

Testing with testthat

Lecture 06

Dr. Colin Rundel

Package checking

R CMD check - What it does

R CMD check is CRAN's comprehensive quality control system that runs dozens of checks:

- **Package structure** - Correct directories, required files (DESCRIPTION, NAMESPACE)
- **Code syntax** - All R code parses without errors
- **Documentation** - All functions documented, all examples run
- **Dependencies** - All used packages exist and in DESCRIPTION
- **Tests** - All test files execute without errors
- **CRAN policy compliance** - Follows all submission guidelines (written and unwritten)

devtools::check() vs R CMD check

`devtools::check()` is a convenient wrapper around `R CMD check`:

```
1 devtools::check()
```

```
1 R CMD build .  
2 R CMD check packagename_1.0.0.tar.gz
```

Benefits of `devtools::check()`:

- Automatically handles building and checking
- Better integrated with RStudio workflow
- Cleaner output formatting
- Automatically installs package first

Interpreting check output

R CMD check produces three levels of issues:

- **ERROR** 🛑 - Must be fixed before CRAN submission
- **WARNING** ⚠️ - Should be fixed (CRAN may reject)
- **NOTE** 📝 - Optional improvements (CRAN usually accepts)

Anything flagged must be addressed in the CRAN submission process.

While the check is running issues are shown inline, *and* summarized at the end.

GitHub Actions for continuous checking

Set up automated checking with:

```
1 usethis::use_github_action("check-standard")
```

This creates `.github/workflows/R-CMD-check.yaml` that runs checks on:

- Latest R on macOS, Windows, Linux (Ubuntu)
- Previous R and R-devel on Linux (Ubuntu)

Package testing

Basic test structure

Package tests live in `tests/`,

- Any R scripts found in the folder will be run when checking
- Generally tests “fail” if an error is thrown, warnings are also tracked
- Testing with base R alone is possible but not recommended (See [Writing R Extensions](#))
- There is functionality for comparing test outputs to expected results, but it has limited functionality
- Note that R CMD check also runs all documentation examples (unless explicitly tagged with don't run)

testthat



testthat fundamentals

testthat is the most widely used testing framework for R packages, it also has excellent integration with RStudio & usethis.

A project can be initialized to use testthat via:

```
1 usethis::use_testthat()
```

This creates the following files and directories:

- `tests/testthat.R` - Entry point for R CMD check
- `tests/testthat/` - Directory for test files
- Adds testthat to `DESCRIPTION`'s `Suggests` field

testthat project structure

```
mypackage/  
├── R/  
│   └── utils.R  
├── tests/  
│   ├── testthat.R  
│   └── testthat/  
│       └── test-utils.R  
└── DESCRIPTION
```

Test file naming:

- Must start with `test-` or `test_`
- Typically tests `test-utils.R` map to scripts `R/utils.R` (use `usethis::use_test()` to create for current open file)
- Can also group related functions: `test-data-processing.R`
- `helper*.R`, `teardown*.R`, and `setup.R` all have special behavior - see [Special files](#)
- All other files are ignored

testthat script structure

Tests are hierarchically organized:

- **File** - Collection of related tests
- **Test** - Group of related expectations (`test_that()`)
- **Expectation** - Single assertion (`expect_equal()`, `expect_error()`, etc.)

```
1 test_that("`+` works correctly", {  
2   expect_equal(`+`(2, 3), 5)  
3   expect_equal(`+`(0, 0), 0)  
4   expect_type(`+`(1, 1), "double")  
5   expect_type(`+`(1L, 1L), "integer")  
6 })
```

Test passed with 4 successes 🏆.

Running tests

There are multiple ways to execute your package's tests:

During development:

- `devtools::test()` - Run all tests
- `devtools::test_file("tests/testthat/test-utils.R")` - Run one file
- **Ctrl/Cmd+Shift+T** (RStudio) - Run all tests
- **Ctrl/Cmd+T** (RStudio) - Run tests for current file

From command line:

- `R CMD check` - Runs tests as part of package check
- `Rscript -e "devtools::test()"` - In scripts / GitHub Actions

Core expectation functions

testthat provides many expectation functions for different scenarios.

