RMarkdown dokument

Vukoje, Leo

Supplement / Prilog

Publication year / Godina izdavanja: **2024**

Permanent link / Trajna poveznica: https://urn.nsk.hr/urn:nbn:hr:195:052438

Rights / Prava: In copyright / Zaštićeno autorskim pravom.

Download date / Datum preuzimanja: **2024-10-15**

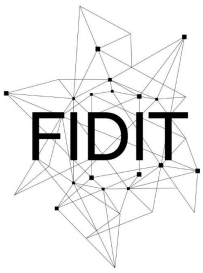

Sveučilište u Rijeci
Fakultet informatike i digitalnih tehnologija

Repository / Repozitorij:

Repository of the University of Rijeka, Faculty of Informatics and Digital Technologies - INFORI Repository

Preparation and exploratory analysis of a dataset with car prices

Ian Kollipara
Last Tested

Required Libraries

```
tidymodels
ggplot2
dplyr
tidyr
readr
purrr
```

Data loading

```
library(tidymodels)
data <- read_csv("data.csv")
```

```
## Rows: 11914
## Columns: 16
## ── Column specification ────────────────────────────────────────────────────────
## Delimiter: ","
## chr (10): Make, Model, Engine Fuel Type, Transmission Type, Driven_Wheels, M...
## dbl  (6): Year, Engine HP, Engine Cylinders, Number of Doors, Highway MPG, ...
##
## ℹ Use `spec()` to retrieve the full column specification for this data.
## ℹ Specify the column types or set `show_col_types = FALSE` to quiet this message.
```

Descriptive statistics

Range, median, mode and quantiles for each variable except the car identifier

```
summary(data[,-c(1)])
```

```
##      Make               Model               Year      Engine Fuel Type
##  Length:11914       Length:11914       Min.   :1990   Length:11914
##  Class :character   Class :character   1st Qu.:2007   Class :character
##  Mode  :character   Mode  :character   Median :2015   Mode  :character
##                                        Mean   :2011
##                                        3rd Qu.:2016
##                                        Max.   :2017
##
##   Engine HP      Engine Cylinders  Number of Doors Transmission Type
##  Min.   : 55.0   Min.   : 0.000   Min.   :2.000   Length:11914
##  1st Qu.:170.0   1st Qu.: 4.000   1st Qu.:2.000   Class :character
##  Median :227.0   Median : 6.000   Median :4.000   Mode  :character
##  Mean   :249.4   Mean   : 5.628   Mean   :3.437
##  3rd Qu.:300.0   3rd Qu.: 6.000   3rd Qu.:4.000
##  Max.   :999.0   Max.   :16.000   Max.   :4.000
##  NA's   :69      NA's   :30       NA's   :6
##  Driven_Wheels      Market Category      Vehicle Size       Vehicle Style
##  Length:11914       Length:11914       Length:11914       Length:11914
##  Class :character   Class :character   Class :character   Class :character
##  Mode  :character   Mode  :character   Mode  :character   Mode  :character
##
##
##
##   Highway MPG      City mpg        Popularity       Price
##  Min.   : 12.00   Min.   : 7.00   Min.   :   2   Min.   :   2000
##  1st Qu.: 22.00   1st Qu.:16.00   1st Qu.:1436   1st Qu.:  21000
##  Median : 26.00   Median :18.00   Median :1436   Median :  29995
##  Mean   : 26.64   Mean   :19.73   Mean   :1554   Mean   :  40594
##  3rd Qu.: 30.00   3rd Qu.:22.00   3rd Qu.:1436   3rd Qu.:  42651
##  Max.   :354.00   Max.   :137.00  Max.   :5657   Max.   :2065902
##
```

Standard Deviation

```
data[,-c(1,2,3,4,5,6,7,8,9,10,11,12,13,14,15,16)] %>% sapply(sd, na.rm = TRUE)
```

```
##  Engine HP Engine Cylinders  Number of Doors     Highway MPG       City mpg
##   109.1916     1.7806        0.8805        8.8634        8.9871
##    Popularity         Price
##   1441.8574    60109.1028
```

Variance

```
sapply(data[,-c(1,2,3,4,5,6,7,8,9,10,11,12,13,14,15,16)], var, na.rm = TRUE)
```

```
##  Engine HP Engine Cylinders  Number of Doors     Highway MPG       City mpg
##   11916.70       3.17         0.78        78.59        80.77
##    Popularity         Price
##  2078760.89   3613142718.46
```

