_openers = c("", ""),
    intv_closers = c("", ""),
    include_oob = TRUE,
    ordered = TRUE
  )

  # Below is just renaming the last age group

  lvls <- levels(age_groups)
  n_lvls <- length(lvls)
  max_ages <- paste0(max(breaks), "+")
  attr(age_groups, "levels") <- c(lvls[-n_lvls], max_ages)
  age_groups
}

age_group(ages, seq(0, 80, 20))
age_group(ages, seq(0, 25, 5))
age_group(ages, 5)

# To closely replicate `cut()` with `as_discrete()` we can use the following

cheapr_cut <- function(x, breaks, right = TRUE,
  include.lowest = FALSE,
  ordered.result = FALSE){
  if (length(breaks) == 1){
    breaks <- get_breaks(x, breaks, pretty = FALSE,
      expand_min = FALSE, expand_max = FALSE)
    adj <- diff(range(breaks)) * 0.001
    breaks[1] <- breaks[1] - adj
    breaks[length(breaks)] <- breaks[length(breaks)] + adj
  }
  as_discrete(x, breaks, left_closed = !right,
    include_endpoint = include.lowest,
    ordered = ordered.result,
```

```

    intv_start_fun = function(x) formatC(x, digits = 3, width = 1),
    intv_end_fun = function(x) formatC(x, digits = 3, width = 1))
  }

x <- rnorm(100)
cheapr_cut(x, 10)
identical(cut(x, 10), cheapr_cut(x, 10))

```

 attrs

Add and remove attributes

Description

Simple tools to add and remove attributes, both normally and in-place. To remove specific attributes, set those attributes to NULL.

Usage

```
attrs_add(x, ..., .set = FALSE, .args = NULL)
```

```
attrs_rm(x, .set = FALSE)
```

Arguments

x	Object to add/remove attributes.
...	Named attributes, e.g 'key = value'.
.set	Should attributes be added in-place without shallow-copying x? Default is FALSE.
.args	An alternative to ... for easier programming with lists.

Value

The object x with attributes removed or added.

See Also

[shallow_copy](#)

bin	<i>A sometimes cheaper but argument richer alternative to .bincode()</i>
-----	--

Description

When `x` is an integer vector, `bin()` is cheaper than `.bincode()` as no coercion to a double vector occurs. This alternative also has more arguments that allow you to return the start values of the binned vector, as well as including out-of-bounds intervals.

Usage

```
bin(  
  x,  
  breaks,  
  left_closed = TRUE,  
  include_endpoint = FALSE,  
  include_oob = FALSE,  
  codes = TRUE  
)
```

Arguments

<code>x</code>	A numeric vector.
<code>breaks</code>	A numeric vector of breaks.
<code>left_closed</code>	Should intervals be left-closed (and right-open)? Default is TRUE. If FALSE they are left-open (and right-closed).
<code>include_endpoint</code>	Equivalent to <code>include.lowest</code> in <code>?bincode</code> .
<code>include_oob</code>	Should out-of-bounds interval be included? Default is FALSE. This is the equivalent of adding <code>Inf</code> as the last value of the breaks, or <code>-Inf</code> as the first value of the breaks if <code>left_closed = FALSE</code> . When TRUE, this essentially becomes <code>findInterval()</code> .
<code>codes</code>	Should an integer vector indicating which bin the values fall into be returned? Default is TRUE. If FALSE the start values of the respective bin intervals are returned, i.e the corresponding breaks.

Value

Either an integer vector of codes indicating which bin the values fall into, or the start of the intervals for which each value falls into.

See Also

[get_breaks as_discrete](#)

case	<i>A cheapr case-when and switch</i>
------	--------------------------------------

Description

case and val_match are cheaper alternatives to dplyr::case_when and dplyr::case_match respectively.

Usage

```
case(..., .default = NULL)
```

```
val_match(.x, ..., .default = NULL)
```

Arguments

...	Logical expressions or scalar values in the case of val_match.
.default	Catch-all value or vector.
.x	Vector used to switch values.

Details

val_match() is a very efficient special case of the case() function when all lhs expressions are scalars, i.e. length-1 vectors. RHS expressions can be vectors the same length as .x. The below 2 expressions are equivalent.

```
val_match(
  x,
  1 ~ "one",
  2 ~ "two",
  .default = "Unknown"
)
case(
  x == 1 ~ "one",
  x == 2 ~ "two",
  .default = "Unknown"
)
```

Value

A vector the same length as .x or same length as the first condition in the case of case, unless the condition length is smaller than the rhs, in which case the length of the rhs is used.

See Also

[cheapr_if_else](#)

cheapr_c	<i>A cheapr version of c()</i>
----------	--------------------------------

Description

cheapr's version of `c()`. It is quite a bit faster for atomic vectors and combines data frame rows instead of cols.

Usage

```
cheapr_c(..., .args = NULL)
```

Arguments

<code>...</code>	Objects to combine.
<code>.args</code>	An alternative to <code>...</code> for easier programming with lists.

Value

Combined objects.

Examples

```
library(cheapr)

# Combine just like `c()`
cheapr_c(1, 2, 3:5)

# It combines rows by default instead of cols
cheapr_c(new_df(x = 1:3), new_df(x = 4:10))

# If you have a list of objects you want to combine
# use `.args` instead of `do.call` as it's more efficient

list_of_objs <- rep(list(0), 10^4)

bench::mark(
  do.call(cheapr_c, list_of_objs),
  cheapr_c(.args = list_of_objs)
)
```

cheapr_if_else	<i>Cheaper version of ifelse()</i>
----------------	------------------------------------

Description

Cheaper version of ifelse()

Usage

```
cheapr_if_else(condition, true, false, na = false[NA_integer_])
```

Arguments

condition	logical A condition which will be used to evaluate the if else operation.
true	Value(s) to replace TRUE instances.
false	Value(s) to replace FALSE instances.
na	Catch-all value(s) to replace all other instances, where is.na(condition).

Value

A vector the same length as condition, using a common type between true, false and default.

See Also

[case_val_match](#)

cheapr_table	<i>Fast frequency tables - Still experimental</i>
--------------	---

Description

This is not a one-to-one copy of `base::table()` as some behaviours differ. It is more flexible as it accepts inputs such as data frames and `vctrs_rcrd` objects.

Usage

```
cheapr_table(
  ...,
  names = TRUE,
  order = FALSE,
  na_exclude = FALSE,
  classed = FALSE
)
```

Arguments

<code>...</code>	<code>>=1</code> objects that can be converted to a factor through <code>cheapr::factor_()</code> .
<code>names</code>	Should level names be kept? Default is <code>TRUE</code> .
<code>order</code>	Should result be ordered by level names? Default is <code>FALSE</code> .
<code>na_exclude</code>	Should NA values be excluded? Default is <code>FALSE</code> .
<code>classed</code>	Should a table object be returned? Default is <code>FALSE</code>

Value

A named integer vector if one object is supplied, otherwise an array.

copy

Copy R objects

Description

`shallow_copy()` and `deep_copy()` are just wrappers to the R C API functions `Rf_shallow_duplicate()` and `Rf_duplicate()` respectively. `semi_copy()` is something in between whereby it fully copies the data but only shallow copies the attributes.

Usage

