

Chapter 3: Linear Regression

Linear regression is a simple approach for supervised learning when the response is quantitative. Linear regression has a long history and we could actually spend most of this semester talking about it.

Although linear regression is not the newest, shiniest thing out there, it is still a highly used technique out in the real world. It is also useful for talking about more modern techniques that are **generalizations** of it. *Ridge regression, Lasso, logistic regression, GAMs...*

We will review some key ideas underlying linear regression and discuss the least squares approach that is most commonly used to fit this model.

Linear regression can help us to answer the following questions about our **Advertising** data:

1. Is there a relationship between advertising and sales?
i.e. should people spend money on ads?
2. How strong is the relationship between ads & sales?
i.e. how well can we predict sales based on ads?
3. Which media contribute to sales?

1 Simple Linear Regression

Simple Linear Regression is an approach for predicting a quantitative response Y on the basis of a single predictor variable X .

It assumes:

Which leads to the following model:

For example, we may be interested in regressing `sales` onto `TV` by fitting the model

Once we have used training data to produce estimates $\hat{\beta}_0$ and $\hat{\beta}_1$, we can predict future sales on the basis of a particular TV advertising budget.

1.1 Estimating the Coefficients

In practice, β_0 and β_1 are **unknown**, so before we can predict \hat{y} , we must use our training data to estimate them.

Let $(x_1, y_1), \dots, (x_n, y_n)$ represent n observation pairs, each of which consists of a measurement of X and Y .

Goal: Obtain coefficient estimates $\hat{\beta}_0$ and $\hat{\beta}_1$ such that the linear model fits the available data well.

The most common approach involves minimizing the *least squares* criterion.

The least squares approach results in the following estimates:

$$\hat{\beta}_1 =$$

$$\hat{\beta}_0 =$$

We can get these estimates using the following commands in R:

```
## load the data in
ads <- read_csv("../data/Advertising.csv")

## fit the model
model <- lm(sales ~ TV, data = ads)

summary(model)

##
## Call:
## lm(formula = sales ~ TV, data = ads)
##
## Residuals:
##      Min       1Q   Median       3Q      Max
## -8.3860 -1.9545 -0.1913  2.0671  7.2124
##
## Coefficients:
##              Estimate Std. Error t value Pr(>|t|)
## (Intercept)  7.032594   0.457843   15.36  <2e-16 ***
## TV           0.047537   0.002691   17.67  <2e-16 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Residual standard error: 3.259 on 198 degrees of freedom
## Multiple R-squared:  0.6119, Adjusted R-squared:  0.6099
## F-statistic: 312.1 on 1 and 198 DF,  p-value: < 2.2e-16
```

1.2 Assessing Accuracy

Recall we assume the *true* relationship between X and Y takes the form

If f is to be approximated by a linear function, we can write this relationship as

and when we fit the model to the training data, we get the following estimate of the *population model*

But how close this this to the truth?

In general, σ^2 is not known, so we estimate it with the *residual standard error*,
 $RSE = \sqrt{RSS/(n - 2)}$.

We can use these standard errors to compute confidence intervals and perform hypothesis tests.

Once we have decided that there is a significant linear relationship between X and Y that is captured by our model, it is natural to ask

To what extent does the model fit the data?

The quality of the fit is usually measured by the *residual standard error* and the R^2 statistic.

RSE: Roughly speaking, the RSE is the average amount that the response will deviate from the true regression line. This is considered a measure of the *lack of fit* of the model to the data.

R^2 : The RSE provides an absolute measure of lack of fit, but is measured in the units of Y . So, we don't know what a "good" RSE value is! R^2 gives the proportion of variation in Y explained by the model.

```
summary(model)
```

```
##
## Call:
## lm(formula = sales ~ TV, data = ads)
##
## Residuals:
##      Min       1Q   Median       3Q      Max
## -8.3860 -1.9545 -0.1913  2.0671  7.2124
##
## Coefficients:
##              Estimate Std. Error t value Pr(>|t|)
## (Intercept)  7.032594   0.457843   15.36  <2e-16 ***
## TV           0.047537   0.002691   17.67  <2e-16 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Residual standard error: 3.259 on 198 degrees of freedom
## Multiple R-squared:  0.6119, Adjusted R-squared:  0.6099
## F-statistic: 312.1 on 1 and 198 DF,  p-value: < 2.2e-16
```

2 Multiple Linear Regression

Simple linear regression is useful for predicting a response based on one predictor variable, but we often have **more than one** predictor.