Visualization

Price Range Frequency Histogram

```
library(ggplot2)
ggplot(data, aes(x = Price)) + geom_histogram(bins = 100, fill = "blue", na.rm = TRUE)
```

```
# Add a title and axis labels
ggplot(data, aes(x = Price)) +
  geom_histogram(bins = 100, fill = "blue", na.rm = TRUE) +
  labs(title = "Frequency of Cars in Price Ranges",
       x = "Price",
       y = "Frequency")
```


Average Price by Car Length Line Chart

```
# Calculate the average price for each length
average_prices_by_length <- data %>%
  group_by(Length = cut(Year, breaks = seq(min(Year), max(Year), by = 1))) %>%
  summarise(Average_Price = mean(Price, na.rm = TRUE))

# Create the line chart
ggplot(average_prices_by_length, aes(x = Length, y = Average_Price)) +
  geom_line(group = 1, color = "blue") +
  geom_point(color = "red") +
  labs(title = "Average Price by Car Length",
       x = "Length",
       y = "Average Price") +
  theme_minimal() +
  theme(axis.text.x = element_text(angle = 45, hjust = 1))
```


Average Price by Car Width Line Chart

```
# Calculate the average price for each width
average_prices_by_width <- data %>%
  group_by(Width = cut(Year, breaks = seq(min(Year), max(Year), by = 1))) %>%
  summarise(Average_Price = mean(Price, na.rm = TRUE))

# Create the line chart
ggplot(average_prices_by_width, aes(x = Width, y = Average_Price)) +
  geom_line(group = 1, color = "blue") +
  geom_point(color = "red") +
  labs(title = "Average Price by Car Width",
       x = "Width",
       y = "Average Price") +
  theme_minimal() +
  theme(axis.text.x = element_text(angle = 45, hjust = 1))
```


Average Price by Curb Weight Line Chart

```
# Calculate the average price for each curb weight
average_prices_by_curb_weight <- data %>%
  group_by(Curb_Weight = cut(Year, breaks = seq(min(Year), max(Year), by = 1))) %>%
  summarise(Average_Price = mean(Price, na.rm = TRUE))

# Create the line chart
ggplot(average_prices_by_curb_weight, aes(x = Curb_Weight, y = Average_Price)) +
  geom_line(group = 1, color = "blue") +
  geom_point(color = "red") +
  labs(title = "Average Price by Curb Weight",
       x = "Curb Weight",
       y = "Average Price") +
  theme_minimal() +
  theme(axis.text.x = element_text(angle = 45, hjust = 1))
```


Average Price by Engine Size Line Chart

```
# Calculate the average price for each engine size
average_prices_by_engine_size <- data %>%
  group_by(Engine_Size = cut(Year, breaks = seq(min(Year), max(Year), by = 1))) %>%
  summarise(Average_Price = mean(Price, na.rm = TRUE))

# Create the line chart
ggplot(average_prices_by_engine_size, aes(x = Engine_Size, y = Average_Price)) +
  geom_line(group = 1, color = "blue") +
  geom_point(color = "red") +
  labs(title = "Average Price by Engine Size",
       x = "Engine Size",
       y = "Average Price") +
  theme_minimal() +
  theme(axis.text.x = element_text(angle = 45, hjust = 1))
```


Average Price by Horsepower Line Chart

```
# Calculate the average price for each horsepower
average_prices_by_horsepower <- data %>%
  group_by(Horsepower = cut(Year, breaks = seq(min(Year), max(Year), by = 1))) %>%
  summarise(Average_Price = mean(Price, na.rm = TRUE))

# Create the line chart
ggplot(average_prices_by_horsepower, aes(x = Horsepower, y = Average_Price)) +
  geom_line(group = 1, color = "blue") +
  geom_point(color = "red") +
  labs(title = "Average Price by Horsepower",
       x = "Horsepower",
       y = "Average Price") +
  theme_minimal() +
  theme(axis.text.x = element_text(angle = 45, hjust = 1))
```


