Making binary predictions with regression

SUPERVISED LEARNING IN R: CLASSIFICATION

Brett Lantz
Instructor

datacamp

Introducing linear regression

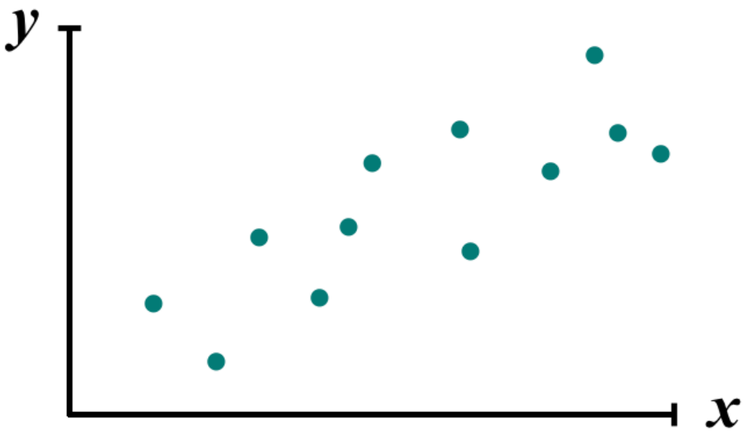

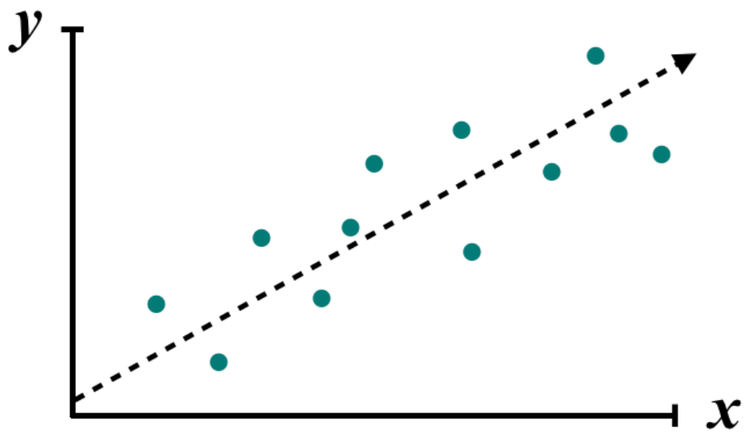

Regression for binary classification

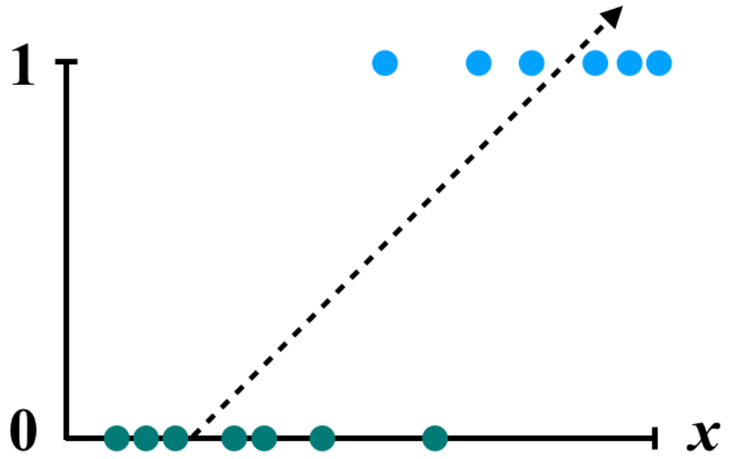

Introducing logistic regression

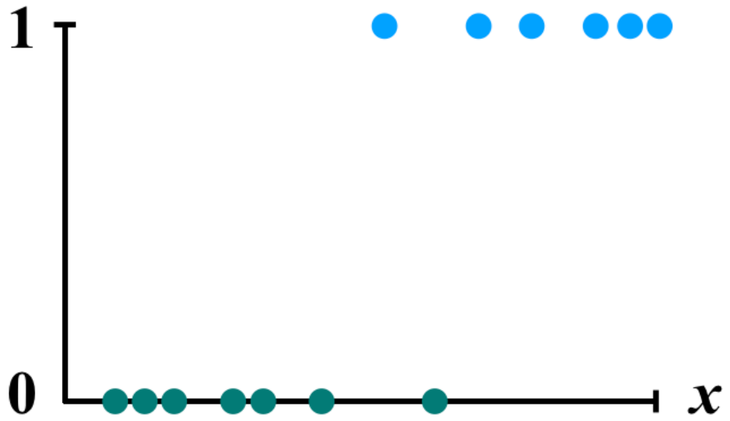

