AGTR: Adversarial Generation of Target Review for Rating Prediction

Huilin Yu¹ · Tieyun Qian¹ · Yile Liang¹ · Bing Liu¹

Received: 28 April 2020 / Revised: 17 August 2020 / Accepted: 24 August 2020 / Published online: 17 September 2020
© The Author(s) 2020

Abstract

Recent years have witnessed a growing trend of utilizing reviews to improve the performance and interpretability of recommender systems. Almost all existing methods learn the latent representations from the user's and the item's historical reviews and then combine these two representations for rating prediction. The fatal limitation in these methods is that they are unable to utilize the most predictive review of the target user for the target item since such a review is not available at test time. In this paper, we propose a novel recommendation model, called AGTR, which can *generate the unseen target review with adversarial training for rating prediction*. To this end, we develop a unified framework to combine *the rating tailored generative adversarial nets* for synthetic review generation and *the neural latent factor module* using the generated target review along with historical reviews for rating prediction. Extensive experiments on four real-world datasets demonstrate that our model achieves the state-of-the-art performance in both rating prediction and review generation tasks.

Keywords Recommender systems · Review aware recommendation · Generative adversarial network

1 Introduction

A user's rating indicates his/her attitude toward an purchased item. Rating prediction aims to predict the user's ratings on unrated items which may reflect his/her potential interests on these items. Collaborative filtering (CF) approaches, which mainly depend on historical ratings, have aroused great research interests and become the dominant method in recommender systems. As a typical CF technique, matrix factorization (MF) learns the latent features of users and items by decomposing the user-item rating matrix and then uses these two feature vectors to predict the rating that the user would assign to the item.

MF is the most widely used technique for rating prediction. However, MF-based methods suffer from the data

sparsity problem and the predicted rating lacks the interpretability on why the user gives high or low scores. To tackle these issues, textual reviews have become a key complementary data source to enhance the performance and interpretation of the rating prediction task [1, 8, 26, 41]. In particular, due to the power of nonlinear combination of different types of information, impressive progress has been made by applying deep neural networks to this problem [3, 4, 6, 24, 35, 42].

The pioneering work by Zheng et al. [42] proposed a DeepCoNN model to represent both users and items in a joint manner using all the reviews of users and items. As proven in [3], *the target review*, which is written by the target user for the target item, provides much of the predictive value for rating prediction. The performance of the DeepCoNN model [42] drops severely when the target reviews are omitted. Indeed, the target review usually contains the target user's preference on the target item's attributes or properties and is closely related to the rating score. However, the target review will not be available at test time in real-world recommendation settings. The hereafter studies along this line do not access the target reviews in the validation and test set at any time to simulate a real-world scenario. Clearly, the inherent limitation in these methods is that they are unable to utilize the most predictive target review.

In light of this, we propose a novel framework, namely AGTR, to generate the target review for rating prediction.

✉ Tieyun Qian
qty@whu.edu.cn

Huilin Yu
huilin_yu@whu.edu.cn

Yile Liang
liangyile@whu.edu.cn

Bing Liu
liub@uic.edu

¹ School of Computer Science, Wuhan University, Wuhan, Hubei, China

Our model has two distinguishing characteristics. Firstly, we *generate the target review with rating tailored generative adversarial nets (RTGAN)* which incorporates the rating into its objective function in addition to the user's and the item's historical reviews. Secondly, we *develop a neural latent factor module (NLFM) to accurately predict the rating score* by learning from the generated target review which encodes the user's specific preference on the item. In such a way, the target review naturally provides guidance for the rating prediction task beyond the above-mentioned review-aware deep recommendation approaches [3, 4, 6, 24, 35]. Meanwhile, the rating drives the RTGAN module to produce a target review conveying consistent sentiment with the rating score.

We are aware of a few existing studies for generating reviews [5, 25, 37] or abstractive tips [21]. However, our AGTR model is fundamentally different from the NRT [21], MT [25], and CAML [5] models, in the sense that all these approaches do not directly utilize the target review for rating prediction. Although the neural memory (NM) model proposed by Wang and Zhang [37] also integrates the target review in their prediction step, we distinguish our model with NM in both the review generation and rating prediction modules. We present a conditional GAN architecture for review generation, whereas NM [37] uses the sequence-to-sequence (seq2seq) [33] generative model. More importantly, we design a novel neural latent factor model to stress the target review to make good use of its predictive ability, while NM simply feeds the target review as the input of rating prediction in the last layer.

We have implemented the proposed AGTR model in both the rating prediction and review generation tasks. Empirical evaluation on four real world datasets proves that our model achieves the state-of-the-art performance on both tasks.

The rest of this paper is organized as follows. Section 2 reviews the related work. Section 3 introduces the preliminary and problem definition. Section 4 presents our AGTR model in detail. Section 5 gives the experimental results. Finally Sect. 6 concludes the paper.

2 Related Work

We summarize the research progress in review-aware rating prediction, categorized by the traditional methods and deep learning-based methods. We omit the classic collaborative filtering methods which do not use text reviews.

2.1 Traditional Methods

When integrating review texts, the traditional methods can be roughly classified into three categories. The first one is to extract useful textual information such as topics or aspects from review texts and learn latent factors from

ratings and then link the textual information and latent factors together using linear [2, 26, 34, 40] or Bayesian combination [23, 38]. The second one is by extending the latent factor model [7, 13, 28, 29, 41] to encode the textual influence. The third one is to modify graphic models to include latent factors from ratings [1, 8, 11, 36].

2.2 Deep Learning-Based Methods

The first type of deep learning-based methods only uses historical reviews without generating the target review. These approaches differ mainly in how they combine reviews with ratings. For example, NARRE [4] jointly learns hidden latent features for users and items using two parallel neural networks with the attention mechanism [4]. TARMF [24] adopts a neural network for mutual learning between reviews and ratings, where the features from reviews are optimized by an attention-based GRU network. A³NCF [6] extracts features from reviews using topic models and fuses them with the embeddings from ratings, and it then captures a user's attention on the item with an attention network. MPCN [35] presents a pointer-based co-attention mechanism which can extract multiple interactions between user and item reviews.

The second type of deep learning-based methods generates the target review, but not all of them exploit the predictive ability of the target review. As we have illustrated this issue in the introduction section, here we discuss these methods on how they generate target reviews. NRT [21] is mainly for the purpose of enhancing explainability by generating tips based on a standard generative model with the GRU architecture. NM [37] adopts the seq2seq modeling [33] technique for review generation. Meanwhile, MT [25] uses an adversarial training process which helps overcome the problem of exposure bias in seq2seq models.