```
shallow_copy(x)
```

```
semi_copy(x)
```

```
deep_copy(x)
```

Arguments

`x` An object to shallow, semi, or deep copy.

Details

Shallow duplicates are mainly useful for adding attributes to objects in-place as well assigning vectors to shallow copied lists in-place.

Deep copies are generally useful for ensuring an object is fully duplicated, including all attributes associated with it. Deep copies are generally expensive and should be used with care.

`semi_copy()` deep copies everything except the attributes. This is experimental but in theory should be much more efficient and generally preferred to `deep_copy()`.

To summarise:

- `shallow_copy` - Shallow copies data and attributes
- `semi_copy` - Deep copies data and shallow copies attributes
- `deep_copy` - Deep copies both data and attributes

It is recommended to use these functions only if you know what you are doing.

Value

A shallow, semi or deep copied R object.

Examples

```
library(cheapr)
library(bench)
df <- new_df(x = sample.int(10^4))

# Note the memory allocation
mark(shallow_copy(df), iterations = 1)
mark(deep_copy(df), iterations = 1)

# In both cases the address of df changes

address(df);address(shallow_copy(df));address(deep_copy(df))

# When shallow-copying attributes are not duplicated

address(attr(df, "names"));address(attr(shallow_copy(df), "names"))

# They are when deep-copying

address(attr(df, "names"));address(attr(deep_copy(df), "names"))

# Adding an attribute in place with and without shallow copy
invisible(attrs_add(df, key = TRUE, .set = TRUE))
attr(df, "key")

# Remove attribute in-place
invisible(attrs_add(df, key = NULL, .set = TRUE))

# With shallow copy
invisible(attrs_add(shallow_copy(df), key = TRUE, .set = TRUE))

# 'key' attr was only added to the shallow copy, and not the original df
attr(df, "key")
```

factor_

A cheaper version of factor() along with cheaper utilities

Description

A fast version of factor() using the collapse package.

There are some additional utilities, most of which begin with the prefix 'levels_', such as as_factor() which is an efficient way to coerce both vectors and factors, levels_factor() which returns the levels of a factor, as a factor, levels_used() which returns the used levels of a factor, levels_unused()

which returns the unused levels of a factor, `levels_add()` adds the specified levels onto the existing levels, `levels_rm()` removes the specified levels, `levels_add_na()` which adds an explicit NA level, `levels_drop_na()` which drops the NA level, `levels_drop()` which drops unused factor levels, `levels_rename()` for renaming levels, `levels_lump()` which returns top n levels and lumps all others into the same category, `levels_count()` which returns the counts of each level, and finally `levels_reorder()` which reorders the levels of x based on y using the ordered median values of y for each level.

Usage

```
factor_(
  x = integer(),
  levels = NULL,
  order = TRUE,
  na_exclude = TRUE,
  ordered = is.ordered(x)
)

as_factor(x)

levels_factor(x)

levels_used(x)

levels_unused(x)

levels_rm(x, levels)

levels_add(x, levels, where = c("last", "first"))

levels_add_na(x, name = NA, where = c("last", "first"))

levels_drop_na(x)

levels_drop(x)

levels_reorder(x, order_by, decreasing = FALSE)

levels_rename(x, ..., .fun = NULL)

levels_lump(
  x,
  n,
  prop,
  other_category = "Other",
  ties = c("min", "average", "first", "last", "random", "max")
)

levels_count(x)
```

Arguments

<code>x</code>	A vector.
<code>levels</code>	Optional factor levels.
<code>order</code>	Should factor levels be sorted? Default is TRUE. It typically is faster to set this to FALSE, in which case the levels are sorted by order of first appearance.
<code>na_exclude</code>	Should NA values be excluded from the factor levels? Default is TRUE.
<code>ordered</code>	Should the result be an ordered factor?
<code>where</code>	Where should NA level be placed? Either first or last.
<code>name</code>	Name of NA level.
<code>order_by</code>	A vector to order the levels of <code>x</code> by using the medians of <code>order_by</code> .
<code>decreasing</code>	Should the reordered levels be in decreasing order? Default is FALSE.
<code>...</code>	Key-value pairs where the key is the new name and value is the name to replace that with the new name. For example <code>levels_rename(x, new = old)</code> replaces the level "old" with the level "new".
<code>.fun</code>	Renaming function applied to each level.
<code>n</code>	Top n number of levels to calculate.
<code>prop</code>	Top proportion of levels to calculate. This is a proportion of the total unique levels in <code>x</code> .
<code>other_category</code>	Name of 'other' category.
<code>ties</code>	Ties method to use. See <code>?rank</code> .

Details

This operates similarly to `collapse::qF()`.

The main difference internally is that `collapse::funique()` is used and therefore s3 methods can be written for it.

Furthermore, for date-times `factor_` differs in that it differentiates all instances in time whereas `factor` differentiates calendar times. Using a daylight savings example where the clocks go back: `factor(as.POSIXct(1729984360, tz = "Europe/London") + 3600 * (1:5))` produces 4 levels whereas `factor_(as.POSIXct(1729984360, tz = "Europe/London") + 3600 * (1:5))` produces 5 levels.

`levels_lump()` is a cheaper version of `forcats::lump_n()` but returns levels in order of highest frequency to lowest. This can be very useful for plotting.

Value

A factor or character in the case of `levels_used` and `levels_unused`. `levels_count` returns a data frame of counts and proportions for each level.

Examples

```
library(cheapr)

x <- factor_(sample(letters[sample.int(26, 10)]), 100, TRUE), levels = letters)
x
```

```

# Used/unused levels

levels_used(x)
levels_unused(x)

# Drop unused levels
levels_drop(x)

# Top 3 letters by frequency
lumped_letters <- levels_lump(x, 3)
levels_count(lumped_letters)

# To remove the "other" category, use `levels_rm()`
levels_count(levels_rm(lumped_letters, "Other"))

# We can use levels_lump to create a generic top n function for non-factors too

get_top_n <- function(x, n){
  f <- levels_lump(factor_(x, order = FALSE), n = n)
  levels_count(f)
}

get_top_n(x, 3)

# A neat way to order the levels of a factor by frequency
# is the following:

levels(levels_lump(x, prop = 1)) # Highest to lowest
levels(levels_lump(x, prop = -1)) # Lowest to highest

```

gcd

Greatest common divisor and smallest common multiple

Description

Fast greatest common divisor and smallest common multiple using the Euclidean algorithm.

gcd() returns the greatest common divisor.

scm() returns the smallest common multiple.

gcd2() is a vectorised binary version of gcd.

scm2() is a vectorised binary version of scm.

Usage

