Forecasting Print Demand with Machine Learning at HP Inc.

Harshvardhan[1], Chuanren Liu[1], Cara Curtland[2], Adam Ghozeil[2]

[1]University of Tennessee, Knoxville; [2]HP Inc.

2025 Impact of Forecasting in Practice Award

Foresight Practitioner Conference 2025 | UNC Charlotte

International Institute of Forecasting

Outline

1. Background and Contributions
2. Methodology and Details
3. Business Process and Operationalization
4. Operational Benefits and Impact
5. Concluding Remarks

HP Annual Report, 2024

Integrated Demand Forecasting System at Scale

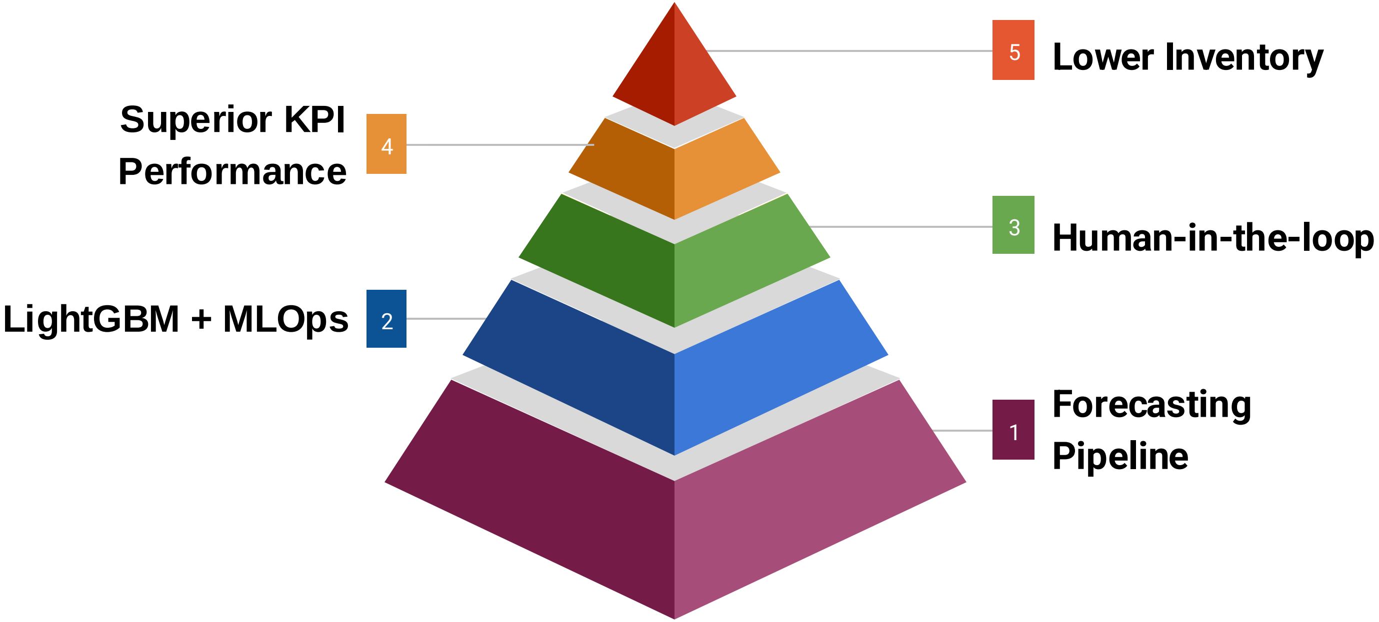

Background and Contributions

Problem

Accurate print demand forecasts increase product availability and profitability

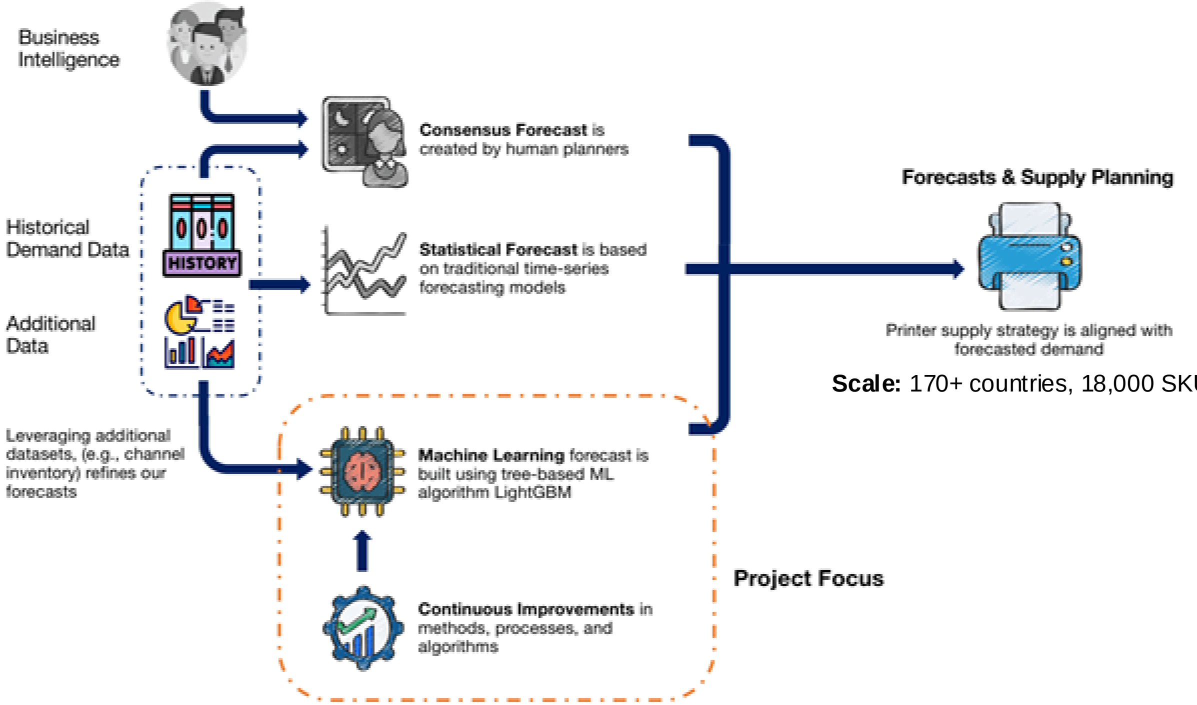

Scale: 170+ countries, 18,000 SKUs

Forecast is used by multiple downstream processes

- Inventory Sizing
- Purchasing
- Capacity sizing
- Contract negotiations
- Portfolio Planning

Challenges

Print demand forecasting is especially hard due to scale and complexity of business

Business Complexity

- **Global market dynamics** vary across regions and product lines
- **Complex interplay** between economic conditions, seasonality, and local trends
- **Niche products** exhibit intermittent demand patterns

Technical Hurdles

- **High data variability** and inconsistency across markets
- **Automated preprocessing** and validation needed at scale
- **Model accuracy and generalization** capability

Organizational Transformation

- **Data-driven forecasting** from human judgement
- Trust in **ML-based predictions**
- **Change management** and process integration

Related Literature

Table 2.1: Summary of related research papers.

Paper	Input	Model	Evaluation Metric
Dodin et al. (2023)	Lagged demands, demand statistics, seasonality components, region and month index, average age of shipped products	Improved LightGBM, Elastic Net	RMSSE
Gardner (1990)	Lagged demand	Exponential-smoothing Model (ETS)	Investment and Delay Time