Our proposed AGTR model falls into the second type of deep learning-based methods. It has the following three advantages. Firstly, the exploiting of target reviews provides a target-dependent modeling of user and item characteristics. Secondly, our model incorporates rating as one of the conditions in both the generator and discriminator when generating the target review with GAN mechanism. This is different from the previous MT [25] method which relies purely on reviews without the guidance of ratings. As we will show in the experimental part, the incorporation of ratings helps the model generate high-quality target reviews and further improves the rating prediction task. Thirdly, when utilizing the target reviews, the previous MT method adopts a traditional MF method for rating prediction, which does not take the target review into consideration. In contrast, our model can fully leverage the target review with a carefully designed neural latent factor model.

Table 1 Notations used in this paper

Variable	Interpretation
$\mathcal{U}$	User set
$\mathcal{I}$	Item set
$\mathcal{R}$	Rating set
$\mathcal{D}$	Review set
$u \in \mathcal{U}$	A user $u \in \mathcal{U}$
$i \in \mathcal{I}$	An item $i \in \mathcal{I}$
$r_{ui} \in \mathcal{R}$	User u 's rating on item i
$d_{ui} \in \mathcal{D}$	User u 's review on item i
$d_u \subset \mathcal{D}$	User u 's all reviews except d_{ui}
$d_i \subset \mathcal{D}$	Item i 's all reviews except d_{ui}
$\hat{r}_{ui}$	User u 's predicted rating on item i
s_{ui}	User u 's generated review on item i
$\mathbf{c}_{ui}$	Condition vector
$\mathbf{f}_t$	Leaked feature vector in MANAGER module
$\mathbf{g}_t$	Guiding goal vector in WORKER module

3 Problem Definition and Preliminary

This section presents the problem definition and preliminary.

3.1 Problem Definition

The goal of rating prediction is to predict the rating of a given user to an item. Additionally, in our task, we also generate a target review in the form of a sentence that the user will write for the item. More formally, let $\mathcal{U}$ be a user set and $\mathcal{I}$ be an item set, and $\mathcal{D}$ be a review set on the items in $\mathcal{I}$ written by a set of users in $\mathcal{U}$. Each review d_{ui} written by user u on item i has an accompanying rating r_{ui} indicating u 's overall satisfaction toward i . We refer to all historical reviews written by the user, i.e., except that on item i , as the *user's historical review document* d_u . Similarly, the set of historical reviews on item i , except the one written by u , is referred to as the *item's historical review document* d_i . Each training instance is denoted as a sextuple $(u, i, d_{ui}, r_{ui}, d_u, d_i)$. The goal is to predict a rating $\hat{r}_{ui}$ and learn a synthetic target review s_{ui} for each item i that u does not interact with.

For ease of presentation, we summarize the notations in Table 1.

3.2 Preliminary

One key property of our model is to utilize the generative adversarial nets (GAN) to generate the target review. GAN is proposed by Goodfellow et al. [9]. GAN is inspired by the two-player zero-sum game. The main components in GAN is a generator G and a discriminator D . These two components are trained simultaneously under the adversarial learning idea. The optimization process of GAN is a mini-max

game process, where the target is to reach Nash equilibrium [30]. Later on, Salimans et al. [32] present several new architectural features and training procedures, and the goal is to improve the stability of training and the perceptual quality of GAN samples. Mirza and Osindero [27] construct the conditional adversarial nets where both the generator and discriminator are conditioned on some auxiliary information like class labels or data from other modalities. Yu et al. [39] further combine sequence to sequence model with the generative adversarial nets. GAN has been applied to the research fields like vision tasks [32], speech and language processing [20, 39].

In this subsection, we give a brief introduction to the original GAN and the closely related conditional GAN and sequence GAN. The generator G in GAN is used to capture the data distribution, while the discriminator D estimates the probability that a sample comes from the training data rather than G . In order to learn the distribution p_g over data x in G , we first define a prior on input noise variables $p_z(z)$, then we define two mapping functions $G(z, \theta_g)$ and $D(x, \theta_d)$ which represents a mapping from a prior noise distribution $p_z(z)$ to data space and outputs a single scalar representing the probability that x comes from training data rather than p_g , respectively. Both G and D can be represented by a multilayer perceptron with parameters θ_d and θ_g .

The objective of D is to maximize the probability of assigning the correct label to both training examples and samples from G . Meanwhile, the objective of G is to minimize $\log(1 - D(G(z)))$. That is to say, D and G play the two-player min-max game with the objective function $L(G, D)$ defined as follows.

$$\min_G \max_D L(D, G) = E_{x \sim p_{\text{data}}(x)} [\log D(x)] + E_{z \sim p_z(z)} [\log(1 - D(G(z)))] \quad (1)$$

The generative models which learn complex distributions via variational inference, e.g., VAE [18] and its variants like [15, 17], suffer from the difficulty of approximating intractable probabilistic computations. In contrast, GAN does not assume the generated data belongs to a predefined distribution. Instead, it uses a discriminative model to guide the training of the generative model has succeeded in generating real-valued data. However, GAN also has the disadvantage that there is no control on the data being generated. Hence Mirza and Osindero [27] extend GAN to a conditional model to incorporate extra information. Formally, the objective function of the conditional GAN can be defined as follows.

$$\min_G \max_D V(D, G) = E_{x \sim p_{\text{data}}(x)} [\log D(x|y)] + E_{z \sim p_z(z)} [\log(1 - D(G(z|y)))] \quad (2)$$

When GAN is introduced into the sequence generation, there are two main challenges. Firstly, the sequence data is not

Fig. 1 The architecture of our AGTR model

continuous and needs a sampling procedure when generating the word; hence, it is hard to pass the gradient update from the discriminative model D to the generative model G . Secondly, D can only assess a complete sequence. However, sequence generation is performed one by one, and it is nontrivial to balance the current score and the future one on a partially generated sequence. To address the above problems, Yu et al. [39] propose a sequence generation GAN framework which adopts the reinforcement learning and Monte Carlo search strategy, and then, Guo et al. [10] further present a LeakGAN model for generating long sequence.

Since our task is to generate the target review which is usually a long sequence, we adopt the LeakGAN as the basic framework and we extend it to combine the conditional GAN into the model. We will present the detail later in our model.

4 Our Proposed AGTR Model

In this section, we introduce our proposed AGTR model. We begin with the overall architecture and then go to the details of two modules.

4.1 Model Overview

Our AGTR model consists of two modules. One is the rating tailored GAN (RTGAN), which takes the rating as an important condition in the generator and the discriminator of GAN for review generation. The other is the neural latent factor module (NLFM) that leverages the generated target review along with the historical reviews for rating prediction using a neural network. The overall architecture of our model is shown in Fig. 1.

4.2 Rating Tailored GAN (RTGAN) Module

We have two basic assumptions for generating the synthetic target review s_{ui} . Firstly, s_{ui} should reflect the user u 's preferences and the item i 's features. Secondly, the sentiment expressed in s_{ui} should be consistent with the rating score r_{ui} . Following these assumptions, we design our rating tailored GAN (RTGAN) module conditioned on three types of information: (1) the user's historical review document d_u to capture u 's preferences, (2) the item's historical review document d_i to represent i 's features, and (3) the rating r_{ui} of the user u to the item i to serve as a constraint. During training, we learn a generator G using three types of condition information to produce a synthetic review, and a discriminator D to distinguish it with the real one.