```

gcd(
  x,
  tol = sqrt(.Machine$double.eps),
  na_rm = TRUE,

```

```

    round = TRUE,
    break_early = TRUE
  )

scm(x, tol = sqrt(.Machine$double.eps), na_rm = TRUE)

gcd2(x, y, tol = sqrt(.Machine$double.eps), na_rm = TRUE)

scm2(x, y, tol = sqrt(.Machine$double.eps), na_rm = TRUE)

```

Arguments

x	A numeric vector.
tol	Tolerance. This must be a single positive number strictly less than 1.
na_rm	If TRUE the default, NA values are ignored.
round	If TRUE the output is rounded as <code>round(gcd, digits)</code> where <code>digits</code> is <code>ceiling(abs(log10(tol))) + 1</code> . This can potentially reduce floating point errors on further calculations. The default is TRUE.
break_early	This is experimental and applies only to floating-point numbers. When TRUE the algorithm will end once <code>gcd > 0 && gcd < 2 * tol</code> . This can offer a tremendous speed improvement. If FALSE the algorithm finishes once it has gone through all elements of x. The default is TRUE. For integers, the algorithm always breaks early once <code>gcd > 0 && gcd <= 1</code> .
y	A numeric vector.

Details

Method:

GCD (Greatest Common Divisor):

The GCD is calculated using a binary function that takes input `GCD(gcd, x[i + 1])` where the output of this function is passed as input back into the same function iteratively along the length of x. The first gcd value is `x[1]`.

Zeroes are handled in the following way:

`GCD(0, 0) = 0`

`GCD(a, 0) = a`

This has the nice property that zeroes are essentially ignored.

SCM (Smallest Common Multiple):

This is calculated using the GCD and the formula is:

$SCM(x, y) = (abs(x) / GCD(x, y)) * abs(y)$

If you want to calculate the gcd & lcm for 2 values or across 2 vectors of values, use `gcd2` and `scm2`.

A note on performance:

A very common solution to finding the GCD of a vector of values is to use `Reduce()` along with a binary function like `gcd2()`.

e.g. `Reduce(gcd2, seq(5, 20, 5))`.

This is exactly identical to `gcd(seq(5, 20, 5))`, with `gcd()` being much faster and overall cheaper as it is written in C++ and heavily optimised. Therefore it is recommended to always use `gcd()`.

For example we can compare the two approaches below,

```
x <- seq(5L, length = 10^6, by = 5L)
bench::mark(Reduce(gcd2, x), gcd(x))
```

This example code shows `gcd()` being ~200x faster on my machine than the `Reduce + gcd2` approach, even though `gcd2` itself is written in C++ and has little overhead.

Value

A number representing the GCD or SCM.

Examples

```
library(cheapr)
library(bench)

# Binary versions
gcd2(15, 25)
gcd2(15, seq(5, 25, 5))
scm2(15, seq(5, 25, 5))
scm2(15, 25)

# GCD across a vector
gcd(c(0, 5, 25))
mark(gcd(c(0, 5, 25)))

x <- rnorm(10^5)
gcd(x)
gcd(x, round = FALSE)
mark(gcd(x))
```

get_breaks

Pretty break-points for continuous (numeric) data

Description

The distances between break-points are always equal in this implementation.

Usage

```
get_breaks(x, n = 10, ...)

## Default S3 method:
get_breaks(x, n = 10, ...)

## S3 method for class 'numeric'
```

```

get_breaks(
  x,
  n = 10,
  pretty = TRUE,
  expand_min = FALSE,
  expand_max = pretty,
  ...
)

## S3 method for class 'integer64'
get_breaks(x, n = 10, ...)

```

Arguments

x	A numeric vector.
n	Number of breakpoints. You may get less or more than requested.
...	Extra arguments passed onto methods.
pretty	Should pretty break-points be prioritised? Default is TRUE. If FALSE bin-widths will be calculated as <code>diff(range(x)) / n</code> .
expand_min	Should smallest break be extended beyond the minimum of the data? Default is FALSE. If TRUE then <code>min(get_breaks(x))</code> is ensured to be less than <code>min(x)</code> .
expand_max	Should largest break be extended beyond the maximum of the data? Default is TRUE. If TRUE then <code>max(get_breaks(x))</code> is ensured to be greater than <code>max(x)</code> .

Value

A numeric vector of break-points.

See Also

[bin as_discrete](#)

Examples

```

library(cheapr)

set.seed(123)
ages <- sample(0:80, 100, TRUE)

# Pretty
get_breaks(ages, n = 10)
# Not-pretty
# bin-width is diff(range(ages)) / n_breaks
get_breaks(ages, n = 10, pretty = FALSE)

# `get_breaks()` is left-biased in a sense, meaning that
# the first break is always <= min(x) but the last break
# may be < max(x)

```

```
# To get right-biased breaks we can use a helper like so..

right_breaks <- function(x, ...){
  -get_breaks(-x, ...)
}

get_breaks(4:24, 10)
right_breaks(4:24, 10)

# Use `rev()` to ensure they are in ascending order
rev(right_breaks(4:24, 10))
```

int_sign	<i>A fast and integer-based sign()</i>
----------	--

Description

A fast and integer-based sign()

Usage

```
int_sign(x)
```

Arguments

x	Integer or double vector.
---	---------------------------

Value

An integer vector denoting the sign, -1 for negatives, 1 for positives and 0 for when $x == 0$.

is_na	<i>Efficient functions for dealing with missing values.</i>
-------	---

Description

is_na() is a parallelised alternative to is.na().
 num_na(x) is a faster and more efficient sum(is.na(x)).
 which_na(x) is a more efficient which(is.na(x))
 which_not_na(x) is a more efficient which(!is.na(x))
 row_na_counts(x) is a more efficient rowSums(is.na(x))
 row_all_na() returns a logical vector indicating which rows are empty and have only NA values.
 row_any_na() returns a logical vector indicating which rows have at least 1 NA value.
 The col_ variants are the same, but operate by-column.

Usage

```
is_na(x)

## Default S3 method:
is_na(x)

## S3 method for class 'POSIXlt'
is_na(x)

## S3 method for class 'vctrs_rcrd'
is_na(x)

## S3 method for class 'data.frame'
is_na(x)

num_na(x, recursive = TRUE)

which_na(x)

which_not_na(x)

any_na(x, recursive = TRUE)

all_na(x, recursive = TRUE)

row_na_counts(x, names = FALSE)

col_na_counts(x, names = FALSE)

row_all_na(x, names = FALSE)

col_all_na(x, names = FALSE)

row_any_na(x, names = FALSE)

col_any_na(x, names = FALSE)
```

Arguments

x	A vector, list, data frame or matrix.
recursive	Should the function be applied recursively to lists? The default is TRUE. Setting this to TRUE is actually much cheaper because when FALSE, the other NA functions rely on calling <code>is_na()</code> , therefore allocating a vector. This is so that alternative objects with <code>is.na</code> methods can be supported.
names	Should row/col names be added?

Details

These functions are designed primarily for programmers, to increase the speed and memory-efficiency of NA handling.

Most of these functions can be parallelised through `options(cheapr.cores)`.

Common use-cases:

To replicate `complete.cases(x)`, use `!row_any_na(x)`.

To find rows with any empty values, use `which_(row_any_na(df))`.

To find empty rows use `which_(row_all_na(df))` or `which_na(df)`. To drop empty rows use `na_rm(df)` or `sset(df, which_(row_all_na(df), TRUE))`.

is_na:

`is_na` is an S3 generic function. It will internally fall back on using `is.na` if it can't find a suitable method. Alternatively you can write your own `is_na` method. For example there is a method for `vctrs_rcrd` objects that simply converts it to a data frame and then calls `row_all_na()`. There is also a `POSIXlt` method for `is_na` that is much faster than `is.na`.

Lists:

When `x` is a list, `num_na`, `any_na` and `all_na` will recursively search the list for NA values. If `recursive = F` then `is_na()` is used to find NA values.

`is_na` differs to `is.na` in 2 ways:

- List elements are counted as NA if either that value is NA, or if it's a list, then all values of that list are NA.
- When called on a data frame, it returns TRUE for empty rows that contain only NA values.

Value

Number or location of NA values.

Examples

