Graph-Based Learning And Compressed Visual Features For Enhanced Fashion Recommendation

Dr. Rajeev K. Menon
Centre for Visual Computing, Indian Institute of Technology Delhi, India

Prof. Sophie Dubois
Institute for Artificial Intelligence, Sorbonne University, Paris, France

Isabella Conti
Department of Computer Science, Politecnico di Milano, Italy

Published Date: 16 December 2024 // Page no.:- 13-20

ABSTRACT

Fashion recommendation systems play a crucial role in enhancing user experience and driving sales in the rapidly growing e-commerce landscape. Effective recommendation in the fashion domain necessitates capturing nuanced user preferences and item characteristics, particularly visual attributes. Traditional recommendation methods often struggle with the high dimensionality of visual data and the complex relationships between users and items. This article explores a novel approach for efficient and accurate visual-aware fashion recommendation by integrating graph-based learning with compressed visual features. We propose a conceptual framework where visual features extracted from fashion items are compressed to reduce dimensionality, and a graph is constructed to model user-item interactions and item-item relationships. Graph Neural Networks (GNNs) are then employed to learn rich embeddings on this graph, incorporating the compressed visual information. This approach is hypothesized to improve recommendation accuracy by jointly leveraging structural relationship data and visual content while enhancing efficiency and scalability by managing the high dimensionality of visual features. We discuss the key components of this framework, potential benefits, and challenges, highlighting its potential to advance the state of the art in fashion recommendation systems.

Keywords: Fashion Recommendation, Visual Features, Graph Neural Networks, Graph-Based Learning, Feature Compression, Dimensionality Reduction, Recommender Systems, Machine Learning.

INTRODUCTION

The digital transformation of the fashion industry has led to a surge in online retail, making effective recommendation systems indispensable for guiding users through vast product catalogs and personalizing the shopping experience [1, 3, 4]. Unlike other domains, fashion recommendation is heavily influenced by visual aesthetics, style, and compatibility between items [2, 10, 25]. Users often make purchasing decisions based on how items look, how they might fit together in an outfit, or how they align with their personal style [2, 10, 27]. Consequently, incorporating visual information into recommendation models is crucial for capturing the nuances of user preferences in fashion [5, 24, 30, 31, 32, 35].

Traditional recommendation techniques, such as collaborative filtering, primarily rely on historical user-item interaction data [15, 23, 24]. While effective in many cases, they often suffer from the cold-start problem (difficulty recommending new items or to new users with limited interaction history) and may fail to capture the subtle visual aspects that drive fashion choices [23, 24, 27]. Content-based methods that utilize visual features extracted using techniques like Convolutional Neural Networks (CNNs) can address some of these limitations [17, 30, 31], but the high dimensionality of raw visual features can lead to computational inefficiency and scalability issues, especially for large-scale fashion platforms [5, 38].