4.2.1 Condition Information Encoder

We first introduce the condition information encoder (the left gray part in Fig. 1). It maps three types of condition information into user's general preference embedding $\mathbf{g}_u$, item's feature embedding $\mathbf{g}_i$, and the rating embedding $\mathbf{h}_{ui}$.

We take the process of mapping user's review document d_u to his/her preference embedding $\mathbf{g}_u$ as an example. Each word in d_u is randomly initialized as a d dimensional vector, and each review in d_u is transformed into a matrix with the fixed length T (padded with 0 if necessary). Since the text processing is not the focus of this study, we take the same TextCNN [4] approach to encode each review in d_u . Essentially, TextCNN can be summarized as a CNN structure followed by an attention mechanism. The convolution layer consists of m neurons. Each neuron is associated with a filter $\mathbf{K} \in \mathbb{R}^{r \times d}$ which produces features by applying convolution operator on word vectors. Let $\mathbf{V}_{ui}$ be the embedding matrix

corresponding to the l th review in d_u , the j th neuron in CNN produces its feature as:

$$\mathbf{z}_j = \sigma(\mathbf{V}_{ul} * \mathbf{K}_j + b_j), \quad (3)$$

where $*$ is convolution operator, b_j is bias term and σ is a nonlinear RELU activation function. We then apply a max-pooling operation to obtain the output feature $\mathbf{o}_j$ corresponding to this neuron. By concatenating the output from all m neurons, the convolution layer can produce the embedding $\mathbf{o}_{ul}$ of the review d_{ul} as:

$$\mathbf{o}_{ul} = [\mathbf{o}_1, \mathbf{o}_2, \mathbf{o}_3, \dots, \mathbf{o}_m], \quad (4)$$

After getting the embedding for each review in d_u , the attention mechanism is adopted to get the weights for these reviews. The attention a_{ul} for review d_{ul} is defined as:

$$a_{ul}^* = \mathbf{h}_a^T \text{ReLU}(\mathbf{W}_O \mathbf{o}_{ul} + \mathbf{W}_i \mathbf{i}_{ul} + b_1) + b_2, \quad (5)$$

where $\mathbf{h}_a \in \mathbb{R}^t$, $\mathbf{W}_O \in \mathbb{R}^{t \times k_1}$, $\mathbf{W}_i \in \mathbb{R}^{t \times k_2}$, $b_1 \in \mathbb{R}^t$, $b_2 \in \mathbb{R}^1$ are model parameters, $\mathbf{i}_{ul} \in \mathbb{R}^K$ is the embedding of the item which the user write this review for.

A softmax function is used to normalize the above a_{ul}^* to get the final attention a_{ul} . The user's u general preference embedding $\mathbf{g}_u$ is then calculated as the attention weighted sum of all reviews $d_{ul} \in d_u$, i.e.,

$$\mathbf{g}_u = \sum_{l=1, \dots, |d_u|} a_{ul} \mathbf{o}_{ul} \quad (6)$$

The mapping of the item's review document d_i to its feature embedding $\mathbf{g}_i$ is in the same way. We first encode each review d_{il} into an embedding $\mathbf{o}_{il}$ using convolution and pooling operation in TextCNN. We then employ the attention mechanism to combine the review embeddings with different weights.

$$\mathbf{g}_i = \sum_{l=1, \dots, |d_i|} a_{il} \mathbf{o}_{il} \quad (7)$$

The mapping from the original rating r_{ui} to an one-hot embedding $\mathbf{h}_{ui}$ is straight-forward. We simply discretize the rating r_{ui} into a m -dimension vector ($m = 5$ in our case). If the value falls into an interval, the corresponding dimension is set to 1 and other dimensions are set to 0. For example, a rating $r_{ui} = 3.78$ will be mapped into a $\mathbf{h}_{ui}$ as $(0, 0, 0, 1, 0)^T$. Note that the rating r_{ui} is known only in training. During validation or test, we will use a basic rating from NLFM module instead. The detail will be given later.

4.2.2 RTGAN for Target Review Generation

A good number of generative methods have been proposed for text generation in recent years, such as seq2seq [33]-based models, SeqGAN [39], and RankGAN [22]. Since

the reviews are usually long (with average length > 40), we adopt the state-of-the-art LeakGAN [10] model to generate reviews in this paper and extend it by incorporating three types of condition information into both the generator and the discriminator.

Conditional generator Starting from the random state, LeakGAN generates texts via the adversarial generation of synthetic texts against real texts. This implies that, if simply adopting LeakGAN in our model, the generated reviews are only ensured to be written in a human-like style. However, we need to generate the target review that is written by a specific user for a specific item.

In order to provide additional information for guiding the target review generation, we incorporate LeakGAN with the conditional GAN by taking three types of information as the condition of the generator in LeakGAN. We call the combination of these three types of information as a condition vector $\mathbf{c}_{ui}$ and define it as:

$$\mathbf{c}_{ui} = \mathbf{g}_u \oplus \mathbf{g}_i \oplus (\mathbf{W}_r * \mathbf{h}_{ui}), \quad (8)$$

where $\mathbf{W}_r$ is a mapping matrix to transform the sparse $\mathbf{h}_{ui}$ to a dense vector.

Similar to many text generation methods [10, 25], we employ a decoder GRU to iteratively generate a review word by word. Different from these methods, the decoder layer in our RTGAN module is conditioned on $\mathbf{c}_{ui}$, which is the combination of three types of information. By doing so, our generator produces a synthetic target review that reflects not only the user u 's preferences but also the item i 's features. Moreover, the sentiment contained in the synthetic review is also forced to match the rating score.

To ensure that the condition information is maintained during the generation process, the condition vector $\mathbf{c}_{ui}$ is concatenated with the word vector before it is fed into the decoder GRU at each time step. Suppose $\mathbf{x}_t$ is the embedding for the current word being processed at time step t , the concatenated vector $\mathbf{x}_t' = \mathbf{c}_{ui} \oplus \mathbf{x}_t$ is input into the decoder GRU to get the hidden state $\mathbf{h}_t$. And then, the hidden state $\mathbf{h}_t$ is multiplied by an output projection matrix and passed through a softmax over all the words in the vocabulary to obtain the probability of each word in the current context. Finally, the output word y_t at time t is sampled from the multi-nominal distribution through a softmax layer.

The difference between the generator in our RTGAN module and that in LeakGAN is that our generator is conditioned on the additional information as discussed above. For learning, we follow the generator training method in LeakGAN [10] by adopting a hierarchical architecture to effectively generate long texts.

