Inductive learning for product assortment graph completion

H. Dukić¹, G. Deligiorgis², P. Sepe¹, D. Bacciu¹, and M. Trincavelli²

1- Universitá di Pisa - Dipartimento di Informatica Largo B. Pontecorvo 3, 56127 Pisa - Italy

2- H&M Group - Business Tech, AI, Analytics, and Data Mäster Samuelsgatan 46A - Sweden

Abstract. Global retailers have assortments that contain hundreds of thousands of products that can be linked by several types of relationships like style compatibility, "bought together", "watched together", etc. Graphs are a natural representation for assortments, where products are nodes and relations are edges. Relations like style compatibility are often produced by a manual process and therefore do not cover uniformly the whole graph. We propose to use inductive learning to enhance a graph encoding style compatibility of a fashion assortment, leveraging rich node information comprising textual descriptions and visual data. Then, we show how the proposed graph enhancement improves substantially the performance on transductive tasks with a minor impact on graph sparsity.

1 Introduction

Fashion data are interesting for research because of their polymorphism and the complexity of the relations that can be defined among them, i.e. compatibility, transactional, similarity, substitution, etc. Fashion items are considered compatible if they can be worn simultaneously, meaning that the clothing items are part of an outfit. Our work develops on fashion data assembled by H&M. In this context, the compatibility of fashion items is manually determined by experts on item pairs. The assortment, however, is composed by tens of thousands articles and the number of pairs grows quadratically with the number of articles, making exhaustive manual labelling highly impractical. Furthermore, when new products enter in the assortment, they stay disconnected for a rather long period. The lack of exhaustive indications of item compatibility can considerably impact the performance of recommendation systems that are leveraging such information to provide personalized and style-coherent advice to customers. Motivated by this, we tackle the problem of augmenting such sparse item compatibility information with newly discovered compatibility relationships. Existing works have addressed the problem with recurrent models [1, 2, 3] or with contrastive learning [4, 5]. Our approach, instead, leverages inductive learning on graphs [6]. Inductive link prediction, as opposed to transductive link prediction that assumes all nodes to be present at training time, aims at predicting links for new, unobserved, nodes. However, inductive link prediction usually obtains a lower performance on existing nodes. The method we propose is using inductive link prediction to enrich the graph with new links and then train a transductive

model on the new graph to maximize the link prediction performance, getting in this way the best of both worlds. Following [7], we represent items as nodes and compatibility as edges of the graph, together with their associated information (node and edge labels). In particular, our items are bound to rich textual and visual information, for which we define an appropriate encoding as node features. We then put forward an inductive learning approach based on the DEAL model [8], that has been extended to exploit the richness of the multimodal node features available in our industrial case study. As a first result, we show how these features positively contribute to relationship inference. The trained inductive model is then applied to produce an enriched graph for a second transductive task, modelling clothing pairing suggestions as a link prediction problem. The empirical analysis shows that the enriched graph yields to substantially improved link prediction performance over the original graph, at the cost of a minor decrease in graph sparsity. This second result is particularly interesting as it shows the effectiveness and efficiency of a pipeline of inductive-transductive methods when dealing with predictive tasks over large-scale sparse graphs.

