# Mutating Joins

Max Turgeon

SCI 2000–Introduction to Data Science

# Lecture Objectives

- Understand the difference between the different types of mutating joins
- Be able to choose and select the appropriate mutating join

# Motivation

- For the first part of the course, we focused on the *data analysis* part of data science.
    - It probably felt like a sped-up version of STAT 1150
- The datasets we analyzed were provided either as part of packages, or as a single CSV file.
- In the second part of the course, we will discuss some strategies for collecting and combining datasets.

# Data science life cycle

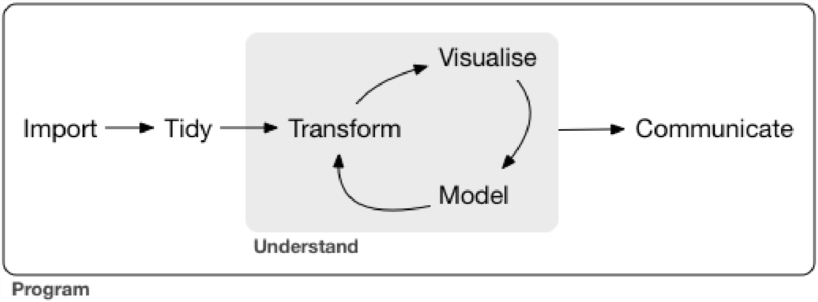


Source: `https://r4ds.had.co.nz/introduction.html`

# Relational data i

- So far we've looked at data that fits neatly into a `data.frame`.
    - Each row is an observation, and for each observation we collected the **same** variables.
- This is not the only way to store data. Let's look at an example: university course enrollment data.
    - For every student we need to collect personal information.
    - For every course we need to collect specific information.
- Clearly these datasets should be separate; you can think of them as two different `data.frame`s.
- **Question**: How should we store information about which courses students are taking?

# Relational data ii

- Should we add the name of courses to the student `data.frame` as new variables? How many variables should we create?
- Should we add the name of students to the course `data.frame` as new variables? How many variables should we create?
- **A better solution**: Create a *new* dataset, where each row corresponds to a pair (student, course).
- **Why does this work?** Each student has a **unique** identifier, and so does each course.

# Relational data iii

- To create a class list:
    - Filter the (student, course) `data.frame` to only keep pairs for a given course.
    - Look up which students appear in the filtered dataset
    - Keep relevant personal information (e.g. student number, major, degree)
- The process of “looking up” is called a **mutating join**.

# Example i

- This dataset is separated into two CSV files:
  - One contains a list of 2,410 US craft beers
  - The other contains data on 510 US breweries
- The beers and breweries datasets have a variable in common, called `brewery_id`.

```
library(tidyverse)

df_beers <- read_csv("beers.csv")
df_breweries <- read_csv("breweries.csv")

glimpse(df_beers)
```

# Example ii

```
## Rows: 2,410
## Columns: 7
## $ abv <dbl> 0.050, 0.066, 0.071, 0.090, 0~
## $ ibu <dbl> NA, NA, NA, NA, NA, NA, NA, N~
## $ id <dbl> 1436, 2265, 2264, 2263, 2262,~
## $ name <chr> "Pub Beer", "Devil's Cup", "R~
## $ style <chr> "American Pale Lager", "Ameri~
## $ brewery_id <dbl> 408, 177, 177, 177, 177,
177,~
## $ ounces <dbl> 12, 12, 12, 12, 12, 12, 12, 1~
```

# Example iii

```r
glimpse(df_breweries)
```

```
## Rows: 558
## Columns: 4
## $ brewery_id <dbl> 0, 1, 2, 3, 4, 5, 6, 7, 8,
9,~
## $ name <chr> "NorthGate Brewing", "Against~
## $ city <chr> "Minneapolis", "Louisville", ~
## $ state <chr> "MN", "KY", "MA", "CA", "CA",~
```

# Mutating joins

- **Mutating joins** create a new dataset by combining two datasets and respecting their relationship.
    - This relationship is encoded by a common variable (or set of variables), often a unique identifier.
- The main idea is as follows:
    - Take a row from the first dataset
    - Find a matching row in the second dataset
    - Create a new row by concatenating the two rows
- The different types of mutating joins differ in how they handle cases with no matches.