Equality and identity:

```
1 expect_equal(actual, expected)
2 expect_identical(actual, expected)
3 expect_true(x)
4 expect_false(x)
```

Types and classes:

```
1 expect_type(x, "double")
2 expect_s3_class(df, "data.frame")
```

Conditions:

```
1 expect_error(code, regexp = "...")
2 expect_warning(code, regexp = "...")
3 expect_message(code, regexp = "...")
```

expect_equal() vs expect_identical()

Equal compares values based on a tolerance (not the same as `==`), identity check for exact equivalence.

```
1 test_that("equality", {
2   expect_equal(0.2 + 0.2, 0.4)
3   expect_equal(0.1 + 0.2, 0.3)
4   expect_equal(1L, 1.0)
5 })
```

Test passed with 3 successes 🎉.

```
1 test_that("identity", {
2   expect_identical(0.2 + 0.2, 0.4)
3   expect_identical(0.1 + 0.2, 0.3)
4   expect_identical(1L, 1.0)
5 })
```

```
— Failure: identity —————
Expected `0.1 + 0.2` to be identical to 0.3.
Differences:
Objects equal but not identical
— Failure: identity —————
Expected 1L to be identical to 1.
Differences:
Objects equal but not identical
Error:
! Test failed with 2 failures and 1
  success.
```

Testing function outputs

```
1 calculate_mean_ci = function(x, conf_level = 0.95) {
2   if (length(x) == 0)
3     stop("Cannot calculate CI for empty vector")
4   if (any(is.na(x)))
5     stop("Missing values not allowed")
6
7   n = length(x)
8   mean_x = mean(x)
9   se = sd(x) / sqrt(n)
10  t_val = qt((1 + conf_level) / 2, df = n - 1)
11
12  c(lower = mean_x - t_val * se, upper = mean_x + t_val * se)
13 }
```

Example tests

```
1 test_that("calculate_mean_ci works correctly", {
2   # Test normal case
3   result = calculate_mean_ci(c(1, 2, 3, 4, 5))
4   expect_type(result, "double")
5   expect_length(result, 2)
6   expect_named(result, c("lower", "upper"))
7   expect_true(result["lower"] < result["upper"])
8
9   # Test with known values
10  expect_equal(
11    calculate_mean_ci(c(0, 0, 0)),
12    c(lower = 0, upper = 0)
13  )
14
15  # Test confidence level parameter
16  ci_95 = calculate_mean_ci(c(1, 2, 3), conf_level = 0.95)
17  ci_99 = calculate_mean_ci(c(1, 2, 3), conf_level = 0.99)
18  expect_true(ci_99["upper"] - ci_99["lower"] > ci_95["upper"] - ci_95["lower"])
19 }
```

Test passed with 6 successes 😊.

Testing error conditions

```
1 test_that("calculate_mean_ci error cases",
2   # Empty vector should error
3   expect_error(
4     calculate_mean_ci(numeric(0)),
5     "Cannot calculate CI for empty vector"
6   )
7
8   # Missing values should error
9   expect_error(
10    calculate_mean_ci(c(1, 2, NA)),
11    "Missing values not allowed"
12  )
13
14  # Invalid confidence level should error
15  expect_error(
16    calculate_mean_ci(1:5, conf_level = 1.5),
17    "conf_level must be between 0 and 1"
18  )
19
20  # Single value (edge case to think about)
21  expect_error(
22    calculate_mean_ci(5)
23  )
```

```
— Warning: calculate_mean_ci error cases ———
NaNs produced
Backtrace:
```

```
1. ── testthat::expect_error(...)
2.   └─ testthat::quasi_capture(...)
3.     └─ testthat (local) .capture(...)
4.       └─ base::withCallingHandlers(...)
5.         └─ rlang::eval_bare(quo_get_expr(.quo), q
6.           └─ global calculate_mean_ci(1:5, conf_level =
7.             └─ stats::qt((1 + conf_level)/2, df = n - 1
```

```
— Failure: calculate_mean_ci error cases ———
Expected `calculate_mean_ci(1:5, conf_level = 1.
```

```
— Warning: calculate_mean_ci error cases ———
NaNs produced
Backtrace:
```

```
1. ── testthat::expect_error(calculate_mean_ci(5
2.   └─ testthat::quasi_capture(...)
3.     └─ testthat (local) .capture(...)
4.       └─ base::withCallingHandlers(...)
5.         └─ rlang::eval_bare(quo_get_expr(.quo), q
6.           └─ global calculate mean ci(5)
```

Error:

```
! Test failed with 2 failures and 2
  successes.
```

Testing for errors

Testing for errors is important, but `expect_error()` can be dangerous if you don't check the output. All that the expectation tells you is that *some* error was thrown, not that it was the *right* error.

```
1 calculate_discount = function(price, discount_percent) {
2   if (price < 0) stop("Price cannot be negative")
3   if (discount_percent > 100) stop("Discount cannot exceed 100%")
4
5   price * (1 - discount_pct / 100) # Bug: wrong variable name
6 }
7
8 test_that("demonstrates why checking error messages matters", {
9   # x passes but for the wrong reason!
10  expect_error(calculate_discount(100, -50))
11  # ✓ This correctly tests the price validation
12  expect_error(calculate_discount(-50, 10), "Price cannot be negative")
13 })
```

