

The Mode: Takeaways

by Dataquest Labs, Inc. - All rights reserved © 2021

Syntax

- Coding from scratch a function that computes the mode of a vector:

```
library(dplyr)
compute_mode <- function(vector) {
  counts_df <- tibble(vector) %>%
    group_by(vector) %>%
    summarise(frequency=n()) %>%
    arrange(desc(frequency))
  counts_df$vector[1]
}
```

- Plotting kernel density of a distribution

```
ggplot(data = df,
  aes(x = distribution)) +
  geom_density(alpha = 0.1,
    color='blue',
    fill='blue') +
  geom_vline(aes(xintercept = compute_mode(distribution),
    color = 'Mode'),
    size = 1.2 ) +
  geom_vline(aes(xintercept = median(distribution),
    color = 'Median'),
    size = 1.2 ) +
  geom_vline(aes(xintercept = mean(distribution),
    color = 'Mean'),
    size = 1.2 ) +
  scale_y_continuous(labels = scales::comma) +
  scale_x_continuous(labels = scales::comma,
    lim = c(min(df$distribution),
      max(df$distribution))) +
  scale_colour_manual(values = c("Mode" = "green",
    "Median" = "black",
    "Mean" = "orange"),
    name = "") +
  theme_bw() +
  theme(legend.position='top') +
  xlab("Distribution") +
  ylab("Density")
```

Concepts

- The **most frequent** value in the distribution is called **the mode**.

- A distribution can have:
 - One mode (**unimodal** distribution).
 - Two modes (**bimodal** distribution).
 - More than two modes (**multimodal** distribution).
 - **No mode** (as for a perfectly uniform distribution or the distribution of a continuous variable).
- The mode is an ideal summary metric for:
 - **Nominal** data.
 - **Ordinal** data (especially when the values are represented using words).
 - **Discrete** data (when we need to communicate the average value to a non-technical audience).
- The location of the mean, median, and mode is usually predictable for certain kinds distributions:
 - **Left-skewed** distributions: the mode is on the far right, the median is to the left of the mode, and the mean is to the left of the median.
 - **Right-skewed** distributions: the mode is on the far left, the median is to the right of the mode, and the mean is to the right of the median.
 - **Normal** distributions: the mean, the median, and the mode are all in the center of the distribution.
 - **Uniform** distributions: the mean and the median are at the center, and there's no mode.
 - Any **symmetrical** distribution: the mean and the median are at the center, while the position of the mode may vary, and there can also be symmetrical distributions having more than one mode (see example in the lesson).

Resources

- [The Wikipedia entry](#) on the mode.
- [Paul von Hippel's paper](#) addressing patterns in skewed distributions.
- [Split-apply-combine](#) workflow to compute the frequency of each value in the vector.