Hierarchical Forecasting of Spare Parts Components on a Supply Chain Network

Ravi B.D. Figueiredo[1], **Italo P. Caliari**[1], **Wington L. Vital, Natasha A. Caxias, Fabrício S. Batalha**[1]

[1]Faculdade FPFtech
Av. Danilo de Matos Areosa, 1170 – Distrito Industrial
69.075-351 – Manaus – AM – Brazil

{ravi.figueiredo,italo.caliari}@fpf.br

Abstract. *This study evaluates hierarchical demand forecasting for approximately 900 electronic interchangeable component groups within a Brazilian spare parts supply chain. To overcome limited cloud computational resources, we implemented reconciliation methods—Bottom-Up, Top-Down, Min-Trace, and ERM—using PySpark for large-scale distributed processing. Results show that the Min-Trace (MinT) model, including proposed variations to handle negative values, outperformed other methods in MAE and RMSE. The findings confirm that hierarchical reconciliation improves forecasting accuracy and distribution efficiency across planning and distribution levels.*

1. Introduction

Companies in the electronics and microcomputers sector manufacture and distribute a diverse range of products, including computers, peripherals, and mobile devices, all while navigating a highly competitive and dynamic environment. A significant challenge these companies face is managing the demand for components in technical service centers, particularly in fulfilling Service Level Agreements (SLAs) for repairing products that develop defects during the warranty period. Meeting these SLAs is crucial, as they establish clear expectations regarding response times, resolution deadlines, and the availability of necessary parts.

The logistics of ensuring that service centers have the right materials in the appropriate quantities and at the right time are complex. Accurate demand forecasting is essential for optimizing inventory management and ensuring that service centers are adequately stocked. Inadequate forecasting can lead to stockouts, resulting in service delays and customer dissatisfaction. This dissatisfaction can directly impact a company's reputation and lead to lost customers. Conversely, excess inventory incurs unnecessary costs and resource waste, negatively affecting profitability. Companies that struggle to balance supply and demand face high storage costs and risks of product obsolescence.

Despite the challenges, there is limited research specifically addressing the demand for components in technical assistance companies, although related fields have been explored extensively [Fanoodi et al. 2019, Ameer and Fatahi Valilai 2025, Ji and Wang 2017]. For instance, Mahin *et. al.* (2025) presented an innovative approach to sales prediction through advanced machine learning methods, highlighting the effectiveness of the Voting Regressor in improving forecasting accuracy and optimizing supply chain ope-

rations [Rahman Mahin et al. 2025]. Similarly, Espinel *et al.* (2025) explored the integration of econometric modeling and deep learning strategies to improve forecasting accuracy within the pharmaceutical supply chain, ultimately benefiting patient care and operational efficiency [Sierra Espinel and Suarez Barón 2025]. Andrianakis (2024) *et. al.* developed and implemented PredictionSCMS, an integrated supply chain management system designed to optimize product flow, prevent sales losses due to stock shortages, and enhance resource efficiency through inventory management and supply chain automation [Andrianakis et al. 2024].

Hierarchical demand forecasting requires coherence, ensuring lower-level predictions sum exactly to their upper-level aggregates [Athanasopoulos et al. 2024]. These reconciliation methods have become central to supply chain management, bridging operational SKU-level replenishment with long-term strategic planning [Babai et al. 2022]. Empirical evidence across multi-echelon networks [Villegas and Pedregal 2019] and power hierarchies [Silveira G. and Costa 2020] confirms that optimal reconciliation generally outperforms single-level approaches, though structural heterogeneity in demand may still necessitate tailored modeling for specific subsystems.

Ghelasi and Ziel [Ghelasi and Ziel 2024] conceptualize aggregated supply and demand curves as intrinsic hierarchies to apply forecast reconciliation. Alongside traditional models, they introduce an "aggregated-down"benchmark—a localized top-down approach utilizing proportions from the immediate upper node to minimize error aggregation. Theoretically, they prove optimal reconciliation remains independent of specific curve representations. Empirically, applying the X-Model to the German day-ahead electricity market, they demonstrate that optimal techniques (notably shrinkage, weighted least squares, and lambda) consistently outperform bottom-up baselines, validating the efficacy of hierarchical methods in functional curve forecasting.

