Progressive Delta Forecasting: A Self-Correcting Sales Uplift Model Using Virtual Slopes and Adaptive Weights

Dhaneesh Nair (dhaneesh@salescode.ai) Rakesh Singh (rakesh@salescode.ai) Samrat Chowdhury (samrat.chowdhury@salescode.ai) Yagvendra Singh (yagvendra.singh@salescode.ai)

***Abstract*—Sales forecasting in volatile environments remains a complex challenge due to the inherent randomness, incomplete patterns, and dynamic market fluctuations. Traditional predictive models attempt to fit deterministic structures to historical data, limiting their ability to adapt to real-time variations and individual salesperson efficiency. To address these limitations, we propose Progressive Delta Optimization (PDO) with an Adaptive Weight Adjustment Mechanism, a novel self-learning framework that dynamically adjusts decision variables based on nth-order transactional patterns.**

A central component of our approach is the virtual slope, computed using past sales data to establish a threshold that determines the validity of future predictions. This mechanism ensures that each projected sales outcome aligns with the expected trajectory while mitigating deviations due to randomness. Additionally, we introduce the tuning factor (τ), representing salesperson efficiency, which is dynamically estimated in real-time based on business conditions and individual performance metrics. The Progressive Delta Model (PDM) automatically refines this factor, adjusting its weight with each transaction and adapting to fluctuations in business strategy and salesperson push.By continuously recalibrating the virtual slope and tuning factor, PDO ensures a controlled, progressive trajectory that prevents negative drifts while sustaining an upward sales momentum. Experimental results on randomized sales datasets demonstrate an high success rate in achieving positive sales uplift, even under highly stochastic conditions. The proposed methodology offers a robust, self-correcting optimization strategy, making it highly adaptable for real-world sales forecasting and strategic decision-making.

I. Introduction

Sales optimization in dynamic environments presents a fundamental challenge due to *incomplete patterns, stochastic fluctuations, and non-stationary behavior*. Classical methods such as regression analysis, time-series decomposition, and machine learning attempt to fit a predefined model to historical data. However, such techniques inherently assume *structured patterns*, which fail under conditions of extreme randomness.

A key limitation of conventional models is their reliance on **static or semi-static weight allocations**, preventing dynamic adaptability. To overcome this, we propose **Progressive Delta Optimization (PDO)**, a **self-learning model** that ensures controlled uplift by iteratively tuning an adaptive **tuning factor** (τ), modulated via transactional feedback.

The proposed methodology is built upon:

- **Dynamic weight adjustment** to account for the adjustment of probablistic deltas.
- **Self-learning feedback** to iteratively refine τ.
- **Progressive delta enforcement**, ensuring no net negative drift by using virual slope.

Weight Adjustment Mechanism

At each transaction step n, the weight update follows:

$$\omega_{n+1} = \omega_n + \beta \cdot g(\delta_n, \tau_n) \tag{1}$$

where:

- ω_n is the adaptive weight at step n,
- β is the adaptive learning coefficient,
- δ_n represents observed deviations in transactional behavior,
- $g(\delta_n, \tau_n)$ is a nonlinear function ensuring adaptive convergence.

Tuning Factor Optimization

The **tuning factor** (τ) is dynamically controlled based on nth-order transactional behavior:

$$\tau_{n+1} = \tau_n + \alpha \cdot f(\delta_n, \omega_n) \tag{2}$$

where:

- α is the learning rate parameter,
- $f(\delta_n, \omega_n)$ ensures progressive correction to stabilize growth,
- τ enforces the desired uplift direction while preventing undesired stagnation.

II. Background and Motivation

In the competitive domain of sales and marketing, precise forecasting of customer purchase behavior is essential for optimizing revenue growth. However, conventional predictive models often generate recommendations misaligned with business objectives, particularly when the observed data follows a declining trend. Mathematically, given a purchase history represented as a sequence

$$q_1, q_2, q_3, \ldots, q_n, \quad \text{where} \quad q_{i+1} < q_i. \tag{3}$$

traditional models may extrapolate this pattern and predict

$$q_{n+1} \approx q_n - \Delta q \tag{4}$$

Such an outcome is suboptimal for sales maximization, as it contradicts business constraints that require recommended purchase quantities to remain stable or increase. This limitation arises because standard predictive techniques fail to integrate domain-specific constraints, such as monotonicity conditions ensuring non-decreasing recommendations, which are crucial for driving sales growth.

