Forecasting the Subway Volume using Local Linear Kernel Regression

Abstract

Entrusted by the Kaohsiung Rapid Transit Corporation (KRTC), this study attempts to devise a more effective methodology to forecast the passenger volume of the subway system in the city of Kaohsiung, Taiwan. We propose a local linear kernel model to incorporate different weights for each realized observations. It enables us to capture richer information and improve rate of accuracy. We compare different methodologies, for example, $ARIMA$, Best in-sample fit $ARIMA$, linear model, and their rolling versions with our proposed local linear kernel regression model by examining the in-sample and out-of-sample performances. Our results indicate that the proposed rolling local linear kernel regression model performs the best in forecasting the passenger volume in terms of smaller prediction errors in a wide range of measurements.

Keyword: Subway Volume Forecasting, Local Linear Kernel Regression, $ARIMA$

1 Introduction

Kaohsiung city is the second largest city in Taiwan with a population of approximately 2.77 million. The Kaohsiung Mass Rapid Transit System (hereafter “KMRT”), the city’s subway system, covers the metropolitan areas of the city and is operated by the Kaohsiung Rapid Transit Corporation (hereafter “KRTC”) under the contract with the local government. KMRT consists of two lines with 36 stations covering a distance of 42.7 kilometer. The Red Line and the Orange Line started their operation for services on March 9 and September 14, 2008, respectively. The number of passengers that KMRT served in 2014 reached about 168,093 people per day and the accumulated volume of passengers has been over 200 million people since it is launched in early 2008.

The number of the passengers has been increasing drastically since year 2008. In order to better plan and operate the transportation systems[1], understanding and forecasting the number of passengers are extremely crucial. Prior literature has attempted to develop appropriate models for forecasting passenger traffic flows, e.g. air traffic flows (Carson, Cenesizoglu, and Parker, 2011; Fildes, Wei, and Ismail, 2011), the passenger numbers in trains (Nielsen, Frolich, Nielsen, and Filges, 2014), and the freight markets (Batchelor, Alizadeh, and Visvikis, 2007). KRTC also studies the forecasting of the subway volumes internally and publishes a report titled “The Study of Kaohsiung Subway Volume Forecasting - $ARIMA$ Approach” (hereafter “the REPORT (2014)”).

Carson et al. (2011) propose an aggregating individual markets (AIM) approach to predict air travel demand, while Fildes et al. (2011) examine several popular approaches to forecast short- to medium-term air passenger traffic flow. Nielsen et al. (2014) present an innovative counting technique using the weighting systems installed in trains to predict the passenger numbers in the capital region of Denmark. Batchelor et al. (2007) investigate the performance of prevalent time series models, including $ARIMA$ models, VAR models, and $VECM$ models, in forecasting spot and freight rate. Batchelor et al. (2007) find that VECM models generate the best in-sample fit and that all models beat a random walk benchmark in out-of-sample forecasting. The REPORT (2014) applies the $ARMIA$ model to fit the time series data of the

[1] For example, daily operations optimizing, strategic planning and revenue distributing among the operations and etc.

Kaohsiung subway volumes and tests the in-sample performance over the historical time series.

Due to the unsatisfactory predictive capability of the current $ARIMA$ models, the KRTC entrusts us to conduct further research on forecasting the number of passengers with smaller prediction errors. We first consider the linear models, which are able to capture the additional local economy and weather information, see Woodridge (2002), Jin, Nimalendran and Ray (2014). Fan (1992, 1993) proposes an innovative approach in the estimation of unknown regression functions using kernel weighted local linear methods [2], called the local linear kernel regression. The basic prediction problems are the same for both linear model and local linear kernel regression, however, the coefficients are estimated using the weighted realized information in local linear kernel regression, which helps to capture the nonlinearity in historical data. Li and Racine (2004) further discuss how to select the optimal bandwidths in the above mentioned kernel weighting problem using the cross-validation method. Several applications are studied using Li and Racine (2014)'s methodology, for example, Zhu (2014) shows that the local linear kernel regression can provide a better way to predict the crude oil price with much better prediction accuracy. In this paper, we will also apply the local linear kernel regression method in dealing with the subway volume data.