The remainder of the present paper is organized as follows. Section 2 defines the study methodology, describing the problem formulation, the data set, the forecasting models employed, the hierarchical reconciliation method used, and the metrics performance. Section 3 presents the solution for negative values in the prediction of hierarchical forecasting, 4 presents the results, and 5 presents the conclusions and guidelines for future work.

2. Methodology

2.1. Problem Formulation

Big technology companies need to have a wide technical assistance to guarantee their solutions in a country. Each technical assistance has it own demands on spare parts, and this information is shared with the distribution center, and the forecasting for every spare part in every technical assistance is calculated. Fig 1 shows the hierarchical structure of the spare parts supply chain. The hierarchy is composed of one distribution center (DC), 448 technical assistants (TA), and 2611 spare parts (SP) components. The demand for each TA is not the same, every TA has a different set of SP, which may vary each month. The distribution center is represented as one, even though there are some intermediate distribution centers, because the planning of the distribution is made as if it has just one. Predictions are made at two levels, in DC and TA, with different purposes. The objective at the DC level is to plan the future stock for all the SP in the distribution center in a

forecasting horizon of six months, i.e., the stock has to be replenished with new SP to have the minimum amount to feed the technical assistance, and for that, it is necessary to predict the demand for SP for all TA. The goal for prediction at the TA level is to realize the most efficient distribution among the TAs in the whole country. Without the prediction on TA level, it may lead to a stockout, compromising the product warranty, or leading to an overstock in the TA, which creates additional costs in the whole planning. Predictions at both levels need to be reconciled to ensure that stock planning at the DC level and the distribution of SP to the TA level are consistent. The goal of hierarchical forecasting is to minimize prediction error while maintaining coherence across levels.

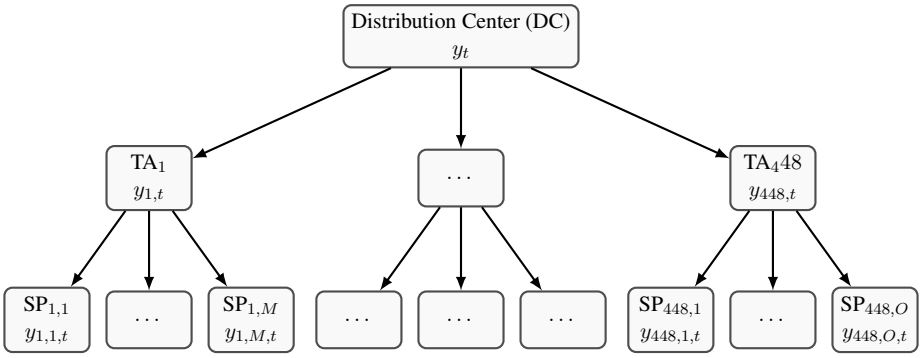


Figura 1. Hierarchical aggregation structure of spare parts demand, illustrating the flow from nine individual Spare Parts (SP) to three Technical Assemblies (TA), and up to the central Distribution Center (DC).

The hierarchical forecasting prediction is calculated in a cloud environment with limited resourcess, the implementation of the hierarchical reconciliation algorithms in this study relies heavily on Just-In-Time (JIT) compilation via the numba library and a C++ backend (coreforecast) to optimize computational throughput across large-scale spare parts data. However, deploying this architecture within a serverless computing paradigm presents critical infrastructure incompatibilities that necessitated the use of a specifically configured, traditional distributed cluster.

Most hierarchical forecast libraries and their associated dependencies necessitate specific system-level native binaries for optimal execution. Traditional distributed clusters permit the execution of initialization scripts to dynamically install missing OS-level dependencies prior to runtime. Conversely, the fully managed nature of serverless architectures explicitly disables cluster-scoped initialization scripts to maintain environment consistency. This limitation precludes the installation of requisite native libraries, rendering the pre-warmed serverless environment incompatible with the low-level requirements of the forecasting algorithms.

2.2. Dataset Description

One of the primary challenges in this study was the sparse historical demand data for individual components, which hindered the development of predictive models. To address this, we aggregated data by creating subgroups of interchangeable components—parts conforming to identical design specifications that can substitute for one another despite minor variations (e.g., different versions of a WXGA+ LCD screen) [Swamidass 2000].

