



Image-Based Fashion Product Recommendation with Deep Learning

Hessel Tuinhof¹(✉), Clemens Pirker², and Markus Haltmeier³

¹ InvoiceFinance, 's-Hertogenbosch, Netherlands
hessel@invoicefinance.nl

² Department of Strategic Management, Marketing and Tourism,
University of Innsbruck, Innsbruck, Austria
clemens.pirker@uibk.ac.at

³ Department of Mathematics, University of Innsbruck, Innsbruck, Austria
markus.haltmeier@uibk.ac.at

Abstract. We develop a two-stage deep learning framework that recommends fashion images based on other input images of similar style. For that purpose, a neural network classifier is used as a data-driven, visually-aware feature extractor. The latter then serves as input for similarity-based recommendations using a ranking algorithm. Our approach is tested on the publicly available Fashion dataset. Initialization strategies using transfer learning from larger product databases are presented. Combined with more traditional content-based recommendation systems, our framework can help to increase robustness and performance, for example, by better matching a particular customer style.

Keywords: Recommendation · Deep learning · Convolutional neural networks · Similarity recommendation

1 Introduction

Identifying products a specific customer likes most can significantly increase the earnings of a company [15]. Clearly, recommending suitable products in E-commerce increases the probability of a customer's purchase. Additionally, offering too many products can reduce the probability that a potential customer performs a purchase at all. Finally, knowing and subsequently targeting customer preferences increases the medium- and long-term commitment of the customer to the company, which is a key factor to profitability [3, 17]. Prior studies demonstrate that recommendation engines help consumers to make better decisions, reduce search efforts and find the most suitable prices [5].

One possibility to infer knowledge about customer preferences is via specific questioning in customer surveys. However, this is not always possible and customer responses may not be correct or sufficient for accurately describing preferences. In this work, we follow a different, data-driven approach, where customer preferences are automatically extracted from available information on the

customer. More specifically, we focus on fashion products and develop a method that only requires a single input image to return a ranked list of similar-style recommendations.

1.1 Proposed Recommendation System

The proposed recommendation system operates in a two-stage mode. In the first step, we train a convolutional neural network (CNN) to solve specific image classification tasks. The trained CNN is then used as a problem-specific feature extractor, where the features serve as inputs for the ranking system. While in this paper we work with fashion products, similar recommendation systems can be employed for other product categories as well.

Image data provides a wealth of information on a visually-aware feature level, e.g. edges and color blobs. Plenty of image processing techniques exist to extract such low-level features [14]. Deep learning provides a technique to extract hidden higher-level features by composing several convolutional layers. Therefore CNNs are a natural choice to provide fashion product recommendations based solely on image data. Compared to classical content-based recommendation, which is mainly based upon descriptive metadata like manually annotated product tags or user reviews, our approach relies on visual information.

1.2 Relation to Previous Work

There are at least two main approaches for product recommendations: collaborative filtering and content-based filtering. Whereas the former relies on historical user-item interactions, the latter tries to relate user profiles and item descriptors. A recent deep learning approach is the neural collaborative filtering framework proposed in [7], which generalizes the matrix factorization technique used extensively in collaborative filtering methods. Others like [6] employ a hybrid approach, where a matrix factorization based predictor is combined with a deep learning model that extracts visual features as well as latent non-visual user features. A recent thorough overview on deep learning-based recommender systems can be found in [20].

The success of CNNs for computer vision tasks like object classification, detection and segmentation [4] gives reason to decouple classical product recommendation solutions from its extensive user-item interaction data usage requirement. Therefore our method uses product image data, which, for example in E-commerce, is readily available. This also allows to mitigate the cold start problem of collaborative filtering and classical content-based recommender systems. Closely related to our approach are the works [1, 16]. Due to the high degree of subjectivity related to fashion articles, general recommender systems usually perform poorly in fashion recommendation tasks. We show that recommendation systems purely relying on visual features are reasonable as they are able to provide highly visually appealing recommendations of similar style. This can also be helpful in the case of new customers, where no historical user data is yet available. It can also be integrated in existing content-based systems, for

example, to account for a particular or desired style of a customer, or to address the cold-start problem.

1.3 Outline

The remainder of the paper is structured as follows. Section 2 presents the proposed product recommendation method. In particular, we give details on the used network architectures, the used ranking algorithm and describe the Fashion dataset. In Sect. 3 we present some numerical results. The paper concludes with a short discussion in Sect. 4.