Briefly, the hierarchical generator G consists of a high-level MANAGER module and a low-level WORKER module. At each step, the MANAGER receives a leaked feature

vector $\mathbf{f}_t$ (which is the last layer in discriminator D), and uses $\mathbf{f}_t$ to form the guiding goal vector $\mathbf{g}_t$ for the WORKER module. Compared to the scalar classification probability of D , the leaked feature vector $\mathbf{f}_t$ is a much more informative guiding signal for G , since it tells what the position of currently generated word is in the extracted feature space.

The loss for the MANAGER module is defined as:

$$L_{G_M} = - \sum_{t=1}^T Q(\mathbf{f}_t, \mathbf{g}_t) * d_{\cos}(\mathbf{f}_{t+c} - \mathbf{f}_t, \mathbf{g}_t), \quad (9)$$

where $Q(\mathbf{f}_t, \mathbf{g}_t)$ is the expected reward (the classification probability output by D) under the current policy, and $d_{\cos}$ represents the cosine similarity between the change of leaked feature representation of discriminator after c -step transition (from $\mathbf{f}_t$ to $\mathbf{f}_{t+c}$) and the goal vector $\mathbf{g}_t$, and T is the maximum sequence length we set for review. The loss function aims to force the goal vector to match the transition in the feature space while achieving high reward. Meanwhile, the loss for the WORKER module is defined as:

$$L_{G_W} = - \sum_{t=1}^T r_t^I \cdot p(y_t | s_{t-1}, \mathbf{c}_{ui}), \quad (10)$$

where $p(y_t | s_{t-1}, \mathbf{c}_{ui})$ denotes the conditional generative probability of the next token y_t given a sequence $s_{t-1} = [y_0, y_1, \dots, y_{t-1}]$ and the condition vector $\mathbf{c}_{ui}$ in WORKER module. r_t^I is the intrinsic reward defined as:

$$r_t^I = \frac{1}{c} \sum_{i=1}^T d_{\cos}(\mathbf{f}_t - \mathbf{f}_{t-i}, \mathbf{g}_{t-i}) \quad (11)$$

The objective in G is to minimize L_{G_M} and L_{G_W} in two modules, which are alternatively trained while fixing the other.

Conditional discriminator The discriminator learns to distinguish the ground-truth review d_{ui} from the synthetic one s_{ui} . We adopt the same CNN structure in the generator to process review texts, and we can get the embedding $\mathbf{d}_{ui}$ for d_{ui} and $\mathbf{s}_{ui}$ for s_{ui} , respectively. Different from the discriminator that only distinguishes between the real and the synthetic one, our discriminator needs to determine whether the review is related to the user and the item, and whether the review is written by the user for this item. Therefore, we take the condition information $\mathbf{c}_{ui}$ into account in the discrimination as well. The loss for the discriminator D is defined as:

$$L_D = -(\log(D(\mathbf{d}_{ui} | \mathbf{c}_{ui})) + \log(1 - D(\mathbf{s}_{ui} | \mathbf{c}_{ui}))), \quad (12)$$

where $D()$ is the probability function computed by applying a softmax layer to the concatenation of $\mathbf{d}_{ui}/\mathbf{s}_{ui}$ and $\mathbf{c}_{ui}$. The objective in D is to maximize the probability of classifying the ground-truth review as positive and to minimize the probability of classifying the synthetic one as authentic.

The training of G and D in RTGAN module is an adversarial process. The goal of generator is to produce the most indistinguishable synthetic reviews to fool the discriminator, while the discriminator aims to distinguish synthetic and ground-truth reviews as much as possible. Hence, we iteratively train G and D to reach an equilibrium. The procedure for generating a target review $\hat{s}_{ui}$ is illustrated in Algorithm 1.

Algorithm 1 Procedure for generating a target review $\hat{s}_{ui}$

Input: the condition vector $\mathbf{c}_{ui}$, the initial word vector x_0 , and the predefined sequence length T ;

Output: the generated target review $\hat{s}_{ui}$;

- 1: **for** each t in T **do**
 - 2: Apply a look-up layer on x_t , get the word embedding $\mathbf{x}_t$;
 - 3: Concatenate $\mathbf{x}_t$ and $\mathbf{c}_{ui}$ to get $\mathbf{x}'_t$;
 - 4: Input $\mathbf{x}'_t$ to the current memory of LSTM in Worker, output $\mathbf{O}_t$;
 - 5: Input $S_t = [x_1, x_2, \dots, x_t]$ to Discriminator and get $\mathbf{f}_t$;
 - 6: Take $\mathbf{f}_t$ as the leakage information, convert it by Manager to get the goal vector $\mathbf{g}_t$
 - 7: Multiply $\mathbf{g}_t$ and $\mathbf{O}_t$, and get the probability $p_{x_{t+1}}$ for the next word x_{t+1}
 - 8: Sample on $p_{x_{t+1}}$ using the multi-nominal function and get the next word x_{t+1}
 - 9: **return** $\hat{s}_{ui} = [x_1, x_2, \dots, x_T]$
-

In Algorithm 1, we generate the words in a sentence one by one. We first get the word embedding $\mathbf{x}_t$ in line 2 and concatenate it with the condition vector $\mathbf{c}_{ui}$ to get a new word embedding $\mathbf{x}'_t$ in line 3. Lines 4–7 follows the standard generation procedure in LeakGAN. Specifically, in line 4, the Worker module takes the new word embedding $\mathbf{x}'_t$ as input and outputs a matrix $\mathbf{O}_t$ which represents the current vector for all words using an LSTM. Line 5 sends the current sentence $S_t = [x_1, x_2, \dots, x_t]$ into the Discriminator and get the output vector $\mathbf{f}_t$. In line 6, the Manager also implemented by an LSTM takes the extracted feature vector $\mathbf{f}_t$ as its input, and outputs a goal vector $\mathbf{g}_t$, which is in turn fed into the Worker module to guide the generation of the next word in line 7. Line 8 is used to get the next word x_{t+1} . Finally line 9 returns the entire sentence.

4.3 Neural Latent Factor Model (NLFM) Module

Inspired by the neural latent factor models in [4, 12], we propose our NLFM module by extending these neural models in the following ways. Firstly, we represent *general latent factors* of user and item merely based on historical reviews without ratings. Secondly, we extend to exploit the *special latent factors* which encode the user's preference on the item in the target review.

Specifically, the embeddings of user preferences and item features, i.e., $\mathbf{g}_u$ and $\mathbf{g}_i$, are passed from the RTGAN module, and then, we map them with a hidden layer to get the general latent factors of user and item. To obtain the special latent factors, we transform the target review d_{ui} (s_{ui} when testing) through a CNN structure and a hidden layer as follows:

$$\mathbf{p}_u = \tan h(\mathbf{W}_{su} * \text{CNN}(d_{ui}) + b_{su}), \quad (13)$$

$$\mathbf{p}_i = \tan h(\mathbf{W}_{si} * \text{CNN}(d_{ui}) + b_{si}), \quad (14)$$

where $\text{CNN}()$ is a convolutional neural network that maps the target review d_{ui} into a feature vector, and $\mathbf{W}_{su}$, $\mathbf{W}_{si}$ are the projection matrices and b_{su} , b_{si} are biases.

