The Pennsylvania State University

The Graduate School

College of Earth and Mineral Sciences

# WIND SPEED PREDICTION VIA TIME SERIES MODELING

A Thesis in

Meteorology

by

Daniel L. Alexander



Submitted in Partial Fulfillment

of the Requirements

for the Degree of

Master of Science

August 2009

The thesis of Daniel L. Alexander was reviewed and approved* by the following:

George S. Young

Professor of Meteorology

Thesis Advisor

Andrew N. Kleit

Professor of Energy and Environmental Economics

Hampton N. Shirer

Associate Professor of Meteorology

Associate Head of the Department of Meteorology

Johannes Verlinde

Associate Professor of Meteorology

*Signatures on file at the graduate school

## ABSTRACT

Projected construction of nearby wind farms motivates this study of statistical forecasting of wind speed, for which accurate prediction is critically important to the fluid integration of wind power into the electricity grid and energy market. An 18-year record of hourly wind speed data from Williamsport, Pa. is used to develop a series of Autoregressive (AR) or Autoregressive Moving Average (ARMA) models. Performance assessments of these advanced persistence models allow for the quantification of baseline skill in wind speed forecasting. Further investigation reveals marked annual and diurnal patterns in the wind speed record, prompting the creation of scaled variables with mixed success. For each method, a persistent skill wall in modeled wind speed is observed, but this threshold is surpassed using Artificial Neural Networks (ANN).

Table of Contents

List of Tables……………………………..…………………………………………………..............v

List of Figures……………………………………………………………………………………vi

Acknowledgements………………………………………………………………………………vii

Chapter 1: INTRODUCTION……………………………………………………………………..1

*1.1 Statistical Modeling*……………………………………………………………..............6

*1.2 Data and Methodology*………………….…………………………………………..............7

*1.3 Wind Speed Cycles*……………….………………………………………………………9

Chapter 2: AUTOREGRESSIVE AND AUTOREGRESSIVE MOVING

AVERAGE MODELS……………………...……………………………….……….................12

*2.1 Introduction*…..…………………………………………………………………....……12

*2.2 Results: Skill of AR and ARMA models as measured by $r^2$ and MSE using*

*raw wind speed as a predictor*…………………………………………………...……...13

2.*3 Wind Volatility as measured by the standard deviation*...……………………….....…..17

*2.4 Scaling of wind by the mean and results of AR and ARMA models measured*

*by $r^2$ using scaled wind as predictor*……………………………………………..…..…18

*2.5 Heteroskedasticity in wind speed models*……………………………………………….21

*2.6 Scaling of wind by the standard deviation and results of AR and ARMA*

*models measured by $r^2$ and MSE using this scaled wind as predictor*………..……..27

Chapter 3: NEURAL NET RESULTS……………………………………………………………..30

Chapter 4: LITERATURE REVIEW…………………………………………………...………..32

Chapter 5: CONCLUSIONS AND FUTURE WORK……………………………..……......34

References...………………....…………………………………………………………………..36

List of Tables

Table 1: Six-hour sample of wind (m/s) with five lagged predictors……………………………..…8

Table 2: $r^2$ of AR and ARMA models using raw wind speed as a predictor……………………13

Table 3: Mean Squared Error of AR and ARMA models using raw wind speed as
a predictor……………………………………………………………………………..16

Table 4: AR and ARMA models using $W_s$ as a predictor……………………………………….20

Table 5: $r^2$ using $W_{std}$ as predictor……………………………………………………………….28

Table 6: MSE using $W_{std}$ as predictor…………………………………………………………...28

List of Figures


Figure 1: Worldwide wind capacity since 1998…………………………………………..……1

Figure 2: US wind capacity since 1999……………………………………………….......2

Figure 3: Mean wind at Williamsport by month and hour……………………………....10

Figure 4: Standard Deviation (m/s) of wind…………………………………………..17

Figure 5: Standard deviation of $W_s$……………………………………………………..19

Figure 6: Heteroskedasticity in Wind Speed Modeling……………….………….......22

Figure 7: Mean error (m/s) resulting from an ARMA(1,1) model………………………………22

Figure 8: Mean Hourly Wind Speed Changes at Williamsport………………………….23

Figure 9: Variance of the model residual by month and hour…………………………...25

Figure 10: Mean residual variance by month……………………………………………25

Figure 11: Mean residual variance by hour………………………………………….....26

Acknowledgements

I am deeply grateful to Dr. George Young, my advisor, for his infinite wisdom, patience and flexibility during the thesis process. I would also like to thank my committee, Drs. Hampton Shirer and Andrew Kleit, for their help and guidance. Moreover, I would like to thank friends and family for their support during this stressful period. Lastly, I thank Storm Exchange, Inc., for their contributions to this research.

Chapter 1: INTRODUCTION

Wind power in 2009 is the fastest growing energy resource globally, nationally, and regionally (1,3,7). Over the last decade, the total amount of harvested wind energy worldwide has grown exponentially (Figure 1). During that same period, installed wind power capacity in the United States has increased more than tenfold, with the US becoming the global leader in wind energy production in 2008 (Figure 2) (1).