2 Methods

2.1 Fashion Dataset

Throughout this paper, we work with a subset of the publicly available Fashion¹ dataset [12]. In order to obtain high-quality ground-truth labels for category type and texture attributes, we design a labeling questionnaire on the crowdsourcing platform CrowdFlower². Every image is labeled by a total maximum of five human operators. To be a valid label at least three human operators have to agree. Each labeling task consists of five images to be labelled, one of which is a simple test image. If a human operator fails a test more than twice, she is no longer allowed to continue. Separate datasets for category and texture classification have been created.

The used class labels for category types are blouse, dress, pants, pullover, shirt, shorts, skirt, top, T-shirt. For the texture attributes we use the labels graphic, plaid, plain, spotted, striped. Figure 1 shows the frequency distributions for the two datasets. The category type dataset contains 11 851 and the texture attributes dataset 7342 images. Further characteristics can be found in Table 1.

Table 1. Summary of the datasets created from the Fashion dataset. The third column indicates the total amount of class labels for the respective dataset.

Dataset	Classification		Samples		
	Type	No.	Total	Train	Val
Fashion category	Multinomial	9	11,851	9,480	2,371
Fashion texture	Multinomial	5	7,342	5,873	1,469

¹ http://imagedlab.ing.unimore.it/fashion_dataset.asp.

² www.crowdflower.com.

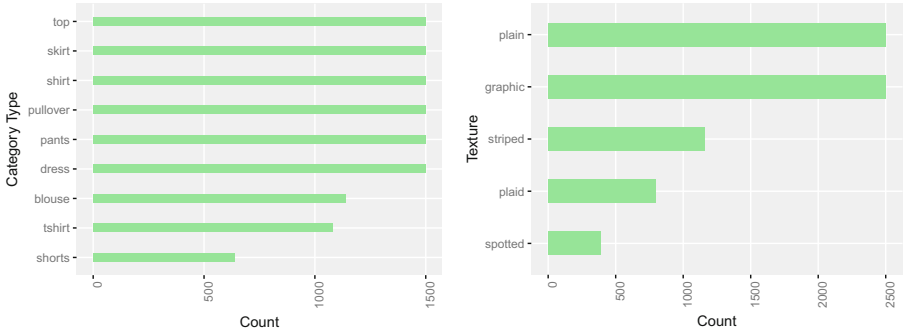


Fig. 1. Frequency distributions for the category type and texture attributes datasets created from the Fashion dataset.

2.2 Proposed Framework

Our method composes of a trained CNN classifier used as image feature extractor and a modification of the k -nearest neighbors (k -NN) algorithm used for ranking in feature space.

■ **Classification via CNNs:** In the first step, we train separate CNNs to predict the category and texture type. Each of the CNNs can be written as

$$\mathcal{N}_i(\mathbb{W}_i, \cdot) \triangleq \mathcal{S}_i(\mathbb{V}_i, \mathcal{F}_i(\mathbb{U}_i, \cdot)) \quad \text{for } i = 1, 2. \quad (1)$$

Here $\mathbb{W}_i = (\mathbb{U}_i, \mathbb{V}_i)$ are weight vectors, $\mathcal{S}_i(\mathbb{V}_i, \cdot)$ are fully connected softmax output layers that actually perform classification and $\mathcal{F}_i(\mathbb{U}_i, \cdot)$ are the CNNs without the last layer. The latter are used as feature extractor.

■ **Ranking in feature space:** After training and evaluating the performance of these classifiers, we remove the softmax output layer \mathcal{S}_i of each model. The remaining CNNs are then concatenated and $\mathcal{F} = [\mathcal{F}_1, \mathcal{F}_2]$ is used to extract the feature vector $\mathcal{F}(\mathbf{X})$ of any input image $\mathbf{X} \in \mathbb{R}^{N \times N}$. We then use the k -NN algorithm to search for the closest items to $\mathcal{F}(\mathbf{X})$ in feature space.

Details on the employed CNNs and the k -NN algorithm for ranking are presented below.

2.3 Network Architectures

A wealth of CNN architectures are available today. In this section we briefly discuss the two architectures that we use in our work: AlexNet and batch-normalized Inception (BN-Inception). The AlexNet and BN-Inception are both standard architectures and well established. AlexNet has been chosen as a benchmark to compare against deeper, more complex networks like the BN-Inception. AlexNet consists of 8 layers and BN-Inception of 34. Both use an image of size 224×224 as input.