An engineering team manually reviewed datasheets to group these interchangeable parts. This laborious, ongoing process yielded a dataset of replacement histories for 5,936 subgroups comprising 9,452 total components from Brazilian technical assistance. The monthly data tracks the number of components under warranty and the number of damaged items replaced.

Tabela 1. Statistics of the data

	Quantity	Mean	Std. Dev.	Min	Max
Subgroup	5936	1.61	4.71	1	334

As shown in Table 1, subgroups are generally small (averaging 1.6 components), underscoring the difficulty of the manual grouping process. Furthermore, Figure 2 illustrates the distribution of these parts clustered by similarity for easier visualization. Motherboards and assembly parts are the most common, followed by covers, LCDs, full product replacements (FinalizedProduct), and minor categories.

2.3. Base Forecasting Models

This section describes the prediction models used for forecasting using the reconciliation methods. This study is based on a real scenario, in which the prediction models have already been selected, and their performance is not validated in the sense of finding the best model, but to use the prediction outputs to find the best hierarchical conciliation method.

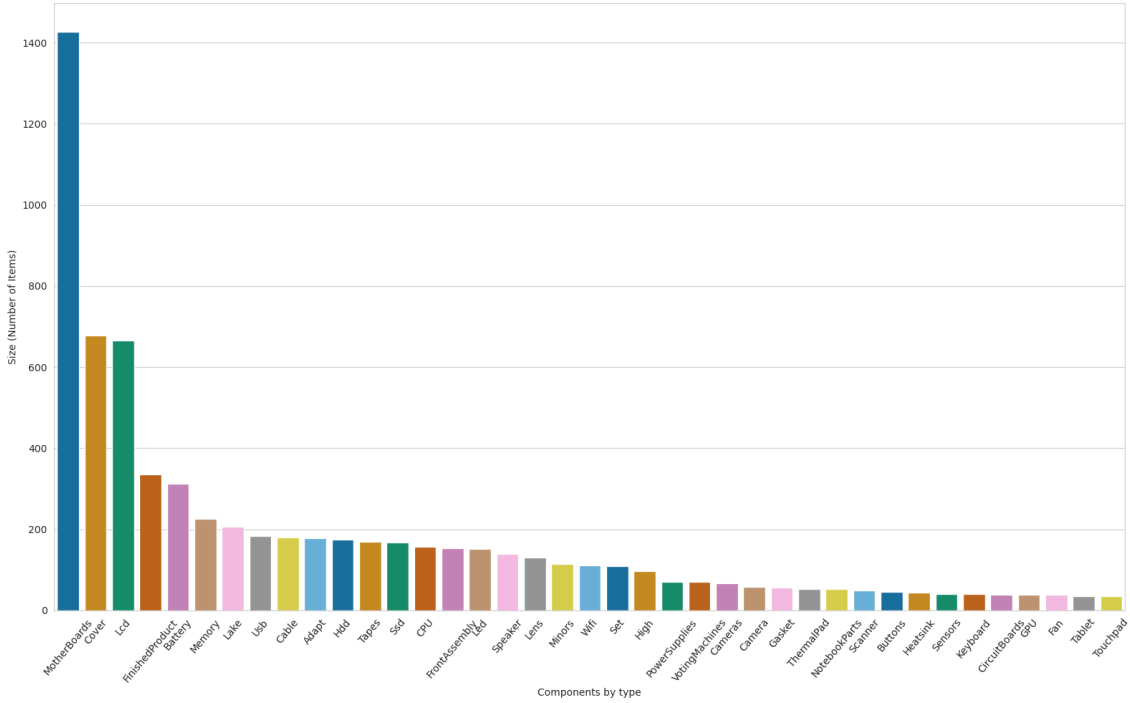


Figura 2. Sufficiency results for the six-month forecast for models with different pre-processing techniques.

2.3.1. Simple Moving Average (SMA)

The Simple Moving Average (SMA) is a statistical method employed in time series analysis and forecasting. It calculates the average of a selected set of data points over a specified period. In this approach, each data point in the time series is given equal weight, meaning that no weighting factors are applied to any of the observations [Hansun 2013]. It can be expressed as:

$$SMA_t = \frac{Y_{t-1} + Y_{t-2} + \cdots + Y_{t-n}}{n}, \tag{1}$$

where SMA_t is a simple Moving Average forecast for observation t, Y_t denotes the data point at observation t, and n indicates the number of observations within the moving window.