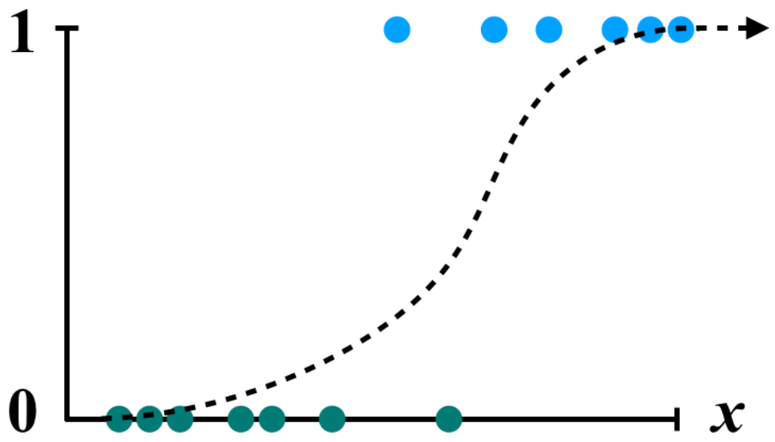

Making predictions with logistic regression

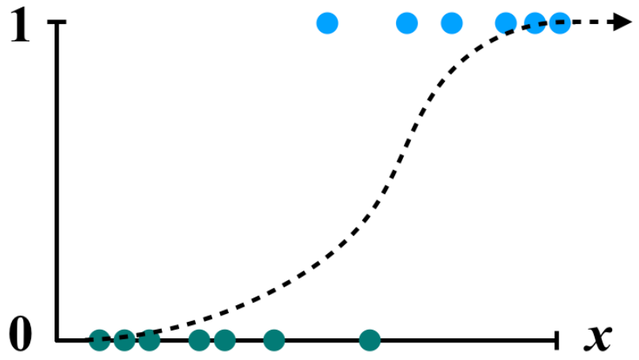


```
m <- glm(y ~ x1 + x2 + x3,
         data = my_dataset,
         family = "binomial")
```

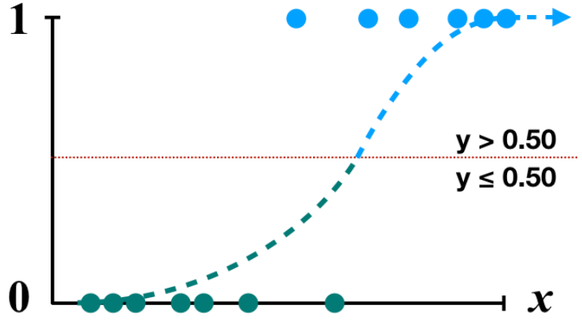


```
prob <- predict(m, test_dataset,
                type = "response")
```

```
pred <- ifelse(prob > 0.50, 1, 0)
```

Let's practice!