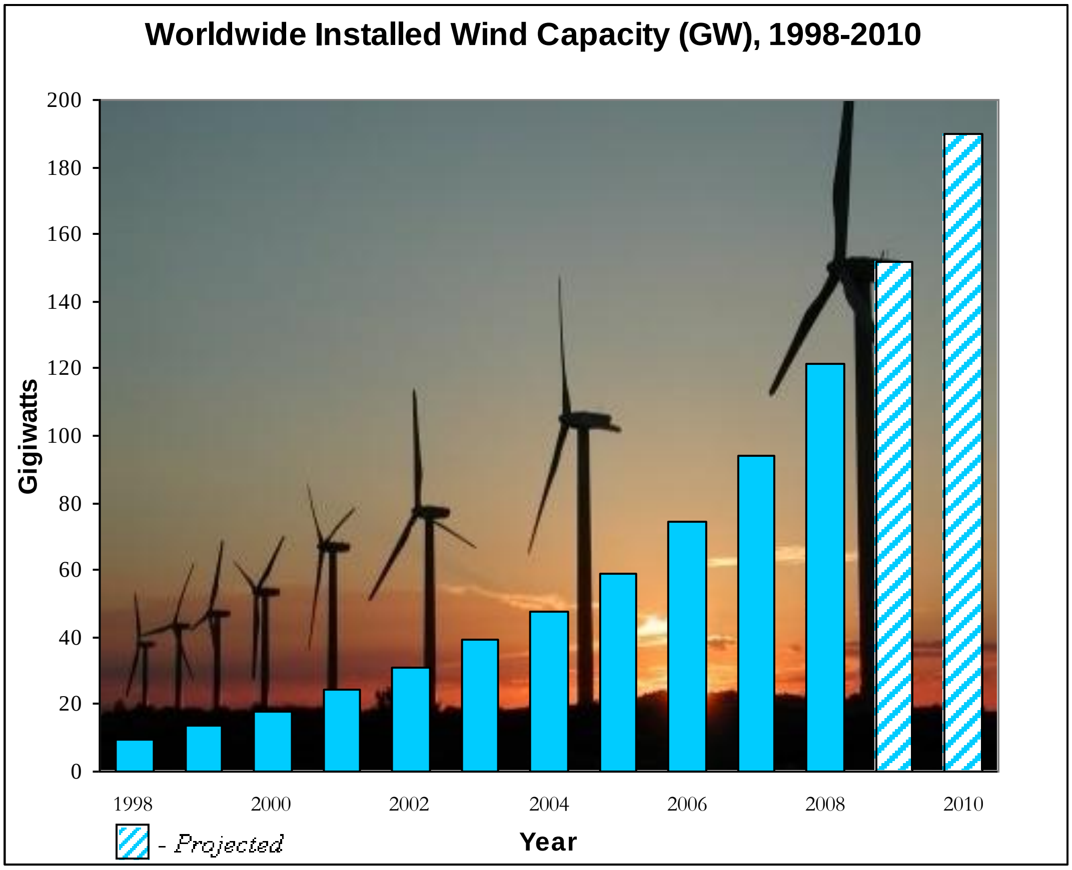


**Figure 1: Worldwide wind capacity since 1998.**

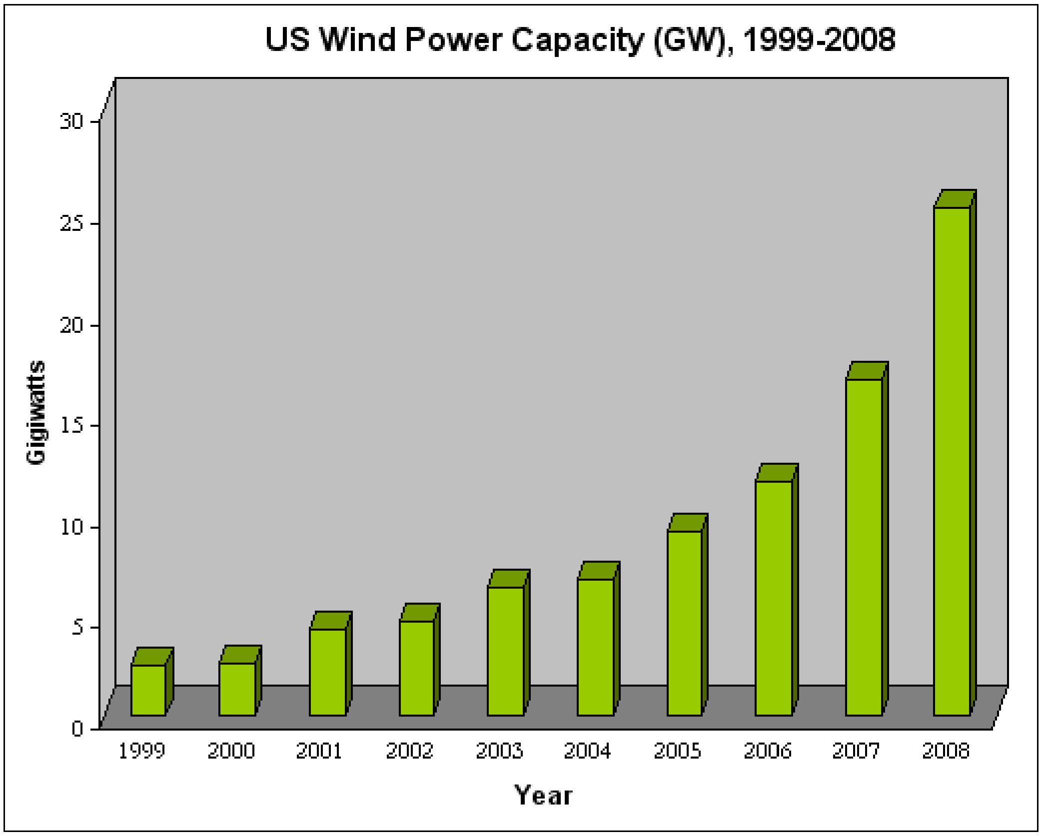


**Figure 2: US wind capacity since 1999.**

These dramatic gains in US wind capacity likely will continue well into this century; a 2006 proposal from the Department of Energy established a goal of 20% of US electricity from wind power by the year 2030, implying another tenfold increase from current levels to 300 GW per year (8). Political and economic developments at the time of this investigation, particularly the $70 billion allocated for renewable energy technologies in the 2009 US Stimulus Package, should ensure the continuation of this upward trend (9).

The site for this investigation is the Ridge and Valley Province of Pennsylvania where a similar explosion in wind power capacity has occurred in recent years. Wind power in Pa. has nearly tripled from 129 MW at the start of 2006 to 360 MW by the end of 2008 (7). In addition, several proposed wind farm development projects will boost statewide wind capacity by 234

MW in the next few years, with many of these anticipated installations slated to occur within the Ridge and Valley region (7).

Unlike conventional power plant based sources of electricity, wind is an intermittent resource in that it is highly variable across many different timescales (5). Therefore, the rapid penetration of wind power into the electricity grid presents numerous challenges for those involved in its production and distribution. Additionally, for market participants, the growing percentage of electricity supplied by wind warrants increased real-time monitoring and short-term forecasting of this fluctuating resource.

At any time on the grid, electricity consumption and generation must be in balance (4). The Transmission System Operator (TSO) is responsible for maintaining this balance, relying on a system of Independent Power Producers (IPPs) and other utilities to match demand (4). In response to load demand forecasts, typically the TSO will try to schedule production in advance (4). Any IPP exceeding or failing to meet production may be fined by the TSO, who in turn faces the potentially expensive proposition of buying (selling) on the spot market to counteract surpluses (deficits) (4). A wind farm is one member of this interdependent chain of producers. Thus, accurate wind power forecasts at grid-appropriate time scales are required to ensure smooth integration of the wind farm output into both the grid and market.

Essentially, reliable wind prediction systems lessen the risk associated with unanticipated and costly supply-side discrepancies. Speculative buyers/sellers of wind energy are keenly interested not only in the forecasts of wind power but also in the estimation of their uncertainties (14). These uncertainty values may provide clues to the volatility of energy prices – specifically as pertains to shocks created by unforeseen developments in the market and, for traders, the lucrative prospect of timing such occurrences (14). Hence, this study explores several mechanisms used in the prediction of wind power and the uncertainty related to such predictions. Lastly, an improved understanding of wind power variability on longer time scales indirectly mitigates risk for wind farm developers because accurate estimates of future wind power production enable successful wind farm operation and may lead to additional development opportunities (5).

The power produced via wind passing perpendicular to a circular area is:

$$P = C \tfrac{1}{2} \rho \pi r^2 v^3$$

where $P$ denotes total power in Watts, $\rho$ is air density, $r$ is the circle radius, and $v$ is the wind velocity (2). $C$ is 0.59, a quantity known as the Betz limit, named for the German physicist Albert Betz (2). It is the maximum theoretical fraction of energy that can be captured from the wind using a perfectly efficient turbine (2). To excellent approximation, the amount of potential wind energy varies directly by the cube of the wind speed (2). Although power is linearly proportional to air density, this quantity does not change drastically across typical surface temperature ranges; thus, wind speed variability is by far the dominant driver of variability in wind power production.

Wind prediction methods are classified broadly as either physical or statistical. Physical methods of wind power prediction utilize current atmospheric observations and the equations governing fluid motion to forecast the future state of the atmosphere. The last few decades has witnessed the emergence of many such Computational Fluid Dynamics (CFD) models (11). Typically, owing to the complexity of the equation set and to the number of calculations required to cover large areas, CFD models are run on supercomputers. The modern weather forecaster is confronted with a plethora of CFD models covering a wide spectrum of spatial and temporal scales. Two examples of medium range global models are the US-based Global Forecasting

System (GFS) and the European Center for Medium Range Forecasts (ECMWF). The GFS, which runs four times daily with output to fifteen days, and ECMWF, which runs twice daily out to ten days, are best suited for predicting the evolution of synoptic or large-scale atmospheric patterns. Yet these models are not ideal for addressing the problem of wind prediction at sufficient spatial and temporal detail to make precise forecasts of local wind speed at time intervals that are pertinent to grid operation.

Meeting this challenge requires the use of high-resolution mesoscale models such as the Weather Research and Forecasting (WRF) (13), the PSU/NCAR Mesoscale Model (MM5) (18) and the Rapid Update Cycle (RUC) (15). Relative to their global predecessors, these regional models are superior at resolving small-scale atmospheric patterns through a denser network of observations and the incorporation of topographical information onto a finer-meshed computational grid. Also, these mesoscale models are run more frequently and make projections at smaller time intervals, justifying their use in the short-term forecasting of hourly wind power. They have the added advantage of resolving the small-scale flow perturbations induced by terrain.

Although these complex physical models provide the backbone for modeling and forecasting the weather, applying their output at a given location requires calibration to correct for atmospheric processes and terrain effects that are not resolved on the model grid, a procedure known as downscaling (12). This post-processing of the forecast data using regression-based techniques creates a set of Model Output Statistics (MOS) for sites of interest within the model domain (12). These same statistical methods also are used to correct for model biases. In the arena of wind power forecasting, statistical methods are fundamental as a pathway to synthesizing complex physical model predictions into high-accuracy site-specific forecasts.

*1.1: Statistical Modeling*

A rather naïve prediction system, simple persistence – i.e., forecasting that the next hour wind will equal this hour wind – has shown fairly good accuracy for short-term wind forecasts of between 1 and 6 hours (6). This short-term success occurs largely because a time series of wind speed, like that of many meteorological variables, has the tendency to be somewhat recurrent. In fact, 24.5% of hourly wind measurements at Williamsport, Pa., our study site, were exactly the same as the previous hour, and 76.4% of wind speeds at a given hour were within 10% of that in the previous hour. Given this degree of autocorrelation (0.82), persistence alone forms a robust benchmark skill for short-term wind prediction.

Simple statistical forecast methods such as persistence have a long history in the field of weather forecasting. Such statistical methods are sometimes referred to as Classical, because their use predates the era of fluid-dynamical NWP models (6). Indeed, many modern weather and climate forecast models incorporate statistical post-processing components (6). For instance, MOS predictions from most-medium range forecast models provide bias removal, calibration, and downscaling as described above. Moreover, MOS weights the NWP forecasts towards local climatology at larger time horizons as uncertainty increases, thereby ensuring that the forecast does not include variability which is not backed by model skill (12).

