ORIE 3120: Practical Tools for OR, DS, and ML

Forecasting

Professor Udell

Operations Research and Information Engineering
Cornell

May 5, 2020

Announcements

- submit recitation by 4:30pm ET Friday (last recitation!)
- logistic regression homework due 2:30pm ET Wednesday (tomorrow)
- have a question about your grade? (eg, "should I go S/U or GRV?")
  $\rightarrow$ ask in TA office hours!
  (we use breakout rooms so this discussion will be private)
- ask questions about project after class, or in office hours

Project milestone II rubric

- Is the project driven by asking and answering interesting questions?
- How well does the report answer the questions posed?
- Are the visualizations easy to understand? Do they add value?
- Is the report well-written and interesting?
- Does the project use at least 3 tools from class?
  - Linear regression
  - Logistic regression
  - Checking assumptions of linear regression to ensure validity of pvalues
  - Cross-validation or out-of-sample validation
  - Model selection
  - Assessing collinearity
  - Forecasting (with trend / with seasonality)
- How creative are the analyses? Did this project surprise you? Did you learn something?
- Are the techniques they use well-explained and easy to understand?
- Does the project comply with the technical requirements (eg, page limit)? Is it well-formatted and pretty?

Outline

Forecasting time series

A time series, $x_1, x_2, x_3, \ldots$ is a data sequence observed over time, for example,

- demand for parts
- sales of a product
- unemployment rate

In this segment of the course we study special methods for forecasting time series.

- the idea is to develop an algorithm to track the time series and to extrapolate into the future

Outline

Constant mean model: introduction

Suppose demand for a product follows the (very) simple model

$$x_n = a + \epsilon_n$$

Here

- x_n = demand for time period n
- a is the expected demand – which is constant in this simple model
- $\epsilon_1, \epsilon_2, \ldots$ are independent with mean 0
- the best forecast of a future value of x_n is a
- we want to estimate a and update the estimate as each new x_n is observed

Constant mean model: forecasts

- Let $\widehat{x}_n(\ell)$ be the ℓ-step ahead forecast at time period n
- Stated differently, $\widehat{x}_n(\ell)$ is the forecast at time n of demand at time $n+\ell$
- Let

$$\widehat{a}_n = \frac{x_1 + \cdots + x_n}{n}$$

- Then, in this simple model, the best forecasts at time n are

$$\widehat{x}_n(\ell) = \widehat{a}_n, \text{ for all } \ell > 0$$

Constant mean model: updating $\widehat{a}_n$

- In this simple model, a does not change, but our estimate of a does
- Here is a simple way to update $\widehat{a}_n$ to $\widehat{a}_{n+1}$

$$
\begin{aligned}
\widehat{a}_{n+1} &= \frac{(x_1 + \cdots + x_n) + x_{n+1}}{n+1} \\
&= \frac{n}{n+1}\widehat{a}_n + \frac{1}{n+1}x_{n+1} \\
&= \widehat{a}_n + \frac{1}{n+1}(x_{n+1} - \widehat{a}_n)
\end{aligned}
$$

Advantages of the updating formula

The simple updating formula

$$\widehat{a}_{n+1} = \widehat{a}_n + \frac{1}{n+1}(x_{n+1} - \widehat{a}_n)$$

has several advantages:

- reduced storage
  - we only store $\widehat{a}_n$
- computational speed
  - the mean need not be recomputed each time
- suggests ways to handle a slowly changing mean
  - coming soon

Outline

Forecasting: overview

Constant mean model

Simple exponential smoothing

Holt-Winters

- Holt's nonseasonal model
- Winters' seasonal methods

Lake Huron level – example with a slowly changing mean

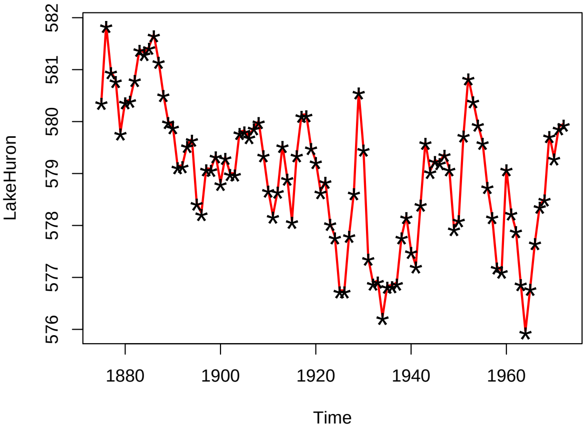

Slowly changing mean model: introduction

- Now suppose that

$$x_n = a_n + \epsilon_n$$

where a_n is slowly changing

- The forecast is the same as for the constant mean model:

$$\widehat{x}_n(\ell) = \widehat{a}_n, \text{ for all } \ell > 0$$

- What changes is the way $\widehat{a}_n$ is updated
  - We need $\widehat{a}_n$ to track a_n

