

Personalized Compatibility Metric Learning

Meet Taraviya
UCLA
Los Angeles, USA

Yen-Liang Lin
Amazon
Palo Alto, USA

Anurag Beniwal
Amazon
Seattle, USA

Larry Davis
Amazon
New York, USA

ABSTRACT

Recommending sets of items that include both personalized and compatible items is crucial to personalized styling programs such as Amazon’s Personal Shopper. There is both an extensive literature on learning generic fashion compatibility and also on personalization in fashion. However, recommending pairs of items that the customer would like to wear together is still less studied as it involves learning a compatibility metric personalized to each customer. We propose a new framework (PSA-Net) to learn compatibility that is personalized to the customer - a customer dependent subspace learning framework where attention weights of subspaces are learnt using customer representations. We evaluate our approach on compatibility data provided directly by customers. Our approach outperforms the non-personalized approach in predicting compatibility preferences of customers. In other words, an approach that learns a common compatibility metric for all customers. In addition, we compare the significance of feedback collected directly from customers to that of data collected from human stylists in predicting compatibility for Amazon customers.

ACM Reference Format:

Meet Taraviya, Anurag Beniwal, Yen-Liang Lin, and Larry Davis. 2021. Personalized Compatibility Metric Learning. In *Proceedings of KDD 2021 (KDD’21)*. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/1122445.1122456>

1 INTRODUCTION

Personal shopper by Prime Wardrobe (PSPW) [1] is a subscription service that provides personalized styling using machine learning with humans in the loop to recommend sets of items to customers. Within the Personal shopper experience, *Love It or Leave It* is a fun and frictionless way for any Amazon customer to share their style likes and dislikes through a simple thumbs up/down interaction. Customer feedback is used to improve understanding of their style, resulting in better recommendations. Customers can vote items individually (Figure 1) as well as pairs of items (Figure 2). Votes on individual items are used to learn representations of customers and items to predict their preferences on individual items. Votes on pairs of items are used to learn compatibility of pairs of items for each customer. The feedback collected from Love it or leave it, in addition to other data sources, is used to train recommendation models. These models along with inputs from fashion experts are used to curate a personalized box for the customers. Customers can

choose what they want to try and are only charged for items they keep after trying. (See [1]).

Fashion compatibility is subjective and depends on several factors including the preference of customer for the individual items, and what they like to wear together. Depending on various factors like profession, location and age compatibility could mean different things to different customers, for e.g. for a particular customer black top on a white jeans is more compatible than a pink top on a white jeans and for another customer the opposite could be true. There is significant literature that addresses the problem of recommending compatible items but most learn a common compatibility metric for everyone ignoring personal tastes of individual customers. We propose a subspace attention based architecture that projects item images into multiple embedding subspaces. Different subspaces represent different notions of compatibility and the importance of each subspace in learning a customer’s compatibility metric is determined by its attention weight. The attention weights are computed using customer embeddings learnt from matrix completion-like algorithms. The number of subspaces is a hyperparameter in the model.

Our contributions are two fold 1) A novel architecture that takes into account the customer side information to project item embedding into multiple subspaces using attention to learn a personalized metric of compatibility for each customer 2) A two stage training process that learns customer representations from item level preferences in the first stage and then use these learnt customer representations along with pairwise preference data to predict pairwise compatibility in the second stage.