Two important contributions of the AlexNet [10] are popularizing usage of the non-saturating rectified linear unit activation function, $\text{ReLU}(x) \triangleq \max(0, x)$, and introducing a normalization layer after the ReLU activation. Empirical results show that the normalization layer improves the generalization ability of the network. The BN-Inception [8] is an extension of the GoogLeNet architecture [18], which allows deeper and wider CNNs by mapping the output of a layer to several layers at once. The output of these parallel layers is then again concatenated. The proposed batch normalization extension addresses the internal covariate shift problem. The latter describes the problem that the latent input distribution of every hidden layer constantly changes, because every training iteration updates the weight vector \mathbb{W}_i . Batch normalization also has a regularization effect.

2.4 Network Training

In order to adjust $\mathcal{N}_i(\mathbb{W}_i, \cdot)$ to the particular classification task, the weight vector \mathbb{W}_i is selected depending on a set of training data $\mathcal{T}_i \triangleq \{(\mathbf{X}_n, \mathbf{Y}_n)\}_{n=1}^{N_i}$. For this purpose, the weights are adjusted in such a way, that the overall error of $\mathcal{N}_i(\mathbb{W}_i, \cdot)$ made on the training set is small. This is achieved by minimizing the error function

$$E(\mathbb{W}_i) \triangleq \sum_{n=1}^{N_i} d(\mathcal{N}_i(\mathbb{W}_i, \mathbf{X}_n), \mathbf{Y}_n) + \lambda \|\mathbb{W}_i\|^2, \quad (2)$$

where d is a distance measure that quantifies the error made by the network function $\mathcal{N}_i(\mathbb{W}_i, \cdot)$ for classifying the n -th training sample.

To stabilize the weight computation in (2), we add a L^2 -regularization term $\lambda \|\mathbb{W}_i\|^2$ with regularization parameter $\lambda \geq 0$. As is common for classification with neural networks, we use the cross entropy for the loss function d . The actual minimization of (2) is performed by stochastic gradient descent.

2.5 Ranking by k -NN

The k -NN algorithm can be used as simple ranking algorithm. For that purpose, consider the feature space \mathbb{R}^p and denote with $d_2(\mathbf{f}, \mathbf{g}) = \|\mathbf{f} - \mathbf{g}\|_2$ the Euclidean distance of two feature vectors. Let $\{\mathbf{f}_1, \dots, \mathbf{f}_m\}$ be a training set of feature vectors. A k -NN algorithm then solves some regression or classification task at $\mathbf{f} \in \mathbb{R}^p$ using the k closest training features. This can be implemented by first computing an enumeration $\pi(\mathbf{f}): \{1, 2, \dots, m\} \rightarrow \{1, 2, \dots, m\}$ satisfying $d_2(\mathbf{f}, \mathbf{f}_{\pi(\mathbf{f})(i)}}) \leq d_2(\mathbf{f}, \mathbf{f}_{\pi(\mathbf{f})(i+1)})$. We use the permutation $\pi(\mathbf{f})$ as ranking output for the input feature \mathbf{f} . To reduce memory requirements of the k -NN ranking, we use an implementation that employs a balltree search [13].

3 Results

In this section we present results for the image classification and similarity recommendation with the proposed framework.

Table 2. Summary of the datasets created from the DeepFashion Attribute Prediction dataset used for pretraining. The third column indicates the total amount of class labels for the respective dataset.

Dataset	Classification		Samples		
	Type	No.	Total	Train	Val
DeepFashion category	Multinomial	46	289,222	231,377	57,845
DeepFashion texture	Multinomial	156	111,405	89,124	22,281

3.1 Pretraining

To overcome difficulties arising from the relative small size of the Fashion dataset, we use the concept of transfer learning [4, 19]. For that purpose, we pretrain the classification models on a larger dataset (namely, the DeepFashion Attribute Prediction³ dataset, [11]) containing 289 222 garment images. A full summary of the dataset can be found in Table 2.