In this article, using the data provided by KRTC and the Kaohsiung City Government, we compare several different prediction methods, including $ARIMA$ models as in the REPORT (2014), linear model, local linear kernel regression and their rolling versions. Both in-sample and out-of-sample tests are conducted and the rolling local linear kernel regression demonstrates the best prediction abilities in forecasting the future subway volumes with considerably small prediction errors. We suggest that the KRTC consider our proposed model to conduct future KMRT demand predictions.

The remainder of the paper is organized as follows. The next section presents our proposed forecasting methodologies. Section 3 describes the data and pre-treatment. We run the unit root test and construct the predictive models. The comparison of different prediction models are then analyzed in Section 4. Finally, Section 5 provides discussions and conclusions.

[2] For further reference, see Ruppert and Wand (1994), Fan and Gijbels (1995) and etc.

2 The Model

2.1 The Basic Prediction Problem

Consider the basic prediction problem with the following general regression form:

$$y = g(x) + \mu$$

where x is the vector of explanatory variables, y is the response variable and μ is the noise term.

In the current application of $ARIMA$ model, only the information of the time series $\{y_i\}_{i=1}^n$ is captured. Following the REPORT (2014), we also include the local gasoline price, the local unemployment rate, the logarithm of the rainfall level, and the temperature of Kaohsiung and others as the explanatory variables, thus the vector x contains not only the lagged information of y, but also the other explanatory variables.

2.1.1 Linear Model

When $g(\cdot)$ is a linear function, the prediction problem is

$$y = a + x'b + \mu$$

and the parameters are simply estimated by the ordinary least squire method,

$$min_{a,b} \sum_{i=1}^{n} (Y_i - a - X_i'b)^2$$

where (Y_i, X_i) are sample realizations. In the linear model, the coefficients a, b are fixed, independent of the input vector x.

2.1.2 Local Linear Kernel Regression

Following Fan (1992, 1993), the information for each realized observation (Y_i, X_i) should not be equally weighted in estimating the coefficients of a, b for different input vector x, if considering the prediction model as a local linear function

$$y = a(x) + x'b(x) + \mu,$$

and the information is weighted by the kernel $K(\frac{X_i - x}{h})$, the parameters $a(x), b(x)$ are estimated by

$$min_{a,b} \sum_{i=1}^{n} (Y_i - a - (X_i - x)'b)^2 K(\frac{X_i - x}{h})$$

where (Y_i, X_i) are sample realizations, $K(\cdot)$ is the kernel function, h is the bandwidth.

The optimal bandwidth h can be selected using cross validation methods (Li and Racine 2004) as below:

$$CV_f(h_1, ..., h_s) = \frac{1}{n^2} \sum_{i=1}^{n} \sum_{j=1}^{n} \bar{K}_h(X_i, X_j) - \frac{2}{n(n-1)} \sum_{i=1}^{n} \sum_{j \neq i, j=1}^{n} K_h(X_i, X_j)$$

when using the integrated squared error as the loss function.

Zhu (2014) shows that the local linear kernel regression can provide a better way to predict the crude oil price with much better prediction accuracy. In this paper, we also apply the local linear kernel regression method in dealing with the subway volume data.

3 Data

3.1 Data Description and Pre-treatment

Our data is provided by Kaohsiung Rapid Transit Corporation and Kaohsiung City Government. The database includes five time series: the logarithm of the monthly subway volume, the local gasoline price, the local unemployment rate, the logarithm of the rainfall level, and the temperature of Kaohsiung[3]. The database spans from April 2008 to December 2013, including a total of 69 observations for each time series.

[Table 1]

We impose three unit root tests on the logarithm of the monthly volume and other explanatory variables (Augmented Dickey-Fuller test, KPSS test and Philipps-Perron test). Table 1 demonstrates the three unit root test statistics and their p-values. Only the $ln(volume)$ cannot reject the existence of unit root at 1% level, while there are insufficient evidence to argue the

[3] We follow the REPORT (2014), which suggests the local gasoline price, the local unemployment rate, the logarithm of the rainfall level, and the temperature of Kaohsiung may affect the subway volume, and all the times series are collected from Kaohsiung City Government.

four explanatory variables also have the unit root[4]. Thus we take the first order differences of $ln(volume)$ and keep the other explanatory variables in the predictive model.