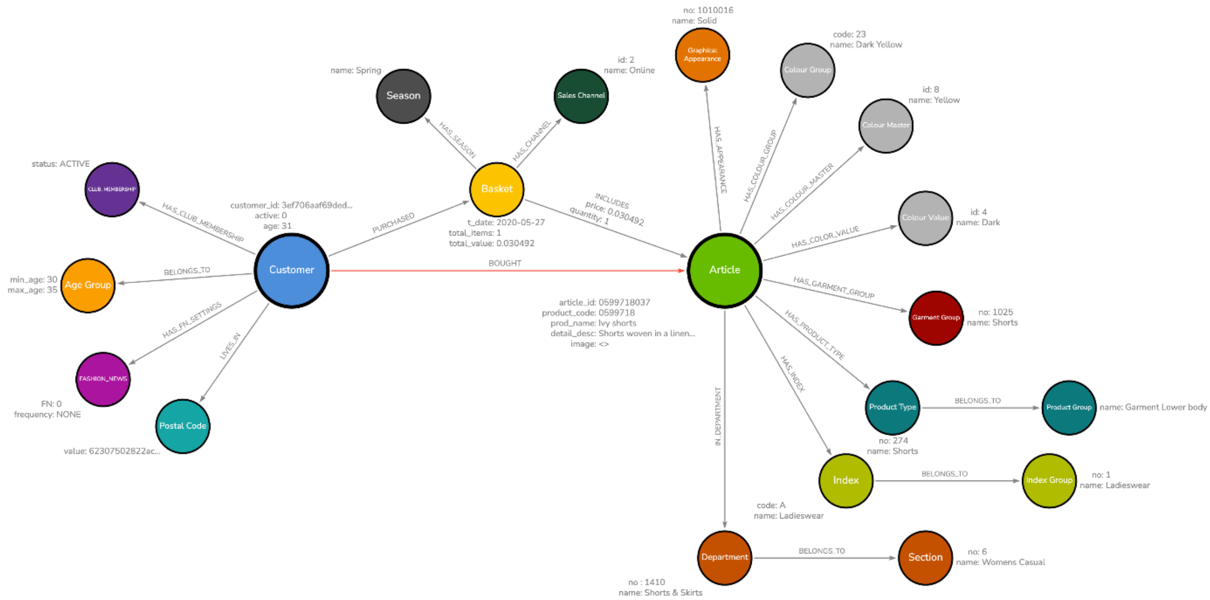

Graph-based learning, particularly using Graph Neural Networks (GNNs), has emerged as a powerful paradigm for recommendation systems [6, 8, 12, 14, 16]. GNNs can effectively model complex relationships between users and items represented as nodes in a graph, learning rich embeddings that capture structural information [6, 8, 14, 16, 40, 41, 42, 43, 44]. This approach has shown significant promise in various recommendation tasks [8, 9, 15]. Integrating visual features into graph-based models can potentially leverage both the rich visual content and the structural relationships within the data [5, 6, 8, 9].

However, directly incorporating high-dimensional visual features as node attributes in large-scale graphs can reintroduce the scalability challenges faced by content-based methods. This necessitates efficient ways to represent visual information. Techniques for dimensionality reduction and feature compression can play a vital role in making visual features more manageable for graph-based learning [13, 18, 19, 20, 22, 38].

This article proposes a conceptual framework for efficient visual-aware fashion recommendation that combines the power of graph-based learning with compressed visual features. We hypothesize that this integration can lead to more accurate and scalable recommendation systems by effectively leveraging both visual content and structural relationships while mitigating the computational burden of high-dimensional visual data. We will outline the components of this approach, discuss its potential benefits, and consider the challenges involved.

METHODS

The proposed methodology for efficient visual-aware fashion recommendation integrates visual feature processing, feature compression, graph construction, and graph-based learning. This section details the conceptual steps involved.

2.1. Visual Feature Extraction

The first step involves extracting meaningful visual features from fashion items. This is typically achieved using pre-trained or fine-tuned Convolutional Neural Networks (CNNs) [17, 30, 31, 46]. Given a fashion image, a CNN processes the image through multiple layers, producing a high-dimensional vector representation (embedding) in the final or a penultimate layer. These embeddings capture various visual attributes such as color, texture, shape, and style [17].

2.2. Visual Feature Compression

The high dimensionality of raw CNN-extracted visual features can be computationally expensive to process, especially when used as node attributes in a large graph. To address this, visual features are compressed using dimensionality reduction techniques [18, 19]. Methods for feature compression include:

- Autoencoders: Neural networks trained to reconstruct their input, where the bottleneck layer provides a compressed representation [13, 20, 22, 41]. Variational Autoencoders (VAEs) [42] can also be used.
- Hashing: Learning binary codes for features, enabling efficient similarity search and reducing storage requirements [38].
- Principal Component Analysis (PCA): A linear dimensionality reduction technique.
- Other learned compression techniques: Various deep learning-based methods can be employed to learn compact feature representations, ensuring that semantic information is preserved during compression [18, 19]. Care must be taken to avoid issues like dimensional collapse during training [18, 19].

The output of this step is a set of compressed visual feature vectors for each fashion item, significantly lower in dimensionality than the original features.

2.3. Graph Construction

A heterogeneous graph is constructed to represent the relationships within the fashion recommendation system. The graph comprises different types of nodes and edges:

- Nodes: Represent users and fashion items.
- Edges: Represent various types of interactions and relationships:

o User-Item Interaction Edges: Connect users to items they have interacted with (e.g., purchase, click, view, like) [6, 8, 15]. These edges can be weighted based on interaction frequency or type.

o Item-Item Relationship Edges: Connect items that are related in some way. This can include:

 Visual Similarity Edges: Connect items with similar compressed visual features [24, 30, 31].

