ISSN No:-2456-2165

Harnessing Deep Learning for Enhanced Demand Forecasting in Wooden Pallet Manufacturing

Swarna Chaithanya Kollipara[1], Satya Krishna M B[2], Sai Vishal Golem[3], Purvaja Fursule[4], Bharani Kumar Depuru[5]
[1,2,3]Business Analytics Student, ISB Institute of Data Science
[4]Team Leader, Research and Development, Innodatatics, Hyderabad, India
[5]Director, Innodatatics, Hyderabad, India

*Corresponding Author: Bharani Kumar Depuru – research@innodatatics.com.
ORC ID:0009-0003-4338-8914

Abstract:- Wooden pallet manufacturers contend with erratic demand patterns, impeding optimal resource allocation and operational performance. In the dynamic industry of wooden pallet manufacturing, the imperative for precise demand forecasting arises from this inherent variability in customer demand, demanding accuracy in inventory management, warehouse capacity utilization, and production planning. This study harnesses deep learning models for enhancing demand forecasting in the wooden pallet manufacturing industry because conventional forecasting methodologies encounter difficulties adapting to these dynamic conditions, resulting in inaccuracies and consequential inventory mismanagement, which incur substantial costs.

A comprehensive evaluation of 14 deep learning models, including Autoformer, Informer, Neural Basis Expansion Analysis for Interpretable Time Series Forecasting (N-BEATS), Neural Basis Expansion Analysis for Interpretable Time Series Forecasting with Exogenous Variables (N-BEATSx), Neural Hierarchical Interpolation for Time Series Forecasting (N-HiTS), PatchTST, Prophet, Temporal Convolutional Network (TCN), Temporal Fusion Transformer (TFT), TimeGPT, TimesNet, TSmixer, and the AutoTS library, culminated in the identification of AutoTS as the most effective for consistently accurate predictions. AutoTS library autonomously analyzes customer data, tests approximately 800 models for each customer, and adeptly selects the most suitable model from its expansive library, ensuring optimal forecasting accuracy tailored to each unique customer. The amalgamation of multiple models through AutoTS mitigates risks associated with reliance on a singular algorithm, contributing to producing more robust and reliable forecasts. Rigorous testing on historical data from 3,710 unique customers across India revealed AutoTS's capability to generate precise weekly, bi-weekly, and monthly forecasts, surpassing an accuracy benchmark of 80%.

Integrating an interactive dashboard in the study facilitates real-time data analysis and visualization, fostering informed decision-making in critical operational domains of our client. By delivering highly accurate demand forecasts, this approach empowers wooden pallet manufacturers to efficiently manage inventory, optimize production schedules, and ultimately enhance operational efficiency and profitability.

Keywords:- *Deep Learning Models, AutoTS Library, Machine Learning, Predictive Modelling, Demand Forecasting, Supply Chain Optimization, Inventory Management.*

I. INTRODUCTION

Demand forecasting constitutes the anticipation of future product or service demand, playing a pivotal role in effective supply chain management by ensuring optimal inventory levels aligned with customer and business requirements. Despite the challenges posed by diverse influencing factors such as historical sales data, economic trends, competitor activities, and seasonal patterns [2], the accuracy of demand predictions can be enhanced through various forecasting techniques. While traditional univariate models exhibit limitations in retail demand forecasting, recent advancements in machine learning and deep learning, particularly when integrated with big data and supply chain data, present promising avenues [1].

Wooden pallets, defined as flat wooden structures supporting goods for efficient transportation and storage, are integral components within industrial and logistics processes. This study focuses on a prominent wooden pallet manufacturer (the client) strategically positioned across India, catering to diverse sectors, including 3rd Party Logistics (3PL), Beverages, e-commerce, Manufacturing, and Retail and wholesale. The study emphasizes the critical role of wooden pallets in logistics and distribution, necessitating an optimized production-to-delivery process. Avoiding operational costs due to irregular transportation and distribution hinges on precise demand prediction, streamlined inventory management, and accurate production forecasting [4].

Between January 2019 and August 2023, the client engaged with 448 unique customers, constituting 3710 unique interactions if each transaction is considered a separate customer engagement. The client faced challenges from volatile demand changes during this period, affecting warehouse storage and disrupting pallet production planning. Addressing these challenges through precise forecasting mechanisms becomes imperative to mitigate inefficiencies like stockouts, overstocking, and non-optimized resource allocation, thereby enhancing operational efficiency and minimizing costs. Additionally, solving such complications positively impacts sales and marketing strategies, contributing to potential enhanced returns on investment.

At the core of this study is the utilization of the CRISP-ML(Q) methodology, available as open-source on the 360DigiTMG website (ak.1) [Fig.1][12]. The objective is to develop an innovative demand forecasting system rooted in advanced machine learning and deep learning techniques. Implementing this solution holds transformative potential for the client's business operations, extending beyond sales forecasting to encompass holistic improvements in resource allocation, inventory management [4], marketing effectiveness, and overall decision-making capabilities. While relying on historical sales data for demand prediction, the project acknowledges a limitation—the forecasting model's accuracy is contingent on the quality of the training data. Unforeseen changes, such as gaining or losing a significant customer, pose potential impacts on demand dynamics, influencing model predictions.

