
ContentNCF: Content Based Neural Collaborative Filtering

Anup Anand Deshmukh

School of Computer Science
University of Waterloo
aa2deshmukh@uwaterloo.ca

Anupa Murali

School of Computer Science
University of Waterloo
a23mural@uwaterloo.ca

Udhav Sethi

School of Computer Science
University of Waterloo
udhav.sethi@uwaterloo.ca

Abstract

1 Some recent works have shown promising results in using deep learning, in partic-
2 ular convolutional neural networks (CNN), for collaborative filtering (CF) [1, 2].
3 These approaches employ deep neural networks to learn high order implicit cor-
4 relations between users and items.

5 Up until now, deep learning based methods in collaborative filtering have focused
6 on learning user-item interactions solely based on IDs as item representations.
7 Modeling user-item interactions based on more extensive item features can pro-
8 vide more context with which we may recommend items to a user. By including
9 content features, we can implicitly learn properties of items that similar users may
10 prefer.

11 In this work, we extend Neural Collaborative Filtering (NCF) [1], to content-
12 based recommendation scenarios and present a CNN based collaborative filter-
13 ing approach tailored to image recommendation. We build upon the Pinterest
14 ICCV dataset used in [1] so as to include image features, and use it to make
15 content-based image recommendations. This content-based approach, Content-
16 NCF, proves successful in predicting user item interactions on our new Pinterest
17 Image 2019 dataset.

18 1 Introduction

19 There is a large amount of visual information available in today’s online world. Users are constantly
20 faced with the dilemma of sifting through a large volume of data to find relevant and novel con-
21 tent. Recommender Systems (RS) aim to solve this by predicting the user’s rating on an item, or
22 recommending items to users based on their preferences.

23 Based on a user’s profile and their interactions with past items, recommender systems find new
24 items which may be of relevance to them. In general, these recommendations are generated based
25 on user preferences, item features and past user-item interactions. Many RSs exploit textual data to
26 recommend items and content to users, even in the presence of image content. Due to the abundance
27 of visual information online today, there has been recent success in making recommendations based
28 on visual features in addition to user features and textual data for items [3].

29 The problem of data sparsity that often plagues recommender systems is due to insufficient infor-
30 mation on a user’s preferences. Content based image recommendations may solve this problem by

31 extracting latent information about the user’s preferences from images. Our work will build upon
 32 existing literature on CNN based collaborative filtering [2, 4] and extend them to make content based
 33 image recommendations.

34 It is known that a convolutional neural network (CNN) is good at learning local features, and the
 35 number of feature maps can reveal multiple aspects of local dimension correlations. ConvNCF [4]
 36 and SECNCF [2] employ CNNs to learn high-order correlations among embedding dimensions.
 37 Both these works have shown promising capability to handle content information and mention the
 38 possibility of future work in this direction.

39 2 Related Work

40 Among different techniques used by RSs, collaborative filtering (CF) [5] is one of the most popular.
 41 The goal of a CF algorithm is to suggest new items or to predict the utility of a certain item for a
 42 particular user based on the user’s previous likes and the opinions of other like-minded users. The
 43 opinions of users can be obtained explicitly from the users, i.e., ratings and reviews, or through im-
 44 plicit feedback, which indirectly reflects users’ preference through behaviours like watching videos
 45 or clicking items. In the next section we will discuss few CF based methods which motivate our
 46 project idea.

47 2.1 Matrix Factorization (MF)

48 Many areas in machine learning use the idea of learning latent (hidden) factors. The MF represents
 49 both items and users by vectors which are obtained from their interaction matrix [6]. The high
 50 correlation between these user and item factors lead to a recommendation. In these set of methods,
 51 users and items are mapped to a joint latent space of dimensionality k .