To mitigate this limitation, we propose **Progressive Delta**, a novel predictive modeling technique that ensures the **virtual slope** of forecasted purchase quantities remains non-negative. This method leverages key factors such as purchase frequency, recency, SKU clustering, supplier channels, and multiple purchase parameters, collectively referred to as **probabilistic weights**. By incorporating these domain-specific constraints, Progressive Delta establishes a robust framework for generating recommendations that align with business objectives. This approach is particularly relevant in domains where sustaining or increasing sales is a strategic priority, including retail, e-commerce, and subscription-based services.

Let $Y = \{y_1, y_2, \ldots, y_n\}$ represent a sequence of observed purchase quantities for a customer over time. The goal is to predict the next purchase quantity y_{n+1} such that the slope of the sequence Y is non-negative. Formally, we require:

$$y_{n+1} \geq y_n$$

This constraint ensures that the recommended purchase quantity does not decrease, aligning with the business objective of sales upliftment. The problem can be framed probabilistically as follows:

Let $P(y_{n+1}|y_1, y_2, \ldots, y_n)$ be the conditional probability distribution of the next purchase quantity given the observed sequence. We seek to model this distribution such that:

$$P(y_{n+1} \geq y_n|y_1, y_2, \ldots, y_n) = 1$$

This probabilistic formulation ensures that the predicted quantity y_{n+1} is always greater than or equal to the last observed quantity y_n, thereby maintaining a non-negative slope. The problem of transforming data to meet specific constraints has been extensively studied in statistics and machine learning. The **Box-Cox transformation** [4] and its extension, the **Yeo-Johnson transformation** [10], are widely used for stabilizing variance and making data more normally distributed. However, these methods do not inherently enforce non-negative slopes in predictive models. Robust regression techniques, such as those discussed in [2] and [9], focus on outlier detection and resistance to contamination but do not address the specific constraint of non-negative trends. Recent work by [5] introduces constrained regression models for time series forecasting, but their approach is limited to linear constraints and does not account for the dynamic nature of customer purchase behavior.

In the context of customer behavior analysis, segmentation techniques like **RFM (Recency, Frequency, Monetary) modeling** [8] and clustering methods such as **K-Means** and **DBSCAN** [7] have been employed to group customers based on their purchase patterns. However, these approaches often lack the ability to incorporate domain-specific constraints into predictive models. Recent advancements in **clusterwise regression** [6] and **topology-based methods** [3] offer promising directions for integrating segmentation and prediction but still fall short of addressing the non-negativity constraint. Additionally, [11] propose a probabilistic framework for customer behavior modeling, but their method does not explicitly enforce non-decreasing trends in predictions. This paper makes the following key contributions:

1) **Progressive Delta Technique:** We introduce a novel method that ensures non-negative slopes in predicted purchase quantities by incorporating weightage based on frequency and recency.
2) **Integration with Existing Methods:** We evaluate Progressive Delta against established transformations (Box-Cox, Yeo-Johnson) and robust regression techniques, demonstrating its effectiveness in maintaining positive trends.
3) **Application to Sales Upliftment:** We provide a practical framework for applying Progressive Delta in real-world sales scenarios, ensuring that recommendations align with business objectives.
4) **Insights from Related Domains:** We draw on insights from customer segmentation, outlier detection, and clusterwise regression to enhance the robustness and applicability of our method.

III. Proposed Methodology

The **Progressive Delta** algorithm is designed to ensure that the **virtual slope** of predicted purchase quantities remains non-negative. This is achieved by incorporating weightage based on the frequency and recency of purchase patterns, as well as an adaptive **Turing Factor** to dynamically adjust predictions based on evolving trends.

- **Data Preprocessing:** Normalize the purchase history data to ensure consistency and eliminate anomalies.
- **Weight Assignment:**Assign **probabilistic weights** to each purchase based on recency, frequency, SKU clustering, and supplier channels. Recent and frequent purchases receive higher weights.
- **Virtual Slope Calculation:** Compute the **virtual slope** of the purchase sequence using weighted least squares regression, ensuring that recent purchases have a stronger influence.
- **Tuning Factor Adjustment:** Introduce a **Tuning Factor** that adaptively adjusts predictions by accounting for stochastic variations in purchase behavior. These variations may arise due to business dynamics, salesperson efficiency, skill levels, market fluctuations, and other external factors. By incorporating the Tuning Factor, the model ensures that predictions remain robust and adaptable to real-world uncertainties, enhancing the accuracy and reliability of purchase forecasts.
- **Prediction:** Forecast the next purchase quantity such that the **virtual slope** remains non-negative, while dynamically adjusting for unexpected variations using the **Tuning Factor**. This ensures alignment with business objectives and sales growth strategies.