Makridakis et al. (2018)	M-3 data	MLP, BNN, RBF, GRNN, KNN, CART, SVR, GP, RNN, LSTM, SES, ETS	sMAPE, MASE
Qi et al. (2023)	Lagged demand, inventory	End-to-end Model (Dynamic Programming, RNN, MLP)	Stockout rate, turnover rate, total inventory management, holding, and stockout costs
Deng et al. (2023)	Lagged demand, inventory, among others	DeepAR, N-BEATS, Prophet	WMAPE
Sagaert et al. (2018)	Lagged demand, macroeconomic indicators	LASSO Regression	MAPE

Research Gap:

- ❑ Limited studies address the practical challenges of **large-scale** demand forecasting implementation
- ❑ Critical role of **human expertise** working with **machine learning** is underexplored

We bridge current practices with analytical innovations, enabling integration of new advances while preserving valuable existing methodologies

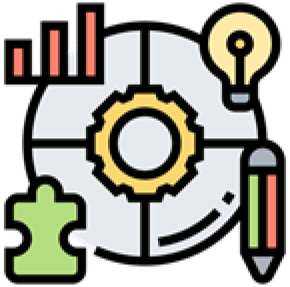

Methodology and Details

Solution

We create enterprise-level forecasting pipeline to go from data to decisions

Prepare Data

Historical **Demand**

Product Hierarchy

Location Hierarchy

Channel Data

Product **Lifecycle**

Create Model

LightGBM: tree-based Learner

Dynamic features

Hyperparameters Tuning

Feature Engineering

Run Process

Train model to create predictions

Predict future demand

Back-test models for performance

Store data for future

Publish Results

Create and Share **Forecast Accuracy Dashboard**

Disaggregate demand for each SKU, every GOC

Publish forecasts into SAP planning tool

Problem Formulation

We now describe our problem setting for print demand forecasting

- Our work addresses the problem of predicting demand for a product p in a specific country c at time t.
- $X_{t,c,p}$: Input features for period t, product p in country c
- $y_{t,c,p}$: Actual demand for period t, product p in country c
- Training dataset:
  - $D = \{(X_{t,c,p}, y_{t+1,c,p}) : \forall c, p, t_{first} \leq t \leq t_{now}\}$
  - where t_{first} is first time-period
- Loss function:
  - $\ell(f|D) = \sqrt{E_{X,y \in D}(f(X) - y)^2}$ (RMSE)