2.3.2. The Exponential Moving Average (EMA)

The Exponential Moving Average (EMA) is a variant of the moving average that gives greater weight to more recent data points, making it more responsive to new information compared to the Simple Moving Average (SMA). Unlike traditional moving averages, the EMA is mathematically defined as a limit of an exponential average with a specific weighting factor [Klinker 2011]. The weights assigned to older data points decrease exponentially, ensuring that they never reach zero. It can be expressed as:

$$EMA_t = \frac{(1-\alpha)Y_{t-1} + (1-\alpha)^2 Y_{t-2} + \cdots + (1-\alpha)^n Y_{t-n}}{(1-\alpha) + (1-\alpha)^2 + \cdots + (1-\alpha)^n}, \tag{2}$$

where EMA_t denotes the Exponential Moving Average forecast for observation t, Y_t denotes the data point at observation t, n indicates the number of observations within the moving window, and α represents the degree of weighting decrease, a constant smoothing factor between 0 and 1.

2.4. Seasonal AutoRegressive Integrated Moving Average with Exogenous Variables - SARIMAX

The SARIMAX model [Box et al. 2015] extends the traditional ARIMA framework by incorporating both seasonal patterns (S) and exogenous variables (X). It is formally denoted as $SARIMAX(p,d,q)(P,D,Q,M)$, where the non-seasonal parameters (p,d,q) represent the autoregressive order, degree of differencing, and moving average order, respectively. Correspondingly, (P,D,Q) capture these components for the seasonal cycle, while M specifies the number of periods per season. This structure allows the model to leverage past values, forecast errors, and external drivers to produce stationary, seasonally-adjusted predictions.

A $SARIMAX(p,d,q)(P,D,Q)_s$ model is mathematically expressed using the backshift operator B (where $B^k y_t = y_{t-k}$) as follows:

$$\phi_p(B)\Phi_P(B^s)(1-B)^d(1-B^s)^D y_t = \theta_q(B)\Theta_Q(B^s)\epsilon_t + \sum_{i=1}^{m} \beta_i X_{i,t}, \tag{3}$$

where s is the seasonal period length, $\phi_p(B)$ and $\Phi_P(B^s)$ are the regular and seasonal autoregressive (AR) polynomials of orders p and P, respectively. $\theta_q(B)$ and $\Theta_Q(B^s)$ are the regular and seasonal moving average (MA) polynomials of orders q and Q, d and D are the regular and seasonal differencing orders required to achieve stationarity, ϵ_t represents the white noise error term, and $X_{i,t}$ represents the i-th exogenous variable at time t (e.g., scheduled maintenance flags or equipment age), and β_i is its corresponding coefficient.

2.5. Hierarchical Reconciliation Methods

2.5.1. Bottom-Up (BU)

The Bottom-Up approach is the most straightforward reconciliation method. It entirely discards the base forecasts generated at the upper levels of the hierarchy and instead relies exclusively on the bottom-level forecasts, aggregating them upward using the structural matrix S. The projection matrix for BU is defined as:

$$P_{BU} = [\mathbf{0}_{m\times(n-m)} \mid \mathbf{I}_m] \tag{4}$$

where $\mathbf{0}$ is a null matrix and $\mathbf{I}_m$ is an identity matrix of size m. While BU guarantees coherency without information loss at the granular level, it can perform poorly if the bottom-level data is highly sparse or plagued by intermittent demand, as the aggregation of pure noise propagates upward.

2.5.2. Top-Down (TD)

Conversely, the Top-Down method focuses on the most aggregated level of the hierarchy, leveraging the statistical stability found in larger volumes of data. It takes the base forecast of the top-level node and disaggregates it down the hierarchy using a defined set of proportions. The projection matrix is formulated as:

$$P_{TD} = [\mathbf{p} \mid \mathbf{0}_{m\times(n-1)}] \tag{5}$$

where $\mathbf{p} \in \mathbb{R}^m$ is a vector of disaggregation proportions. These proportions can be calculated using historical averages or forecasted proportions. While TD excels at capturing macro-level trends, it completely ignores granular dynamics and assumes that lower-level consumption distributions remain stationary over time.

