Lecture 15: Time Series Modeling

Steven Skiena

Department of Computer Science
State University of New York
Stony Brook, NY 11794–4400

http://www.cs.sunysb.edu/∼skiena

Modeling and Forecasting

We seek to construct a *model* of a function so as to *forecast* future values of it.
Certain models are proposed on theoretical grounds from first principles.
If stock prices are defined by an unbiased random walk, the best forecast is the last observation, $F_n = S_{n-1}$.
Other models are constructed by fitting observed time series data.

Statisticians vs. Computer Scientists

There is tension between the way different disciplines view the world of modeling/forecasting.

Statisticians are concerned with the limits of what can rigorously be inferred from data.

Computer Scientists are concerned with building the best model they can, regardless of whether it can be justified.

We will first explore standard statistical time series models before introducing data mining/pattern recognition methods.

A Model Building Strategy

- Model selection / specification identifies the right class of models from observation or a priori knowledge of the time series
- Model fitting finds the best possible estimates for the parameters underlying the model.
- Model diagnosis / evaluation assesses how well the model fits the original data and whether it can be improved.

The *principle of parsimony* states a model should use the smallest number of parameters to represent it adequately. Anything can be fit by an nth degree polynomial.

Estimating Volatility

Volatility is a measure of the uncertainty of prices in a given market.

The deviation from expectation in a series of observations is typically measured by its *standard deviation* σ.

The more points we observe, the more accurate our estimation of σ, *if* it isn't changing.

If it is changing, the more recent observations should receive higher weight.

EWMA models for Volatility Prediction

Exponential moving average models are often used for volatility prediction,

$$V_n = \lambda V_{n-1} + (1-\lambda) V'_{n-1}$$

where V is the predicting volatility and V' the observed volatility.

Volatility is fairly stable, as shown by the fact that the last 30 to 90 days still has predictive power. Thus the exponential decay must be small.

The RiskMetrics model uses $\lambda \approx 0.944$ for volatility estimation.

Estimating the Parameter

Such exponenetial weighting can be used to smooth a series of observations, once one has selected the parameter.

The least squares method / regression can be generally be used to find the coefficients best defining a linear function of history, as will be discussed with respect to autoregressive functions.

However, the "coefficients" here are not independent, λ^i.

Binary search / trial-and-error (for $MA(1)$ models) against some gold standard or more sophisticated numerical methods are needed.

Volatility Index

The Chicago Board Options Exchange Volatility Index (VIX) is an important measure of the market's expectation of volatility over the next 30 day period.

It is computed as a weighted blend of prices for options on a given stock index (primarily the S&P 500 index).

Recall that all parameters of the Black-Scholes formula are directly observable *except* volatility. The *implied volatility* can be computed given the other parameters, including the current options price.

The VIX is quoted in percentage points, and translates to the expected movement in the index over the next 30-day period, on an annualized basis.

VIX Securities

VIX-based derivatives include actively traded future contracts and options.

The price of a contract is 1,000 times the index; call and put options are offered at a variety of strike prices.

A VIX at 15 implies an annual 15% change, or $15\%/\sqrt{12} = 4.33\%$ over the next 30 day period, meaning that there is a 68% likelihood (one standard deviation) that the absolute value of the 30-day return will be $\leq 4.33\%$.

VIX Index

The VIX hit a record of 81.17 on October 16, 2008.

Looking at the graph, one can conclude that (1) volatility shows a high degree of short-term correlation, (2) is subject to sudden shocks, and (3) reverts to a fairly consistent mean level in quiet times.

Autoregressive (AR) Models

A simple *autoregressive* model to capture a significant lag-1 autocorrelation is

$$r_t = \phi_0 + \phi_1 r_{t-1} + a_t$$

where $\{a_t\}$ is assumed to be a *white noise* series with mean zero and ${\sigma_a}^2$.
The autocorrelation function of white noise should be near zero.
The *order* of the model is the number of terms of history; this is an $AR(1)$ model, which can readily be generalized into $AR(p)$ models for arbitrary p.

Building Autoregressive Models

For a given return series and desired history, the ϕ_i parameters can be found using the least squares method.

Build an $(n-p) \times (p+1)$ matrix where the lth column is the lag-$(l-1)$ return series.

If $p \approx n$, we have a complete linear system, and Gaussian elimination will set all of the parameters.

When $p \leq n$, we can find the ϕ coefficients which minimize least-square errors.

A model has likely captured enough history if the autocorrelation function of the residuals is small.

Capturing Seasonality

We seek to limit the history in an autoregressive model to minimize the number of parameters needed.

Seasonal effects can be captured using a few appropriate longer terms:

$$r_t = \phi_0 + \phi_1 r_{t-1} + \phi_2 r_{t-7} + \phi_3 r_{t-365} + a_t$$

Forecasting with Autoregressive Models

The value of the model at the next time step can be easily predicted by plugging in terms.

$$r_t = \phi_0 + \phi_1 r_{t-1}$$

The observed variance of the residual terms provides error bounds on the reliability of our forecast.
By plugging the predicted next value into the model and repeating, we can forecast indefinitely into the future but the error bounds on our predictions will deteriorate.