Average Price by City MPG Line Chart

```
# Calculate the average price for each city MPG
average_prices_by_city_mpg <- data %>%
  group_by(City_MPG = cut(Year, breaks = seq(min(Year), max(Year), by = 1))) %>%
  summarise(Average_Price = mean(Price, na.rm = TRUE))

# Create the line chart
ggplot(average_prices_by_city_mpg, aes(x = City_MPG, y = Average_Price)) +
  geom_line(group = 1, color = "blue") +
  geom_point(color = "red") +
  labs(title = "Average Price by City MPG",
       x = "City MPG",
       y = "Average Price") +
  theme_minimal() +
  theme(axis.text.x = element_text(angle = 45, hjust = 1))
```


Average Price by Highway MPG Line Chart

```
# Calculate the average price for each highway MPG
average_prices_by_highway_mpg <- data %>%
  group_by(Highway_MPG = cut(Year, breaks = seq(min(Year), max(Year), by = 1))) %>%
  summarise(Average_Price = mean(Price, na.rm = TRUE))

# Create the line chart
ggplot(average_prices_by_highway_mpg, aes(x = Highway_MPG, y = Average_Price)) +
  geom_line(group = 1, color = "blue") +
  geom_point(color = "red") +
  labs(title = "Average Price by Highway MPG",
       x = "Highway MPG",
       y = "Average Price") +
  theme_minimal() +
  theme(axis.text.x = element_text(angle = 45, hjust = 1))
```


Correlational Matrix

```
numeric_data <- data[sapply(data, is.numeric)]
cor_matrix <- cor(numeric_data, use = "complete.obs")
```

```
## Warning in cor(numeric_data, use = "complete.obs"): the standard deviation is
## zero
```


Inferential Statistics

Hypothesis Testing

T-test

```
t_test_result <- t.test(data$Price ~ data$Year)
print(t_test_result)
```

```
##
##  Welch Two Sample t-test
##
## data:  data$Price by data$Year
## t = -0.61, df = 1.1, p-value = 0.6
## alternative hypothesis: true difference in means between group 2011 and group 2017 is not equal to 0
## 95 percent confidence interval:
##  -1.0e+08  1.0e+08
## sample estimates:
## mean in group 2011 mean in group 2017
##           40594.74           40594.74
```

ANOVA test

```
anova_result <- aov(Price ~ Year, data = data)
summary(anova_result)
```

```
##                Df    Sum Sq   Mean Sq F value Pr(>F)
## Year            1 1.000e+12 1.000e+12   0.016  0.899
## Residuals   11912 7.000e+14 5.877e+10
```

Chi-square test

```
chi_square_result <- chisq.test(data$Year, data$Price)
```

```
## Warning in chisq.test(data$Year, data$Price): Chi-squared approximation may be
## incorrect
```

```
print(chi_square_result)
```

```
##
##  Pearson's Chi-squared test
##
## data:  data$Year and data$Price
## X-squared = 1.2e+08, df = 1.2e+08, p-value = 1
```

Categorical table

```
table_data <- table(data$Year, data$Price)
```

```
##
##      2000  2001  2002  2003  2004  2005  2006  2007  2008  2009  2010  2011
##   1     0     0     0     0     0     0     0     0     0     0     0     0
##   2     0     0     0     0     0     0     0     0     0     0     0     0
##   3     0     0     0     0     0     0     0     0     0     0     0     0
##   4     0     0     0     0     0     0     0     0     0     0     0     0
##   5     0     0     0     0     0     0     0     0     0     0     0     0
##   6     0     0     0     0     0     0     0     0     0     0     0     0
##   7     0     0     0     0     0     0     0     0     0     0     0     0
##   8     0     0     0     0     0     0     0     0     0     0     0     0
##   9     0     0     0     0     0     0     0     0     0     0     0     0
##   10    0     0     0     0     0     0     0     0     0     0     0     0
##   11    0     0     0     0     0     0     0     0     0     0     0     0
##   12    0     0     0     0     0     0     0     0     0     0     0     0
##   13    0     0     0     0     0     0     0     0     0     0     0     0
##   14    0     0     0     0     0     0     0     0     0     0     0     0
##   15    0     0     0     0     0     0     0     0     0     0     0     0
##   16    0     0     0     0     0     0     0     0     0     0     0     0
##   17    0     0     0     0     0     0     0     0     0     0     0     0
##
##      2012  2013  2014  2015  2016  2017
##   1     0     0     0     0     0     0
##   2     0     0     0     0     0     0
##   3     0     0     0     0     0     0
##   4     0     0     0     0     0     0
##   5     0     0     0     0     0     0
##   6     0     0     0     0     0     0
##   7     0     0     0     0     0     0
##   8     0     0     0     0     0     0
##   9     0     0     0     0     0     0
##   10    0     0     0     0     0     0
##   11    0     0     0     0     0     0
##   12    0     0     0     0     0     0
##   13    0     0     0     0     0     0
##   14    0     0     0     0     0     0
##   15    0     0     0     0     0     0
##   16    0     0     0     0     0     0
##   17    0     0     0     0     0     0
```