52 Each item i is associated with a vector $\mathbf{q}_i \in \mathbb{R}^k$, and each user u is associated with a vector
 53 $\mathbf{p}_u \in \mathbb{R}^k$. The elements of a \mathbf{q}_i would capture the extent of latent factors of that particular item.
 54 The user’s interest in items is captured by the elements of \mathbf{p}_u . Finally, the dot product between \mathbf{q}_i^T
 55 and \mathbf{p}_u captures the overall interest of user u in the item i . Following equation captures the above
 56 mentioned steps.

$$\hat{r}_{ui} = \mathbf{q}_i^T \mathbf{p}_u \quad (1)$$

57 To learn the latent factors, \mathbf{p}_u and \mathbf{q}_i , the method involves following minimization of the regularized
 58 squared error on the set of known ratings where L is the training set.

$$\min_{\mathbf{q}^*, \mathbf{p}^*} \sum_{(u,i) \in L} (r_{ui} - \mathbf{q}_i^T \mathbf{p}_u)^2 + \lambda(\|\mathbf{q}_i\|_2^2 + \|\mathbf{p}_u\|^2) \quad (2)$$

59 Recent works have shown limitation of MF caused by the use of simple inner product to capture the
 60 complex user-item interaction in the latent space. Recently, this has motivated researchers to use
 61 Deep Neural Networks (DNNs) to model this interaction function and we are going to exploit the
 62 same for the task of content based image recommendation.

63 2.2 Neural Collaborative Filtering (NCF)

64 Deep Learning (DL) has shown to be promising for problems in recommender systems. This in-
 65 cludes, in a broader sense, an opportunity to reinvent the way we handle the user-item interactions.
 66 Mainly, DL methods are able to capture the non-linear relationship between users and items ef-
 67 fectively. Xiangnan He et al. [1] proposed one such seminal method called Neural Collaborative
 68 Filtering (NCF) which effectively captures the non-linearities in user-item relationships by adopting
 69 a multilayer representation. NCF’s predictive model can be formalized as equation 3,

$$\hat{y}_{ui} = f(\mathbf{P}^T \mathbf{v}_u^U, \mathbf{Q}^T \mathbf{v}_i^I | \mathbf{P}, \mathbf{Q}, \Theta_f) \quad (3)$$

70 where $\mathbf{P} \in \mathbb{R}^{M \times K}$ and $\mathbf{Q} \in \mathbb{R}^{N \times K}$, denote the latent factor matrix for users and items. Also, \mathbf{v}_u^U
 71 and \mathbf{v}_i^I represent feature vectors for user u and item i respectively. f is a neural network and Θ_f

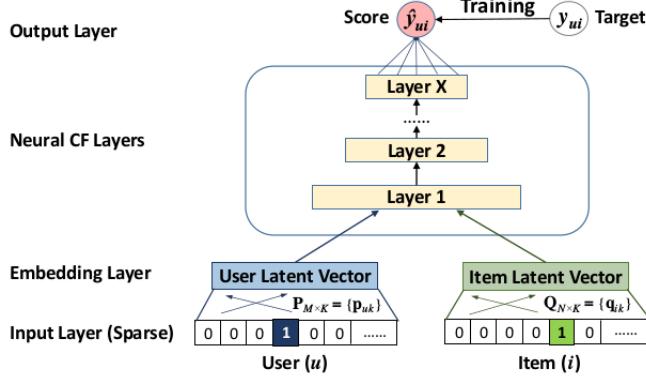


Figure 1: Framework for Neural Collaborative Filtering [1]

72 is the model parameter. Since ϕ_x is the mapping function for the x -th layer as given in Fig 1, f
 73 can be formulated as equation 4. To learn the model parameters authors of NCF use straightforward
 74 regression with squared loss.

$$f(\mathbf{P}^\top \mathbf{v}_u^U, \mathbf{Q}^\top \mathbf{v}_i^I) = \phi_{out}(\phi_X(\dots \phi_2(\phi_1(\mathbf{P}^\top \mathbf{v}_u^U, \mathbf{Q}^\top \mathbf{v}_i^I))) \quad (4)$$

75 It is interesting to note that, stacking more layers in NCF is proven to be beneficial to the recom-
 76 mendation accuracy which also encourages the effectiveness of using deep models for collaborative
 77 recommendation. There are two drawbacks in employing NCF to the task of image recommenda-
 78 tions. Firstly, NCF simply stacks the user and item embeddings and feeds it to the DNN i.e it does
 79 not explicitly consider the correlation among embedding dimensions [4]. On the other hand, in the
 80 context of content based image recommendations, explicitly modelling the pairwise correlation be-
 81 tween embedding dimensions becomes very important. Secondly, it does not make use of content
 82 information while performing CF and relies on user and item IDs instead. This motivates us to use
 83 the appropriate image representation along with the Convolutional Neural Network (CNN) within
 84 the NCF framework. Next section discusses our project idea in detail.

