

## Generalized Linear Models II

Today you use GLMs to explore a data set on ozone pollution:

<http://sciences.ucf.edu/biology/d4lab/wp-content/uploads/sites/125/2017/11/ozonedata.txt>

People chronically exposed to ground-level ozone (this is not about ozone “holes” in the troposphere) are more likely to suffer from asthma, bronchitis, and cardiopulmonary problems. In principle, ozone levels are increased by sunlight (here radiation, or “rad”) and warmer temperatures, but might be diluted by winds.

Your Mission: Use the techniques from the GLMs I class and prior classes to obtain the most plausible model of ozone levels, where the assumption of the statistical distribution is also most legitimate.

*Two jobs: most plausible and most legit.*

Remember these tricks (though you can choose to use them or not):

- **pairs** to squint at all variables in a data set. This can help you decide if you should use:
  - interactive terms (e.g., temp\*wind), or
  - simple ( $y = a + bx$ ), or
  - quadratic (e.g., quadratic;  $y = a + bx + cx^2$ ) functions for predictors.

And yes, a quadratic function is a linear model, because it’s all about the coefficients, and  $x$  is not raised here to some solved-for coefficient (e.g.,  $y = ax^2$ ).

- **vif** in the car package to evaluate multicollinearity of model terms
- **scale** to make all predictors in units of SD (and thus more comparable), despite having different units and ranges
- **AICc** to compare alternative models based on weights, using the bbmle package.
- **glm** is already available in the base stats package, but to run **glm.nb** you will first need to load the MASS package

Having obtained your most plausible and legit model, here is a nice tool to visually compare predictors. A **classification and regression tree** (aka CART) recursively partitions the response variable into subsets based on its relationship to predictor variables. The predictor variable at the first split yields the greatest change in explained deviance (like minimizing SSE in an ANOVA).

Install and load the **tree** package and then run this command on your final, “best” model (where you fill in the predictors):

```
mytree <- tree(ozone ~ predictors)
plot(mytree)
text(mytree)
```

Does this make sense when compared to the scaled coefficients in your output? A limitation: a CART cannot represent interactive terms.