

Enhanced Demand Forecasting: A Dynamic Iterative Regression Approach for Time Series Data

Chenghuan Qiang

School of Mechanical Engineering, Shanghai Jiao Tong University, Shanghai, 200240, China

qiang_ch@sjtu.edu.cn

Abstract. Accurate demand forecasting is a cornerstone of efficient resource allocation and data-driven decision-making, especially in operations management. Traditional linear regression models often need help to handle the complexities of time series data, characterized by evolving patterns and unpredictable fluctuations. This study introduces a novel demand forecasting method that departs from the conventional linear regression paradigm. The proposed method focuses on capturing and emphasizing overarching trends and long-term patterns within time series data while minimizing the impact of historical data that may no longer reflect the current state of affairs. The study comprehensively explores the novel approach, providing insights into its underlying principles, features, and algorithmic intricacies. The method's effectiveness is rigorously validated using real-world time series data, exemplified by the Global Average Temperature Anomaly dataset, in direct comparison with traditional linear regression and established techniques. The findings underscore the method's proficiency in capturing and replicating long-term trends, rendering it well-suited for datasets characterized by relatively stable fluctuations. However, it exhibits limited sensitivity to rapid data fluctuations, necessitating alternative strategies for datasets marked by pronounced high-frequency variations. The research delves into the broader implications of these results, emphasizing potential applications in disentangling seasonal variations from overall trend predictions and offering insights into future avenues for improvement and development.

Keywords: Demand Forecasting; Time Series Analysis; Point-by-Point Forecasting; Linear Regression.

1. Introduction

In the realm of operations management, the art of demand forecasting plays a pivotal role in effective resource allocation, strategic planning, and data-driven decision-making. Accurate forecasting not only ensures that organizations meet customer demands efficiently but also aids in the optimization of inventory levels, production schedules, and supply chain operations. In other words, accurate demand forecasting is critical to resource allocation and decision-making, as Fildes and Goodwin highlight the need for improved forecasting methods to enhance organizational decision processes [1]. As organizations across various industries recognize the significance of demand forecasting, it becomes increasingly evident that more than traditional methods may be needed to address modern data's complexities. The need for sophisticated and adaptable forecasting approaches has never been more pronounced.

Many of the datasets we aim to predict are time series data, including the demand data we focus on. Time series data typically exhibit a combination of underlying trends and stochastic noise, making accurate forecasting challenging [2]. As explained in detail by Chatfield (2019), time series analysis plays a crucial role in understanding and predicting temporal patterns, making it a fundamental tool for researchers across various fields [3]. Time series forecasting finds significant applications in diverse domains, including energy and finance [4, 5]. This analytical discipline provides the foundation for modeling and deciphering intricate time-dependent phenomena, ranging from financial markets to climate patterns. Within operations management, the application of time series analysis is indispensable. It equips decision-makers with the ability to grasp the underlying dynamics of demand fluctuations, enabling them to devise more informed strategies for resource allocation.

Traditional linear regression has long served as a workhorse for demand forecasting. Its simplicity and ease of implementation have made it a widely adopted approach. However, linear regression

faces significant limitations when confronted with time series data characterized by evolving patterns and unpredictable fluctuations. In scenarios where historical data exhibit substantial variations over short intervals and lack clear linear relationships, the predictive power of linear regression diminishes. Thus, the quest for improved forecasting methods, as emphasized by Fildes and Goodwin, has become a cornerstone of modern business strategy.

This necessitates the development of innovative forecasting methodologies capable of adapting to the dynamic nature of time series data. In this context, we present a novel demand forecasting method that departs from the traditional linear regression paradigm. Our approach focuses on capturing and emphasizing the overarching trends and long-term patterns embedded within time series data while minimizing the impact of historical data that may no longer reflect the current state of affairs. By doing so, we aim to provide organizations with a more accurate and adaptable tool for demand forecasting that aligns with contemporary data's complexities.