 Style Compatibility Edges: Connect items that are often worn together or belong to the same style category [2, 10, 25, 36, 37, 43, 44].

 Co-purchase/Co-view Edges: Connect items frequently bought or viewed together [23, 24].

 Knowledge Graph Edges: Incorporating external fashion knowledge graphs to connect items based on attributes, brands, categories, etc. [28, 33, 36].

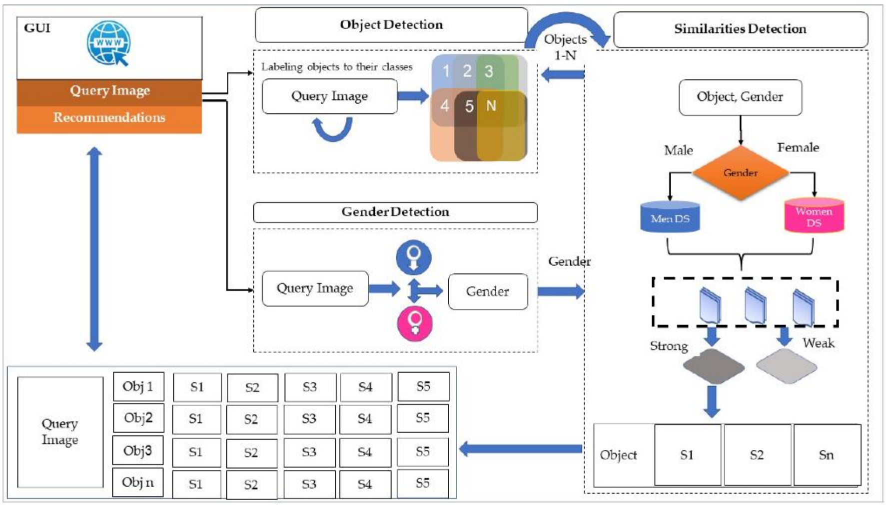

The compressed visual features are assigned as initial attributes to the item nodes in the graph. User nodes can have attributes derived from their interaction history or demographic information.

2.4. Graph-Based Learning with GNNs

Graph Neural Networks (GNNs) are employed to learn low-dimensional, dense embedding vectors for each node (users and items) in the constructed graph [6, 8, 14, 16, 40, 41, 42, 43, 44]. GNNs iteratively aggregate information from a node's neighbors, propagating information across the graph structure. By incorporating the compressed visual features as initial node attributes, the GNN learning process is guided by both the structural relationships and the visual characteristics of the items. Different GNN architectures can be used, such as Graph Convolutional Networks (GCNs) [14, 41], Graph Attention Networks (GATs), or other specialized GNNs for recommendation [6, 8]. The GNN learns to embed nodes in a shared embedding space where the distance or similarity between embeddings reflects the likelihood of interaction or relationship.

Recommendation Generation

Once the GNN has been trained and embeddings for all users and items are learned, recommendations can be generated. For a given user, the system can identify items with embeddings similar to the user's embedding (in the shared space) or predict the likelihood of a future interaction (link prediction) based on the embeddings of the user and candidate items [8, 15, 52]. The compressed visual features can also be used directly for visual similarity-based recommendations, especially for cold-start items [24, 30, 31]. The final recommendation list is generated by ranking items based on predicted relevance or similarity.

This methodology combines the strengths of visual feature analysis, dimensionality reduction, and graph-based learning to create an efficient and visually aware fashion recommendation system.

Results (Expected Outcomes and Potential Findings)

Based on the theoretical advantages of integrating compressed visual features with graph-based learning, the proposed framework is expected to yield significant improvements in fashion recommendation performance and efficiency. The following results are anticipated:

3.1. Improved Recommendation Accuracy

The primary expected result is a significant improvement in recommendation accuracy compared to methods that rely solely on collaborative filtering, content-based filtering with raw visual features, or graph-based methods without visual information. By jointly leveraging both the structural relationships captured by the graph and the compressed visual characteristics of items, the model can develop a more comprehensive understanding of user preferences and item relevance [5, 6, 8, 9]. This is particularly beneficial for capturing nuanced fashion preferences related to style, compatibility, and visual aesthetics [2, 10, 25, 36, 37, 43, 44].