SUPERVISED LEARNING IN R: CLASSIFICATION

Model performance tradeoffs

SUPERVISED LEARNING IN R: CLASSIFICATION

Brett Lantz
Instructor

Understanding ROC curves

XOXOXOXOXOXOXOXOXOXO

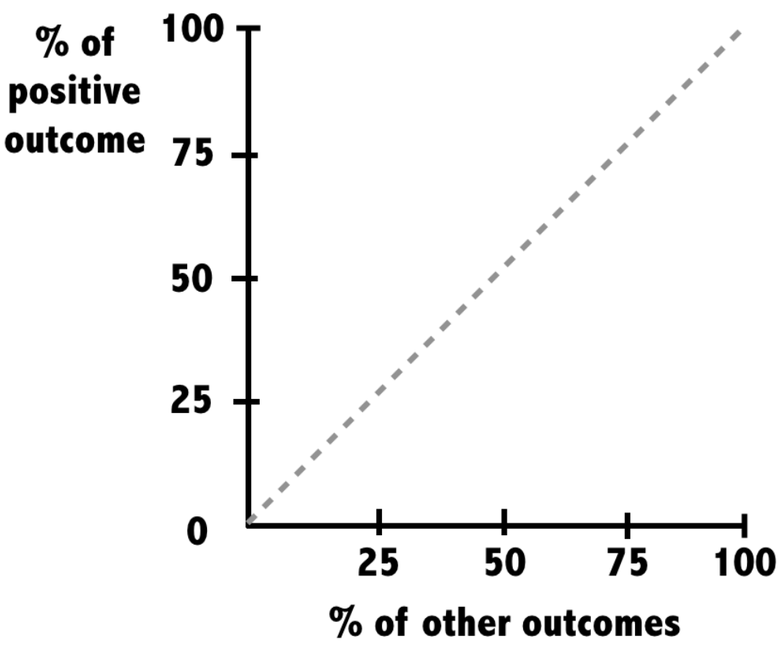

XXOXXXXOXXOOXOXOOOOO

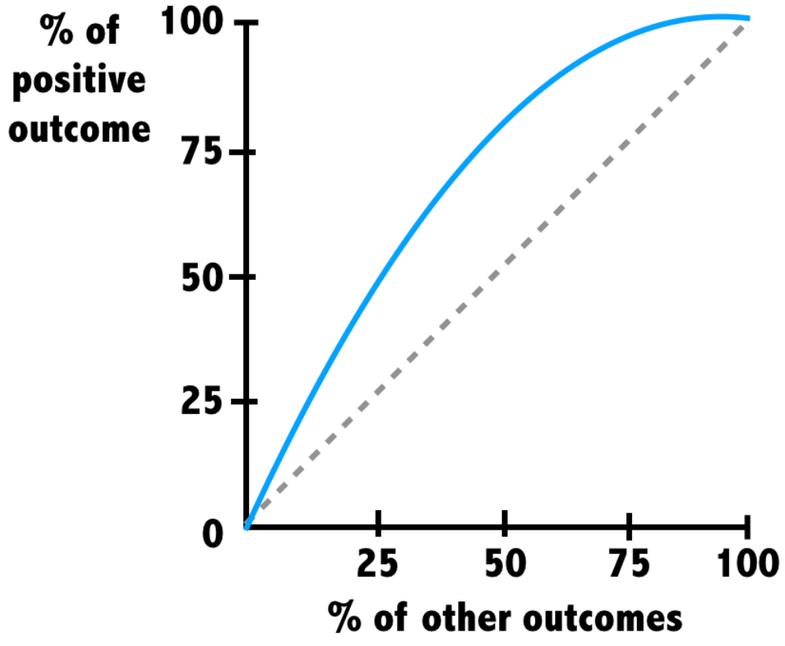

Area under the ROC curve