This paper comprehensively explores the new demand forecasting method, adaptive weighted regression forecasting, highlighting its underlying principles, features, and algorithmic details. We evaluate the method's effectiveness compared to traditional linear regression and other established forecasting techniques through a rigorous validation process utilizing real-world time series data. Furthermore, we discuss the implications of our findings and offer insights into potential applications and future directions in demand forecasting. Our research endeavors to bridge the gap between the growing importance of demand forecasting in operations management and the need for advanced forecasting methodologies to navigate modern data patterns' intricate landscape.

2. Literature Review

Pursuing accurate demand forecasting techniques has been a longstanding endeavor in operations research, supply chain management, and economics. This section reviews the literature to provide context for developing our novel forecasting method.

Linear regression has historically been a stalwart method in demand forecasting. Its foundational principles, rooted in statistical analysis, have made it a valuable tool for modeling relationships between independent and dependent variables. However, its applicability is constrained by the assumption of linear relationships, which may not hold in complex, dynamic systems characterized by time-dependent data.

Alternative methods such as moving averages (MA), weighted moving averages (WMA), and exponential smoothing (SES) have emerged as contenders in the field of demand forecasting. These methods often perform well, and all of them are easy to use. The state-of-the-art exponential smoothing, as outlined by Gardner, sheds light on the evolution of forecasting techniques and their relevance in capturing data patterns [6]. These techniques better capture short-term fluctuations and seasonality within time series data. While they excel in certain scenarios, they may only partially exploit long-term trends and patterns.

Autoregressive Integrated Moving Average (ARIMA) models have become versatile and powerful tools for forecasting time series. ARIMA models incorporate autoregressive (AR) components, differencing (I) to achieve stationarity, and moving average (MA) components to capture short-term dependencies. They have proven effective in modeling various time series data, including financial market trends, climate patterns, and demand forecasting.

In addition to ARIMA models, an emerging approach in time series analysis involves decomposing data into its constituent parts: the overall trend and seasonal fluctuations. ARIMA models have been widely adopted in time series analysis due to their ability to capture both short-term dependencies through moving averages and long-term trends through autoregressive components [7]. By decomposing the data, analysts can gain deeper insights into the underlying dynamics and improve forecasting accuracy.

Time series data often exhibit underlying trends and short-term fluctuations, challenging accurate forecasting [8]. Seasonal decomposition of time series data involves regularly identifying and

isolating the recurring patterns. As introduced by Cleveland et al., this method has proven effective in extracting underlying trends and cyclic patterns from time series data, enabling separate modeling of seasonal fluctuations and overall trends [9]. These patterns can be daily, weekly, monthly, or even annual, depending on the nature of the data. Once the seasonal components are extracted, the remaining data represents the underlying trend, which may gradually increase or decrease over time.

The decomposition approach allows for developing forecasting models that separately consider the seasonal variations and the overall trend. This can be particularly valuable in industries where understanding and predicting long-term demand trends and seasonal peaks are crucial, such as retail, agriculture, and energy management.

Our new forecasting method draws inspiration from the principles of data decomposition while incorporating iterative adjustments to regression models. By progressively placing greater weight on recent data and dynamically adapting to changing data characteristics, our method seeks to capture the overarching trends and short-term fluctuations within time series data. It aims to provide a practical and interpretable alternative to complex decomposition methods, bridging the gap between traditional linear regression and more advanced time series techniques.

In the following sections, we delve into the specifics of our novel demand forecasting method and present empirical results to support its efficacy in addressing the challenges posed by time series data with evolving characteristics and rapid fluctuations while drawing inspiration from the capabilities of decomposition-based approaches like seasonal decomposition of time series data.

3. Algorithm Overview

3.1. Principle

Traditional linear regression may need help to produce accurate predictions when applied to time series data lacking clear linear patterns. In cases where the time series exhibits evolving characteristics and no apparent seasonal fluctuations, it is essential to adapt forecasting methods to emphasize the relevance of recent data while still considering the historical context. This new approach enhances demand forecasting by iteratively adjusting regression models, progressively placing greater weight on more recent data, and dynamically adapting to the changing nature of the time series.

The enhanced demand forecasting algorithm improves upon linear regression-based predictions. Instead of treating all historical data equally, it iteratively refines regression models by reducing the influence of early data points and giving more weight to recent observations. This iterative process ultimately results in a composite regression equation considering early and recent data, providing more accurate and adaptive forecasts.