```
library(cheapr)
library(bench)

x <- 1:10
x[c(1, 5, 10)] <- NA
num_na(x)
which_na(x)
which_not_na(x)

row_nas <- row_na_counts(airquality, names = TRUE)
col_nas <- col_na_counts(airquality, names = TRUE)
row_nas
col_nas

df <- sset(airquality, j = 1:2)

# Number of NAs in data
num_na(df)
```

```
# Which rows are empty?
row_na <- row_all_na(df)
sset(df, row_na)

# Removing the empty rows
sset(df, which_(row_na, invert = TRUE))
# Or
na_rm(df)
# Or
sset(df, row_na_counts(df) < ncol(df))
```

lag_

Lagged operations.

Description

Fast lags and leads optionally using dynamic vectorised lags, ordering and run lengths.

Usage

```
lag_(x, n = 1L, fill = NULL, set = FALSE, recursive = TRUE)
```

```
lag2_(
  x,
  n = 1L,
  order = NULL,
  run_lengths = NULL,
  fill = NULL,
  recursive = TRUE
)
```

Arguments

x	A vector or data frame.
n	Number of lags. Negative values are accepted. lag2_ accepts a vector of dynamic lags and leads which gets recycled to the length of x.
fill	Value used to fill first n values. Default is NA.
set	Should x be updated by reference? If TRUE no copy is made and x is updated in place. The default is FALSE.
recursive	Should list elements be lagged as well? If TRUE, this is useful for data frames and will return row lags. If FALSE this will return a plain lagged list.
order	Optionally specify an ordering with which to apply the lags. This is useful for example when applying lags chronologically using an unsorted time variable.

run_lengths Optional integer vector of run lengths that defines the size of each lag run. For example, supplying `c(5, 5)` applies lags to the first 5 elements and then essentially resets the bounds and applies lags to the next 5 elements as if they were an entirely separate and standalone vector. This is particularly useful in conjunction with the `order` argument to perform a by-group lag. See the examples for details.

Details

For most applications, it is more efficient and recommended to use `lag_()`. For anything that requires dynamic lags, lag by order of another variable, or by-group lags, one can use `lag2_()`. To do cyclic lags, see the examples below for an implementation.

lag2_:

`lag2_` is a generalised form of `lag_` that by default performs simple lags and leads. It has 3 additional features but does not support updating by reference or long vectors.

These extra features include:

- `n` - This shares the same name as the `n` argument in `lag_` for consistency. The difference is that `lag_` accepts a lag vector of length 1 whereas this accepts a vector of dynamic lags allowing for flexible combinations of variable sized lags and leads. These are recycled to the length of the data and will always align with the data, meaning that if you supply a custom `order` argument, this ordering is applied both to `x` and the recycled lag vector `n` simultaneously.
- `order` - Apply lags in any order you wish. This can be useful for reverse order lags, lags against unsorted time variables, and by-group lags.
- `run_lengths` - Specify the size of individual lag runs. For example, if you specify `run_lengths = c(3, 4, 2)`, this will apply your lags to the first 3 elements and then reset, applying lags to the next 4 elements, to reset again and apply lags to the final 2 elements. Each time the reset occurs, it treats each run length sized 'chunk' as a unique and separate vector. See the examples for a showcase.

Table of differences between `lag_` and `lag2_`:

Description	<code>lag_</code>	<code>lag2_</code>
Lags	Yes	Yes
Leads	Yes	Yes
Long vector support	Yes	No
Lag by reference	Yes	No
Dynamic vectorised lags	No	Yes
Data frame row lags	Yes	Yes
Alternative order lags	No	Yes

Value

A lagged object the same size as `x`.

Examples

```

library(cheapr)
library(bench)

# A use-case for data.table
# Adding 0 because can't update ALTREP by reference
df <- data.frame(x = 1:10^5 + 0L)

# Normal data frame lag
sset(lag_(df), 1:10)

# Lag these behind by 3 rows
sset(lag_(df, 3, set = TRUE), 1:10)

df$x[1:10] # x variable was updated by reference!

# The above can be used naturally in data.table to lag data
# without any copies

# To perform regular R row lags, just make sure set is `FALSE`

sset(lag_(as.data.frame(EuStockMarkets), 5), 1:10)

# lag2_ is a generalised version of lag_ that allows
# for much more complex lags

x <- 1:10

# lag every 2nd element
lag2_(x, n = c(1, 0)) # lag vector is recycled

# Explicit Lag(3) using a vector of lags
lags <- lag_sequence(length(x), 3, partial = FALSE)
lag2_(x, n = lags)

# Alternating lags and leads
lag2_(x, c(1, -1))

# Lag only the 3rd element
lags <- integer(length(x))
lags[3] <- 1L
lag2_(x, lags)

# lag in descending order (same as a lead)

lag2_(x, order = 10:1)

# lag that resets after index 5
lag2_(x, run_lengths = c(5, 5))

# lag with a time index
years <- sample(2011:2020)

```



```

lag2_(x, order = order(years))

# Example of how to do a cyclical lag
n <- length(x)

# When k >= 0
k <- min(3, n)
lag2_(x, c(rep(-n + k, k), rep(k, n - k)))
# When k < 0
k <- max(-3, -n)
lag2_(x, c(rep(k, n + k), rep(n + k, -k)))

# As it turns out, we can do a grouped lag
# by supplying group sizes as run lengths and group order as the order

set.seed(45)
g <- sample(c("a", "b"), 10, TRUE)

# NOTE: collapse::flag will not work unless g is already sorted!
# This is not an issue with lag2_()
collapse::flag(x, g = g)
lag2_(x, order = order(g), run_lengths = collapse::GRP(g)$group.sizes)

# For production code, we can of course make
# this more optimised by using collapse::radixorderv()
# Which calculates the order and group sizes all at once

o <- collapse::radixorderv(g, group.sizes = TRUE)
lag2_(x, order = o, run_lengths = attr(o, "group.sizes"))

# Let's finally wrap this up in a nice grouped-lag function
grouped_lag <- function(x, n = 1, g = integer(length(x))) {
  o <- collapse::radixorderv(g, group.sizes = TRUE, sort = FALSE)
  lag2_(x, n, order = o, run_lengths = attr(o, "group.sizes"))
}

# And voila!
grouped_lag(x, g = g)

# A method to extract this information from dplyr

## We can actually get this information easily from a `grouped_df` object
## Uncomment the below code to run the implementation
# library(dplyr)
# library(timeplyr)
# eu_stock <- EuStockMarkets |>
#   ts_as_tibble() |>
#   group_by(stock_index = group)
# groups <- group_data(eu_stock) # Group information
# group_order <- unlist(groups$.rows) # Order of groups
# group_sizes <- lengths_(groups$.rows) # Group sizes
#

```

```

# # by-stock index lag
# lag2_(eu_stock$value, order = group_order, run_lengths = group_sizes)
#
# # Verifying this output is correct
# eu_stock |>
#   ungroup() |>
#   mutate(lag1 = lag_(value), .by = stock_index) |>
#   mutate(lag2 = lag2_(value, order = group_order, run_lengths = group_sizes)) |>
#   summarise(lags_are_equal = identical(lag1, lag2))

# Let's compare this to data.table

library(data.table)
default_threads <- getDTthreads()
setDTthreads(1)
dt <- data.table(x = 1:10^5,
                 g = sample.int(10^4, 10^5, TRUE))

bench::mark(dt[, y := shift(x), by = g][[["y"]],
             grouped_lag(dt$x, g = dt$g),
             iterations = 10)
setDTthreads(default_threads)

```

list_lengths

List utilities

Description

Functions to help work with lists.

Usage