Confidence Intervals

```
t_test_result <- t.test(data$Price)
print(t_test_result)
```

```
##
##  One Sample t-test
##
## data:  data$Price
## t = 73.6, df = 11913, p-value < 2.2e-16
## alternative hypothesis: true mean is not equal to 0
## 95 percent confidence interval:
##  39410.29 41778.19
## sample estimates:
## mean of x
##  40594.24
```

```
desired_confidence_level <- 0.95
t_test_result <- t.test(data$Price, conf.level = desired_confidence_level)
print(t_test_result)
```

```
##
##  One Sample t-test
##
## data:  data$Price
## t = 73.6, df = 11913, p-value < 2.2e-16
## alternative hypothesis: true mean is not equal to 0
## 95 percent confidence interval:
##  39410.29 41778.19
## sample estimates:
## mean of x
##  40594.24
```

```
desired_confidence_level <- 0.99
t_test_result <- t.test(data$Price, conf.level = desired_confidence_level)
print(t_test_result)
```

```
##
##  One Sample t-test
##
## data:  data$Price
## t = 73.6, df = 11913, p-value < 2.2e-16
## alternative hypothesis: true mean is not equal to 0
## 99 percent confidence interval:
##  38973.31 42215.17
## sample estimates:
## mean of x
##  40594.24
```

Manual Insertion of NA Values

```
data[1:5, "Price"] <- NA
```

```
## # A tibble: 5 × 16
##   Make  Model  Year `Engine Fuel Type` `Engine HP` `Engine Cylinders`
##   <chr> <chr> <dbl> <chr>                    <dbl>              <dbl>
## 1 BMW   1 Se…  2011 premium unleaded …         335                  6
## 2 BMW   1 Se…  2011 premium unleaded …         300                  6
## 3 BMW   1 Se…  2011 premium unleaded …         300                  6
## 4 BMW   1 Se…  2011 premium unleaded …         230                  6
## 5 BMW   1 Se…  2011 premium unleaded …         230                  6
## # ℹ 10 more variables: `Number of Doors` <dbl>, `Transmission Type` <chr>,
## #   Driven_Wheels <chr>, `Market Category` <chr>, `Vehicle Size` <chr>,
## #   `Vehicle Style` <chr>, `Highway MPG` <dbl>, `City mpg` <dbl>,
## #   Popularity <dbl>, Price <dbl>
```

Determination of Numeric Variables

```
numeric_columns <- data[sapply(data, is.numeric)]
print(numeric_columns)
```

```
## # A tibble: 11,914 × 9
##    Year `Engine HP` `Engine Cylinders` `Number of Doors` `Highway MPG`
##   <dbl>       <dbl>              <dbl>            <dbl>         <dbl>
## 1  2011         NA                  6                2            26
## 2  2011         NA                  6                2            28
## 3  2011         NA                  6                2            28
## 4  2011         NA                  6                2            28
## 5  2011         NA                  6                2            28
## 6  2011         NA                  6                2            26
## 7  2011         NA                  6                2            27
## 8  2011         NA                  6                2            28
## 9  2011         NA                  6                2            28
## 10 2011         NA                  6                2            28
## # ℹ 11,904 more rows
## # ℹ 4 more variables: `City mpg` <dbl>, Popularity <dbl>, Price <dbl>,
## #   Horsepower <chr>
```