The primary goal of this investigation is to quantify a baseline performance of persistence-type wind speed forecast methods using only statistical techniques, i.e. those based on the preceding time series of wind speed observations, commonly called Box-Jenkins models. Thus, the models tested do not incorporate any dynamical information about the weather pattern or physical information regarding local topography other than that implicit in the wind speed time series. We will also be testing some nonlinear time series models using artificial intelligence software in WEKA (Waikato Environment for Knowledge Analysis) (20). This software provides a diverse collection of machine learning algorithms for many statistical applications, including time series analysis (20).

Identifying baseline skill in hourly wind speed prediction serves several purposes. First, the evaluation of even simple statistical models is necessary for determining which model(s) to include in a more comprehensive, ensemble-based wind prediction system. More crucially, this baseline skill establishes a yardstick upon which more complex models can be gauged. A more advanced prediction system that consistently does not outperform these straightforward statistical models is obviously not a viable alternative. Further, many NWP forecast products only are available every 3, 6 or 12 hours and so require statistical augmentation to produce the hourly forecasts required by the wind power industry (13,15,18). Statistical time series methods based on a sequence of NWP forecasts could bridge these gaps in available guidance and thus could fulfill the demand for hourly wind speed forecasts.

### *1.2 Data and methodology*

The data used for this study were collected by the first-order weather station at the Williamsport, PA. airport and were obtained from the Pennsylvania Climate Office. A period of nearly 18 years was obtained spanning January 1991 - September 2008. Along with the date and time, hourly wind speed is the only variable used in this analysis. In particular, no information on wind direction or other meteorological variables is used.

A data matrix on which to train and test various Box-Jenkins models is created by assembling an array consisting of the time series of current wind speed along with the lagged time series of wind speeds at previous hours. Models were built with lags going back five hours, thus resulting in the following data structure, created using MATLAB:

**Wind Speed (m/s)**

| (t) | (t-1) | (t-2) | (t-3) | (t-4) | (t-5) |
|---|---|---|---|---|---|
| 3.4 | 6.9 | 5.8 | 4.7 | 4.7 | 3.4 |
| 8.1 | 3.4 | 6.9 | 5.8 | 4.7 | 4.7 |
| 9.2 | 8.1 | 3.4 | 6.9 | 5.8 | 4.7 |
| 12.8 | 9.2 | 8.1 | 3.4 | 6.9 | 5.8 |
| 9.2 | 12.8 | 9.2 | 8.1 | 3.4 | 6.9 |
| 6.9 | 9.2 | 12.8 | 9.2 | 8.1 | 3.4 |

**Table 1: Six-hour sample of wind (m/s) with five lagged predictors.**

Depending on the order of the Box-Jenkins model, between one and five of these lagged predictors are used, with current wind speed at time t always being the predictand.

Within the approximately 150,000-hour data set, there were a few hundred missing wind values at various hours. Normally, these missing times would simply be discarded prior to analysis. Because this investigation involved lagged values as well as current ones, each missing entry contaminated the next one to five hours. To resolve this issue, a procedure was developed to eliminate these missing values, ensuring that only wind speed values with corresponding lags were included in the data set. This procedure was repeated when building each of the five arrays.

### *1.3 Wind Speed Cycles*

As depicted in Figure 3, the mean wind speed at Williamsport exhibits strong diurnal and seasonal cycles. During all months, the surface wind speed peaks in the mid afternoon, in association with solar heating, the growth of the Convective Boundary Layer and the downward mixing of higher momentum air from aloft. In contrast, radiational cooling at night creates a shallow stable layer near the ground, typically leading to the decoupling of winds from the surface. This nocturnal inversion is responsible for the well-defined overnight minimum in wind speeds. On average, wind speeds in the afternoon are more than twice as large as they are overnight. In addition, the greatest relative spread occurs during meteorological summer (June, July, August), when maximum wind speeds are about three times the corresponding minimum values. The months of January, February and March show the least relative variability, with afternoon winds only 70% larger than overnight winds.

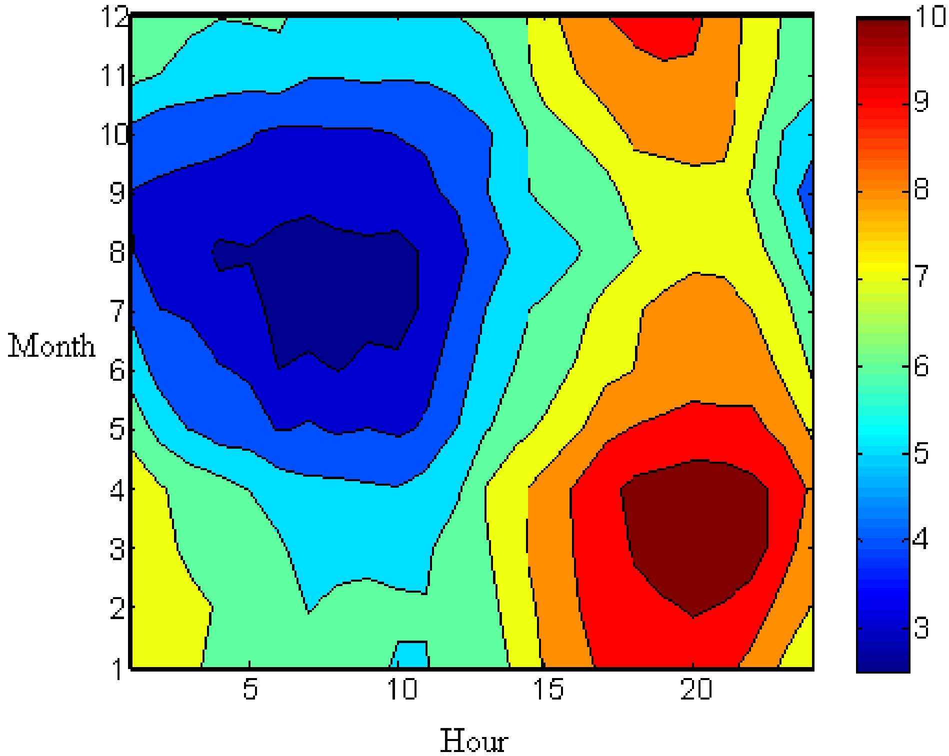


**Figure 3: Mean wind at Williamsport by month and hour.**

A distinct seasonal oscillation is also observed in Figure 3, with fastest winds occurring during the winter months and slower winds in the summer. This variation can be attributed to enhanced wintertime mid-latitude baroclinic instability and the resultant increased frequency of synoptic-scale disturbances. This seasonal cycle is not as large as the diurnal cycle, with mean wintertime winds at most hours ranging between 1.5 and 2 times the speed of the corresponding summertime winds. The magnitude of seasonal variability itself varies diurnally. Overnight, at 6 and 7 UTC, mean wintertime winds are more than twice as strong as those in summer whereas late afternoon winds are, on average, only 40% greater.

One thing to note is that the behavior of winds at a valley location such as Williamsport airport (elevation 529 ft.) may not necessarily be representative of many of the more elevated regions targeted for wind farm construction in Pennsylvania. Topographical features in the ridge and valley province can impact local wind speed greatly (16). Recalling the cubic relationship of wind speed to wind power, even relatively slight topography induced changes in wind speed will have drastic implications on harvested wind energy. Also, as an inland station, the daily and annual pattern of winds at Williamsport may not be relevant for proposed coastal wind farm sites. At these locations, sea breezes during the warm season may be the leading cause of surface winds. Unfortunately, in the ridge and valley province of Pennsylvania, as in most places, long-term observations are available primarily from airports that are typically located in valley floors.

Estimating the pattern of diurnal and seasonal oscillations of surface winds is critical not only to understanding the nature of wind but also to the way in which it is modeled.

Chapter 2: AUTOREGRESSIVE AND AUTOREGRESSIVE MOVING AVERAGE MODELS:

*2.1 Introduction*

Two types of iterative statistical models were applied to the wind speed time series data, the autoregressive model (AR) and autoregressive moving average model (ARMA). These are sometimes referred to as Box-Jenkins models (11), named for the statisticians George Box and Gwilym Jenkins who first applied these methods to various econometric time series data (11). AR and ARMA models use information about past values of the forecast variable to understand current values and make predictions of future values. The AR model can be written as follows:

$$X_t = c + \sum_{i=1}^{p} \varphi_{t-i} X_{t-i} + \varepsilon_t$$

where $\varphi_1, \ldots, \varphi_p$ are the autoregressive parameters of the model, $c$ is a constant, and $\varepsilon_t$ refers to the random error, or residual term in ordinary least squares regression. . Also, $p$ denotes the order of the autoregression, i.e. the number of lagged predictors, $\boldsymbol{X}_{t-1}, \ldots, \boldsymbol{X}_{t-p}$ included in the model. For this study, the order of the autoregressive models ranged between one and five; that is, between one and five previous hourly wind speeds were used as predictors.