2.5.3. Minimum Trace (MinT)

Minimum Trace is an optimal reconciliation technique that utilizes the base forecasts from all levels of the hierarchy. It adjusts the base forecasts by minimizing the trace of

the covariance matrix of the coherent forecast errors, effectively minimizing the overall forecast variance. Under the assumption that the base forecasts are unbiased, the MinT projection matrix is given by:

$$P_{MinT} = (S^{\intercal} W^{-1} S)^{-1} S^{\intercal} W^{-1} \tag{6}$$

where $W \in \mathbb{R}^{n \times n}$ represents the covariance matrix of the base forecast errors. In practice, W is often estimated using the in-sample residual variance. MinT is highly effective in complex hierarchies, but the inversion of the large covariance matrix can be computationally intensive when dealing with high-dimensional item catalogs, such as extensive spare parts inventories.

2.5.4. Empirical Risk Minimization (ERM)

Empirical Risk Minimization shifts the reconciliation paradigm from statistical assumptions to a machine learning framework. Instead of deriving P from a presumed error covariance structure, ERM directly optimizes the projection matrix by minimizing a predefined loss function (e.g., L_1 norm for robustness against outliers or L_2 norm) calculated on a validation set. The optimal matrix is found by solving:

$$P_{ERM} = \arg\min_{P} \sum_{t=1}^{T} \mathcal{L}(y_t, SP\hat{y}_t) \tag{7}$$

where $\mathcal{L}$ is the chosen empirical loss function, and T represents the time steps in the validation period. This data-driven approach allows the reconciliation process to adapt to specific demand structures, such as extreme intermittency, without strictly requiring the unbiasedness assumptions inherent in optimal reconciliation methods like MinT.

2.6. Evaluation Metrics

The most used accuracy metrics for evaluating forecast accuracy, are mean absolute percentage error (MAPE), mean absolute error (MAE), mean absolute scaled error (MASE), or root mean square error (RMSE) [Silveira G. and Costa 2020], with the following definitions:

$$MAE = \frac{1}{T} \sum_{t=1}^{T} |y_t - \hat{y}_t| \,, \tag{8}$$

$$RMSE = \sqrt{\frac{1}{T} \sum_{t=1}^{T} |y_t - \hat{y}_t|^2}, \tag{9}$$

$$BIAS = \frac{1}{T} \sum_{t=1}^{T} (y_t - \hat{y}_t) \,, \tag{10}$$

where y_t is the real value at time t, and $\hat{y}_t$ is the predicted value for the same time t.

3. Handling Negative Values in Hierarchical Forecasting

Additive hierarchical reconciliation methods, such as Min-Trace, frequently encounter boundary constraint issues when applied to strictly non-negative quantities, such as spare parts inventory. The fundamental limitation arises because the standard implementation utilizes an unconstrained additive adjustment that does not inherently respect a zero-lower-bound. When the top-level aggregate forecast is significantly smaller than the sum of the independent bottom-level forecasts, the resulting downward adjustment can force bottom-level node predictions into negative territory.

3.1. Unconstrained Additive Adjustments

The reconciliation process distributes the discrepancy between the aggregate forecast and the sum of base forecasts according to historical variance weights. The reconciled forecast is calculated as:

$$y_{reconciled,i} = y_{base,i} + w_i \times \left(Y_{top} - \sum y_{base,j} \right) \tag{11}$$

Consider a scenario where the top-down forecast (Y_{top}) predicts 100 total units, but the sum of the independent bottom-up forecasts ($\sum y_{base}$) equals 150 units. The system must reconcile a negative gap of 50 units. This penalty is distributed according to the variance weights (w_i). If a highly volatile spare part has a base forecast of 5 units but receives a historical variance weight of 20%, the algorithm applies an adjustment of $-50 \times 0.20 = -10$.

Consequently, the final reconciled forecast for this node drops to -5 units. Because the underlying solver minimizes variance without recognizing the physical impossibility of negative stock, it mathematically balances the hierarchy at the expense of logical domain boundaries. Addressing this limitation requires modifying the reconciliation framework, with three primary methodologies offering different trade-offs between computational efficiency, strict hierarchical summation, and variance minimization.

3.2. Post-Reconciliation Clipping

The most computationally pragmatic approach involves post-reconciliation clipping. In this methodology, the standard Min-Trace algorithm is executed, and any resulting negative values are subsequently bounded to zero.