2 Inductive-Transductive Graph Processing Pipeline

For inductive learning, we consider DEAL [8], an architecture leveraging two encoders, an attribute-oriented encoder H_a and a structure-oriented encoder H_s , as well as an alignment mechanism. The aim of the attribute-oriented encoder is to project a node's feature vector from the high-dimensional feature space into a low-dimensional embedding space, while the structure-oriented encoder generates an embedding vector of the node, by considering only the structural information of the graph (no node features). If two graph nodes are connected (positive samples), then their H_a and H_s embedding vectors will have high similarity. To this end, we measure similarity of the embedding vectors by cosine similarity. A Tight Alignment mechanism [8] is used to maximize the similarity between the embedding vectors produced by both H_a and H_s for each node. Both encoders are updated during the training process and the embeddings are kept aligned. The attribute-oriented encoder can be realized by MLP or GCNlike modules [8]. In our implementation we use the personalized ranking loss in [8]. Transductive learning is implemented with consolidated deep graph networks. In particular, in our empirical analysis, we confront the performance of three popular methods that well represent three different families of neural approaches for graphs that are Graph SAGE [9] model, GCN [10] and GAT [11]. As anticipated, the focus of this work is to propose and assess inductive learning as a preliminary step to improve transductive task performance on sparsely connected, large-scale graphs. The process comprises a first step where the DEAL [8] inductive model is trained and the best performing model (in validation) is selected. As a second step, we run the best inductive model selected on the original graph to enrich it with the introduction of new edges. For this second step, we define two thresholds: the maximum node degree and a probability of link existence among the nodes. As a final step, we train the transductive models on the new structure. In addition to the pipeline above, we extend DEAL [8] to work on textual, visual, or concatenation of textual and visual features, instead of the tabular features used in [8]. The embedding of the textual and visual information attached to each product in our case study has been obtained by a BERT model [12] pre-trained on English Wikipedia and by a ResNet512 [13] pre-trained on ImageNet [14], respectively.

3 Item Compatibility Graph

Our study considers two novel industrial proprietary datasets provided by H&M, where each node is associated with a fashion item and the presence of an edge between two nodes denotes style compatibility between the items. These graphs have been compiled from a list of pairwise fashion item compatibility statements compiled by H&M domain experts. This information has been used to build two separate graphs, one for Men and one for Women clothing. Both contain a large number of products, represented by an image, text description, colour, and other tabular data. These fashion graphs have been assembled specifically for this work and this is the first graph-based predictive analysis being performed on such data. We complement our analysis on proprietary data with a publicly available dataset, the Computers network [15], with structural characteristics that are akin to those of our industrial use case. All networks have been represented using the Open Graph Benchmark (OGB) [16] format, and the relevant properties of the aforementioned datasets are given in Table 1. The challenging aspect shared by all datasets is the high level of edge sparsity, nearing 100%, and the nontrivial proportion of disconnected nodes (i.e. with zero degree). The latter is particularly true for the fashion data. This is the key motivational aspect for our approach, as we would like to be able to enrich the graph edges by inductive learning before fitting the target transductive task to the data. With respect to this, Table 1 already reports an anticipation of the results of the inductive enrichment of the graph (marked in bold). One can clearly see a considerable drop in disconnected nodes, with a minor change to the edge sparsity. From a model selection perspective, we split differently the graph adjacency matrix A depending on the task to be performed (inductive or transductive link prediction). In particular, for the inductive link prediction task, we needed to assure that one or both nodes, seen during the training process are not seen during the evaluation process. For this reason, \mathbf{A} is split on a node basis. For the transductive task, instead, we partition the network on an edge basis. It is important to mention that negative training edges are sampled uniformly during the training phase, while the validation ones are sampled in advance and are kept fixed for the duration of the model assessment. After the final inductive model is chosen, we set different thresholds for the maximum node degree and a link existence probability. In the case of graph for Men, graph for Women and Computers [15], the thresholds for maximum node degree are set to 5, 2, and 20, respectively, while the thresholds for the probability of link existence are set to 0.85, 0.99 and 0.60, respectively.

Graphs' properties	Men	Men*	Women	Women*	Computers	$Computers_*$
Number of nodes	22912	22912	57447	57447	13752	13752
Number of zero-degree nodes	5420	16	11720	2811	281	0
Percentage of zero degree nodes	23.65%	0.06%	20.40%	4.89%	2.04%	0.00%
Number of edges	290514	4265230	642090	4013554	491722	3461572
Sparsity	99.94%	$\boldsymbol{99.18\%}$	99.98%	$\mathbf{99.87\%}$	99.73%	98.16%

Table 1: Dataset summarization: with * we denote enriched graphs.

4 Experiments & Results

The data used for training the baselines and our proposed method are images and text descriptions of H&M's assortments. For each dataset, configuration, and task, hyperparameter selection has been performed by using Optuna [17], an hyperparameter optimization software framework. For the inductive link prediction task, we trained DEAL-based [8] models with different architectures and configurations. The results are reported in Table 2. In particular, we consider two attribute-oriented encoder mechanisms: an MLP and a trainable Embedding layer [18]. The performances for both inductive (Table 2) and transductive (Table 3) link prediction are highly improved when using visual or concatenation of visual and textual node features. The best performances on the graph for Men and graph for Women achieved DEAL_{MLP}, while DEAL_{EMB} performed better in the case of Computers [15] for two out of three metrics. The best performing configuration for each graph is used to perform graph enrichment for the successive transductive analysis.