ARMA models include additional information about prior values by inclusion of a moving average (MA) component that can be expressed as follows:

$$X_t = \varepsilon_t + \sum_{i=1}^{q} \theta_i \varepsilon_{t-i}$$

where $\varepsilon_t$ are the error terms, $\theta_1, \ldots, \theta_q$ are the moving average parameters, and $\varepsilon_{t-i}$ represents the residual from the last hour model prediction. As with *p* in the AR model, *q* is the order of the moving average such that the full ARMA(p,q) model takes the following notation:

$$X_t = c + \varepsilon_t + \sum_{i=1}^{p} \varphi_{t-i} X_{t-i} + \sum_{i=1}^{q} \theta_i \varepsilon_{t-i}$$

*2.1 Results: Skill of AR and ARMA models as measured by $r^2$ and MSE using raw wind speed as a predictor:*

**Autoregressive Terms**

| | | **1** | **2** | **3** | **4** | **5** |
|---|---|---|---|---|---|---|
| **Moving Average Terms** | **0** | 0.6315 | 0.6461 | 0.6469 | 0.6465 | 0.6463 |
| | **1** | 0.6459 | 0.6468 | 0.6465 | 0.6464 | 0.6464 |
| | **2** | 0.6460 | 0.6466 | 0.6463 | 0.6461 | 0.6461 |
| | **3** | 0.6460 | 0.6462 | 0.6462 | 0.6461 | 0.6461 |
| | **4** | 0.6457 | 0.6460 | 0.6460 | 0.6459 | 0.6459 |
| | **5** | 0.6454 | 0.6457 | 0.6457 | 0.6457 | 0.6456 |

**Pure Persistence: $r^2$ = 0.6315**

**Table 2: $r^2$ of AR and ARMA models using raw wind speed as a predictor.**

Displayed in Table 2 are the performances of 30 distinct AR and ARMA models using raw wind speed as a predictor. The columns denote the number p of autoregressive terms and the rows indicate the number q of moving average terms. Models occupying the top row of the chart contain no moving average terms and are thus AR models. Skill is quantified by the coefficient of determination, which is the square of the correlation coefficient between the predicted wind speed of the model and the actual wind speed. This is sometimes referred to as the $r^2$ value, and this descriptor is used hereafter. It should be noted that this measure does not adjust for the number of degrees of freedom, which is known as the adjusted $r^2$. Using adjusted $r^2$, however, would not change the results much, since our data set had so many values.

Several significant features are depicted by the table. First, the skill obtained from the AR(1) model matches the persistence method. Although perhaps not obvious from the AR notation defined above, the AR(1) model uses identical information as the persistence method, namely the last hour wind data. However, the two prediction schemes contain subtle differences in implementation. Persistence simply takes the last hour wind speed as the next-hour prediction. While the AR(1) model also employs only last hour wind in its prediction, it scales that input using the autocorrelation coefficient between past and current values. Similarly, as the least skilled of the models, these provide a benchmark for the others.

It is also striking to note the similarity between the results of the other 29 models. Each model skill falls between 0.6454 and 0.6469, implying that beyond an incremental jump in skill from persistence, the inclusion of additional lagged predictors does not appreciably alter the model performance. Moreover, the incorporation of moving average terms does not necessarily improve model performance; beyond two such terms, there was no improvement in the model, and having more than three MA components had a detrimental effect across the model suite. The best model is the AR(3), followed closely by the ARMA(2,1), ARMA(2,2), AR(4) and ARMA(3,1). For simple statistical modeling purposes, having between 3 or 4 components is preferred; more complex representations of the data bring about declines in model accuracy.

The slight declines in skill with increasing model complexity may be an artifact of our iterative model fitting method. More complex models are harder to fit and so may not have been tuned to convergence.

Two important conclusions can be gleaned from this analysis. The first is that all Box-Jenkins models universally beat simple persistence although the improvements are fairly unremarkable. Second, there appears to be a plateau in skill with fairly low-order autoregression and that greater model complexity is not necessarily desirable. These findings, particularly the marginal performance gains compared with persistence, align well with prior study results. Miranda and Dunn (2006) used a Bayesian approach to build an AR(6) model, using only wind speed time series to forecast the next hour wind (21). In their study, the difference between the prediction errors for the AR model and persistence model is around 1%, which corresponds with the improvements of this study (21). As future work, they asserted that ARMA models would improve the model performance substantially, a claim that is refuted by our results (21).

| | | Autoregressive Terms | | | | |
|---|---|---|---|---|---|---|
| | | **1** | **2** | **3** | **4** | **5** |
| | **0** | 11.068 | 10.369 | 10.279 | 10.269 | 10.266 |
| **Moving** | **1** | 10.280 | 10.278 | 10.273 | 10.269 | 10.266 |
| **Average** | **2** | 10.280 | 10.266 | 10.265 | 10.265 | 10.265 |
| **Terms** | **3** | 10.282 | 10.265 | 10.265 | 10.264 | 10.262 |
| | **4** | 10.278 | 10.262 | 10.262 | 10.262 | 10.252 |
| | **5** | 10.271 | 10.255 | 10.255 | 10.255 | 10.252 |

**Table 3: Mean Squared Error of AR and ARMA models using raw wind speed as a predictor.**

Another measure of model performance is mean squared error (MSE), displayed in Table 3. A decrease in MSE is indicative of a better model, which is opposite the interpretation of decreasing $r^2$. Using this alternative metric, we observe several patterns that were present in the $r^2$ table as well as one subtle difference. First, the same improvement over simple persistence occurs when an additional AR or MA term is added. Moreover, there is a distinct plateau in skill at higher ordered models. Lastly, unlike the $r^2$ table, adding complexity to any model results in a reduction, however slight, in MSE.

*2.3 Wind Volatility as measured by the Standard Deviation*

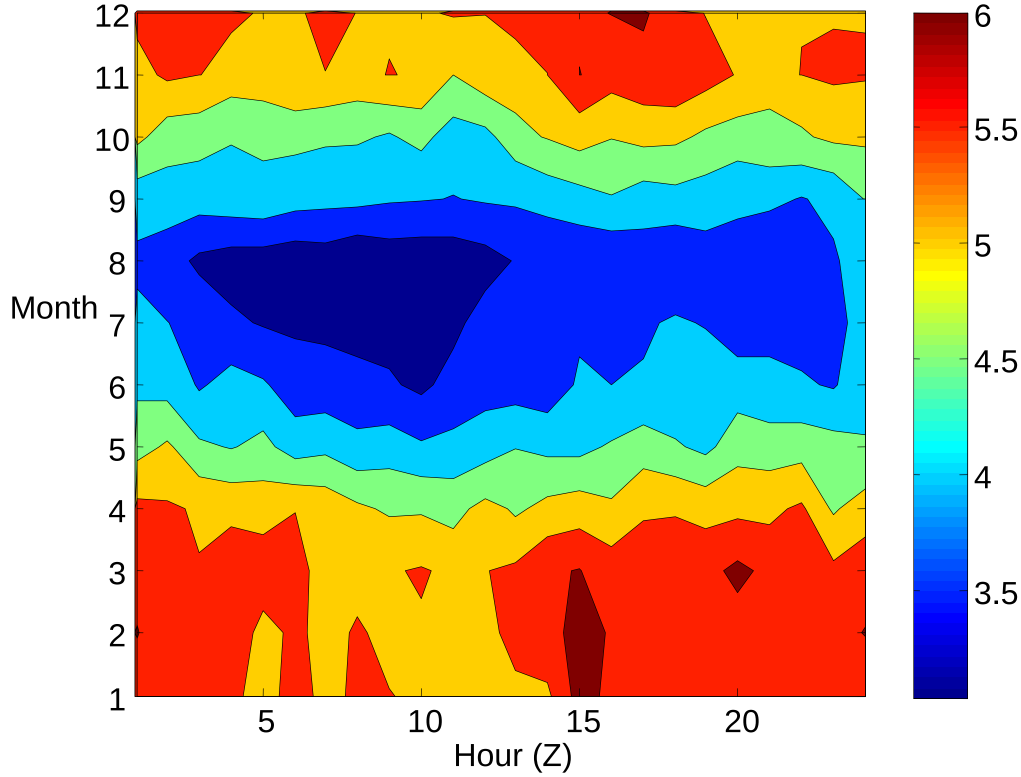


**Figure 4: Standard Deviation (m/s) of wind.**

Before proceeding to other modeling techniques, we must examine some other critical features of wind. To assess the volatility of wind, we sorted the wind data by month and hour to identify annual and diurnal patterns in variability as defined by the standard deviation, σ:

$$\sigma = \sqrt{\frac{1}{N}\sum_{i=1}^{N}(x_i - \mu)^2},$$

where $\sigma$ is the square root of the squared residual. Figure 4 is a plot of the standard deviation as computed on this monthly and hourly basis. The most prominent aspect is the greater variability of the wind speed during the winter. This seasonal pattern in $\sigma$ indicates that the expected spread in wind speeds is large in winter, with summertime wind speeds falling in relatively narrow ranges. No definite diurnal pattern is observed, indicating that although the mean wind speed fluctuates from hour to hour in a given month, the standard deviation, a measure of variability from the mean, does not change appreciably throughout the day.

