Graph Neural Network Models for Retail Basket Recommendation Prediction

COLUMBIA UNIVERSITY

ORACLE

Frank Shi, Mika Zhu

Advisor: Professor Ali Hirsa, John Bible, Bat-Orgil Batjargal, Miao Wang

Background & Motivation

- Accurately predicting customer purchase behavior is a **central challenge** in retail analytics. Basket prediction can drive better recommendations, smarter inventory management, and personalized marketing. Yet the task is difficult: baskets often contain diverse items, and user preferences evolve over time.
- Our motivation is to build models that capture complex item–item and user–item relationships while accounting for temporal dynamics. By improving model architectures and scaling strategies, we aim to **advance basket prediction** toward more robust and practical real-world applications.

Data & Model Introduction

Datasets

- **Groceries**: smaller and simpler, only basic purchase histories
- **Multimodal**: larger and more complex, includes product descriptions, prices, and additional attributes

Models

- **MITGNN**: A multi-intent graph neural network that represents each basket as a combination of latent “shopping intents.” This allows the model to capture diverse co-purchase patterns.
- **MITGNN + GRNN**: Intergrated with recurrent module (e.g., LSTM) to model how user preferences evolve across multiple baskets.
- **MITGNN + TGN**: Intergrated with Temporal Graph Networks, adding a memory module and time-aware message passing. This design explicitly models temporal evolution of user–item interactions.

Metric culculation

- **Recall**: measures how many of the items a user actually bought appear in the top-K recommendations.
- **Precision**: measures accuracy of prediction.
- **NDCG**: measuring ranking quality, rewarding models that place relevant items higher in the list.
- **Hit Ratio**: checks whether at least one relevant item appears in the top-K list.

Model Prediction & Time Embedding

- **Model Prediction**: Our framework focuses on two key tasks: within-basket prediction and next-basket prediction. In within-basket prediction, the model predicts the remaining items in a user’s current basket based on the items already selected. In next-basket prediction, the model forecasts the items a user will purchase in their next shopping trip.
- **Improvements of Time Embedding**: The prior team used constant time embeddings, which limits temporal representation. We introduced a more expressive encoding that captures when purchases happen, enabling the model to better learn evolving user–item patterns.
- **Sliding Window Adoption**: We further introduced a configurable sliding window mechanism for sequence construction. It prevents information leakage and allows testing different context lengths, leading to more stable training and more robust predictions.

Configurations

- **Early Stopping**: stops training when validation performance plateaus, preventing overfitting.
- **Sliding Window**: constructs basket sequences with fixed window length and stride, enabling sequential training.
- **Prediction Analysis**: exports training/testing metrics, user-level prediction history (JSON/CSV), and learning curves for visibility.
- **Learning Rate Schedulers**:
  - Linear Warmup: gradually increases LR to stabilize early training.
  - Warmup Cosine Annealing: smooth decay to a minimum value after warmup.
  - Cosine Warmup Restart: cyclic cosine schedule with periodic restarts.
  - Polynomial Decay: reduces LR following a polynomial function, flexible for different decay speeds.

Experiments Result

- result of mitgnn, groceries data, next basket prediction

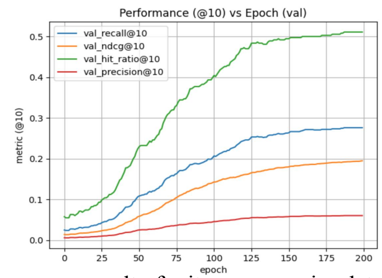

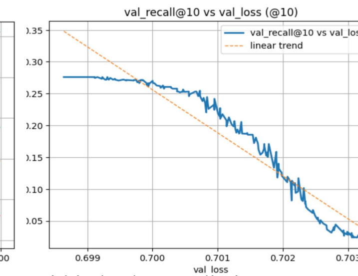

- result of mitgnn, groceries data, within basket prediction

	Recall	NDCG	Hit Ratio	Precision
@10	44.22%	31.06%	44.22%	4.42%
@20	47.53%	33.60%	47.53%	2.38%

- result of mitgnn, groceries data, next basket prediction

	Recall	NDCG	Hit Ratio	Precision
@10	32.69%	24.80%	64.50%	8.94%
@20	46.77%	29.33%	78.90%	6.50%

- result of tgn, groceries data, within basket prediction

	Recall	NDCG	Hit Ratio	Precision
@10	23.53%	13.27%	23.53%	2.35%
@20	40.03%	17.55%	40.03%	2.00%

- result of tgn, groceries data, next basket prediction

	Recall	NDCG	Hit Ratio	Precision
@10	6.01%	2.35%	16.15%	1.65%
@20	27.91%	9.32%	57.54%	3.81%

- **Performance Gains**: Our models achieve higher Recall, NDCG, and Precision compared with the prior baseline, showing clear improvements in both within-basket and next-basket tasks.
- **Training Dynamics**: Recall steadily increases as Loss decreases, confirming that the model is learning meaningful co-purchase and sequential patterns rather than overfitting.

Future Plan

- **Model Design**: Auto Neural Architecture Search; Task optimization; New Embedding Injection: flexibly integrate more features.
- **Practical Application:** Explore pretraining on representative stores and fine-tuning for new ones, using global item embeddings for consistency and store-specific user/basket embeddings for personalization, save resources and time.