# Inner join

- In inner joins, we only create a new row if we can match rows from both datasets.

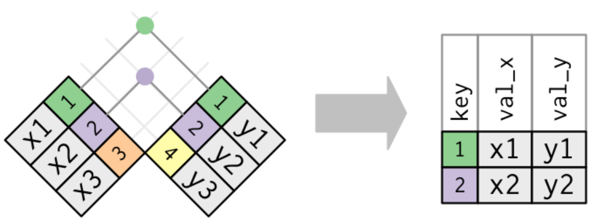

# Example i

```
library(tidyverse)

dataset <- inner_join(df_beers,
                      df_breweries,
                      by = "brewery_id")
glimpse(dataset)
```

# Example ii

```
## Rows: 2,410
## Columns: 10
## $ abv <dbl> 0.050, 0.066, 0.071, 0.090, 0~
## $ ibu <dbl> NA, NA, NA, NA, NA, NA, NA, N~
## $ id <dbl> 1436, 2265, 2264, 2263, 2262,~
## $ name.x <chr> "Pub Beer", "Devil's Cup", "R~
## $ style <chr> "American Pale Lager", "Ameri~
## $ brewery_id <dbl> 408, 177, 177, 177, 177,
177,~
## $ ounces <dbl> 12, 12, 12, 12, 12, 12, 12, 1~
## $ name.y <chr> "10 Barrel Brewing Company", ~
## $ city <chr> "Bend", "Gary", "Gary", "Gary~
```

# Example iii

```
## $ state <chr> "OR", "IN", "IN", "IN", "IN",~
```

```r
# dataset and df_beers have the same # of rows
nrow(dataset) == nrow(df_beers)
```

```
## [1] TRUE
```

```r
# dataset has one less than the sum of # cols
c(ncol(dataset), ncol(df_beers), ncol(df_breweries))
```

```
## [1] 10  7  4
```

# Exercise

Find the state with the highest average of alcohol by volume (`abv`) per beers.

# Solution i

- Now that the datasets are joined, we can use `group_by` and `summarise`.

```
dataset %>%
  group_by(state) %>%
  summarise(avg_abv = mean(abv)) %>%
  filter(avg_abv == max(avg_abv))
```

```
## # A tibble: 0 x 2
## # ... with 2 variables: state <chr>, avg_abv
<dbl>
```

# Solution ii

```r
# What went wrong?
# Let's look at the data
dataset %>%
  filter(is.na(abv)) %>%
  glimpse
```

```
## Rows: 62
## Columns: 10
## $ abv <dbl> NA, NA, NA, NA, NA, NA, NA, N~
## $ ibu <dbl> NA, NA, NA, NA, NA, NA, NA, N~
## $ id <dbl> 1541, 1025, 2490, 2489, 2488,~
```

# Solution iii

```
## $ name.x <chr> ”Double Play Pilsner”, ”N Str~
## $ style <chr> ”American Pilsner”, ”American~
## $ brewery_id <dbl> 380, 380, 77, 77, 77, 77,
10,~
## $ ounces <dbl> 12, 12, 12, 12, 12, 12, 16, 1~
## $ name.y <chr> ”Blue Blood Brewing Company”,~
## $ city <chr> ”Lincoln”, ”Lincoln”, ”Austin~
## $ state <chr> ”NE”, ”NE”, ”TX”, ”TX”, ”TX”,~
```

