Proceedings of the 2025 Winter Simulation Conference
E. Azar, A. Djanatliev, A. Harper, C. Kogler, V. Ramamohan, A. Anagnostou, and S. J. E. Taylor, eds.

COLD-START FORECASTING OF NEW PRODUCT LIFE-CYCLES VIA CONDITIONAL DIFFUSION MODELS

Ruihan Zhou[1,2]

[1]Guanghua School of Management, Peking University, Beijing, CHINA
[2]Wuhan Institute of Artificial Intelligence, Peking University, Beijing, CHINA

ABSTRACT

Accurately forecasting the demand trajectories of newly launched products is a fundamental challenge due to sparse or nonexistent historical data. This paper introduces the Conditional Diffusion Life-cycle Forecaster (CDLF), a generative modeling framework designed for cold-start scenarios. CDLF integrates product attributes, reference trajectories of comparable products, and early post-launch sales signals into a unified conditional representation that guides a denoising diffusion process. By doing so, the framework produces realistic and uncertainty-aware forecasts even when sales data are absent or extremely limited. Empirical studies on the Intel microprocessor dataset show that CDLF outperforms classical diffusion models, Bayesian updating approaches, and state-of-the-art machine learning baselines. The results demonstrate that CDLF provides more accurate forecasts and calibrated uncertainty, highlighting its potential to support inventory planning and decision-making for new product launches.

1 INTRODUCTION

Forecasting the life-cycle demand trajectories of newly launched products is vital yet highly uncertain, as managers must plan inventory and production without historical data, while early forecasts also inform marketing decisions—such as pricing, promotion, and channel strategy—that shape long-term product adoption. The product life cycle typically evolves through five stages, but forecasting challenges are concentrated in the pre-launch and early post-launch phases. In the pre-launch phase, no sales data are available, forcing reliance on static product features and reference trajectories. In the early post-launch phase, limited and volatile signals emerge, but they remain insufficient for stable forecasts. Poor performance at these stages often leads to overstocks, stockouts, and high product failure rates (Lei et al. 2023; Ban et al. 2019; Hu et al. 2019). Traditional methods such as the Bass diffusion model (Bass 1969) impose rigid functional forms, while analogy-based methods or machine learning pipelines still struggle to capture multimodal uncertainty. To formalize the forecasting task, let $\mathbf{s} \in \mathbb{R}^p$ denote static product attributes and $\mathbf{x}_{1:T} = \{\mathbf{x}_1, \ldots, \mathbf{x}_T\}$ represent the life-cycle trajectory over T periods. The objective is to model the conditional distribution of future outcomes given evolving information sets $\mathscr{F}_t$: $q(\mathbf{x}_{t_0:T} \mid \mathscr{F}_{t_0-1}) = \prod_{t=t_0}^{T} q(\mathbf{x}_t \mid \mathscr{F}_{t-1})$, where $\mathscr{F}_0 = \sigma(\mathscr{X}, \mathbf{s})$ combines static attributes and complete trajectories from comparable products $\mathscr{X} = \{\mathbf{x}_{1:T}^{(k)}\}_{k=1}^{K}$. When $t_0 = 1$, no history exists and forecasts depend entirely on context. As observations $\mathbf{x}_{1:t_0-1}$ accumulate, $\mathscr{F}_t$ expands to incorporate them, enabling adaptive refinement. The key challenge lies in constructing a robust prior for $\mathbf{x}_1$, as early misspecification propagates errors across the entire trajectory.

2 METHOD

We formulate cold-start life-cycle forecasting as conditional generative modeling. Given static attributes $s \in \mathbb{R}^p$, a set of reference trajectories $\mathscr{X} = \{\mathbf{x}_{1:T}^{(k)}\}_{k=1}^{K}$, and an observed prefix $\mathbf{x}_{1:t_0-1}$, the task is to learn the conditional distribution $p_\theta(\mathbf{x}_{t_0:T} \mid \mathbf{s}, \mathscr{X}, \mathbf{x}_{1:t_0-1})$. This formulation simultaneously covers pre-launch ($t_0 = 1$) when no sales data are observed, and early post-launch when only sparse prefixes are available.

Our proposed Conditional Diffusion Life-cycle Forecaster (CDLF) integrates heterogeneous information into a unified embedding and then generates trajectories via a denoising diffusion process. Static features are encoded as a fixed-length vector $\mathbf{h}_s$. Reference trajectories are weighted according to a softmax similarity and aggregated into $\mathbf{h}_X$ through a recurrent encoder. Early sales are autoregressively encoded to hidden states $\mathbf{h}_t = \mathrm{RNN}_\theta(\mathrm{concat}(\mathbf{x}_{t-1}, \mathbf{h}_s, \mathbf{h}_X), \mathbf{h}_{t-1})$, ensuring that static and reference information is preserved across time.