Problem Scale:

- ❑ 170+ countries
- ❑ 18,000 SKUs

THE UNIVERSITY OF TENNESSEE KNOXVILLE

Iterative Forecasting Algorithm

LightGBM algorithm used as predictive engine; it is adaptable

Algorithm 1 Enhanced training and forecasting algorithm with LightGBM

1: **Preprocess the data:** Data cleaning and feature engineering.
2: **Determine optimal hyperparameters:** Use grid search or random search for the LightGBM model.
3: **Initialize forecast horizon** T (e.g., 7).
4: **for** t_α in $(t_{\text{first}} : t_{\text{now}})$ **do**
5: Create the training data:

$$D_\alpha = \{(X_{t,c,p}, y_{t,c,p}) : \forall c, p, t_{\text{first}} \leq t \leq t_\alpha\}$$

6: Perform time-series cross-validation on D_α and train the LightGBM model $f(\cdot)$ with optimal hyperparameters, minimizing loss (RMSE):

$$\ell(f|D) = \sqrt{\mathbb{E}_{X,y\in D}\,(f(X) - y)^2}$$

7: With the fitted model, create T forecasts for $t_\alpha + 1$ to $t_\alpha + T$:

$$F^T_{t_\alpha,c,p} = \left(f(\hat{X}_{t_\alpha+1,c,p}), f(\hat{X}_{t_\alpha+2,c,p}), \cdots, f(\hat{X}_{t_\alpha+T,c,p})\right)$$

8: Update the LightGBM model incrementally by warm starting from last month's best results if possible, or retrain it from scratch.
9: **end for**
10: **Perform Backtesting:** Apply the trained model to a historical dataset $D_{\text{historical}}$ to simulate past predictions. Evaluate its performance using appropriate metrics (e.g., RMSE, MAE).
11: **Store Forecasts:** Save the generated forecasts $F^T_{t_\alpha,c,p}$ to a dedicated database or file storage for future evaluation, comparison, or direct usage.
12: **Log Model:** Serialize the LightGBM model, hyperparameters, and performance metrics for future reference or retraining using MLFlow.