*2.4 Scaling of wind by the mean and results of AR and ARMA models measured by $r^2$ using scaled wind as predictor.*

In light of this seasonally dependent variance in wind speed, we attempted to scale the data to reduce or eliminate this characteristic. One method is to use a new wind speed metric:

$$W_s = \frac{wind\,speed}{mean\,wind\,speed}$$

Here, the dimensionless scaled wind, $W_s$, is the actual wind divided by the mean wind speed for that month and hour.

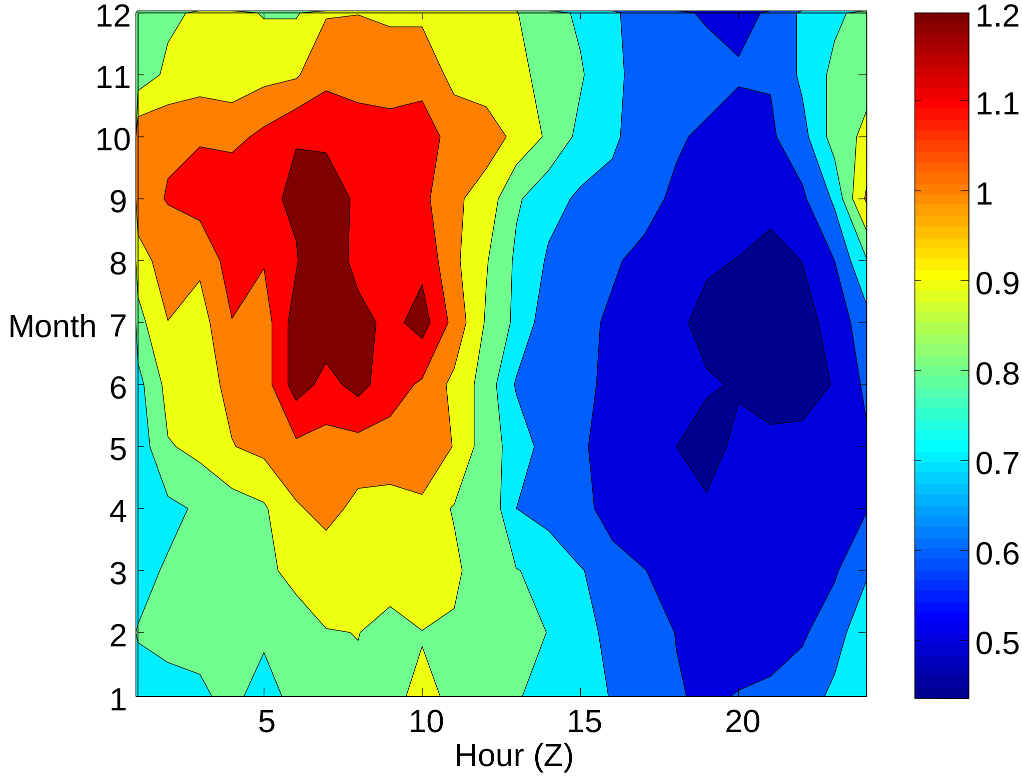


**Figure 5: Standard deviation of $W_s$.**

Scaling the wind by its expected mean proved ineffective with regards to reducing variability as depicted in figure 5. Although we have diminished much of the differential month-to-month variability, we have introduced a diurnal oscillation in the standard deviation of the scaled wind. In fact, the pattern in the figure – highest values in summer mornings and low values in the afternoon – is nearly opposite to the pattern in mean wind speed. This occurs because the scaling includes the mean and the resultant σ plot thus reflects the mean wind pattern. Such behavior occurs because the mean wind speed is quite low during the early mornings of summer, and so any disturbance may generate surface winds that are many times the mean. When the average wind speed is 10 m/s during the winter afternoon it is comparatively much rarer for the wind to be double or triple its mean value.

We repeated the building of AR and ARMA models discussed above for different numbers of terms, using the scaled wind as the predictor, with the following results:

| Moving Average Terms | Autoregressive Terms: 1 | 2 | 3 | 4 | 5 |
|---|---|---|---|---|---|
| 0 | 0.5634 | 0.5726 | 0.5724 | 0.5722 | 0.5719 |
| 1 | 0.5712 | 0.5737 | 0.5728 | 0.5720 | 0.5716 |
| 2 | 0.5737 | 0.5737 | 0.5728 | 0.5720 | 0.5715 |
| 3 | 0.5737 | 0.5736 | 0.5728 | 0.5720 | 0.5715 |
| 4 | 0.5734 | 0.5733 | 0.5729 | 0.5721 | 0.5715 |
| 5 | 0.5733 | 0.5732 | 0.5728 | 0.5722 | 0.5717 |

**Table 4: AR and ARMA models using Ws as a predictor.**

Our scaled predictor did not improve the performance of the linear least squares autoregressive model suite as measured by $r^2$. Likewise, several features of the raw wind model skill matrix are still observed. Most notably, after an initial jump in skill from beyond that of persistence, little further improvement was observed. Once again, more complex representations were not associated with improvements in skill. Perhaps the most important observation is the existence of a distinct and pervasive plateau in skill using AR and ARMA models. This plateau in skill shown here and in previous tables using different predictors may be attributed to the iterative method used to fit our linear model. The use of many predictors poses major challenges for an iterative process, as the actual coefficient values oscillates with each iteration. There is no guarantee that the weights applied to each lagged predictor are truly representative of the degree to which each impacts the actual value. Beyond a certain number of predictors, the model is essentially adjusting the fractional importance of each predictor which may explain the leveling in skill.

### *2.5 Heteroskedasticity in wind speed prediction*

An alternative approach to the issue of seasonal volatility in wind speed prediction is to employ an Autoregressive Conditional Heteroskedasticity (ARCH) model (17). The motivation for ARCH modeling is to adjust for the heteroskedasticity, a measure of uneven predictability, provided by time series models. Heteroskedasticity refers to the characteristic of changing variances in the residuals that describe the expected spread of the error terms. ARCH models have their roots in econometric analysis of financial markets, in which quiescent, relatively predictable periods are sometimes interrupted by wild fluctuations (17). The central premise of ARCH models is that forecast errors are likely to be conditionally dependent on prior errors (17). These models can be applied to many atmospheric variables such as wind speed that are inherently volatile. Figure 4 is an example of heteroskedasticity in wind speed forecasting – it is a 40-hour period of wind speed at Williamsport and the corresponding ARMA(1,2) model errors. During that span, dramatic swings exist in both the wind speed and the model forecast errors. Furthermore, sudden spikes in wind speed are associated with large underestimates by the model with sizable overestimates occurring once the wind speed returns to more typical levels.