For pretraining we use AlexNet and BN-Inception architectures. For the AlexNet we minimize (2) with stochastic gradient descent using batch size of 64, regularization parameter $\lambda = 0.0005$, learning rate 0.01 and momentum 0.9. For training the BN-Inception we use the ADAM [9] algorithm with batch size of 32, $\lambda = 0$, and learning rate 0.001. Following [4], we use early stopping as an efficient regularization technique to prevent overfitting. We therefore stop training AlexNet/BN-Inception after 9/8 and 17/13 epochs for the category and texture classification, respectively.

Additional to the cross entropy loss, we use the evaluation metrics accuracy,

$$\text{accuracy}(\mathbf{y}, \hat{\mathbf{y}}) \triangleq \frac{1}{N} \sum_{n=1}^N \mathbf{1}_{y_n}(\hat{y}_n), \quad (3)$$

and top- K accuracy, which is defined as in Eq. (3) with a slightly modified indicator function such that top- K predicted classes are incorporated. Table 3 shows accuracy, top- K -accuracy and loss evaluated on the test set for both AlexNet and BN-Inception. The BN-Inception achieves higher accuracy and better generalization ability. Therefore, we only use the BN-Inception architecture for classification on the Fashion dataset.

3.2 Classification

For the final classification models we train the BN-Inception by minimizing (2) on the Fashion dataset with ADAM, where the weights are initialized using the ones from the pretraining stage. Due to the small size of the Fashion dataset, we add L^2 -regularization with $\lambda = 0.0001$ to the loss function and also reduce the batch size to 16.

³ <http://mmlab.ie.cuhk.edu.hk/projects/DeepFashion/AttributePrediction.html>.

Table 3. Pretraining results: the left table depicts results for the category classification and the right table for the texture classification.

Category AlexNet BN-Inception			Texture AlexNet BN-Inception		
Accuracy	0.57	0.63	Accuracy	0.28	0.32
Top-5	0.79	0.84	Top-3	0.62	0.66
Loss	1.48	1.27	Loss	3.00	2.82

Several image augmentation techniques are applied in order to effectively increase dataset size. These include random rotations with a maximum rotation angle of ± 3 for the category type model and ± 8 for the texture attributes model, random changes of HSL color channels within a range of $[-6, 6]$, a shear transformation with random shear factor within $[-0.25, 0.25]$, random aspect ratio changes within a range of $[0.875, 1.125]$ and random vertical flips. The random augmentations are applied to the training set every epoch anew. This allows to train longer without overfitting too fast. Following early stopping regularization, we stop training the category type and texture attributes classification models after 15 and 4 epochs respectively. Table 4 summarizes the final training results. The top- K accuracy metric is however excluded due the smaller number of class labels in the Fashion datasets.

Table 4. Final BN-Inception classification results on the Fashion datasets for category and texture.

	Category	Texture
Accuracy	0.87	0.80
Loss	0.42	0.61

3.3 Similarity Recommendation

The CNN classifiers are used as feature extractors and return feature vectors $\mathcal{F}_i(\mathbf{X})$ of size $d = 1024$ for any input image. The feature extractors are applied to a set of $n = 19422$ test images. These corresponding feature vectors are concatenated and stacked to obtain a $n \times 2d$ feature matrix. The k -NN ranking algorithm is applied to the feature matrix. For the recommendation task, it is now sufficient to extract the features from an input image, submit them to the k -NN ranking algorithm and return the top- k matching style recommendations. In Fig. 2 we present several query images and corresponding top-5 recommendations. Subjectively, the top-5 recommendations indeed look quite similar to the query images. In the top row a query image from the dataset itself is used. This corresponding top-5 recommendations demonstrate that if the image appears in the dataset it is actually most similar to itself. Similar results have been




Fig. 2. k -NN recommendation ranking. First column displays the query images and columns 2–6 display the predicted five nearest neighbors, where column 2 is the most similar.

obtained in other performed tests. Other than that, an implicit objective metric for recommendation quality can be found by means of the classification accuracies reported in Table 4. The definition of a precise objective evaluation criterion, however, remains difficult due to the inherent subjectivity of recommendation quality. This also makes comparison with other methods quite challenging. The computationally most time-consuming part in the application of the proposed recommendation system is the evaluation of the CNN classifiers.

In our implementation, we have implemented the CNNs in MXNet [2] using its Python API. Running on a desktop PC with an Intel i7-6850K CPU and a NVIDIA 1080Ti GPU, the whole image processing pipeline applied to a given input image only takes fractions of a second. Note that the potentially time-consuming network training is done before a new input image is provided to the recommendation system, which therefore allows fast online product recommendation.