# Solution iv

```
# NA in abv trickles down to the average
dataset %>%
  group_by(state) %>%
  summarise(avg_abv = mean(abv))
```

```
## # A tibble: 51 x 2
##    state avg_abv
##  * <chr>   <dbl>
##  1 AK     0.0556
##  2 AL     0.062
##  3 AR     0.052
```

# Solution v

```
##  4 AZ    NA
##  5 CA    NA
##  6 CO    NA
##  7 CT      0.0611
##  8 DC      0.0656
##  9 DE    NA
## 10 FL    NA
## # ... with 41 more rows
```

# Solution vi

```r
# Solution: na.rm = TRUE in mean
dataset %>%
  group_by(state) %>%
  summarise(avg_abv = mean(abv, na.rm = TRUE)) %>%
  filter(avg_abv == max(avg_abv))
```

```
## # A tibble: 1 x 2
##   state avg_abv
##   <chr>   <dbl>
## 1 NV     0.0669
```

# Left/right join i

- But what if we want to keep rows from a dataset that don't have a matching row in the other dataset?
- Left and right (outer) joins will do just that and replace the non-matching row with `NA`s.
- Left and right refer to the dataset from which we want to keep rows.
    - `left_join(x, y)` will keep rows of `x`
    - `right_join(x, y)` will keep rows of `y`

# Left/right join ii

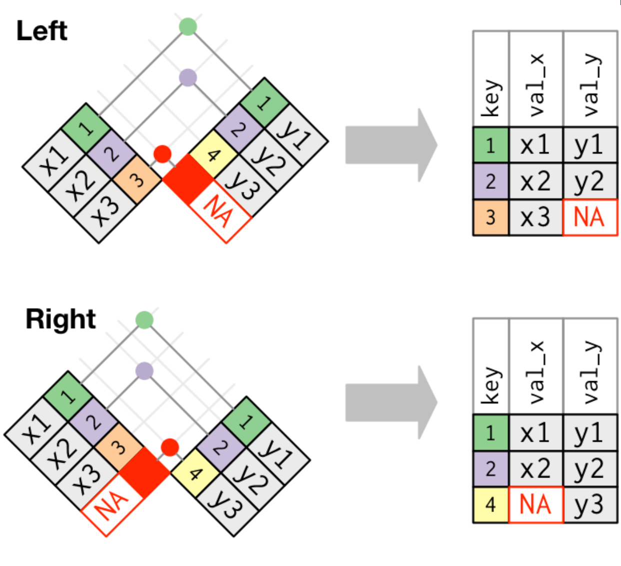

# Example i

```r
library(nycflights13)

# Information about flights
glimpse(flights)
```

```
## Rows: 336,776
## Columns: 19
## $ year <int> 2013, 2013, 2013, 2013, 2~
## $ month <int> 1, 1, 1, 1, 1, 1, 1, 1, 1~
## $ day <int> 1, 1, 1, 1, 1, 1, 1, 1, 1~
## $ dep_time <int> 517, 533, 542, 544, 554, ~
```

# Example ii

```
## $ sched_dep_time <int> 515, 529, 540, 545,
600, ~
## $ dep_delay <dbl> 2, 4, 2, -1, -6, -4, -5, ~
## $ arr_time <int> 830, 850, 923, 1004, 812,~
## $ sched_arr_time <int> 819, 830, 850, 1022,
837,~
## $ arr_delay <dbl> 11, 20, 33, -18, -25, 12,~
## $ carrier <chr> "UA", "UA", "AA", "B6", "~
## $ flight <int> 1545, 1714, 1141, 725, 46~
## $ tailnum <chr> "N14228", "N24211", "N619~
## $ origin <chr> "EWR", "LGA", "JFK", "JFK~
## $ dest <chr> "IAH", "IAH", "MIA", "BQN~
```