3.2 Predictive Models

After taking care of the unit root issues, we set up the basic predictive model as below:

$$\begin{aligned}\Delta ln(Volume)_t = \alpha + \beta_1 \Delta ln(Volume)_{t-1} + \beta_2 ln(Volume)_{t-1} + \beta_3 Gas_{t-1} \\ + \beta_4 Unemployment_{t-1} + \beta_5 Ln(Rainfall)_{t-1} + \beta_6 Temperature_{t-1} + \epsilon_t\end{aligned} \tag{1}$$

where $ln(Volume)_t$ is the logarithm of the subway volume at time t, Gas_{t-1} is the local gasoline price at time $t-1$, $Unemployment_{t-1}$ is the local unemployment rate at time $t-1$, $Ln(Rainfall)_{t-1}$ is the logarithm of the rainfall level at time $t-1$, $Temperature_{t-1}$ is the temperature of Kaohsiung at time $t-1$. Thus given the information at time $t-1$, we would like to forecast the difference in logarithm of the subway volume at time t, then obtain the subway volume prediction at time t.

In the next section, we would like to apply $ARIMA$, linear model and local linear kernel regression to demonstrate the predictive capability in both training and holdout samples.

4 Empirical Results

With the objective of providing a better forecasting method, we compare four different methodologies including the REPORT $ARIMA$, the best in-sample fitted $ARIMA$, the linear model and the local linear kernel model. We divide the sample into two groups: 1. we use the first 50 observations as the training sample, which is used to construct the fitted model; 2. the remaining 19 observations as the holdout sample, which is used to test performance of the constructed model .

In the REPORT (2014), the internal statisticians only apply the $ARIMA$ model and the best fitted model is $ARIMA(0,1,1)$. However, they only consider the in-sample fitting and testing. Therefore in this paper, we consider both $ARIMA(0,1,1)$ and the best in-sample fitted $ARIMA$ using the in-sample observations (50 observations). Further we also utilize the

[4]Because the unit root test results may be driven by the "Pseudo Long Memory Phenomenon" in a piecewise stationary time series as documented in Jin and Yau (2012), therefore we apply multiple unit root tests here.

linear model and the local linear kernel model to capture additional information provided by Kaohsiung Rapid Transit Corporation and the local government.

4.1 Performance Measures

We implement several performance measures which are common in literature (e.g. Hyndman and Koehler, 2006; Steyerberg et al., 2010; Zhu, 2014; Zhong et al., 2015 and etc.). Our performance measures include Mean Squared Prediction Error (MSPE), Mean Squared Prediction Percentage Error (MSPE), Mean Absolute Percentage Error (MAPE), R^2, Directional Accuracy Ratio and etc.

Mean Squared Prediction Error (MSPE)

Mean Squared Prediction Error (MSPE) is a widely used performance measure (Hyndman and Koehler, 2006) to test the validity of a prediction model with the following form:

$$MSPE = \frac{1}{N}\sum_{i=1}^{N}(Y_i - \hat{Y}_i)^2$$

where N is the number of observations in the training/holdout sample, Y_i is the realized observation and $\hat{Y}_i$ is the prediction value. MSPE is to measure the dispersion of the realized observations and the prediction values.

Mean Squared Prediction Percentage Error (MSPPE)

Mean Squared Prediction Percentage Error (MSPPE) is a percentage version of the dispersion between the realized observations and the prediction values, with the following form:

$$MSPPE = \frac{1}{N}\sum_{i=1}^{N}(\frac{Y_i - \hat{Y}_i}{Y_i})^2.$$

Mean Absolute Percentage Error (MAPE)

Stutzer (1996) and Zhong et al. (2015) argue that the MAPE is a more robust dispersion measure in testing the pricing errors. Thus we also include MAPE as an alternative dispersion measure as below:

$$MAPE = \frac{1}{N}\sum_{i=1}^{N}|\frac{Y_i - \hat{Y}_i}{Y_i}|.$$

ρ (Correlation)

ρ is the correlation between the prediction values and the true observations, which not only measures the level of co-movement but also measures the direction of co-movement (the sign of ρ).

R^2 (Explained Variation)

R^2 is a widely used performance measure for continuous outcomes in both in-sample test and out-of-sample test. Steyerberg et al. (2010) address that R^2 is an overall performance measure, with the following alternative definition:

$$R^2 = \rho^2$$

where ρ is the correlation.