Test passed with 2 successes 🎉.

```
1 calculate_discount(100, -50)
```

```
Error in `calculate_discount()`:  
! object 'discount_pct' not found
```

```
1 calculate_discount(-50, 10)
```

```
Error in `calculate_discount()`:  
! Price cannot be negative
```

In this case the issue would likely be caught by other tests,

```
1 test_that("Calculation test", {  
2   expect_equal(calculate_discount(100, 20), 80)  
3 })
```

— Error: Calculation test _____

```
Error in `calculate_discount(100, 20)`: object 'discount_pct' not found
```

Backtrace:

```
1. ──  
2. │ ── testthat::expect_equal(calculate_discount(100, 20), 80)  
3. │ │ ── testthat::quasi_label(enquo(object), label)  
4. │ │ │ ── rlang::eval_bare(expr, quo_get_env(quo))  
5. ── global calculate_discount(100, 20)
```

Error:

```
! Test failed with 1 failure and 0  
  successes.
```

Skipping tests

Skip tests when certain conditions aren't met:

```
1 test_that("database connection works", {
2   skip_if_not_installed("RPostgreSQL")
3   skip_if(Sys.getenv("TEST_DB_URL") == "", "Database URL not set")
4   skip_on_cran()
5   skip_on_ci()
6
7   # Database tests ...
8 })
9
10 test_that("internet-dependent test", {
11   skip_if_offline()
12
13   # Test that requires internet connection
14   result = download_data("https://example.com/api")
15   expect_type(result, "list")
16 })
```

Snapshot tests

Snapshot tests capture the output of your functions and compare against previously saved results:

- **First run:** Snapshot is created and saved
- **Subsequent runs:** Current output compared against saved snapshot
- **When output changes:** Test fails, you review and accept/reject the change

Snapshot tests are best for:

- Error messages and warnings
- Complex data structure outputs
- Printed output from functions
- Any output where exact specification is difficult

expect_snapshot() for output

Test printed output and messages:

```
1 print_summary = function(data) {
2   cat("Data summary:\n")
3   cat("Rows:", nrow(data), "\n")
4   cat("Columns:", ncol(data), "\n")
5   cat("Column names:", paste(names(data), collapse = ", "), "\n")
6 }
7
8 test_that("print_summary produces consistent output", {
9   df = data.frame(x = 1:3, y = letters[1:3])
10
11   expect_snapshot({
12     print_summary(df)
13   })
14 })
```

Snapshot output

Creates `tests/testthat/_snaps/print-summary.md`:

```
# print_summary produces consistent output
```

Code

```
print_summary(df)
```

Output

```
Data summary:
```

```
Rows: 3
```

```
Columns: 2
```

```
Column names: x, y
```

Managing snapshots

Accepting changes:

```
1 snapshot_accept()  
2 snapshot_accept("test-myfunction.R")
```

Reviewing changes

```
1 snapshot_review()  
2 snapshot_review("test-myfunction.R")
```

Some best practices:

- Review snapshot changes carefully in code review
- Don't commit snapshot updates without understanding why they changed
- Use descriptive test names for easier snapshot identification

Why testing matters

Testing is a fundamental part of creating reliable, maintainable R packages (and code in general):

- Catch bugs early - Find problems before they reach users
- Document behavior - Tests serve as executable specifications
- Prevent regressions - Ensure new changes don't break existing functionality
- Enable refactoring - Change implementation with confidence

Testing as documentation

Well-written tests serve multiple purposes:

```
1 test_that("mean() behaves as expected", {
2   # Basic calculations
3   expect_equal(mean(c(1, 2, 3)), 2)
4   expect_equal(mean(c(3, 2, 1)), 2)
5
6   # Missing values
7   expect_true(is.na(mean(c(1, 2, NA))))
8   expect_equal(mean(c(1, 2, NA), na.rm = TRUE), 1.5)
9   result = mean(numeric(0))
10  expect_true(is.na(result))
11  expect_true(is.nan(result))
12 })
```

Tests make your intentions clear to future maintainers / contributors (including yourself!)

Code coverage

Code coverage measures the % of your code that is executed during testing.