```

list_lengths(x, names = FALSE)

lengths_(x, names = FALSE)

unlisted_length(x)

new_list(length = 0L, default = NULL)

list_assign(x, values)

list_combine(..., .args = NULL)

list_drop_null(x)

```

Arguments

<code>x</code>	A list.
<code>names</code>	Should names of list elements be added? Default is FALSE.
<code>length</code>	Length of list.
<code>default</code>	Default value for each list element.
<code>values</code>	A named list
<code>...</code>	Objects to combine into a list.
<code>.args</code>	An alternative to <code>...</code> for easier programming with lists.

Value

`list_lengths()` returns the list lengths.

`unlisted_length()` is a fast alternative to `length(unlist(x))`.

`new_list()` is like `vector("list", length)` but also allows you to specify a default value for each list element. This can be useful for initialising with a catch-all value so that when you unlist you're guaranteed a list of length \geq to the specified length.

`list_assign()` is vectorised version of `[[<-` that concatenates values to `x` or modifies `x` where the names match. Can be useful for modifying data frame variables.

`list_combine()` combines each element of a set of lists into a single list. If an element is not a list, it is treated as a length-one list. This happens to be very useful for combining data frame cols.

`list_drop_null()` removes NULL list elements very quickly.

Examples

```
library(cheapr)
l <- list(1:10,
         NULL,
         list(integer(), NA_integer_, 2:10))

lengths_(l) # Faster lengths()
unlisted_length(l) # length of vector if we unlist
paste0("length: ", length(print(unlist(l))))

unlisted_length(l) - na_count(l) # Number of non-NA elements

# We can create and initialise a new list with a default value
l <- new_list(20, 0L)
l[1:5]
# This works well with vctrs_list_of objects
```

named_list	<i>Turn dot-dot-dot (...) into a named list</i>
------------	---

Description

A fast and useful function for always returning a named list from ...

Usage

```
named_list(..., .keep_null = TRUE)
```

Arguments

- ... Key-value pairs.
- .keep_null Should NULL entries be kept? Default is TRUE.

Value

A named list.

new_df	<i>Cheap data frame utilities</i>
--------	-----------------------------------

Description

Cheap data frame utilities

Usage

```
new_df(..., .nrows = NULL, .recycle = TRUE, .name_repair = TRUE, .args = NULL)

as_df(x)

fast_df(..., .args = NULL)

list_as_df(x)

name_repair(x, dup_sep = "_", empty_sep = "col_")

unique_name_repair(x, dup_sep = "_", empty_sep = "col_")

col_c(..., .recycle = TRUE, .name_repair = TRUE, .args = NULL)

row_c(..., .args = NULL)
```

Arguments

<code>...</code>	Key-value pairs.
<code>.nrows</code>	<code>[integer(1)]</code> - (Optional) number of rows. Commonly used to initialise a 0-column data frame with rows.
<code>.recycle</code>	<code>[logical(1)]</code> - Should arguments be recycled? Default is TRUE.
<code>.name_repair</code>	<code>[logical(1)]</code> - Should duplicate and empty names repaired and made unique? Default is TRUE.
<code>.args</code>	An alternative to <code>...</code> for easier programming with lists.
<code>x</code>	An object to coerce to a <code>data.frame</code> or a character vector for <code>unique_name_repair()</code> .
<code>dup_sep</code>	<code>[character(1)]</code> A separator to use between duplicate column names and their locations. Default is <code>'_'</code>
<code>empty_sep</code>	<code>[character(1)]</code> A separator to use between the empty column names and their locations. Default is <code>'col_'</code>

Details

`fast_df()` is a very fast bare-bones version of `new_df()` that performs no checks and no recycling or name tidying, making it appropriate for very tight loops.

Value

A `data.frame`.
`name_repair` takes a character vector and returns unique strings by appending duplicate string locations to the duplicates. This is mostly used to create unique col names.

 overview

An alternative to `summary()` inspired by the `skimr` package

Description

A cheaper `summary()` function, designed for larger data.

Usage

```
overview(x, digits = getOption("cheapr.digits", 2), ...)

## Default S3 method:
overview(x, digits = getOption("cheapr.digits", 2), ...)

## S3 method for class 'logical'
overview(x, digits = getOption("cheapr.digits", 2), ...)

## S3 method for class 'integer'
overview(x, digits = getOption("cheapr.digits", 2), hist = TRUE, ...)
```

```
## S3 method for class 'numeric'
overview(x, digits = getOption("cheapr.digits", 2), hist = TRUE, ...)

## S3 method for class 'integer64'
overview(x, digits = getOption("cheapr.digits", 2), hist = TRUE, ...)

## S3 method for class 'character'
overview(x, digits = getOption("cheapr.digits", 2), ...)

## S3 method for class 'factor'
overview(x, digits = getOption("cheapr.digits", 2), ...)

## S3 method for class 'Date'
overview(x, digits = getOption("cheapr.digits", 2), ...)

## S3 method for class 'POSIXt'
overview(x, digits = getOption("cheapr.digits", 2), ...)

## S3 method for class 'ts'
overview(x, digits = getOption("cheapr.digits", 2), ...)

## S3 method for class 'zoo'
overview(x, digits = getOption("cheapr.digits", 2), ...)

## S3 method for class 'data.frame'
overview(x, digits = getOption("cheapr.digits", 2), hist = TRUE, ...)
```

Arguments

<code>x</code>	A vector or data frame.
<code>digits</code>	How many decimal places should the summary statistics be printed as? Default is 2.
<code>...</code>	Further arguments passed onto methods. Currently unused.
<code>hist</code>	Should in-line histograms be returned? Default is FALSE.

Details

No rounding of statistics is done except in printing which can be controlled either through the `digits` argument in `overview()`, or by setting the option `options(cheapr.digits)`.

To access the underlying data, for example the numeric summary, just use `$numeric`, e.g. `overview(rnorm(30))$numeric`.

Value

An object of class "overview". Under the hood this is just a list of data frames. Key summary statistics are reported in each data frame.

Examples

```
library(cheapr)
```

```

overview(iris)

# With histograms
overview(airquality, hist = TRUE)

# Round to 0 decimal places
overview(airquality, digits = 0)

# We can set an option for all overviews
options(cheapr.digits = 1)
overview(rnorm(100))
options(cheapr.digits = 2) # The default

```

reconstruct

Reconstruct an object from a template

Description

Reconstruct an object from a template

Usage

```

reconstruct(x, template, ...)