85 3 Our Contribution

86 While CF based recommendations exploit similar users' ratings, content-based recommenders rec-
 87 ommend an item to a user based on the content and description of the item. For multimedia items
 88 like images and videos, extracted visual features have richer semantics which can capture more
 89 potential user preferences. With these richer semantics, it becomes more important to find the cor-
 90 relation between user and item embedding dimensions. CNNs have been shown to be beneficial for
 91 this, as the correlation among multiple dimensions could be represented by a convolutional filter [4].
 92 To summarize, this project aims to build a content based image recommender system which exploits
 93 advantages of a CNN based NCF framework. Our contribution is two-fold as follows.

94 3.1 Pinterest Image Dataset 2019

95 Our initial plan was to use the Pinterest data used in the NCF paper to train and evaluate ContentNCF.
 96 However, it did not work out to be feasible since for our models, it is important to preserve the image
 97 content while training. The Pinterest data used in NCF simply maintains the ID numbers for images
 98 and completely ignores the URL or content information. To deal with this problem, we create our
 99 own dataset (Pinterest Image 2019) as described in the following sections.

100 3.1.1 Preprocessing

101 To create the training and testing data to suit our needs, we use the Pinterest ICCV dataset [7]. The
 102 dataset contains a mapping of boards, which represent users, to their pins, which represent images.

103 Users “pin” images to their own boards, showing their preferences of these images. In our dataset,
 104 each image is represented by a unique ID, as well as a URL. During training, this maintained URL
 105 for every user item interaction is used to extract the features. Each user may have consumed multiple
 106 images, and each image may have been consumed by multiple users. Also, each such user-image
 107 entry is marked as 1 or 0 indicating whether the user has consumed the item or not. This way we
 108 have four columns in our training file where each row takes the form: user ID, image ID, binary
 109 interaction, image URL. Processing all of the Pinterest ICCV data spread across multiple bson files
 110 led us to arrive at a very large dataset, surpassing our computing power. As such, we prune our
 111 dataset to arrive at a filtered dataset containing 500 users and 24498 user-image interactions.

112 **3.1.2 Feature Extraction**

113 With focus on the pure collaborative filtering setting, NCF uses only the identity (ID) of an item
 114 as the input feature, transforming it to a sparse binary vector with one-hot encoding. The sparse
 115 representation, $\mathbf{v}_i^I = [0, 0, 1, 0, 0, 0, \dots]$, is mapped to a dense vector, $\mathbf{Q}^\top \mathbf{v}_i^I$, that represents an
 116 item embedding. This is then fed into a multi-layer neural architecture in NCF. We extract feature
 117 representations from the images themselves as described below.

118 **CNN** With ContentNCF, our goal is to leverage image features to get a richer feature representation
 119 than the item IDs. To obtain the content information for each image, we download the image using
 120 its unique URL and save it in a file named after its ID. Next, we face the non-trivial problem of
 121 representing the image as a vector that describes it effectively and accurately. As we know, deep
 122 learning methods have achieved the most success in computer vision, and many powerful deep
 123 models based on CNN have shown promising results in learning features from 2D image data. We
 124 use VGGNet with pre-trained weights to extract features from the image and represent it as a feature
 125 vector. VGGNet uses only 3×3 convolutional layers stacked on top of each other in increasing
 126 depth. Along with the convolutional layer there are max pooling layers which handle the reducing
 127 volume size. Two fully-connected layers, each with 4,096 nodes are then followed by a softmax
 128 classifier. We pop the VGG layers and remove the last softmax layer. This allows us to take the
 129 output of last fully connected layer as an image representation.

130 **PCA** The above approach leaves us with a 4096 dimensional representation of each image. Next,
 131 we use Principal Component Analysis (PCA) to reduce the dimensionality of image representations
 132 to the number of latent factors configured in our algorithm. Thus, an image is finally represented by
 133 a k -vector, which is then fed into the ContentNCF architecture described in the following section.