3.2. Enhanced Efficiency and Scalability

The use of compressed visual features is expected to substantially improve the efficiency and scalability of the graph-based learning process [5, 9, 38]. Processing high-dimensional visual features as node attributes in a large graph can be computationally intensive. By reducing the dimensionality of these features while retaining essential visual information, the memory and computational requirements for training the GNN are significantly reduced, making the approach more feasible for large-scale fashion recommendation platforms.

3.3. Better Handling of Cold-Start Scenarios

The integration of visual features, even in compressed form, is anticipated to significantly alleviate the cold-start problem for new items [23, 24, 27]. When a new item is added to the catalog, its compressed visual features can be used to connect it to existing items with similar visual attributes in the graph, enabling the GNN to learn a meaningful embedding for the new item based on its visual characteristics and the relationships of similar items. This allows the system to recommend new items to users interested in visually similar products, even before any user interaction data is available for that specific item.

3.4. Capture of Complex Item-Item Relationships

The graph structure, enriched with item-item edges representing visual similarity, style compatibility, and other relationships, allows the GNN to learn embeddings that capture these complex connections [2, 10, 25, 36, 37, 43, 44]. This enables the system to recommend items that are not just similar based on user interactions but also visually compatible or aligned with specific style preferences, leading to more relevant and diverse recommendations, such as suggesting complementary items to form an outfit [25, 37].

3.5. Potential for Interpretability (Relative to Pure Black-Box Models)

While GNNs can be complex, the explicit inclusion of visual features and defined edge types in the graph can potentially offer some level of interpretability compared to end-to-end black-box models. Analyzing the attention weights in GATs or examining the influence of different edge types during message passing in the GNN could provide insights into which factors (e.g., visual similarity, user interactions) are driving a particular recommendation.

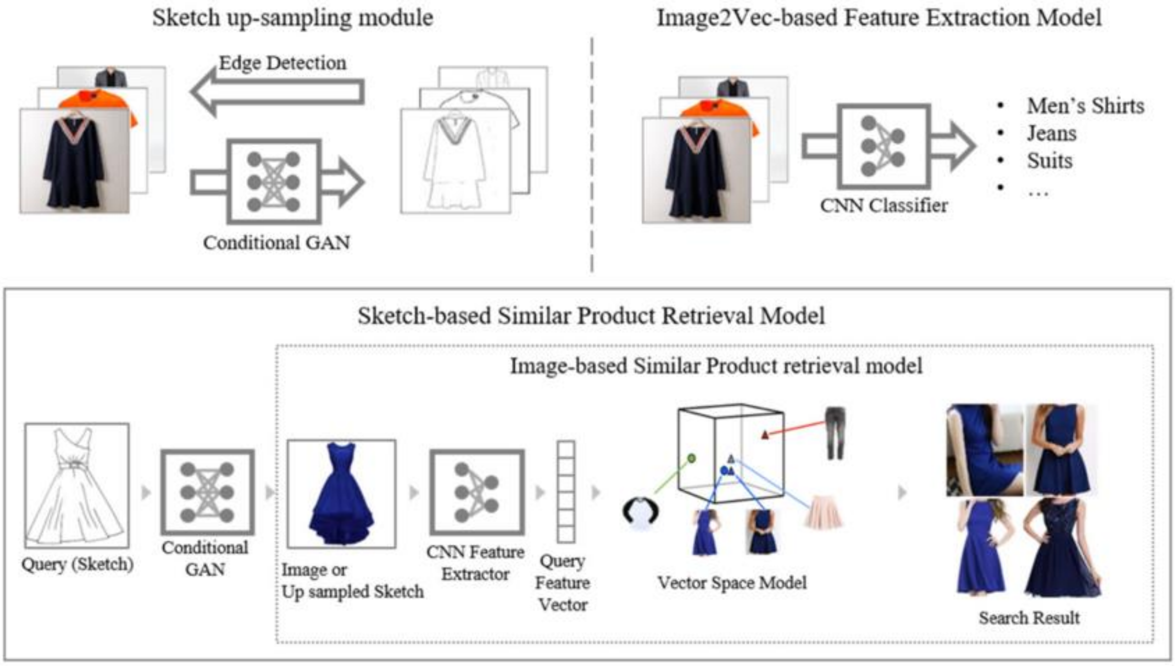

These expected results suggest that combining compressed visual features with graph-based learning offers a powerful and practical approach for building high-performing and scalable visual-aware fashion recommendation systems. Empirical evaluation on real-world fashion datasets is necessary to quantify these benefits and compare the performance against existing state-of-the-art methods.