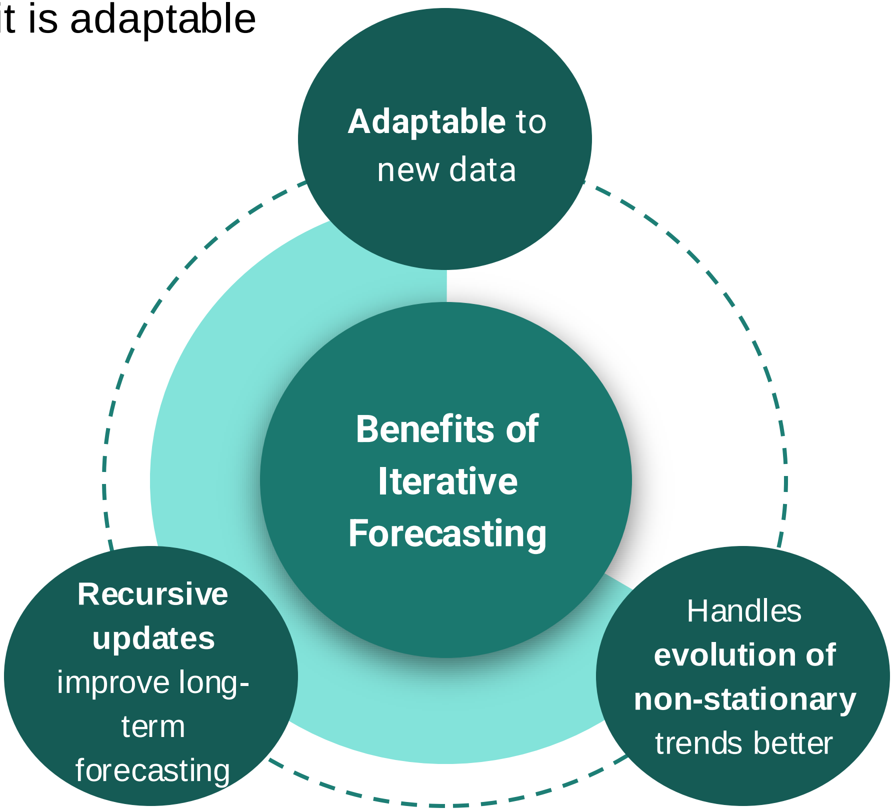

THE UNIVERSITY OF TENNESSEE KNOXVILLE

Model Input Features

Feature Name	Description	Granularity	Utility for Forecasting
Lagged Demand	Size of demand from previous m months, m varies per product group	Month (t)	Captures influence of past on future trends
Rolling Demand Features	Statistics of demands within n-month rolling window (mean, coefficient of variation, outliers)	Month (t)	Assesses recent trend, variability
Product-based Statistics	Mean and coefficient of variation of lagged demand and rolling features, per product category	SKU (p)	Specific trends in product categories
Geography-based Statistics	Mean and coefficient of variation of lagged demand and rolling features, per country	Country (c)	Location-specific trends
Seasonal Fluctuation	Binary indicator for each fiscal quarter and integer month within a quarter	Month (t)	Captures seasonal effects
Product Life Cycle	Proportion of product life cycle left, calculated as $(M - m)/M$	SKU, Country (p, c)	Stage of the product in its life cycle
Channel Inventory	Inventory reported by distribution channel partners	SKU, Country, Month (p, c, t)	Indicates potential reordering
Sell-through	Sales to distribution channel partners	SKU, Country, Month (p, c, t)	Reflection of downstream demand (to channel partners' customers)

100+ Input Features

- ❑1 thru 15 month lagged demand for trends
- ❑Product Life Cycle for growth and promotions
- ❑Channel data provides early demand sensing

Feature Selection and Engineering

Methods

- ❑Fast AI Method (Howard 2019)
- ❑QPFS (Rodriguez-Lujan et al., 2010)

Trade-offs

- ❑Stability vs Guarantee → Explainability
- ❑Speed of execution

Hyperparameter Tuning

- ❑Hyperopt Tree-Structured Prazen Estimator (TPE)

Project Management with MLOps

Handling data, experiments and results at scale required creation of new systems

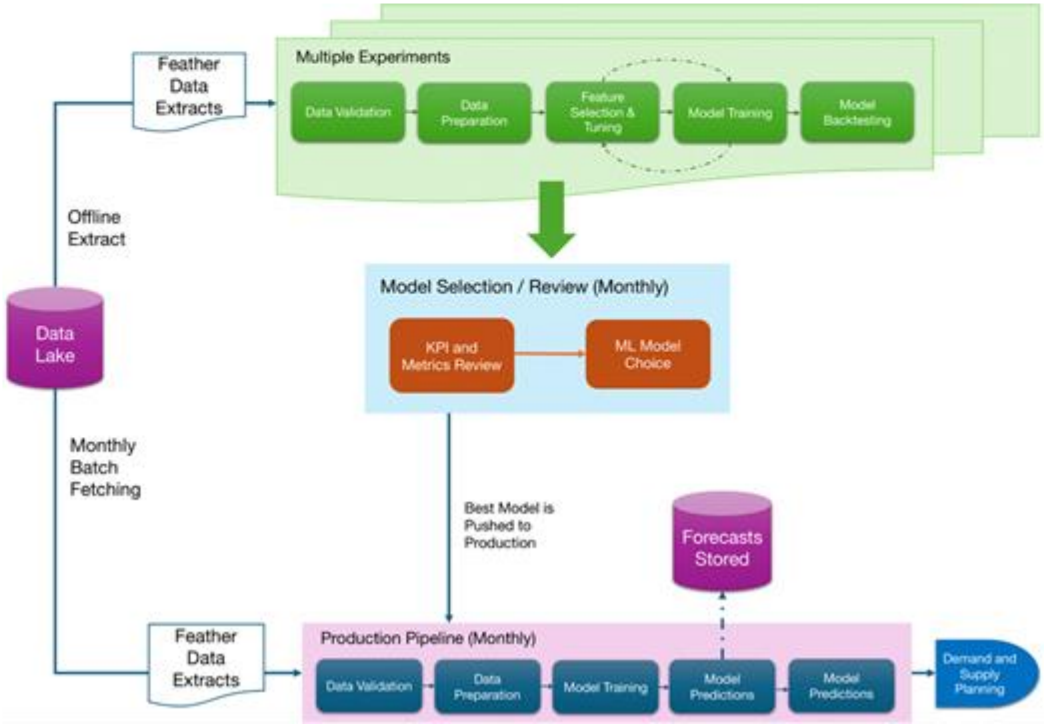

- MLFlow: Experimentation
- Jupyter and Papermill: Documentation
- Model Serialization for Warm-starting
- Efficient Data Storage
- Rapid Testing with FLAML