134 **3.2 Content Neural Collaborative Filtering - ContentNCF**

135 **3.2.1 Architecture**

136 See Figure 2 for an architecture diagram of ContentNCF. For the purpose of our discussion of the
 137 algorithm, let the latent space used in the Generalized Matrix Factorization (GMF) layer have di-
 138 mension k , and let the first layer of the Multi-Layer Perceptron (MLP) take input with dimension
 139 d .

140 **User Features** We obtain two dense representations of the input from the sparse input vector
 141 $\mathbf{v}_u^U = [0, 1, 0, 0, 0, \dots]$ for user u . We call these two representations $\mathbf{v}_u^{GMF} \in \mathbb{R}^k$, which we will
 142 use as input into the GMF layer, and $\mathbf{P}^\top \mathbf{v}_i^I = \mathbf{v}_u^{MLP} \in \mathbb{R}^{d/2}$, which we will input into the MLP.
 143 \mathbf{v}_u^{GMF} is “MF User Vector” and \mathbf{v}_u^{MLP} is “MLP User vector” from 2.

144 **Dense Layers** We use two densely connected layers, as shown in 2, to learn rich vector represen-
 145 tations of the input image. The output of each dense layer is then passed into the GMF Layer and
 146 the MLP.

147 As noted in previous sections, we are using the image itself as input for item i in ContentNCF, rather
 148 than an image ID as in NCF [1]. We obtain a rich vector representation, $\mathbf{v}_{ik} \in \mathbb{R}^k$, of the image
 149 using a CNN and PCA, as described in the previous section on feature extraction.

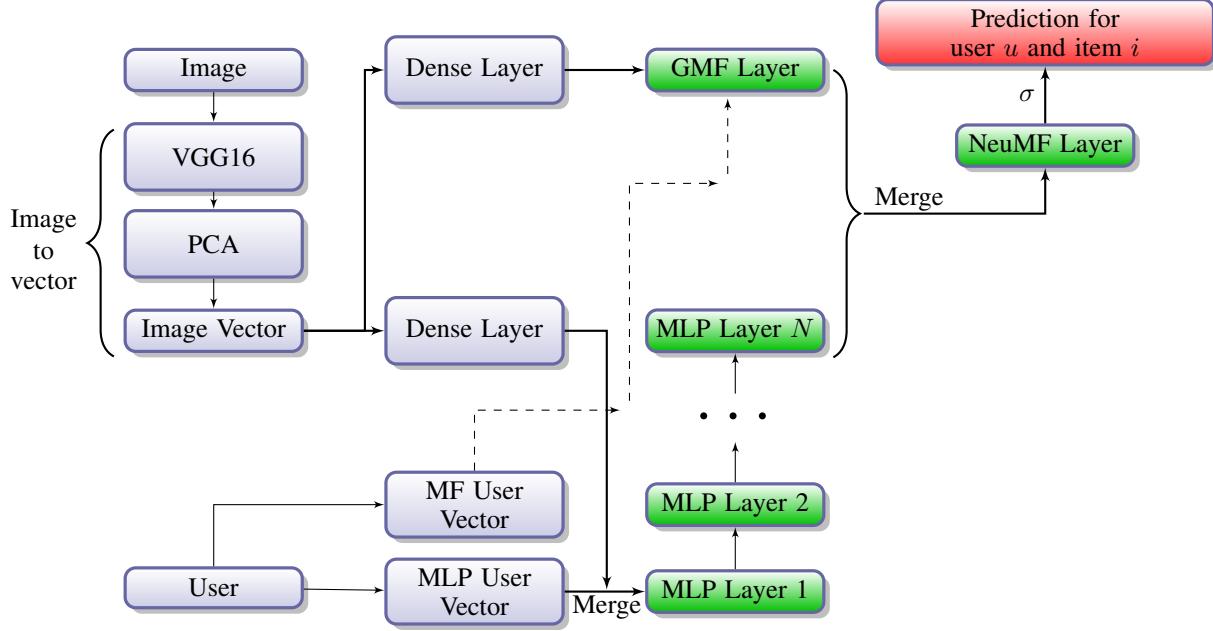


Figure 2: Framework for Content Based Neural Collaborative Filtering (ContentNCF)

150 To learn rich representations from our already reduced image vector, \mathbf{v}_{ik} , we can use densely connected layers. Note that a densely connected layer will feed each of the k features in \mathbf{v}_{ik} into each of its neurons. This could ensure that we capture all of the information and patterns encoded in \mathbf{v}_{ik} .
 151
 152

153 For the dense layer feeding into the GMF layer, for input \mathbf{v}_{ik} we let the output of this dense layer be
 154 $\mathbf{v}_i^{GMF} \in \mathbb{R}^k$.

