

# Enhancing Recommender Systems with Cross-Attribute Matrix Factorization and Shared User Embeddings

**Wen Liang**

University of California, San Diego  
La Jolla, CA 92093 wel245@eng.ucsd.com

**Youzhi Liang**

Department of Computer Science  
Stanford University  
Stanford, CA 94305 youzhi1@stanford.edu

**Zeng Fan**

University of California, San Diego  
La Jolla, CA 92093 zengfan1109@gmail.com

**Jianguo Jia**

Department of Computing  
Hong Kong Polytechnic University  
Hong Kong,, China  
jianguo1.jia@connect.polyu.hk

## Abstract

This paper introduces an advanced model for recommender systems, combining cross-attribute matrix factorization with shared user embeddings. This hybrid approach is designed to improve recommendation quality, particularly in scenarios with sparse data and for new users or items. By incorporating both user and item attributes into the matrix factorization process, our model effectively addresses the cold-start problem and increases the robustness of recommendations. Extensive evaluations on the MovieLens and Pinterest datasets confirm the superiority of our model, highlighting its potential for practical applications in diverse recommender system environments.

## Introduction

Recommender systems have become an indispensable component of many online platforms, guiding users through a plethora of content to find what best matches their preferences and needs. Traditional methods, such as collaborative filtering [1] and matrix factorization [2], have been successful in capturing the latent preferences of users based on historical interactions [3]. The rise of deep learning has revolutionized numerous domains, including computer vision where it has enabled breakthroughs in image recognition [4, 5], natural language processing [6, 7], and speech recognition [8]. Variational techniques and hybrid architectures of deep learning with data augmentations [9] elevate the performance of modeling to a higher standard [10, 11]. Leveraging these advancements, models such as NeuMF and NCF have incorporated neural networks in the realm of recommender systems,

Manuscript in submission 2023, do not distribute.

$$M = U V^T$$

Figure 1: Matrix Factorization

advancing beyond the linearity of matrix factorization to allow for more intricate representations of user-item interactions [12]. However, a gap remains in considering the intrinsic attributes of users and items, which can be pivotal in addressing challenges such as the cold-start problem and ensuring robustness, especially for long-tail items and users.

While the NeuMF and NCF model mark significant advancements in the field, they primarily focus on the interactions derived from user-item matrices without tapping into the wealth of information contained within intrinsic attributes of users and items. By overlooking these attributes, the models potentially miss out on crucial signals that could enhance recommendation quality, especially in scenarios where interaction data is sparse or nonexistent, such as in the cold-start problem. To address this, the Attribute-aware Deep CF model incorporates item and user features into the matrix factorization (MF) model [13]. It transforms these features into learnable embeddings, subsequently

integrating them into the multi-layer perceptron (MLP). However, we believe there is potential to further enhance this approach. Specifically, there’s an opportunity to exploit more features by cross-referencing user attributes with item attributes and vice versa.

Our *Cross-Attribute Matrix Factorization Model with Shared User Embedding* offers several distinctive contributions to the realm of recommender systems:

1. Incorporation of a Shared User Embedding: One of the major contributions of our model is the introduction of a shared user embedding. In situations where we encounter coldstart users, this shared embedding facilitates a basic recommendation process. It mitigate the robustness problems posed by relying on either randomly-initialized embeddings or immaturely learned user embeddings, which often aren’t reliable for providing meaningful recommendations.
2. Enhancement of the Existing MF Model with Cross-Attribute Interactions: Building on the foundation of the current matrix factorization (MF) model, we’ve added cross-attribute matrix factorization capabilities. To elaborate, each user embedding interacts with every item attribute embedding, and vice versa; every item interacts with user attribute embeddings. This intricate design ensures that we harness the full potential of existing data and features, fostering richer and more insightful recommendation outcomes.

Our model are further validated through a series of experiments. We used two benchmark datasets - Movielens and Pinterest [13] - consistently underscore the enhanced performance of our model. Notably, CAMF shines with pronounced superiority in environments marked by heightened dataset sparsity, thereby showcasing its practical applicability in real-world settings where user-item interactions may be limited.

