This CVPR Workshop paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

A Fashion Item Recommendation Model in Hyperbolic Space

Ryotaro Shimizu^{1,2}, Yu Wang¹, Masanari Kimura³, Yuki Hirakawa², Takashi Wada², Yuki Saito², Julian McAuley¹ ¹University of California, San Diego ² ZOZO Research ³ The University of Melbourne {r2shimizu, yuw164}@ucsd.edu,m.kimura@unimelb.edu.au,

{yuki.hirakawa,takashi.wada,yuki.saito}@zozo.com,jmcauley@ucsd.edu

Abstract

In this work, we propose a fashion item recommendation model that incorporates hyperbolic geometry into user and item representations. Using hyperbolic space, our model aims to capture implicit hierarchies among items based on their visual data and users' purchase history. During training, we apply a multi-task learning framework that considers both hyperbolic and Euclidean distances in the loss function. Our experiments on three data sets show that our model performs better than previous models trained in Euclidean space only, confirming the effectiveness of our model. Our ablation studies show that multi-task learning plays a key role, and removing the Euclidean loss substantially deteriorates the model performance.

1. Introduction

In hyperbolic space, the distance from the origin increases exponentially as one moves towards the outside of the space. This property makes it well-suited for modeling hierarchical data by representing the root node around the origin and leaf nodes near the surface of the space. In practice, hyperbolic space has been used to represent various types of instances that potentially possess hierarchical structures, such as sentences and words [25, 46], questions and answers [44], and objects and scenes [18]. In computer vision, its effectiveness has been confirmed on many tasks [15, 36], including classification [10, 12, 31], retrieval [14, 28, 53, 56], segmentation [3, 8, 18], and generation [27].

In this paper, we propose a new fashion item recommendation model that incorporates hyperbolic geometry into user and item representations; we name it Hyperbolic Visual Attentive Collaborative Filtering (HVACF). Using hyperbolic space, our model aims to capture implicit hierarchies among fashion items based on their image data and users' purchase history. To train our model, we apply a



Figure 1. Overview of our proposed model.

multi-task learning framework that combines the hyperbolic distance with the Euclidean one, which was found effective in previous studies for training recommender systems without visual features and also for learning representations of objects and scenes [18, 50]. Our experiments show that our model outperforms existing models trained in Euclidean space only, confirming the effectiveness of our model. We also conduct ablation studies and find that multi-task learning plays a key role, and removing the Euclidean loss results in poor performance, even worse than baseline systems without visual features. This result suggests that multi-task learning is crucial for training fashion item recommender systems, in contrast to the previous findings on recommendation tasks without visual features that hyperbolic models perform better than the Euclidean counterparts [7, 49].

2. Related Work

2.1. Visually-Aware Recommender Systems

Visually-aware recommendation models leverage item visual data to capture characteristics of items and users [9, 22, 24, 30]. Visual information is particularly valuable for recommending items whose appearances and styles heavily

affect users' preferences, such as clothing [24, 30, 40], cosmetic products [1, 55], and hair-style [29, 54] [11]. In this work, we incorporate hyperbolic geometry into a fashion item recommender system with visual features and demonstrate its effectiveness.

2.2. Recommender Systems in Hyperbolic Space

Hyperbolic geometry differs from Euclidean geometry in that it measures distances on curves rather than straight lines. In hyperbolic space, the distance from the origin increases exponentially as one moves towards the surface of the space. This property is well-suited for modeling hierarchies and has been used for representation learning on various tasks [2, 5, 6, 16, 19, 35, 38, 45, 47]. Previous work also showed that it is effective in capturing the complex interactions between items and user preferences in recommender systems [7, 36, 39, 50, 52]. For instance, Tran et al. [49] proposed a method based on hyperbolic space, called Hyperbolic Bayesian Personalized Ranking (BPR), and Benjamin et al. [7] extended this model using largescale data. Lucas et al. [50] proposed Hyperbolic Metric Learning (HML), which incorporates the hyperbolic distance into metric-learning models [23, 43]. Other studies proposed recommendation models that incorporate hyperbolic geometry into graph-based methods [26, 41, 51, 57]; temporal point process [58]; or autoencoders [34]. However, to the best of our knowledge, hyperbolic space has not been used on recommendation tasks where the visual data of items provides critical information, such as fashion item recommendation.