On top of the context encoder, CDLF applies a conditional denoising diffusion model. In the forward process, each clean variable $\mathbf{x}_t^0$ is gradually perturbed by Gaussian noise according to a variance schedule $\{\beta_n\}_{n=1}^N$, where $\beta_n \in (0,1)$ controls the noise magnitude at step n. The cumulative product $\alpha_n = 1-\beta_n$ and $\bar{\alpha}_n = \prod_{i=1}^n \alpha_i$ summarize the noise schedule. This yields the closed-form marginal distribution $q(\mathbf{x}_t^n \mid \mathbf{x}_t^0) = \mathcal{N}(\sqrt{\bar{\alpha}_n}\,\mathbf{x}_t^0, (1-\bar{\alpha}_n)\mathbf{I})$. The reverse process then reconstructs clean data from noise, conditioned on the hidden state h_{t-1}. At each step it samples from $p_\theta(\mathbf{x}_t^{n-1} \mid \mathbf{x}_t^n, \mathbf{h}_{t-1}) = \mathcal{N}(\frac{1}{\sqrt{\alpha_n}}(\mathbf{x}_t^n - \frac{\beta_n}{\sqrt{1-\alpha_n}} s_\theta(\mathbf{x}_t^n, \mathbf{h}_{t-1}, n)),\ \sigma_n^2 I)$, where the score network s_θ predicts the injected noise. Training reduces to minimizing the expected squared error between predicted and true noise, while inference proceeds autoregressively: the model starts from pure noise, samples $\mathbf{x}_1^0$ given context $(\mathbf{s}, \mathbf{X})$, updates its hidden state, and continues step by step. When early sales are available, they are injected directly into the hidden state, enabling refined forecasts without retraining.

3 EMPIRICAL STUDY AND DISCUSSION

We evaluate CDLF on the Intel microprocessor dataset (Manary and Willems 2022) and benchmark it against several representative forecasting approaches from both classical and modern paradigms. Forecasting is tested under pre-launch, where predictions rely solely on product attributes and reference trajectories, and early post-launch, where limited observed sales are incorporated. Using mean absolute error (MAE) and continuous ranked probability score (CRPS), we find that CDLF consistently achieves the best performance across both horizons. In the pre-launch setting, CDLF reduces MAE by more than 15% relative to the strongest benchmark and improves CRPS by nearly 20%, highlighting its ability to generate reliable priors without sales data. In the early post-launch setting, as a few weeks of observations become available, CDLF further improves accuracy, lowering MAE by over 20% and delivering sharper and better-calibrated predictive intervals than competing models.

Our findings suggest that conditional diffusion models provide a promising direction for cold-start forecasting. By combining available contextual information, CDLF generates life-cycle forecasts that are both accurate and uncertainty-aware, even under extreme data scarcity. The model's ability to seamlessly transition from pre-launch to post-launch settings indicates practical value for managers facing high-stakes decisions with limited information. Beyond product demand, this approach highlights broader opportunities for applying diffusion-based generative models to other domains where cold-start and uncertainty dominate.

REFERENCES

Ban, G.-Y., J. Gallien, and A. J. Mersereau. 2019. "Dynamic Procurement of New Products with Covariate Information: The Residual Tree Method". *Manufacturing & Service Operations Management* 21(4):798–815 https://doi.org/10.1287/msom.2018.0725.

Bass, F. M. 1969. "A New Product Growth for Model Consumer Durables". *Management Science* 15(5):215–227 https://doi.org/10.1287/mnsc.15.5.215.

Hu, K., J. Acimovic, F. Erize, D. J. Thomas, and J. A. Van Mieghem. 2019. "Forecasting New Product Life Cycle Curves: Practical Approach and Empirical Analysis". *Manufacturing & Service Operations Management* 21(1):66–85 https://doi.org/10.1287/msom.2017.0691.

Lei, D., H. Hu, D. Geng, J. Zhang, Y. Qi, S. Liu *et al.* 2023. "New product life cycle curve modeling and forecasting with product attributes and promotion: A Bayesian functional approach". *Production and Operations Management* 32(2):655–673 https://doi.org/https://doi.org/10.1111/poms.13892.

Manary, M. P., and S. P. Willems. 2022. "Data Set: 187 Weeks of Customer Forecasts and Orders for Microprocessors from Intel Corporation". *Manufacturing & Service Operations Management* 24(1):682–689.