# Example iii

```
## $ air_time <dbl> 227, 227, 160, 183, 116, ~
## $ distance <dbl> 1400, 1416, 1089, 1576, 7~
## $ hour <dbl> 5, 5, 5, 5, 6, 5, 6, 6, 6~
## $ minute <dbl> 15, 29, 40, 45, 0, 58, 0,~
## $ time_hour <dttm> 2013-01-01 05:00:00, 201~
```

```r
# Information about airplanes
glimpse(planes)
```

# Example iv

```
## Rows: 3,322
## Columns: 9
## $ tailnum <chr> ”N10156”, ”N102UW”, ”N103US~
## $ year <int> 2004, 1998, 1999, 1999, 200~
## $ type <chr> ”Fixed wing multi engine”, ~
## $ manufacturer <chr> ”EMBRAER”, ”AIRBUS
INDUSTRI~
## $ model <chr> ”EMB-145XR”, ”A320-214”, ”A~
## $ engines <int> 2, 2, 2, 2, 2, 2, 2, 2, 2, ~
## $ seats <int> 55, 182, 182, 182, 55, 182,~
## $ speed <int> NA, NA, NA, NA, NA, NA, NA,~
## $ engine <chr> ”Turbo-fan”, ”Turbo-fan”, ”~
```

# Example v

```r
# How many flights? How many planes?
c(nrow(flights), nrow(planes))
```

```
## [1] 336776   3322
```

```r
# How many flights have matching plane?
inner_join(flights, planes, by = "tailnum") %>%
  nrow
```

```
## [1] 284170
```

# Example vi

```r
# With left_join, we keep all flights
left_join(flights, planes, by = "tailnum") %>%
  glimpse
```

```
## Rows: 336,776
## Columns: 27
## $ year.x <int> 2013, 2013, 2013, 2013, 2~
## $ month <int> 1, 1, 1, 1, 1, 1, 1, 1, 1~
## $ day <int> 1, 1, 1, 1, 1, 1, 1, 1, 1~
## $ dep_time <int> 517, 533, 542, 544, 554, ~
## $ sched_dep_time <int> 515, 529, 540, 545,
```

# Example vii

```
600, ~
## $ dep_delay <dbl> 2, 4, 2, -1, -6, -4, -5, ~
## $ arr_time <int> 830, 850, 923, 1004, 812,~
## $ sched_arr_time <int> 819, 830, 850, 1022,
837,~
## $ arr_delay <dbl> 11, 20, 33, -18, -25, 12,~
## $ carrier <chr> "UA", "UA", "AA", "B6", "~
## $ flight <int> 1545, 1714, 1141, 725, 46~
## $ tailnum <chr> "N14228", "N24211", "N619~
## $ origin <chr> "EWR", "LGA", "JFK", "JFK~
## $ dest <chr> "IAH", "IAH", "MIA", "BQN~
## $ air_time <dbl> 227, 227, 160, 183, 116, ~
```

# Example viii

```
## $ distance <dbl> 1400, 1416, 1089, 1576, 7~
## $ hour <dbl> 5, 5, 5, 5, 6, 5, 6, 6, 6~
## $ minute <dbl> 15, 29, 40, 45, 0, 58, 0,~
## $ time_hour <dttm> 2013-01-01 05:00:00, 201~
## $ year.y <int> 1999, 1998, 1990, 2012, 1~
## $ type <chr> "Fixed wing multi engine"~
## $ manufacturer <chr> "BOEING", "BOEING",
"BOEI~
## $ model <chr> "737-824", "737-824", "75~
## $ engines <int> 2, 2, 2, 2, 2, 2, 2, 2, 2~
## $ seats <int> 149, 149, 178, 200, 178, ~
## $ speed <int> NA, NA, NA, NA, NA, NA, N~
```

# Example ix

```
## $ engine <chr> ”Turbo-fan”, ”Turbo-fan”,~
```

# Full join

- The full join allows us to keep unmatched rows from **both** datasets.