3. Methodology

The overview of our proposed model is illustrated in Figure 1. The training data of our model is constructed as a set of triplets of users and positive/negative items, and is defined as follows:

$$D_S = \{(i, j, k) | i \in \mathcal{U}, j \in \mathcal{I}_i^+, k \in \mathcal{I} \setminus \mathcal{I}_i^+\}, \qquad (1)$$

where \mathcal{U}, \mathcal{I} , and \mathcal{I}_i^+ are the sets of users, items, and items purchased by the *i*-th user, respectively.

3.1. Training a Model in Hyperbolic Space

Among several isometric models for modeling hyperbolic space, this paper adopts the Poincaré ball model, one of the most common models in computer vision [15, 18, 20, 33]. In this paper, the induced geodesic distance between \boldsymbol{x} and \boldsymbol{y} in hyperbolic space $\mathbb{D}_c^n \coloneqq \{\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^n : c \|\boldsymbol{x}\|^2, c \|\boldsymbol{y}\|^2 < 1\}$ for $c \ge 0$ is defined as follows:

$$d_c(\boldsymbol{x}, \boldsymbol{y}) \coloneqq \frac{2}{\sqrt{c}} \tanh^{-1} \left(\sqrt{c} \| (-\boldsymbol{x}) \oplus_c \boldsymbol{y} \| \right), \quad (2)$$

$$\boldsymbol{x} \oplus_{c} \boldsymbol{y} \coloneqq \frac{(1 + 2c\langle \boldsymbol{x}, \boldsymbol{y} \rangle + c \|\boldsymbol{y}\|^{2})\boldsymbol{x} + (1 - c \|\boldsymbol{x}\|^{2})\boldsymbol{y}}{1 + 2c\langle \boldsymbol{x}, \boldsymbol{y} \rangle + c^{2} \|\boldsymbol{x}\|^{2} \|\boldsymbol{y}\|^{2}}.$$
(3)

In addition, the following function maps a point $z \in T_q \mathbb{D}_c^n$ from a *n*-dimensional vector space $T_q \mathbb{D}_c^n$, which is a tangent space of \mathbb{D}_c^n at a point $q \in \mathbb{D}_c^n$, to the hyperbolic space \mathbb{D}_c^n .

$$h(\boldsymbol{z}) = \exp_{\boldsymbol{q}}^{c}(\boldsymbol{z}) \coloneqq \boldsymbol{q} \oplus_{c} \Big(\tanh\left(\sqrt{c} \frac{\lambda_{\boldsymbol{q}}^{c} \|\boldsymbol{z}\|^{2}}{2}\right) \frac{\boldsymbol{z}}{\sqrt{c} \|\boldsymbol{z}\|} \Big),$$
(4)

where $\lambda_{q}^{c} = \frac{2}{1-c||q^{2}||}$ is the conformal factor [17]. Detailed definitions can be found in the appendix or previous studies such as [15, 36].

3.2. Neighbor-Attentive Aggregation

Our proposed model extracts user preferences based on their purchase history and image data of each item. Here, $\mathbf{U} = \{[\boldsymbol{u}_1; ...; \boldsymbol{u}_n; ...; \boldsymbol{u}_{N_u}] | \boldsymbol{u}_n \in \mathbb{R}^D\} \in \mathbb{R}^{N_u \times D}$ denotes the user embeddings, and $\mathbf{V} = \{[\boldsymbol{v}_1; ...; \boldsymbol{v}_n; ...; \boldsymbol{v}_{N_v}] | \boldsymbol{v}_n \in \mathbb{R}^D\} \in \mathbb{R}^{N_v \times D}$ and $\mathbf{P} = \{[\boldsymbol{p}_1; ...; \boldsymbol{p}_n; ...; \boldsymbol{p}_{N_v}] | \boldsymbol{p}_n \in \mathbb{R}^D\} \in \mathbb{R}^{N_v \times D}$ denote separate embeddings associated with the items. N_u and N_v are the numbers of users and items, respectively; D is the embedding dimension; and $[\boldsymbol{a}; \boldsymbol{b}]$ denotes the concatenation operation between the vectors \boldsymbol{a} and \boldsymbol{b} .