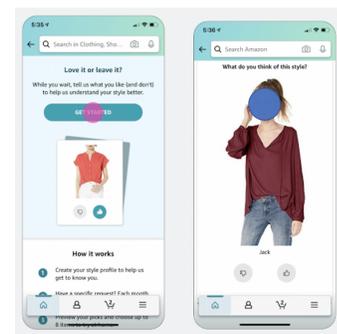


Figure 1: Item level feedback

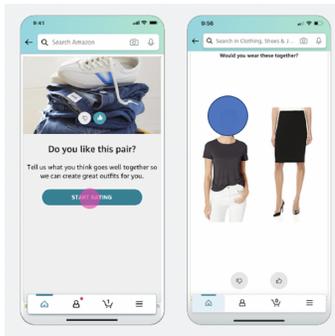


Figure 2: Feedback on pairs

2 RELATED WORK

Existing approaches to complementary item recommendation consider the pairwise similarities between fashion items [5, 8, 10]. Veit et al. [10] proposed a Siamese network to learn item compatibility using co-purchase or co-view information. Instead of computing the similarity in a single space, several recent approaches [5, 8, 9] explored learning sub-space embeddings to capture different notions of similarities. Lin et al. [5] designed a scalable sub-space attention approach for large-scale retrieval. All of these methods consider a generic compatibility model and don't take user preferences into account to personalize compatibility. In contrast, our method utilizes a two stage approach where we first learn customer representation based on millions of customer feedbacks on individual items. We then use these customer representations along with customer preferences on pairs of items to learn a personalized metric of fashion compatibility for each customer.

Various personalized complementary methods have been proposed to incorporate personalized information into the complementary model. McAuley et al. [7] learned a personalized distance function (a weighting matrix) to measure the pair-wise item distances for a user. The matrix is learned using co-purchase data, which is often noisy. Some recent work learned personalized complementary models at the outfit-level [2, 6, 12]. Lu et al. [6] learned binary codes for each item type and user embeddings. The outfit score for a user is computed based on the compatibility score between each pair of fashion items in an outfit and the user's preference for each fashion item (i.e., distance between user embedding and a fashion item). Zheng et al. [12] utilized user information (e.g., selfie posts) from social-media to recommend outfits. Chen et al. [2] collected a large-scale user clicked set of outfits, and used user clicked items as user preference signals to generate outfits. These methods often use random negative images (or outfits) when learning the embeddings. However, these negative images may not be the "true" negatives to the anchor images (outfits). In practice, good negative images are critical to achieve good performance. Therefore, we collect user ratings on pairs of items (Figure 2) to obtain positive and negative image pairs for learning embeddings.

3 PROPOSED APPROACH

Figure 3 illustrates the system overview of our framework. Our framework has three inputs: images of top and bottom items forming a pair, the customer's provided rating on the pair, and the

customer's embedding. The input images are passed through a CNN to extract image feature vectors. A set of masks are applied to the image feature vector to generate multiple subspace embeddings. They correspond to different notions of compatibility that are dependent on the customer. The customer embedding is fed into a sub-network to predict the attention weights, which in turn select the appropriate subspace embeddings for the final embedding computation. The selected subspace embeddings for a particular customer would encode the compatibility relationship that is personalized to that customer. The details of the subspace attention network are presented in a later section.

Our network is trained using contrastive loss on pairs of items. It takes in a training sample $\{(A_i, B_i, c, r_i) | i \in [n]\}$, where r_i is the customer c 's rating on compatibility of items corresponding to the pair (A_i, B_i) of images. The loss of the i^{th} data point is calculated using the distance $d(A_i, B_i)$ and the label r_i .

3.1 Customer embeddings

We learn customer representation using a class of algorithms called matrix completion. We experiment with different algorithms including Partial Dual (a matrix completion algorithm) and neural matrix factorization to learn customer style (Stynet). We use customer preferences on individual items (upvotes/down votes on "Love it or Leave it"), item features (visual features, brand, color etc.) and other customer side information as input to these algorithms. These embeddings are used to compute attention weights in the personalized compatibility model (Figure 3).

3.1.1 Partial Dual. This is a class of matrix completion algorithm where a low rank representation \mathbf{W} of the rating matrix \mathbf{Y} is learnt. \mathbf{W} can be factorized into customer and item embedding matrix. The customer embedding matrix is used to compute subspace attention weights in PSA-Net. $\mathbf{Y} \in \mathbb{R}^{d \times T}$ is the rating matrix on individual items and $L : \mathbb{R}^{d \times T} \times \mathbb{R}^{d \times T} \rightarrow \mathbb{R}$ is a convex loss function, $\|\cdot\|_*^2$ is a nuclear norm regularizer, $\lambda > 0$ is a cost parameter and D is a set of constraints. The \mathbf{W} matrix is a low rank matrix that is learnt by solving the following optimization problem as described in [4]:

$$\begin{aligned} \min_{\mathbf{W} \in \mathbb{R}^{d \times T}} \quad & \lambda L(\mathbf{Y}, \mathbf{W}) + \|\mathbf{W}\|_*^2, \\ \text{subject to } \quad & \mathbf{W} \in D \end{aligned} \quad (1)$$

\mathbf{W} can be factorized into customer and item embedding matrix denoted by \mathbf{E} and \mathbf{M} respectively. Embedding of a customer c is given by \mathbf{E}_c . The customer embeddings learnt using this method are then used to learn attention weights for compatibility subspaces for the personalized compatibility model. The loss function used is mean squared loss between the predicted ratings matrix \mathbf{W} and actual rating matrix \mathbf{Y} and is optimized using Riemannian conjugate gradient and trust-region algorithms.