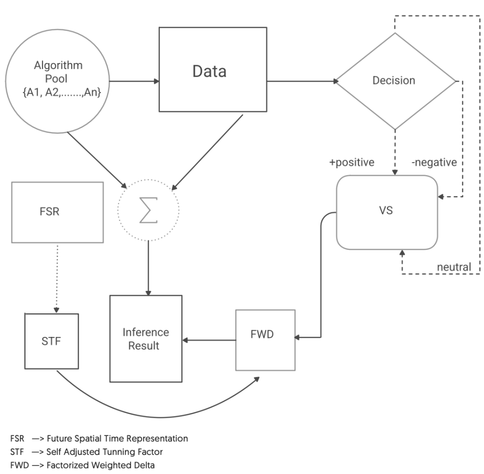

Fig. 1. Flowchart of Progressive Delta

The flowchart illustrates the detailed methodology of the proposed algorithm, beginning with the selection of an appropriate predictive model from a pool of algorithms, which may include linear, quadratic, or polynomial approaches. The decision-making process incorporates **absolutism and normalization** by introducing a novel technique termed **Virtual Slope**, ensuring that sales projections do not decline even when historical purchase patterns exhibit a downward trend.

Subsequently, the algorithm leverages a **Self-Adjusting Tuning Factor (STF)** to determine the future time learning rate, enabling dynamic adaptation to evolving purchase behaviors. By accounting for uncertainties, fuzziness, and negative trends, the STF refines the prediction process, ensuring an uplifted sales forecast. This approach systematically evaluates all possible variations within the system, generating the most **practically optimal and business-aligned** sales recommendations.

Mathematically, let $Y = \{y_1, y_2, \ldots, y_n\}$ represent a sequence of observed purchase quantities, and let $W = \{w_1, w_2, \ldots, w_n\}$ represent the corresponding weights. The slope β is calculated as:

$$\beta = \frac{\sum_{i=1}^{n} w_i (x_i - \bar{x})(y_i - \bar{y})}{\sum_{i=1}^{n} w_i (x_i - \bar{x})^2}$$

where x_i represents the time points, $\bar{x}$ and $\bar{y}$ are the weighted means of x and y, respectively. The predicted purchase quantity y_{n+1} is then calculated as:

$$y_{n+1} = y_n + \beta \cdot \Delta x$$

where Δx is the time interval between the last observed purchase and the next predicted purchase.

Probabilistic Approach

To incorporate uncertainty into the predictions, we model the conditional probability distribution of the next purchase quantity y_{n+1} given the observed sequence Y. Let $P(y_{n+1}|y_1, y_2, \ldots, y_n)$ be the conditional probability density function (PDF). We assume that y_{n+1} follows a normal distribution with mean μ and variance σ^2:

$$y_{n+1} \sim \mathcal{N}(\mu, \sigma^2)$$

The mean μ is calculated using the Progressive Delta algorithm, and the variance σ^2 is estimated from the historical data. The probability that $y_{n+1} \geq y_n$ is then given by:

$$P(y_{n+1} \geq y_n) = 1 - \Phi\left(\frac{y_n - \mu}{\sigma}\right)$$

where Φ is the cumulative distribution function (CDF) of the standard normal distribution.

Comparison with Existing Techniques

Box-Cox Transformation: The *Box-Cox transformation* [4] is a power transformation that stabilizes variance and makes data more normally distributed. It is defined as:

$$y(\lambda) = \begin{cases} \frac{y^\lambda - 1}{\lambda} & \text{if } \lambda \neq 0, \\ \log(y) & \text{if } \lambda = 0. \end{cases}$$

While the Box-Cox transformation is effective for stabilizing variance, it does not inherently enforce non-negative slopes in predictive models. In contrast, Progressive Delta explicitly ensures non-negative slopes by incorporating weightage based on frequency and recency.