Normalization and Standardization of Numeric Variables

```
numeric_columns <- data[sapply(data, is.numeric)]
print(numeric_columns)
```

```
## # A tibble: 11,914 × 9
##    Year `Engine HP` `Engine Cylinders` `Number of Doors` `Highway MPG`
##   <dbl>       <dbl>              <dbl>            <dbl>         <dbl>
## 1  2011         NA                  6                2            26
## 2  2011         NA                  6                2            28
## 3  2011         NA                  6                2            28
## 4  2011         NA                  6                2            28
## 5  2011         NA                  6                2            28
## 6  2011         NA                  6                2            26
## 7  2011         NA                  6                2            27
## 8  2011         NA                  6                2            28
## 9  2011         NA                  6                2            28
## 10 2011         NA                  6                2            28
## # ℹ 11,904 more rows
## # ℹ 4 more variables: `City mpg` <dbl>, Popularity <dbl>, Price <dbl>,
## #   Horsepower <chr>
```

```
normalized_data <- as.data.frame(lapply(numeric_columns, function(x) (x - min(x)) / (max(x) - min(x))))
print(normalized_data)
```

```
## # A tibble: 11,914 × 9
##    Year `Engine HP` `Engine Cylinders` `Number of Doors` `Highway MPG`
##   <dbl>       <dbl>              <dbl>            <dbl>         <dbl>
## 1  0.84          NA              0.375              0.5         0.068
## 2  0.84          NA              0.375              0.5         0.074
## 3  0.84          NA              0.375              0.5         0.074
## 4  0.84          NA              0.375              0.5         0.074
## 5  0.84          NA              0.375              0.5         0.074
## 6  0.84          NA              0.375              0.5         0.068
## 7  0.84          NA              0.375              0.5         0.071
## 8  0.84          NA              0.375              0.5         0.074
## 9  0.84          NA              0.375              0.5         0.074
## 10 0.84          NA              0.375              0.5         0.074
## # ℹ 11,904 more rows
## # ℹ 4 more variables: `City mpg` <dbl>, Popularity <dbl>, Price <dbl>,
## #   Horsepower <chr>
```

```
standardized_data <- as.data.frame(lapply(numeric_columns, function(x) (x - mean(x)) / sd(x)))
print(standardized_data)
```

```
## # A tibble: 11,914 × 9
##    Year `Engine HP` `Engine Cylinders` `Number of Doors` `Highway MPG`
##   <dbl>       <dbl>              <dbl>            <dbl>         <dbl>
## 1 -0.11          NA              0.107            -0.54        -0.067
## 2 -0.11          NA              0.107            -0.54         0.157
## 3 -0.11          NA              0.107            -0.54         0.157
## 4 -0.11          NA              0.107            -0.54         0.157
## 5 -0.11          NA              0.107            -0.54         0.157
## 6 -0.11          NA              0.107            -0.54        -0.067
## 7 -0.11          NA              0.107            -0.54         0.045
## 8 -0.11          NA              0.107            -0.54         0.157
## 9 -0.11          NA              0.107            -0.54         0.157
## 10 -0.11         NA              0.107            -0.54         0.157
## # ℹ 11,904 more rows
## # ℹ 4 more variables: `City mpg` <dbl>, Popularity <dbl>, Price <dbl>,
## #   Horsepower <chr>
```

Mean and Standard Deviation of Numeric Variables

```
mean_values <- sapply(numeric_columns, mean, na.rm = TRUE)
print(mean_values)
```

```
##         Year    Engine HP Engine Cylinders Number of Doors   Highway MPG
##  2010.384338   249.386070     5.628829787     3.436093190    26.637485970
##     City mpg   Popularity        Price
##    19.733255  1554.911197 40594.737032
```

```
sd_values <- sapply(numeric_columns, sd, na.rm = TRUE)
print(sd_values)
```

```
##         Year    Engine HP Engine Cylinders Number of Doors   Highway MPG
##  7.579740221  68.997672296     1.780836517     0.881314988     8.863401793
##     City mpg   Popularity        Price
##    8.987435405  1441.855347 60109.103604
```

Standardization and Normalization with dplyr

```
standardized_data <- numeric_columns %>% mutate_all(~(scale(.) %>% as.vector))
print(standardized_data)
```

```
## # A tibble: 11,914 × 9
##    Year `Engine HP` `Engine Cylinders` `Number of Doors` `Highway MPG`
##   <dbl>       <dbl>              <dbl>            <dbl>         <dbl>
## 1 -0.11          NA              0.107            -0.54        -0.067
## 2 -0.11          NA              0.107            -0.54         0.157
## 3 -0.11          NA              0.107            -0.54         0.157
## 4 -0.11          NA              0.107            -0.54         0.157
## 5 -0.11          NA              0.107            -0.54         0.157
## 6 -0.11          NA              0.107            -0.54        -0.067
## 7 -0.11          NA              0.107            -0.54         0.045
## 8 -0.11          NA              0.107            -0.54         0.157
## 9 -0.11          NA              0.107            -0.54         0.157
## 10 -0.11         NA              0.107            -0.54         0.157
## # ℹ 11,904 more rows
## # ℹ 4 more variables: `City mpg` <dbl>, Popularity <dbl>, Price <dbl>,
## #   Horsepower <chr>
```