The new demand forecasting method combines iterative demand deviation calculations with fixed-intercept regression to predict demand for a time series dataset. It allows for the dynamic adjustment of demand predictions, accommodating trends and deviations from the overall trend. The method refines its predictions by iterating through the data segments, ultimately offering accurate and adaptable demand forecasts. This approach is suitable for applications where demand forecasting needs to consider the evolving nature of the data.

3.2. Features

At its core, the enhanced demand forecasting algorithm recognizes that traditional linear regression models may need to capture the complexities of a time series dataset adequately.

The new method takes data heterogeneity into account. The algorithm acknowledges that time series data can exhibit evolving patterns and trends. As a result, it does not treat all data points equally but focuses on recognizing shifts in data behavior over time.

The new method emphasizes progressive data. To adapt to changing data patterns, the algorithm starts with an initial linear regression model on the entire dataset. However, it recognizes that early

data points might have different characteristics from recent ones. To address this, it iteratively reduces the influence of early data points by excluding them from consideration in each iteration.

The new method refines the iterative regression. Each iteration creates a new linear regression model using the data subset remaining after the reduction step. The key innovation is that this regression model is constrained to pass through a specific point, typically the latest data point's timestamp, where the y-coordinate is set to 0. This constraint ensures that the model considers the overall trend while giving more weight to recent observations.

The new method is an adaptive composite model. The composite regression model is obtained by combining the results of all iterations. This composite model represents a piecewise linear function across the entire domain of the time series. Importantly, it adapts to the evolving characteristics of the data by incorporating both early and recent data patterns.

The final composite model generated by the algorithm provides enhanced forecasting capabilities. It leverages the strength of early and recent data while addressing the limitations of linear regression models that might need to emphasize early data points.

3.3. Algorithm

The new demand forecasting method is designed to predict demand for a time series dataset comprising a total of m time points. It leverages an iterative approach and a user-defined hyperparameter, n , to partition the time series into n segments, effectively creating n nodes (Including starting point, does not include ending point). The method proceeds iteratively from the first node to the m -th time point, performing linear regression and demand prediction for each segment. The method iteratively forecasts demand values, considering n segments of the dataset.

3.3.1 First Iteration

In the first iteration, linear regression is applied to the data from the first node (starting point) to the n -th last time point, resulting in first-order predictions for all time points. These predictions are subtracted from the actual values to obtain first-order demand deviations.

3.3.2 Subsequent Iterations

A fixed-intercept regression is performed in each subsequent iteration using the demand deviations obtained from the previous iteration. The regression line is forced to pass through a fixed point, where the abscissa is the n th node position, and the ordinate is 0. This yields second-order predictions for the segment. These second-order predictions are then subtracted from the corresponding demand deviations to calculate second-order demand deviations. This process is repeated for n iterations, progressively increasing the order of demand deviations.

3.3.3 Final Demand Prediction

After n iterations, the method collects all n regression equations representing a segment's demand prediction. The user inputs the corresponding node position to predict the demand for a specific time point, and the method evaluates all n regression equations. The results are summed to provide the final demand prediction for that time point.

4. Method

4.1. Data Source

Accurate demand forecasting is a pivotal aspect of efficient resource allocation and decision-making in operations management. To verify the feasibility of the new demand forecast model, it is necessary to use a time series data set to test the performance of this forecast model. The Global Average Temperature Anomaly (GCAG) data set is a good data set that can be used to test the demand forecast model.

The GCAG dataset is a comprehensive repository of global temperature anomaly data. It is derived from NOAA's Global Surface Temperature Analysis (NOAAGlobalTemp), which merges data from

two fundamental sources, Global Historical Climatology Network-Monthly (GHCN-M) and Extended Reconstructed Sea Surface Temperature (ERSST). The first covers land surface temperature measurements, offering a robust and extensive collection of historical climate data. At the same time, the other one contributes ocean surface temperature information, enriching the dataset's global coverage.

GCAG provides a detailed temporal resolution, enabling near real-time analysis of monthly and annual temperature anomalies at a global scale. This temporal granularity is particularly valuable for studying climate variations and trends.