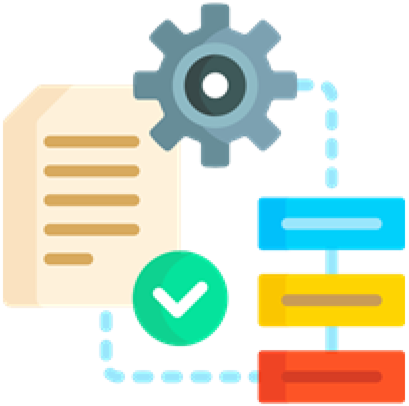

Business Process and Operationalization

Implementation Journey

Increasing business value driven by deeper integration of ML forecasts into decision making

Business KPI Dashboard	ML Forecast Pilot	ML Forecast Adoption	SKU-level ML Forecast	Automated Ensembling
Single integrated KPI dashboard for the entire Print Business	ML visibility in dashboard led to **directional guidance** in forecast Bias reduction	**Manual ML Forecast use based on superior FVA** performance relative to Statistical Forecast	**Full integration into decision making pipeline** within Integrated Business Planning (SAP)	**Auto-Ensembling of ML and Statistical Forecasts** passed as Analytical Forecast to planners

- Executive KPI alignment
- Monthly review process
- Forecast Value Add (FVA) informs decisions

Dashboard, ML Forecast, and Pipeline integration increasingly empower Human-in-loop Decisions

THE UNIVERSITY OF TENNESSEE KNOXVILLE

Human-in-the-loop Ensembling

Final forecast is optimal balance of machine precision and human insight

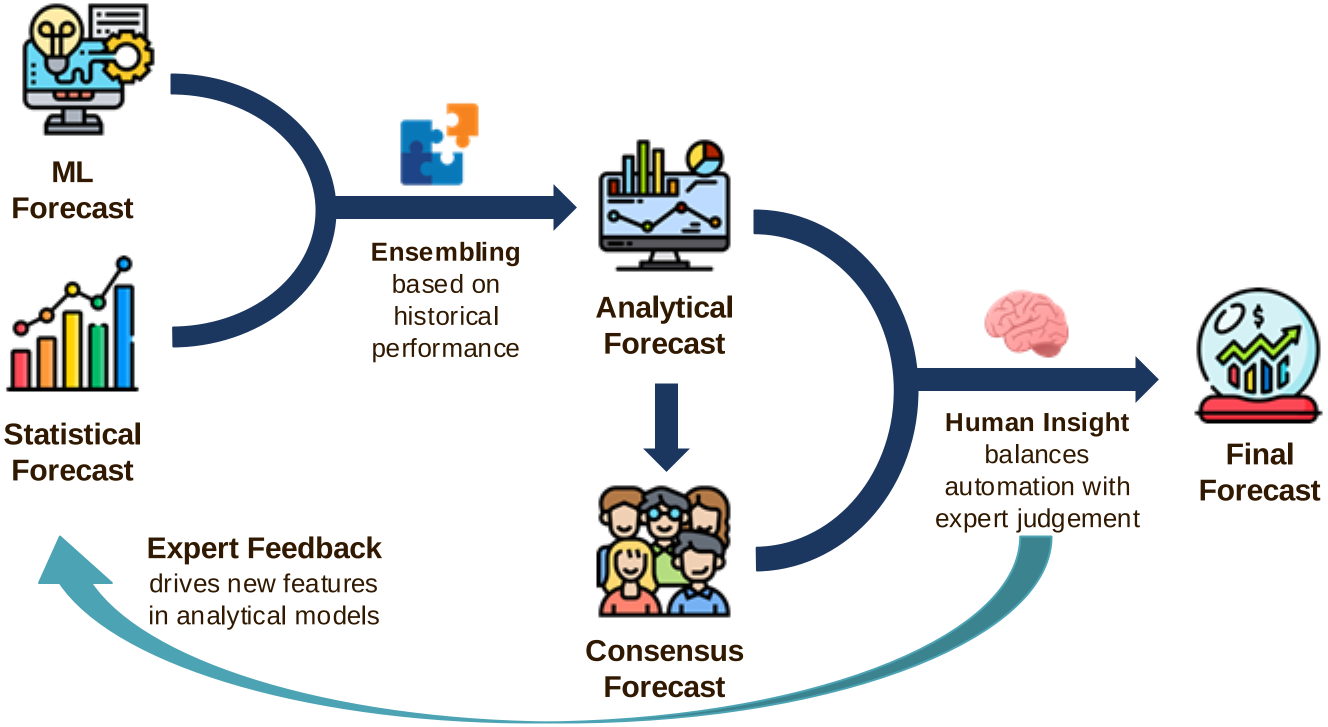

Forecast Value Add Drives Model Choice

- Analytical
- Consensus
- Hybrid

Closed-loop feedback guides new feature development for ML and Statistical models