155 For the dense layer feeding into the MLP, for input \mathbf{v}_{ik} we let the output of this dense layer be
 156 $\mathbf{v}_i^{MLP} \in \mathbb{R}^{d/2}$. The rationale for this is that we wish to concatenate the latent vector representation
 157 of the image with user vector $\mathbf{v}_u^{MLP} \in \mathbb{R}^{d/2}$ (as per 2) to create a valid input into the first layer of
 158 the MLP with dimensions d .

159 **GMF Layer** The GMF layer computes user item interactions by computing the element wise
 160 product of its inputs, as in NCF [1]. It takes in both \mathbf{v}_u^{GMF} (“MF User Vector” in 2) and \mathbf{v}_i^{GMF} (the
 161 output of the dense layer which feeds into the GMF layer, as shown in 2). Note that $\mathbf{v}_u^{GMF}, \mathbf{v}_i^{GMF} \in$
 162 \mathbb{R}^k . The output of this layer will be the element wise product of vectors \mathbf{v}_u^{GMF} and \mathbf{v}_i^{GMF} . Let us
 163 call the output of the GMF layer $\mathbf{v}_u^{GMF} \odot \mathbf{v}_i^{MLP} = \mathbf{v}_{GMF} \in \mathbb{R}^k$.

164 **MLP Layers** The MLP in our model has multiple layers can be passed in as an argument to our
 165 implementation. By using multiple layers, it learns higher order patterns and correlations between
 166 the user and image. Recall that the dimensions of the input of the first layer of the MLP (“MLP Layer
 167 1” in 2) is $d \times 1$. We create the input into the first layer of the MLP by concatenating $\mathbf{v}_u^{MLP} \in \mathbb{R}^{d/2}$
 168 (i.e. “MLP User Vector” in 2) and $\mathbf{v}_i^{MLP} \in \mathbb{R}^{d/2}$ (the output of the dense layer which feeds into the
 169 MLP, as shown in 2).

170 We concatenate \mathbf{v}_i^{MLP} and \mathbf{v}_u^{MLP} to create a d -dimensional vector and feed it into “MLP Layer 1”.
 171 Let us call the final output of the MLP $\mathbf{v}_{MLP} \in \mathbb{R}^{d_N}$, where d_N is the dimension of the output of
 172 the N^{th} layer.

173 **Prediction Layer** Finally we concatenate \mathbf{v}_{GMF} and \mathbf{v}_{MLP} and pass the resulting vector into the
 174 prediction layer (i.e. “NeuMF Layer” in 2). The prediction layer is densely connected and uses a
 175 sigmoid activation function in order to restrict the output to be in $(0, 1)$. This final layer will output
 176 a prediction for the interaction between user u and item i , $\hat{y}_{ui} \in (0, 1)$. \hat{y}_{ui} denotes how likely i is
 177 salient to u , and therefore should be recommended to u .

178 **3.2.2 Learning with Log Loss**

179 Traditional squared loss methods to learn parameters may not work well with the implicit data. This
 180 is because $y_{ui} \in \{0, 1\}$ (i.e. is binary), indicating whether or not there exists a user-item interaction,
 181 for implicit data; if we were to use squared loss, however, we are to assume that y_{ui} comes from a
 182 Gaussian distribution [8]. We use the approach used in the NCF paper by formulating ContentNCF
 183 as a probabilistic model and optimize it with the log loss:

$$\mathcal{L} = - \sum_{(u,i) \in y} y_{ui} \log \hat{y}_{ui} + (1 - y_{ui}) \log (1 - \hat{y}_{ui}) \quad (5)$$

184 We use the above objective function which we minimize using the Adam optimization algorithm.
 185 During training, we randomly sample negative instances from unobserved user-item interactions.
 186 The number of negative samples per positive sample can be passed in as an argument to our im-
 187 plementation. Our results from varying the number of negative samples per positive sample are
 188 described in the Results section.