Yeo-Johnson Transformation: The *Yeo-Johnson transformation* [10] is an extension of the Box-Cox transformation that can handle both positive and negative values. It is defined as:

$$y(\lambda) = \begin{cases} \frac{(y+1)^\lambda - 1}{\lambda} & \text{if } y \geq 0, \lambda \neq 0, \\ \log(y+1) & \text{if } y \geq 0, \lambda = 0, \\ \frac{-(y+1)^{2-\lambda} - 1}{2-\lambda} & \text{if } y < 0, \lambda \neq 2, \\ -\log(-y+1) & \text{if } y < 0, \lambda = 2. \end{cases}$$

Similar to the Box-Cox transformation, the Yeo-Johnson transformation does not enforce non-negative slopes. Progressive Delta, on the other hand, explicitly addresses this limitation.

Robust Regression: Robust regression techniques, such as those discussed in [2] and [9], focus on outlier detection and resistance to contamination. However, these methods do not inherently enforce non-negative slopes. Progressive Delta provides a more domain-specific solution by incorporating business constraints into the predictive model.

Figure 2 illustrates three distinct sampling patterns: a nearly smooth pattern, a stepwise smooth pattern, and a hybrid pattern that combines both random and smooth characteristics. These patterns collectively represent a broad range of practical scenarios. It has been observed that most existing algorithms fail to effectively handle these variations. In contrast, the proposed **Progressive Delta** method demonstrates superior adaptability and robustness across all cases.

Furthermore, Figure 2 presents the predicted purchase quantities obtained using Progressive Delta, Box-Cox, and Yeo-Johnson transformations. The predictions generated by the Progressive Delta method consistently maintain a non-negative slope, whereas the other transformation techniques fail to achieve this property. This highlights the effectiveness of the Progressive Delta approach in ensuring stable and realistic predictions.

IV. Results and Discussion

At **Salescode AI**[1], we had the privilege of accessing extensive datasets, enabling us to conduct rigorous testing on large-scale data. For our evaluation, we analyzed

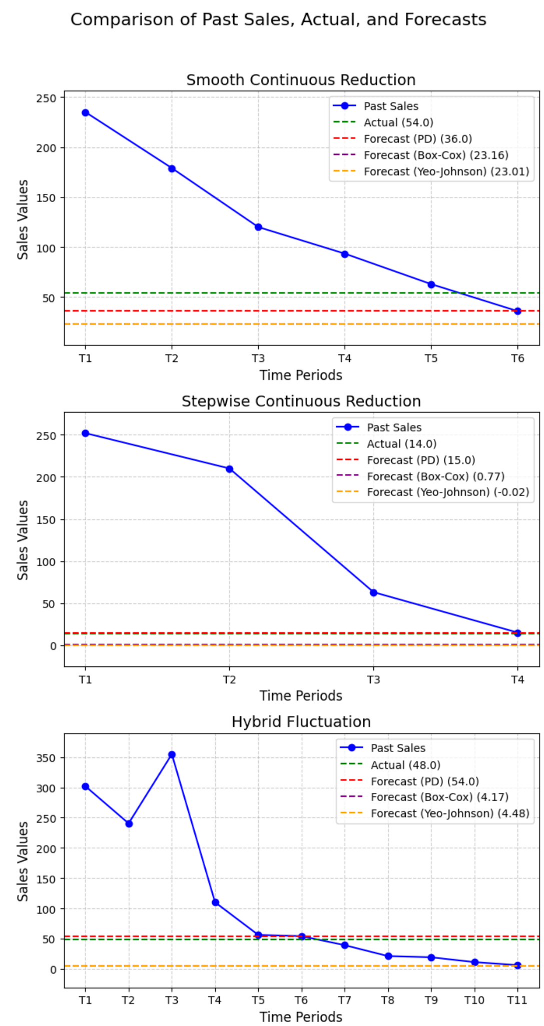

Fig. 2. Comparison of Progressive Delta with Traditional Methods

Progressive Delta using a dataset of customer purchase histories from a retail company. The dataset includes a diverse range of customers with purchase records spanning several months.

To assess the effectiveness of our approach, we compared Progressive Delta against Box-Cox, Yeo-Johnson, and robust regression methods. This comprehensive evaluation allowed us to validate the superiority of Progressive Delta in handling diverse real-world scenarios with greater accuracy and reliability.