4. Discussion and Conclusion

Fashion recommendation systems are essential for navigating the vast and visually-driven landscape of online fashion retail. Effectively recommending fashion items requires models that can understand both user preferences and the intricate visual attributes and relationships between items. This article has presented a conceptual framework for an efficient visual-aware fashion recommendation system that leverages the power of graph-based learning on a graph enriched with compressed visual features.

The proposed approach addresses key challenges in fashion recommendation, namely the high dimensionality of visual data and the need to capture complex relationships. By compressing visual features [13, 18, 19, 20, 22, 38], the framework mitigates the computational burden associated with high-dimensional inputs, enhancing scalability and efficiency [5, 9]. Simultaneously, constructing a heterogeneous graph that incorporates user-item interactions and various item-item relationships (including visual similarity and style compatibility) provides a rich structure for graph-based learning [6, 8, 14, 16, 40, 41, 42, 43, 44]. GNNs, by operating on this graph and utilizing the compressed visual features as node attributes, can learn powerful embeddings that capture both structural and visual information, leading to potentially more accurate and relevant recommendations [5, 6, 8, 9].

The anticipated benefits include improved recommendation accuracy, particularly in capturing visually-driven preferences and item compatibility [2, 10, 25, 36, 37, 43, 44]. The approach is also expected to enhance efficiency and scalability, making it suitable for large-scale e-commerce platforms. Furthermore, the explicit inclusion of visual features and graph structure is likely to improve the handling of cold-start scenarios for new items [23, 24, 27].

However, several challenges need to be addressed for the successful implementation of this framework. Determining the optimal compression techniques and the degree of compression for visual features is crucial to ensure that essential information is retained while achieving efficiency. Constructing and maintaining a large, heterogeneous graph with diverse edge types can be complex, requiring robust data pipelines. Training GNNs on large graphs can be computationally intensive, although advancements in distributed graph processing and GNN architectures are continuously improving scalability [35, 51]. While potentially more interpretable than some deep learning models, understanding the exact reasoning behind GNN recommendations can still be challenging [5, 58, 59], necessitating further research into explainability techniques for graph-based recommendation [55, 56, 57]. Ensuring fairness in recommendations, avoiding biases that might be present in the data or learned by the model, is also a critical consideration [45].

Future research directions include exploring advanced visual feature compression techniques specifically optimized for graph-based learning, investigating dynamic graph construction and GNN models that can adapt to evolving user preferences and item trends, and incorporating other modalities such as text descriptions or user reviews into the graph representation. Developing robust evaluation methodologies for visual-aware recommendation that go beyond traditional metrics to assess aspects like style consistency and outfit compatibility is also important.

In conclusion, the integration of compressed visual features and graph-based learning using GNNs offers a promising direction for developing efficient and highly accurate visual-aware fashion recommendation systems. By effectively managing the dimensionality of visual data and leveraging the power of graph structures to model complex relationships, this approach has the potential to significantly enhance the user experience in online fashion retail and drive business value. Addressing the associated implementation challenges will be key to unlocking the full potential of this framework in real-world applications.

REFERENCES

Chen, H.J.; Shuai, H.H.; Cheng, W.H. A survey of artificial intelligence in fashion. IEEE Signal Process. Mag. 2023, 40, 64–73.

Hidayati, S.C.; Goh, T.W.; Chan, J.S.G.; Hsu, C.C.; See, J.; Wong, L.K.; Hua, K.L.; Tsao, Y.; Cheng, W.H. Dress with style: Learning style from joint deep embedding of clothing styles and body shapes. IEEE Trans. Multimed. 2020, 23, 365–377.

Ding, Y.; Lai, Z.; Mok, P.; Chua, T.S. Computational Technologies for Fashion Recommendation: A Survey. ACM Comput. Surv. 2023, 56, 121.

Zanker, M.; Rook, L.; Jannach, D. Measuring the impact of online personalisation: Past, present and future. Int. J. Hum.-Comput. Stud. 2019, 131, 160–168.

Markchom, T.; Liang, H.; Ferryman, J. Scalable and explainable visually-aware recommender systems. Knowl.-Based Syst. 2023, 263, 110258.