Directional Accuracy Ratio

Directional Accuracy Ratio is a descriptive measure, which measures the percentage accuracy in direction in the prediction values:

$$\text{Directional Accuracy Ratio} = \frac{\#(sign(Y_i) = sign(\hat{Y}_i))}{N}.$$

4.2 Empirical Results

In the REPORT (2014), only $ARIMA$ model and in-sample performance are discussed. However, in a forecasting problem, testing the performance in the holdout sample is much more important than just testing in the training sample. Therefore in this section we will discuss both in-sample and out-of-sample performances of different methods. There are four basic prediction models addressed in our prediction comparisons:

1. REPORT $ARIMA$

The REPORT $ARIMA$ is the prediction model mentioned in the REPORT (2014), i.e. $ARIMA(0,1,1)$. The coefficients are estimated by the training sample.

2. Best In-sample fitted $ARIMA$

The best in-sample fitted $ARIMA$ is the prediction model with the best AIC, AICc and BIC performance in running the training sample, i.e. $ARIMA(1, 1, 0)$[5].

3. Linear Model

The linear model is the model mentioned in Section 2.1.1.

4. Local Linear Kernel Regression

The local linear kernel regression is the model mentioned in Section 2.1.2.

4.2.1 In-Sample Comparison

In-sample comparison examines the in-sample performance on the training sample after calibrating the prediction models by the training sample. We first fit the models using training data, then test their performance. Please notice that in-sample testing does not examine the predictive capability but shows how much variation (in the training sample) can be explained by the fitted models. Figure 1 shows the in-sample fitted time series ($ln(Volume)$) for for four different fitting methods (REPORT $ARIMA$, Best-fit $ARIMA$, Linear Model, Local Linear Kernel Regression) and all the models provide satisfactory performance in in-sample fitting.

[Figure 1]

Table 2 demonstrates the different prediction performance measures (Mean Squared Prediction Error (MSPE), Mean Squared Prediction Percentage Error (MSPPE), Mean Absolute Percentage Error (MAPE), R^2 (Explained Variation), ρ (Correlation) and DAR (Directional Accuracy Ratio)) of the dependent variable $\Delta ln(Volume)$ for four different in-sample fitting methods (REPORT $ARIMA$, Best-fit $ARIMA$, Linear Model, Local Linear Kernel Regression) . Two $ARIMA$ models provide a better overall performance and a higher directional accuracy ratio. They can explain the variation by around 80%. However, the linear model

[5] $ARIMA(0, 1, 1)$: AIC=-69.18; AICc=-68.92; BIC=-65.4. $ARIMA(1, 1, 0)$: AIC=-66.38; AICc=-66.12; BIC=-62.6.

and the local linear kernel regression model have much smaller mean squared errors and two percentage error measures.[6]

[Table 2]

4.2.2 Out-of-Sample Comparison

We mainly focus on the out-of-sample comparison because it demonstrates the prediction ability in a *Holdout Sample*. Notice that the $ARIMA$ models have limited long term forecasting abilities, thus we propose a rolling version for all the four prediction methods: at any time t in holdout sample, we will use all the observations from 1 to t as the training sample, and forecast the value at time $t+1$. On one hand, we can benefit from all up-to-date information; on the other hand, we improve the prediction ability of $ARIMA$ model a lot.[7] Therefore, we will have a total of eight prediction models to compare in this section.

Figure 2 shows the in-sample fitted time series ($ln(Volume)$) for eight different prediction methods (REPORT $ARIMA$, Best-fit $ARIMA$, Linear Model, Local Linear Kernel Regression and their rolling versions). Overall, the rolling versions provide a better predictive capacity in fitting the out-of-sample time series than the non-rolling models.