Operational Benefits and Impact

Business Impact: 28% Inventory Reduction

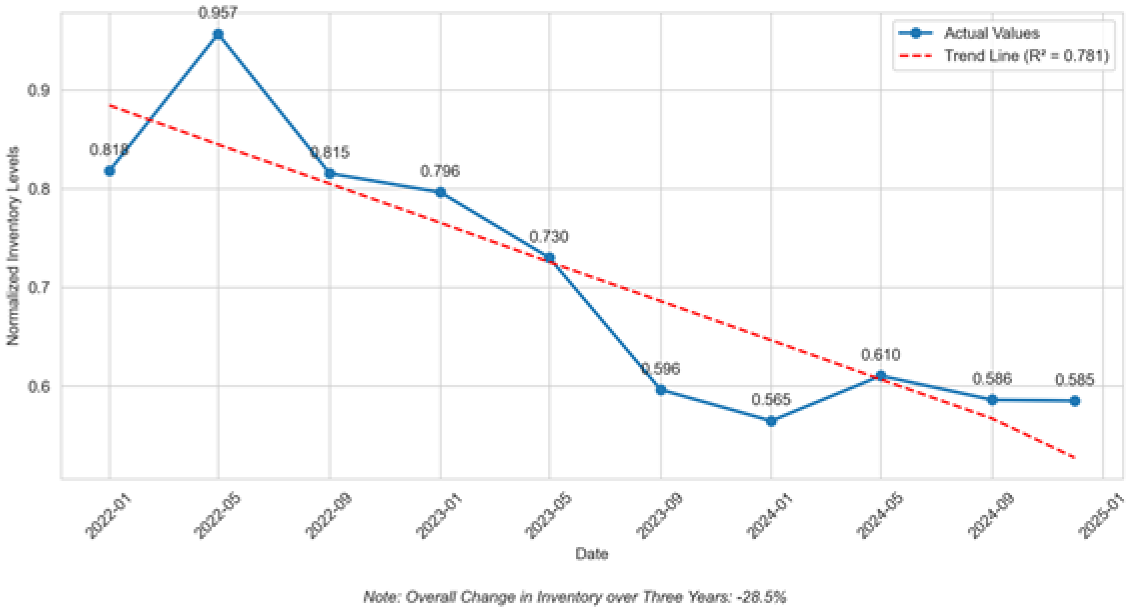

Values are normalized relative to peak inventory level (1.0 represents the maximum monthly inventory across all periods) to protect sensitive data while preserving trend patterns

Sustained inventory decline demonstrates the effectiveness of our integrated approach

- ❑ 28.5% 3-yr reduction
- ❑ R^2 = 0.781
- ❑ Customer Service Levels maintained

Results for a select business segment of HP Print (1,484 products)

Forecast Accuracy Metrics

RMSE, Bias and wMAPE

Bias measures the weighted percentage error in forecasts, signified by a positive or negative value indicating over or underforecasting, respectively. Bias is calculated using the formula:

$$Bias = \sum_{i=1}^{n} \frac{y_i - \hat{y}_i}{y_i}.$$

wMAPE represents the weighted mean of absolute percentage errors, a metric easily understood even by non-technical stakeholders as percentage deviation from actuals. It is expressed as:

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

RMSE (Root Mean Squared Error) is defined as:

$$RMSE = \sqrt{\frac{\sum_{i=1}^{n} (y_i - \hat{y}_i)^2}{n}}.$$

Bias:

- ❑ Easiest to address
- ❑ Consistently positive = overforecasting
- ❑ Consistently negative = underforecasting

wMAPE:

- ❑ Percentage error; easy to interpret
- ❑ Business KPI

RMSE:

- ❑ Used for ML model training
- ❑ Symmetric and always positive
- ❑ Continuously differentiable
- ❑ Sensitive to outliers