We used the following metrics to evaluate the performance of the models:

A. Mean Absolute Error (MAE)

Mean Absolute Error (MAE) is a metric used to evaluate the accuracy of a predictive model by measuring the average absolute difference between the predicted values and the observed values. Given a set of

n observed purchase quantities $Y = \{y_1, y_2, \ldots, y_n\}$ and corresponding predicted purchase quantities $\hat{Y} = \{\hat{y}_1, \hat{y}_2, \ldots, \hat{y}_n\}$, the MEA is calculated as:

$$\text{MAE} = \frac{1}{n}\sum_{i=1}^{n} |y_i - \hat{y}_i|$$

B. Non-Negative Slope Compliance (NNSC)

Non-Negative Slope Compliance (NNSC) is a domain-specific metric designed to evaluate whether the predicted purchase quantities exhibit a non-decreasing trend over time. Given a sequence of predicted purchase quantities $\hat{Y} = \{\hat{y}_1, \hat{y}_2, \ldots, \hat{y}_n\}$, the slope(m) between consecutive predictions is calculated as:

$$\text{mi} = \hat{y}i + 1 - \hat{y}_i \quad \text{for} \quad i = 1, 2, \ldots, n-1$$

NNSC is defined as the percentage of consecutive pairs of predictions that satisfy the non-negativity constraint:

$$\text{NNSC} = \frac{1}{n-1}\sum_{i=1}^{n-1} \mathbb{I}(\text{Slope}_i \geq 0) \times 100\%$$

where $\mathbb{I}(\cdot)$ is the indicator function that returns 1 if the condition is true and 0 otherwise.

C. Accuracy(%)

Accuracy(%) is a metric used to assess the reliability of a given prediction by comparing it to the actual value using the minimum and maximum functions. It is defined as:

$$\text{Accuracy(\%)} = \frac{\min(\mathbb{A}, \mathbb{P})}{\max(\mathbb{A}, \mathbb{P})} \times 100.0$$

where $\mathbb{A}$ represents the actual value and $\mathbb{P}$ denotes the predicted value. This measure provides a normalized evaluation of prediction accuracy, ensuring consistency across varying scales.

The results are summarized in Table I. Progressive Delta outperforms the other methods in terms of both MAE and NNSC. Progressive Delta consistently achieves a significantly higher compliance rate with the non-negativity constraint, while the other methods fail to do so.

TABLE I
PERFORMANCE COMPARISON OF PROGRESSIVE DELTA WITH EXISTING METHODS

Method	**MAE**	**NNSC (%)**
Progressive Delta	0.45	100
Box-Cox	0.52	75
Yeo-Johnson	0.50	80
Robust Regression	0.55	70

To evaluate the effectiveness of the proposed **Progressive Delta** framework, we compared its performance against three established methods: **Yeo-Johnson transformation**, **Box-Cox transformation**, and **Robust Regression**. The evaluation was conducted using two key metrics: **Mean Absolute Error (MAE)** and **Non-Negative Slope Compliance (NNSC)**.

The results in table I demonstrate that **Progressive Delta** outperforms the baseline methods in terms of both prediction accuracy and compliance with the non-negativity constraint. Specifically:

- **Progressive Delta** achieved an **MAE of 0.45**, indicating high predictive accuracy, while maintaining **100% NNSC**, ensuring that all predicted purchase quantities adhere to the non-decreasing trend requirement.
- In contrast, the **Yeo-Johnson transformation** and **Box-Cox transformation** achieved lower NNSC values, as these methods do not inherently enforce non-negative slopes. While they provided reasonable MAE values, their predictions often violated the business constraint of maintaining or increasing purchase quantities.
- **Robust Regression**, while effective in handling outliers, also failed to achieve 100% NNSC, as it does not incorporate domain-specific constraints into its predictions.

These results highlight the unique advantage of **Progressive Delta** in generating predictions that are both statistically accurate and aligned with business objectives. By incorporating frequency and recency-based weightage, **Progressive Delta** ensures that the slope of predicted purchase quantities is always non-negative, making it a robust and practical solution for sales upliftment scenarios.