This operation is highly efficient and preserves the underlying variance-weighted distribution of error, ensuring that highly volatile components absorb the majority of the reconciliation penalty. However, it introduces a slight disruption to the strict hierarchy: by clipping negative values up to zero without redistributing the resultant volume, the sum of the clipped bottom forecasts will marginally exceed the top-level aggregate. In many practical supply chain applications, this minor aggregation discrepancy is an acceptable trade-off for execution speed.

3.3. Proportional Distribution

To maintain strict hierarchical summation without violating non-negativity constraints, researchers can abandon the additive adjustment in favor of a multiplicative proportional distribution. This method recalculates the bottom-level nodes by distributing the top-level

forecast proportionally, based on the ratio of each base forecast to the total sum of base forecasts:

$$y_{reconciled,i} = Y_{top} \times \left(\frac{y_{base,i}}{\sum y_{base,j}} \right) \tag{12}$$

This approach guarantees non-negative results and perfectly reconciled nodes. However, it fundamentally discards the core advantage of the Min-Trace approach. High-variance and low-variance components are treated equally and scaled purely by their base forecast magnitude rather than their historical error profile, potentially penalizing highly stable components unnecessarily.

3.4. Constrained Optimization

While proportional allocation and post-reconciliation clipping offer computationally lightweight heuristics, they inherently compromise either the variance-minimizing properties of the Min-Trace estimator or the strict aggregation constraints of the hierarchy. To satisfy both conditions simultaneously—ensuring that adjustments reflect historical reliability while strictly respecting the zero-lower-bound—the reconciliation process must be formally defined as a constrained optimization problem.

Specifically, the objective is to determine the reconciled bottom-level forecasts, y_i, that minimize the variance-weighted squared deviations from the original base forecasts, $\hat{y}_i$. This formulation constitutes a Quadratic Programming (QP) problem, expressed mathematically as:

$$\min_{\mathbf{y}} \sum_{i=1}^{n} \frac{(y_i - \hat{y}_i)^2}{\sigma_i^2} \tag{13}$$

Subject to:

$$\sum_{i=1}^{n} y_i = Y_{top} \quad \text{(Strict Hierarchical Aggregation)} \tag{14}$$

$$y_i \geq 0 \quad \forall i \quad \text{(Non-Negativity Constraint)} \tag{15}$$

Where σ_i^2 represents the historical variance of the forecast errors for node i.

In the unconstrained Min-Trace model, high volatility can lead to negative forecasts if the required adjustment exceeds the base value. To prevent this, the constrained Quadratic Programming (QP) framework enforces a non-negativity constraint ($y_i \geq 0$). When a node's unconstrained trajectory drops below zero, the solver binds it at exactly zero and redistributes the residual penalty among remaining positive nodes. This redistribution is optimized by variance weights, ensuring highly stable components are shielded until more volatile nodes are fully exhausted.

4. Results

The hierarchical reconciliation methods are analyzed with the RMSE, MAE and Bias metrics in both hierarchical levels. Table 2 presents the metrics of predictions on the

distribution level (bottom level) with predictions for the three prediction models using all the reconciliation methods, and also the mean value for the reconciliation method. The best value for MAE is for the *Bottom-up* reconciliation method with good results for *Min-trace* and *ERM*. For RMSE metrics, *Min-trace* has the best results followed by *ERM* and *Bottom-up*. The Bias metrics is studied because one of the objectives is to guarantee that the service center has the right amount of stock, and the best results are the positive values closer to zero.

Table 3 presents the results for the planning level (Top), bests results are from *Min-trace* reconciliation method, while the *Bottom-up* has the worst results. The results from both tables show that the bottom-up reconciliation methods don't have value information to use for forecasting purposes alone, deriving results from bottom to up has an enormous error, showing that the bottom time series levels are very noisy or have intermittent values. The best method is the ones that get information in both levels to use for forecasting purposes, like *Min-trace* and *ERM*, so the best reconciliation method for this problems is the *Min-trace*.

Tabela 2. Hierarchical Forecast Metrics by Hierarchical Reconciliation Method in the Distribution level.

