

STAT 436 / 536 - Lecture 11:

Forecasting and Prediction

- Model selection or model choice focuses on decisions about which covariates and/or model structure should be used.
- **Q:** What are common model selection procedures that you have implemented?
- For this class we will focus on a prospective, predictive framework where the goal is to predict the next k observations given some observed data.

Decision Theory: Risk and Loss Functions

- A loss function defines a structure for evaluating predictions where a penalty is assigned to a prediction based on the distance from the true value.
- Formally, a loss function is defined as a penalty for a prediction, or decision, a and a true value θ .
- Many predictive modeling scenarios, such as those on Kaggle, have a defined loss function.
- For continuous predictions, two common loss functions are
- More exotic loss functions can be devised that account for uncertainty in the prediction or penalize over or under estimates in a more severe manner.
- In a predictive setting, the loss is defined for a single prediction, typically
- A model, or an estimator, can be evaluated using a risk function. A risk function is defined as $R(\theta, \delta(y)) = E_{\theta} L(\theta, \delta(y))$, where $\delta(y)$ is an model, or estimator, of θ .
- The interpretation is that the risk function is

- A common use of the a risk function is known as the mean squared error (MSE) loss. MSE is defined as $MSE(\theta) =$
- Furthermore, MSE can be decomposed as $MSE(\theta) =$
- In practice, risk is evaluated using a hold-out sample or a cross-validation setting.

Holdout Samples and Cross-Validation

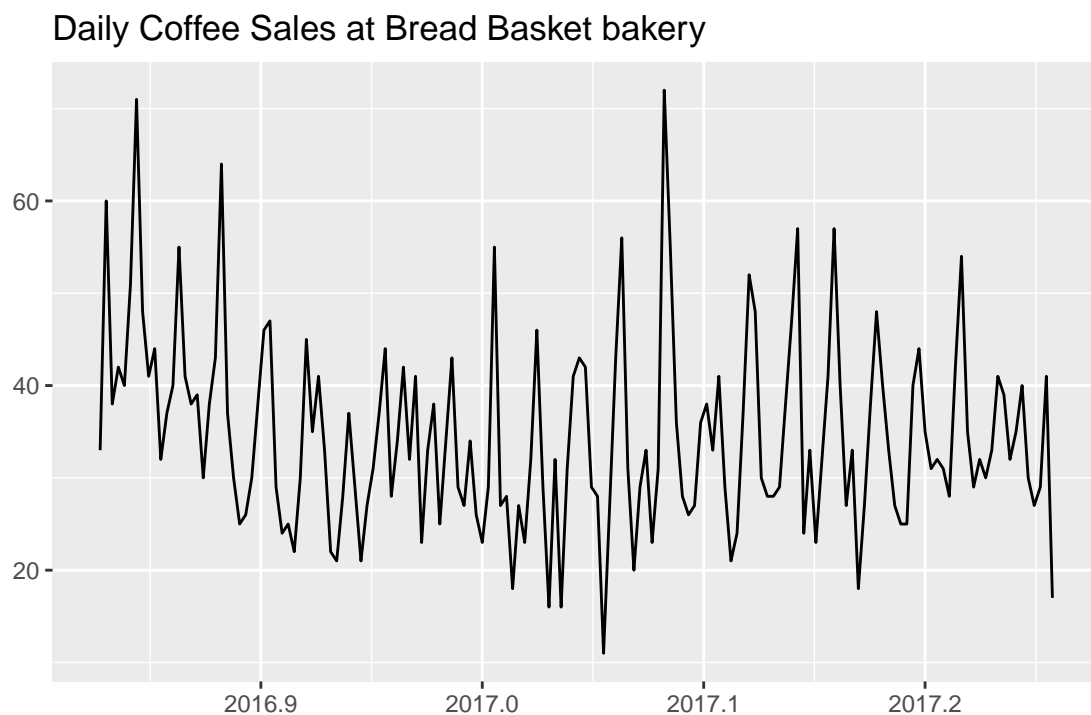
- In many predictive modeling scenarios, test and training sets are constructed. Generally, 70 percent of the data is used for training (learning the model) and the evaluated on the remaining thirty percent (the test set).
- **Q:** why does a test set need to be constructed rather than just testing the predictive ability on the training set?
- Cross validation divides the data into chunks, often called folds, and then one fold is set aside for testing and the others are used to train the model. In an iterative fashion each fold, and each data point inside the fold, are used for testing once.
- **Q:** would cross-validation, test-training set, or some other approach be more more appropriate for testing in a time series framework?