Combining the general and special latent factors together, we can obtain the user's and item's overall representations:

$$\mathbf{f}_u = \tan h(\mathbf{W}_{gu} * \mathbf{g}_u) + \tan h(\mathbf{W}_{pu} * \mathbf{p}_u), \quad (15)$$

$$\mathbf{f}_i = \tan h(\mathbf{W}_{gi} * \mathbf{g}_i) + \tan h(\mathbf{W}_{pi} * \mathbf{p}_i), \quad (16)$$

where $\mathbf{W}_{gu}$, $\mathbf{W}_{pu}$, $\mathbf{W}_{gi}$, $\mathbf{W}_{pi}$ are weight matrices.

We then pass these two overall representations $\mathbf{f}_u$ and $\mathbf{f}_i$ to a prediction layer to get a real-valued rating $\hat{r}_{ui}$:

$$\hat{r}_{ui} = \mathbf{f}_u^T \mathbf{f}_i + b_u + b_i + b, \quad (17)$$

where b_u , b_i , and b denotes the user bias, item bias and global bias, respectively. Clearly, our predicted rating $\hat{r}_{ui}$ encodes the general user interests and item features as well as the user's specific interest on this item.

Since rating prediction is actually a regression problem, a commonly used squared loss is adopted as the objective function for our NLFM module:

$$L_r = \sum_{u,i \in \mathcal{U}, \mathcal{I}} (\hat{r}_{ui} - r_{ui})^2, \quad (18)$$

where $\mathcal{U}$, $\mathcal{I}$ denotes the user and item set, respectively, and r_{ui} is the ground-truth rating assigned by u on i .

4.4 Training and Prediction

We iteratively train the RTGAN and NLFM modules. Since these two modules share the parameters in the historical reviews encoder layer, the parameters will be iteratively updated.

At the time of validation and testing, we first get a basic rating using the user's and item's embeddings saved in NLFM after training. We then input this basic rating as a condition to RTGAN to generate the synthetic target review. Finally, the generated review is fed into NLFM to get the final rating score. Note that though we add the RTGAN module in order to generate and utilize the synthetic review, the rating prediction task in our AGTR model can be performed offline like MF methods. The procedure for the entire AGTR model is presented in Algorithm 2.

Table 2 Statistics of the datasets

Datasets	Users	Items	Ratings	Sparsity
Garden	1686	962	13,272	0.9918
Automotive	2928	1834	20,473	0.9962
Grocery	14,679	8711	151,254	0.9988
Yelp2017	29,406	39,643	1,239,518	0.9990

Algorithm 2 Procedure for the entire AGTR model**Input:** $X_{train} = (u, i, d_{ui}, r_{ui}, s_{ui}, d_u, d_i)_{i=1}^{n_{train}}$;**Output:** the generated target review $\hat{s}_{ui}$, the predicted rating $\hat{r}_{ui}$;

```

1: Initialize the parameters, including  $\Theta_{en}$  in the encoder,  $\Theta_G$  in  $G$ ,  $\Theta_D$  in  $D$ ,  $\Theta_{nl}$  in
   NLFM;
2: Pre-train the generator using MLE, update  $\Theta_{en}$  and  $\Theta_G$ ;
3: Generator  $G$  generates a target review  $\hat{s}_{ui}$ ;
4: Train Discriminator  $D$  using  $s_{ui}$  and  $\hat{s}_{ui}$  as positive and negative samples;
5: Update  $\Theta_D$ ;
6: repeat
7:   for each  $d_{ui}$  and  $r_{ui}$  in  $X_{train}$  do
8:     Optimize  $L_r$  for NLFM;
9:     Update  $\Theta_{en}$  and  $\Theta_{nl}$ ;
10:  for each  $d_u$  or  $d_i$  in  $X_{train}$  do
11:    Optimize  $L_{G_M}$  and  $L_{G_W}$  for Generator  $G$  in RTGAN;
12:    Generate a target review  $\hat{s}_{ui}$ ;
13:    Update  $\Theta_{en}$  and  $\Theta_G$ ;
14:  for each  $d_u$  or  $d_i$  in  $X_{train}$  do
15:    Optimize  $L_D$  for Discriminator  $D$  in RTGAN;
16:    Update  $\Theta_D$ ;
17: until Reach the maximum number of iteration;
18: Input the generated  $\hat{s}_{ui}$  into NLFM to get the final prediction  $\hat{r}_{ui}$ ;
19: return  $\hat{s}_{ui}$  and  $\hat{r}_{ui}$ 

```

In Algorithm 2, line 1 initializes the parameters and line 2 pre-trains the generator. Lines 6–17 train the AGTR module. Specifically, lines 3–5 train the NLFM module, and lines 10–16 train the RTGAN module. The generator G and the discriminator D is trained in lines 10–13 and lines 14–16, respectively. Line 18 inputs the generated $\hat{s}_{ui}$ into NLFM module to get the final prediction $\hat{r}_{ui}$. Finally, line 19 returns the generated target reviews and the predicted ratings.

Finally, we analyze the time complexity of our AGTR model. For the NLFM module, the time complexity for calculating the predicted rating is $O(Td)$, where T and d are the fixed review length and embedding size, respectively. For the RTGAN module, we encode the user's and item's historical reviews to generate target review. Hence, its time complexity is $O(CTd)$, where $C = \max\{\|d_u\|, \|d_i\|\}$, where $\|d_u\|$ and $\|d_i\|$ is the number of reviews of the user and the item, respectively. We train the AGTR model over $|\mathcal{R}|$ training instances, where $|\mathcal{R}|$ denotes the number of user-item pairs. Therefore, the overall time complexity for training is $O(|\mathcal{R}|CTd)$.

5 Experiments

5.1 Experimental Setup

Datasets We conduct experiments on two publicly accessible data sources: Amazon product review¹ and Yelp 2017.² We use three of product categories in Amazon: *Patio, Lawn and Garden, Automotive*, and *Grocery and Gourmet Food*. We take the 5-core version for experiments following the previous studies [4, 6, 35]. In this version, each user or item has at least 5 interactions. For all datasets, we extract the textual reviews as well as the numerical ratings to conduct experiments. The basic statistics of the datasets are shown in Table 2.

Evaluation metrics For rating prediction, we employ MAE [21] and MSE [25, 35, 37] as evaluation metrics. For review generation, we report the results in terms of negative log-likelihood (NLL) [10, 39] and ROUGE-1 [21, 37]. All these metrics are widely used in text generation and recommendation systems.

5.2 Evaluation Metrics

For rating prediction, we employ Mean Absolute Error (MAE) [21] and Mean Squared Error (MSE) [25, 35, 37]

¹ jmcauley.ucsd.edu/data/amazon/html.

² www.yelp.com/datasetchallenge/.