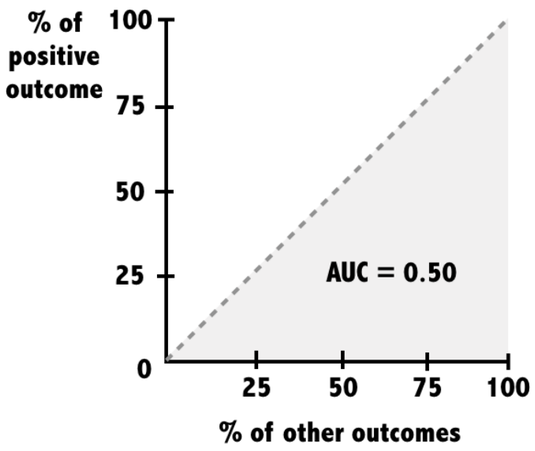

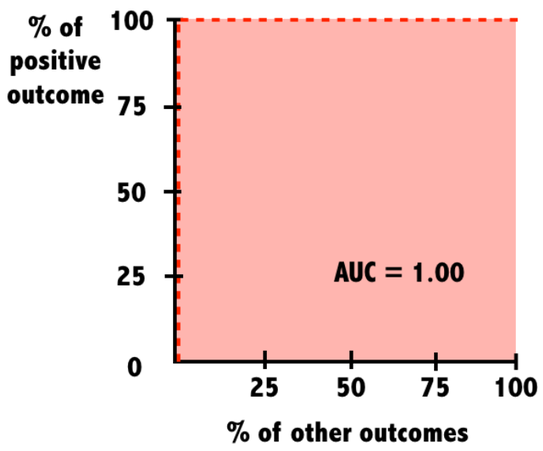

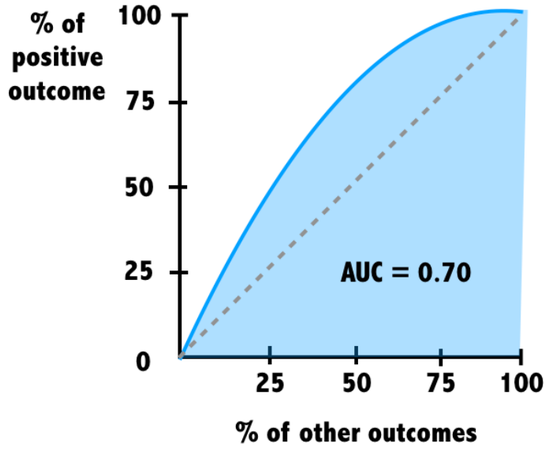

Using AUC and ROC appropriately

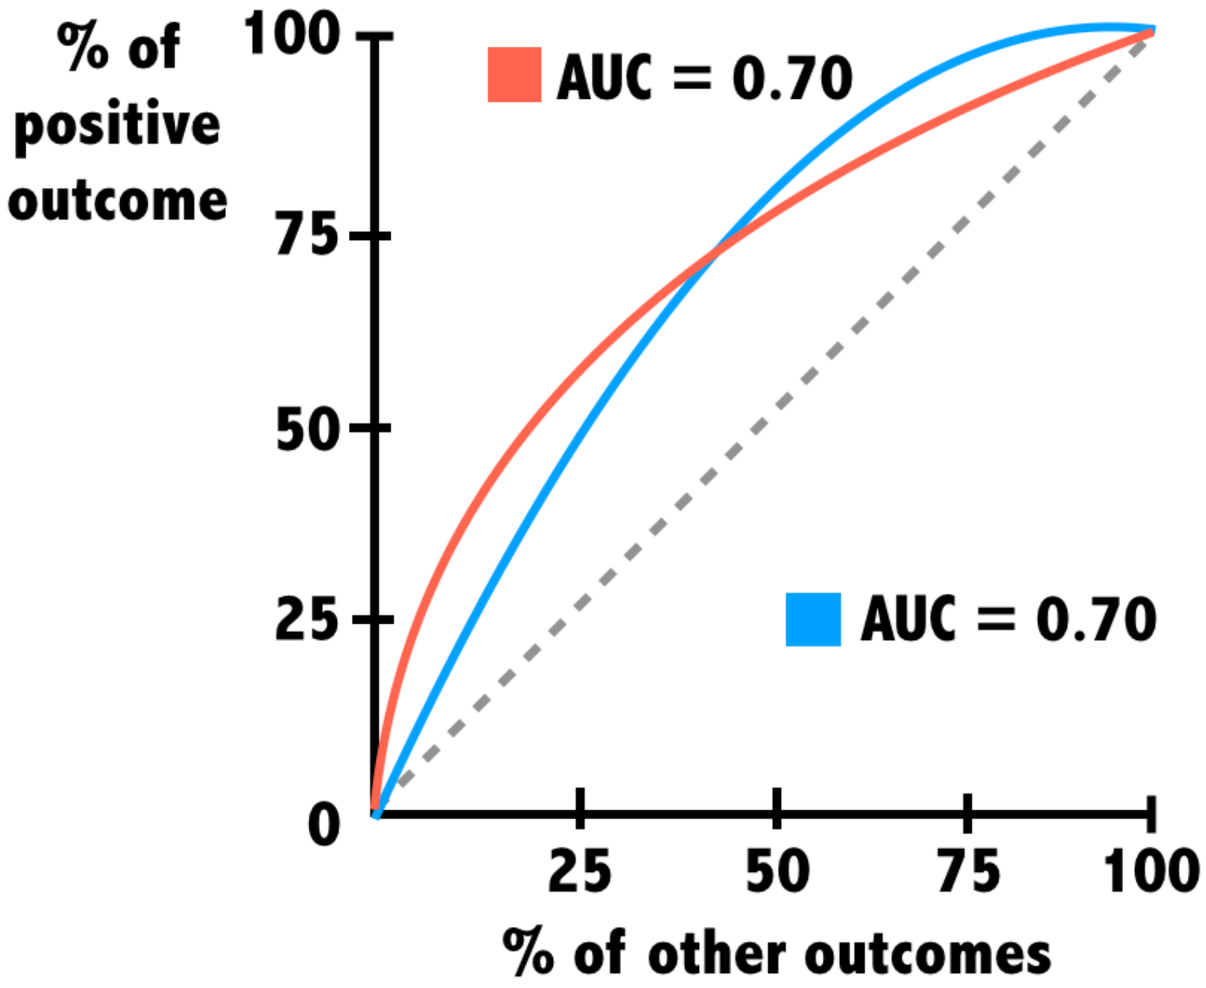

Let's practice!

SUPERVISED LEARNING IN R: CLASSIFICATION

Dummy variables, missing data, and interactions

SUPERVISED LEARNING IN R: CLASSIFICATION

Brett Lantz
Instructor

datacamp

Dummy coding categorical data

```
# create gender factor
my_data$gender <- factor(my_data$gender,
                         levels = c(0, 1, 2),
                         labels = c("Male", "Female", "Other"))
```