When applied to our study, the CRISP-ML(Q) methodology helped design a robust system to extract data from our client’s SAP system, preprocess the data, select only required features and create a data table suitable for applying deep learning models. Later, several traditional and deep learning models were applied to the cleaned dataset, and the results were meticulously monitored and visualized to avoid biases. Further, the data forecasted was made accessible to various teams of our client, such as the CEO’s office, operations, marketing and business development teams through our innovative front-end deployed on Streamlit. The detailed architecture diagram is depicted in [Fig.2].

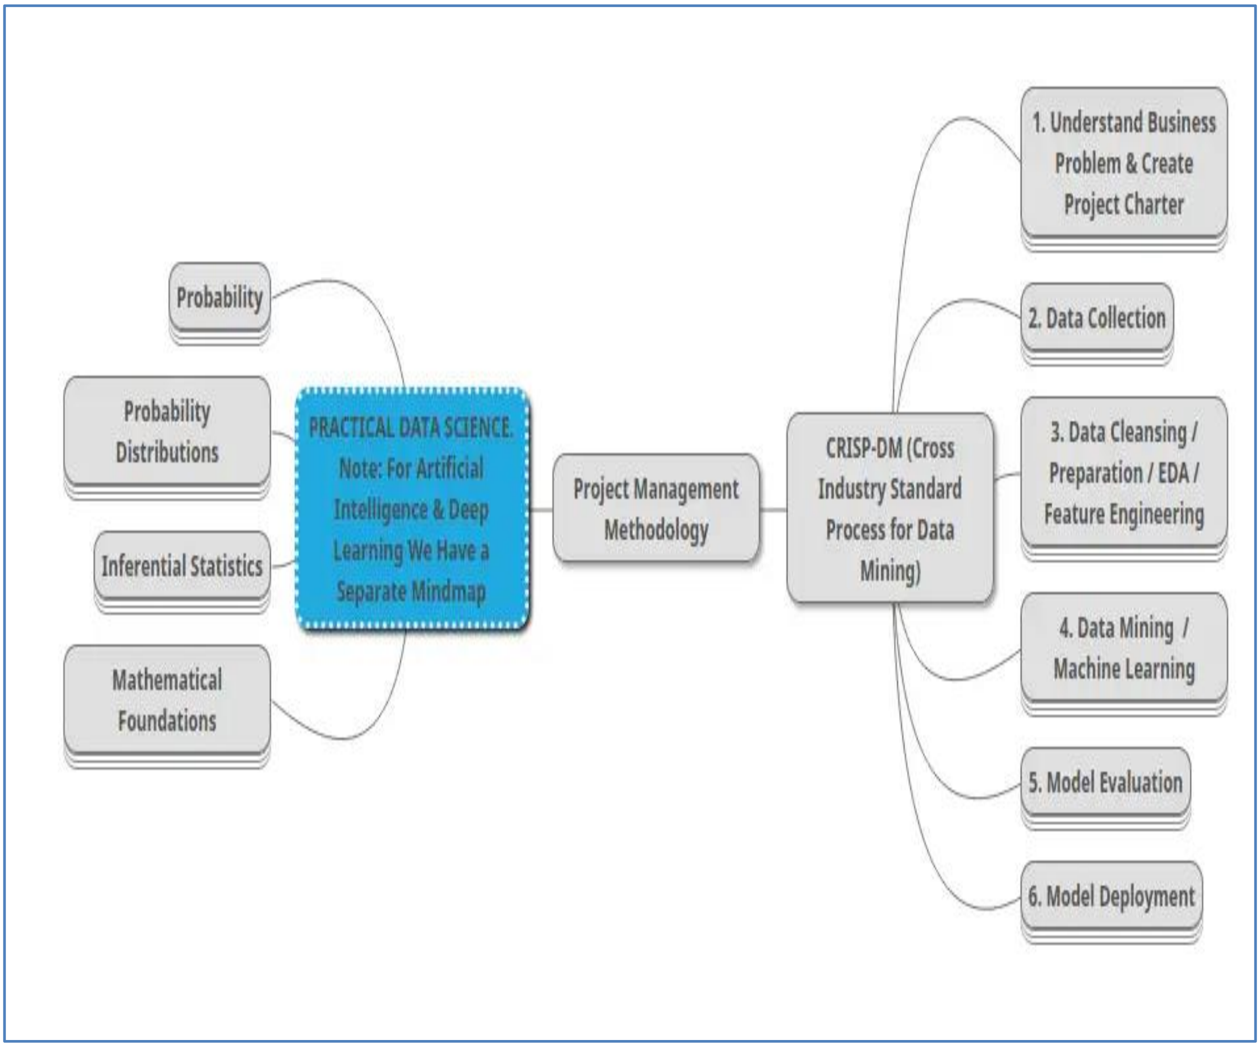

Fig. 1: CRISP-ML(Q) Methodology was Followed to Complete the Project
Source: https://360digitmg.com/mindmap