Table 3 Rating prediction performance in terms of MAE and MSE

	Garden		Automotive		Grocery		Yelp	
	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE
SentiRec	0.833*	1.067*	0.637*	0.824*	0.742*	1.014*	0.926*	1.371*
MPCN	0.852*	1.166*	0.576*	0.815*	0.821*	1.904*	0.902*	1.286*
A ³ NCF	0.793*	1.035*	0.696*	0.823*	0.777*	1.020*	0.846*	1.137*
ALFM	0.749*	0.984*	0.631*	0.772*	0.746*	1.001*	0.828*	1.096*
NARRE	0.772*	0.990*	0.621*	0.781*	0.743*	0.997*	0.819*	1.105*
TARMF	0.832*	1.103*	0.730*	0.868*	0.775*	1.073*	0.849*	1.196*
MT	0.848*	1.112*	0.747*	0.879*	0.769*	1.015*	0.852*	1.191*
MT-lg	0.799*	1.074*	0.701*	0.851*	0.762*	1.005*	0.855*	1.148*
NM	0.810 [−]	1.181 [−]	0.602 [−]	0.829 [−]	0.724*	1.020*	0.819*	1.116*
NRT	0.874*	1.109*	0.769*	0.814*	0.868*	1.174*	0.912*	1.127*
CAML	0.742	1.023*	0.625*	0.775*	0.704	0.979	0.815 [−]	1.089*
AGTR-r	0.750	0.972	0.602	0.767	0.737	0.994	0.821	1.091
AGTR	0.743	0.955	0.566	0.754	0.706	0.981	0.808	1.073

The best results are in bold and the second best ones (except those in our AGTR-r variant) are italicized – and *Significant difference according to paired t test between our model and each baseline for $p < 0.05$ and $p < 0.01$, respectively

as evaluation metrics. For review generation, we report the results in terms of negative log-likelihood (NLL) [10, 39], perplexity [25], and Recall-Oriented Understudy for Gisting Evaluation (ROUGE) [21, 37]. All these metrics are widely used in text generation and recommendation systems.

MAE and MSE are used to measure accuracy for continuous variables (in our case, ratings), they are defined as follows.

$$\text{MAE} = \frac{1}{m} \sum_{i=1}^m |(r_i - \hat{r}_i)| \quad (19)$$

$$\text{MSE} = \frac{1}{m} \sum_{i=1}^m (r_i - \hat{r}_i)^2, \quad (20)$$

where m is the number of interactions between user and item, r_i is the ground-truth rating, and $\hat{r}_i$ is the predicted rating. For MAE and MSE, smaller values indicate better prediction performance.

NLL is defined as follows:

$$\text{NLL} = - \sum_{i=1}^T p(y_i) * \log(q(x_i)), \quad (21)$$

where T is the length of the generated review, $p(y_i)$ is the distribution probability of word y_i in the real dataset, $q(x_i)$ is the generation probability of word x_i .

The perplexity metric is defined as the exponent of the average negative log-likelihood per word. Below is the formula for computing perplexity:

$$\text{Perplexity} = e^{\frac{-\sum \log(p(w))}{N}}, \quad (22)$$

where $p(w)$ is the probability of the word in the test set and N the total number of words. Lower perplexity implies higher log likelihood, and it indicates a better language model.

ROUGE is a classic metric in text summarization and machine translation. We adopt it to evaluate the quality of the generated reviews. There are three commonly used metrics including ROUGE-1, ROUGE-2, and ROUGE-L. The definitions for these metrics are as follows.

$$\text{ROUGE} - N(\hat{s}_{ui}) = \frac{\sum_{ng \in s_{ui}} \text{Count}_{\text{match}}(ng(s_{ui}, \hat{s}_{ui}))}{\sum_{ng \in s'} \text{Count}(ng(s'))}, \quad (23)$$

where ng is the n -gram, $\text{Count}(ng(s'))$ is the number of n -grams in s' (either $\hat{s}_{ui}$ or s_{ui}). For ROUGE-1 and ROUGE-2, $\text{Count}_{\text{match}}(ng(s_{ui}, \hat{s}_{ui}))$ is the number of uni-grams and bi-grams co-occurring in s_{ui} and $\hat{s}_{ui}$. For ROUGE-L, $\text{Count}_{\text{match}}(ng(s_{ui}, \hat{s}_{ui}))$ denotes the longest common subsequence in s_{ui} and $\hat{s}_{ui}$. We report the scores in terms of ROUGE-1, ROUGE-2, and ROUGE-L to evaluate the quality of the generated reviews in different granularities.

Table 4 Review generation performance in terms of NLL

	Garden	Automotive	Grocery	Yelp
MT	5.74	4.01	4.28	5.19
MT-lg	5.61	3.96	4.30	5.14
NRT	6.24	4.34	4.57	5.43
NM	5.63	4.06	4.81	5.84
CAML	5.45	3.28	3.51	4.84
AGTR-r	5.68	3.99	4.34	5.08
AGTR	5.51	3.86	3.25	4.99

Table 5 Review generation performance in terms of Perplexity

	Garden	Automotive	Grocery	Yelp
MT	1916.2	1435.5	1625.5	2556.2
MT-lg	874.4	669.6	1243.3	1659.8
NRT	11439.3	13220.5	620.8	1032.7
NM	4256.8	9357.9	8234.8	2785.4
CAML	3291.4	2417.9	2683.6	3253.7
AGTR-r	805.3	654.7	677.1	1143.0
AGTR	775.8	633.3	420.8	893.5

5.3 Compared Methods

We compare our AGTR model with the following state-of-the-art methods.

SentiRec [14] first encodes each review into a fixed-size vector using CNN and then generates recommendations using vector-encoded reviews.

MPCN [35] exploits review-level co-attention mechanism to determine the most informative reviews and gets the representations of users and items.

A³NCF [6] designs a new topic model to extract user preferences and item characteristics from review texts and then feeds them into a neural network for rating prediction.

ALFM [7] develops an aspect-aware latent factor model where a new topic model is integrated to model user preferences and item features from different aspects.

NARRE [4] processes each review using CNN and adopts attention mechanism to build the recommendation model and select useful reviews simultaneously.

TARMF [24] adopts attention-based RNN to extract textual features and maximizes the similarity between latent factors and textual features.

MT [25] jointly learns to perform rating prediction and recommendation explanation by combining MF for rating prediction and SeqGan [39] for review generation.

NRT [21] uses MF and generation networks to combine ratings, reviews, and tips for rating prediction and abstractive tips generation.

NM [37] uses a single neural network to model users and products, and generates customized product representations using a deep memory network, from which customized ratings and reviews are constructed jointly.

CAML [5] uses an encoder–selector–decoder architecture to model the cross-knowledge transferred for both the recommendation task and the explanation task using a multi-task framework.

In addition to the above baselines, we propose two variants for MT and our AGTR models. Specifically, **MT-lg** replaces SeqGan [39] in the review generation module of MT [25] with LeakGan [10] in our model to exclude the potential influence caused by using different generation models. **AGTR-r** removes the rating condition from the generation module in our AGTR model to investigate the effects of our rating tailored GAN.