### Related Works

#### Matrix Factorization

MF convert each user and item to a real-valued vector of latent variables. The MF estimates the interaction between user and item by using the inner product of two vectors:

$$r(u, i) = \gamma_u \cdot \gamma_i, \tag{1}$$

where the  $\gamma_u$  and  $\gamma_i$  denotes the latent factor for user  $u$  and item  $i$ , the  $r(u, i)$  denotes the interaction between this user and item. This method is very popular by combining good scalability with predictive accuracy. In addition, they offer much flexibility for modeling various real-life situations. However, the matrix factorization has several limitations that it is prone to recommend popular items and is vulnerable to sparse dataset.

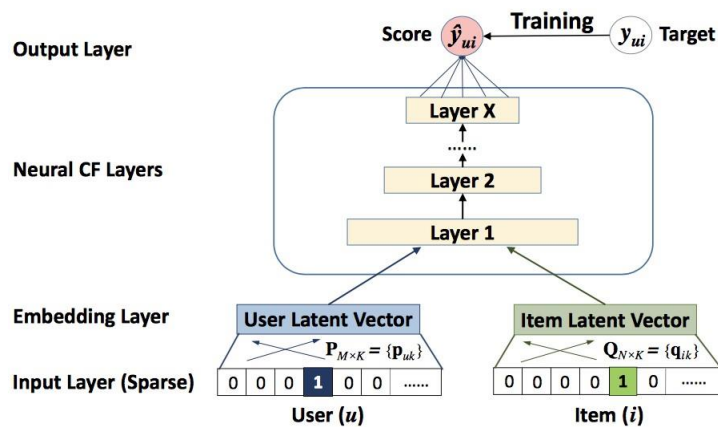


Figure 2: Neural collaborative filtering framework

**NCF, GMF and NeuMF Model**

Neural Collaborative Filtering (NCF) [12] serves as a foundational framework to model the interaction between users and items using neural networks. The general frame work of NCF is shown in Figure 2. Within the NCF paradigm, two primary models emerge: GMF (Generalized Matrix Factorization) and MLP (Multi-Layer Perceptron). GMF can be seen as a neural analog to traditional matrix factorization techniques, capturing linear patterns in user-item interactions. On the other hand, MLP is designed to discover intricate non-linear patterns. The NeuMF model blends the linear capabilities of GMF with the rich expressiveness of MLP to further improve recommendation performance. We mainly focus on GMF model and NeuMF model in our work. The mapping function of the GMF layer is defined as following:

$$\phi(pu, qi) = pu \odot qi \tag{2}$$

where  $\odot$  denotes the element-wise product of vectors. Then this vector can be projected to the output layer:

$$y^{ui} = \mathbf{aout}(hT(pu \odot qi)). \tag{3}$$

In the work presented by He et al. [12], the GMF (Generalized Matrix Factorization) layer, characterized by its activation function  $a_{out}$  and edge weights  $h$ , offers a pathway to revert back to the traditional matrix factorization model. As depicted in Figure 3, the NeuMF (Neural Matrix Factorization) model is an integration of both the GMF and the MLP (Multi-Layer Perceptron) models. Within this architecture, the GMF and MLP layers independently process the input embeddings. Subsequently, their outputs are concatenated, forming the basis for the NeuMF layer.