- Useful as a rough indicator of how well tested your code is
- The `covr` package provides coverage tooling for R packages
- Most CI services (e.g. GitHub Actions) can track coverage over time




Coverage has important limitations:

- Measures execution, not correctness
- Does not measure the code your code uses
- 100% coverage does not mean bug-free code
- Incentivizes writing tests that touch lines rather than tests that verify behavior
- Edge cases and input validation can be missed even at high coverage

Test-Driven Development

The TDD cycle: Red-Green-Refactor

Test-Driven Development follows a simple cycle:

1.  **Red:** Write a failing test for the functionality you want to implement
2.  **Green:** Write the minimal code to make the test pass
3.  **Refactor:** Clean up the code while keeping tests green
4. **Repeat:** Move on to the next piece of functionality

This approach ensures:

- You only write code that's actually needed
- Every line of code is covered by tests
- Your design is driven by actual usage

TDD example

Let's implement a `is_palindrome()` function using TDD:

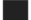
Step 1 - Write the test(s) first

```
1 test_that("is_palindrome works correctly", {
2   expect_true(is_palindrome(c(1, 2, 3, 2, 1)))
3   expect_true(is_palindrome(c("a", "b", "a")))
4   expect_false(is_palindrome(c(1, 2, 3)))
5   expect_true(is_palindrome(c(5))) # Single element
6   expect_true(is_palindrome(numeric(0))) # Empty vector
7 })
```

— Error: `is_palindrome works correctly` —————

Error in `is_palindrome(c(1, 2, 3, 2, 1))`: could not find function "is_palindrome"

Backtrace:

1.  `testthat::expect_true(is_palindrome(c(1, 2, 3, 2, 1)))`
2. `testthat::quasi_label(enquo(object), label)`
3. `rlang::eval_bare(expr, quo_get_env(quo))`

Error:

! Test failed with 1 failure and 0 successes.

Step 2

Write minimal code to pass:

```
1 is_palindrome = function(x) {  
2   all(x == rev(x))  
3 }
```

Which we then check with our existing tests:

```
1 test_that("is_palindrome works correctly", {  
2   expect_true(is_palindrome(c(1, 2, 3, 2, 1)))  
3   expect_true(is_palindrome(c("a", "b", "a")))  
4   expect_false(is_palindrome(c(1, 2, 3)))  
5   expect_true(is_palindrome(c(5))) # Single element  
6   expect_true(is_palindrome(numeric(0))) # Empty vector  
7 })
```

Test passed with 5 successes 🎉.

Step 3: Refactor

We can consider a slightly improved implementation:

```
1 is_palindrome = function(x) {  
2   identical(x, rev(x))  
3 }
```

Which we again verify with the tests:

```
1 test_that("is_palindrome works correctly", {  
2   expect_true(is_palindrome(c(1, 2, 3, 2, 1)))  
3   expect_true(is_palindrome(c("a", "b", "a")))  
4   expect_false(is_palindrome(c(1, 2, 3)))  
5   expect_true(is_palindrome(c(5))) # Single element  
6   expect_true(is_palindrome(numeric(0))) # Empty vector  
7 })
```

Test passed with 5 successes 🐱.

Step 4: Repeat

We can consider additional functionality, such as input validation by expanding our tests:

```
1 test_that("is_palindrome errors for non-atomic input", {
2   expect_error(is_palindrome(list(1, 2, 1)))
3 })
```

— Failure: is_palindrome errors for non-atomic input —
Expected `is_palindrome(list(1, 2, 1))` to throw an error.

Error:

! Test failed with 1 failure and 0 successes.

```
1 is_palindrome = function(x) {
2   stopifnot("Input must be an atomic vector" = is.atomic(x))
3   identical(x, rev(x))
4 }
```

```
1 test_that("is_palindrome errors for non-atomic input", {
2   expect_error(is_palindrome(list(1, 2, 1)))
3 })
```

Test passed with 1 success 🎉.

TDD in the real world

In practice, TDD may not be followed strictly, but the principles remain valuable:

- Tests should guide your design and implementation
- Tests should not be an afterthought once your code is “done”
- Refactoring is easier and safer with a solid test suite
- Writing tests 2nd can lead to missing edge cases / faulty assumptions

Why Packages?

Benefits of packages

Organizing your projects as a package provides many advantages:

- Benefit from the existing infrastructure for package development
- Easier to share and distribute your code (dependencies, installation, documentation, etc.)
- Easier to bundle and document data sets
- Better support for testing and documentation
- Tends to lead to better organized, modular code and overall better design

Packages and LLMs

We will go into this more on Thursday, but packages are also a great way to structure your code to work with LLMs:

- Prescribed structure makes it easier for the LLMs to understand your codebase
- Better context management
- Better grounding and easier iteration through tests and checks