Slowly changing mean: updating

- For a constant mean, the update is

$$\widehat{a}_{n+1} = \widehat{a}_n + \frac{1}{n+1}(x_{n+1} - \widehat{a}_n)$$

- For a slowly changing mean, the update is

$$\widehat{a}_{n+1} = \widehat{a}_n + \alpha(x_{n+1} - \widehat{a}_n) = (1-\alpha)\widehat{a}_n + \alpha x_{n+1}$$

  for a constant α

- α is adjusted depending on how fast a_n is changing
  - $0 < \alpha < 1$
  - faster changes in a necessitate larger α

Demo: Exponential smoothing

https://github.com/madeleineudell/orie3120-sp2020/blob/master/demos/forecasting.ipynb

Exponential weighting

Start with the updating equation and iterate backwards:

$$
\begin{aligned}
\widehat{a}_{n+1} &= (1-\alpha)\widehat{a}_n + \alpha x_{n+1} \\
&= (1-\alpha)\{\widehat{a}_{n-1}(1-\alpha) + \alpha x_n\} + \alpha x_{n+1} \\
&= (1-\alpha)^2\widehat{a}_{n-1} + (1-\alpha)\alpha x_n + \alpha x_{n+1} \\
&= (1-\alpha)^3\widehat{a}_{n-2} + (1-\alpha)^2\alpha x_{n-1} + (1-\alpha)\alpha x_n + \alpha x_{n+1} \\
&\approx \alpha\Big\{ x_{n+1} + (1-\alpha)x_n + (1-\alpha)^2 x_{n-1} \\
&\quad + (1-\alpha)^3 x_{n-2} + \cdots + (1-\alpha)^n x_1 \Big\}
\end{aligned}
$$

Exponential weighted moving average

Use previous page + summation formula for geometric series (see next page):

$$\widehat{a}_{n+1} \approx$$

$$\alpha\left\{(1-\alpha)^0 x_{n+1} + (1-\alpha)^1 x_n + (1-\alpha)^2 x_{n-1} + \cdots + (1-\alpha)^n x_1\right\}$$

$$\approx \frac{\left\{(1-\alpha)^0 x_{n+1} + (1-\alpha)^1 x_n + (1-\alpha)^2 x_{n-1} + \cdots + (1-\alpha)^n x_1\right\}}{1 + (1-\alpha) + \cdots + (1-\alpha)^n}$$

Hence $\widehat{a}_{n+1}$ is an exponentially weighted moving average

Large values of α mean faster discounting of the past values.

Summing a geometric series

Assume $|\gamma| < 1$ so $\gamma^n \to 0$ as $n \to \infty$

$$1 + \gamma + \gamma^2 + \cdots + \gamma^n = \frac{1 - \gamma^{n+1}}{1 - \gamma} \approx \frac{1}{1 - \gamma} \text{ (if } n \text{ is large enough)}$$

Now let $\gamma = 1 - \alpha$. Then

$$1 + (1 - \alpha) + \cdots + (1 - \alpha)^n \approx \frac{1}{\alpha}$$

since $1 - (1 - \alpha) = \alpha$.