Model	mean_MAE	media	stddev_MAE	mean_RMSE	media	stddev_RMSE	Viés
Bottom-up-ARIMA	7.07	**8.613**	1.41	138.79	171.68	24.67	5.85
Bottom-up-EMA	9.29		2.63	184.12		66.70	5.07
Bottom-up-SMA	9.48		2.72	192.14		69.49	4.90
Top-down CV-AutoARIMA	24.35	25.843	2.11	297.13	330.813	43.97	-23.04
Top-down CV-EMA	25.56		3.25	321.32		114.95	-23.96
Top-down CV-SMA	27.62		5.10	373.99		185.10	-21.89
Min-trace-ARIMA	10.55	10.833	0.88	186.96	**150.76**	32.78	-7.61
Min-trace-EMA	1.42		0.071	8.23		1.73	-1.13
Min-trace-SMA	20.53		15.92	257.11		240.75	-0.27
ERM-ARIMA	10.5486	11.44	0.8362	185.71	159.58	28.86	-7.96
ERM-EMA	1.4160		0.0730	8.23		1.81	-1.15
ERM-SMA	22.3624		18.255	284.8281		277.455	6.97

Tabela 3. Hierarchical Forecast Metrics by Hierarchical Reconciliation Method in the planning level

Model	mean_MAE	media	stddev_MAE	mean_RMSE	media	stddev_RMSE	Viés
Bottom-up-ARIMA	49,402.31	51,215.97	8,922.83	50,368.27	53,218.00	9,214.44	49,398,00
Bottom-up-EMA	51,943.60		13,029.52	54,810.42		12,535.28	51,341.00
Bottom-up-SMA	52,302.00		13,707.28	54,475.31		13,128.72	52,299.00
Top-down CV-ARIMA	91.10	98.46	7.08	608.52	824.42	84.37	-86.11
Top-down CV-EMA	96.53		12.12	667.61		240.71	-90.00
Top-down CV-SMA	107.76		22.91	1,197.14		1,047.53	-78.70
Min-trace-ARIMA	54.85	**45.01**	0.88	465.14	**250.96**	71.39	15.31
Min-trace-EMA	35.68		0.071	76.92		5.68	-5.30
Min-trace-SMA	44.51		15.92	210.83		144.92	-0.18
ERM-ARIMA	55.01	75.80	3.49	463.10	540.19	67.95	-45.52
ERM-EMA	35.16		5.60	74.16		1.81	-33.11
ERM-SMA	137.23		96.46	1,083.33		1,017.58	84.36

As discussed in Section 3, the standard **Min-trace** method can yield illogical negative forecasts in supply chain contexts. To address this, we evaluated three modified non-negative solutions: Clipping, Proportional distribution, and Constrained Optimization

(Tables 5 and 4). Overall, these adjustments outperform the original Min-trace by reducing distance-to-zero errors. The Proportional method achieved the best overall results in the study, particularly at the planning level, while both Clipping and Proportional performed well at the distribution level. Conversely, the Constrained method yielded the worst results across both levels, driven by the poor performance of the underlying EMA forecasts.

Tabela 4. Forecast Metrics for Min-trace without negative values in the Planning level.

Model	MAE	MAE_{mean}	MAE_{std}	RMSE	$RMSE_{mean}$	$RMSE_{std}$	BIAS
Clipping-ARIMA	57.06			587.73			-25.43
Clipping-EMA	16.81	34.85	20.44	110.49	401.1	255.04	-9.40
Clipping-SMA	30.72			505.08			6.53
Proportional-ARIMA	46.23			509.81			-46.23
Proportional-EMA	16.55	27.78	16.1	120.93	283.10	202.30	-16.53
Proportional-SMA	20.56			218.57			-11.63
Constrained-ARIMA	35.85			390.13			-35.85
Constrained-EMA	8966.69	3,005.4	5,162.56	1,479,139.34	493,244.32	853,810.13	8,950.41
Constrained-SMA	13.93			203.50			-1.83

Tabela 5. Forecast Metrics for Min-trace without negative values in the Distribution level.