Brand Values Consistency Fix

```
data[!is.na(data$Year) & data$Year == 2011, "Year"] <- "2011"
```

```
## # A tibble: 11,914 × 16
##    Make  Model  Year `Engine Fuel Type` `Engine HP` `Engine Cylinders`
##    <chr> <chr> <dbl> <chr>                    <dbl>              <dbl>
##  1 BMW   1 Se…  2011 premium unleaded …         335                  6
##  2 BMW   1 Se…  2011 premium unleaded …         300                  6
##  3 BMW   1 Se…  2011 premium unleaded …         300                  6
##  4 BMW   1 Se…  2011 premium unleaded …         230                  6
##  5 BMW   1 Se…  2011 premium unleaded …         230                  6
##  6 BMW   1 Se…  2011 premium unleaded …         300                  6
##  7 BMW   1 Se…  2011 premium unleaded …         300                  6
##  8 BMW   1 Se…  2011 premium unleaded …         230                  6
##  9 BMW   1 Se…  2011 premium unleaded …         230                  6
## 10 BMW   1 Se…  2011 premium unleaded …         230                  6
## # ℹ 11,904 more rows
## # ℹ 10 more variables: `Number of Doors` <dbl>, `Transmission Type` <chr>,
## #   Driven_Wheels <chr>, `Market Category` <chr>, `Vehicle Size` <chr>,
## #   `Vehicle Style` <chr>, `Highway MPG` <dbl>, `City mpg` <dbl>,
## #   Popularity <dbl>, Price <dbl>
```

```
data[!is.na(data$Year) & data$Year == 2017, "Year"] <- "2017"
```

```
## # A tibble: 11,914 × 16
##    Make  Model  Year `Engine Fuel Type` `Engine HP` `Engine Cylinders`
##    <chr> <chr> <dbl> <chr>                    <dbl>              <dbl>
##  1 BMW   1 Se…  2011 premium unleaded …         335                  6
##  2 BMW   1 Se…  2011 premium unleaded …         300                  6
##  3 BMW   1 Se…  2011 premium unleaded …         300                  6
##  4 BMW   1 Se…  2011 premium unleaded …         230                  6
##  5 BMW   1 Se…  2011 premium unleaded …         230                  6
##  6 BMW   1 Se…  2011 premium unleaded …         300                  6
##  7 BMW   1 Se…  2011 premium unleaded …         300                  6
##  8 BMW   1 Se…  2011 premium unleaded …         230                  6
##  9 BMW   1 Se…  2011 premium unleaded …         230                  6
## 10 BMW   1 Se…  2011 premium unleaded …         230                  6
## # ℹ 11,904 more rows
## # ℹ 10 more variables: `Number of Doors` <dbl>, `Transmission Type` <chr>,
## #   Driven_Wheels <chr>, `Market Category` <chr>, `Vehicle Size` <chr>,
## #   `Vehicle Style` <chr>, `Highway MPG` <dbl>, `City mpg` <dbl>,
## #   Popularity <dbl>, Price <dbl>
```

Frequency of Cars per Brand

```
brand_counts <- table(data$Make)
print(brand_counts)
```

```
##
##     Alfa Romeo            Audi             BMW           Bentley
##             78             504             787              74
##       Bugatti           Buick         Cadillac       Chevrolet
##              3             196             397            1123
##      Chrysler           Dodge           Ferrari            Fiat
##            187             627              69             808
##         Ford           Genesis           GMC           Honda
##           1475              49             515            1252
##      Hyundai          Infiniti          Kia         Lamborghini
##           1103             330             555              52
##   Land Rover           Lincoln       Lotus           Maserati
##            143             164              29             133
##        Maybach           Mazda        McLaren   Mercedes-Benz
##             16             535              40            1349
##        Mercury          Mitsubishi        Nissan         Oldsmobile
##             11             626            1559              54
##       Plymouth          Pontiac         Porsche          Rolls-Royce
##             19             669             536              31
##          Saab           Scion           Spyker         Subaru
##            111              60               3             708
##         Suzuki          Tesla          Toyota       Volkswagen
##            785              18            1436             809
##          Volvo
##            281
```

```
barplot(brand_counts, main = "Count of Cars by Brands")
```