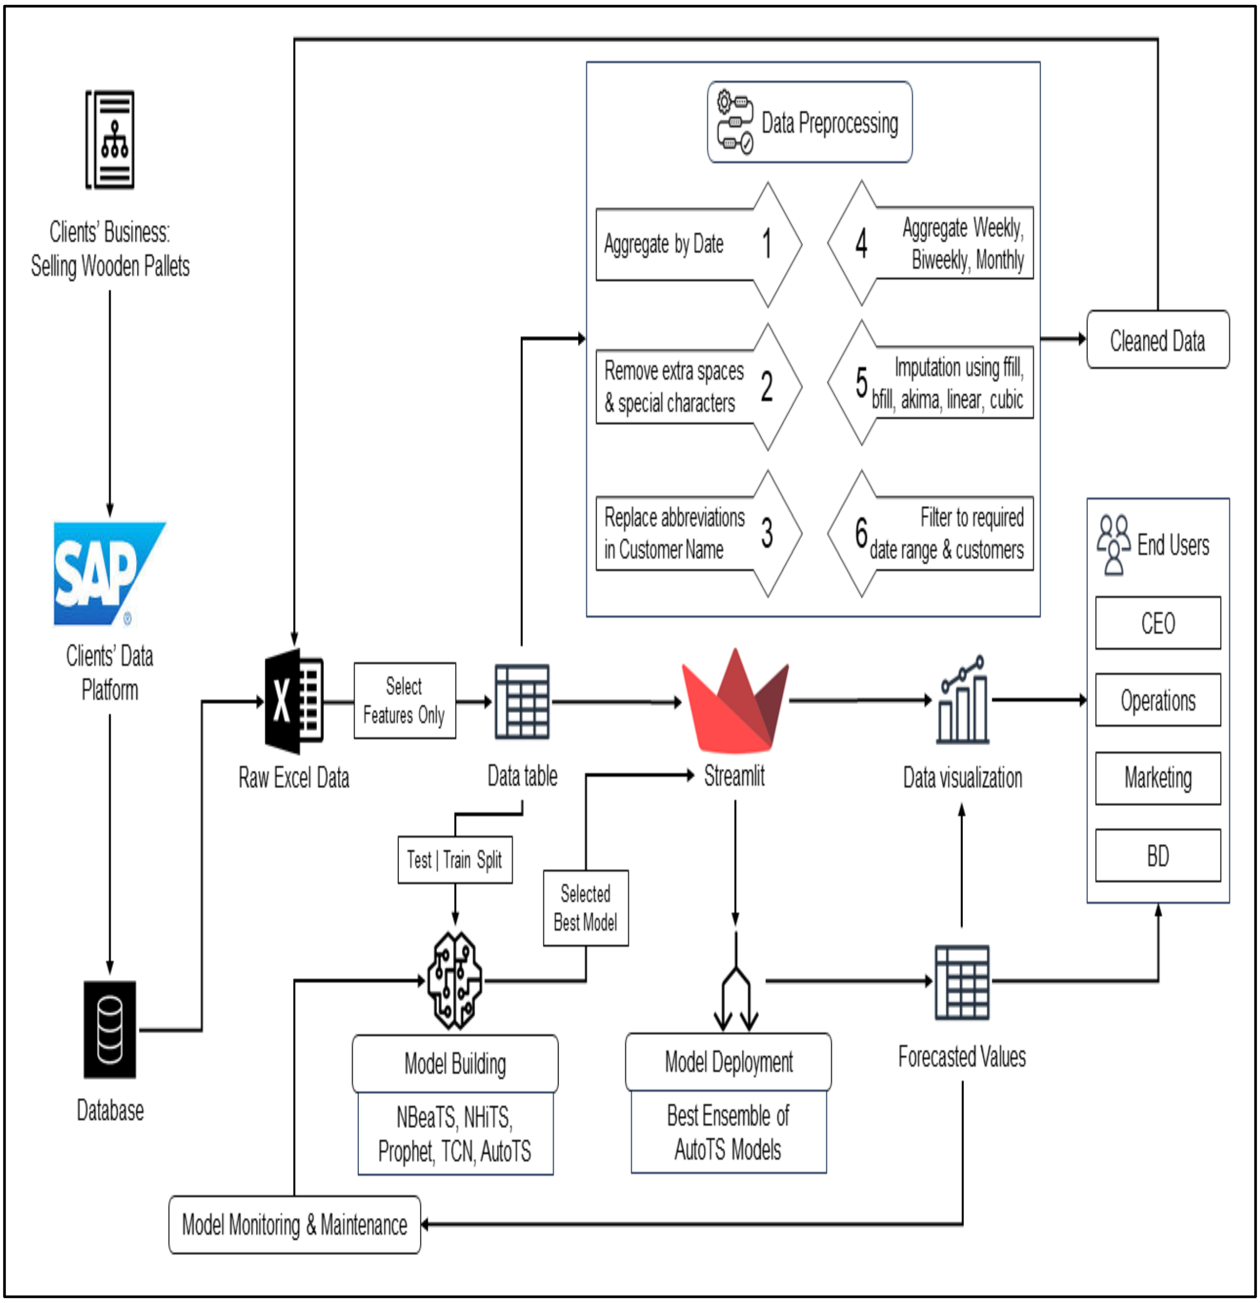

Fig. 2: Architecture Diagram Showing the Flow of the Entire Project with Detailed Information
Source: https://360digitmg.com/ml-workflow

II. METHODS AND TECHNIQUES

A. Data Collection

The client furnished an Excel sheet extracted from their SAP system, as denoted in [Fig.2]. This Excel sheet encapsulated pallet sales data spanning from 1 January 2019 to 23 August 2023. The dataset comprised 80962 rows and 20 columns, offering a comprehensive depiction of the transactions. This structured collection of data provided valuable insights into various transaction details, serving as the primary source for the analysis.