First, we calculate the embedding \tilde{u}_i that combines the *i*-th user embedding and item embeddings purchased by the user, inspired by [9], as follows:

$$\tilde{\boldsymbol{u}}_i = \frac{\boldsymbol{u}_i + \sum_{l \in A(\mathcal{I}_i^+ \setminus j, L)} \alpha_{il} \boldsymbol{p}_l}{2}, \quad (5)$$

where $A(\mathcal{I}_i^+ \setminus j, L)$ is an item set obtained by randomly sampling L > 0 items from $\mathcal{I}_i^+ \setminus j$.¹ Here, α_{il} represents the contribution (attention) score of the *l*-th item to the preference profile of the *i*-th user, and is defined as follows:

$$\alpha_{il} = \frac{\exp(\alpha'_{il}/\tau)}{\sum_{n \in A(\mathcal{I}_i^+ \setminus j, L)} \exp(\alpha'_{in}/\tau)},$$
(6)

$$\alpha_{il}' = \boldsymbol{w}_{2}^{\top} \phi(\mathbf{W}_{u} \boldsymbol{u}_{i} + \mathbf{W}_{v} \boldsymbol{v}_{l} + \mathbf{W}_{p} \boldsymbol{p}_{l} + \mathbf{W}_{f} E(\boldsymbol{I}_{l}) + \boldsymbol{b}_{1}) + \boldsymbol{b}_{2}, \qquad (7)$$

where τ is a scaling factor (the larger the value is, the sharper the attention score becomes), which we set $\tau = \sqrt{D}$, and $\phi(x) = \max(0, x)$ is the ReLU function. $E(\cdot)$ is a pre-trained backbone encoder model and I_l is the *l*-th item's image data. $\mathbf{W}_u, \mathbf{W}_v, \mathbf{W}_p \in \mathbb{R}^{D \times D}$, and $\mathbf{W}_f \in \mathbb{R}^{D \times D_{\text{pool}}}$ denote transformation matrices for the users, items, and item visual features extracted by the image encoder, respectively (D_{pool} is dimension size of the final pooling layer in the backbone model); $b_1 \in \mathbb{R}^D$ is the bias vector of the first layer; and $w_2 \in \mathbb{R}^D$ and $b_2 \in \mathbb{R}$ are the weighting vector and bias of the second layer.

¹If
$$L > |\mathcal{I}_i^+ \setminus j|$$
, we retrieval all the items from $\mathcal{I}_i^+ \setminus j$.

Table 1. Statistics of the data sets used in our experiments on recommendation tasks.

Dataset	#users	#items	#feedbacks
Amazon Women [21]	30,631	285,269	849,801
Amazon Men [21]	27,892	120,758	451,932
TOWN Women [59] ²	37,357	187,866	800,000

Table 2. Comparison of our baselines; "Vis" and "Hyp" denote models with visual features and models in hyperbolic space, respectively.

Method	Vis	Нур	Method	Vis	Нур
BPRMF [37]	X	X	VBPR [22]	1	X
HBPR [49]	X	1	DVBPR [24]	\checkmark	X
HRec [7]	X	1	DeepStyle [30]	\checkmark	X
LFM [52]	X	1	ACF [9]	\checkmark	X
HML [50]	X	1	HVACF	1	1

3.3. Objective Function

Following [50], the target parameters $\Theta = \{\mathbf{U}, \mathbf{V}, \mathbf{P}, \mathbf{W}_u, \mathbf{W}_v, \mathbf{W}_p, \mathbf{W}_f, b_1, w_2, b_2, q\}$ are optimized as follows:

$$\arg\min_{\boldsymbol{\Theta}} \mathcal{L} = \mathcal{L}_{hyp} + \gamma \mathcal{L}_{adj} + \lambda \|\boldsymbol{\Theta}\|^2, \qquad (8)$$

where $\gamma > 0$ is a scalar hyperparameter that balances the loss based on the hyperbolic distance (\mathcal{L}_{hyp}) and the adjustment loss based on the Euclidean distance (\mathcal{L}_{adj}), and $\lambda > 0$ determines the degree of regularizing the model parameters (Θ). Each loss function is defined as follows:

$$\mathcal{L}_{\text{hyp}} \coloneqq \sum_{(i,j,k)\in D_S} \left[m + d_c^2(h(\tilde{\boldsymbol{u}}_i), h(\boldsymbol{v}_j)) - d_c^2(h(\tilde{\boldsymbol{u}}_i), h(\boldsymbol{v}_k)) \right]_+,$$
(9)

$$\mathcal{L}_{\mathrm{adj}} \coloneqq \sum_{(i,j)\in D_S} f_{\mathrm{adj}}(\tilde{\boldsymbol{u}}_i, \boldsymbol{v}_j) + \sum_{(i,k)\in D_S} f_{\mathrm{adj}}(\tilde{\boldsymbol{u}}_i, \boldsymbol{v}_k), \quad (10)$$

$$f_{\rm adj}(\boldsymbol{x}, \boldsymbol{y}) \coloneqq \left[\frac{|d_c(h(\boldsymbol{x}), h(\boldsymbol{y})) - d_{\rm euc}(\boldsymbol{x}, \boldsymbol{y})|}{d_{\rm euc}(\boldsymbol{x}, \boldsymbol{y})}\right]_+, \quad (11)$$

where $[a]_+$ denotes $\max(0, a)$; m > 0 is the margin in the hinge function; and $d_{\text{euc}}(a, b)$ is the Euclidean distance. During inference, the model recommends the items with the smallest hyperbolic distances to the user vector \tilde{u}_i .

4. Experiments

4.1. Experimental Settings

We perform experiments using three data sets for fashion item recommendation tasks (shown in Table 1). Each data

Table 3. Results of our experiments (*DStyle = DeepStyle). The numbers denote the AUC scores of each baseline and our model.

Method	Amazom (W)	Amazon (M)	TOWN (W)
Rand	0.499 ± 0.002	0.499 ± 0.001	0.499 ± 0.002
BPRMF	0.702 ± 0.002	0.717 ± 0.003	0.741 ± 0.001
HBPR	0.711 ± 0.001	0.743 ± 0.001	0.716 ± 0.001
HRec	0.762 ± 0.000	0.808 ± 0.002	0.784 ± 0.002
LFM	0.679 ± 0.001	0.703 ± 0.002	0.685 ± 0.002
HML	0.777 ± 0.002	0.818 ± 0.000	$0.795 \ {\pm} 0.001$
VBPR	0.730 ± 0.000	0.756 ± 0.002	0.728 ± 0.001
DVBPR	0.776 ± 0.001	0.818 ± 0.001	0.790 ± 0.000
DStyle	0.728 ± 0.001	0.730 ± 0.004	0.700 ± 0.004
ACF	0.767 ± 0.001	0.797 ± 0.002	0.772 ± 0.000
HVACF	0.804 ±0.001	$\textbf{0.830} \pm 0.002$	0.810 ±0.001

Table 4. Results of our ablation studies.

Variants	Amazon (W)	Amazon (M)	TOWN (W)
(complete)	$\textbf{0.804} \pm 0.001$	$\textbf{0.830} \pm 0.002$	$\textbf{0.810} \pm 0.001$
$d_c(\cdot) \Rightarrow d_{\rm euc}(\cdot)$	$0.755 \ {\pm} 0.001$	0.790 ± 0.001	$0.775 \ {\pm} 0.002$
w/o $\mathcal{L}_{\mathrm{adj}}$	$0.701 \ {\pm} 0.001$	$0.749 \ {\pm} 0.001$	0.722 ± 0.001
w/o aggregation	$0.745 \ {\pm} 0.004$	0.809 ± 0.002	0.779 ± 0.003
w/o attention	0.798 ± 0.000	0.827 ± 0.001	0.807 ± 0.001
attention w/o $E(I_l)$	0.798 ± 0.001	0.826 ± 0.001	0.808 ± 0.001
attention w/o v_l	0.799 ± 0.001	0.828 ± 0.001	0.808 ± 0.001
attention w/o p_l	0.798 ± 0.000	0.827 ± 0.002	0.807 ± 0.001



Figure 2. Performance with different hyperparameters.

set is chronologically sorted and then divided into the train, validation, and test sets with a ratio of 7:1:2, respectively. Besides, the ten previous models listed in Table 2 plus a random recommendation model are used for our baselines. These baselines are chosen from the "visually-aware recommendation models" listed in [11] and the "hyperbolic recommender systems" listed in [36]. The area under the ROC curve (AUC) is used for evaluation, following previous work [22, 37].

