AI FOR RETAIL AND ECOMMERCE DEMAND PREDICTION

- Mithun Thakkar (ID: 300681732)
- Shehmeet Kaur (ID:300684888)
- Suhasini Cherukuri(ID: 300680883)
- Aayush Meena (ID: 300684718)
- Jeevan Bhusal (ID :300674411)
- Jude Zhang (ID: 300664488)
- Anshu Maskey (ID: 300680279)

INTRODUCTION

WHY DEMAND FORECASTING MATTERS

- Critical for supply chain efficiency, inventory management, and profitability
- Impacts customer satisfaction and operational resilience
- Evolving with big data, AI, and dynamic market conditions

WHAT TO EXPECT

- Diverse methods: Statistical models, machine learning, deep learning, and graph neural networks
- Data sources: Sales, weather, macroeconomic indicators, consumer sentiment
- Applications: Perishable goods, e-commerce, real-time demand sensing
- Goal: Understand strengths, limitations, and future directions

COMMON THREADS IN MODERN FORECASTING

- Integration of exogenous factors (e.g., weather, COVID-19, CPI)
- Hybrid models combining temporal and covariate data
- Emerging tech: AI, big data, and spatial analytics
- Focus on flexibility for short- and long-term predictions

DEMAND MODELS FOR SUPERMARKET DEMAND FORECASTING

TRADITIONAL STATISTICAL MODELS

- Model-based approaches depend on theoretical assumptions.
- Example: Poisson-type processes.

  1) Purchases happen randomly at a certain rate.

  2) Interarrival times follow an exponential distribution.

- This method provides clear and mathematically manageable results.

NEGATIVE BINOMIAL DISTRIBUTION (NBD)

- Also called the Gamma-Poisson distribution.
- Why NBD?
  - It extends the Poisson by **adding customer differences**.
  - Each customer has their **own buying rate**, which is Gamma-distributed.
  - Purchases then follow a Poisson process.
- This model captures more **variability** in demand compared to the pure Poisson.