How can we extend our approach to accommodate additional predictors?

We can give each predictor a separate slope coefficient in a single model.

We interpret β_j as the “average effect on Y of a one unit increase in X_j , *holding all other predictors fixed*”.

In our `Advertising` example,

2.1 Estimating the Coefficients

As with the case of simple linear regression, the coefficients $\beta_0, \beta_1, \dots, \beta_p$ are unknown and must be estimated. Given estimates $\hat{\beta}_0, \hat{\beta}_1, \dots, \hat{\beta}_p$, we can make predictions using the formula

The parameters are again estimated using the same least squares approach that we saw in the context of simple linear regression.

```
# model_2 <- lm(sales ~ TV + radio + newspaper, data = ads)
model_2 <- lm(sales ~ ., data = ads[, -1])

summary(model_2)

##
## Call:
## lm(formula = sales ~ ., data = ads[, -1])
##
## Residuals:
##      Min       1Q   Median       3Q      Max
## -8.8277 -0.8908  0.2418  1.1893  2.8292
##
## Coefficients:
##              Estimate Std. Error t value Pr(>|t|)
## (Intercept)  2.938889   0.311908   9.422  <2e-16 ***
## TV           0.045765   0.001395  32.809  <2e-16 ***
## radio        0.188530   0.008611  21.893  <2e-16 ***
## newspaper   -0.001037   0.005871  -0.177    0.86
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Residual standard error: 1.686 on 196 degrees of freedom
## Multiple R-squared:  0.8972, Adjusted R-squared:  0.8956
## F-statistic: 570.3 on 3 and 196 DF,  p-value: < 2.2e-16
```


2.2 Some Important Questions

When we perform multiple linear regression we are usually interested in answering a few important questions:

- 1.
- 2.
- 3.
- 4.

2.2.1 Is there a relationship between response and predictors?

We need to ask whether all of the regression coefficients are zero, which leads to the following hypothesis test.

$$H_0 :$$

$$H_a :$$

This hypothesis test is performed by computing the F -statistic

$$F =$$

2.2.2 Deciding on Important Variables

After we have computed the F -statistic and concluded that there is a relationship between predictor and response, it is natural to wonder

Which predictors are related to the response?

We could look at the p -values on the individual coefficients, but if we have many variables this can lead to false discoveries.

Instead we could consider *variable selection*. We will revisit this in Ch. 6.

2.2.3 Model Fit

Two of the most common measures of model fit are the RSE and R^2 . These quantities are computed and interpreted in the same way as for simple linear regression.

Be careful with using these alone, because R^2 will **always increase** as more variables are added to the model, even if it's just a small increase.

```
# model with TV, radio, and newspaper
summary(model_2)

##
## Call:
## lm(formula = sales ~ ., data = ads[, -1])
##
## Residuals:
##      Min       1Q   Median       3Q      Max
## -8.8277 -0.8908  0.2418  1.1893  2.8292
##
## Coefficients:
##              Estimate Std. Error t value Pr(>|t|)
## (Intercept)  2.938889   0.311908   9.422  <2e-16 ***
## TV           0.045765   0.001395  32.809  <2e-16 ***
## radio        0.188530   0.008611  21.893  <2e-16 ***
## newspaper   -0.001037   0.005871  -0.177    0.86
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Residual standard error: 1.686 on 196 degrees of freedom
## Multiple R-squared:  0.8972, Adjusted R-squared:  0.8956
## F-statistic: 570.3 on 3 and 196 DF, p-value: < 2.2e-16
```

```
# model without newspaper
summary(lm(sales ~ TV + radio, data = ads))

##
## Call:
## lm(formula = sales ~ TV + radio, data = ads)
##
## Residuals:
##      Min       1Q   Median       3Q      Max
## -8.7977 -0.8752  0.2422  1.1708  2.8328
##
## Coefficients:
##              Estimate Std. Error t value Pr(>|t|)
## (Intercept)  2.92110    0.29449   9.919  <2e-16 ***
## TV           0.04575    0.00139  32.909  <2e-16 ***
## radio        0.18799    0.00804  23.382  <2e-16 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Residual standard error: 1.681 on 197 degrees of freedom
## Multiple R-squared:  0.8972, Adjusted R-squared:  0.8962
## F-statistic: 859.6 on 2 and 197 DF, p-value: < 2.2e-16
```

It may also be useful to plot residuals to get a sense of the model fit.

```
ggplot() +
  geom_point(aes(model_2$fitted.values, model_2$residuals))
```