Categories of Categorical Variables

```
unique_fuel_types <- unique(data$`Engine Fuel Type`)
print(unique_fuel_types)
```

```
## [1] "premium unleaded (required)"
## [2] "regular unleaded"
## [3] "premium unleaded (recommended)"
## [4] "flex-fuel (unleaded/E85)"
## [5] "diesel"
## [6] "electric"
## [7] "flex-fuel (premium unleaded required/E85)"
## [8] "natural gas"
## [9] "flex-fuel (premium unleaded recommended/E85)"
## [10] "flex-fuel (unleaded/natural gas)"
## [11] NA
```

```
unique_transmission_types <- unique(data$`Transmission Type`)
print(unique_transmission_types)
```

```
## [1] "MANUAL"           "AUTOMATIC"        "AUTOMATED_MANUAL" "DIRECT_DRIVE"
## [5] "UNKNOWN"
```

```
unique_driven_wheels <- unique(data$Driven_Wheels)
print(unique_driven_wheels)
```

```
## [1] "rear wheel drive"  "front wheel drive" "all wheel drive"
## [4] "four wheel drive"
```

```
unique_market_categories <- unique(data$`Market Category`)
print(unique_market_categories)
```

```
##  [1] "Factory Tuner,Luxury,High-Performance"
##  [2] "Luxury,Performance"
##  [3] "Luxury,High-Performance"
##  [4] "Luxury"
##  [5] "Performance"
##  [6] "Flex Fuel"
##  [7] "Hatchback"
##  [8] "Hatchback,Performance"
##  [9] "Hatchback,Luxury"
## [10] "Luxury,Performance,Hybrid"
## [11] "Hybrid"
## [12] "Crossover,Hatchback"
## [13] "Crossover"
## [14] "Luxury,Hybrid"
## [15] "Diesel"
## [16] "Luxury,Diesel"
## [17] "Flex Fuel,Performance"
## [18] "Flex Fuel,Luxury,Performance"
## [19] "Flex Fuel,Luxury"
## [20] "Flex Fuel,Hybrid"
## [21] "Exotic,Luxury,High-Performance"
## [22] "Exotic,Luxury,Performance"
## [23] "Exotic"
## [24] "Crossover,Luxury"
## [25] "Crossover,Performance"
## [26] "Flex Fuel,Crossover"
## [27] "Crossover,Diesel"
## [28] "Flex Fuel,Crossover,Hybrid"
## [29] "Flex Fuel,Crossover,Luxury"
## [30] "Flex Fuel,Crossover,Performance"
## [31] "Flex Fuel,Crossover,Luxury,Performance"
## [32] "Flex Fuel,Crossover,Luxury,Hybrid"
## [33] "Flex Fuel,Crossover,Performance,Hybrid"
## [34] "Flex Fuel,Crossover,Luxury,Performance,Hybrid"
## [35] NA
```

```
unique_vehicle_sizes <- unique(data$`Vehicle Size`)
print(unique_vehicle_sizes)
```

```
## [1] "Compact"  "Midsize"  "Large"
```

```
unique_vehicle_styles <- unique(data$`Vehicle Style`)
print(unique_vehicle_styles)
```

```
##  [1] "Coupe"                "Convertible"          "Sedan"
##  [4] "Wagon"                "4dr Hatchback"        "2dr Hatchback"
##  [7] "4dr SUV"              "Crew Cab Pickup"      "Extended Cab Pickup"
## [10] "Regular Cab Pickup"   "2dr SUV"              "Convertible SUV"
## [13] "Passenger Minivan"    "Cargo Minivan"        "Cargo Van"
## [16] "Passenger Van"        "Minivan"
```