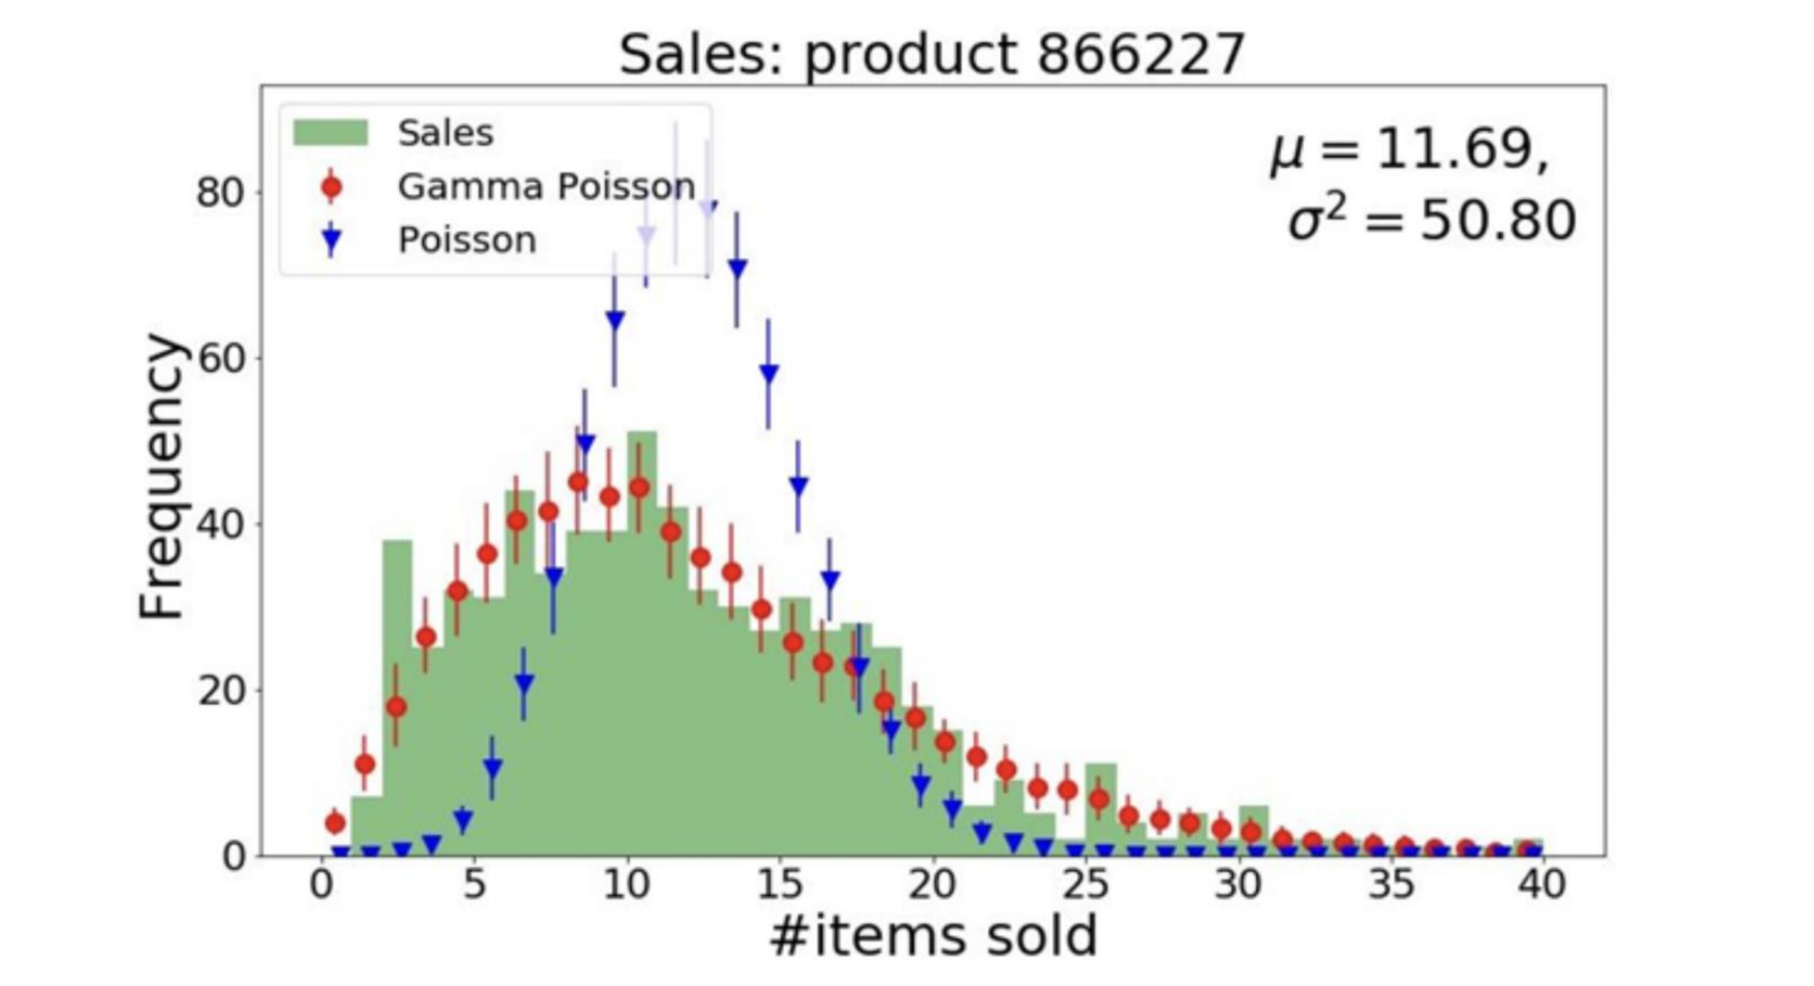

Count data for an individual product well described by a NBD model but overdispersed for a Poisson model

LIMITATIONS OF NBD

- Fails to capture **scheduled shopping patterns**
  - e.g., peaks at 7-day, 14-day intervals (weekly routines)
- Struggles with **different product categories**:
  - Perishables (milk, deli foods) vs. durable goods (cleaning items)