Exponential weights: examples

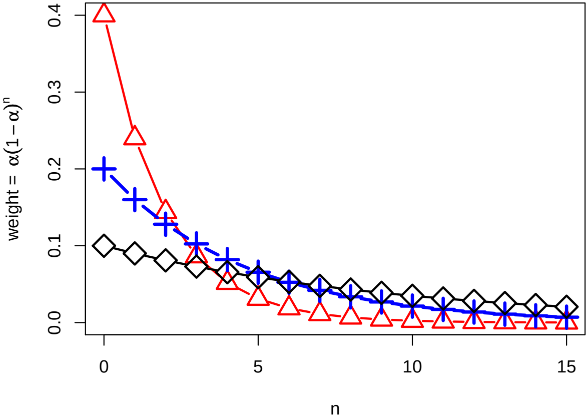

Note: weights start at α when $n = 0$

Outline

Forecasting with trends: example

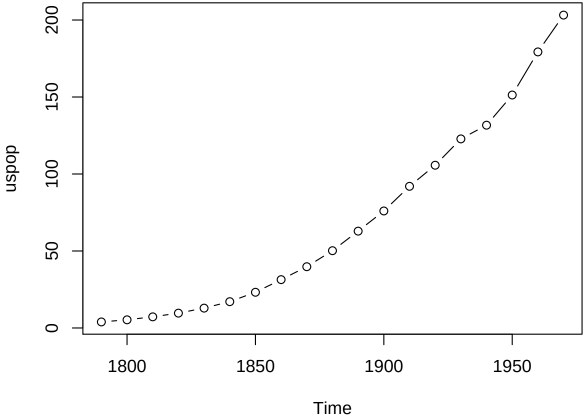

Census counts of US population

Forecasting with trends and seasonality: example

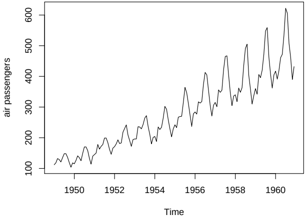

Airline passenger miles

Note the seasonal pattern and trend in this example

- these are typical of much business data

Holt method: forecasting with trend

For now, assume data has trend but no seasonality

Holt's forecasting method uses a linear trend

$$\text{estimate at time } n \text{ of } x_{n+\ell} := \widehat{x}_n(\ell) = \widehat{a}_n + \widehat{b}_n \ell$$

- n is "origin" – the point in time when forecasts are being made
- ℓ is the "lead" – how far ahead one is forecasting
- $\widehat{a}_n$ is called the level
- $\widehat{b}_n$ is called the slope

Both $\widehat{a}_n$ and $\widehat{b}_n$ are updated as we make more observations n

Holt method: Updating the level

In the Holt model, the level $\widehat{a}_n$ is updated by the equation:

$$\widehat{a}_{n+1} = (1-\alpha)(\widehat{a}_n + \widehat{b}_n) + \alpha x_{n+1}$$

or, equivalently,

$$\widehat{a}_{n+1} = \widehat{a}_n + (1-\alpha)\widehat{b}_n + \alpha(x_{n+1} - \widehat{a}_n)$$

- $\widehat{a}_n + \widehat{b}_n$ is predicted value at time $n+1$ = lead time 1
- α is for updating the level and β for the slope (next)

Compare with previous update equation (for no-trend model):

$$\widehat{a}_{n+1} = \widehat{a}_n + \alpha(x_{n+1} - \widehat{a}_n) = (1-\alpha)\widehat{a}_n + \alpha x_{n+1}$$

Holt model: updating the slope

In the Holt model, the slope $\widehat{b}_n$ is updated by the equation:

$$\widehat{b}_{n+1} = (1-\beta)\widehat{b}_n + \beta(\widehat{a}_{n+1} - \widehat{a}_n)$$

or, equivalently,

$$\widehat{b}_{n+1} = \widehat{b}_n + \beta\left\{(\widehat{a}_{n+1} - \widehat{a}_n) - \widehat{b}_n\right\}$$

Demo: Holt's method

https://github.com/madeleineudell/orie3120-sp2020/blob/master/demos/forecasting.ipynb

Winters' additive seasonal method

Winters extended Holt's method to include seasonality. The method is usually called Holt-Winters forecasting

Let s be the period length:

- $s = 4$ for quarterly data
- $s = 12$ for monthly data
- $s = 52$ for weekly data
- $s = 13$ for data collected over 4-week periods
- $s = 24$ for hourly data

Holt-Winters updating

Holt-Winters forecasting can use either of two types of updating

- additive
- multiplicative

These refer to how the trend and seasonal components are put together

- the trend and seasonal components can be added or multiplied

Holt-Winters additive seasonal method

The forecasts are periodic

With the additive methods they are:

$$\begin{aligned} \widehat{x}_n(\ell) &= \widehat{a}_n + \widehat{b}_n \ell + \hat{S}_{n+\ell-s}, \text{ for } \ell = 1, 2, \ldots, s \\ &= \widehat{a}_n + \widehat{b}_n \ell + \hat{S}_{n+\ell-2s}, \text{ for } \ell = s+1, \ldots, 2s \end{aligned}$$

and so forth

Winters' additive seasonal model: updating

$$\begin{aligned}
\widehat{a}_{n+1} &= \alpha(x_{n+1} - \hat{S}_{n+1-s}) + (1-\alpha)(\widehat{a}_n + \widehat{b}_n) \\
\widehat{b}_{n+1} &= \beta(\widehat{a}_{n+1} - \widehat{a}_n) + (1-\beta)\widehat{b}_n \\
\hat{S}_{n+1} &= \gamma(x_{n+1} - \widehat{a}_{n+1}) + (1-\gamma)\hat{S}_{n+1-s}
\end{aligned}$$