We do not compare our model with other methods like DeepCoNN [42] and TransNet [3] using reviews for rating prediction, neither with the traditional methods like NMF [19], FM [31], and NeuMF [12] which do not use reviews. These methods have been shown to be weaker than the baselines [7, 24, 35] used in our experiments; thus, we only show improvements over the baselines.

5.4 Parameter Settings

Each dataset is divided into 80%/10%/10% splits for training, validation, and testing, respectively. We train the model on the training set and tune the hyper-parameters on the validation set. The ground-truth reviews in the training set are used for training the model. Note that those in validation or testing sets are never accessed. Instead, only the generated target reviews are used for validation or testing.

The parameters of all baselines are the same as those in the corresponding original papers. For our AGTR model, we set dimensionality to 32 for all embeddings of users, items, and word latent factors. In review generation, the maximum review length T is set to 40 words, and other parameters such as the kernel size of CNN are the same as those in LeakGAN. We use Adam [16] for optimization. We set learning rate=0.002, minibatch size=64, and dropout ratio=0.5 for all the datasets.

Table 6 Review generation performance in terms of ROUGE

	Garden			Automotive			Grocery			Yelp		
	<i>R</i> − 1	<i>R</i> − 2	<i>R</i> − <i>L</i>	<i>R</i> − 1	<i>R</i> − 2	<i>R</i> − <i>L</i>	<i>R</i> − 1	<i>R</i> − 2	<i>R</i> − <i>L</i>	<i>R</i> − 1	<i>R</i> − 2	<i>R</i> − <i>L</i>
MT	3.22	0.03	2.11	2.95	0.04	2.89	5.20	0.06	3.68	5.22	0.05	3.64
MT-lg	3.25	0.03	2.10	2.94	0.04	3.05	5.01	0.06	3.71	5.31	0.06	3.72
NRT	0.52	0.00	0.49	1.72	0.00	1.56	7.51	0.14	3.25	6.21	0.07	3.78
NM	3.33	0.02	2.63	2.98	0.01	2.46	4.40	0.04	3.39	6.25	0.07	3.75
CAML	3.33	0.06	2.43	4.96	0.10	3.48	7.57	0.31	4.76	7.38	0.15	4.49
AGTR-r	3.42	0.04	2.16	2.74	0.06	2.80	4.54	0.06	3.05	5.74	0.07	3.55
AGTR	3.49	0.06	2.43	3.01	0.08	3.07	7.73	0.17	3.84	6.43	0.09	3.88

R-* refers to ROUGE-*

Table 7 Examples of the predicted ratings and the generated reviews (Ref. denotes the ground-truth review and rating)

	Rating	Review
Ref.	5.0	Last very long time stainless steel very good quality i not buy another sure use alot
MT	4.25	Good want use like another days earth hubby metal very activity...
MT-lg	4.47	Think nice product use buy but still want again skin cool cold rarely does like...
NRT	4.17	Shaped nice seldom introduced so sneak transplanting still momentum ...
NM	4.33	Activity absolutely very well down won cool quality skin sheath ...
CAML	4.84	Bought happy grill test propane ignition roast mind what built ...
AGTR-r	4.68	Good product use still some lot not very operate only so middle ...
AGTR	4.85	Worked very well very easy use still from some quality rain not sure good value few days ...

5.5 Rating Prediction

The results of all methods for rating prediction are presented in Table 3. (1) The upper six rows from SentiRec to TARMF are *the first type* of review-aware rating prediction methods which do not generate target reviews. (2) The middle five rows from MT to CAML are *the second type* which generates target reviews/tips. (3) The last two rows are our AGTR model and its variant. From Table 3, we have the following important observations.

Firstly, our AGTR model statistically significantly outperforms all baselines in terms of MAE and MSE metrics on three of the four datasets. The baselines' performances fluctuate among different datasets. MPCN, ALFM, and CAML once becomes the second best in some cases. This shows that it is hard to get the consistently better performance for one method due to the characteristics of the different datasets. In contrast, our model achieves the best performance on Garden, Automotive, and Yelp datasets. CAML is the best on Grocery. However, the difference between our model and CAML on this dataset is not significant. All these results clearly demonstrate the effectiveness of our model.

Secondly, among six methods in the first type, ALFM and NARRE are generally better than other methods. Both these methods differentiate the importance of each review or each aspect. This infers that a fine-grained analysis on the reviews has great impacts on the related rating prediction task. Among five methods in the second type, CAML benefits a lot from the joint training of two tasks under the multi-task framework. Moreover, NM performs better than MT and NRT which only generate but do not integrate target reviews for rating prediction. Both these clearly show the predictive ability of target reviews. Our AGTR model's superior performance over NM can be due to our carefully designed NLFM module, which makes the best use of the target review. The other reason is that the quality of our generated reviews is higher than that of NM with the help of rating tailored adversarial learning.

Thirdly, MT-lg is better than the original MT, suggesting the importance of generative model. On the other hand, GRT-r performs worse than AGTR, showing that rating condition plays a critical role in generating reviews consistent with rating scores. However, the enhanced MT-lg is still worse than our simplified version GRT-r. This indicates that our NLFM module performs much better the matrix

Fig. 2 Performance of different size of latent factor and max length of review

factorization model in MT. NRT is designed for abstractive tips generation, which results in its inferior performance.

5.6 Review Generation

This section evaluates the performance of our AGTR model on review generation by comparing it with the second-type baselines. The results in terms of NLL, Perplexity, and ROUGE are presented in Tables 4, 5, and 6, respectively. Remember that for NLL and Perplexity, smaller scores indicate better performance. In contrast, larger scores indicate better performance for ROUGE. The best results are in bold, and the second best ones (except those in our AGTR-r variant) are italicized.

From Tables 4, 5, and 6, it is clear that our AGTR model can generate the best or the second best reviews in terms of NLL, Perplexity, and ROUGE metrics on all datasets. Moreover, AGTR-r's results are not as good as AGTR. This, once again, demonstrates that our strategy of taking rating as the condition in GAN helps generate high-quality reviews.

Among the baselines, CAML is the second best in terms of NLL and ROUGE metrics in most cases. This infers that CAML can generate good reviews with the help of supervision from the rating subtask under the multi-task learning

framework. However, when the evaluation metric is Perplexity, the second best ones are MT-Ig and NRT. We can conclude that none of these methods always performs the best in terms of all metrics on all datasets. This is due to that the metrics focus on different factors. Also note that the perplexity values are quite large. This is probably caused by the removal of stop words from reviews, which results in the lack of continuity between two words. Finally, we find that both NRT and NM perform relatively poorly. The reason might be that they only adopt the maximum likelihood estimation to generate reviews without exploiting the adversarial network. On the other hand, MT-Ig is better than MT, indicating that LeakGAN performs better than SeqGAN.