3.1.2 Stynet. StyNet is a deep neural-network based matrix factorization method motivated from [11] which is able to generalize to both cold-start customers and items by learning the interactions between customer and item features in a latent embedding space. In this architecture, customer embeddings and item embeddings are learnt through customer sub-network and an item sub-network respectively. Customer features and customer's past preferences on the "Love it or leave it" serve as the input layer to the customer

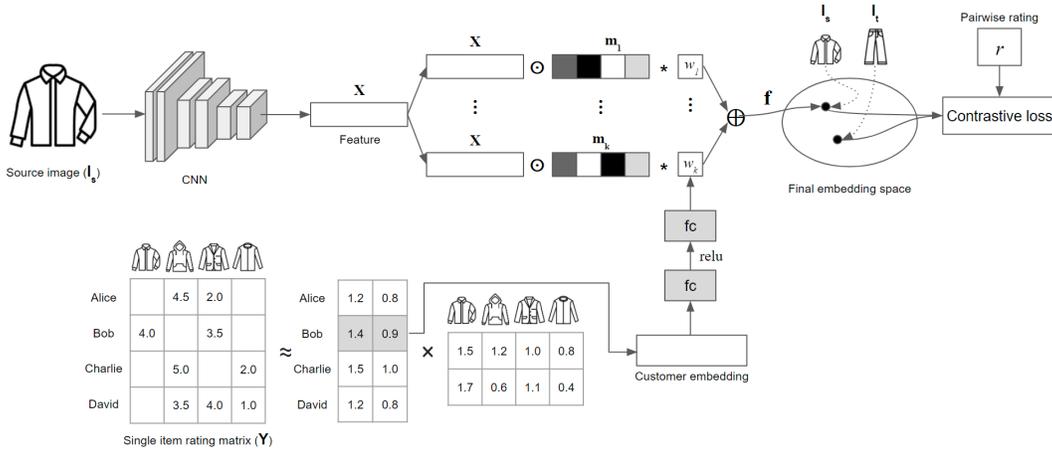


Figure 3: PSA-Net: An overview of our proposed approach

sub-network while item catalog features like itemId, color, category, brand etc. serve as the input layer to the item sub-network. Two embeddings are concatenated and a sigmoid function is applied to the concatenated embedding to transform it into the predicted rating of an item. The overall network is optimized in an end to end fashion using a mean squared loss.

3.2 Subspace attention network

In this section, we describe our personalized subspace attention network (PSA-Net). Instead of computing similarity in a single space, we utilize style subspaces [5, 8, 9] to capture multiple dimensions of compatibility. This is important in learning a personalized notion of compatibility for each customer.

The network learns a non-linear function $\psi(I_s, c)$ that takes a source image I_s and a customer embedding E_c as inputs and generates a feature embedding f . The image is first passed through a CNN to extract its visual feature vector (denoted as X). The customer embedding is used to predict the subspace attention weights (w_1, \dots, w_k) (k is the number of subspaces) using a sub-network. Then, a set of learnable masks (m_1, \dots, m_k) that have the same dimensionality as the image feature vector are applied to the image feature vector via Hadamard product. These masks are learned so as to project the features into a set of subspaces that encode different compatibility substructures. The final embedding constructed by the network is a weighted sum of the subspace embeddings:

$$f = \sum_{i=1}^k w_i * (X \odot m_i), \quad (2)$$

where k is the number of subspaces, X is the image feature vector after the base CNN, m_i are learnable masks, w_i are attention weights, and f is the final embedding.

4 DATA

While we have items from multiple categories in Amazon catalogue, we consider only Tops (Shirts, Sweaters) and Bottoms (Pants, Jeans, Skirts, Shorts) categories for this paper.