The dataset also supports spatial analysis, allowing for exploring temperature patterns across various geographical regions. It provides data on a global scale and data segregated by hemisphere and distinguishes between land and ocean surface components. Researchers can leverage its interactive mapping tool to investigate spatial patterns of global temperature anomalies and conduct regional temperature assessments.

It's essential to note that recent data within the GCAG dataset are preliminary and subject to thorough quality control procedures. This stringent quality assurance process ensures that the dataset maintains high levels of accuracy and reliability.

It can be seen from the GCAG images (Figure 1) that the global annual average temperature data lacks obvious seasonal fluctuations and focuses on long-term trends. This simplifies the isolation and analysis of overarching long-term trends in the data, which is vital for evaluating how well our forecasting model captures and predicts demand trends over extended periods. Moreover, seasonal fluctuations in data can introduce noise that complicates model validation. The GCAG dataset's focus on anomalies relative to long-term averages minimizes this noise, enabling us to assess our model's performance under more realistic and less noisy conditions.




Fig. 1 GCAG data line chart

4.2. Model Validation Experiments

Hyndman and Athanasopoulos (2018) emphasize the importance of robust forecasting principles and practices, which are essential in addressing the challenges posed by complex time series data.[2] Therefore, forecasting principles and practices must be given sufficient attention. The experimental design for model validation is crucial in assessing the performance of our novel demand forecasting model under the influence of global annual mean temperature data. To validate the feasibility of our new demand forecasting model, we employed a point-by-point forecasting approach without the need for a training and testing dataset split. The detailed experimental procedure is as follows.

4.2.1 Point-by-Point Forecasting

To demonstrate the feasibility of the new model using the GCAG dataset. We need to intercept part of the data in the data set, use the new model to predict subsequent data, conduct multiple experiments, and evaluate the prediction results. When using a new model, consider specifying a

length of time as a specified hyperparameter, called unit of time (abbreviated as UOT), and then divide the intercepted time by this hyperparameter and round to the nearest integer to obtain the number of iterations to be performed. This hyperparameter is necessary for the following reasons. First, it can be observed from the GCAG image that although there is no obvious pattern in the data fluctuations, the up and down fluctuations have a relatively fixed frequency. Iterating n times means dividing the intercepted data set into n segments. If segmented, there may be an impact between the frequency and the frequency of fluctuations. We can choose a better value to improve the prediction results by studying this hyperparameter. Secondly, the data set intercepted during regression should be just a short time. In a word, selecting an appropriate hyperparameter, here units of time (UOT), is critical to the performance of our forecasting model because it allows us to strike a balance between the granularity of data segments and the frequency of fluctuations [10]. For different lengths, It is not reasonable for the data set to be divided into the same number of periods using the same number of iterations. During the experiment, selecting as many intercepted data sets as possible for multiple experiments is necessary, so all data sets starting from the starting point with at least 50 data are used.

4.2.2 Error Calculation

For each point-by-point forecast, we computed the prediction error. The error was determined by measuring the difference between the predicted value and the actual observed value.

We calculated the Mean Squared Error (MSE) to evaluate the model's performance comprehensively. The MSE is a critical performance metric, quantifying the model's predictive accuracy. It was computed as follows:

$$MSE = \frac{1}{n} \sum_{i=1}^n y_i - \hat{y}_i \quad (1)$$

Here, n represents the number of forecast points, y_i denotes the actual observed values, and $\hat{y}_i$ corresponds to the model's predicted values.

5. Results

5.1. Determine Hyperparameters

This section will present the results of our model validation experiments. This will encompass the MSE values obtained from our point-by-point forecasting approach and comparative analysis against MSE values from other models, such as weighted moving average (WMA), simple exponential smoothing (SES), and linear regression. These experimental findings aim to comprehensively assess the model's feasibility and performance when applied to demand forecasting.

Through this meticulous process of point-by-point forecasting and MSE evaluation, we seek to establish the predictive accuracy of our traditional model while comparing its performance against other forecasting methods. This comprehensive analysis will yield insights into the model's potential for accurate demand forecasting without a traditional training and testing dataset division.

Try using all integers from 2 to 136 as the hyperparameter UOT and record the predicted MSE, you can get the following results.