189 **4 Experiments**

190 **4.1 Evaluation Methodology**

191 The model is trained on the user-image interaction matrix. For every positive interaction of a user
 192 we add certain number (num_negative) of negative examples. To evaluate the performance of our
 193 ContentNCF we use the leave-one-out strategy. For each user, on average we have 5 interactions
 194 in the test set and rest of the interactions are utilized for training. For each such test image, we
 195 randomly sample 100 images that the user has not interacted with. The goal of ContentNCF is now
 196 to rank the positive test image among these 101 images. This common strategy is also followed in
 197 [9]. The algorithm then generates the list of top-K recommendations which is evaluated by hit ratio
 198 (HR) and Normalized Discounted Cumulative Gain (NDCG) [10]. HR checks if the test image is
 199 present in the top-K list and NDCG measures the position of the hit by assigning higher scores to
 200 hits at top ranks. We calculated both metrics for each test user and reported the average score. We
 201 have made our implementation open source and it can be found here.¹

202 **4.2 Results**

203 This section demonstrates the recommendation performance of ContentNCF with different param-
 204 eter settings. In the first set of experiments we show the working of ContentNCF with top-K = 5,
 205 $k = 20$ and num_negative = 4 where k is the number of latent factors used while representing a user
 206 and an image.

Table 1: ContentNCF (Top-K = 5, $k = 20$, num_negative = 4)

Iterations	Evaluation Metrics		
	HR	NDCG	Loss
init	0.049	0.0293	
1	0.4759	0.2968	0.4659
2	0.4824	0.3092	0.2831
3	0.5127	0.3295	0.2758
4	0.5135	0.323	0.2688
5	0.5118	0.3222	0.2649

207 Table 1 and Table 2 show the performance of ContentNCF after every iteration. The top-K = 5
 208 recommendations task is more difficult than top-K = 10 and hence we get lower HR of 0.51 opposed
 209 to HR of 0.74 for top-K = 10. In further experiments we set top-K to 10.

210 In the next set of experiments we compare the performance by changing the number of latent factors
 211 (k). The latent factor influences the final image representation we get from PCA and higher values of

¹https://github.com/udhavsethi/neural_collaborative_filtering

Table 2: ContentNCF (Top-K = 10, $k = 20$, num_negative = 4)

Iterations	Evaluation Metrics		
	HR	NDCG	Loss
init	0.1216	0.0534	
1	0.7282	0.3867	0.453
2	0.7273	0.3885	0.2814
3	0.7257	0.4041	0.272
4	0.7412	0.4023	0.2658
5	0.742	0.4015	0.2629

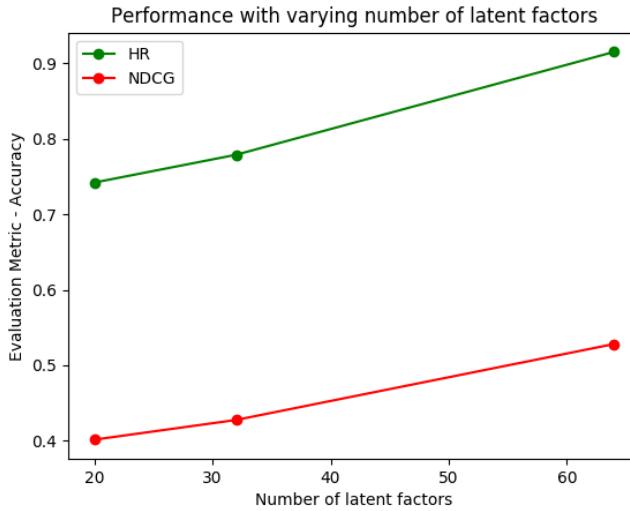


Figure 3: ContentNCF (Top-K = 10, varying k , num_negative = 4)

212 k would give us succinct representation of an image. Both, HR and NDCG increases as we increase

213 k . HR jumps from 0.74 ($k = 20$) to 0.91 ($k = 64$) and NDCG reaches 0.53 for $k = 64$ (Figure 3).