Table.1: Detailed description of the columns in the dataset along with data description, type, units, count, unique count and missing values

S. No.	Descriptor	Description	Type	Units	Count	Unique Values	Missing Values
1	DOC ID	Document Identifier	String	N/A	80962	75358	No
2	Date	Transaction date	Date	DD/MM/YYYY	80962	1608	No
3	BP Code	Business Partner Code	Integer	N/A	80962	4127	No
4	Customer Name	Name of the customer	String	N/A	80962	4183	No
5	City	City of transaction	String	N/A	80962	736	No
6	Region	Region of transaction	String	N/A	80962	4	No
7	State	State of transaction	String	N/A	80962	33	No
8	Product Name	Name of the product	String	N/A	80962	1	No
9	Product Code	Code of the product	String	N/A	80962	70	No
10	Product Code 2	Another code for the product	String	N/A	80962	44	No
11	Description	Description of the product	String	N/A	80962	46	No
12	Transaction Type	Type of transaction	String	N/A	80962	2	No
13	Model	Model type	String	N/A	80962	2	No
14	Model 2	Another model type	String	N/A	80962	2	No
15	QTY	Quantity involved in the transaction	Integer	Pieces / Units	80962	961	No
16	Client Category	Category of the client	String	N/A	80962	1	No
17	SO ID	Sell Order ID	String	N/A	46170	21945	Yes
18	SO Date	Sell Order Date	Date	DD/MM/YYYY	46170	4101	Yes
19	WH Code	Warehouse Code	String	N/A	80918	73	Yes
20	WH Name	Warehouse Name	String	N/A	80918	86	Yes

➢ *Data Preprocessing*

The Data cleaning or preprocessing phase involved careful review of the attributes of the data. The product code was analyzed and over 90% of the transactions had a single product code. A separate analysis for each product code did not make business sense. Hence, we considered every product to be a wooden pallet and proceeded with analysis. The WH code (Warehouse Code) assigned to certain warehouses had inconsistencies in the allocation, which were accordingly modified to be unique to a warehouse location. The BP codes (Business Partner Codes) were assigned to the individual customers, yet, certain codes were duplicated and/or assigned to multiple customers [Table.1]. The warehouse cities and the respective states have also been mismatched and we have worked around the data to match the warehouses with the states. Finally, even the customer names had inconsistencies with certain descriptions of the type of company (Pvt ltd, private LTD, Private Limited etc.,). Hence, we cleaned the data to be consistent. The cleaned data was used to identify the patterns of the customer sales, key locations, seasonality and the distribution of the customers per state and further to the warehouse.

From a business stand-point, we felt segmenting the individual customers into distinct sectors would help in identifying inherent patterns in the data divided by sector. The segmentation of customers into child and parent sectors involved heavy ground work into the business and the market in which the customer operates. We have individually identified the sector of each customer since FMCG was the only sector classification given by the client. The parent sectors namely Beverages, 3PL, E-commerce, Retail and Wholesale, Manufacturing and others have been shortlisted based on the child sectors.

B. Exploratory Data Analysis (EDA):

The EDA for the project had special emphasis on the demographic classification and the segmentation of the customers. Through the demographic classification such as region, state and city, we have identified that south of India had the highest number of orders with Karnataka leading the chart across the states. In addition to this Bengaluru, India, had the highest sales among the cities followed by Ahmedabad, India. We have further identified that covid waves had significant effect on the sales of the pallets and the distribution of pallets per order had been distorted, inferring that the pallets have been ordered in bulk to tackle any adversities due to COVID-19.

The customer segmentation has allowed us to explore the trend and seasonality among the sectors, where the highest number of customers were from the beverage sector. We have done STL decomposition of QTY over time for all the parent sectors, however in the analysis of the 3PL sector, a comprehensive exploration of segment-level STL decomposition was conducted over multiple timeframes. For the 6-month period, a cyclical pattern emerged, revealing growth followed by decline. As the time frame narrowed to 3 months, quarterly fluctuations became more evident, emphasizing pronounced growth in specific quarters. On a monthly scale, granular trends such as short-term spikes and drops were visible, indicating varying demand. Further insights were gained at the 15-day and 1-week intervals, capturing fluctuations related to specific events. The beverages sector displayed similar cyclical behaviors, with semi-annual, quarterly, monthly, bi-weekly, and weekly timeframes revealing distinct patterns [3,4]. The visualizations confirmed consistent seasonality, more pronounced with longer timeframes, providing valuable insights for forecasting and decision-making in both sectors.