**Attributed-Aware Deep Collaborative Filtering Model**

The NCF framework takes into account the user-item interaction information, but it neglects the intrinsic attribute information of users/items. Wang et al. [14] proposed an Attribute-aware deep CF model, as depicted in Figure 4, which incorporates the attributes of users/items. They introduced a pairwise pooling layer following the embedding layer to effectively capture the relationship between users/items and their corresponding attributes. The operation for pairwise pooling is defined as:

$$p_u = \varphi_{\text{pairwise}}(u, g_u^t) = \sum_{t=1}^{V_u} u \odot g_u^t + \sum_{t=1}^{V_u} \sum_{t'=t+1}^{V_u} g_u^t \odot g_u^{t'} \tag{4}$$

$$q_i = \varphi_{\text{pairwise}}(i, g_i^t) = \sum_{t=1}^{V_i} i \odot g_i^t + \sum_{t=1}^{V_i} \sum_{t'=t+1}^{V_i} g_i^t \odot g_i^{t'} \tag{5}$$

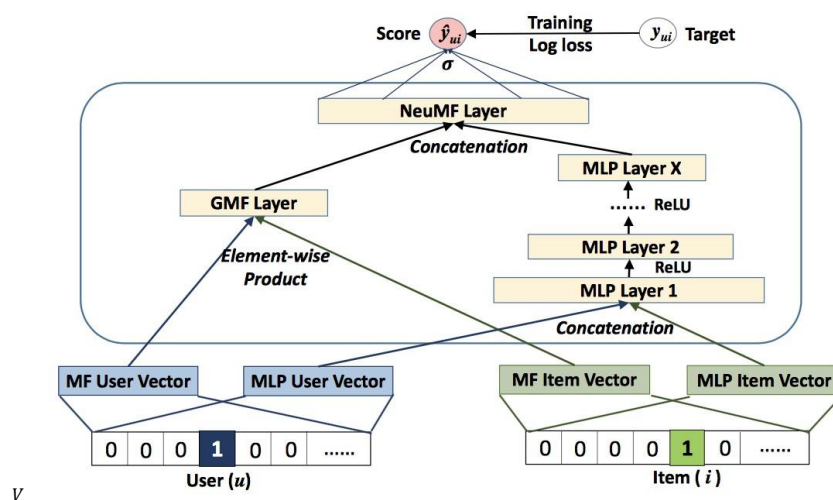


Figure 3: Neural Matrix Factorization Model

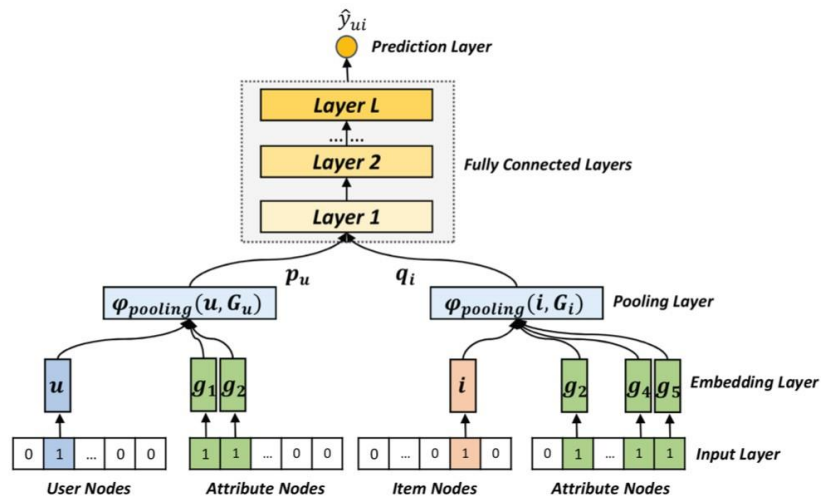


Figure 4: Attribute-aware Neural CF Model

After calculate  $p_u$  and  $q_i$ , the model takes the element-wise product of  $p_u$  and  $q_i$  into MLP layers for the final prediction.

## Method

### Shared User Embedding

Our experiments with GMF, NeuMF, and Attribute-aware Deep CF models showed that the GMF model trained faster and had nearly the same performance as the more complex NeuMF. The detailed results are discussed later. Due to its efficiency and effectiveness, we chose to focus on the GMF model.