4.2. Implementation Details

We use Riemannian Adam [4] as the optimizer for HBPR, HRec, and LFM models based on the performance on the validation set, and use Adam [13] with decoupled weight decay [32] for the other models. We set the embedding dimension D to 50 and the batch size

²ZOZOTOWN is the largest fashion e-commerce application in Japan.



Figure 3. The embedding-norm distribution on Amazon Women.

Figure 4. User and item 2D-representations compressed by tSNE [48] on Amazon Women.

to 512 for all models. We tune the learning rate $r \in \{0.01, 0.001, 0.0001\}$ and the hyperparameter for regularization $\lambda \in \{0.1, 0.01, 0.001\}$ for each model based on the performance on the validation set. Inception V3 [42] is used as the backbone image encoder model for all models with visual features, and its parameters are freezed during training. For HVACF, we set $\gamma = 0.5$, c = 1.0, L = 32, and m = 0.5 based on the performance on the validation set. We repeat each experiment three times with different random seeds, and report the mean and standard deviation values.

4.3. Results

Table 3 shows the results. It demonstrates that our proposed model outperforms all baselines on three data sets. These results highlight the effectiveness of using hyperbolic space for fashion item recommender systems.

4.4. Ablation Studies

We perform ablation studies and Table 4 summarizes the results. It shows that each component of our proposed model is indispensable to achieve high accuracy. Notably, the performance of "w/o L_{adj} " is lower than that of " $d_c(\cdot) \Rightarrow d_{euc}(\cdot)$ ": a model trained in Euclidean space only. The model "w/o \mathcal{L}_{adj} " even underperforms baselines without visual features. These results indicate the importance of considering both Euclidean and hyperbolic distances to train fashion item recommendation models.

4.5. Impacts of Hyperparameters

Figure 2a and 2b illustrate how the hyperparameters of our model affect its performance. Figure 2a shows that setting γ to a small value harms the performance, indicating that it is important to consider both the hyperbolic and Euclidean distances during training. Figure 2b shows that when c is set too large, the performance drops sharply as a result of the radius of the embedding space (ball) $1/\sqrt{c}$

getting too small, which causes both users and items to be placed near the ball's surface during training, and results in inefficient training.

4.6. Analysis of Embeddings

Figure 3 shows the norms of the user and item embeddings trained on the Amazon Women data set in hyperbolic space (i.e., $h(\tilde{u})$ and h(v), respectively), which correspond to the distance distributions from the ball's origin of the space. It clearly indicates that items and users are separated in the embedding space, with the items being placed near the origin and the users mapped near the ball's surface.

Next, we analyse the 2D-compressed representations of the users and items in 4a-4c, where each plot is colored based on the logarithmically transformed value of the number of the purchases or purchases made. We can see that items with many purchases are mapped near the user embeddings. We also calculate the correlation coefficient between the norms of item embeddings and the number of purchases made, and observe that it is -0.360. This indicates that popular items tend to be mapped near the origin, and our proposed model separates items largely based on how popular they are among users.

5. Conclusion

In this paper, we proposed HVACF, a fashion item recommendation model trained in hyperbolic space. Our experiments on three data sets showed that it performs better than previous models trained in Euclidean space, confirming the effectiveness of our model. We also performed ablation studies and showed that it is crucial to consider both hyperbolic and Euclidean distances in the objective function. Lastly, we analysed the trained embeddings of users and items and found that items and users are separated in hyperbolic space. We also found that popular and unpopular items are also separated in the space, with the popular ones being centered near the origin.