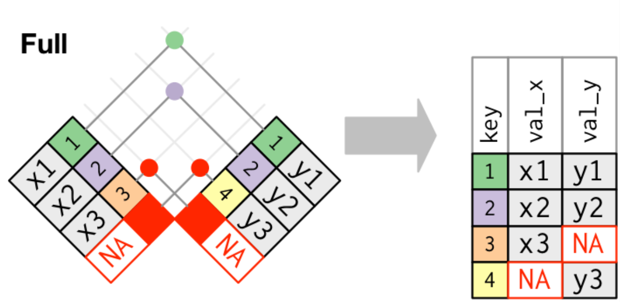

The `flights` dataset contains information about departure and arrival delays (`dep_delay` and `arr_delay`). Compute the average delays for each manufacturing year (i.e. the year the plane was manufactured). Plot the relationship between these two quantities. Do you see any evidence of an association?

## Solution i

```r
# First we combine the two datasets
# and create tot_delay variable
# Note: Could also use inner_join
dataset <- left_join(flights,
                     planes,
                     by = "tailnum") %>%
  mutate(tot_delay = dep_delay + arr_delay)
```

# Solution ii

```r
# Next group by year and summarise
data_avg <- dataset %>%
  group_by(year) %>%
  summarise(avg_delay = mean(tot_delay, na.rm = TRUE))
```

```
## Error: Must group by variables found in `.data`.
## * Column `year` is not found.
```

```r
# What happened?
# Both flights and planes have a variable year
# year.y refers to the one from planes
names(dataset)
```

# Solution iii

```
## [1] ”year.x” ”month” ”day” ”dep_time”
## [5] ”sched_dep_time” ”dep_delay” ”arr_time”
”sched_arr_time”
## [9] ”arr_delay” ”carrier” ”flight” ”tailnum”
## [13] ”origin” ”dest” ”air_time” ”distance”
## [17] ”hour” ”minute” ”time_hour” ”year.y”
## [21] ”type” ”manufacturer” ”model” ”engines”
## [25] ”seats” ”speed” ”engine” ”tot_delay”
```

## Solution iv

```
# Try again
data_avg <- dataset %>%
  group_by(year.y) %>%
  summarise(avg_delay = mean(tot_delay, na.rm = TRUE))
```

```
data_avg %>%
  ggplot(aes(x = year.y,
             y = avg_delay)) +
  geom_point()
```

# Solution v

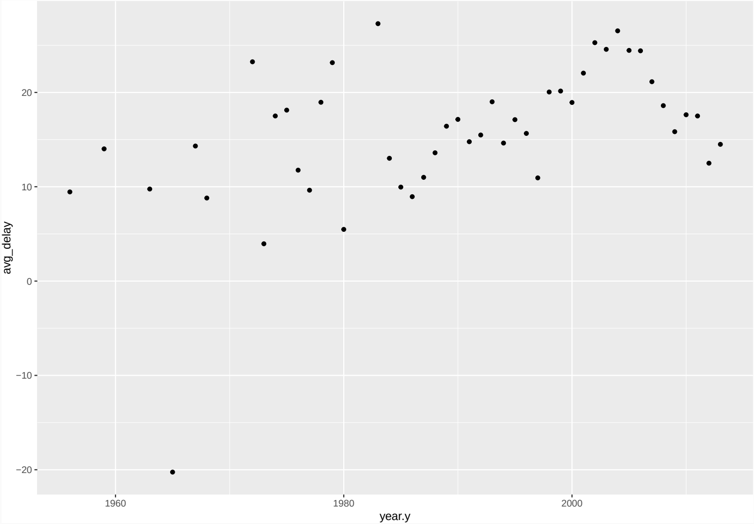

# Summary

- Not all data is neatly packaged into CSV files.
- Often the data we need is spread over multiple datasets.
- If these datasets have a matching variable, we can create a new dataset with matching rows using **mutating joins**.
- Choosing between an inner join, left/right join or full join depends on what we want to do with unmatched rows.
    - Do we keep all of them? Only those from one of the two datasets?