PROBABILISTIC DEMAND FORECASTING WITH GRAPH NEURAL NETWORKS

ECML PKDD 2023
ML4ITS

Nikita Kozodoi[1], Elizaveta Zinovyeva[1], Simon Valentin[2], João Pereira[3], Rodrigo Agundez[3]
[1]Amazon Web Services, [2]University of Edinburgh, [3]adidas AG

INTRODUCTION

- Demand forecasting is instrumental to optimize stock planning and minimize waste.
- Demand of a product depends on the demand of other similar products.
- Most current forecasting approaches do not account for product relationships at inference time.

CONTRIBUTIONS

- End-to-end demand forecasting solution:
  - Based on **DeepAR** (sequence model).
  - Integrates **Graph Neural Network** encoder to account for **product relationships**.
  - Enables **probabilistic** forecasting.
- All-purpose, domain-agnostic graph construction approach based on product attribute similarity.
- Scales to thousands of products and product relations.

EXPERIMENTAL SETUP

DATA

	Dataset	# products	# weeks	# features
Public	Retail	629	148	12
Public	eCom	8,810	128	5
Private	adidas eCom	80,838	140	20

EVALUATION

- Benchmarked the proposed GraphDeepAR model against the DeepAR baseline.
- Metrics: RMSE, MAE, wMAPE, and financial metric.