The above EDA has allowed us to deep dive into the customer sales, demographic and market segment data. However, to proceed with further analysis we have to extract features from the data since the current segmentation did not provide any impactful insights on a significant volume of the data.

C. Feature Engineering

In our dataset selection process, we intricately focused on 16 customers, a task of notable complexity given the extensive pool of 3710 unique customers available for consideration.

Guided by discussions incorporating client perspectives and analytical rigor, we systematically identified and prioritized specific customers based on diverse criteria. These criteria encompassed transaction volume, the temporal scope of the data, the strategic significance of forecasting for each customer, and the distinctive characteristics exhibited by these entities. This methodical selection process held paramount importance in ensuring that our subsequent analysis was both manageable and in alignment with the primary business objectives of the study. The meticulous selection allowed us to derive meaningful and actionable insights from the dataset.

The dataset encompassed daily sales information for 3,710 unique customers. For analytical convenience, we aggregated the sales data into daily, followed by weekly, bi-weekly, and monthly intervals. Weekly aggregation emerged as particularly informative, revealing that among the 247 weeks of sales data, approximately 32 customers consistently ordered pallets for at least 50 weeks [3,4]. The presence of zeroes in the data varied from 7.29% to 99.60%. Subsequent refinement of the dataset focused on these 32 customers, leading to the application of various deep learning models. Eventually, 16 customers exhibited promising forecasts.

Our strategic decision to concentrate on this subset of customers during the period from November 1, 2021, to August 23, 2023, ensured robust forecasts, encompassing 65% of total orders. In subsequent phases, we refined our forecasting methodology by exploring different aggregation techniques and testing various neural network models [5,8]. The incorporation of feature engineering played a pivotal role in establishing a robust foundation for subsequent analyses, underscoring the significance of tailored approaches in time series forecasting within the wooden pallet manufacturing sector.

D. Model Development

The central aim is to select a model that optimizes the precision of demand predictions while minimizing errors in sales forecasts. This objective is geared towards furnishing our client with a resilient sales prediction solution, utilizing historical sales data and exploring the incorporation of external factors, if feasible.

➢ *Pre-Model Data Assumptions*

Our assumptions encompass the correctness of the provided pallet data, with temporal considerations emphasizing the examination of stable patterns for predicting future sales under the assumption of a relatively stable market and consideration of any seasonality effects. Regarding customer behavior, we presume that customers within the same industry may exhibit different buying habits, yet anticipate loyalty from existing customers. Product-specific assumptions posit uniformity among all wooden pallets, with no substitutes expected to significantly reduce demand. External factors, such as pandemics or sudden government regulations, are not anticipated to exert a major impact. Operational continuity relies on the assumption of smooth supply chain [3] and warehouse operations without major disruptions. Modeling assumptions involve statistical tests assuming independent data points with the same distribution, and certain models assume constant residual variance across different levels of independent variables.

The methodological approach involves meticulous data collection and preparation to construct a reliable dataset. A user-friendly visualization dashboard will be devised for stakeholders to analyze past sales data and derive meaningful insights. Various time series forecasting models, including ARIMA [8,9,10], LSTM, Prophet, regression models, NHiTS, N-BEATS, and libraries such as AutoTS, will be explicated in detail below, encompassing discussions on training, hyperparameter tuning, and performance evaluation.

➢ *Models and Metrics Considered*

With carefully preprocessed data excluding the External factors, aggregating to three different levels namely 15 Days, weekly & Bi-weekly [3,4,6,9] also imputing, we opted for different models to test the performance, and faced specific challenges with each. PatchTST is good at using transformers for forecasting but struggles with different scales and patching strategies. TimesNet is customized for time series but finds it hard to handle noise and optimize patterns. TimeGPT is strong in forecasting but needs careful computer management. TSmixer mixes data in a cool way but gets complicated when blending different time series. N-BEATS is clear but balancing simplicity and complexity is tricky. N-HiTS uses hierarchies but crafting them is tough. Autoformer use transformers but fine-tuning is hard.

Informer captures long patterns but balancing is tricky. TFT is tailored for forecasting but tuning settings is demanding. AUTO TS [11] automates selection but ensuring it works well with different data is a challenge. Prophet is good at seasonal forecasting but adjusting for irregular patterns is tough. Our study brings together these models to tackle diverse forecasting challenges

- *Model Chosen:*

To tackle the challenge of trying different models, we use AutoTS [11], an ensemble approach that automatically picks the best model for accurate forecasts. AutoTS library assesses models individually and together, selecting the most effective one. Key settings, like forecast length, frequency, and prediction range, help optimize the model's performance. With parameters like ensemble and max_generations, AutoTS evolves and improves, ensuring reliable predictions for diverse time series data [Table.2][6].

Evaluation of ensembles in AutoTS is based on the lowest MAPE [3,6,8,9], Mean Absolute Percentage Error is the most suitable measure for assessing the accuracy of forecasting models, especially in demand forecasting.