Imputing missing data

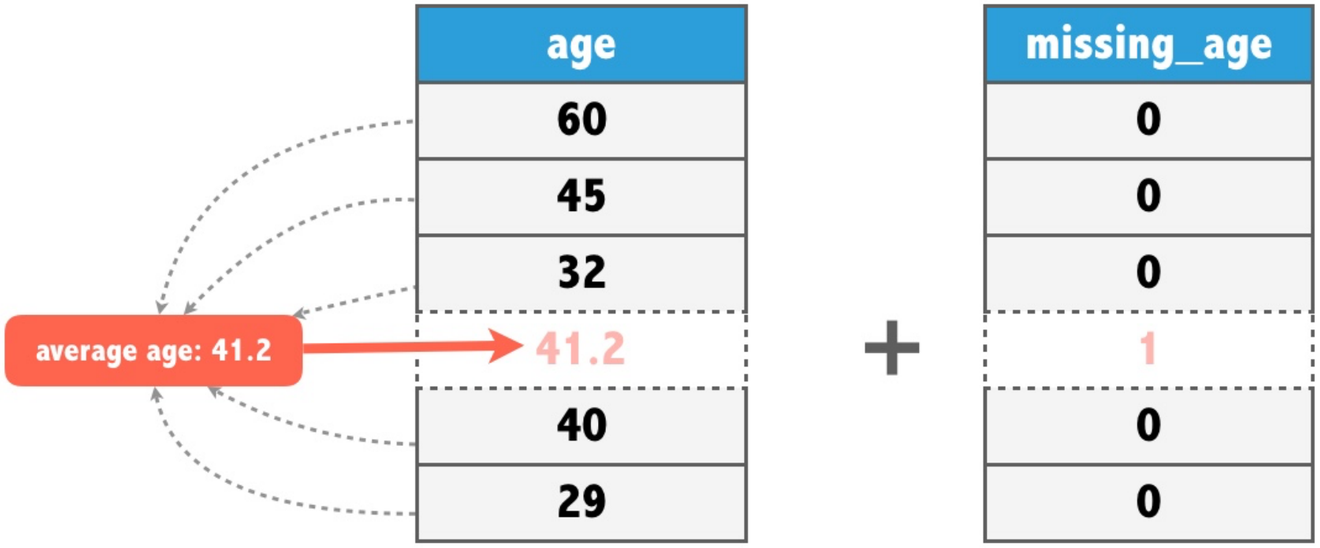

Interaction effects

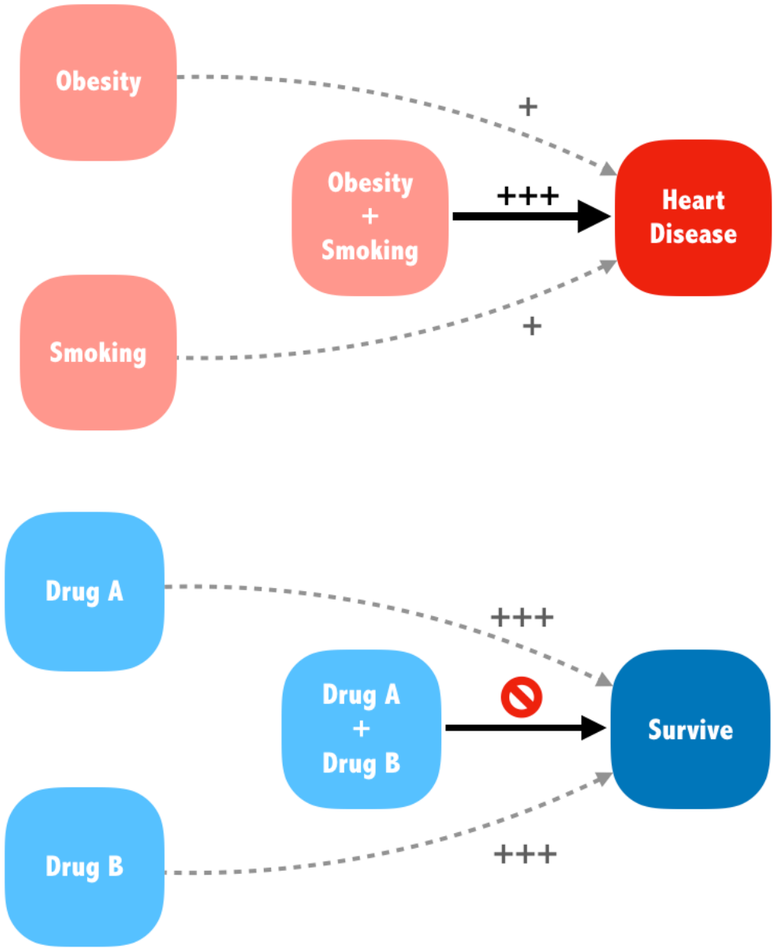


```
# interaction of obesity and smoking
glm(disease ~ obesity * smoking,
      data = health,
      family = "binomial")
```

Let's practice!

