

Lab 13: Logistic Regression

Spam Emails

Today we will be working with a corpus of emails received by a single gmail account over the first three months of 2012. Just like any other email address this account received and sent regular emails as well as receiving a large amount of spam, unsolicited bulk email. We will be using what we have learned about logistic regression models to see if we can build a model that is able to predict whether or not a message is spam based on a variety of characteristics of the email (e.g. inclusion of words like winner, inherit, or password, the number of exclamation marks used, etc.) While the spam filters used by large corporations like Google and Microsoft are quite a bit more complex the fundamental idea is the same - binary classification based on a set of predictors.

Template for lab report

Write your report, or at least run the code and create the plots, as you go so that if you get errors you can ask your TA to help on the spot. Knit often to more easily determine the source of the error.

```
download.file("http://stat.duke.edu/courses/Spring13/sta102.001/Lab/lab13.Rmd", destfile = "lab13.Rmd")
```

The data

```
download.file("http://stat.duke.edu/courses/Spring13/sta102.001/Lab/email.Rdata",  
destfile = "email.Rdata")  
load("email.Rdata")
```

List of variables:

1. `spam`: Indicator for whether the email was spam.
2. `to_multiple`: Indicator for whether the email was addressed to more than one recipient.
3. `from`: Whether the message was listed as from anyone (this is usually set by default for regular outgoing email).
4. `cc`: Indicator for whether anyone was CCed.
5. `sent_email`: Indicator for whether the sender had been sent an email in the last 30 days.
6. `image`: Indicates whether any images were attached.
7. `attach`: Indicates whether any files were attached.
8. `dollar`: Indicates whether a dollar sign or the word 'dollar' appeared in the email.
9. `winner`: Indicates whether "winner" appeared in the email.
10. `inherit`: Indicates whether "inherit" (or an extension, such as inheritance) appeared in the email.
11. `password`: Indicates whether "password" appeared in the email.
12. `num_char`: The number of characters in the email, in thousands.
13. `line_breaks`: The number of line breaks in the email (does not count text wrapping).
14. `format`: Indicates whether the email was written using HTML (e.g. may have included bolding or active links) or plaintext.

15. `re_subj`: Indicates whether the subject started with “Re:”, “RE:”, “re:”, or “rE”:
16. `exclaim_subj`: Indicates whether there was an exclamation point in the subject.
17. `urgent_subj`: Indicates whether the word “urgent” was in the email subject.
18. `exclaim_mess`: The number of exclamation points in the email message.
19. `number`: Factor variable saying whether there was no number, a small number (under 1 million), or a big number.

Exercise 1 In this lab we will focus on predicting whether an email is spam or not. How many emails make up this data set? What proportion of the emails were spam?

A simple spam filter

We will start with a simple spam filter that will only use a single predictor `to_multiple` to classify a message as spam or not. To do this we will fit a logistic regression model between `spam` and `to_multiple` using the `glm` function. This is done in the same way that a simple or multiple regression model is fit in R, except we use the `glm` function instead of `lm`, and we must indicate that we wish to fit a logistic model by include the argument `family=binomial`.

```
g_simple = glm(spam ~ to_multiple, data = email, family = binomial)
summary(g_simple)

##
## Call:
## glm(formula = spam ~ to_multiple, family = binomial, data = email)
##
## Deviance Residuals:
##   Min       1Q   Median       3Q      Max
## -0.477  -0.477  -0.477  -0.477   2.809
##
## Coefficients:
##              Estimate Std. Error z value Pr(>|z|)
## (Intercept)   -2.1161     0.0562  -37.67 < 2e-16 ***
## to_multipleyes -1.8092     0.2969   -6.09 1.1e-09 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## (Dispersion parameter for binomial family taken to be 1)
##
##   Null deviance: 2437.2  on 3920  degrees of freedom
## Residual deviance: 2372.0  on 3919  degrees of freedom
## AIC: 2376
##
## Number of Fisher Scoring iterations: 6
```

Exercise 2 Based on the results of this logistic regression does the inclusion of multiple recipients make a message more or less likely to be spam? Explain your reasoning.

Exercise 3 Using these results calculate the probability that a message is spam if it has multiple recipients, what is the probability if it does not have multiple recipients.

Exercise 4 Pick one of the other 17 remaining predictors that you think is most likely to predict

an email's spam status and fit a `glm` model using it. Describe how this predictor affects the probability of an email being spam.