STRUGGLES WITH DIFFERENT PRODUCT CATEGORIES

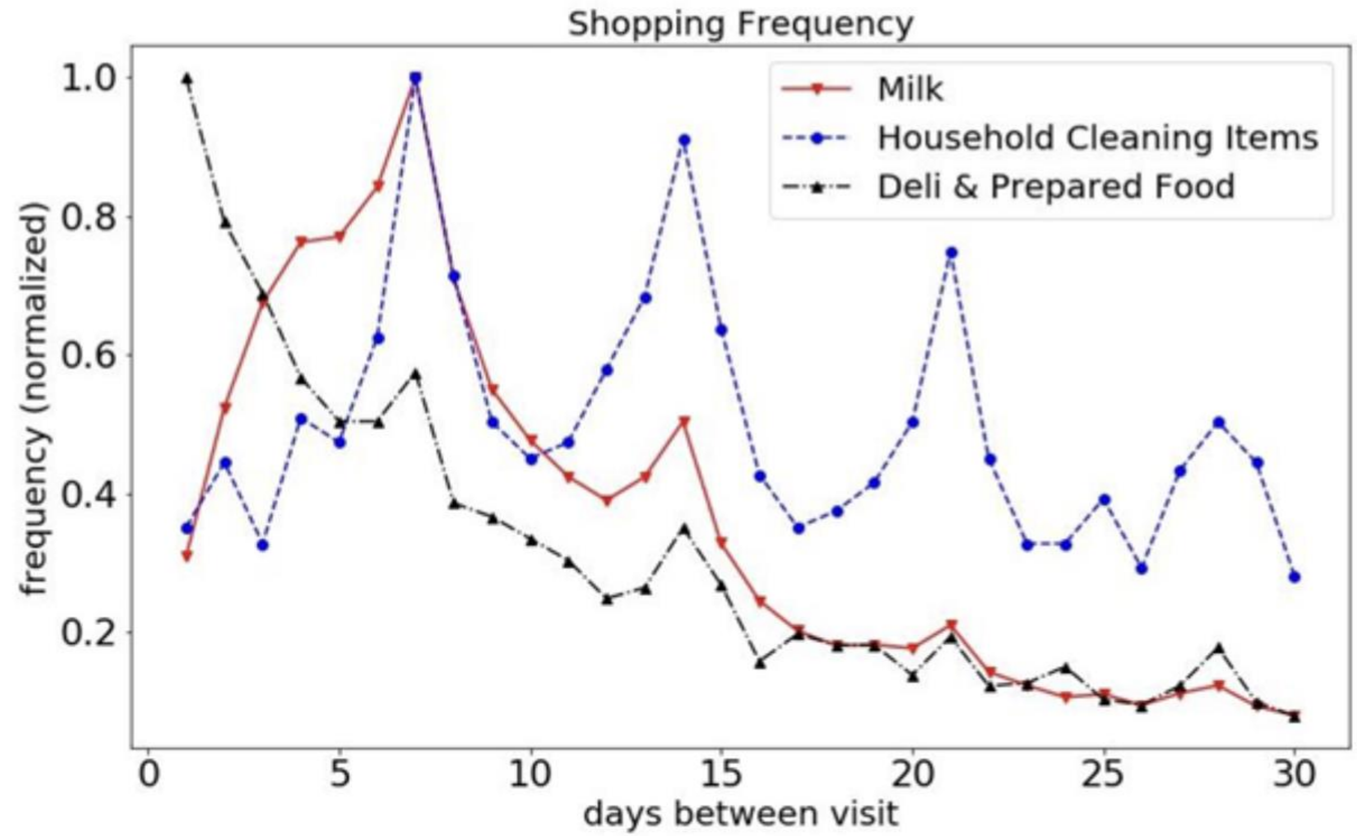

Inter-arrival times for different Product sub-groups for individual households.

TRANSITION TO ADVANCED APPROACHES

- Traditional models (like NBD) provide baseline insights
- But… limited in handling:
  - Complex seasonality
  - Multi-product interactions
  - Customer-specific behaviours
- We need more powerful models, such as **Deep Learning**.

PREDICTIVE ANALYTICS FOR DEMAND FORECASTING A DEEP LEARNING-BASED DECISION SUPPORT SYSTEM

PREDICTIVE ANALYTICS FOR DEMAND FORECASTING A DEEP LEARNING-BASED DECISION SUPPORT SYSTEM

- Introduces a novel deep learning-based ensemble model for demand forecasting that combines
  - a sequence modeling method (LSTM) - to capture temporal(linear and nonlinear)
  - a machine learning method (Random Forest) to capturecovariate-based variations in demand data.
- Also introduces a new demand sensing algorithm for demand forecasting in real-time

PREDICTIVE ANALYTICS FOR DEMAND FORECASTING A DEEP LEARNING-BASED DECISION SUPPORT SYSTEM

The model was tested on a large dataset of 4,235 demand series from packaged food products

- It incorporated both structured and unstructured data such as:
  - point-of-sales
  - Promotions
  - Weather
  - economic indicators, and
  - internet media.
- The proposed method outperformed benchmark models across multiple error metrics (meanerror, mean absolute error, mean squared error).

MODELLING

- Dimensionality reduction using PCA
- Ensemble Models:
  - LSTM Networks - Used to model the univariate sales time-series data
  - Random Forest (RF) - Used to model the multivariate relationships between the sales data and the principal components
- Prediction Aggregation
  - The predictions from the LSTM and the Random Forest are combined into a single, final forecast using Genetic Algorithm.
  - a Genetic Algorithm (GA) is used to optimize the weights assigned to each model's prediction.
  - Fitness Function: minimize the Mean Absolute Error

FORECASTING FOR DIFFERENT TIME HORIZONS

- The main model generates weekly forecasts
- Long-term forecasts:
  - Weekly data is aggregated into lower-frequency data.
  - The proposed ensemble model is applied to this aggregated data to generate separate forecasts for these longer horizons.
  - A hierarchical reconciliation algorithm is then used to ensure these forecasts are "coherent" (e.g., the monthly forecasts sum correctly to the quarterly forecasts).
- Short-term forecasts
  - To generate daily forecasts from the weekly forecasts, a temporal disaggregation (TD) method is used.
  - This algorithm leverages real-time, high-frequency daily data as an indicator to "disaggregate" the low-frequency weekly forecast into a high-frequency daily sales forecast.