Our goal was to improve recommendations in scenarios with sparse data and for new users or items. To address the cold-start and sparsity challenges, we use a shared user embedding embedding might encapsulate overall user traits. As an illustrative case, encountering a novel user or item could harness this shared embedding, furnishing more generic and resilient recommendation outcomes. To combine this shared user embedding with a specified user embedding, we introduce a weight  $\alpha$  ranging from 0 to 1 to adjust the balance between them by controlling the activation of the specified user embedding:

$$u_{merged} = \alpha u_{shared} + (1 - \alpha) u_{embedded} \quad (6)$$

$$\phi(u_{merged}, i) = u_{merged} \odot i \quad (7)$$

In alignment with the GMF paradigm delineated earlier, the weighted user vector undergoes an element-wise multiplication with the item vector. The balancing weight,  $\alpha$ , is determined by a singular hidden layer, taking another ensemble of item and user attributes as inputs:

$$\mathbf{z} = \begin{bmatrix} u \\ u_{attr} \\ i \\ i_{attr} \end{bmatrix} \quad (8)$$

$$\alpha = \text{Sigmoid}(\mathbf{w}^T \mathbf{z} + \mathbf{b}) \quad (9)$$

Herein,  $\mathbf{z}$  symbolizes the layer input, which incorporates attributes like item popularity and user historical interaction frequency. For instance, if attributes indicate an new user, the model might assign greater emphasis on the shared user vector, recommending universally acclaimed movies. Conversely, for universally appealing movies, the model could diminish the weight on specified user embedding, offering recommendations less tailored to individual tastes.

### Cross-Attribute Matrix Factorization Model

Another pivotal enhancement in our approach lies in the explicit incorporation of attribute information. Specifically, we multiplied the user vectors with item attributes and item vectors with user attributes. To accommodate this, we introduced dedicated embeddings tailored for attributes. This setup is not just a mere inclusion of attributes; it fundamentally changes the representation learning. By allowing these interactions, we aim to learn the underlying patterns and relationships between users, items, and their inherent attributes.

After facilitating these interactions, the resultant vectors are concatenated. This consolidated representation captures both the standalone and cross-attribute information. To process this rich representation further, we direct it through several neural network layers. These layers are designed to abstract higher-level patterns from the combined input.

Our comprehensive model, which encompasses the previously discussed Shared User Embedding and the current Cross-Attribute interactions, is illustrated in Figure 5.

### Experiments

In our evaluations, we leverage two prominent datasets: MovieLens and Pinterest. Both have been previously used in [12, 13] due to their rich attributes and feasibility for preprocessing.

#### MovieLens Dataset

For MovieLens, we employ the version containing 1 million ratings, ensuring that each user has rated at least 20 movies. Explicit ratings are transformed: a rating from a user denotes a label of 1 for the movie. To supplement this, we randomly sample 99 unobserved entries, tagging them as negative samples with a label of 0.

The dataset also provides user metadata encompassing gender, age, and occupation. This metadata is used as the user’s attributes, denoted as  $g_u^t$ , both in the Attribute-aware Deep CF model and our proposed architecture.

Movie data primarily comprise genres. Recognizing that movies can span multiple genres, each genre is treated as a distinct attribute. These are used as  $g_t^i$  during our experiments.