A more complex spam filter

We will now fit the fully specified model by including all the predictors in our model. Note that we use `'.'` as short hand to indicate that we wish to include all predictors in our model. For the time being we will ignore the warning that is produced when we fit the model.

```
g_full = glm(spam ~ ., data = email, family = binomial)
## Warning: glm.fit: fitted probabilities numerically 0 or 1 occurred
summary(g_full)

##
## Call:
## glm(formula = spam ~ ., family = binomial, data = email)
##
## Deviance Residuals:
##      Min       1Q   Median       3Q      Max
## -2.001  -0.437  -0.175   0.000   3.390
##
## Coefficients:
##              Estimate Std. Error z value Pr(>|z|)
## (Intercept)    2.00e+01  5.92e+03   0.00  0.99731
## to_multipleyes -2.64e+00  3.31e-01  -7.96  1.7e-15 ***
## fromyes        -2.06e+01  5.92e+03   0.00  0.99722
## ccyes          -5.98e-01  3.41e-01  -1.75  0.07947 .
## sent_emailyes -1.71e+01  2.84e+02  -0.06  0.95195
## imageyes       -1.72e+00  8.01e-01  -2.14  0.03210 *
## attachyes      1.28e+00  2.71e-01   4.72  2.3e-06 ***
## dollaryes      1.67e-01  1.81e-01   0.92  0.35516
## winneryes      1.92e+00  3.61e-01   5.33  1.0e-07 ***
## inherityes     3.08e-02  3.40e-01   0.09  0.92785
## passwordyes   -1.79e+00  5.42e-01  -3.30  0.00097 ***
## num_char       3.03e-02  2.42e-02   1.25  0.21065
## line_breaks    -4.77e-03  1.35e-03  -3.53  0.00041 ***
## formatPlain    6.40e-01  1.51e-01   4.25  2.1e-05 ***
## re_subjyes     -1.40e+00  4.11e-01  -3.41  0.00064 ***
## exclaim_subjyes -1.41e-02  2.41e-01  -0.06  0.95343
## urgent_subjyes  3.58e+00  1.24e+00   2.88  0.00396 **
## exclaim_mess    9.64e-03  1.78e-03   5.41  6.3e-08 ***
## numberssmall   -1.23e+00  1.58e-01  -7.82  5.3e-15 ***
## numberbig      -4.62e-01  2.30e-01  -2.01  0.04430 *
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## (Dispersion parameter for binomial family taken to be 1)
##
##      Null deviance: 2437.2  on 3920  degrees of freedom
## Residual deviance: 1668.5  on 3901  degrees of freedom
## AIC: 1708
##
## Number of Fisher Scoring iterations: 18
```

Exercise 5 Imagine we are using this logistic regression model as a spam filter, every new message you receive is analyzed and its characteristics calculated for each of the 18 predictor variables. If the message contained an image is it more or less likely to be flagged as spam? What if it only had an attachment? What if it had an image and an attachment?

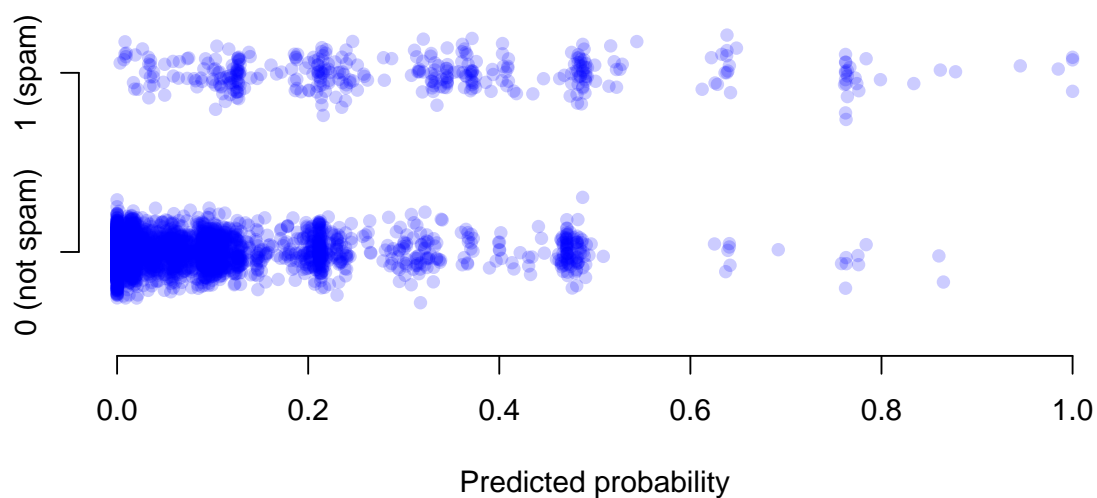
Exercise 6 Which variables appear to be meaningful for identifying spam? Describe how each of these variables affects the probability of an email being spam, numerical answers are not needed here. (For categorical predictors be sure to indicate the reference level and how this affects the interpretation)

Exercise 7 Which predictor appears to have the largest effect? Does this agree with what you have seen in the spam you receive?

Assessing our spam filter


While not quite the same as the residual plots we saw in simple and multiple linear regression we can create a plot that shows how well out logistic regression model is doing by plotting our predicted probability against the response variable (spam or not spam). We jitter the y coordinates slightly so that emails with similar probabilities are not directly on top of one another in the plot.

```
set.seed(1)
jitter = rnorm(nrow(email), sd = 0.08)
plot(g_full$fitted.values, email$spam + jitter, xlim = 0:1, ylim = c(-0.5, 1.5), axes =
FALSE,
      xlab = "Predicted probability", ylab = "", col = adjustcolor("blue", 0.2), pch = 16)
axis(1)
axis(2, at = c(0, 1), labels = c("0 (not spam)", "1 (spam)"))
```



We can also see the difference in the distribution of probabilities for the two classes by plotting side by side box plots.

```
plot(factor(email$spam, labels = c("not spam", "spam")), g_full$fitted.values)
```



From both plots it is clear that in general spam messages have higher probabilities, which is what we expect to see if our spam filter is working well.

Exercise 8 Fit another `glm` model with the subset of the predictors that you think will result in the best possible spam filter. Your criteria should be based on both what you know about spam and the results of fitting the full model. Describe why you chose this particular model and how its results differ from the full model.

Exercise 9 Recreate the above plots for your new model, based on these plots discuss if your model appears to be a better or worse spam filter than the full model. What other information do you think you would need to better make this decision?

Extra Credit In the dot plot above, there appears to be a banding pattern in the predicted probabilities, many points falling at the same or very nearly the same predicted probabilities. Explain why this pattern is occurring.