ML Forecast Consistently Improves Business KPI

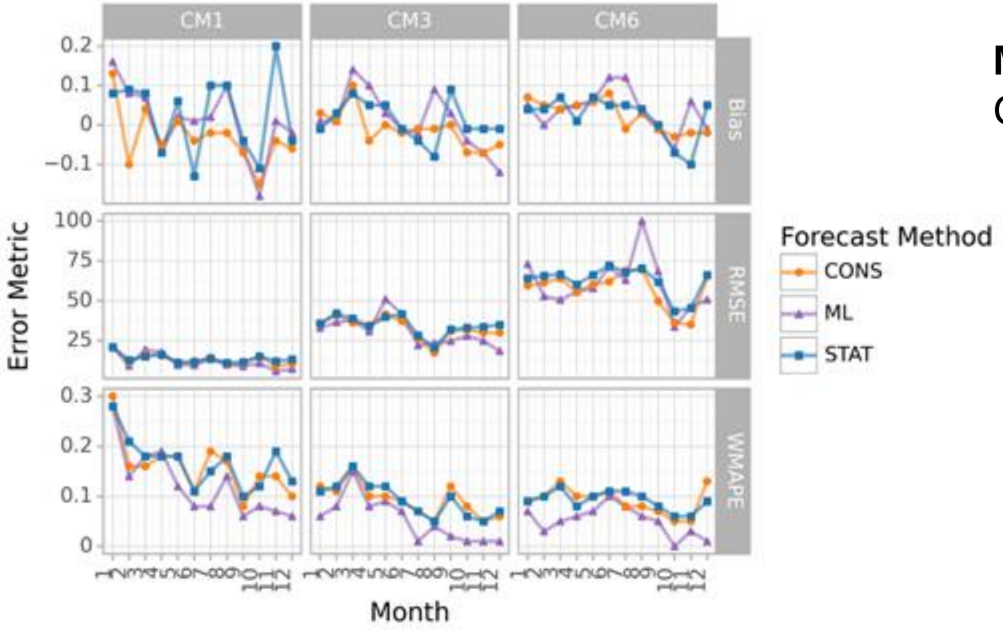

ML model wMAPE outperforms alternatives
Comparable performance in other metrics

Table 2.4: Comparative statistical analysis of forecasting accuracy metrics—namely Bias, wMAPE, and RMSE—is presented for consensus (CONS), machine learning (ML), and statistical (STAT) methods across cumulative forecast horizons (CM1, CM3, CM6). The table delineates t-statistics along with the corresponding p-values (enclosed in brackets) to facilitate a comprehensive evaluation.

Cumulative	Comparison	Bias	RMSE	WMAPE
CM1	CONS vs ML	-1.295 (0.209)	0.716 (0.482)	1.421 (0.169)
	STAT vs ML	0.385 (0.704)	1.128 (0.272)	1.832 (0.080)
CM3	CONS vs ML	-0.929 (0.363)	0.518 (0.610)	2.528 (0.019)
	STAT vs ML	-0.065 (0.949)	1.089 (0.288)	2.541 (0.019)
CM6	CONS vs ML	-0.701 (0.490)	-0.507 (0.617)	3.526 (0.002)
	STAT vs ML	-0.767 (0.451)	0.399 (0.694)	4.074 (0.001)

Results for a select business segment of HP Print (1,484 products)

Comparing Performance between Methods

No model is consistently better or worse than other; which is why human expertise is valued

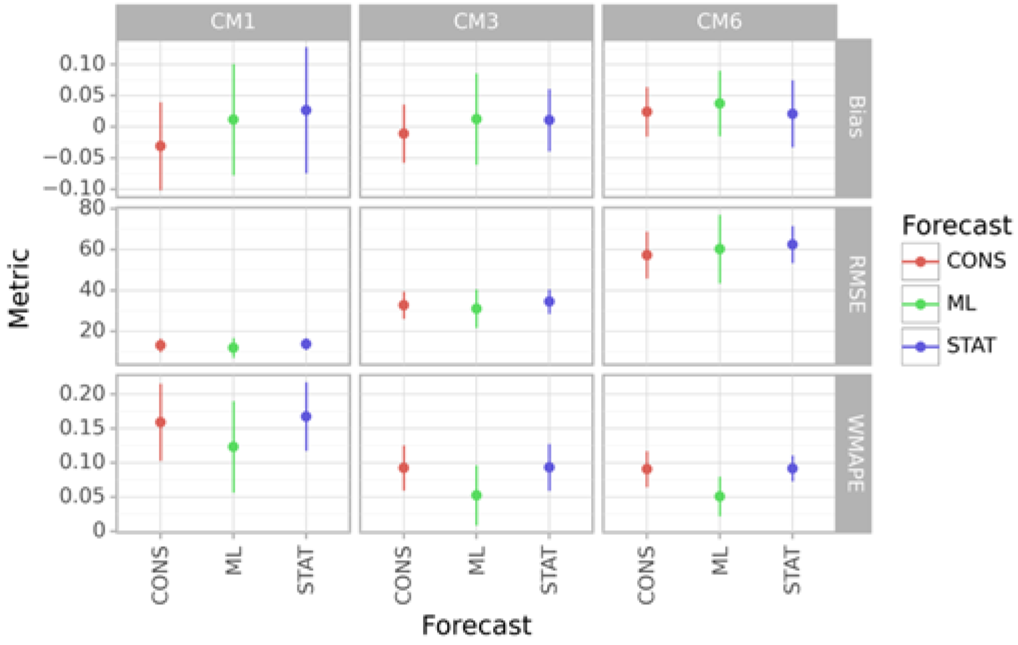

Mean as dot and one standard deviation as whiskers.
Lower is better.