## S3 method for class 'data.frame'
reconstruct(x, template, shallow_copy = TRUE, ...)

## S3 method for class 'data.table'
reconstruct(x, template, shallow_copy = TRUE, ...)

```

Arguments

x	An object in which carefully selected attributes will be copied into from template.
template	A template object used to copy attributes into x.
...	Further arguments passed onto methods.
shallow_copy	Should x be shallow copied before reconstructing? Default is TRUE.

Details

In R attributes are difficult to work with. One big reason for this is that attributes may or may not be independent of the data. Date vectors for example have attributes completely independent of the data and hence if the attributes are removed at any point, they can easily be re-added without any calculations. Factors have almost data-independent attributes with an exception being when factors are combined. In some cases it is not possible to reconstruct attributes from the data alone.

You can add your own reconstruct method for an object not covered by the methods here.

Value

An object similar to `template`.

<code>recycle</code>	<i>Recycle objects to a common size</i>
----------------------	---

Description

A convenience function to recycle R objects to either a common or specified size.

Usage

```
recycle(..., length = NULL, .args = NULL)
```

Arguments

<code>...</code>	Objects to recycle.
<code>length</code>	Optional length to recycle objects to.
<code>.args</code>	An alternative to <code>...</code> for easier programming with lists.

Details

Data frames are recycled by recycling their rows.
`recycle()` is optimised to only recycle objects that need recycling.
 NULL objects are ignored and not recycled or returned.

Value

A list of recycled R objects.

Examples

```
library(cheapr)

# Recycles both to size 10
recycle(Sys.Date(), 1:10)

# Any vectors of zero-length are all recycled to zero-length
recycle(integer(), 1:10)

# Unless length is supplied
recycle(integer(), 1:10, length = 10)

# Data frame rows are recycled
recycle(sset(iris, 1:3), length = 9)

# To recycle objects in a list, use `.args`
my_list <- list(from = 1L, to = 10L, by = seq(0.1, 1, 0.1))
recycle(.args = my_list)
```

rep	<i>cheapr style repeat functions</i>
-----	--------------------------------------

Description

cheapr style repeat functions

Usage

```
cheapr_rep(x, times)
```

```
cheapr_rep_len(x, length)
```

```
cheapr_rep_each(x, each)
```

Arguments

x	A vector or data frame.
times	[integer(n)] A vector of times to repeat elements of x. Can be length 1 or the same length as vector_length(x).
length	[integer(1)] - Length of the recycled result.
each	[integer(1)] - How many times to repeat out each element of x.

Value

Repeated out object.

sequence_	<i>Utilities for creating many sequences</i>
-----------	--

Description

sequence_ is an extension to [sequence](#) which accepts decimal number increments.
seq_id can be paired with sequence_ to group individual sequences.
seq_ is a vectorised version of [seq](#).
window_sequence creates a vector of window sizes for rolling calculations.
lag_sequence creates a vector of lags for rolling calculations.
lead_sequence creates a vector of leads for rolling calculations.

Usage

```

sequence_(size, from = 1L, by = 1L, add_id = FALSE)

seq_id(size)

seq_(from = 1L, to = 1L, by = 1L, add_id = FALSE)

seq_size(from, to, by = 1L)

window_sequence(size, k, partial = TRUE, ascending = TRUE, add_id = FALSE)

lag_sequence(size, k, partial = TRUE, add_id = FALSE)

lead_sequence(size, k, partial = TRUE, add_id = FALSE)

```

Arguments

size	Vector of sequence lengths.
from	Start of sequence(s).
by	Unit increment of sequence(s).
add_id	Should the ID numbers of the sequences be added as names? Default is FALSE.
to	End of sequence(s).
k	Window/lag size.
partial	Should partial windows/lags be returned? Default is TRUE.
ascending	Should window sequence be ascending? Default is TRUE.

Details

sequence_() works in the same way as sequence() but can accept non-integer by values. It also recycles from and to, in the same way as sequence().

If any of the sequences contain values > .Machine\$integer.max, then the result will always be a double vector.

from can be also be a date, date-time, or any object that supports addition and multiplication.

seq_() is a vectorised version of seq() that strictly accepts only the arguments from, to and by.

Value

A vector of length sum(size) except for seq_ which returns a vector of size sum((to - from) / (by + 1))

Examples

```

library(cheapr)
sequence(1:3)
sequence_(1:3)

```

```

sequence(1:3, by = 0.1)
sequence_(1:3, by = 0.1)

# Add IDs to the sequences
sequence_(1:3, by = 0.1, add_id = TRUE)
# Turn this quickly into a data frame
seqs <- sequence_(1:3, by = 0.1, add_id = TRUE)
new_df(name = names(seqs), seq = seqs)

sequence(c(3, 2), by = c(-0.1, 0.1))
sequence_(c(3, 2), by = c(-0.1, 0.1))

# Vectorised version of seq()
seq_(1, 10, by = c(1, 0.5))
# Same as below
c(seq(1, 10, 1), seq(1, 10, 0.5))

# Programmers may use seq_size() to determine final sequence lengths

sizes <- seq_size(1, 10, by = c(1, 0.5))
print(paste(c("sequence sizes: (", sizes, ") total size:", sum(sizes)),
            collapse = " "))

# We can group sequences using seq_id

from <- Sys.Date()
to <- from + 10
by <- c(1, 2, 3)
x <- seq_(from, to, by, add_id = TRUE)
class(x) <- "Date"
x

# Utilities for rolling calculations

window_sequence(c(3, 5), 3)
window_sequence(c(3, 5), 3, partial = FALSE)
window_sequence(c(3, 5), 3, partial = TRUE, ascending = FALSE)
# One can for example use these in data.table::frollsum

```

setdiff_

Extra utilities

Description

Extra utilities

Usage

```

setdiff_(x, y, dups = TRUE)

intersect_(x, y, dups = TRUE)

cut_numeric(
  x,
  breaks,
  labels = NULL,
  include.lowest = FALSE,
  right = TRUE,
  dig.lab = 3L,
  ordered_result = FALSE,
  ...
)

x %in_% table

x %!in_% table

enframe_(x, name = "name", value = "value")

deframe_(x)

sample_(x, size = vector_length(x), replace = FALSE, prob = NULL)

val_insert(x, value, n = NULL, prop = NULL)

na_insert(x, n = NULL, prop = NULL)

vector_length(x)

cheapr_var(x, na.rm = TRUE)

cheapr_rev(x)

with_local_seed(expr, .seed = NULL, ...)

```

Arguments

x	A vector or data frame.
y	A vector or data frame.
dups	Should duplicates be kept? Default is TRUE.
breaks	See ?cut.
labels	See ?cut.
include.lowest	See ?cut.
right	See ?cut.

dig.lab	See ?cut.
ordered_result	See ?cut.
...	Further arguments passed onto cut or set.seed.
table	See ?collapse::fmatch
name	The column name to assign the names of a vector.
value	The column name to assign the values of a vector.
size	See ?sample.
replace	See ?sample.
prob	See ?sample.
n	Number of scalar values (or NA) to insert randomly into your vector.
prop	Proportion of scalar values (or NA) values to insert randomly into your vector.
na.rm	Should NA values be ignored in cheapr_var() Default is TRUE.
expr	Expression that will be evaluated with a local seed that is independent and has absolutely no effect on the global RNG state.
.seed	A local seed to set which is only used inside with_local_seed(). After the execution of the expression the original seed is reset.

Value

enframe()_ converts a vector to a data frame.
 deframe()_ converts a 1-2 column data frame to a vector.
 intersect_() returns a vector of common values between x and y.
 setdiff_() returns a vector of values in x but not y.
 cut_numeric() places values of a numeric vector into buckets, defined through the breaks argument and returns a factor unless labels = FALSE, in which case an integer vector of break indices is returned.
 %in_% and %!in_% both return a logical vector signifying if the values of x exist or don't exist in table respectively.
 sample_() is an alternative to sample() that natively samples data frame rows through sset(). It also does not have a special case when length(x) is 1.
 val_insert inserts scalar values randomly into your vector. Useful for replacing lots of data with a single value.
 na_insert inserts NA values randomly into your vector. Useful for generating missing data.
 vector_length behaves mostly like NROW() except for matrices in which it matches length().
 cheapr_var returns the variance of a numeric vector. No coercion happens for integer vectors and so is very cheap.
 cheapr_rev is a much cheaper version of rev().
 with_local_seed offers no speed improvements but is extremely handy in executing random number based expressions like rnorm() without affecting the global RNG state. It allows you to run these expressions in a sort of independent 'container' and with an optional seed for that 'container' for reproducibility. The rationale for including this in 'cheapr' is that it can reduce the need to set many seed values, especially for multiple output comparisons of RNG expressions. Another way of thinking about it is that with_local_seed() is a helper that allows you to write reproducible code without side-effects, which traditionally cannot be avoided when calling set.seed() directly.

Examples