Table 2: Hyper Parameters Available within the AutoTS Library, Along with Tuning Parameters Available and Selected

Parameter	**Options Available**	**Chosen for Model**
forecast_length	Integer (number of periods to forecast)	20 periods
frequency	'15D', 'W', 'M', etc. (time intervals)	User-selected (e.g., '15D')
prediction_interval	Float (0.80 to 0.99 representing 80% to 99% confidence intervals)	User-selected (e.g., 0.95)
ensemble	'simple', 'distance', 'horizontal', etc.	'simple'
max_generations	Integer (number of algorithm iterations)	5
num_validations	Integer (number of validation splits)	2
validation_method	'backwards', 'walk_forward', 'block', etc.	'backwards'
imputation_method	'None', 'ffill', 'bfill', 'linear', 'akima', 'cubic', etc.	User-selected (e.g., 'linear')

- *Model Fitting and Forecasting:*

The AutoTS model [11] is fitted to the preprocessed and aggregated data. It automatically evaluates various time series models and selects the one best suited for our dataset. The model fitting is visualized through a comparison of original and imputed data over time.

The model runs the forecasting process, outputting predictions along with the upper and lower confidence intervals. These forecasts are presented both in tabular form and graphically, Showcasing historical data with the predicted values and their confidence intervals.

- *Model Deployment:*

We deployed AutoTS model [11] by leveraging the Streamlit open-source framework [Fig.3]. Development code is hosted on a GitHub repository, integrating seamlessly with Streamlit. This strategic move will enable end users to access the website easily, allowing them to run the model on their data. This integration represents a crucial step in enhancing the accessibility and usability of the model, fostering a user-friendly environment where clients can directly interact with the tool and apply it to their specific datasets.

E. Software and Tools

During the project, we used different tools to make things easier. Python [8,10] was the main language, and we used it for pre-processing data, aggregating, imputation, model selection & deployment. Within Python, the packages used were Pandas for data manipulation, Numpy for numerical computations, and Matplotlib and Seaborn for data visualization, including the creation of trendline plots for order analysis. Scikit-Learn was used for machine learning tasks, hyperparameter tuning, and cross-validation, while Library AutoTS [11] was specifically chosen for automated time series forecasting.

To transform our data scripts into user-friendly web applications, our choice was the Streamlit library [Fig.3], providing an open-source platform for hosting interactive applications. Users can upload data, choose specific customers, aggregate data, set confidence intervals, and apply imputation techniques to streamline the data [Fig.4]. Finally, they can push the refined data into the AutoTS Library for a 20-week forecasting [4,7]. Versions we used included Streamlit (1.29.0), AutoTS, Matplotlib (3.4.3), Numpy (1.21.2), Pandas (1.3.2), Openpyxl (3.0.10), and Plotly (5.3.0). utilized Visual Studio Code for coding, and GitHub for keeping track of different versions.

III. RESULTS AND DISCUSSION

Min of MAPE				
Customer	**NBEATS**	**NHiTS**	**Prophet**	**TCNModel**
Customer 1	77.27036274	69.04260029	103.3293346	63.80743708
Customer 2	44.82977226	47.41346898	52.80026749	34.0606392
Customer 3	24.48914317	22.96470138	41.03558419	36.5528239
Customer 4	16.01931425	14.52695216	13.2120438	30.44483043
Customer 5	115.3725487	92.12317116	107.9940539	75.65974989
Customer 6	87.1247304	74.20267487	103.9817524	69.61153334
Customer 7	28.48530968	25.38444532	34.58462849	35.9462726
Customer 8	28.90332605	32.73511726	45.95143086	42.60632996
Customer 9	16.34408723	14.57317049	50.130839	35.75015927
Customer 10	179.0772186	162.5881972	161.6615986	75.54547229
Customer 11	28.96020155	24.4703672	29.14768173	43.178706
Customer 12	57.67485256	69.55427955	79.79882086	76.82517332
Customer 13	23.12198874	17.80155173	18.59398714	17.1413089
Customer 14	125.3977291	102.0842376	173.9128529	80.56029758
Customer 15	92.55648671	70.25664085	106.6908995	80.42910428
Customer 16	459.6526311	362.3025841	144.4583114	94.59718518

Fig.3: Above figure shows the MAPE value comparison for various deep learning models on a monthly time frame

Row Labels	akima	bfill	cubic	ffill	linear
Monthly	4.240053446	3.865563459	5.160878828	4.90935105	2.332060422
Customer 1	6.892626452	9.470206776	6.892626452	10.22296416	8.677005379
Customer 2	6.085139162	7.204892108	10.9173108	13.61488404	9.514434194
Customer 3	15.10203586	16.73416718	12.2424006	16.73416718	14.8580566
Customer 4	7.906427805	10.44031467	8.879164863	9.070508796	9.070508796
Customer 5	16.33779021	18.94857871	23.23529297	23.83388361	9.096602413
Customer 6	4.240053446	8.500525254	7.911951356	8.500525254	2.332060422
Customer 7	19.15498115	20.39475147	19.48834387	19.62555688	19.4407025
Customer 8	18.69686302	20.78307458	8.520832418	16.91651219	8.520832418
Customer 9	5.93724881	3.865563459	6.413710945	4.90935105	8.122501529
Customer 10	16.14017581	18.55439599	18.55439599	21.29892279	18.55439599
Customer 11	9.277551539	17.26461095	9.277551539	17.26461095	9.277551539
Customer 12	27.38555258	23.83028221	31.30995859	41.34832944	34.1970317
Customer 13	5.160878828	5.016762583	5.160878828	6.316989154	5.947015301
Customer 14	20.68187271	22.58160939	20.15342565	27.88584151	24.92522437
Customer 15	6.341908031	12.12754767	27.46802726	13.44077219	13.44077219
Customer 16	34.60455708	29.12230503	35.07772556	31.89107656	35.54782472