ML model has comparable accuracy (and sometimes better) vs. other methods

Table 2.3: Forecasting Accuracy Metrics. The table shows bias, RMSE, and wMAPE for cumulative forecast horizons (CM1, CM3, CM6) with Mean (Standard Deviation).

Model	CM1			CM3			CM6		
Metric	Bias	RMSE	wMAPE	Bias	RMSE	wMAPE	Bias	RMSE	wMAPE
Value (SD)	-3.08% (7.05%)	13.09 (3.38)	15.92% (5.62%)	-1.08% (4.68%)	32.76 (6.66)	9.25% (3.28%)	2.42% (3.96%)	57.29 (11.51)	9.08% (2.64%)
	1.17% (8.92%)	11.87 (4.87)	12.33% (6.69%)	1.25% (7.34%)	31.03 (9.43)	5.25% (4.39%)	3.75% (5.26%)	60.28 (16.83)	5.08% (2.91%)
	2.67% (10.14%)	13.71 (2.89)	16.75% (4.99%)	1.08% (5.00%)	34.55 (6.03)	9.33% (3.42%)	2.08% (5.38%)	62.47 (9.08)	9.17% (1.90%)

Results for a select business segment of HP Print (1,484 products)

THE UNIVERSITY OF TENNESSEE KNOXVILLE

Concluding Remarks

Integrated Demand Forecasting System at Scale

Proven Business Impact and Extensible

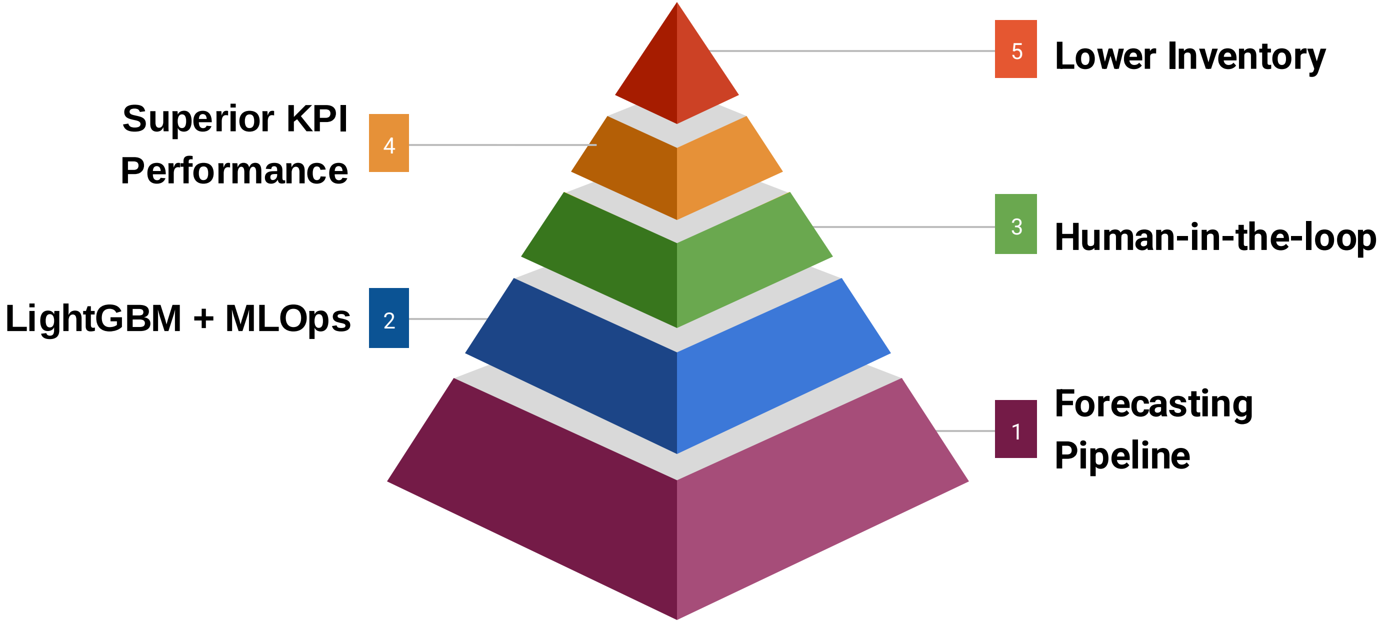

Framework Enables

New Models & Features
Process & architecture enables seamless upgrades

Prescriptive analytics
Refinement from human-in-the loop supports new changes

Expansion across HP
Forecasts beyond Print demand

Thank you!

Any questions?

Harshvardhan, *harshvar@utk.edu*
Chuanren Liu, *cliu89@utk.edu*
Cara Curtland, *cara.curtland@hp.com*
Adam Ghozeil, *adam.ghozeil@hp.com*