Ordinal Encoding

```
data$`Vehicle Size` <- factor(data$`Vehicle Size`, levels = c("Small", "Medium", "Large"), ordered = TRUE)
```

```
## # A tibble: 11,914 × 16
##    Make  Model  Year `Engine Fuel Type` `Engine HP` `Engine Cylinders`
##    <chr> <chr> <dbl> <chr>                    <dbl>              <dbl>
##  1 BMW   1 Se…  2011 premium unleaded …         335                  6
##  2 BMW   1 Se…  2011 premium unleaded …         300                  6
##  3 BMW   1 Se…  2011 premium unleaded …         300                  6
##  4 BMW   1 Se…  2011 premium unleaded …         230                  6
##  5 BMW   1 Se…  2011 premium unleaded …         230                  6
##  6 BMW   1 Se…  2011 premium unleaded …         300                  6
##  7 BMW   1 Se…  2011 premium unleaded …         300                  6
##  8 BMW   1 Se…  2011 premium unleaded …         230                  6
##  9 BMW   1 Se…  2011 premium unleaded …         230                  6
## 10 BMW   1 Se…  2011 premium unleaded …         230                  6
## # ℹ 11,904 more rows
## # ℹ 10 more variables: `Number of Doors` <dbl>, `Transmission Type` <chr>,
## #   Driven_Wheels <chr>, `Market Category` <chr>, `Vehicle Size` <ord>,
## #   `Vehicle Style` <chr>, `Highway MPG` <dbl>, `City mpg` <dbl>,
## #   Popularity <dbl>, Price <dbl>
```

```
data$`Vehicle Size` <- as.numeric(data$`Vehicle Size`)
```

```
## # A tibble: 11,914 × 16
##    Make  Model  Year `Engine Fuel Type` `Engine HP` `Engine Cylinders`
##    <chr> <chr> <dbl> <chr>                    <dbl>              <dbl>
##  1 BMW   1 Se…  2011 premium unleaded …         335                  6
##  2 BMW   1 Se…  2011 premium unleaded …         300                  6
##  3 BMW   1 Se…  2011 premium unleaded …         300                  6
##  4 BMW   1 Se…  2011 premium unleaded …         230                  6
##  5 BMW   1 Se…  2011 premium unleaded …         230                  6
##  6 BMW   1 Se…  2011 premium unleaded …         300                  6
##  7 BMW   1 Se…  2011 premium unleaded …         300                  6
##  8 BMW   1 Se…  2011 premium unleaded …         230                  6
##  9 BMW   1 Se…  2011 premium unleaded …         230                  6
## 10 BMW   1 Se…  2011 premium unleaded …         230                  6
## # ℹ 11,904 more rows
## # ℹ 10 more variables: `Number of Doors` <dbl>, `Transmission Type` <chr>,
## #   Driven_Wheels <chr>, `Market Category` <chr>, `Vehicle Size` <dbl>,
## #   `Vehicle Style` <chr>, `Highway MPG` <dbl>, `City mpg` <dbl>,
## #   Popularity <dbl>, Price <dbl>
```

One-hot Encoding

```
library(fastDummies)
data <- dummy_cols(data, select_columns = c("Transmission Type", "Driven_Wheels", "Vehicle Size", "Vehicle Style"))
```

```
## # A tibble: 11,914 × 32
##    Make  Model  Year `Engine Fuel Type` `Engine HP` `Engine Cylinders`
##    <chr> <chr> <dbl> <chr>                    <dbl>              <dbl>
##  1 BMW   1 Se…  2011 premium unleaded …         335                  6
##  2 BMW   1 Se…  2011 premium unleaded …         300                  6
##  3 BMW   1 Se…  2011 premium unleaded …         300                  6
##  4 BMW   1 Se…  2011 premium unleaded …         230                  6
##  5 BMW   1 Se…  2011 premium unleaded …         230                  6
##  6 BMW   1 Se…  2011 premium unleaded …         300                  6
##  7 BMW   1 Se…  2011 premium unleaded …         300                  6
##  8 BMW   1 Se…  2011 premium unleaded …         230                  6
##  9 BMW   1 Se…  2011 premium unleaded …         230                  6
## 10 BMW   1 Se…  2011 premium unleaded …         230                  6
## # ℹ 11,904 more rows
## # ℹ 26 more variables: `Number of Doors` <dbl>, `Transmission Type` <chr>,
## #   Driven_Wheels <chr>, `Market Category` <chr>, `Vehicle Size` <dbl>,
## #   `Vehicle Style` <chr>, `Highway MPG` <dbl>, `City mpg` <dbl>,
## #   Popularity <dbl>, Price <dbl>, `Transmission Type_AUTOMATED_MANUAL` <int>,
## #   `Transmission Type_AUTOMATIC` <int>, `Transmission Type_DIRECT_DRIVE` <int>,
## #   `Transmission Type_MANUAL` <int>, `Transmission Type_UNKNOWN` <int>, …
```

Persisting Preprocessed Dataset

```
write.csv(data, file = "preprocessed_dataset.csv", row.names = FALSE)
```