

R Toolkit Digest



*...dose of hidden gems, packages,
and practical R tricks...*

Do More with Less
Code in the
Tidyverse



Tidyverse Tools That Do Less

Modern data work is not about writing more code.

It is about:

- Writing **less repetitive code**
- Making pipelines easier to read
- Scaling transformations across many variables
- Reducing human error

Today we focus on three tools:

- `across()`
- `map()`
- `reduce()`

The Problem with Repetition

Consider this common pattern in a dataset:

```
1 df %>%  
2   mutate(  
3     date_of_birth = as.Date(date_of_birth),  
4     date_enrolled = as.Date(date_enrolled),  
5     date_completed = as.Date(date_completed)  
6   )
```

Problems:

- Repetitive code
- Hard to scale
- Easy to make mistakes

1: `across()` as the solution

What it does:

Applies a function to multiple columns at once.

```
1 df %>%  
2   mutate(across(c(var1, var2, var3), mean, na.rm = TRUE))  
3
```

Even better (pattern-based selection):

```
1 df %>%  
2   mutate(  
3     across(  
4       starts_with("date"), # other helpers like `ends_with`, `contains` can be used.  
5       as.Date  
6     )  
7   )
```

across() with Multiple Functions

You can apply multiple transformations at once:

Code:

```
1 library(dplyr)
2
3 # Using the built-in mtcars dataset to illustrate
4
5 mtcars %>%
6   summarise(
7     across(
8       c(mpg, hp, wt), # Selecting columns of interest
9       list(
10        mean = ~mean(.x, na.rm = TRUE),
11        sd   = ~sd(.x, na.rm = TRUE),
12        min  = ~min(.x, na.rm = TRUE),
13        max  = ~max(.x, na.rm = TRUE)
14      )
15    )
16  )
```

Output:

mpg_mean	mpg_sd	mpg_min	mpg_max	hp_mean	hp_sd	hp_min	hp_max	wt_mean	wt_sd	wt_min	wt_max
20.09062	6.026948	10.4	33.9	146.6875	68.56287	52	335	3.21725	0.9784574	1.513	5.424

2: Repeat Tasks Without Loops :

map()

What it does:

Applies the same operation to many objects.

Instead of writing:

```
1 summary(df1)
2 summary(df2)
3 summary(df3)
4 summary(df4)
```

Use:

```
1 library(purrr)
2
3 list(df1, df2, df3, df4) %>%
4   map(summary)
```

Benefits:

map() with Real Data Wrangling

Suppose you receive monthly beneficiary datasets that need the same cleaning steps.

Instead of:

```
1 jan_clean <- jan %>%
2   filter(!is.na(unique_id)) %>%
3   mutate(sex = str_to_title(sex))
4
5 feb_clean <- feb %>%
6   filter(!is.na(unique_id)) %>%
7   mutate(sex = str_to_title(sex))
8
9 mar_clean <- mar %>%
10  filter(!is.na(unique_id)) %>%
11  mutate(sex = str_to_title(sex))
12
13 apr_clean <- apr %>%
14  filter(!is.na(unique_id)) %>%
15  mutate(sex = str_to_title(sex))
```

Use:

```
1 monthly_data <- list(
2   jan,
3   feb,
4   mar,
5   apr
6 )
7
8 monthly_data %>%
9   map(
10    ~ .x %>%
11      filter(!is.na(unique_id)) %>%
12      mutate(sex = str_to_title(sex))
13  )
```

Benefits:

- Apply the same cleaning logic consistently
- Avoid duplicating code
- Scale easily to many datasets

map() for Data Quality Checks

```
1 datasets <- list(  
2   households,  
3   children,  
4   caregivers  
5 )  
6  
7 datasets %>%  
8   map(~ summarise(  
9     .x,  
10    missing_ids = sum(is.na(unique_id)),  
11    records = n()  
12  ))
```

This produces the same quality checks for every dataset with a single pipeline.

Understanding `.x` in `map()`

When using `map()`, `.x` represents the current item being processed.

Think of it as a placeholder for each element in a list.

Example:

```
1 numbers <- list(1, 2, 3)
2
3 numbers %>%
4   map(~ .x * 2)
```

Output:

```
1 [[1]]
2 [1] 2
3
4 [[2]]
5 [1] 4
6
7 [[3]]
8 [1] 6
```

Here:

- First iteration: `.x = 1`
- Second iteration: `.x = 2`
- Third iteration: `.x = 3`

In Our Data Wrangling Example

Code:

```
1 monthly_data %>%
2   map(
3     ~ .x %>%
4       filter(!is.na(unique_id)) %>%
5       mutate(sex = str_to_title(sex))
6   )
```

Output::

Iteration	.x Represents
1	jan
2	feb
3	mar
4	apr

So map() effectively runs:

```
1 jan %>%
2   filter(!is.na(unique_id)) %>%
3   mutate(sex = str_to_title(sex))
4
5 feb %>%
6   filter(!is.na(unique_id)) %>%
7   mutate(sex = str_to_title(sex))
8
9 mar %>%
10  filter(!is.na(unique_id)) %>%
11  mutate(sex = str_to_title(sex))
12
13 apr %>%
14  filter(!is.na(unique_id)) %>%
15  mutate(sex = str_to_title(sex))
```

automatically.

3: `reduce()`: When You Have Many Tables

Problem: multiple left joins

```
1 df1 %>%  
2   left_join(df2) %>%  
3   left_join(df3) %>%  
4   left_join(df4)
```

Solution: `purrr::reduce()`

```
1 library(purrr)  
2  
3 list(df1, df2, df3, df4) %>%  
4   reduce(left_join, by = "id")
```

Benefits:

- Scales to many tables
- Cleaner pipelines
- Easier automation

Real-World Use Case

Imagine program data:

- Household table
- Child table
- Service table
- Location table

Instead of chaining joins manually,
reduce works in this way:

```
1 reduce(list(hh, child, service, location),  
2         left_join, by = "unique_id")
```

When to Use Each Tool

Tool	Purpose	Together they help you write code that is:
<code>across()</code>	Apply functions to many columns.	<ul style="list-style-type: none">• Shorter• More scalable• Easier to maintain• Less error-prone
<code>map()</code>	Apply functions to many objects, lists, e.g. reading multiple files.	
<code>reduce()</code>	Combine many objects into one, avoiding repeated joins.	



Further Reading

- Tidyverse documentation: Column-wise operations with `across()`
- R for Data Science: Data Transformation
- purrr documentation: `map()` Family
- R for Data Science: Iteration
- purrr documentation: `reduce()`
- Advanced R: Functionals