Table 1. Impact of UOT on MSE

UOT	MSE	UOT	MSE	UOT	MSE	UOT	MSE
		35	0.080549	69	0.150593	103	0.285439
2	0.153305	36	0.080162	70	0.152792	104	0.290563
3	0.06538	37	0.080627	71	0.154712	105	0.29895
4	0.047249	38	0.0807	72	0.157089	106	0.308012
5	0.043935	39	0.081184	73	0.159534	107	0.316522
6	0.044638	40	0.082798	74	0.161923	108	0.322726
7	0.044502	41	0.084624	75	0.164315	109	0.329211
8	0.045501	42	0.087355	76	0.166709	110	0.338137
9	0.046162	43	0.089724	77	0.168827	111	0.343938
10	0.046793	44	0.092686	78	0.171648	112	0.351118
11	0.048179	45	0.096003	79	0.174401	113	0.362994
12	0.048598	46	0.099266	80	0.177398	114	0.375237
13	0.050373	47	0.099365	81	0.180559	115	0.387011
14	0.05037	48	0.099295	82	0.183732	116	0.395464
15	0.051752	49	0.099313	83	0.186989	117	0.410028
16	0.053312	50	0.099487	84	0.190302	118	0.417447
17	0.053837	51	0.100901	85	0.193531	119	0.41828
18	0.054127	52	0.102385	86	0.197207	120	0.431299
19	0.05533	53	0.103637	87	0.201141	121	0.446902
20	0.056472	54	0.104561	88	0.205242	122	0.456726
21	0.057776	55	0.106304	89	0.2095	123	0.463438
22	0.060211	56	0.10816	90	0.213773	124	0.470144
23	0.061365	57	0.109853	91	0.218391	125	0.481453
24	0.062192	58	0.11199	92	0.223107	126	0.485802
25	0.063388	59	0.114284	93	0.228162	127	0.496866
26	0.064893	60	0.117173	94	0.232842	128	0.510729
27	0.066055	61	0.119813	95	0.238263	129	0.537863
28	0.067367	62	0.122687	96	0.244074	130	0.556789
29	0.068129	63	0.125996	97	0.250019	131	0.567593
30	0.069578	64	0.129505	98	0.255426	132	0.61411
31	0.072144	65	0.132178	99	0.261816	133	0.670293
32	0.075135	66	0.136098	100	0.267496	134	0.745027
33	0.078109	67	0.140562	101	0.272995	135	0.843149
34	0.079996	68	0.146123	102	0.278225	136	0.876658

When the value of the hyperparameter UOT is greater than the amount of data in the data set, it isn't very sensible, so this part of the data has been eliminated.




Fig. 2 UOT-Related MSE Variation

It can be seen from the image of the change of MSE with UOT that when UOT is 5, MSE reaches the minimum value. The farther the value of UOT is from 5, the greater the MSE will be. Therefore, we can determine the hyperparameter UOT to be 5. The value of the minimum MSE is 0.043935.

5.2. MSE Comparison

Firstly, we compare based on Mean Squared Error (MSE), a widely used metric for assessing predictive accuracy. Below are the calculated MSE values for each of the methods.

Table 2. MSE of different methods

Methods	MSE
New Method	0.043935
Moving Averages (considering the first five data points)	0.012408664
Weighted Moving Averages (weights: 0, 0.198865, 0.083102, 0, 0.718033)	0.009297668
Exponential Smoothing	0.008008798
Linear Regression	0.121693605

From the MSE above comparisons, we can draw the following conclusions.

The New Method demonstrates a significant performance improvement compared to Linear Regression, with an MSE of only 0.043935, whereas Linear Regression has an MSE of 0.121693605. This suggests that the New Method is more effective in handling this dataset. However, the new method's performance could be better compared with other methods. However, the new method has been significantly improved compared to linear regression.




Fig. 3 GCAG-Time when the length of the dataset is 60, 80, 100 or 120.