Gao, C.; Zheng, Y.; Li, N.; Li, Y.; Qin, Y.; Piao, J.; Quan, Y.; Chang, J.; Jin, D.; He, X.; et al. A survey of graph neural networks for recommender systems: Challenges, methods, and directions. ACM Trans. Recomm. Syst. 2023, 1, 3.

Wang, L.; Guo, D.; Liu, X. Research on Intelligent Recommendation Technology for Complex Tasks. In Proceedings of the 2023 4th IEEE International Conference on Computer Engineering and Application (ICCEA), Hangzhou, China, 7–9 April 2023; pp. 353–360.

Dossena, M.; Irwin, C.; Portinale, L. Graph-based recommendation using graph neural networks. In

Proceedings of the 2022 21st IEEE International Conference on Machine Learning and Applications (ICMLA), Nassau, Bahamas, 12–14 December 2022; pp. 1769–1774.

Zzh; Zhang, W.; Wentao. Industrial Solution in Fashion-Domain Recommendation by an Efficient Pipeline Using GNN and Lightgbm. In Proceedings of the Recommender Systems Challenge, Seattle, WA, USA, 18–23 September 2022; pp. 45–49.

Vasileva, M.I.; Plummer, B.A.; Dusad, K.; Rajpal, S.; Kumar, R.; Forsyth, D. Learning type-aware embeddings for fashion compatibility. In Proceedings of the European Conference on Computer Vision (ECCV), Munich, Germany, 8–14 September 2018; pp. 390–405.

Kuang, Z.; Gao, Y.; Li, G.; Luo, P.; Chen, Y.; Lin, L.; Zhang, W. Fashion retrieval via graph reasoning networks on a similarity pyramid. In Proceedings of the IEEE/CVF International Conference on Computer Vision, Seoul, Republic of Korea, 27 October–2 November 2019; pp. 3066–3075.

Marcuzzo, M.; Zangari, A.; Albarelli, A.; Gasparetto, A. Recommendation systems: An insight into current development and future research challenges. IEEE Access 2022, 10, 86578–86623.

Ferreira, D.; Silva, S.; Abelha, A.; Machado, J. Recommendation system using autoencoders. Appl. Sci. 2020, 10, 5510.

Xie, Y.; Yao, C.; Gong, M.; Chen, C.; Qin, A.K. Graph convolutional networks with multi-level coarsening for graph classification. Knowl.-Based Syst. 2020, 194, 105578.

Papadakis, H.; Papagrigoriou, A.; Panagiotakis, C.; Kosmas, E.; Fragopoulou, P. Collaborative filtering recommender systems taxonomy. Knowl. Inf. Syst. 2022, 64, 35–74.

Cai, H.; Zheng, V.W.; Chang, K.C.C. A comprehensive survey of graph embedding: Problems, techniques, and applications. IEEE Trans. Knowl. Data Eng. 2018, 30, 1616–1637.

Georgiou, T.; Liu, Y.; Chen, W.; Lew, M. A survey of traditional and deep learning-based feature descriptors for high dimensional data in computer vision. Int. J. Multimed. Inf. Retr. 2020, 9, 135–170.

Jing, L.; Vincent, P.; LeCun, Y.; Tian, Y. Understanding dimensional collapse in contrastive self-supervised learning. In Proceedings of the 10th International Conference on Learning Representations, ICLR 2022, Virtual, 25–29 April 2022.

Tao, Y.; Guo, K.; Zheng, Y.; Pan, S.; Cao, X.; Chang, Y. Breaking the curse of dimensional collapse in graph contrastive learning: A whitening perspective. Inf. Sci. 2024, 657, 119952.

Tran, B.; Tran, D.; Nguyen, H.; Ro, S.; Nguyen, T. scCAN: Single-cell clustering using autoencoder and network fusion. Sci. Rep. 2022, 12, 10267.

Yan, C.; Malhi, U.S.; Huang, Y.; Tao, R. Unsupervised deep clustering for fashion images. In Proceedings of the Knowledge Management in Organizations: 14th International Conference, KMO 2019, Zamora, Spain, 15–18 July 2019; Proceedings 14. Springer: Berlin/Heidelberg, Germany, 2019; pp. 85–96.