```

library(cheapr)

# Using `with_local_seed()`

# The below 2 statements are equivalent

# Statement 1
set.seed(123456789)
res <- rnorm(10)

# Statement 2
res2 <- with_local_seed(rnorm(10), .seed = 123456789)

# They are the same
identical(res, res2)

# As an example we can see that the RNG is unaffected by generating
# random uniform deviates in batches between calls to `with_local_seed()`
# and comparing to the first result

set.seed(123456789)
batch1 <- rnorm(2)

with_local_seed(runif(10))
batch2 <- rnorm(2)
with_local_seed(runif(10))
batch3 <- rnorm(1)
with_local_seed(runif(10))
batch4 <- rnorm(5)

# Combining the batches produces the same result
# therefore `with_local_seed` did not interrupt the rng sequence
identical(c(batch1, batch2, batch3, batch4), res)

# It can be useful in multiple comparisons
out1 <- with_local_seed(rnorm(5))
out2 <- with_local_seed(rnorm(5))
out3 <- with_local_seed(rnorm(5))

identical(out1, out2)
identical(out1, out3)

```

set_abs

Math operations by reference - Experimental

Description

These functions transform your variable by reference, with no copies being made. It is advisable to only use these if you know what you are doing.

Usage

```
set_abs(x)

set_floor(x)

set_ceiling(x)

set_trunc(x)

set_exp(x)

set_sqrt(x)

set_change_sign(x)

set_round(x, digits = 0)

set_log(x, base = exp(1))

set_pow(x, y)

set_add(x, y)

set_subtract(x, y)

set_multiply(x, y)

set_divide(x, y)
```

Arguments

x	A numeric vector.
digits	Number of digits to round to.
base	Logarithm base.
y	A numeric vector.

Details

These functions are particularly useful for situations where you have made a copy and then wish to perform further operations without creating more copies.

NA and NaN values are ignored though in some instances NaN values may be replaced with NA. These functions will **not work** on **any** classed objects, meaning they only work on standard integer and numeric vectors and matrices.

When a copy has to be made:

A copy is only made in certain instances, e.g. when passing an integer vector to `set_log()`. A warning will always be thrown in this instance alerting the user to assign the output to an object

because `x` has not been updated by reference.
 To ensure consistent and expected outputs, always assign the output to the same object,
 e.g. `x <- set_log(x)` (**do this**)
`set_log(x)` (**don't do this**)
`x2 <- set_log(x)` (Don't do this either)

No copy is made here unless `x` is an integer vector.

Value

The exact same object with no copy made, just transformed.

Examples

```
library(cheapr)
library(bench)

x <- rnorm(2e05)
options(cheapr.cores = 2)
mark(
  base = exp(log(abs(x))),
  cheapr = set_exp(set_log(set_abs(x)))
)
options(cheapr.cores = 1)
```

sset

Cheaper subset

Description

Cheaper alternative to `[]` that consistently subsets data frame rows, always returning a data frame.
 There are explicit methods for enhanced data frames like tibbles, data.tables and sf.

Usage

```
sset(x, ...)
```

S3 method for class 'data.frame'

```
sset(x, i = NULL, j = NULL, ...)
```

S3 method for class 'POSIXlt'

```
sset(x, i = NULL, j = NULL, ...)
```

S3 method for class 'sf'

```
sset(x, i = NULL, j = NULL, ...)
```


Arguments

<code>x</code>	Vector or data frame.
<code>...</code>	Further parameters passed to <code>[]</code> .
<code>i</code>	A logical or vector of indices.
<code>j</code>	Column indices, names or logical vector.

Details

`sset` is an S3 generic. You can either write methods for `sset` or `[]`.
`sset` will fall back on using `[]` when no suitable method is found.

To get into more detail, using `sset()` on a data frame, a new list is always allocated through `new_list()`.

Difference to base R:

When `i` is a logical vector, it is passed directly to `which_()`.

This means that NA values are ignored and this also means that `i` is not recycled, so it is good practice to make sure the logical vector matches the length of `x`. To return NA values, use `sset(x, NA_integer_)`.

ALTREP range subsetting:

When `i` is an ALTREP compact sequence which can be commonly created using e.g. `1:10` or using `seq_len`, `seq_along` and `seq.int`, `sset` internally uses a range-based subsetting method which is faster and doesn't allocate `i` into memory.

Value

A new vector, data frame, list, matrix or other R object.

Examples

```
library(cheapr)
library(bench)

# Selecting columns
sset(airquality, j = "Temp")
sset(airquality, j = 1:2)

# Selecting rows
sset(iris, 1:5)

# Rows and columns
sset(iris, 1:5, 1:5)
sset(iris, iris$Sepal.Length > 7, c("Species", "Sepal.Length"))

# Comparison against base
x <- rnorm(10^4)

mark(x[1:10^3], sset(x, 1:10^3))
```

```
mark(x[x > 0], sset(x, x > 0))

df <- data.frame(x = x)

mark(df[df$x > 0, , drop = FALSE],
      sset(df, df$x > 0),
      check = FALSE) # Row names are different
```

sset_df

Fast functions for data frame subsetting

Description

These functions are for developers that need minimal overhead when filtering on rows and/or cols.

Usage

```
sset_df(x, i = NULL, j = NULL, ...)

sset_row(x, i = NULL)

sset_col(x, j = NULL)
```

Arguments

x	A data.frame.
i	Rows - If NULL all rows are returned.
j	Cols - If NULL all cols are returned.
...	Unused.

Details

If you are unsure which functions to use then it is recommended to use `sset()`. These low-overhead helpers do not work well with `data.tables` but should work well with basic data frames and basic tibbles. The only real difference between `sset_df` and `sset_row/sset_col` is that `sset_df` attempts to return a similar type of data frame as the input, whereas `sset_row` and `sset_col` always return a plain data frame.

Value

A data frame subsetting on rows `i` and cols `j`.

val_count	<i>Efficient functions for counting, finding, replacing and removing scalars</i>
-----------	--

Description

These are primarily intended as very fast scalar-based functions for developers. They are particularly useful for working with NA values in a fast and efficient manner.

Usage