- One strategy would be to use data

- Sketch out pseudocode to implement this procedure.

- Consider a dataset on bakery sales from a bakery in Edinburgh.

```
bakery_sales <- read_csv('http://math.montana.edu/ahoegh/teaching/timeseries/data/BreadBasket.csv')
bakery_sales %>% filter(Item == 'Coffee') %>% group_by(Date) %>% tally() %>% select(n) %>%
  ts(frequency = 365, start=(c(2016,303))) %>% autoplot() +
  ggtitle('Daily Coffee Sales at Bread Basket bakery')
```



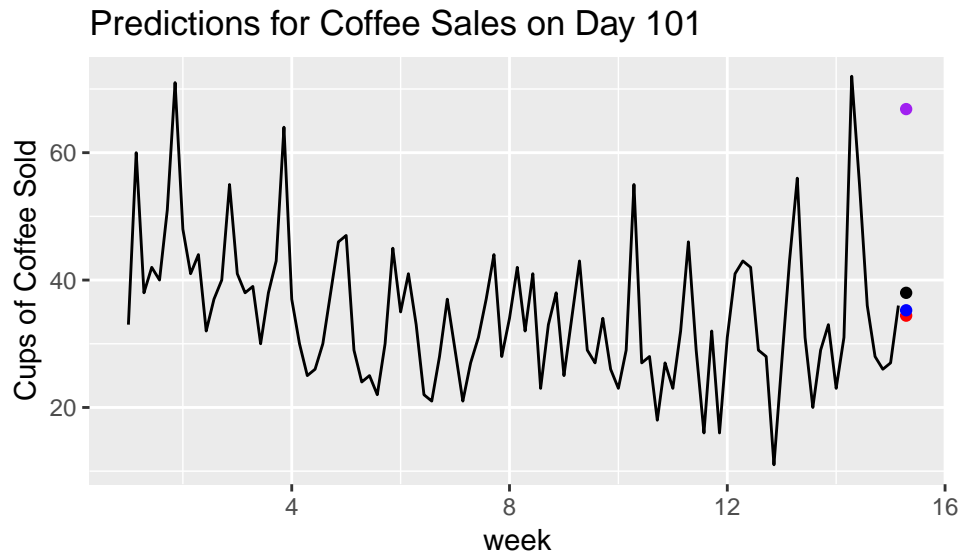


Figure 1: Black is the observed point, purple is 3-parameter HW, blue is 2-parameter HW, red is 1-parameter HW

- **Q:** Do you anticipate seasonality in coffee sales? If so, on what scale and set up a model to evaluate this.

- What would be a reasonable loss function in this situation? Lets just focus on the point estimate for now.

- There are 158 days in the time series dataset. Lets make our first prediction on the 101st time point.

```
num.days <- nrow(coffee.df)
coffee.df <- coffee.df %>% mutate(day.numb = 1:num.days)
tmp <- coffee.df %>% filter(day.numb < 101) %>% select(coffee) %>% ts(frequency = 7)
hw.pred1 <- HoltWinters(tmp, gamma=FALSE, beta=FALSE)
hw.pred2 <- HoltWinters(tmp, gamma=FALSE)
hw.pred3 <- HoltWinters(tmp)
autoplot(tmp) + annotate('point',x=15 +2/7, y=predict(hw.pred1) %>% as.numeric, color='red') +
  annotate('point',x=15 +2/7, y=predict(hw.pred2) %>% as.numeric, color='blue') +
  annotate('point',x=15 +2/7, y=predict(hw.pred3) %>% as.numeric, color='purple') +
  annotate('point',x=15 +2/7, y=coffee.df[101,'coffee'], color='black') +
  ggtitle('Predictions for Coffee Sales on Day 101') + xlab('week') + ylab('Cups of Coffee Sold')
```

- We can evaluate each of the predictions from the three models

```
abs1 <- abs.loss(predict(hw.pred1) %>% as.numeric,coffee.df[101,'coffee']); abs1
```

```
## [1] 3.584554
```

```
abs2 <- abs.loss(predict(hw.pred2) %>% as.numeric,coffee.df[101,'coffee']); abs2
```

```
## [1] 2.745026
```

```
abs3 <- abs.loss(predict(hw.pred3) %>% as.numeric,coffee.df[101,'coffee']); abs3
```

```
## [1] 28.8451
```

- **Q:** Is this enough information to make a decision on the models?

- **Q:** Why do you think the purple prediction is so large?

- Now modify the code to make predictions at the remaining time points. Include a visualization of the three predictions and a comparison of the absolute loss for the three models.