		Accuracy			ROC - AUC			AP		
	Model	Text	Image	Text+Image	Text	Image	Text+Image	Text	Image	Text+Image
Men	$\mathrm{DEAL}_{\mathrm{EMB}}$	0.6432	0.7177	0.7287	0.7117	0.8010	0.8208	0.6966	0.7463	0.7746
	$\mathbf{DEAL}_{\mathbf{MLP}}$	0.7083	0.7432	0.7551	0.7728	0.7999	0.8370	0.7459	0.7438	0.7987
Women	$\mathrm{DEAL}_{\mathrm{EMB}}$	0.6573	0.7258	0.7361	0.7349	0.8183	0.8413	0.7231	0.7812	0.8088
	$\mathbf{DEAL}_{\mathbf{MLP}}$	0.7276	0.7718	0.7851	0.8064	0.8402	0.8622	0.7750	0.7963	0.8223
mp.	DEALEMB	DEAL _{EMB} 0.7425			0.9072			0.9019		
Co	$\mathrm{DEAL}_{\mathrm{MLP}}$	0.7606			0.8808			0.8624		

Table 2: The results of inductive link prediction on Graph For Men, Women, and Computers.

In Table 3 we can see how the enriched graph improves substantially the link prediction performance for all three different types of GNN considered (SAGE, GAT, and GCN), with respect to the metrics we considered, namely Accuracy, Receiver Operating Characteristic (ROC) Area Under Curve (AUC) and Average Precision (AP). This improvement in performance may be interpreted by

		Accuracy			ROC - AUC			AP		
	Model	Text	Image	Text + Image	Text	Image	Text + Image	Text	Image	Text + Image
Men	SAGE	0.9133	0.9256	0.9323	0.9646	0.9710	0.9742	0.9563	0.9649	0.9674
	SAGE*	0.9706	0.9740	0.9736	0.9932	0.9939	0.9944	0.9900	0.9907	0.9915
	GCN	0.9180	0.9206	0.9323	0.9687	0.9718	0.9757	0.9649	0.9682	0.9737
	GCN*	0.9704	0.9751	0.9772	0.9939	0.9951	0.9956	0.9923	0.9935	0.9945
	GAT	0.9024	0.9133	0.9154	0.9565	0.9658	0.9649	0.9483	0.9604	0.9588
	GAT*	0.9720	0.9715	0.9750	0.9940	0.9934	0.9941	0.9913	0.9889	0.9907
nen	SAGE	0.9214	0.9402	0.9392	0.9678	0.9788	0.9790	0.9627	0.9750	0.9756
	SAGE*	0.9700	0.9717	0.9711	0.9941	0.9946	0.9950	0.9932	0.9929	0.9943
	GCN	0.9170	0.9401	0.9462	0.9695	0.9813	0.9839	0.9683	0.9807	0.9834
Μo	GCN*	0.9690	0.9722	0.9745	0.9932	0.9952	0.9951	0.9933	0.9947	0.9953
-	GAT	0.9238	0.9360	0.9344	0.9695	0.9759	0.9752	0.9703	0.9735	0.9728
	GAT_*	0.9722	0.9756	0.9748	0.9935	0.9946	0.9948	0.9916	0.9921	0.9929
	SAGE	0.9457			0.9832			0.9788		
rs	SAGE*	0.9692			0.9935			0.9904		
ute	GCN	0.9421			0.9835			0.9818		
dui	GCN*	0.9647			0.9928			0.9910		
ő	GAT	0.9364			0.9770			0.9706		
	GAT*	0.9669			0.9919			0.9864		

the fact that the enrichment process effectively completes the original graph, making the patterns more regular, general, and therefore easier to learn.

Table 3: The results of transductive link prediction on Graph For Men, Women, and Computers. With * are denoted the results of the models after performing the graph enrichment.