```
val_count(x, value, recursive = TRUE)

val_find(x, value, invert = FALSE)

which_val(x, value, invert = FALSE)

val_replace(x, value, replace, recursive = TRUE)

na_replace(x, replace, recursive = TRUE)

val_rm(x, value)

na_count(x, recursive = TRUE)

na_find(x, invert = FALSE)

na_rm(x)
```

Arguments

x	A vector, list, data frame or matrix.
value	A scalar value to count, find, replace or remove.
recursive	Should values in a list be counted or replaced recursively? Default is TRUE and very useful for data frames.
invert	Should which_val find locations of everything except specified value? Default is FALSE.
replace	Replacement scalar value.

Details

The val_ functions allow you to very efficiently work with scalars, i.e length 1 vectors. Many common common operations like counting the occurrence of NA or zeros, e.g. `sum(x == 0)` or `sum(is.na(x))` can be replaced more efficiently with `val_count(x, 0)` and `na_count(x)` respectively.

At the moment these functions only work for integer, double and character vectors with the exception of the NA functions. They are intended mainly for developers who wish to write cheaper code and reduce expensive vector operations.

- `val_count()` - Counts occurrences of a value
- `val_find()` Finds locations (indices) of a value
- `val_replace()` - Replaces value with another value
- `val_rm()` - Removes occurrences of value from an object

There are NA equivalent convenience functions.

- `na_count() == val_count(x, NA)`
- `na_find() == val_find(x, NA)`
- `na_replace() == val_replace(x, NA)`
- `na_rm() == val_rm(x, NA)`

`val_count()` and `val_replace()` can work recursively. For example, when applied to a data frame, `na_replace` will replace NA values across the entire data frame with the specified replacement value.

In 'cheapr' function-naming conventions have not been consistent but going forward all scalar functions (including the NA convenience functions) will be prefixed with 'val_' and 'na_' respectively. Functions named with the older naming scheme like `which_na` may be removed at some point in the future.

Value

`val_count()` returns the number of times a scalar value appears in a vector or list.

`val_find()` returns the index locations of that scalar value.

`val_replace()` replaces a specified scalar value with a replacement scalar value. If no instances of said value are found then the input `x` is returned as is.

`na_replace()` is a convenience function equivalent to `val_replace(x, NA, ...)`.

`val_rm()` removes all instances of a specified scalar value. If no instances are found, the original input `x` is returned as is.

which_	<i>Memory-efficient alternative to which()</i>
--------	--

Description

Exactly the same as `which()` but more memory efficient.

Usage

```
which_(x, invert = FALSE)
```

Arguments

x	A logical vector.
invert	If TRUE, indices of values that are not TRUE are returned (including NA). If FALSE (the default), only TRUE indices are returned.

Details

This implementation is similar in speed to `which()` but usually more memory efficient.

Value

An unnamed integer vector.

Examples

```
library(cheapr)
library(bench)
x <- sample(c(TRUE, FALSE), 1e05, TRUE)
x[sample.int(1e05, round(1e05/3))] <- NA

mark(which_(TRUE), which(TRUE))
mark(which_(FALSE), which(FALSE))
mark(which_(logical()), which(logical()))
mark(which_(x), which(x), iterations = 20)
mark(base = which(is.na(match(x, TRUE))),
      collapse = collapse::whichv(x, TRUE, invert = TRUE),
      cheapr = which_(x, invert = TRUE),
      iterations = 20)
```

Index

`%!in_(setdiff_)`, 35

`%in_(setdiff_)`, 35

`address`, 3

`all_na(is_na)`, 19

`any_na(is_na)`, 19

`as_df(new_df)`, 28

`as_discrete`, 3, 7, 18

`as_factor(factor_)`, 12

`attrs`, 6

`attrs_add(attrs)`, 6

`attrs_rm(attrs)`, 6

`bin`, 4, 7, 18

`case`, 8, 10

`cheapr(cheapr-package)`, 2

`cheapr-package`, 2

`cheapr_c`, 9

`cheapr_if_else`, 8, 10

`cheapr_rep(rep)`, 33

`cheapr_rep_each(rep)`, 33

`cheapr_rep_len(rep)`, 33

`cheapr_rev(setdiff_)`, 35

`cheapr_table`, 10

`cheapr_var(setdiff_)`, 35

`col_all_na(is_na)`, 19

`col_any_na(is_na)`, 19

`col_c(new_df)`, 28

`col_na_counts(is_na)`, 19

`copy`, 11

`cut_numeric(setdiff_)`, 35

`deep_copy(copy)`, 11

`deframe_(setdiff_)`, 35

`enframe_(setdiff_)`, 35

`factor_`, 12

`fast_df(new_df)`, 28

`gcd`, 15

`gcd2(gcd)`, 15

`get_breaks`, 4, 7, 17

`int_sign`, 19

`intersect_(setdiff_)`, 35

`is_na`, 19

`lag2_(lag_)`, 22

`lag_`, 22

`lag_sequence(sequence_)`, 33

`lead_sequence(sequence_)`, 33

`lengths_(list_lengths)`, 26

`levels_add(factor_)`, 12

`levels_add_na(factor_)`, 12

`levels_count(factor_)`, 12

`levels_drop(factor_)`, 12

`levels_drop_na(factor_)`, 12

`levels_factor(factor_)`, 12

`levels_lump(factor_)`, 12

`levels_rename(factor_)`, 12

`levels_reorder(factor_)`, 12

`levels_rm(factor_)`, 12

`levels_unused(factor_)`, 12

`levels_used(factor_)`, 12

`list_as_df(new_df)`, 28

`list_assign(list_lengths)`, 26

`list_combine(list_lengths)`, 26

`list_drop_null(list_lengths)`, 26

`list_lengths`, 26

`logical`, 10, 45

`na_count(val_count)`, 43

`na_find(val_count)`, 43

`na_insert(setdiff_)`, 35

`na_replace(val_count)`, 43

`na_rm(val_count)`, 43

`name_repair(new_df)`, 28

`named_list`, 28

`new_df`, 28

`new_list(list_lengths)`, 26
`num_na(is_na)`, 19
`numeric`, 16

`overview`, 29

`reconstruct`, 31
`recycle`, 32
`rep`, 33
`row_all_na(is_na)`, 19
`row_any_na(is_na)`, 19
`row_c(new_df)`, 28
`row_na_counts(is_na)`, 19

`sample_(setdiff_)`, 35
`scm(gcd)`, 15
`scm2(gcd)`, 15
`semi_copy(copy)`, 11
`seq`, 33
`seq_(sequence_)`, 33
`seq_id(sequence_)`, 33
`seq_size(sequence_)`, 33
`sequence`, 33
`sequence_`, 33
`set_abs`, 38
`set_add(set_abs)`, 38
`set_ceiling(set_abs)`, 38
`set_change_sign(set_abs)`, 38
`set_divide(set_abs)`, 38
`set_exp(set_abs)`, 38
`set_floor(set_abs)`, 38
`set_log(set_abs)`, 38
`set_multiply(set_abs)`, 38
`set_pow(set_abs)`, 38
`set_round(set_abs)`, 38
`set_sqrt(set_abs)`, 38
`set_subtract(set_abs)`, 38
`set_trunc(set_abs)`, 38
`setdiff_`, 35
`shallow_copy`, 6
`shallow_copy(copy)`, 11
`sset`, 40
`sset_col(sset_df)`, 42
`sset_df`, 42
`sset_row(sset_df)`, 42

`unique_name_repair(new_df)`, 28
`unlisted_length(list_lengths)`, 26

`val_count`, 43

`val_find(val_count)`, 43
`val_insert(setdiff_)`, 35
`val_match`, 10
`val_match(case)`, 8
`val_replace(val_count)`, 43
`val_rm(val_count)`, 43
`vector_length(setdiff_)`, 35

`which_`, 44
`which_na(is_na)`, 19
`which_not_na(is_na)`, 19
`which_val(val_count)`, 43
`window_sequence(sequence_)`, 33
`with_local_seed(setdiff_)`, 35