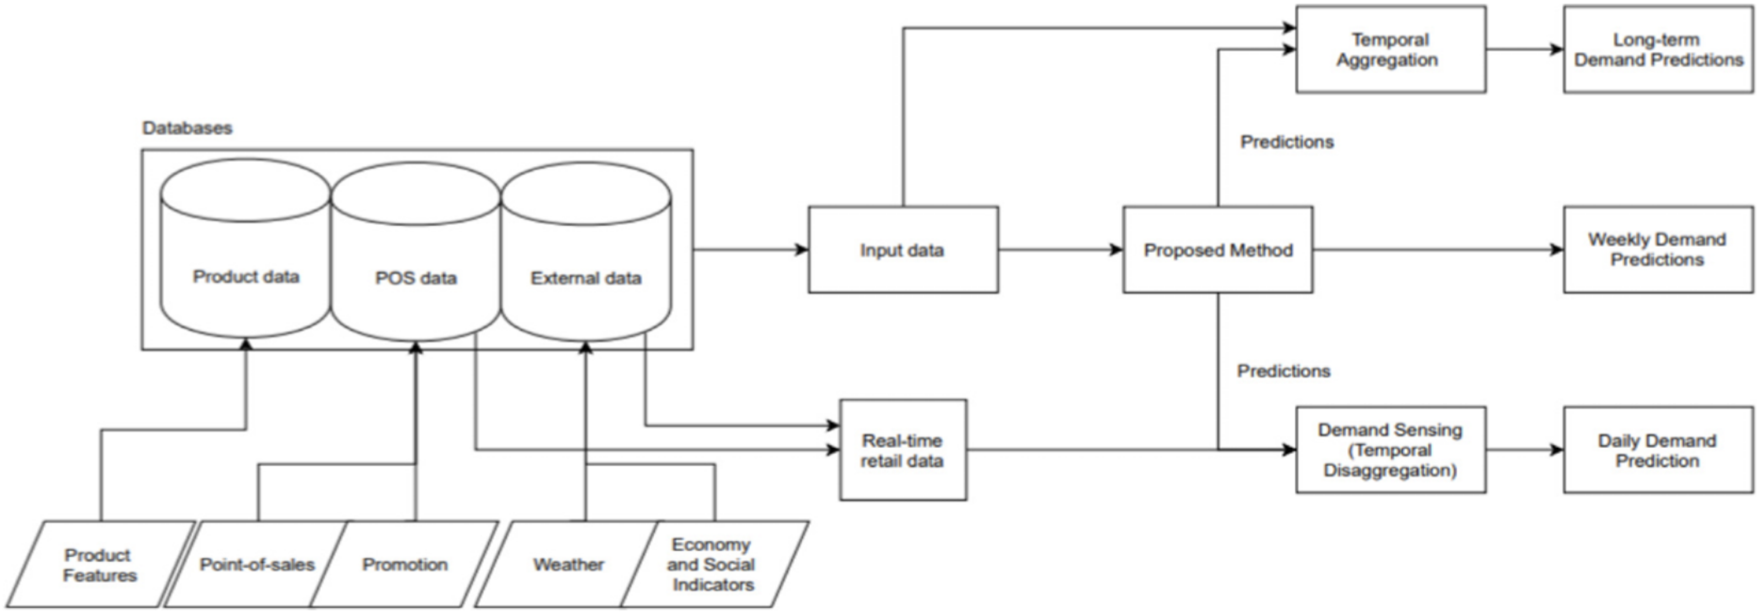

FLOW DIAGRAM FOR PROPOSED DEMAND FORECASTING FRAMEWORK

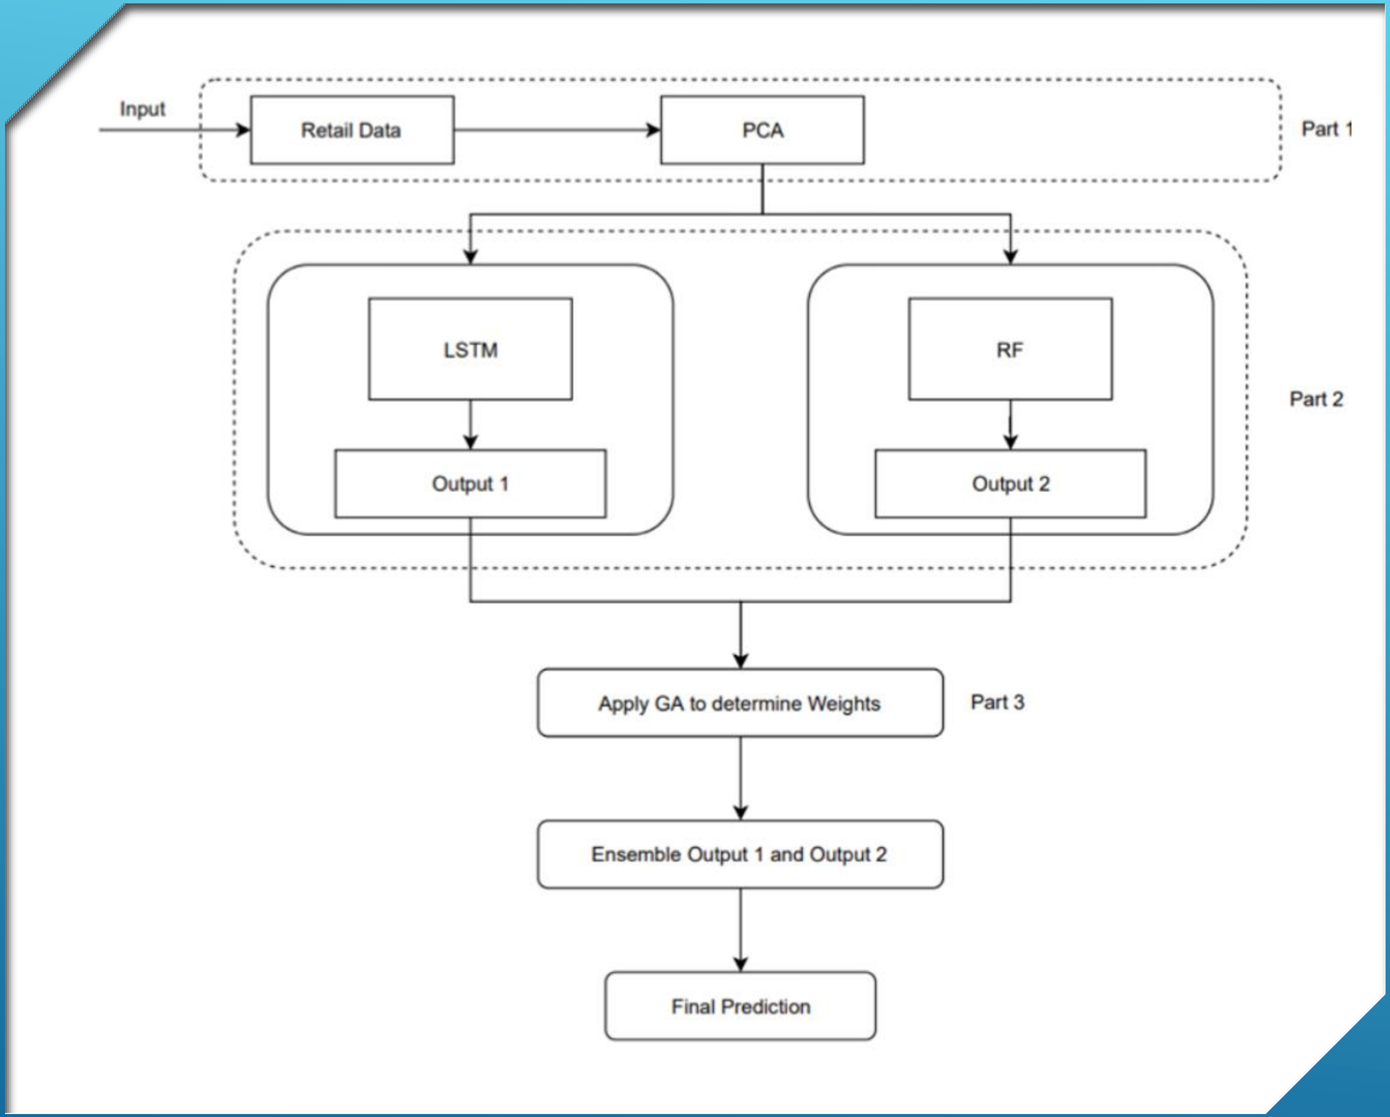

FLOWCHART OF THE PROPOSED ENSEMBLE METHOD

RELATIVE ERRORS FOR ONE-WEEK AHEAD PREDICTIONS (WITH RANKING IN BRACKETS).