[Figure 2]

Table 3 demonstrates the different prediction performance measures (Mean Squared Prediction Error (MSPE), Mean Squared Prediction Percentage Error (MSPPE), Mean Absolute Percentage Error (MAPE), R^2 (Explained Variation), ρ (Correlation) and DAR (Directional Accuracy Ratio)) of the dependent variable $\Delta ln(Volume)$ for eight different prediction methods (REPORT $ARIMA$, Best-fit $ARIMA$, Linear Model, Local Linear Kernel Regression and their rolling versions). For non-rolling versions, two $ARIMA$ models have smaller MSPE, MSPPE and MAPE. However, their prediction points have almost zero variation explanatory power in either R^2 or Directional Accuracy Ratio; this is mainly because $ARIMA$ model has no long term prediction power thus there is no variation in the prediction values when using the

[6] We delete the two observations with the true value 0 in calculating the percentage error measures.

[7] For example, the $ARIMA(0,1,1)$ has very weak prediction ability, because according to the formula, $Y_t = Y_{t-1} + \rho\epsilon_{t-1} + \epsilon_t$, we cannot forecast $t+2$ value with any variations. Similarly in our results, the non-rolling version predicted values are constant for $t+1$ to $t+19$.

$ARIMA$ model. Among all the four methods, local linear kernel regression provides the highest R^2 and considerably small MSPE. Linear model provides slightly weaker, but still satisfactory performance in predicting the subway volume in the holdout sample.

[Table 3]

The rolling methods provide better prediction performances for all four methods than the non-rolling versions. Compared with the $ARIMA$ models, our two proposed methodologies demonstrate lower in prediction error measures (MSPE, MSPPE and MAPE), higher variation explanation measure R^2, larger positive correlation ρ and better directional prediction. Among all the eight methods comparison, the rolling local linear kernel regression dominates the other seven in terms of the out-of-sample performance, followed by the rolling linear model. Though the $ARIMA$ models can explain the in-sample variation reasonably well, they have very limited predictive capability even we adjust the model using the rolling method.

5 Conclusion

Entrusted by the KRTC, this study attempts to devise a more effective methodology to forecast the passenger volume of the subway system in the city of Kaohsiung, Taiwan. For a newly-built subway system in a metropolitan area, it is utterly crucial to accurately understand the demand before the service is put into operation. Not only does it facilitate flow of passenger, public security, and operation efficiency, it is also of vital importance in future planning and development. Previous study by the KRTC has applied the $ARIMA$ model to predict the passenger volume in a linear fashion by various variables, for instance, gasoline price, unemployment rate, temperate and so on. However, there is still room for improvement on the predictive capability.

In this study, we propose a local linear kernel model to incorporate different weights for each realized observations. It enables us to capture richer information and improve rate of accuracy. To this end, we compare different methodologies, for example, $ARIMA$, Best in-sample fit $ARIMA$, linear model, and their rolling versions with our proposed local linear kernel regression model by testing the in-sample and out-of-sample performances. Our results indicate that the proposed rolling local linear kernel regression model performs the best in forecasting the passenger volume in terms of smaller prediction errors in a wide range of measurements. Because of the adjusted weights through the kernel function, the predictive capability of the rolling local linear kernel regression model outperforms all the other tested models. In conclusion, we suggest that the KRTC adopt the proposed model for future demand predictions.

References

Batchelor, R., Alizadeh, A., & Visvikis, I. (2007). Forecasting spot and forward prices in the international freight market. *International Journal of Forecasting*, 23(1), 101-114.

Carson, R. T., Cenesizoglu, T., & Parker, R. (2011). Forecasting (aggregate) demand for US commercial air travel. *International Journal of Forecasting*, 27(3), 923-941.

Fan, J. (1992). Design-adaptive nonparametric regression. *Journal of the American statistical Association*, 87(420), 998-1004.

Fan, J. (1993). Local linear regression smoothers and their minimax efficiencies. *The Annals of Statistics*, 196-216.

Fan, J., & Gijbels, I. (1995). Data-driven bandwidth selection in local polynomial fitting: variable bandwidth and spatial adaptation. *Journal of the Royal Statistical Society. Series B (Methodological)*, 371-394.

Fildes, R., Wei, Y., & Ismail, S. (2011). Evaluating the forecasting performance of econometric models of air passenger traffic flows using multiple error measures. *International Journal of Forecasting*, 27(3), 902-922.

Hyndman, R. J., & Koehler, A. B. (2006). Another look at measures of forecast accuracy. *International Journal of Forecasting*, 22(4), 679-688.