MODEL

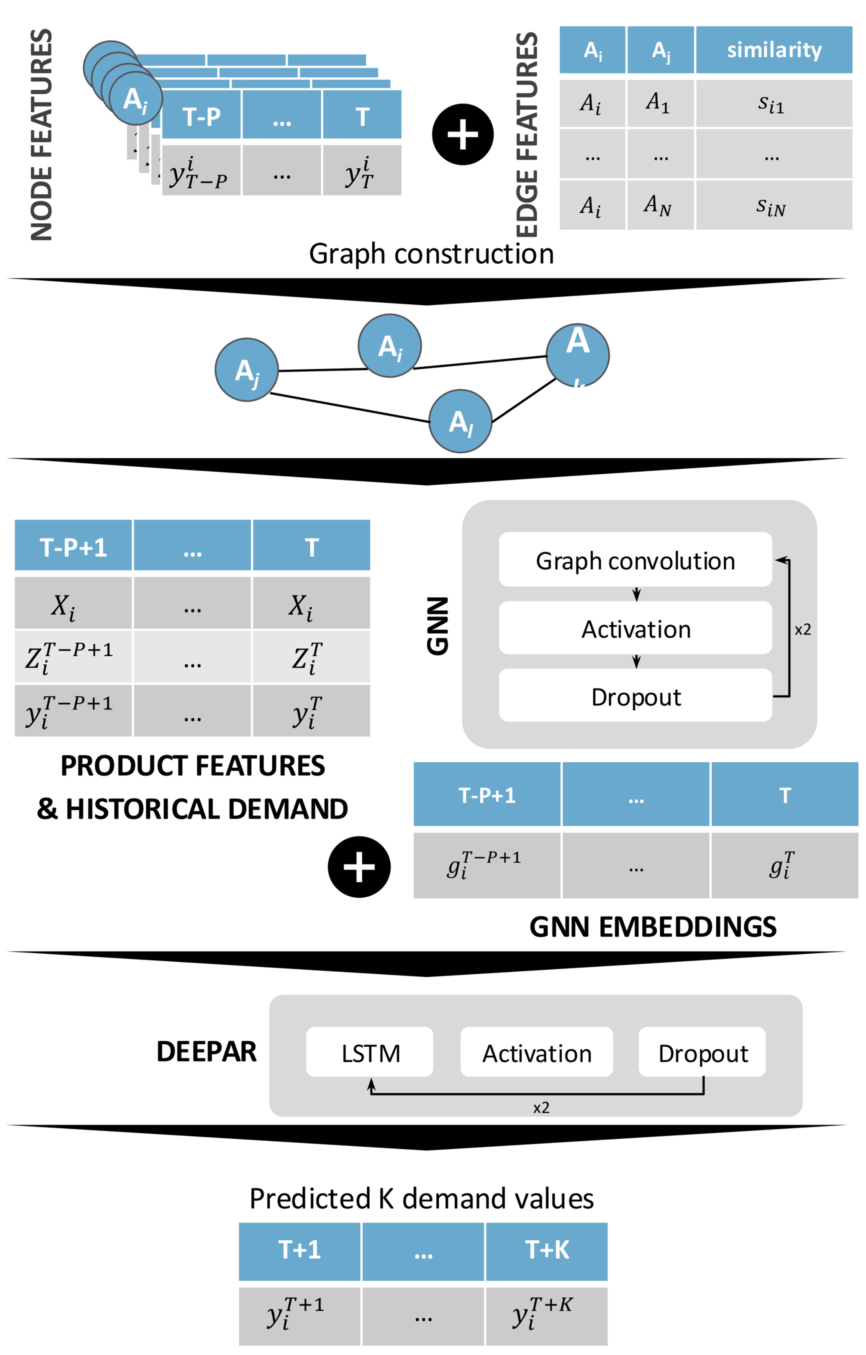

GRAPH CONSTRUCTION AND SAMPLING

- Each **node** represents a **product**, while **edges** are based on the **cosine similarity** between products.
- Similarity is computed using the **product attributes**: color, category, etc.
- Edges with similarity > similarity threshold are kept.
- Graph contains thousands of products that may have a very high number of connections - neighbors.
- **Random sampling** is used to reduce costs, time, and space complexity.

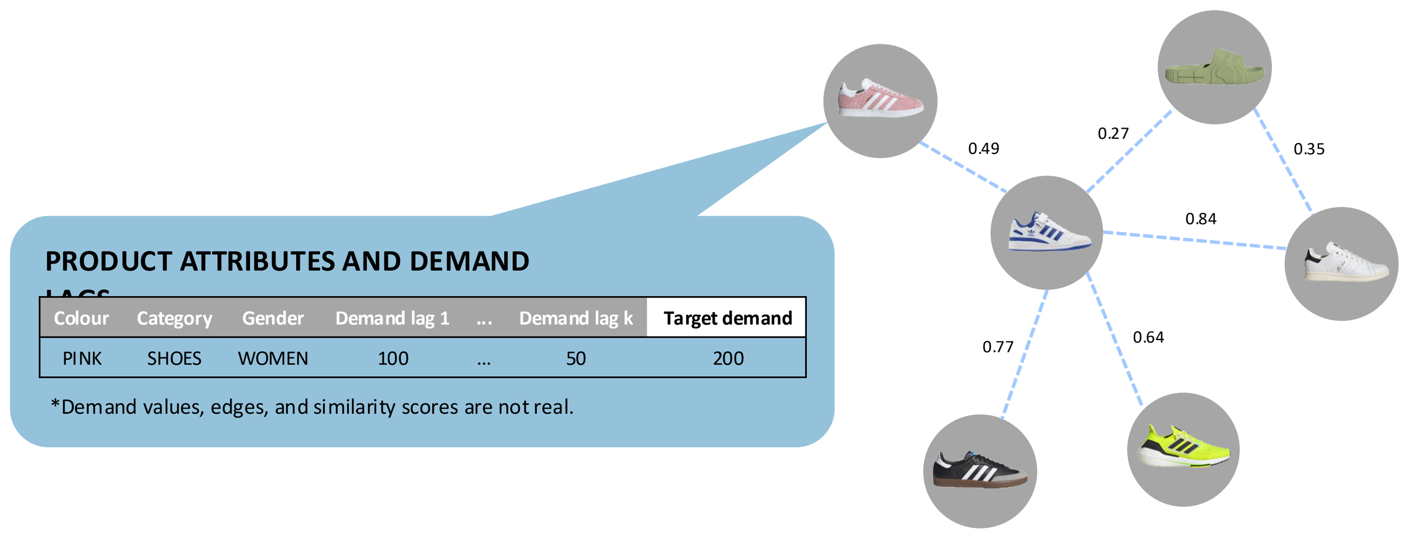

RESULTS

PERFORMANCE

- Adding the graph structure boosts performance in terms of predictive accuracy as well as financial gain:
  - Reduction of 4% in RMSE on public Retail dataset.
  - Reduction of 32% in RMSE on public eCom dataset.
  - 2% financial uplift on adidas eCom dataset.
- GraphDeepAR consistently outperforms the DeepAR baseline on both public datasets on connected products.
- Analyzed time-varying **node embeddings**.
- **Larger datasets** profit to the greater extent from the graph structure.