	ME		MAE		RMSE	
OLS	1.000		1.0000		1.0000	
ARIMA	0.1553	(2)	0.6757	(7)	0.7918	(8)
ARIMAX	0.4018	(6)	0.6367	(5)	0.7009	(6)
RF	0.3146	(5)	0.5106	(4)	0.5796	(4)
NN	−0.4479	(8)	0.7511	(8)	0.7551	(7)
LSTM	−0.1950	(4)	0.4954	(3)	0.5588	(2)
ARIMA+NN	−0.2452	(7)	0.6586	(6)	0.6686	(5)
ARIMA+RF	0.1704	(3)	0.4221	(2)	0.5689	(3)
LSTM+RF	−0.1216	(1)	0.3569	(1)	0.4638	(1)

APPLYING MACHINE LEARNING IN RETAIL DEMAND PREDICTION

A COMPARISON OF TREE-BASED ENSEMBLES AND LONG SHORT-TERM MEMORY-BASED DEEP LEARNING

- THIS STUDY EVALUATES AI TECHNIQUES (RANDOM FOREST, XGBOOST, LSTM)
- FOR DEMAND PREDICTION IN RETAIL AND E-COMMERCE.

CONTEXT & IMPORTANCE

- Forecasting demand helps manage inventory, optimize supply chains, reduce costs, and improve satisfaction.
- Traditional methods struggle with seasonal, non-linear patterns.
- AI models handle complexity better.

DATA & SCOPE

- Real-world retail sales data:
- - Time-series demand
- - Seasonality & promotions
- - Product metadata
- Feature engineering captures temporal effects.

AI TECHNIQUES COMPARED

- **Tree-Based Ensembles (Random Forest, XGBoost):**
- Handle tabular data well
- Capture non-linear relationships
- Explainable, fast
- Feature importance metrics
- **LSTM (Deep Learning):**
- Designed for sequential data
- Captures long-term dependencies
- Requires tuning & computation
- Less interpretable

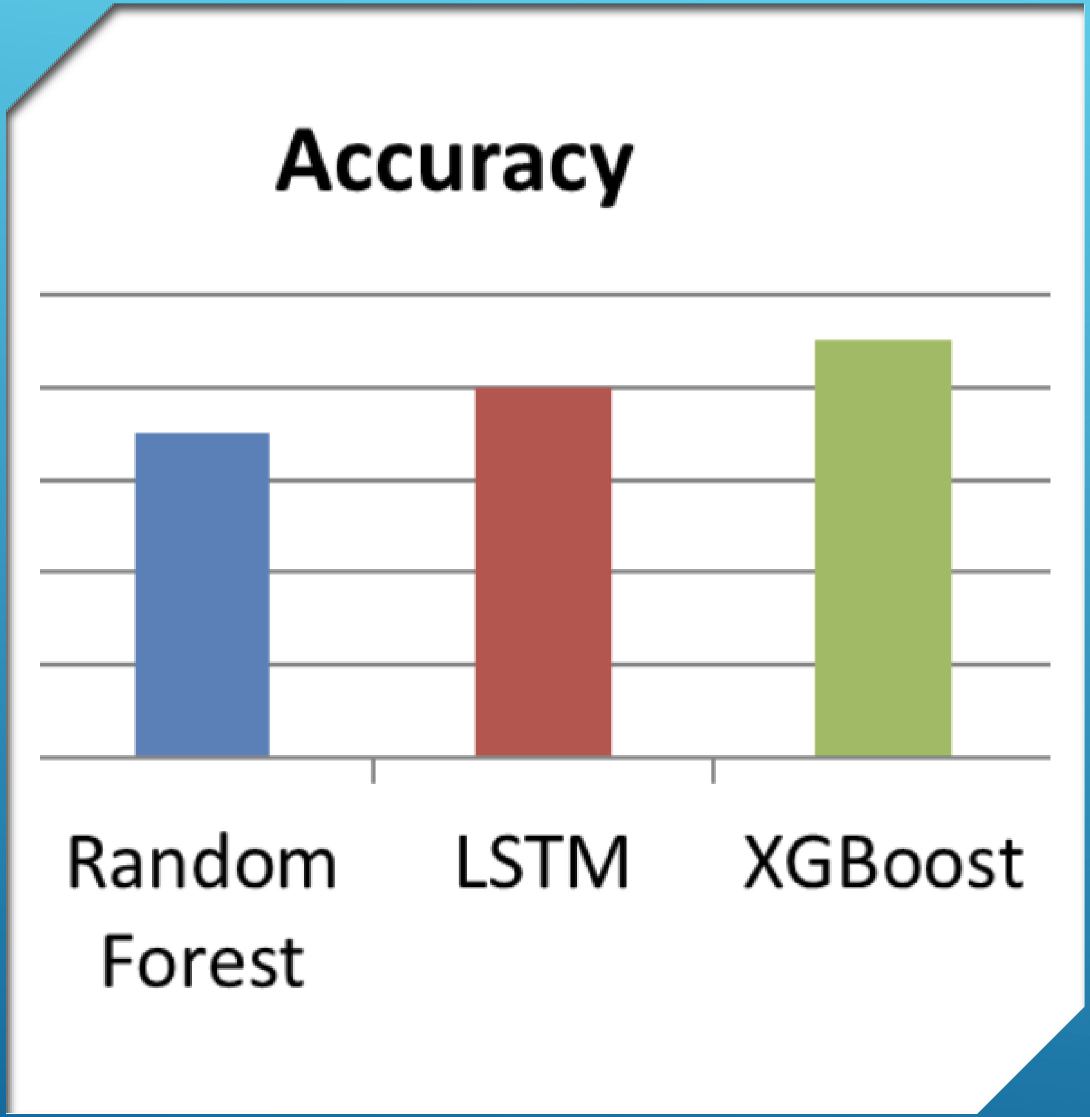

FINDINGS

- XGBoost outperformed LSTM and Random Forest in accuracy.
- LSTM good for strong time patterns.
- Tree-based models: robust, efficient.