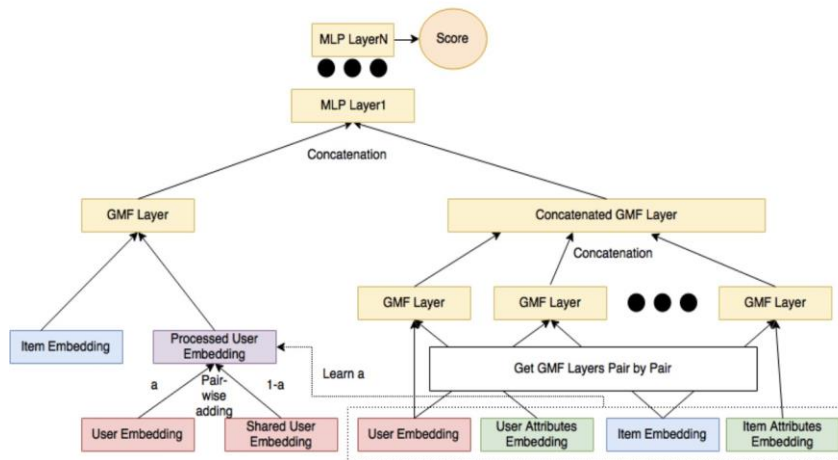


Figure 5: Proposed Attribute-aware Model

#### Pinterest Dataset

The Pinterest dataset, characterized by its enormity and sparsity, presents a unique challenge. With over 20% of the users having a singular pin, assessing algorithmic efficacy becomes intricate. Thus, we refined the dataset to include only users with at least 10 pins, which are labeled as 1. Like the MovieLens dataset, 99 unobserved entries are sampled as negatives labeled as 0.

User data in this dataset provide insights into pins and page categories. We group every 40 pins together, such that a user with, say, 35 pins falls into the first group and a user

with 41 pins falls into the second. The myriad of page categories, initially totaling 468, are consolidated into 45 main categories, given the observation that several categories possess scant samples. This structured data is then taken as the user's attribute input.

Items, represented by images, lack explicit attributes. Notably, an image may be pinned by diverse users. To derive an item's 'pin attribute', we employ the formula:

$$Pins_i = \sum_{(u,i) \in D} Pins_{ui} \quad (10)$$

In the equation,  $V_u$  signifies the user set in the dataset, and  $D$  encapsulates the user-item pairs.  $Pins_u$  stands for the user's pin count. This weighted summation offers insights into an item's popularity beyond a mere pin count. Analogous to user pins, items are grouped every 50, and this processed data serves as the item's attribute input.

### Evaluation and Metrics

We use the leave-one-out evaluation in our experiment. For each user, we randomly choose one of user's interaction and the 99 negative samples mentioned in Data Preprocessing section as test set. We rank the test item among the 100 items. The performance of the ranked list is evaluated by first 10 Hit Ratio (HR@10) and Normalized Discounted Cumulative Gain (NDCG@10). In other word, HR measures whether the positive test item is within the top-10 ranked list, and the NDCG will assign higher scores to hits at top ranks. These 2 metrics are calculated for each user and the average score is reported.

### Parameters Setting

We don't use the pretrained model from the paper. For each model, we initialize model parameters with a Gaussian distribution (mean = 0, standard deviation = 0.01). The models are trained with Adam optimizer with the batch size of 256. For stacked MLP layers, we use [32,16,8] as layer dimensions.




Figure 6: Performance of HR@10 on MovieLens

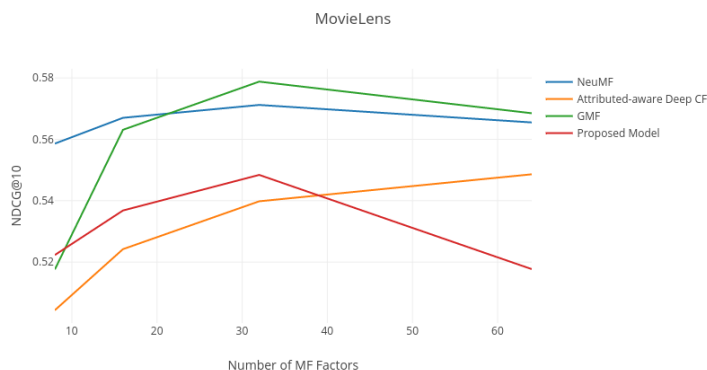


Figure 7: Performance of NDCG@10 on MovieLens



During training process, we will randomly generate 4 negatives instance for each positive instance. He et al. term the last hidden layer of NCF as predictive factors. We evaluate the factors of [8,16,32].

**Results**

The following Figure 6, Figure 7, Figure 8 and Figure 9 show the performance of HR@10 and NDCG@10 with respect to the number of matrix factorization factors.

We optimize all the models with log loss. The Figure 10 shows the training loss vs. epochs on MovieLens dataset for MLP, GMF, NeuMF, Attribute-aware Deep CF, and our proposed model.