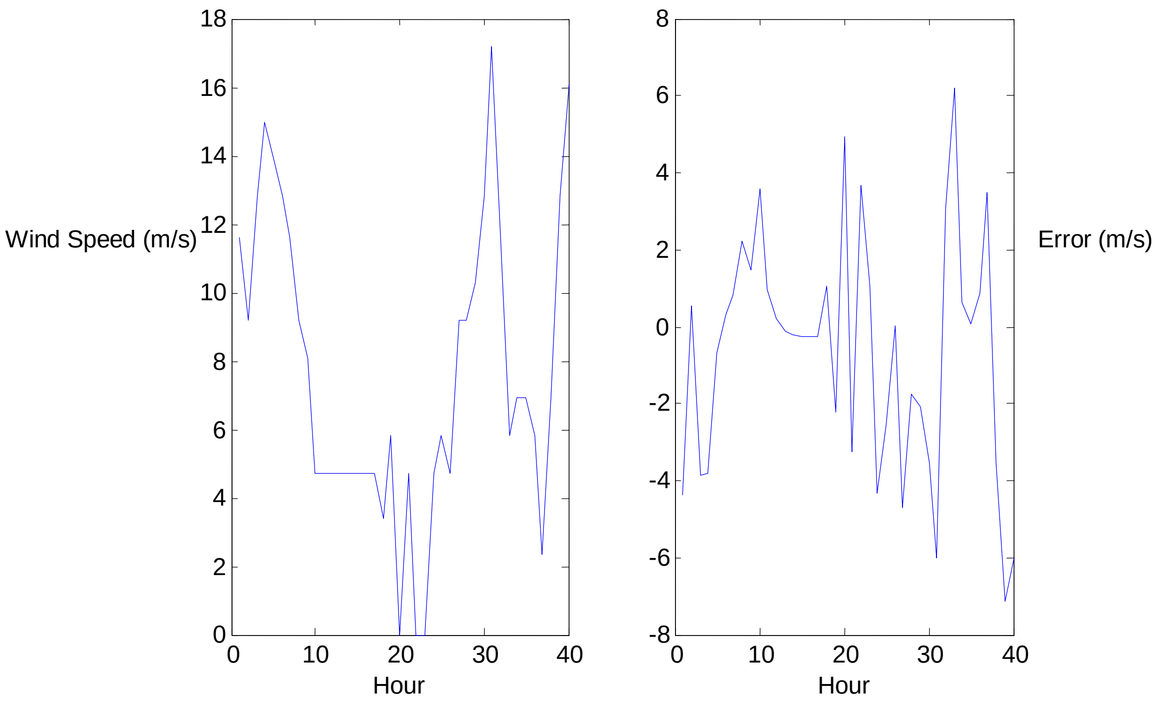


**Figure 6: Heterokedasticty in Wind Speed Modeling.**

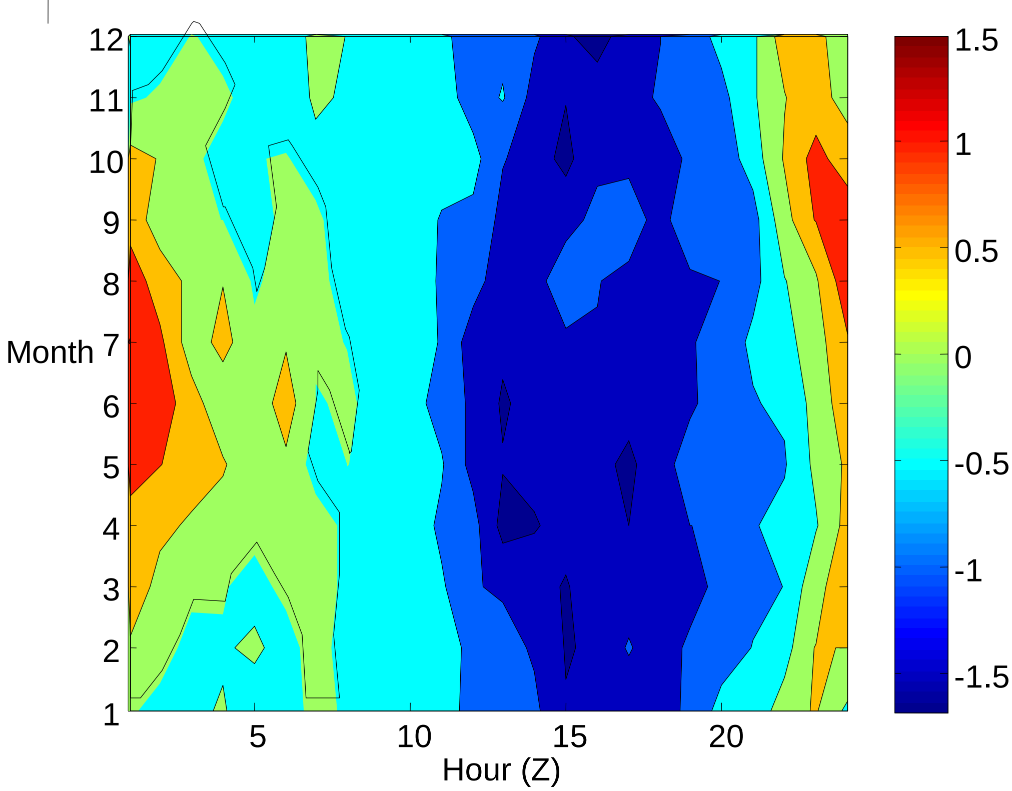


**Figure 7: Mean error (m/s) resulting from an ARMA(1,1) model**

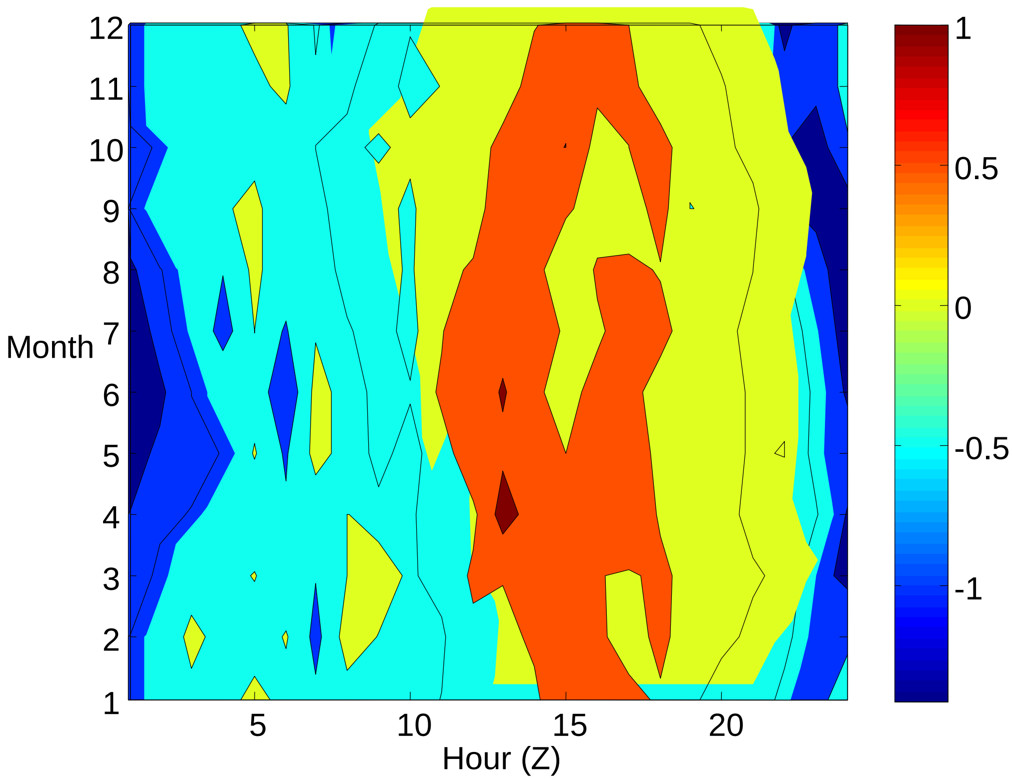


**Figure 8: Mean Hourly Wind Speed Changes at Williamsport.**

The above mentioned tendency for increases(decreases) in wind speed to be associated with model forecasts of wind speed below(above) the actual wind speed is a logical consequence of using any low-ordered AR model, because it draws the majority of its information from the prior hour wind speed. Figure 7 displays a contour plot of mean residuals for the ARMA(1,1) model on a monthly and hourly basis and Figure 8 offers a climatological explanation for such model behavior. As discussed when analyzing the mean wind, there is a strong diurnal cycle caused by the daytime development and growth of the convective boundary layer, and the corresponding rise in downward momentum flux to the surface leads to an increase in mean surface wind speed after sunrise. Mean hourly surface wind speed peaks in the afternoon before decreasing, sometimes dramatically, after sunset, when the boundary layer detaches from the surface. This pattern is captured in Figure 8. Also observable is the earlier onset of increasing winds in summer associated with earlier sunrise, and the corresponding later decline in wind speeds centered upon the summer solstice. The mean error mirrors this pattern, with mean model underestimates occurring during times when the hourly wind speed is increasing with overestimates occurring when the wind typically slackens.

Although providing considerable information on the forecasting process, the diurnal patterns of AR(1) wind errors do not yield direct insight into the possible heteroskedasticity, which is determined not by the residual values themselves, but their spread as measured by the variance of the residuals. Figure 9, which plots the variance of the model residual on a monthly and hourly basis, demonstrates that forecasts of wind speed are unmistakably heteroskedastic with the variance differing significantly over both monthly and hourly time intervals. In particular, there is a belt of large variances covering the wintertime and springtime months that is concentrated over the afternoon and evening hours. Moreover, the model errors are more tightly distributed in summer and early autumn, especially during the early morning hours.