IMPLICATIONS FOR RETAIL

- XGBoost: high accuracy, efficiency, interpretability
- LSTM: useful for long, sequential data
- Hybrid/ensemble models may provide best performance
- Choose model based on data, constraints, horizon

SALES PREDICTION SCHEME USING RFM BASED CLUSTERING AND REGRESSOR MODEL FOR E-COMMERCE COMPANY

OVERVIEW

- Focuses on predicting e-commerce sales using customer behavior data combined with product, region, and shipping data.
- Uses RFM clustering to capture customer habits:
- Recency (R) → How recently a customer made a purchase
- Frequency (F) → How often purchases occur
- Monetary (M) → Total spending of the customer
- Helps create customer segments that reflect buying patterns
- Goal: Provide actionable insights for marketing, inventory, and pricing

METHODOLOGY

- Customers divided into three clusters:
- Inactive Customers
- New Customers
- Loyal Customers
- Cluster labels are added to the main dataset (product/region/shipping info)
- Regression models applied to forecast sales:
- XGBoost used as primary model
- Handles multiple variables simultaneously (categorical + continuous)

KEY RESULTS

- Forecasts produced for all customer clusters with high accuracy:
  - $R^2 \approx 99\%$ → model explains most of the variation in sales
  - $RMSE \approx 9.8$ → low prediction error
- Predictions closely match actual sales trends
- Cluster-based approach captures customer-specific buying behavior, improving precision over using only product/time data
- Works across different types of products and regions

BUSINESS APPLICATIONS

- Marketing: Target promotions for loyal or new customers based on cluster insights
- Inventory: Plan stock levels for different regions/products using forecasted demand
- Pricing: Adjust dynamically for each customer segment
- Can be integrated into BI dashboards for monitoring and real-time decisions
- Potential extensions:
- Test with other datasets and industries
- Automate clustering for ongoing updates
- Combine with other AI techniques for even more granular predictions

EMERGING TECHNOLOGIES IN E-COMMERCE DEMAND FORECASTING TRENDS, PRACTICES, AND FUTURE PERSPECTIVES

- Based on Aye, da Silva, & Pereira Mafia (2024)

BACKGROUND

- Demand forecasting is critical for retail & e-commerce
- Traditional models: ARIMA, regression
- Limitations: rigid, less adaptive to rapid market changes
- AI & data-driven approaches now leading the field

RESEARCH SCOPE

- Systematic review of 42 studies
- Distinction: low-frequency (quarterly/seasonal) vs high-frequency (daily/real-time)
- Long-term strategy vs short-term operations
- Importance of combining both for resilience

AI & ML ALGORITHMS

- Machine Learning: Random Forests, Gradient Boosting → capture nonlinear patterns
- Deep Learning: LSTM, GRU → sequential dependencies
- Generative AI: GANs, Transformers → replicate unusual demand
- Hybrid models (ARIMA + LSTM) emerging

METRICS & EVALUATION

- Metrics: MAE, RMSE, MAPE, R^2
- AI consistently outperforms traditional approaches
- Deep learning = lowest errors, best for volatile data
- Hybrid models balance accuracy & stability
- External data (weather, holidays, economy) → +20% accuracy

KEY CHALLENGES

- Data quality: missing values, noise reduce accuracy
- Computational costs: DL expensive for SMEs
- Interpretability: 'black-box' AI not trusted by managers
- Gaps in research for emerging markets

REAL-WORLD APPLICATIONS

- Amazon: anticipatory shipping using ML/DL
- Combines browsing behavior, sales history, and external signals
- Enables predictive inventory & faster delivery
- Demonstrates AI as a requirement, not an option

FUTURE DIRECTIONS & CONCLUSION

- Explainable AI for manager trust
- Federated learning for privacy-preserving collaboration
- Reinforcement learning for adaptive forecasting
- Sustainable AI forecasting practices
- Conclusion: AI = strategic necessity in e-commerce

COMPARING MACHINE LEARNING MODELS WITH AND WITHOUT MACROECONOMIC VARIABLES FOR RETAIL DEMAND PREDICTION.