References

- Taleb Alashkar, Songyao Jiang, Shuyang Wang, and Yun Fu. Examples-rules guided deep neural network for makeup recommendation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 31, 2017. 2
- [2] Mark Crovella Andrej Cvetkovski. Multidimensional scaling in the poincaré disk. *Applied Mathematics Information Sciences*, 10(1):125–133, 2016. 2
- [3] Mina Ghadimi Atigh, Julian Schoep, Erman Acar, Nanne van Noord, and Pascal Mettes. Hyperbolic image segmentation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 4453–4462, 2022. 1
- [4] Gary Becigneul and Octavian-Eugen Ganea. Riemannian adaptive optimization methods. In *Proceedings of the International Conference on Learning Representations*, 2019.
 3
- [5] Thomas Bläsius, Tobias Friedrich, Anton Krohmer, and Sören Laue. Efficient embedding of scale-free graphs in the hyperbolic plane. *IEEE/ACM Transactions on Networking*, 26(2):920–933, 2018. 2
- [6] Adrian E. Raftery Bolun Liu, Shane Lubold and Tyler H. McCormick. Bayesian hyperbolic multidimensional scaling. *Journal of Computational and Graphical Statistics*, 0:1–32, 2024. 2
- [7] Benjamin Paul Chamberlain, Stephen R Hardwick, David R Wardrope, Fabon Dzogang, Fabio Daolio, and Saúl Vargas. Scalable hyperbolic recommender systems. *arXiv preprint arXiv:1902.08648*, 2019. 1, 2, 3
- [8] Bike Chen, Wei Peng, Xiaofeng Cao, and Juha Röning. Hyperbolic uncertainty aware semantic segmentation. *IEEE Transactions on Intelligent Transportation Systems*, 25(2):1275–1290, 2024.
- [9] Jingyuan Chen, Hanwang Zhang, Xiangnan He, Liqiang Nie, Wei Liu, and Tat-Seng Chua. Attentive collaborative filtering: Multimedia recommendation with item- and component-level attention. In *Proceedings of the International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 335–344, 2017. 1, 2, 3
- [10] Jindou Dai, Yuwei Wu, Zhi Gao, and Yunde Jia. A hyperbolic-to-hyperbolic graph convolutional network. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 154–163, 2021. 1
- [11] Yashar Deldjoo, Fatemeh Nazary, Arnau Ramisa, Julian McAuley, Giovanni Pellegrini, Alejandro Bellogin, and Tommaso Di Noia. A review of modern fashion recommender systems. ACM Computing Surveys, 2023. 2, 3
- [12] Karan Desai, Maximilian Nickel, Tanmay Rajpurohit, Justin Johnson, and Shanmukha Ramakrishna Vedantam. Hyperbolic image-text representations. In *Proceedings of the International Conference on Machine Learning*, volume 202, pages 7694–7731, 2023. 1
- [13] John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12(61):2121–2159, 2011. 3

- [14] Jianping Fan, Daniel A. Keim, Yuli Gao, Hangzai Luo, and Zongmin Li. Justclick: Personalized image recommendation via exploratory search from large-scale flickr images. *IEEE Transactions on Circuits and Systems for Video Technology*, 19(2):273–288, 2009. 1
- [15] Pengfei Fang, Mehrtash Harandi, Trung Le, and Dinh Phung. Hyperbolic geometry in computer vision: A survey. arXiv preprint arXiv:2304.10764, 2023. 1, 2
- [16] Octavian Ganea, Gary Becigneul, and Thomas Hofmann. Hyperbolic entailment cones for learning hierarchical embeddings. In *Proceedings of the International Conference* on Machine Learning, volume 80, pages 1646–1655, 2018.
 2
- [17] Octavian Ganea, Gary Becigneul, and Thomas Hofmann. Hyperbolic neural networks. In Advances in Neural Information Processing Systems, volume 31, 2018. 2
- [18] Songwei Ge, Shlok Mishra, Simon Kornblith, Chun-Liang Li, and David Jacobs. Hyperbolic contrastive learning for visual representations beyond objects. In *Proceedings of* the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 6840–6849, 2023. 1, 2
- [19] Caglar Gulcehre, Misha Denil, Mateusz Malinowski, Ali Razavi, Razvan Pascanu, Karl Moritz Hermann, Peter Battaglia, Victor Bapst, David Raposo, Adam Santoro, and Nando de Freitas. Hyperbolic attention networks. In *Proceedings of the International Conference on Learning Representations*, 2019. 2
- [20] Yunhui Guo, Xudong Wang, Yubei Chen, and Stella X. Yu. Clipped hyperbolic classifiers are super-hyperbolic classifiers. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 1–10, 2022. 2
- [21] Ruining He and Julian McAuley. Ups and downs: Modeling the visual evolution of fashion trends with one-class collaborative filtering. In *Proceedings of the International Conference on World Wide Web*, pages 507–517, 2016. 3
- [22] Ruining He and Julian McAuley. Vbpr: Visual bayesian personalized ranking from implicit feedback. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 144– 150, 2016. 1, 3
- [23] Cheng-Kang Hsieh, Longqi Yang, Yin Cui, Tsung-Yi Lin, Serge Belongie, and Deborah Estrin. Collaborative metric learning. In *Proceedings of the International Conference on World Wide Web*, pages 193–201, 2017. 2
- [24] Wang-Cheng Kang, Chen Fang, Zhaowen Wang, and Julian McAuley. Visually-aware fashion recommendation and design with generative image models. In *Proceedings of the IEEE International Conference on Data Mining*, pages 207– 216, 2017. 1, 2, 3
- [25] Matthias Leimeister and Benjamin J. Wilson. Skip-gram word embeddings in hyperbolic space. In *Proceedings of* the International Conference on Learning Representations, 2019. 1
- [26] Anchen Li, Bo Yang, Huan Huo, Hongxu Chen, Guandong Xu, and Zhen Wang. Hyperbolic neural collaborative recommender. *IEEE Transactions on Knowledge and Data Engineering*, 35(9):9114–9127, 2023. 2
- [27] Lingxiao Li, Yi Zhang, and Shuhui Wang. The euclidean space is evil: Hyperbolic attribute editing for few-shot image

generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 22714–22724, 2023. 1

- [28] Feng Liu and Weiwei Guo. Personalized recommendation algorithm for interactive medical image using deep learning. *Mathematical Problems in Engineering*, 2022. 1
- [29] Luoqi Liu, Hui Xu, Junliang Xing, Si Liu, Xi Zhou, and Shuicheng Yan. "wow! you are so beautiful today!". In Proceedings of the ACM International Conference on Multimedia, pages 3–12, 2013. 2
- [30] Qiang Liu, Shu Wu, and Liang Wang. Deepstyle: Learning user preferences for visual recommendation. In *Proceedings* of the International ACM SIGIR Conference on Research and Development in Information Retrieval, pages 841–844, 2017. 1, 2, 3
- [31] Shaoteng Liu, Jingjing Chen, Liangming Pan, Chong-Wah Ngo, Tat-Seng Chua, and Yu-Gang Jiang. Hyperbolic visual embedding learning for zero-shot recognition. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2020. 1
- [32] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In Proceedings of the International Conference on Learning Representations, 2019. 3
- [33] Rongkai Ma, Pengfei Fang, Tom Drummond, and Mehrtash Harandi. Adaptive poincaré point to set distance for few-shot classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 1926–1934, 2022.
 2
- [34] Leyla Mirvakhabova, Evgeny Frolov, Valentin Khrulkov, Ivan Oseledets, and Alexander Tuzhilin. Performance of hyperbolic geometry models on top-n recommendation tasks. In *Proceedings of the ACM Conference on Recommender Systems*, pages 527–532, 2020. 2
- [35] Maximillian Nickel and Douwe Kiela. Poincaré embeddings for learning hierarchical representations. In Advances in Neural Information Processing Systems, volume 30, 2017.
 2
- [36] Wei Peng, Tuomas Varanka, Abdelrahman Mostafa, Henglin Shi, and Guoying Zhao. Hyperbolic deep neural networks: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(12):10023–10044, 2022. 1, 2, 3
- [37] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. Bpr: Bayesian personalized ranking from implicit feedback. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence*, pages 452–461, 2009.
 3
- [38] Frederic Sala, Chris De Sa, Albert Gu, and Christopher Re. Representation tradeoffs for hyperbolic embeddings. In *Proceedings of the International Conference on Machine Learn-ing*, volume 80, pages 4460–4469, 2018. 2
- [39] Timothy Schmeier, Joeseph Chisari, Sam Garrett, and Brett Vintch. Music recommendations in hyperbolic space: An application of empirical bayes and hierarchical poincaré embeddings. In *Proceedings of the ACM Conference on Recommender Systems*, pages 437–441, 2019. 2
- [40] Ryotaro Shimizu, Takuma Nakamura, and Masayuki Goto. Fashion-specific ambiguous expression interpretation with

partial visual-semantic embedding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 3497–3502, 2023. 2