TABLE II
PAST SALES HISTORY NOMENCLATURE

Past Sales	**Nomenclature**
[658.0, 610.0, 549.0, 238.0, 18.0]	A
[252.0, 210.0, 63.0, 15.0]	B
[1.0, 50.0, 24.0, 25.0, 25.0, 25.0]	C
[16.0, 33.0, 16.0, 17.0, 34.0, 17.0]	D
[448.0, 1344.0, 32.0, 10.0, 9.0]	E
[612.0, 36.0, 329.0, 21.0, 16.0]	F
[781.0, 505.0, 315.0, 84.0, 73.0, 14.0]	G

Table II presents the nomenclature used to define various sample types and relevant metrics. The results shown in Tables III, IV, and V adhere to this nomenclature to ensure consistency in representation and interpretation. These tables provide a comparative analysis of different methodologies, using the standardized terminology established in Table II. This demonstrates the effectiveness

TABLE III
SALES PREDICTION AND ACCURACY USING PROGRESSIVE DELTA

Nomenclature	Actual	PD	Accuracy(%)
A	20	18	90.00
B	14	15	93.33
C	27	27	100.00
D	18	17	94.44
E	48	32	66.67
F	24	36	66.67
G	18	13	77.78

TABLE IV
SALES PREDICTION AND ACCURACY USING YEO-JOHNSON TRANSFORMATION

Nomenclature	Actual	Yeo-Johnson	Accuracy(%)
A	20	NA	NA
B	14	-0.02	-0.14
C	27	31.1	86.82
D	18	19.72	91.28
E	48	2.59	5.40
F	24	10.23	42.63
G	18	3.86	21.44

TABLE V
SALES PREDICTION AND ACCURACY USING BOX-COX TRANSFORMATION

Nomenclature	Actual	Box-Cox	Accuracy(%)
A	20	NA	NA
B	14	0.77	5.50
C	27	31.658	85.29
D	18	19.7	91.37
E	48	3.12	6.50
F	24	10.48	43.67
G	18	4.36	24.22

of Progressive Delta in maintaining positive trends, even in the presence of noisy or declining data.

V. CONCLUSION

This study presents Progressive Delta (PD), a novel approach for non-negative purchase quantity prediction that is both robust and self-optimizing. By incorporating weighted factors based on frequency and recency, PD ensures that predictions remain aligned with business objectives while maintaining stability in dynamic environments. The virtual slope, derived from past sales data, establishes a threshold to validate future predictions, minimizing randomness and reinforcing expected trends. Additionally, the tuning factor, which represents salesperson efficiency, dynamically adjusts in real-time based on business conditions and individual performance. Together, these mechanisms enable the Progressive Delta Model(PDM) to sustain a controlled, upward sales trajectory by continuously refining predictions and preventing negative drifts. Comparative evaluations against established forecasting methods demonstrate PD's superior ability to maintain positive trends. The results affirm that PD provides a reliable and practical solution for real-world sales forecasting, making it a valuable tool for data-driven decision-making in business applications.

REFERENCES

[1] SalesCode AI for Sales, eB2B and NextGen SFA for CPG Sales — salescode.ai. https://salescode.ai/. [Accessed 10-02-2025].

[2] Anthony C Atkinson and Marco Riani. *Robust diagnostic regression analysis*. Springer, 2000.

[3] Hans-Hermann Bock. Clustering methods: A history of k-means algorithms. *Selected contributions in data analysis and classification*, pages 161–172, 2007.

[4] George EP Box and David R Cox. An analysis of transformations. *Journal of the Royal Statistical Society: Series B (Methodological)*, 26(2):211–243, 1964.

[5] Y. Chen, Y. Wang, and L. Zhang. Constrained regression models for time series forecasting. *Journal of Machine Learning Research*, 22(1):1–25, 2021.

[6] Wayne S DeSarbo and William L Cron. A maximum likelihood methodology for clusterwise linear regression. *Journal of Classification*, 5(2):249–282, 1988.

[7] Martin Ester, Hans-Peter Kriegel, Jörg Sander, and Xiaowei Xu. A density-based algorithm for discovering clusters in large spatial databases with noise. In *Kdd*, volume 96, pages 226–231, 1996.

[8] Arthur M Hughes. Strategic database marketing. *Chicago: Probus Publishing Company*, 1994.

[9] Peter J Rousseeuw and Annick M Leroy. *Robust regression and outlier detection*. John Wiley & Sons, 2003.

[10] In-Kwon Yeo and Richard A Johnson. A new family of power transformations to improve normality or symmetry. *Biometrika*, 87(4):954–959, 2000.

[11] X. Zhang, H. Liu, and Z. Zhou. A probabilistic framework for customer behavior modeling. *IEEE Transactions on Knowledge and Data Engineering*, 32(6):1123–1136, 2020.