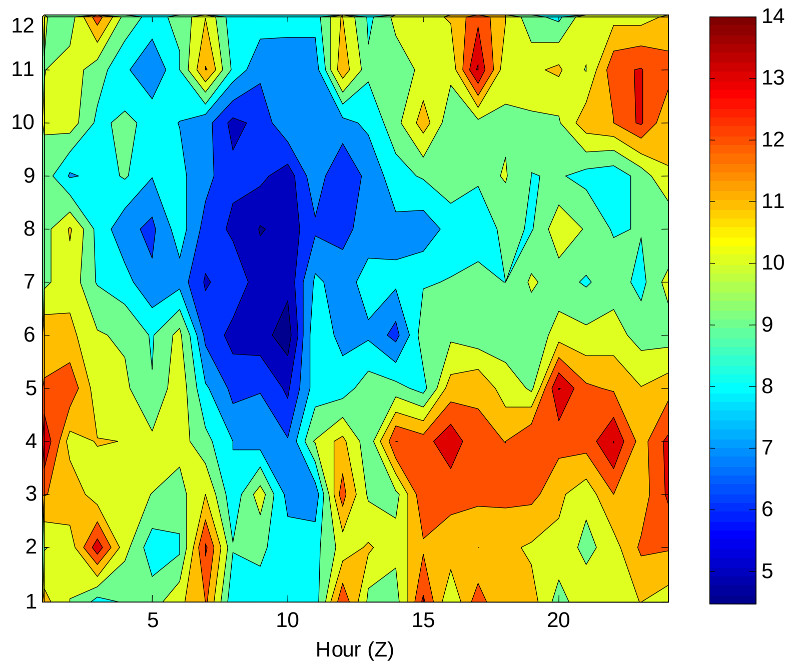


**Figure 9: Variance of the model residual by month and hour.**

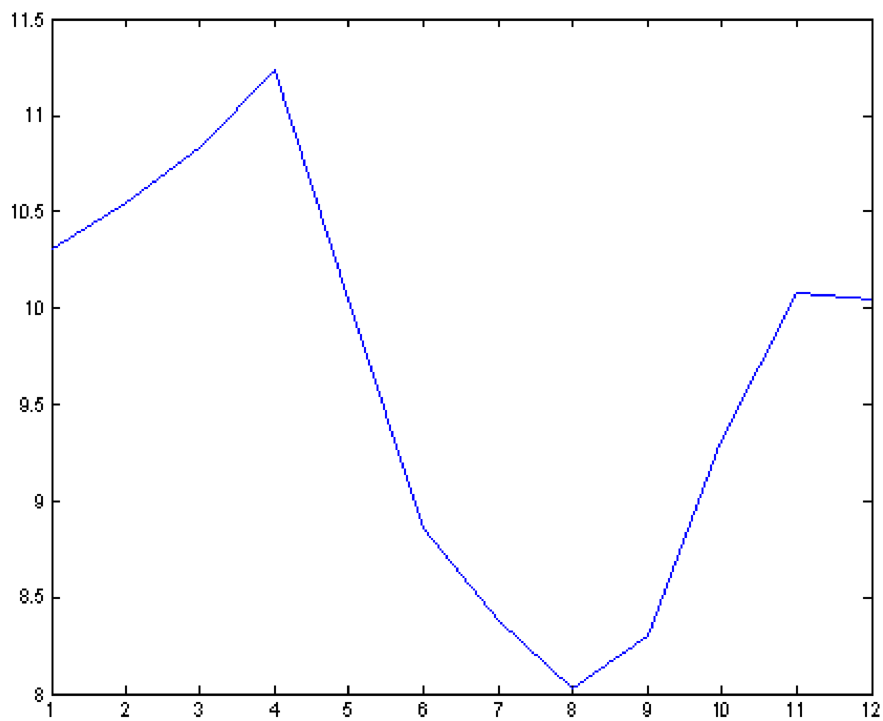

**Figure 10: Mean residual variance by month.**


**Figure 11: Mean residual variance by hour.**

The plots above more explicitly illustrate the annual and diurnal cycles in residual variance. Figure 10, which is the mean variance of residuals by month, captures the oscillation from high variance values during the months of February through April to relatively low variance between July and September. Applying the inverse relationship between predictability and variance, Fgure 10hows that is that with regards to wind speed forecasting, April is the least predictable month whereas August is the most predictable. Figure 11 also shows hourly variance changes; as expected, there is low mean variance during the predawn hours and comparatively high mean variance from late morning through the evening hours.

*2.6 Scaling of wind by the standard deviation and results of AR and ARMA models measured by $r^2$ and MSE using this scaled wind as predictor.*

In lieu of ARCH modeling, we developed a second scaled predictor, *Wstd*:

$$Wstd = \frac{wind\,speed - meanwind\,speed_{(month,\,hour)}}{\sigma_{(month,\,hour)}}$$

where the numerator is the difference between the actual wind speed and the mean wind speed and the denominator is the standard deviation about the mean. The subscript (month, hour) denotes that such quantities are calculated for each month and hour combination.

We decided to scale the predictor by its perturbation from the expected mean divided by its standard deviation to correct for heteroskedasticity. Because the standard deviation of the dependent variable is proportional to the variance in the model error - a quantity not known ahead of time - our hypothesis is that this scaling would wipe out the heteroskedastic signal in model errors. This preprocessing step eventually did create constant residual variance but the resultant skill scores dropped from the previous scaled predictor, as shown in Table 5.

**Autoregressive Terms**

| Moving Average Terms | 1 | 2 | 3 | 4 | 5 |
|---|---|---|---|---|---|
| 0 | 0.4706 | 0.4894 | 0.4927 | 0.4933 | 0.4935 |
| 1 | 0.4930 | 0.4937 | 0.4937 | 0.4938 | 0.4941 |
| 2 | 0.4936 | 0.4937 | 0.4937 | 0.4938 | 0.4941 |
| 3 | 0.4936 | 0.4936 | 0.4937 | 0.4938 | 0.4939 |
| 4 | 0.4936 | 0.4936 | 0.4936 | 0.4935 | 0.4939 |
| 5 | 0.4937 | 0.4937 | 0.4937 | 0.4938 | 0.4937 |

**Table 5: $r^2$ using $W_{std}$ as predictor.**

**Autoregressive Terms**

| Moving Average Terms | 1 | 2 | 3 | 4 | 5 |
|---|---|---|---|---|---|
| 0 | 15.096 | 14.572 | 14.478 | 14.460 | 14.454 |
| 1 | 14.477 | 14.468 | 14.456 | 14.444 | 14.436 |
| 2 | 14.471 | 14.457 | 14.454 | 14.454 | 14.450 |
| 3 | 14.459 | 14.454 | 14.447 | 14.447 | 14.444 |
| 4 | 14.448 | 14.446 | 14.444 | 14.435 | 14.434 |
| 5 | 14.445 | 14.444 | 14.441 | 14.436 | 14.435 |

**Table 6: MSE using $W_{std}$ as predictor.**

The skill scores in Table 5 are incrementally worse than those produced by the mean wind scaling and align with the skill patterns of the prior two charts. This drop off in $r^2$ is not uncommon when scaling by the variance as the presence of heteroskedasticity tends to inflate the

$r^2$ using linear prediction models. In fact, the $r^2$ values and MSE values in Table 7 are a more accurate characterization of correlation in wind speed forecasting using linear methods.

## Chapter 3: NEURAL NET RESULTS:

Several studies have documented the applicability of a wide variety of artificial intelligence methods in the prediction of short-term wind speeds. These sophisticated nonlinear systems represent powerful, fascinating alternatives to traditional approaches to solving optimization problems. Modern Artificial Neural Networks (ANNs) utilizing radial basis functions, support vector machines, back propagation algorithms, and fuzzy logic or genetic algorithm based models, among others, have provided significant improvements over their autoregressive predecessors (22,23,24,25). While an extensive discussion of ANNs is beyond the scope of this study, we trained and tested a variety of of ANNs and found a few results to be rather staggering in light of the skill wall prevalent in our linear models. Specifically, we used the Weka software, importing the same wind speed and lagged predictors data set used in the AR and ARMA models (20). One key difference from our linear models is that these models used 10-fold cross validation. This procedure, also known as jackknifing, repeatedly divides the data into subsets, building the model on one subset and testing it on the other. The results of all trials are then averaged and combined into a single estimate – one purpose of cross validation is to reduce the chance of overfitting the model.

| | **Autoregressive Terms** | | | | |
|---|---|---|---|---|---|
| | 1 | 2 | 3 | 4 | 5 |
| Multi-P | 0.6333 | 0.7396 | 0.6463 | 0.7441 | 0.6450 |
| R Via D | 0.6268 | 0.7363 | 0.6514 | 0.7571 | 0.6874 |

**Table 7: $r^2$ values obtained from two ANNs.**

Two methods that yielded encouraging results are shown in Table 7. The two types of ANNs are a Multilayer Perceptron and a Regression Via Discretization scheme using a Classification Tree (19). Both effectively shattered the skill ceiling of AR/ARMA, especially for 2nd and 4th order AR processes. Although these innovative tools have received a cursory acknowledgement here, future work could investigate these methods in more detail, addressing the question of why even numbered autoregressive lags were particularly successful in the forecasting of wind speed.

# Chapter 4: LITERATURE REVIEW

Although our study did not delve into ARCH modeling, others have taken this approach to successfully counteract heteroskedasticity in wind speed and other atmospheric variables. Tol (1997) invoked a Generalized AutoRegressive Conditional Heteroskedasticity (GARCH) model to account for changes in predictability of wind speed (27). One fundamental difference compared with this study is that he uses a time series of average daily wind speed measurements and included information about wind direction (27). The overarching theme of the study, to build a model that integrated systematic changes in the volatility of wind speed, is entirely germane to our analysis, however (27). In parallel to our study, Tol found that the order of autoregressive (AR) and moving average (MA) parameters had no effect on resolving the inherent volatility present in the data (27). To address this issue, he developed a GARCH model with conditionally changing variances based on the variances of prior time periods (27). The idea behind this heteroskedastic model is that a large observed forecast error at any time is likely to be followed by another large error, a stipulation analogous to autoregression in the mean. Furthermore, the heteroskedastic models were shown to beat their homoskedastic counterparts (27). Future work could be aimed at developing an ARCH or GARCH model to untangle the problem of seasonal volatility in hourly wind speed measurements in our data set.