The following tabel1, tabel2, tabel3 and tabel 4 show the performance of HR@10 and NDCG@10 with respect to the number of matrix factorization factors.

Table 1: Movielens HR@10 Results

#Factors	NeuCF	GMF	AA Deep CF	Our Model
8	0.8121	0.7485	0.7773	0.7923
16	0.8175	0.7774	0.8197	0.8213
32	0.8275	0.7935	0.8303	0.8304

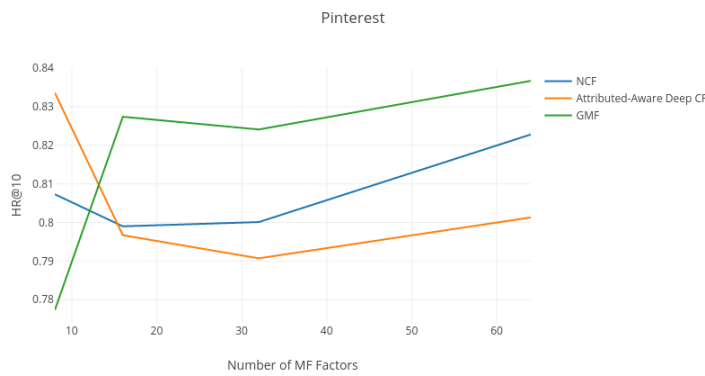


Figure 8: Performance of HR@10 on Pinterest




Figure 9: Performance of NDCG@10 on Pinterest

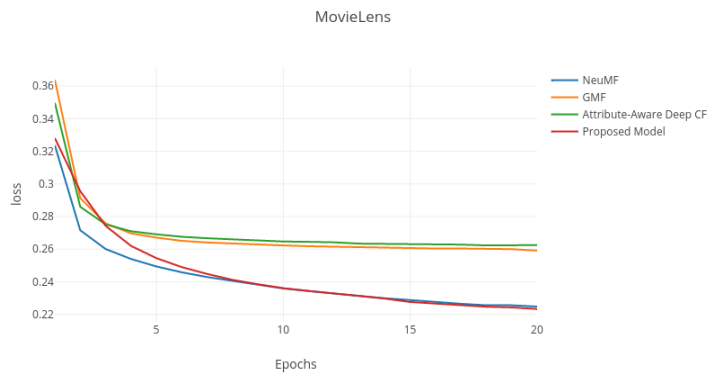


Figure 10: Training loss of all the models

Table 2: Movielens NDCG@10 Results

#Factors	NeuCF	GMF	AA Deep CF	Our Model
8	0.5286	0.5043	0.5176	0.5308
16	0.567	0.5242	0.5631	0.5668
32	0.5712	0.5398	0.5788	0.5784

Table 3: Pinterest HR@10 Results

#Factors	NeuCF	GMF	AA Deep CF	Our Model
8	0.8073	0.8336	0.7773	0.8311
16	0.799	0.7967	0.8274	0.8323
32	0.8001	0.7907	0.8241	0.8610

Table 4: Pinterest NDCG@10 Results

#Factors	NeuCF	GMF	AA Deep CF	Our Model
8	0.5536	0.5574	0.5176	0.5507
16	0.55	0.5425	0.5694	0.5665
32	0.5559	0.5465	0.5703	0.5920

## Discussion

The experimental results strongly underscore the robustness of our proposed model in comparison to existing models. Considering both the MovieLens and Pinterest datasets, our model consistently outperforms or is highly competitive with other methods across the board.

From the tables 1, 2, 3, and 4, it is evident that our model exhibits superior performance, especially with an increase in the number of matrix factorization factors. This superiority in HR@10 metrics is particularly significant given that our model is directly designed to tap into the rich information available in user-item interactions, without compromising on the intrinsic attributes.