5 Conclusions

We proposed an inductive learning approach for completing sparse graphs describing item compatibility information, and we have applied our method to both a publicly available benchmark as well as to a novel industrial use case based on the product assortment of a global fashion retailer. The proposed approach consists of two steps. First, we learn an inductive learning model that we use to generate new links for those nodes of the graph that are disconnected or sparsely connected. We then train a transductive model using the enriched graph, showing that we achieve increased link prediction performance. Our hypothesis is that the inductive learning model manages to learn the patterns of the connected nodes and transfer them to the sparsely connected nodes, making the structure of the graph more regular. This makes sense since we know from the process generating the connections in the graph, which is manual and labor intensive, that many possible connections are missing in the original graph.

Future works will study more thoroughly the graph enrichment step, that in this work has been carried out with a very simple methodology, by selecting the nodes with a maximum number of neighbours and a threshold for the inductive probability prediction. A more principled approach could give further performance improvements. ESANN 2021 proceedings, European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning. Online event, 6-8 October 2021, i6doc.com publ., ISBN 978287587082-7. Available from http://www.i6doc.com/en/.

References

- Xintong Han, Zuxuan Wu, Yu-Gang Jiang, and Larry S Davis. Learning fashion compatibility with bidirectional lstms. In Proceedings of the 25th ACM international conference on Multimedia, pages 1078–1086, 2017.
- [2] Takuma Nakamura and Ryosuke Goto. Outfit generation and style extraction via bidirectional lstm and autoencoder. arXiv preprint arXiv:1807.03133, 2018.
- [3] Federico Monti, Michael Bronstein, and Xavier Bresson. Geometric matrix completion with recurrent multi-graph neural networks. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, Advances in Neural Information Processing Systems, volume 30. Curran Associates, Inc., 2017.
- [4] Mariya I Vasileva, Bryan A Plummer, Krishna Dusad, Shreya Rajpal, Ranjitha Kumar, and David Forsyth. Learning type-aware embeddings for fashion compatibility. In Proceedings of the European Conference on Computer Vision (ECCV), pages 390–405, 2018.
- [5] Reuben Tan, Mariya I Vasileva, Kate Saenko, and Bryan A Plummer. Learning similarity conditions without explicit supervision. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10373–10382, 2019.
- [6] Davide Bacciu, Federico Errica, Alessio Micheli, and Marco Podda. A gentle introduction to deep learning for graphs. *Neural Networks*, 2020.
- [7] Guillem Cucurull, Perouz Taslakian, and David Vazquez. Context-aware visual compatibility prediction. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 12617–12626, 2019.
- [8] Yu Hao, Xin Cao, Yixiang Fang, Xike Xie, and Sibo Wang. Inductive link prediction for nodes having only attribute information. arXiv preprint arXiv:2007.08053, 2020.
- [9] William L Hamilton, Rex Ying, and Jure Leskovec. Inductive representation learning on large graphs. arXiv preprint arXiv:1706.02216, 2017.
- [10] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. arXiv preprint arXiv:1609.02907, 2016.
- [11] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. Graph attention networks. arXiv preprint arXiv:1710.10903, 2017.
- [12] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pretraining of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805, 2018.
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 770–778, 2016.
- [14] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In 2009 IEEE conference on computer vision and pattern recognition, pages 248–255. Ieee, 2009.
- [15] Julian McAuley, Christopher Targett, Qinfeng Shi, and Anton Van Den Hengel. Imagebased recommendations on styles and substitutes. In Proceedings of the 38th international ACM SIGIR conference on research and development in information retrieval, pages 43– 52, 2015.
- [16] Weihua Hu, Matthias Fey, Marinka Zitnik, Yuxiao Dong, Hongyu Ren, Bowen Liu, Michele Catasta, and Jure Leskovec. Open graph benchmark: Datasets for machine learning on graphs. arXiv preprint arXiv:2005.00687, 2020.
- [17] Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. Optuna: A next-generation hyperparameter optimization framework. In Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining, pages 2623–2631, 2019.
- [18] PyTorch Documentation. Embedding. https://pytorch.org/docs/stable/generated/ torch.nn.Embedding.html, 2019.