The issue of heteroskedasticity in wind speed modeling also has direct implications for market-based activities, notably correct pricing of wind energy, proper hedging strategy for wind energy derivatives, and opportunities for arbitrageurs to recognize and benefit from price imbalances (17). Campbell and Diebold (2005), in a time series analysis of daily average temperatures for four US cities, noted that the amplitude of the temperature residuals, or deviations from climatologic normal, varied over the course of the year with the greatest deviations taking place in winter and smallest occurring in summer (26). Although their study was restricted to temperature, the issue of seasonally dependent volatility certainly is related to the similar pattern of volatility in wind speed. As in Tol's analysis, they modeled temperature using a GARCH process, although their primary motivation was not improved modeling but the reduction of weather risk, which is intimately tied to the issue of predictability (26). The increased wintertime frequency of large deviations from temperature norms, so-called weather "surprises" or "shocks", is associated with greater risk that leads to increased hedging demand (26). Their analysis focused on better estimates of probability density forecasts at various lead times, which are more relevant in the energy market (26).

# Chapter 5: CONCLUSIONS AND FUTURE WORK

Spurred by a global and national explosion in wind energy along with the proposed development of wind farms in the Ridge-and-Valley region of Pennsylvania, we have examined an 18-year time series of wind speed at Williamsport, Pa. The primary objectives of this study are to identify key annual and diurnal patterns in wind speed and to develop statistical models for hourly forecasts of this harvestable but intermittent resource. These goals are not unrelated; rather, the characteristics of the wind arising from conditional seasonal and hourly dynamics offer clues to the preferred statistical methods involved in wind modeling as well as explanations for certain shortcomings of these statistical models. Specifically, the tendency of the wind to be recurrent in the short term allows for modest skill using Autoregressive (AR) and Autoregressive Moving Average (ARMA) models to predict the next-hour wind speed, an estimate necessary for successful wind farm operation and efficient electricity grid management. Also, the variability of wind on monthly and hourly timescales has implications for model predictability, an essential measurement in energy market applications.

In this analysis, we test many advanced persistence models and establish similar results using three distinct predictors. The resulting three groups of models all show a relatively large initial gain in skill over simple persistence as well as a leveling in skill with increased complexity. This distinctive skill wall across the AR/ARMA model suite implies an upper limit in the performance of linear models. Additionally, we document the ability of two types of nonlinear models to successfully break this skill wall. Thus, a potential direction for future research is a more detailed exploration of Artificial Neural Networks and other nonlinear models.

Future work might also incorporate high resolution Numerical Weather Prediction (NWP) output as well as detailed topographical information should vastly improve model performance. Indeed, with respect to short-term wind speed prediction, many studies have defined an optimal model as one that includes both physical and statistical constituents and additional research could be aimed at building such a model for operation within potential wind farm locations across the Ridge and Valley province (14). This model could make use of advanced statistical techniques that have been shown to improve accuracy in wind speed forecasting. Such methods include adaptive estimation of time-varying coefficient functions, Kalman filtering of dynamic atmospheric quantities, and real-time evaluation of NWP forecasts, adjusting for recent MOS errors in wind speed (28, 29, 30). These techniques, among others, have proven successful in many projects in Europe, where faster growth of wind power relative to the United States stimulated earlier development of wind speed forecasting systems.

\

References:


1) World Wind Energy Association (February 2009). "World Wind Energy Report 2008." March 18th, 2009. <http://www.wwindea.org/home/images/stories/worldwindenergyreport2008s.pdf>
2) Betz, A. (1966) Introduction to the Theory of Flow Machines. (D. G. Randall, Trans.) Oxford: Pergamon Press.
3) American Wind Energy Association (2009). "Annual Wind Industry Report, Year Ending 2008." March 19th, 2009. <http://www.awea.org/publications/reports/AWEA-Annual-Wind-Report-2009.pdf>
4) R. Doherty and M. O'Malley. A new approach to quantify reserve demand in systems with significant installed wind capacity. *IEEE Transactions on Power Systems* 20(2), pp. 587–595, 2005.
5) L. Landberg, G. Giebel, H.Aa. Nielsen, T.S. Nielsen, H. Madsen. Short-term prediction - an overview. *Wind Energy* 6(3), pp. 273–280, 2003.
6) Wilks, D.S., 1995. Statistical methods in the atmospheric sciences. San Diego, Academic Press.
7) Pa. Windmap. "The complete guide to wind energy resources in Pennsylvania." March 21st, 2009. <http://www.pawindmap.org/index.htm>
8) U.S. Department of Energy (2006). "20% Wind Energy by 2030 Report." March 19th, 2009. <http://www.20percentwind.org/20percent_wind_energy_report_revOct08.pdf>
9) Committee on Appropriations. 2009-02-13. "Summary: American recovery and reinvestment." March 19th, 2009. <http://appropriations.house.gov/pdf/PressSummary02-13-09.pdf.>
10) Box, George and Jenkins, Gwilym (1970) Time series analysis: Forecasting and control. San Francisco: Holden-Day.
11) Elman et al.: A taxonomy and comparison of parallel block multi-level preconditioners for the incompressible Navier-Stokes equations. *Journal of Computational Physics*, vol. 227, 2008.
12) Glahn, H.R., and D.A. Lowry, 1972: The Use of Model Output Statistics (MOS) in Objective Weather Forecasting. *J. Appl. Meteor.*, 11, 1203–1211.
13) <www.wrf-model.org>
14) U.S. Department of Energy: Argonne National Laboratory (June 2009). "Wind power forecasting and electricity market operations." June 30th, 2009. <http://www.usaee.org/usaee2009/submissions/Presentations/SFBotterud.pdf>
15) <www.ruc.noaa.gov>
16) Panofsky, H., and J.A. Dutton, 1984, Atmospheric Turbulence: Models and Methods for Engineers and Scientists. John Wiley and Sons.

17) Robert F. Engle. GARCH 101: The Use of ARCH/GARCH Models in Applied Econometrics. *Journal of Economic Perspectives*, 15(4):157-168, 2001.
18) <www.mmm.ucar.edu/mm5>
19) Haykin, Simon (1998). Neural Networks: A Comprehensive Foundation (2 ed.). Prentice Hall.
20) Ian H. Witten and Eibe Frank (2005) Data Mining: Practical machine learning tools and techniques. 2nd Edition, Morgan Kaufmann, San Francisco, 2005.
21) M.S. Miranda and R.W. Dunn, One-hour-ahead wind speed prediction using a Bayesian methodology. 2006 Power Engineering Society General Meeting.
22) Sideratos, G., Hatziargyriou, N.D., Using Radial Basis Neural Networks to Estimate Wind Power Production. *IEEE Power Engineering Society*, 2007 General Meeting.
23) Enrique Romero and Daniel Toppo, Comparing support vector machines and feed forward neural networks with similar hidden layer weights. *IEEE Trans. on Neural Networks*, vol. 18, no 3. May 2007, pp 959-963.
24) T.G. Barbounis, J.B. Theocharis, Locally recurrent neural networks optimal filtering algorithms: application to wind speed prediction using spatial correlation. *Proc. IEEE International Joint Conf. Neural Networks*, 2005, pp. 2711-2716.
25) Pinson, P., Ranchin, T., Kariniotakis, G., Short-term Wind Power Prediction for Offshore Wind Farms - Evaluation of Fuzzy-Neural Network Based Models, CD-Rom Proceedings of the Global WindPower 2004 Conference, Chicago, Illinois, USA, March 28-31, 2004.
26) Campbell, Sean D. and Diebold, Francis X.,Weather Forecasting for Weather Derivatives (December 2002). PIER Working Paper No. 02-046.
27) Tol, R.S., 1997. Autoregressive conditional heteroscedasticity in daily wind speed measurements. *Theor. Appl. Climatol.*, 56, 113-122.
28) Pinson, P., Kariniotakis, On-line Adaptation of Confidence Intervals based on Weather Stability for Wind Power Forecasting. CD-Rom Proceedings of the Global WindPower 2004 Conference, Chicago, Illinois, USA, March 28-31, 2004.
29) Louka, P., Galanis G., Siebert, N., Kariniotakis, G., Katsafados, P., Kallos, G. Pytharoulis, I., Improvements in wind speed forecasts for wind power prediction purposes using Kalman filtering. CD-Rom Proceedings of the CMMSE 2005, ISBN number 84-609-4844-7.
30) von Bremen, L., Optimal linkage of NWP Models with Neural Networks for Offshore Wind Power Predictions. CD-rom Procd. of the Sixth International workshop on Large-Scale Integration of Wind Power and Transmission Networks for Offshore wind Farms, Delft, Netherlands, 26-28 October 2006.