214 Figure 4 shows the performance of ContentNCF where number of negative samples per positive sam-
215 ple (num_negative) ranges from 4 to 10. It can be seen that, best HR is achieved when num_negative
216 is equal to 10. NDCG consistently increases till num_negative = 8 and then decreases further for
217 num_negative = 10.

218 In our last set of experiments (Figure 5) we compare our ContentNCF with NCF [1] on Pinterest Im-
219 age 2019 dataset. ContentNCF with the best parameter setting achieves HR and NDCG of 0.940 and
220 0.582 resp. Although ContentNCF fails to outperform NCF model in terms of accuracy measures, it
221 has potential to give more diverse recommendations given it is a content based RS.

222 5 Discussion and Future Work

223 Instead of a simple embedding concatenation, ConvNCF [4] and SECNCF [2] employ CNNs to
224 explicitly model the pairwise correlations between embedding dimensions. While ConvNCF applies
225 outer product on the user and item embeddings, SECNCF combines different embeddings together
226 in the same direction to form stacked embeddings. We plan to employ one of these techniques or
227 their variants to model embeddings in our architecture as part of the future work.

228 There has been a lot of consistent effort in increasing the accuracy of recommender systems. A
229 good recommender system should aim to improve user satisfaction apart from increasing traditional
230 accuracy measures. Our goal in this possible extension would be to build on ContentNCF to produce
231 diverse yet relevant recommendations by exploiting content information [11], [12]. There exist

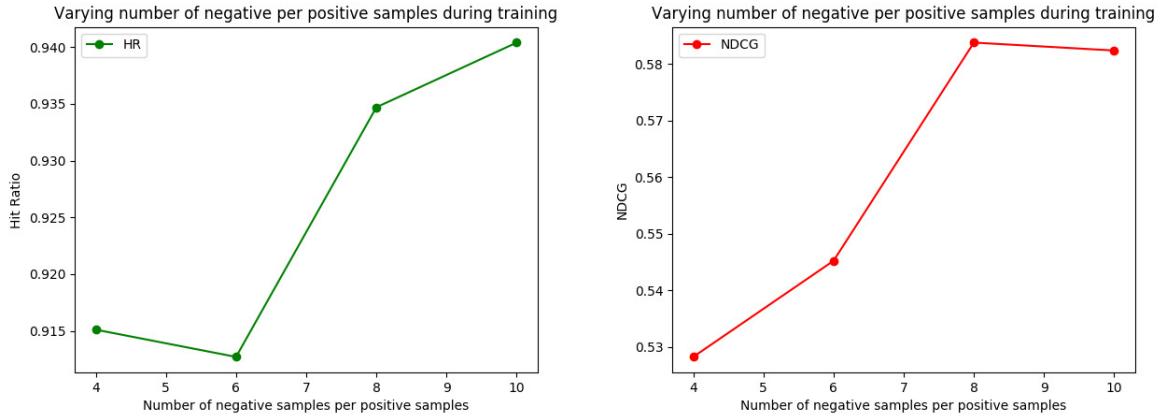


Figure 4: ContentNCF (Top-K = 10, $k = 64$, varying num_negative)

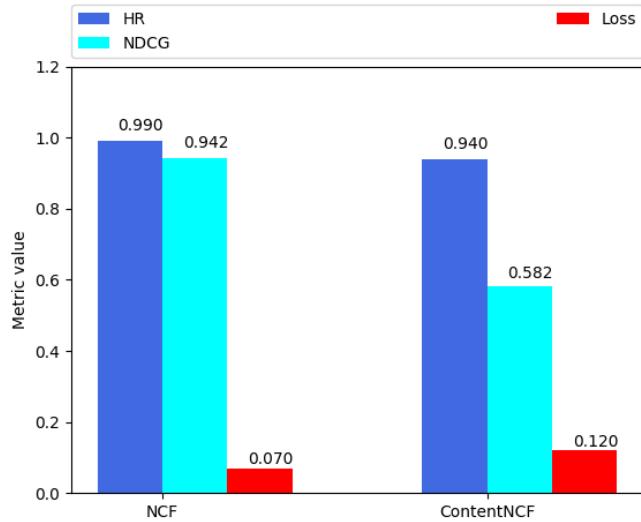


Figure 5: Comparison between NCF and ContentNCF

232 several ways in which we can formalize diversity. For example, in order to measure the diversity in
 233 a list of recommended items, Q_u , we can look at the pairwise distances between items in the list.
 234 We can then calculate the diversity value by averaging the normalized Hamming distance, $H(i, j)$,
 235 for all item pairs i, j in Q_u . Equation 6 can be used to calculate the diversity.