Draw the graph of the composite function model obtained through multiple regressions and the graph of the corresponding real value. It can be found that the new model is generally better than simply doing linear regression (that is, the straight line where the leftmost graph of the new model is) will have better performance. Intuitively, the closer the prominence of the new model is to the right; that is, the closer it is to the time point to be predicted, the better the fit. This means that the new model can well represent the overall changing trend of the data. This is the success of the new model in improving linear regression. However, the new model still retains the problem of linear regression, which is insensitive to rapid fluctuations in data. GCAG data fluctuates with higher frequency in its overall trend, but the new model and linear regression cannot respond in time to such fluctuations.

6. Discussion

The new forecasting method has strength in capturing overall trends but is insensitive to rapid data fluctuations.

The new forecasting method demonstrates a remarkable ability to capture and replicate the overarching trends present in the data. It excels in modeling the long-term patterns and tendencies that underlie the dataset's behavior. This strength is particularly pronounced as the forecast approaches the target prediction point, where the new method's accuracy shines.

Conversely, the new method exhibits limited sensitivity to rapid, high-frequency fluctuations that may occur within the dataset. These fluctuations are common in datasets like GCAG, characterized by quick and frequent variations over short intervals. The new method, like traditional linear regression, needs help to adapt swiftly to such rapid changes.

In conclusion, our research demonstrates that the new forecasting method excels in capturing and replicating the overarching trends and long-term patterns within time series data. It represents a substantial improvement over traditional linear regression, particularly as it approaches the target prediction point. However, it's essential to acknowledge its limited sensitivity to rapid data fluctuations prevalent in datasets like GCAG.

This forecasting method finds its ideal application in datasets characterized by relatively stable fluctuations, particularly when the underlying trend exhibits periods of changing slopes. In such scenarios, the new method showcases remarkable performance improvements compared to linear regression. Nevertheless, it's important to recognize that for datasets featuring pronounced high-frequency fluctuations, alternative methods like Moving Averages (MA), Weighted Moving Averages (WMA), or Exponential Smoothing (SES) may outperform the new method. However, in direct comparison to linear regression, the new method presents a noteworthy enhancement in forecasting accuracy.

The implications of this study extend to various fields, including demand forecasting, time series analysis, and data-driven decision-making. The new forecasting method offers a robust alternative to traditional linear regression, potentially enhancing forecasting accuracy in scenarios where long-term trends are crucial, including applications in separating seasonal variations from overall trend predictions. By incorporating the new forecasting method into models to dissect seasonal changes and overall trends, researchers and practitioners can better capture and predict the overarching trends within time series data, leading to more precise and informed decision-making processes.

There are several promising avenues for further research and development. Future studies could focus on refining the new forecasting method to improve its sensitivity to rapid data fluctuations while preserving its core advantages. Additionally, exploring hybrid approaches that combine the strengths of the new method with other forecasting techniques could lead to even more accurate and adaptable models.

References

- [1] Fildes Robert, Paul Goodwin. Against your better judgment? How organizations can improve their use of management judgment in forecasting. *Interfaces*, 2007, 37(6): 570-576.
- [2] Hyndman Rob J., George Athanasopoulos. *Forecasting: principles and practice*. OTexts, 2018.
- [3] Chatfield C, Xing H. *The analysis of time series: an introduction with R*. CRC Press, 2019.
- [4] Deb C, Zhang F, Yang J, et al. A review on time series forecasting techniques for building energy consumption. *Renewable and Sustainable Energy Reviews*, 2017, 74: 902-924.
- [5] Krollner B, Vanstone B, Finnie G. Financial time series forecasting with machine learning techniques: A survey. *European Symposium on Artificial Neural Networks: Computational Intelligence and Machine Learning*, 2010: 25-30.
- [6] Gardner Jr E S. Exponential smoothing: The state of the art. *Journal of forecasting*, 1985, 4(1): 1-28.
- [7] Box G E P, Jenkins G M, Reinsel G C, et al. *Time series analysis: forecasting and control*. John Wiley & Sons, 2015.
- [8] Enders W. *Applied econometric time series*. John Wiley & Sons, 2008.
- [9] Cleveland R B, Cleveland W S, McRae J E, et al. STL: A seasonal-trend decomposition. *J. Off. Stat*, 1990, 6(1): 3-73.
- [10] Hamilton J D. *Time series analysis*. Princeton University Press, 2020.