Malhi, U.S.; Zhou, J.; Yan, C.; Rasool, A.; Siddeeq, S.; Du, M. Unsupervised Deep Embedded Clustering for High-Dimensional Visual Features of Fashion Images. Appl. Sci. 2023, 13, 2828.

He, R.; McAuley, J. Ups and downs: Modeling the visual evolution of fashion trends with one-class collaborative filtering. In Proceedings of the 25th International Conference on World Wide Web, Montreal, QC, Canada, 11–15 April 2016; pp. 507–517.

McAuley, J.; Targett, C.; Shi, Q.; Van Den Hengel, A. Image-based recommendations on styles and substitutes. In Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval, Santiago, Chile, 9–13 August 2015; pp. 43–52.

Sarkar, R.; Bodla, N.; Vasileva, M.; Lin, Y.L.; Beniwal, A.; Lu, A.; Medioni, G. Outfittransformer: Outfit representations for fashion recommendation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, New Orleans, LA, USA, 18–24 June 2022; pp. 2263–2267.

Deldjoo, Y.; Nazary, F.; Ramisa, A.; Mcauley, J.; Pellegrini, G.; Bellogin, A.; Noia, T.D. A review of modern fashion recommender systems. ACM Comput. Surv. 2023, 56, 1–37.

Cardoso, Â.; Daolio, F.; Vargas, S. Product characterisation towards personalisation: Learning attributes from unstructured data to recommend fashion products. In Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, London, UK, 19–23 August 2018; pp. 80–89.

Hong, W.; Li, S.; Hu, Z.; Rasool, A.; Jiang, Q.; Weng, Y. Improving relation extraction by knowledge representation learning. In Proceedings of the 2021 IEEE 33rd International Conference on Tools with Artificial Intelligence (ICTAI), Washington, DC, USA, 1–3 November 2021; pp. 1211–1215.

Jagadeesh, V.; Piramuthu, R.; Bhardwaj, A.; Di, W.; Sundaresan, N. Large scale visual recommendations from street fashion images. In Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, New York, NY, USA, 24–27 August 2014; pp. 1925–1934.

Deldjoo, Y.; Di Noia, T.; Malitesta, D.; Merra, F.A. A study on the relative importance of convolutional neural networks in visually-aware recommender systems. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Nashville, TN, USA, 20–25 June 2021; pp. 3961–3967.

He, R.; McAuley, J. VBPR: Visual bayesian personalized ranking from implicit feedback. In Proceedings of the AAAI Conference on Artificial Intelligence, Phoenix, AZ, USA, 12–17 February 2016; Volume 30.

Liu, Q.; Wu, S.; Wang, L. Deepstyle: Learning user preferences for visual recommendation. In Proceedings of the 40th International Acm Sigir Conference on Research and Development in Information Retrieval, Shinjuku, Japan, 7–11 August 2017; pp. 841–844.

Yan, C.; Chen, Y.; Zhou, L. Differentiated fashion recommendation using knowledge graph and data augmentation. IEEE Access 2019, 7, 102239–102248.

Yu, W.; He, X.; Pei, J.; Chen, X.; Xiong, L.; Liu, J.; Qin, Z. Visually aware recommendation with aesthetic features. VLDB J. 2021, 30, 495–513.

Dong, M.; Zeng, X.; Koehl, L.; Zhang, J. An interactive knowledge-based recommender system for fashion product design in the big data environment. Inf. Sci. 2020, 540, 469–488.

Li, Y.; Cao, L.; Zhu, J.; Luo, J. Mining fashion outfit composition using an end-to-end deep learning approach on set data. IEEE Trans. Multimed. 2017, 19, 1946–1955.

Liu, L.; Du, X.; Zhu, L.; Shen, F.; Huang, Z. Learning discrete hashing towards efficient fashion recommendation. Data Sci. Eng. 2018, 3, 307–322.

Cohen-Shapira, N.; Rokach, L. Learning dataset representation for automatic machine learning algorithm selection. Knowl. Inf. Syst. 2022, 64, 2599–2635.

Yi, J.; Chen, Z. Multi-modal variational graph auto-encoder for recommendation systems. IEEE Trans. Multimed. 2021, 24, 1067–1079.

Ma, M.; Na, S.; Wang, H. AEGCN: An autoencoder-constrained graph convolutional network. Neurocomputing 2021, 432, 21–31.