4.1 Labelled data from human stylists

We use an existing dataset of human stylists curated outfits within Amazon ($\sim 70K$). Since our task is to learn pairwise compatibility, we produce pairs from the curated outfits dataset. We cannot use rejected outfits to generate negative pairs since some of those pairs can still be compatible even when the entire outfit is not. We use this data to train a pairwise non personalized compatibility model based on the state of the art approach in [5]. This baseline model is compared against a model trained on customer pairwise preference data from *Love it or leave it*.

4.2 Personalized ratings obtained from customers

4.2.1 *Preferences on pairs.* Lack of negatively labelled data and the variance in compatibility labels provided by human stylists encouraged us to collect pairwise compatibility data directly from customers to learn a compatibility metric personalized to them. We get customer feedback on pairs of items through the *Love it or Leave it* experience where a customer upvotes a pair if they think they can wear the two items together and downvote it otherwise (see Figure 2). We get feedback on pairs from each customer (~ 100 on average) and $\sim 500K$ pairs across all customers. This data is used to train the personalized compatibility model.

4.2.2 *Preferences on individual items.* We collect feedback on customer style preferences on individual items through the *Love it or Leave it* experience (see Figure 2). A customer rates approximately 200 items to let us know their style preferences. We use this data to learn customer embeddings that are used in the personalized compatibility model to learn subspace attention weights.

5 PREPROCESSING

Item images sometimes show the model wearing pieces of clothing other than the item of interest, which might add noise to the training. To correct for this, images are cropped before feeding into our models. A YOLO3 object detection model trained on Amazon fashion images that identifies the bounding boxes of all the entities

	Average of metrics over customers(Δ s)			
	ROC AUC	AUPRC	Recall@50%	Precision@10
<i>Baseline 1</i>	-18.27	-15.11	-15	-16.87
<i>Baseline 2</i>	-3.12	-3.65	-1.85	-4.07
<i>Baseline 3</i>	-5.68	-6.25	-3.25	-7.91
<i>PSA-Net</i>	0	0	0	0

Table 1: Comparison of performance of PSA-NET against different baselines. All reported numbers are relative to PSA-NET.

present in the image is used to extract the part of the image which represents the item.

6 EVALUATION

We evaluate our model by computing different offline metrics like customer level ROC, AUPRC, precision@K and recall@K, and compare it with three baseline models, two of them do not take into account the customer embeddings (one trained on fashion expert labelled data and another trained on customer labelled data but which doesn't use customer embeddings to personalize compatibility) and the third one uses customer embeddings but still learns a single compatibility embedding sub-space for all customers. Since most literature on compatibility either does not take personal tastes of the customer into account or collect dataset in a way that is different than ours, there is no external baseline to compare our result against. Works like [3] and [2] include personalization but they work on data-sets consisting of entire outfits and hence are not directly compare-able.

7 EXPERIMENTS

7.1 Comparison with baseline models

7.1.1 Comparison with non-personalized model on customer ratings on pairs of items. For this, we consider the *non-personalized model* wherein the attention weights w_i are directly tunable parameters of our model which do not depend on the customer embeddings. This is equivalent to discarding the customer information associated with each rating item. This model achieves a customer-averaged ROC AUC which is 3.12% less than that of the *personalized model* which uses customer embeddings. The impact on other performance metrics is shown in Table 1. This confirms our hypothesis that leveraging the customer embeddings to predict pairwise compatibility leads to improved offline metrics in compatibility prediction tasks for customers as compared to a model where we do not leverage them on the same dataset.

7.1.2 Comparison with non-personalized model on fashion expert ratings on pairs of items. We also tested the impact of the source of training data on testing performance. We compared the performance of a *non-personalized model* trained using data from *fashion expert labelled data* to PSA-NET, which uses ratings from customers. Using customer provided labels improves the offline recommendation performance metrics over the model trained from labels provided by fashion experts. Specifically, we see an improvement

of 18.26% in customer level ROC AUC. This establishes that collecting compatibility feedback from customers helps in learning a compatibility metric for new customers, whereas data collected from fashion experts is helpful in situations we don't have customer data.