Fig. 4: Above figure shows the MAPE value comparison for AutoTS library on a monthly time frame for various imputation techniques

Row Labels	akima	bfill	cubic	ffill	linear
15D	10.37940013	11.57526769	7.869156185	8.557402863	8.440470723
Customer 1	16.69656468	13.23627296	17.365225	15.69870497	8.915713812
Customer 2	33.82294288	29.28644278	37.72784593	33.757226	34.74511216
Customer 3	19.25070389	18.41128019	17.85559213	20.2270575	15.45282361
Customer 4	18.9779666	13.41299163	19.7336641	18.44187051	17.87205039
Customer 5	23.19952672	15.19260182	21.12507514	24.02071262	25.02146914
Customer 6	16.09745625	16.09745625	7.869156185	16.40054833	8.52313021
Customer 7	30.57670074	30.29987362	35.52392338	35.39286831	30.86335656
Customer 8	33.72168506	45.99462573	45.67435832	46.73023317	43.26423209
Customer 9	13.94032691	12.09305497	14.49575408	13.67893593	12.37161857
Customer 10	21.73889784	20.16684468	20.16684468	17.35792194	20.16684468
Customer 11	12.29151621	14.90830866	14.03357598	14.36558935	13.17997851
Customer 12	36.13248892	29.8292434	49.6398009	45.87020619	22.36248334
Customer 13	10.37940013	11.57526769	7.914745561	8.557402863	8.440470723
Customer 14	24.56606821	28.14135937	28.73384174	27.45489217	24.95369671
Customer 15	21.17026705	36.24841222	37.62903511	42.9072948	40.23072359
Customer 16	45.86552263	46.251227	45.00941517	45.08693945	39.48146287

Fig. 5: Above figure shows the MAPE value comparison for AutoTS library on a 15 Day time frame for various imputation techniques

Row Labels	akima	bfill	cubic	ffill	linear
Weekly	9.561770488	10.58466586	9.921443003	9.777344009	10.56740089
Customer 1	16.50782049	17.74909758	17.82514008	17.03671988	20.29265843
Customer 2	18.43596805	22.64058432	32.635804	12.20406817	14.41104693
Customer 3	22.57979277	23.42133041	34.02531109	22.28355618	26.15676007
Customer 4	27.04442987	25.52310688	25.35775402	25.30732779	27.64121083
Customer 5	33.31720497	32.88807402	32.3044944	37.05273147	26.55396988
Customer 6	32.1801534	30.1212085	34.21917185	30.60953856	30.78726336
Customer 7	26.64881323	36.2994437	29.58528854	34.07542638	33.6355761
Customer 8	36.89508128	45.9795206	41.53897853	34.94522238	35.90265418
Customer 9	33.91231279	31.31575717	42.54017288	35.07614055	29.99709125
Customer 10	37.82627541	39.90267703	41.79053649	36.37947973	40.72672551
Customer 11	22.92574267	22.75461615	24.23798128	26.42008423	24.74774177
Customer 12	37.61398558	42.40616152	57.22648884	36.78140364	37.25536318
Customer 13	9.561770488	10.58466586	9.921443003	9.777344009	10.56740089
Customer 14	20.94030676	25.4620365	23.39702592	27.69307208	21.13413638
Customer 15	25.99632032	19.82262778	35.25584896	28.09556357	25.46617961
Customer 16	24.67298953	15.45430238	40.71750654	19.94097606	26.66829551

Fig. 6: Above figure shows the MAPE value comparison for AutoTS library on a Weekly time frame for various imputation techniques

Our results section encapsulates a thorough examination of various deep learning models aimed at determining the most accurate and effective approach for predicting demand in the wooden pallet industry. Meticulous testing and evaluation were conducted on an extensive array of deep learning models, including NBEATS, NHiTS, Prophet, TCN Model, and others, as illustrated in [Fig.3]. Despite adhering to our machine learning success criterion, which required a minimum accuracy of 80%, only a few models demonstrated success for specific customers, while the majority failed to meet the established criteria.

On the other hand, the AutoTS library, that runs a series of at least 30 models, creates an ensemble of models and finally gives the results after at least running 800 different combinations emerged as the most suitable library for forecasting in our particular case. We had used various imputed techniques on weekly, 15-day & Monthly aggregated data and corresponding MAPE values [Fig.4,5,6] are evident.