4 Conclusion

We presented a visually-aware, data-driven and rather simple but still effective recommendation system for fashion product images. The proposed two-stage approach uses a CNN classifier to extract features that are used as input for similarity recommendations. It can be used, for example, in E-commerce by allowing customers to upload a specific fashion image and then offering similar items based on texture and category type features of the customer's uploaded image. Additional feature extractors, e.g. trained on gender or color classification tasks, can be easily added. Furthermore, generalization to other domains makes sense, e.g. music recommendation based on raw music data, but needs further investigation. Several interesting extensions of our approach are possible. First, it would be promising to integrate the two separate training stages into a single one and provide end-to-end deep learning-based fashion product recommendations. In particular, consideration should be given to Siamese networks. Additionally, hybrid approaches combining image-based and content-based systems will be implemented. Finally, it is important to evaluate the customer impact of our image-based approach and its extensions against other recommender systems through customer surveys.

References

1. Chen, L., Yang, F., Yang, H.: Image-based product recommendation system with convolutional neural networks (2017)
2. Chen, T., et al.: MXNET: a flexible and efficient machine learning library for heterogeneous distributed systems. arXiv preprint [arXiv:1512.01274](https://arxiv.org/abs/1512.01274) (2015)
3. Dick, A.S., Basu, K.: Customer loyalty: toward an integrated conceptual framework. *J. Acad. Market. Sci.* **22**(2), 99–113 (1994)
4. Goodfellow, I., Bengio, Y., Courville, A.: *Deep Learning*. MIT, Cambridge (2016)
5. Häubl, G., Murray, K.B.: Double agents: assessing the role of electronic product recommendation systems. *MIT Sloan Manag. Rev.* **47**(3), 8–12 (2006)
6. He, R., McAuley, J.: VBPR: visual Bayesian personalized ranking from implicit feedback. In: AAAI, pp. 144–150 (2016)
7. He, X., Liao, L., Zhang, H., et al.: Neural collaborative filtering. In: WWW 2017, pp. 173–182 (2017)
8. Ioffe, S., Szegedy, C.: Batch normalization: accelerating deep network training by reducing internal covariate shift. In: ICML, pp. 448–456 (2015)
9. Kingma, D.P., Ba, J.: Adam: a method for stochastic optimization. [arXiv:1412.6980](https://arxiv.org/abs/1412.6980) (2014)
10. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: NIPS, pp. 1097–1105 (2012)
11. Liu, Z., Luo, P., Qiu, S., et al.: DeepFashion: powering robust clothes recognition and retrieval with rich annotations. In: CVPR, pp. 1096–1104 (2016)
12. Manfredi, M., Grana, C., Calderara, S., Cucchiara, R.: A complete system for garment segmentation and color classification. *Mach. Vis. Appl.* **25**(4), 955–969 (2014)
13. Omohundro, S.M.: Bumptrees for efficient function, constraint and classification learning. In: NIPS, pp. 693–699 (1991)

14. Prince, S.: *Computer Vision: Models, Learning, and Inference*. Cambridge University Press, Cambridge (2012)
15. Schafer, J.B., Konstan, J.A., Riedl, J.: E-commerce recommendation applications. *Data Min. Knowl. Discov.* **5**(1–2), 115–153 (2001)
16. Shankar, D., Narumanchi, S., Ananya, H.A., et al.: Deep learning based large scale visual recommendation and search for e-commerce. [arXiv:1703.02344](https://arxiv.org/abs/1703.02344) (2017)
17. Srinivasan, S.S., Anderson, R., Ponnnavolu, K.: Customer loyalty in e-commerce: an exploration of its antecedents and consequences. *J. Retail.* **78**(1), 41–50 (2002)
18. Szegedy, C., Liu, W., Jia, Y., et al.: Going deeper with convolutions. In: *CVPR*, pp. 1–9 (2015)
19. Tajbakhsh, N., Shin, J.Y., Gurudu, S.R., et al.: Convolutional neural networks for medical image analysis: full training or fine tuning? *IEEE Trans. Med. Imag.* **35**(5), 1299–1312 (2016)
20. Zhang, S., Yao, L., Sun, A.: Deep learning based recommender system: a survey and new perspectives. [arXiv:1707.07435](https://arxiv.org/abs/1707.07435) (2017)