7.1.3 Comparison with personalized model on customer ratings on pairs of items, without subspace attention. To test the impact of the subspace attention network, we compare the performance of PSA-NET against one which does not utilize subspace attention and hence can only learn a single notion of compatibility. Removing subspace attention lowers the offline recommendation performance metrics of the model. Specifically, we see a 5.68% reduction in customer level ROC AUC. This establishes the importance of the subspace attention framework and multiple notions of compatibility. We also compared the performance when using different number of subspaces and found that using 3 subspaces is optimal.

7.2 Ablation studies

The model described in Section 3 has several components which can be tweaked. We discuss the effect of changing individual components, while keeping the rest same.

7.2.1 Effect of source of customer embeddings. Customer embeddings are learnt using customer ratings on individual items using two different methods: Partial Dual and Stynet as described in 3.1. We experiment with customer representations learnt from both these methods in computing attention weights of different subspaces. We found that using embeddings learnt from Partial Dual(a state of the art matrix completion algorithm) leads to better performance in predicting personalized compatibility. Using Stynet embeddings instead of Partial Dual embeddings reduced the model's customer level ROC AUC by 5.86%.

7.2.2 Effect of loss function. The customer-averaged ROC AUC of the model when using triplet loss was found to be 19.06% lower than when using contrastive loss. This large difference can be accounted to the fact that triplet loss requires sampling, which introduces noise to the metric.

7.2.3 Effect of masking operation. The masking operation refers to calculating the final embedding by combining different subspace embeddings taking into account attention weights. In addition to the Hadamard masking operation described in a previous section, we also tried using a fully connected layer to calculate the contribution of each subspace embedding. However, that reduced the model's customer level ROC AUC by 3.68%.

REFERENCES

- [1] 2021. Personal Shopper. <https://www.amazon.com/b?node=19190471011>.
- [2] Wen Chen, Pipei Huang, Jiaming Xu, Xin Guo, Cheng Guo, Fei Sun, Chao Li, Andreas Pfadler, Huan Zhao, and Binqiang Zhao. 2019. POG: Personalized Outfit Generation for Fashion Recommendation at Alibaba iFashion. In *KDD*.
- [3] Tong He and Yang Hu. 2018. FashionNet: Personalized Outfit Recommendation with Deep Neural Network.
- [4] Pratik Jawanpuria and Bamdev Mishra. 2018. A Unified Framework for Structured Low-rank Matrix Learning. In *Proceedings of the 35th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 80)*, Jennifer Dy and Andreas Krause (Eds.). PMLR, Stockholmsmässan, Stockholm Sweden, 2254–2263.
- [5] Yen-Liang Lin, Son Tran, and Larry S. Davis. 2020. Fashion Outfit Complementary Item Retrieval. In *CVPR*.
- [6] Zhi Lu, Yang Hu, Yunchao Jiang, Yan Chen, and Bing Zeng. 2019. Learning Binary Code for Personalized Fashion Recommendation. In *CVPR*.
- [7] Julian McAuley, Christopher Targett, Qinfeng Shi, and Anton van den Hengel. 2015. Image-based Recommendations on Styles and Substitutes. In *SIGIR*.
- [8] Reuben Tan, Mariya I. Vasileva, Kate Saenko, and Bryan A. Plummer. 2019. Learning Similarity Conditions Without Explicit Supervision. In *ICCV*.
- [9] Mariya I. Vasileva, Bryan A. Plummer, Krishna Dusad, Shreya Rajpal, Ranjitha Kumar, and David Forsyth. 2018. Learning Type-Aware Embeddings for Fashion Copatibility. In *ECCV*.
- [10] Andreas Veit, Balazs Kovacs, Sean Bell, Julian McAuley, Kavita Bala, and Serge Belongie. 2015. Learning Visual Clothing Style with Heterogeneous Dyadic Co-occurrences. In *ICCV*.
- [11] Shuai Zhang, Lina Yao, Aixin Sun, and Yi Tay. 2019. Deep Learning Based Recommender System. *Comput. Surveys* 52, 1 (Feb 2019), 1–38. <https://doi.org/10.1145/3285029>
- [12] Haitian Zheng, Kefei Wu, Jong-Hwi Park, Wei Zhu, and Jiebo Luo. 2020. Personalized Fashion Recommendation from Personal Social Media Data: An Item-to-Set Metric Learning Approach. In *ACM MM*.