Model	MAE	MAE_{mean}	MAE_{std}	RMSE	$RMSE_{mean}$	$RMSE_{std}$	BIAS
Clipping-ARIMA	2.58			78.06			-1.08
Clipping-EMA	0.48	1.34	0.90	8.01	43.65	28.61	-0.19
Clipping-SMA	0.95			44.86			0.29
Proportional-ARIMA	1.97			76.28			-1.97
Proportional-EMA	0.41	**0.98**	**0.70**	7.93	**37.66**	**28.60**	-0.41
Proportional-SMA	0.55			28.78			-0.26
Constrained-ARIMA	2.65			94.41			-2.65
Constrained-EMA	552.80	185.38	259.81	188355.50	62832.28	88758.33	552.19
Constrained-SMA	0.68			46.91			0.07

5. Conclusion

This study presents the investigation of the planning and distribution of spare parts in a supply chain network using the proper hierarchical information to realize forecasting predictions in a cloud environment with restricted configurations. The motivation is to use hierarchical forecasting to help plan the purchase of new spare parts and the distribution of the spare parts in the stock warehouse. Due to some restrictions in the cloud environment, the reconciliation methods have to be rewritten purely in a distributed and parallel environment, such as PySpark. Our results show that different reconciliation approaches have significant differences in the performance metrics, presenting some characteristics for the time series in the hierarchical structure. Best metrics result are from the Min-trace and variation for solving the negative prediction values, with emphasis on the method of proportional distribution that achieves the lower values for MAE and RMSE.

Future works will be include more recent approaches such as machine learning for the reconciliation methods, create a comparative study on the gains of using purely distributed computing libraries like PySpark, and using another characteristics of the hierarchical structure, like the city and state distribution of the spare parts.

Referências

Ameer, T. and Fatahi Valilai, O. (2025). Predictive exploratory data analysis of shopfloor cnc machine operation through a machine learning model. *Journal of Open Innovation: Technology, Market, and Complexity*, 11(2):100559.

Andrianakis, I., Gkatas, V., Eleftheriadis, N., Ellinidis, A., and Avramidou, E. (2024). Predictionscms: The implementation of an ai-powered supply chain management system. *International Journal of Industrial and Manufacturing Engineering*, 18(3):65 – 71.

Athanasopoulos, G., Hyndman, R. J., Kourentzes, N., and Panagiotelis, A. (2024). Forecast reconciliation: A review. *International Journal of Forecasting*, 40(2):430–456.

Babai, M. Z., Boylan, J. E., and Rostami-Tabar, B. (2022). Demand forecasting in supply chains: a review of aggregation and hierarchical approaches. *International Journal of Production Research*, 60(1):324–348.

Box, G. E. P., Jenkins, G. M., Reinsel, G. C., and Ljung, G. M. (2015). *Time Series Analysis: Forecasting and Control*. John Wiley & Sons, 5 edition.

Fanoodi, B., Malmir, B., and Jahantigh, F. F. (2019). Reducing demand uncertainty in the platelet supply chain through artificial neural networks and arima models. *Computers in Biology and Medicine*, 113:103415.

Ghelasi, P. and Ziel, F. (2024). Hierarchical forecasting for aggregated curves with an application to day-ahead electricity price auctions. *International Journal of Forecasting*, 40(2):581–596.

Hansun, S. (2013). A new approach of moving average method in time series analysis. In *2013 Conference on New Media Studies (CoNMedia)*, pages 1–4.

Ji, W. and Wang, L. (2017). Big data analytics based fault prediction for shop floor scheduling. *Journal of Manufacturing Systems*, 43:187–194.

Klinker, F. (2011). Exponential moving average versus moving exponential average. *Math. Semesterber.*, 58(1):97–107.

Rahman Mahin, M. P., Shahriar, M., Das, R. R., Roy, A., and Reza, A. W. (2025). Enhancing sustainable supply chain forecasting using machine learning for sales prediction. *Procedia Computer Science*, 252:470–479. 4th International Conference on Evolutionary Computing and Mobile Sustainable Networks.

Sierra Espinel, A. I. and Suarez Barón, M. J. (2025). Applying deep learning and forecasting techniques to the pharmaceutical supply chain. *Procedia Computer Science*, 253:2791–2800. 6th International Conference on Industry 4.0 and Smart Manufacturing.

Silveira G., T. and Costa, M. (2020). Forecasting hierarchical time series in power generation. *Energies*, 13:3722.

Swamidass, P. M. (2000). *INTERCHANGEABILITY*, pages 293–294. Springer US, New York, NY.

Villegas, M. A. and Pedregal, D. J. (2019). Forecasting hierarchical time series in supply chains: an empirical investigation. *International Journal of Production Research*, 57(11):3666–3684.