Kipf, T.N.; Welling, M. Variational Graph Auto-Encoders. arXiv 2016, arXiv:1611.07308.

Veit, A.; Kovacs, B.; Bell, S.; McAuley, J.; Bala, K.; Belongie, S. Learning visual clothing style with heterogeneous dyadic co-occurrences. In Proceedings of the IEEE International Conference on Computer Vision, Santiago, Chile, 7–13 December 2015; pp. 4642–4650.

Yin, R.; Li, K.; Lu, J.; Zhang, G. Enhancing fashion recommendation with visual compatibility relationship. In Proceedings of the World Wide Web Conference, San Francisco, CA, USA, 13–17 May 2019; pp. 3434–3440.

Borges, R.; Stefanidis, K. Feature-blind fairness in collaborative filtering recommender systems. Knowl. Inf. Syst. 2022, 64, 943–962.

Chatfield, K.; Simonyan, K.; Vedaldi, A.; Zisserman, A. Return of the devil in the details: Delving deep into convolutional nets. arXiv 2014, arXiv:1405.3531.

Jiang, B.; Zhang, Z.; Lin, D.; Tang, J.; Luo, B. Semi-supervised learning with graph learning-convolutional networks. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Long Beach, CA, USA, 15–20 June 2019; pp. 11313–11320.

Garcia, V.; Bruna, J. Few-shot learning with graph neural networks. arXiv 2017, arXiv:1711.04043.

Xiao, Z.; Deng, Y. Graph embedding-based novel protein interaction prediction via higher-order graph convolutional network. PLoS ONE 2020, 15, e0238915.

Srivastava, N.; Hinton, G.; Krizhevsky, A.; Sutskever, I.; Salakhutdinov, R. Dropout: A simple way to prevent neural networks from overfitting. J. Mach. Learn. Res. 2014, 15, 1929–1958.

Ioffe, S.; Szegedy, C. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In Proceedings of the International Conference on Machine Learning, Lille, France, 6–11 July 2015; pp. 448–456.

Hisano, R. Semi-supervised graph embedding approach to dynamic link prediction. In Complex Networks IX: Proceedings of the 9th Conference on Complex Networks CompleNet 2018; Springer: Berlin/Heidelberg, Germany, 2018; pp. 109–121.

Janssens, A.C.J.; Martens, F.K. Reflection on modern methods: Revisiting the area under the ROC Curve. Int. J. Epidemiol. 2020, 49, 1397–1403.

Goodfellow, I.; Bengio, Y.; Courville, A. Deep Learning; MIT Press: Cambridge, MA, USA, 2016.

Samek, W.; Montavon, G.; Lapuschkin, S.; Anders, S.; Müller, K.R. Explaining Deep Neural Networks and Beyond: A Review of Methods and Applications. Proc. IEEE 2021, 109, 247–278.

Oprescu, A.; Miró-Amarante, G.; García-Díaz, L.; Rey, V.; Chimenea-Toscano, A.; Martínez-Martínez, R.; Romero-Ternero, M. Towards a data collection methodology for Responsible Artificial Intelligence in health: A prospective and qualitative study in pregnancy. Inf. Fusion 2022, 83–84, 53–78.

High-Level Expert Group on Artificial Intelligence. Ethics guidelines for trustworthy AI [Internet]. European Commission. 2019. Available online: https://digital-strategy.ec.europa.eu/en/library/ethics-guidelines-trustworthy-ai (accessed on 28 August 2024).

Oprescu, A.M.; Miró-Amarante, G.; García-Díaz, L.; Beltrán, L.M.; Rey, V.E.; Romero-Ternero, M. Artificial intelligence in pregnancy: A scoping review. IEEE Access 2020, 8, 181450–181484.

Clement, T.; Kemmerzell, N.; Abdelaal, M.; Amberg, M.

XAIR: A Systematic Metareview of Explainable AI (XAI) Aligned to the Software Development Process. Mach. Learn. Knowl. Extr. 2023, 5, 78–108.

Mohammad Amini, M.; Jesus, M.; Fanaei Sheikholeslami, D.; Alves, P.; Hassanzadeh Benam, A.; Hariri, F. Artificial Intelligence Ethics and Challenges in Healthcare Applications: A Comprehensive Review in the Context of the European GDPR Mandate. Mach. Learn. Knowl. Extr. 2023, 5, 1023–1035.