α, β, and γ are "tuning parameters" that we need to adjust

Why we need multiplicative seasonal models

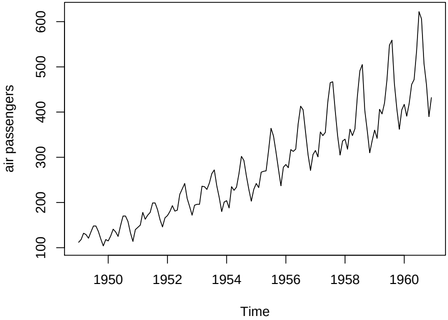

Notice the multiplicative behavior

- the seasonal fluctuations are larger where the trend is larger

Holt-Winters multiplicative seasonal method

$$
\begin{aligned}
\widehat{x}_n(\ell) &= (a_n + \widehat{b}_n \ell)\hat{S}_{n+\ell-s}, \text{ for } \ell = 1, 2, \ldots, s \\
&= (a_n + \widehat{b}_n \ell)\hat{S}_{n+\ell-2s}, \text{ for } \ell = s+1, \ldots, 2s
\end{aligned}
$$

and so forth

Winters' multiplicative seasonal model: updating

$$
\begin{aligned}
\widehat{a}_{n+1} &= \alpha \frac{x_{n+1}}{\hat{S}_{n+1-s}} + (1-\alpha)(\widehat{a}_n + \widehat{b}_n) \\
\widehat{b}_{n+1} &= \beta(\widehat{a}_{n+1} - \widehat{a}_n) + (1-\beta)\widehat{b}_n \\
\hat{S}_{n+1} &= \gamma \frac{x_{n+1}}{\widehat{a}_{n+1}} + (1-\gamma)\hat{S}_{n+1-s}
\end{aligned}
$$

Demo: Exponential smoothing in Python

`https://github.com/madeleineudell/orie3120-sp2020/blob/master/demos/forecasting.ipynb`

Applications:

- Lake Huron
- US population
- CO2
- Airline passengers
- Sales

Outline

Residuals

Selecting the tuning parameters

Forecasting using regression

Residuals

For given values of α, β, and γ:

- $\widehat{a}_n, \widehat{b}_n, \hat{S}_n, \ldots, \hat{S}_{n-s}$ are the level, slope, and seasonalities at time n
- $\widehat{x}_{n+1} = x_n(1) = \widehat{a}_n + \widehat{b}_n + \hat{S}_{n+1-s}$ is the one-step ahead forecast at time n
- $\widehat{\epsilon}_{n+1} = x_{n+1} - \widehat{x}_{n+1}$ is the residual or one-step ahead forecast error

Choosing α, β, and γ

α, β, and γ are called "tuning parameters"

Suppose we have data $x_1, \ldots, x_N$:

- the usual way to select α, β, and γ is to minimize

$$\text{SS}(\alpha, \beta, \gamma) = \sum_{n=N_1+1}^{N} \widehat{\epsilon}_n^2$$

  where the first N_1 residuals are discarded to let the forecasting method "burn-in"
- this technique is used by statsmodels, unless the user specifies the parameters explicitly in the `fit` call

Comparing forecasting methods and diagnosing problem

- Two or more forecasting methods can be compared using

$$\min_{\alpha,\beta,\gamma} \mathsf{SS}(\alpha, \beta, \gamma)$$

- If a forecasting method is working well, then the residuals should not exhibit autocorrelation

Air passengers: additive seasonal method

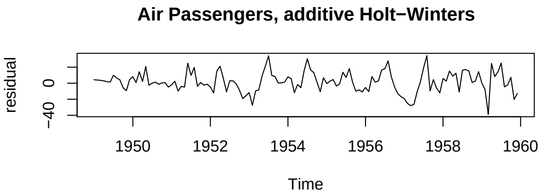

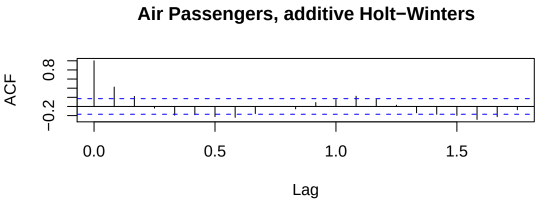

Notice some autocorrelation at small lags.