SUPERVISED LEARNING IN R: CLASSIFICATION

Automatic feature selection

SUPERVISED LEARNING IN R: CLASSIFICATION

Brett Lantz
Instructor

datacamp

Stepwise regression

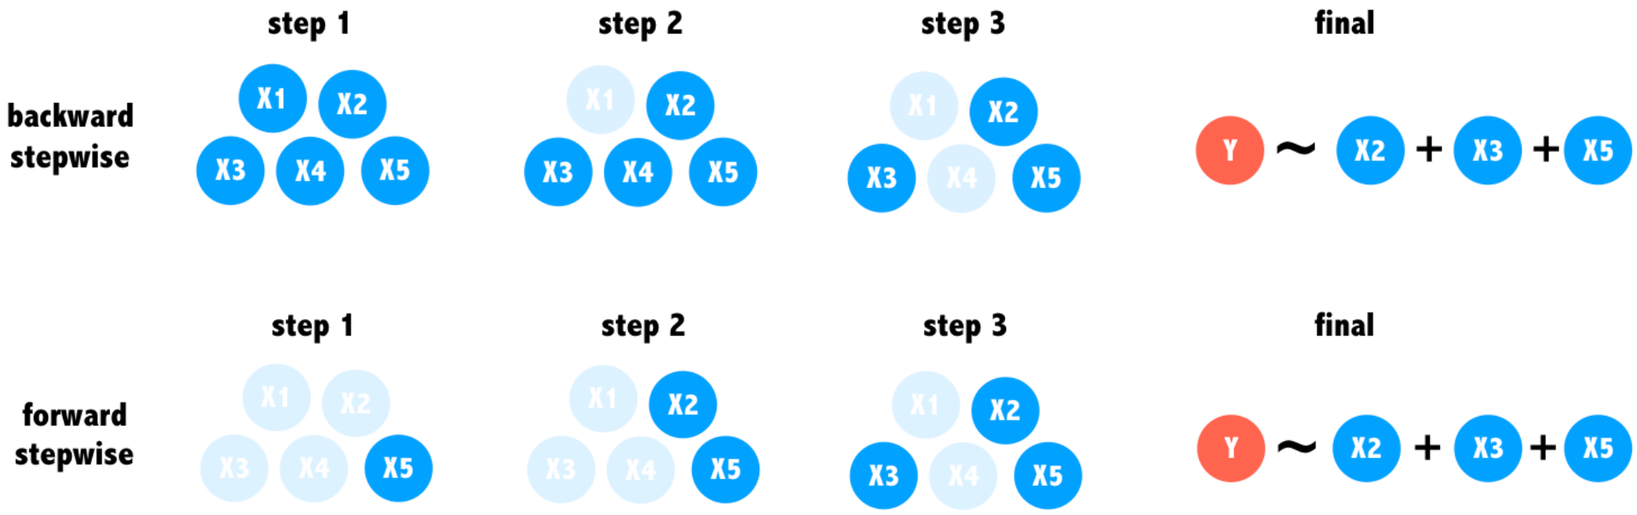

Stepwise regression caveats

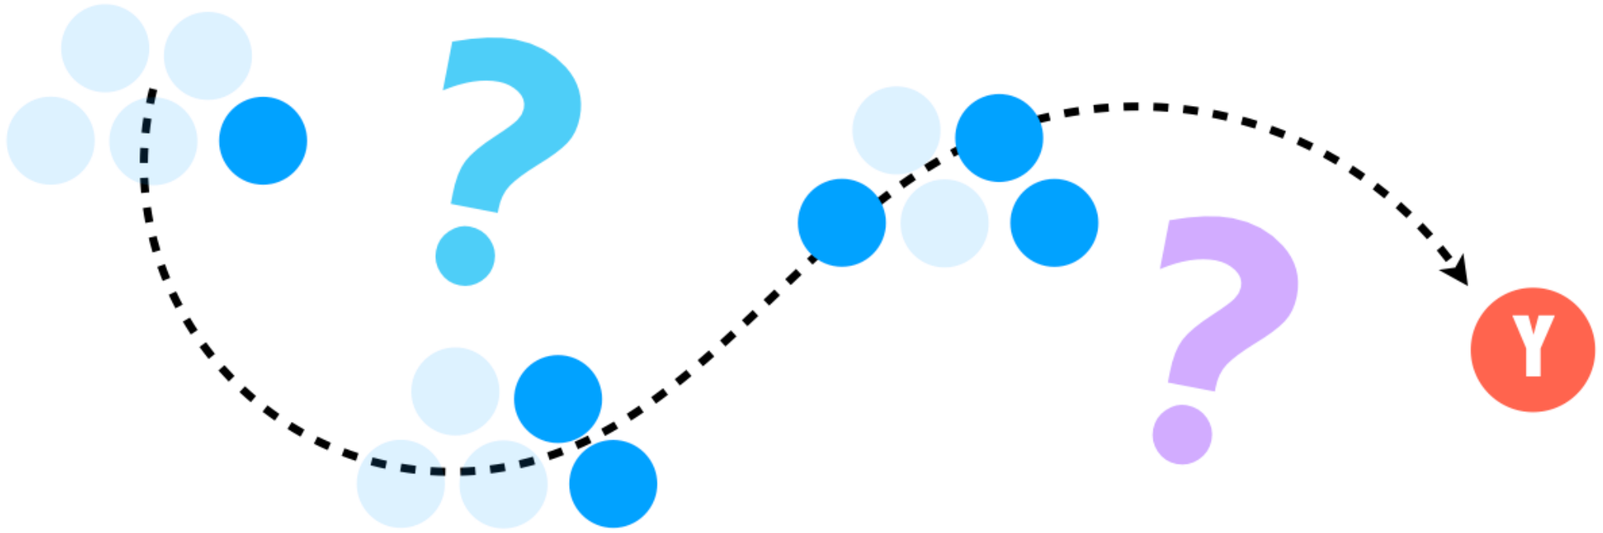

Let's practice!

SUPERVISED LEARNING IN R: CLASSIFICATION