Jin, Y., Nimalendran M. & Ray, S. (2014). Volatility and Directional Information-Based Trading in Options. *Working paper.*

Jin, Y., Yau C. Y. (2012). A new Sequential Change Point Detection Test in Time Series. *Working paper.*

Kaohsiung Rapid Transit Corporate Report (2014). The Study of Kaohsiung Subway Volume Forecasting - $ARIMA$ Approach.

Li, Q., & Racine, J. (2004). Cross-validated local linear nonparametric regression. *Statistica Sinica*, 14(2), 485-512.

Nielsen, B. F., Frlich, L., Nielsen, O. A., & Filges, D. (2014). Estimating passenger numbers in trains using existing weighing capabilities. *Transportmetrica A: Transport Science*, 10(6), 502-517.

Ruppert, D., & Wand, M. P. (1994). Multivariate locally weighted least squares regression. *The Annals of Statistics*, 1346-1370.

Steyerberg, E. W., Vickers, A. J., Cook, N. R., Gerds, T., Gonen, M., Obuchowski, N., ... & Kattan, M. W. (2010). Assessing the performance of prediction models: a framework for some traditional and novel measures. *Epidemiology* (Cambridge, Mass.), 21(1), 128.

Woodridge, J. M. (2002). Econometric Analysis of Cross Section and Panel Data.

Zhong, X., Jin Y., Cao J. & Zheng W. (2015). On the Empirical Likelihood Option Pricing. *Working Paper.*

Zhu Y. (2014). Crude Oil Price Prediction. *Working Paper.*

Table 1: Unit Root Test

Table 1 demonstrates the three unit root test statistics and their p-values. The three unit root tests are Augmented Dickey-Fuller test, KPSS test and Philipps-Perron test.

	ADF test	PP test	KPSS test
ln(volume)	-5.2432	-8.2428	2.9062
	(< 0.01)	(< 0.01)	(< 0.01)
gasoline	-4.3656	-2.8406	1.9320
	(< 0.01)	(0.2335)	(< 0.01)
unemployment	-2.5096	-2.4148	1.5869
	(0.3681)	(0.4066)	(< 0.01)
ln(rainfall)	-3.9612	-4.3683	0.0398
	(0.0167)	(< 0.01)	(> 0.1)
temperature	-5.0849	-3.4333	0.0432
	(< 0.01)	(0.0579)	(> 0.1)

Table 2: In-Sample Performance

Table 2 demonstrates the different prediction performance measures (Mean Squared Prediction Error (MSPE), Mean Squared Prediction Percentage Error (MSPPE), Mean Absolute Percentage Error (MAPE), R^2 (Explained Variation), ρ (Correlation) and DAR (Directional Accuracy Ratio)) of the dependent variable $\Delta ln(Volume)$ for four different in-sample fitting methods (REPORT $ARIMA$, Best-fit $ARIMA$, Linear Model, Local Linear Kernel Regression).

	ARIMA(REPORT)	ARIMA(Best-fit)	LM	LLKR
MSPE	0.013	0.014	0.009	0.007
MSPPE	21.678	14.521	4.019	2.724
MAPE	1.913	1.982	1.253	1.098
R^2	0.799	0.776	0.443	0.570
ρ	0.894	0.881	0.665	0.755
DAR	0.857	0.857	0.735	0.755

Table 3: Out-of-Sample Performance

Table 3 demonstrates the different prediction performance measures (Mean Squared Prediction Error (MSPE), Mean Squared Prediction Percentage Error (MSPPE), Mean Absolute Percentage Error (MAPE), R^2 (Explained Variation), ρ (Correlation) and DAR (Directional Accuracy Ratio)) of the dependent variable $\Delta ln(Volume)$ for eight different prediction methods (REPORT $ARIMA$, Best-fit $ARIMA$, Linear Model, Local Linear Kernel Regression and their rolling versions).