NODE EMBEDDINGS

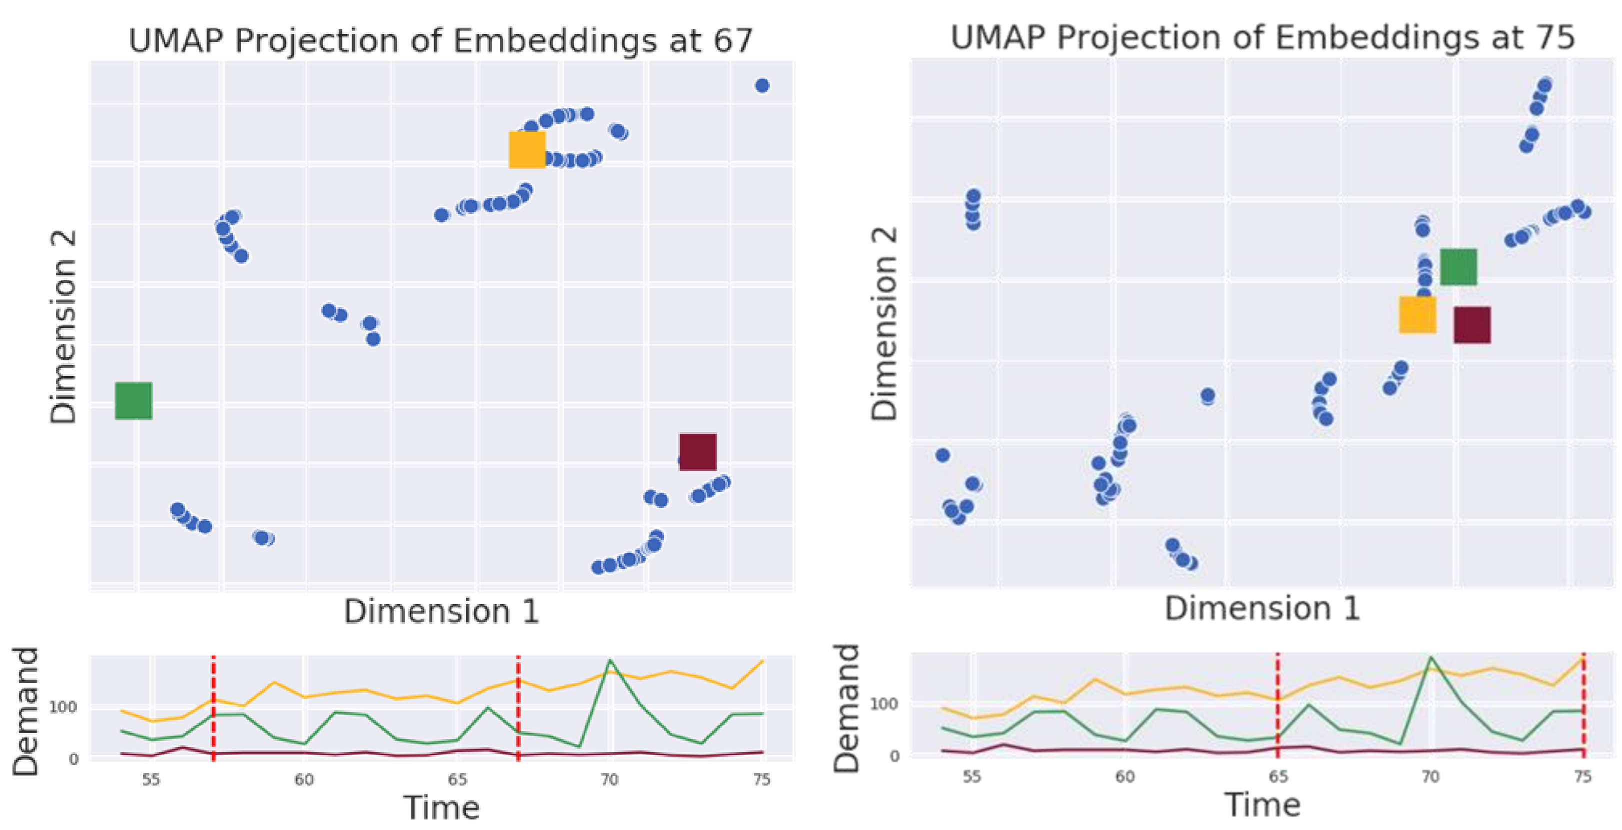

- Embeddings reflect both static product features as well as **demand dynamics** of neighboring products.
- Can serve as a **proxy for substitute and complement** products.

CONCLUSIONS

- GraphDeepAR takes product relationships into account at both training and inference time.
- Forecasting performance improves relative to DeepAR.
- Graph node embeddings are useful for other downstream tasks beyond forecasting.