- [41] Jianing Sun, Zhaoyue Cheng, Saba Zuberi, Felipe Perez, and Maksims Volkovs. Hgcf: Hyperbolic graph convolution networks for collaborative filtering. In *Proceedings of the Web Conference*, pages 593–601, 2021. 2
- [42] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2818–2826, 2016. 4
- [43] Yi Tay, Luu Anh Tuan, and Siu Cheung Hui. Latent relational metric learning via memory-based attention for collaborative ranking. In *Proceedings of the World Wide Web Conference*, pages 729–739, 2018. 2
- [44] Yi Tay, Luu Anh Tuan, and Siu Cheung Hui. Hyperbolic representation learning for fast and efficient neural question answering. In *Proceedings of the ACM International Conference on Web Search and Data Mining*, pages 583–591, 2018.
- [45] Yi Tay, Luu Anh Tuan, and Siu Cheung Hui. Hyperbolic representation learning for fast and efficient neural question answering. In *Proceedings of the ACM International Conference on Web Search and Data Mining*, pages 583–591, 2018.
 2
- [46] Alexandru Tifrea, Gary Becigneul, and Octavian-Eugen Ganea. Poincare glove: Hyperbolic word embeddings. In Proceedings of the International Conference on Learning Representations, 2019. 1
- [47] Alexandru Tifrea, Gary Becigneul, and Octavian-Eugen Ganea. Poincare glove: Hyperbolic word embeddings. In Proceedings of the International Conference on Learning Representations, 2019. 2
- [48] Laurens van der Maaten and Geoffrey Hinton. Visualizing Data using t-SNE. *Journal of Machine Learning Research*, 9(86):2579–2605, 2008. 4
- [49] Tran Dang Quang Vinh, Yi Tay, Shuai Zhang, Gao Cong, and Xiao-Li Li. Hyperbolic recommender systems. arXiv preprint arXiv:1809.01703, 2018. 1, 2, 3
- [50] Lucas Vinh Tran, Yi Tay, Shuai Zhang, Gao Cong, and Xiaoli Li. Hyperml: A boosting metric learning approach in hyperbolic space for recommender systems. In *Proceedings of the International Conference on Web Search and Data Mining*, pages 609–617, 2020. 1, 2, 3
- [51] Hao Wang, Defu Lian, Hanghang Tong, Qi Liu, Zhenya Huang, and Enhong Chen. Hypersorec: Exploiting hyperbolic user and item representations with multiple aspects for social-aware recommendation. ACM Transactions on Information Systems, 40(2), 2021. 2
- [52] Canran Xu and Ming Wu. Learning feature interactions with lorentzian factorization machine. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 6470–6477, 2020. 2, 3
- [53] Jiexi Yan, Lei Luo, Cheng Deng, and Heng Huang. Unsupervised hyperbolic metric learning. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 12465–12474, 2021. 1

- [54] Wei Yang, Masahiro Toyoura, and Xiaoyang Mao. Hairstyle suggestion using statistical learning. In Advances in Multimedia Modeling, pages 277–287, 2012. 2
- [55] Wei Yang, Masahiro Toyoura, and Xiaoyang Mao. Effect of facial makeup style recommendation on visual sensibility. *Multimedia Tools and Applications*, pages 843–853, 2014. 2
- [56] Yun Yue, Fangzhou Lin, Guanyi Mou, and Ziming Zhang. Understanding hyperbolic metric learning through hard negative sampling. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1891–1903, 2024. 1
- [57] Chunmei Zhang, Aoran Zhang, Li Zhang, Yonghong Yu, Weibin Zhao, and Hai Geng. A graph neural networksbased learning framework with hyperbolic embedding for personalized tag recommendation. *IEEE Access*, 12:339– 350, 2024. 2
- [58] Yuchen Zhou, Yanan Cao, Yanmin Shang, Chuan Zhou, Shirui Pan, Zheng Lin, and Qian Li. Explainable hyperbolic temporal point process for user-item interaction sequence generation. ACM Transactions on Information Systems, 41(4), 2023. 2
- [59] ZOZO, Inc. ZOZOTOWN. https://zozo.jp/. 2024 (accessed 1 March 2024). 3