	ARIMA(REPORT)	ARIMA(Best-fit)	LM	LLKR	ARIMA(Rolling, REPORT)	ARIMA(Rolling, Best-fit)	LM(Rolling)	LLKR(Rolling)
MSPE	0.011	0.011	0.016	0.022	0.014	0.016	0.009	0.009
MSPPE	1.000	0.999	11.710	27.893	2.269	25.082	1.895	1.498
MAPE	1.000	1.000	2.168	2.993	1.303	2.358	1.151	0.999
R^2	NA	0.108	0.523	0.557	0.539	0.176	0.421	0.416
ρ	NA	0.329	0.723	0.746	-0.734	-0.420	0.649	0.645
DAR	0.000	0.222	0.500	0.500	0.222	0.389	0.556	0.556

Figure 1: In-Sample fitted Time Series

Figure 1 shows the in-sample fitted time series ($ln(Volume)$) for for four different fitting methods (REPORT $ARIMA$, Best-fit $ARIMA$, Linear Model, Local Linear Kernel Regression).

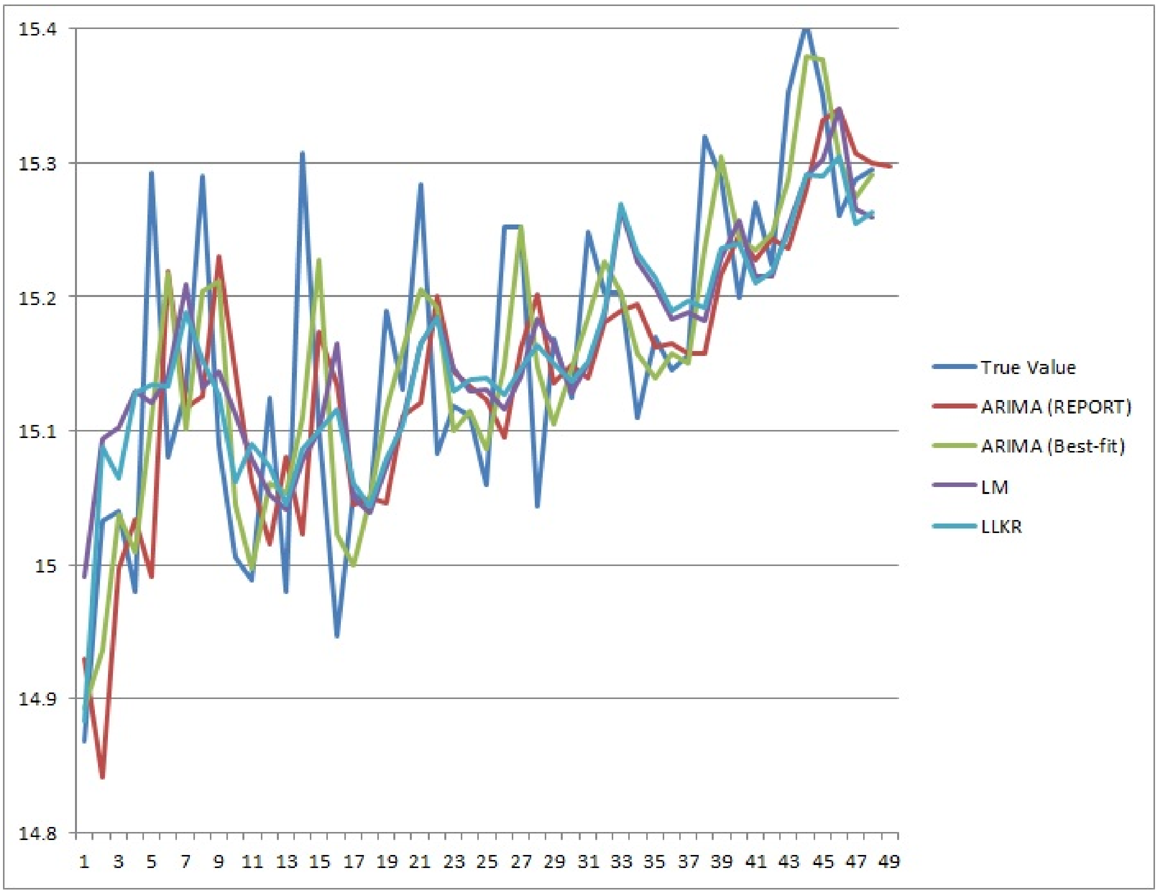

Figure 2: Out-of-Sample fitted Time Series

Figure 2 shows the in-sample fitted time series ($ln(Volume)$) for eight different prediction methods (REPORT $ARIMA$, Best-fit $ARIMA$, Linear Model, Local Linear Kernel Regression and their rolling versions).