The Attribute-Aware Deep CF (AA Deep CF) model's underwhelming performance on these recommendation tasks came as a surprise. One reason for its subpar results might be the pooling layer they adopted. By leveraging user, item, and attribute information, the model may inadvertently discard pivotal details, leading to a dilution of recommendation quality. Additionally, it's noteworthy to mention that the AA Deep CF model was primarily fashioned for social networks and travel recommendation tasks. While it might excel in scenarios with abundant user and item attributes, it evidently struggles when confronted with the data constraints of our datasets. The confluence of the AA Deep CF model with another model optimized for social network relations hints at design decisions which may not necessarily align with our task, and thus, the observed performance disparities.

In summary, our results illuminate the importance of crafting recommendation models tailored to the nuances of the dataset in question. Our model, by astutely leveraging



both interaction data and inherent attributes, solidifies its place as a formidable contender in the recommendation system arena.

## Conclusion

This research highlights two pivotal advancements in recommendation systems. First, we introduced the concept of a *Shared User Embedding*, which offers a robust solution to the cold-start problem by reducing dependence on unstable embeddings. Concurrently, our refined matrix factorization approach, utilizing *Cross-Attribute Interactions*, ensures an in-depth comprehension of user-item dynamics, leveraging every available attribute. Collectively, these innovations represent a substantial progression in providing nuanced and dependable recommendations, indicating an optimistic trajectory for subsequent research in this field.

## References

- [1] Badrul Sarwar, George Karypis, Joseph Konstan, and John Riedl. Item-based collaborative filtering recommendation algorithms. In *Proceedings of the 10th international conference on World Wide Web*, pages 285–295, 2001.
- [2] Yehuda Koren, Robert Bell, and Chris Volinsky. Matrix factorization techniques for recommender systems. *Computer*, 42(8):30–37, 2009.
- [3] Paul Resnick and Hal R Varian. Recommender systems. *Communications of the ACM*, 40(3):56–58, 1997.
- [4] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition, 2015.
- [5] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. Imagenet classification with deep convolutional neural networks. *Commun. ACM*, 60(6):84–90, may 2017.
- [6] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need, 2023.
- [7] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding, 2019.
- [8] Dario Amodei, Rishita Anubhai, Eric Battenberg, Carl Case, Jared Casper, Bryan Catanzaro, Jingdong Chen, Mike Chrzanowski, Adam Coates, Greg Diamos, Erich Elsen, Jesse Engel, Linxi Fan, Christopher Fougner, Tony Han, Awni Hannun, Billy Jun, Patrick LeGresley, Libby Lin, Sharan Narang, Andrew Ng, Sherjil Ozair, Ryan Prenger, Jonathan Raiman, Sanjeev Satheesh, David Seetapun, Shubho Sengupta, Yi Wang, Zhiqian Wang, Chong Wang, Bo Xiao, Dani Yogatama, Jun Zhan, and Zhenyao Zhu. Deep speech 2: End-to-end speech recognition in english and mandarin, 2015.
- [9] Wen Liang, Youzhi Liang, and Jianguo Jia. Miamix: Enhancing image classification through a multi-stage augmented mixed sample data augmentation method. *arXiv preprint arXiv:2308.02804*, 2023.
- [10] Youzhi Liang and Wen Liang. Reswcae: Biometric pattern image denoising using residual wavelet-conditioned autoencoder. *arXiv preprint arXiv:2307.12255*, 2023.
- [11] Youzhi Liang, Wen Liang, and Jianguo Jia. Structural vibration signal denoising using stacking ensemble of hybrid cnn-rnn. *arXiv e-prints*, pages arXiv–2303, 2023.
- [12] Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. Neural collaborative filtering. pages 173–182, 2017.
- [13] Xue Geng, Hanwang Zhang, Jingwen Bian, and Tat-Seng Chua. Learning image and user features for recommendation in social networks. In *2015 IEEE International Conference on Computer Vision (ICCV)*, pages 4274–4282, 2015.
- [14] Xiang Wang, Xiangnan He, Liqiang Nie, and Tat-Seng Chua. Item silk road: Recommending items from information domains to social users. *CoRR*, abs/1706.03205, 2017.