Air passengers: multiplicative seasonal method

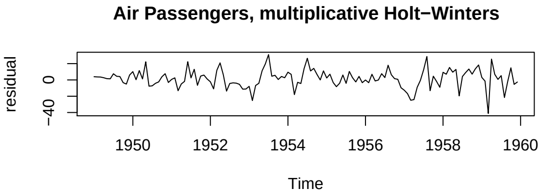

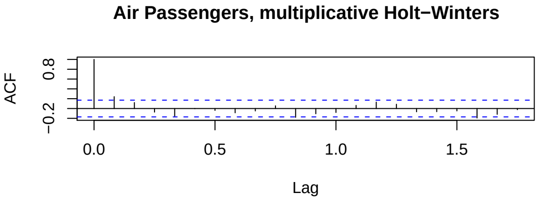

Less autocorrelation than with additive model (good).

Outline

Forecasting using regression

In some situations, regression can be used for forecasting

- in the following example, regression will be used to forecast Stove Top product 285280
  - this is the product that we forecast earlier with Holt-Winters
- the regression model will have seasonal effects but not trend
  - the seasonal effects will be introduced by using `month` as a factor
- regression uses all the data to estimate the level and the seasonal effects
  - so there is no discounting of the past
  - this helps us deal with the small amount of data

Stove Top product 285280 forecasts using regression

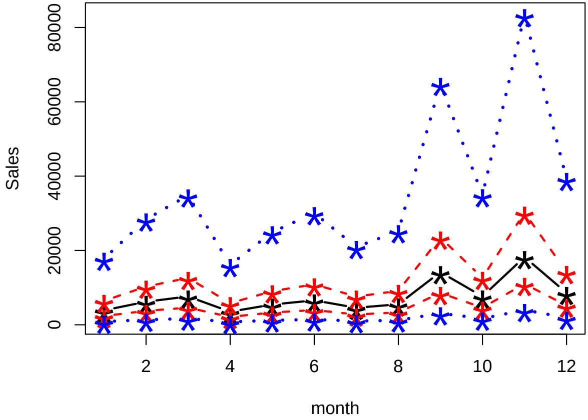

black = predictions, red = 50% pred. int., blue = 95% pred. int.

Holt-Winters product 285280 forecasts: log transformed, zoom in

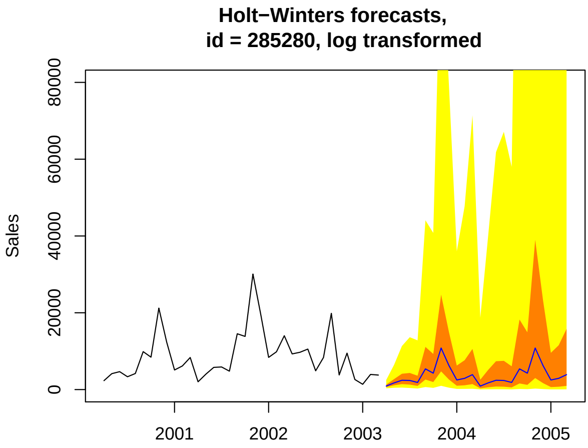

For comparison, here again are the forecasts from Holt-Winters

Why is forecasting so difficult with this product?

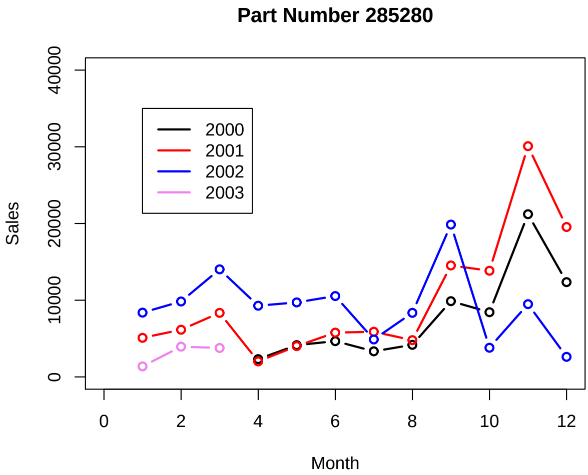

Sales patterns vary across years.