$$Diversity = \frac{2}{|Q_u|(|Q_u| - 1)} \sum_{i \in Q_u} \sum_{j \in Q_u, j \neq i} H(i, j) \quad (6)$$

236 6 Conclusion

237 We present ContentNCF, a content-based RS for learning user-item interactions, as opposed to the
 238 recent deep learning based methods which use only IDs as item representations. We propose the use
 239 of CNNs to leverage visual features from images to learn rich feature representations. We observe
 240 that ContentNCF achieves relevance prediction accuracy that is comparable with that of NCF, but
 241 also has the ability to model content information. With this ability it has potential to give more
 242 diverse yet relevant recommendations which are shown to increase the user satisfaction.

243 **References**

244 [1] X. He, L. Liao, H. Zhang, L. Nie, X. Hu, and T.-S. Chua, “Neural collaborative filtering,” in *Proceedings of the 26th international conference on world wide web*, pp. 173–182, International World Wide Web Conferences Steering Committee, 2017.

245 [2] L. Han, H. Wu, N. Hu, and B. Qu, “Convolutional neural collaborative filtering with stacked embeddings,” in *Asian Conference on Machine Learning*, pp. 726–741, 2019.

246 [3] P. C. F. B. M. Y. Deldjoo, M. Elahi and A. L. E. Caielli, “How to combine visual features with tags to improve movie recommendation accuracy?,” in *Proceedings of International Conference on Electronic Commerce and Web Technologies*, pp. 34–45, 2016.

247 [4] X. He, X. Du, X. Wang, F. Tian, J. Tang, and T.-S. Chua, “Outer product-based neural collaborative filtering,” *arXiv preprint arXiv:1808.03912*, 2018.

248 [5] J. L. Herlocker, J. A. Konstan, A. Borchers, and J. Riedl, “An algorithmic framework for performing collaborative filtering,” in *22nd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 1999*, pp. 230–237, Association for Computing Machinery, Inc, 1999.

249 [6] D. D. Lee and H. S. Seung, “Algorithms for non-negative matrix factorization,” in *Advances in neural information processing systems*, pp. 556–562, 2001.

250 [7] X. Geng, H. Zhang, J. Bian, and T. Chua, “Learning image and user features for recommendation in social networks,” in *2015 IEEE International Conference on Computer Vision (ICCV)*, pp. 4274–4282, Dec 2015.

251 [8] L. Han, H. Wu, N. Hu, and B. Qu, “Convolutional neural collaborative filtering with stacked embeddings,” in *Proceedings of The Eleventh Asian Conference on Machine Learning* (W. S. Lee and T. Suzuki, eds.), vol. 101 of *Proceedings of Machine Learning Research*, (Nagoya, Japan), pp. 726–741, PMLR, 17–19 Nov 2019.

252 [9] A. M. Elkahky, Y. Song, and X. He, “A multi-view deep learning approach for cross domain user modeling in recommendation systems,” in *Proceedings of the 24th International Conference on World Wide Web*, pp. 278–288, International World Wide Web Conferences Steering Committee, 2015.

253 [10] X. He, T. Chen, M.-Y. Kan, and X. Chen, “Trirank: Review-aware explainable recommendation by modeling aspects,” in *Proceedings of the 24th ACM International Conference on Information and Knowledge Management*, pp. 1661–1670, ACM, 2015.

254 [11] A. A. Deshmukh, P. Nair, and S. Rao, “A scalable clustering algorithm for serendipity in recommender systems,” in *2018 IEEE International Conference on Data Mining Workshops (ICDMW)*, pp. 1279–1288, IEEE, 2018.

255 [12] D. Kotkov, J. A. Konstan, Q. Zhao, and J. Veijalainen, “Investigating serendipity in recommender systems based on real user feedback,” in *Proceedings of the 33rd Annual ACM Symposium on Applied Computing*, pp. 1341–1350, ACM, 2018.