The success of our initiative transcends mere model selection, encompassing a holistic approach that allows end-users across production, operations, business development, and marketing teams a comprehensive view of forecasts across varying timeframes. This multifaceted approach enables end-users to optimize inventory, monitor warehouse pallets, and access real-time forecasts. Such alignment with business objectives remains attuned to inherent constraints within the pallet manufacturing industry. Equally significant is the attained enhancement of customer satisfaction through timely deliveries facilitated by our tool. Moreover, our tool facilitates the implementation of improved inventory management procedures, effectively mitigating operational inefficiencies. Consequently, we have achieved the overarching business goal of precise demand forecasting, all within the practical constraints inherent to the wooden pallet manufacturing industry.

IV. CONCLUSION

This study thoroughly examines forecasting demand for wooden pallets, a pivotal element in supply chain management and warehouse operations. Our efforts have resulted in the creation of a user-friendly tool, hosted on Github and enabled by Streamlit, facilitating robust and accurate demand forecasting for clients. Founded on deep learning techniques and advanced data analysis, this tool holds the potential for substantial improvements in the operations of the pallet manufacturing industry. Implementing the CRISP-ML(Q) methodology [Fig.1], complemented by meticulous data preprocessing, exploratory data analysis, and advanced modeling, has reshaped conventional paradigms.

In summary, the proposed strategies, encompassing precise demand forecasting, optimized inventory practices, real-time monitoring, and agile production planning, promise transformative impacts by enhancing customer satisfaction, improving production efficiency, reducing costs and conferring a competitive edge. The comprehensive approach, prioritizing data quality, aims to establish a resilient and responsive supply chain, ensuring enhanced profitability and sustained growth for the client while addressing challenges posed by demand fluctuations in the pallet manufacturing industry.

ACKNOWLEDGMENTS

We acknowledge that with the consent from 360DigiTMG, we have used the CRISP-ML(Q) methodology (ak.1)[12] and the ML Workflow (ak.2), which are available as open-source on the official website of 360DigiTMG.

REFERENCES

[1]. Giri, C and Chen, Y. (2022, June 20). Deep learning for demand forecasting in the fashion and apparel retail Industry. https://www.mdpi.com/2571-9394/4/2/31
[2]. Feizabadi, J. (2020, Aug 04). Machine learning demand forecasting and supply chain performance. Pages 119-142. https://www.tandfonline.com/doi/abs/10.1080/13675567.2020.1803246
[3]. Mediavilla, M A., Dietrich, F., Palm, D. (2022, January 01). Review and analysis of artificial intelligence methods for demand forecasting in supply chain management. https://www.sciencedirect.com/science/article/pii/S2212827122004036
[4]. Deng, C., Liu, Y. (2021, September 21). A deep learning-based Inventory management and demand prediction optimization method for anomaly detection.https://www.hindawi.com/journals/wcmc/2021/9969357/
[5]. Kilimci, Z. H., Akyuz, A. O., Uysal, M., Akyokus, S., Uysal, M. O., Bulbul, B. A., &Ekmis, M. A. (2019, March 19). An improved demand forecasting model using deep learning approach and proposed decision integration strategy for supply chain. https://www.hindawi.com/journals/complexity/2019/9067367/
[6]. Terrada, L., El Khaili, M., & Hassan, O. (May, 2022). Demand forecasting model using deep learning methods for supply chain management. https://thesai.org/Publications/ViewPaper?Volume=13&Issue=5&Code=IJACSA&SerialNo=81
[7]. Demand Forecasting with Supply-Chain Information and Machine Learning: Evidence in the Pharmaceutical Industry by Xiaodan Zhu, Anh Ninh, Hui Zhao and Zhenming Liu. https://econpapers.repec.org/article/blapopmgt/v_3a30_3ay_3a2021_3ai_3a9_3ap_3a3231-3252.htm
[8]. Aguiar-Pérez, J.M.; Pérez-Juárez, M.Á. An Insight of Deep Learning Based Demand Forecasting in Smart Grids. Sensors 2023, 23, 1467. https://doi.org/10.3390/s23031467
[9]. Chung, D. ., Lee, C. G. ., & Yang, S. . (2023). A Hybrid Machine Learning Model for Demand Forecasting: Combination of K-means, Elastic-Net,

and Gaussian Process Regression. International Journal of Intelligent Systems and Applications in Engineering, 11(6s), 325–336. Retrieved from https://ijisae.org/index.php/IJISAE/article/view/2859

[10]. Bansal, Archit. (2020). Role of Machine Learning in Inventory Optimization using Time-series Forecasting. https://www.researchgate.net/publication/352765653_Role_of_Machine_Learning_in_Inventory_Optimization_using_Time-series_Forecasting

[11]. Chunnan Wang, Xingyu Chen, Chengyue Wu, and Hongzhi Wang (2022). AutoTS: Automatic Time Series Forecasting https://doi.org/10.48550/arXiv.2203.14169

[12]. Stefan Studer, Thanh Binh Bui, Christian Drescher, Alexander Hanuschkin, Ludwig Winkler, Steven Peters and Klaus-Robert Muller, Towards CRISP-ML(Q): A Machine Learning Process Model with Quality Assurance Methodology, 2021, Volume 3, Issue 2. https://doi.org/10.3390/make3020020.