RESEARCH GAP & OBJECTIVE

- Gap: Most studies ignore macroeconomic conditions such as inflation, consumer sentiment, and unemployment.
- Objective: To evaluate how including macroeconomic variables improves demand forecasting.
- Models compared: Lasso, Ridge, LightGBM, XGBoost, Decision Tree.

DATASET & PREPROCESSING

- Dataset: Walmart USA, 5 years of data, 3,049 products, 10 stores across 3 states.
- Variables: product details, prices, promotions, calendar events.
- Added external variables: CPI, ICS, unemployment.
- Preprocessing: handled missing values, created lagged features, rolling averages, and removed outdated data.

METHODOLOGY

- Feature Engineering: Added seasonality, trends, and external economic variables.
- Training Strategy: Compared models using only historical sales vs. enriched with macroeconomic data.
- Evaluation Metrics: Mean Absolute Error (MAE) and Root Mean Squared Error (RMSE).

RESULTS OVERVIEW

Model	Without Macro Variables	With Macro Variables	Observation
Lasso	MAE=0.886 / RMSE=1.802	MAE=0.884 / RMSE=1.798	Slight improvement
Ridge	MAE=0.846 / RMSE=1.739	MAE=0.846 / RMSE=1.738	Almost same
LightGBM	MAE=0.849 / RMSE=1.717	MAE=0.847 / RMSE=1.715	Best improvement
XGBoost	MAE=0.841 / RMSE=1.722	MAE=0.839 / RMSE=1.716	Small gain
Decision Tree	MAE=1.000 / RMSE=2.357	MAE=1.002 / RMSE=2.365	Worse

RESULTS (CONTINUED)

- Incorporating macroeconomic data improves forecasting accuracy.
- Findings generalize beyond LSTM models to multiple ML techniques.
- Best model: LightGBM, as it captures complex relationships effectively.
- Retailers can use combined data (sales + economic) for better planning, inventory control, and competitiveness.

AI-DRIVEN DEMAND PREDICTION WITH SPATIAL GRAPH NEURAL NETWORKS(SGNNS)

- Based on Product Demand Prediction with SGNNs – J. Li.

WHAT ARE SGNN?

- Spatial Graph Neural Networks (SGNNs) are powerful AI models. They mix graph learning, which looks at relationships between locations, with temporal learning, which studies time-based trends.
- Nodes represent stores or online regions.
- Edges are connections like demographics, proximity, or purchasing patterns.
- These models are built to capture not only where demand occurs but also when it shifts.

HOW DO SGNNS WORK?

Graph Construction

- Build a network of stores or regions.
- Define links using **geography**, **customer behavior, or product similarity**.

Feature Input

- Demand history, seasonal effects, promotions, and sudden shocks.

HOW DO SGNNS WORK?

Model Processing

- Graph layer captures **spatial dependencies**.
- Sequence layer captures **time-series patterns**.

Prediction Output

- Generates accurate **demand forecasts** that reflect both **location dynamics** and **temporal trends**.

PERFORMANCE COMPARED TO OTHER MODELS

- **ARIMA**: good for simple time-series, but ignores spatial context.
- **LSTM (Deep Learning)**: captures time-based trends but still blind to store-to-store relationships.
- **SGNNs**: integrate both dimensions → leading to more precise, context-aware forecasts.

PERFORMANCE COMPARED TO OTHER MODELS

- Results:
  - **RMSE**: SGNNs cut errors by 10–15% compared to others.
  - **MAE**: consistently lower error values (~0.29).
  - **R^2**: often above **0.90**, showing very close fit with real data.
- **Advantage**: Fewer errors, better reliability, especially in dense retail networks.

BUSINESS IMPACT

- More accurate **inventory planning** → fewer stockouts.
- Reduced **excess inventory costs** → lower operational waste.
- Stronger **supply chain responsiveness** → adapts quickly to demand shifts.
- Smarter **targeted promotions** → higher customer satisfaction and profitability.

FUTURE DIRECTIONS

- Improve **graph construction methods** to better capture hidden store relationships.
- Scale SGNNs for **very large networks** across global retail chains.
- Enhance **real-time adaptability** for sudden demand shocks (e.g., viral trends, crisis events).
- Explore **multi-modal data integration** (social media, weather, economic signals).

CONCLUSION

WHAT WE'VE LEARNED

- Hybrid models enhance accuracy by blending temporal and covariate data
- Exogenous factors (e.g., weather, macroeconomic) are critical
- Emerging tech like GNNs and demand sensing drives e-commerce and retail innovation
- No one-size-fits-all: Model choice depends on context

WHERE FORECASTING IS HEADED

- Blend collaborative forecasting with AI for dynamic markets
- Extend models to emerging markets and long-term horizons
- Improve real-time adaptability for e-commerce and supply chains
- Focus on sustainability through optimized inventory

THANK YOU :)

- Questions?