5.7 Case Study

In order to capture more details, we provide several examples in Table 7 to analyze the relevance between the generated synthetic reviews/ratings and the real ones.

As can be seen, our AGTR model gets the highest rating score, i.e., 4.85, which is very close to the real score. Furthermore, our generated review is suitable to express the strong positive sentiment reflected by the full credit, and it is most similar to the real review. We also need to point out that

the words in the latter half of our generated review are not very accurate. This also happens to other generated reviews. The reason is that the sentences in the training data have varying lengths. However, when training the model, all sentence should have a fixed length. Hence long sentences will be truncated, and short sentences need add some random words to reach the fixed length after the original words. That is to say, the latter part of some training sentences are not semantically related. Consequently, the network is unable to generate accurate words for the latter part of the sentence.

5.8 Parameter Analysis

In this subsection, we investigate the effects of two parameters, i.e., the number of latent factors and the max length of reviews. We first examine the effect of the latent factor size in Fig. 2a, b. We can see that, with the increase number of latent factors, the performance could be enhanced since more latent factors bring better representation capability. For example, the model achieves the best performance when the latent factor number is 32 on almost all datasets. However, too many latent factors, e.g., 64, may cause over-fitting and result in the decrease of performance.

We then study the effects of the max length of reviews in Fig. 2c, d. When the review length is small, the part of texts that exceeds the specified length need to be truncated when preprocessing, which will result in a information loss. In this case, the smaller the specified length, the more information is missing, and thus, the performance will decrease. When the review length increases, the reviews which is shorter than the threshold need to be padded. The irrelevant words padded would bring noises to the model, which will harm the performance of the model. Hence a reasonable max length of review is about 40 words.

6 Conclusion

In this paper, we presented a novel AGTR model to leverage the predictive ability of target reviews. We developed a unified framework to generate target reviews using a rating tailored GAN and to do rating prediction with a neural latent factor model which well exploits the generated target review besides historical reviews. We conducted extensive experiments on four real-world datasets. Results demonstrate our model achieves the state-of-the-art performance in both rating prediction and review generation tasks.

As for future work, one possible direction is to generate target reviews with variable length. The second is to enhance the interaction between two modules under the multi-task framework. The third is to develop new approach instead of extending LeakGAN for review

generation, which might be explored as a separate problem rather than a component in our rating prediction task.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Bao Y, Fang H, Zhang, J (2014) Topicmf: simultaneously exploiting ratings and reviews for recommendation. In: AAAI, pp 2–8
2. Bauman K, Liu B, Tuzhilin A (2017) Aspect based recommendations: recommending items with the most valuable aspects based on user reviews. In: KDD, pp 717–725
3. Catherine R, Cohen W (2017) Transnets: learning to transform for recommendation. In: RecSys, pp 288–296
4. Chen C, Zhang M, Liu Y, Ma, S (2018) Neural attentional rating regression with review-level explanations. In: WWW, pp 1583–1592
5. Chen Z, Xiting W, Xie X, Wu T, Bu G, Wang Y, Chen E (2019) Co-attentive multi-task learning for explainable recommendation. In: IJCAI
6. Cheng Z, Ding Y, He X, Zhu L, Song X, Kankanhalli M (2018) A³ncf: an adaptive aspect attention model for rating prediction. In: IJCAI, pp 3748–3754
7. Cheng Z, Ding Y, Zhu L, Kankanhalli M (2018) Aspect-aware latent factor model: rating prediction with ratings and reviews. In: WWW, pp 639–648
8. Diao Q, Qiu M, Wu CY, Smola AJ, Jiang J, Wang C (2014) Jointly modeling aspects, ratings and sentiments for movie recommendation (jmars). In: KDD, pp 193–202
9. Goodfellow I, Pouget-Abadie J, Mirza M, Xu B, D Warde-Farley, Ozair S, Courville A, Bengio Y (2014) Generative adversarial nets. In: NIPS, pp 2672–2680
10. Guo J, Lu S, Cai H, Zhang W, Yu Y, Wang J (2018) Long text generation via adversarial training with leaked information. In: AAAI
11. He X, Chen T, Kan MY, Chen X (2015) Trirank: review-aware explainable recommendation by modeling aspects. In: CIKM, pp 1661–1670
12. He X, Liao L, Zhang H, Nie L, Hu X, Chua TS (2017) Neural collaborative filtering. In: WWW, pp 173–182
13. Hu L, Sun A, Liu Y (2014) Your neighbors affect your ratings: on geographical neighborhood influence to rating prediction. In: SIGIR, pp 345–354
14. Hyun D, Park C, Yang MC, Song I, Lee JT, Yu H (2018) Review sentiment-guided scalable deep recommender system. In: SIGIR, pp 965–968
15. Im DJ, Ahn S, Memisevic R, Bengio Y (2017) Denoising criterion for variational auto-encoding framework. In: AAAI, pp 2059–2065

16. Kingma D, Ba J (2014) Adam: a method for stochastic optimization. ICLR (Poster)
17. Kingma DP, Mohamed S, Rezende DJ, Welling M (2014) Semi-supervised learning with deep generative models. In: NIPS, pp 3581–3589
18. Kingma DP, Welling M (2014) Auto-encoding variational bayes. In: Proceedings of the 2nd international conference on learning representations
19. Lee DD, Seung HS (2001) Algorithms for non-negative matrix factorization. In: NIPS, pp 556–562
20. Li J, Monroe W, Shi T, Jean S, Ritter A, Jurafsky D (2017) Adversarial learning for neural dialogue generation. In: EMNLP, pp 2151–2156
21. Li P, Wang Z, Ren Z, Bing L, Lam W (2017) Neural rating regression with abstractive tips generation for recommendation. In: SIGIR, pp 345–354
22. Lin K, Li D, He X, Zhang Z, Ting Sun M (2017) Adversarial ranking for language generation. In: NIPS
23. Ling G, Lyu MR, King I (2014) Ratings meet reviews, a combined approach to recommend. In: RecSys, pp 105–112
24. Lu Y, Dong R, Smyth B (2018) Coevolutionary recommendation model: mutual learning between ratings and reviews. In: WWW, pp 773–782
25. Lu Y, Dong R, Smyth B (2018) Why i like it: multi-task learning for recommendation and explanation. In: RecSys, pp 4–12 (2018)
26. McAuley J, Leskovec J (2013) Hidden factors and hidden topics: understanding rating dimensions with review text. In: RecSys, pp 165–172
27. Mirza M, Osindero S (2014) Conditional generative adversarial nets. CoRR, [arXiv:1411.1784](https://arxiv.org/abs/1411.1784)
28. Pappas N, Popescu-Belis A (2013) Sentiment analysis of user comments for one-class collaborative filtering over ted talks. In: SIGIR, pp 773–776
29. Š Pero, Horáth T (2013) Opinion-driven matrix factorization for rating prediction. In: UMAP, pp 1–13
30. Ratliff LJ, Burden SA, Sastry SS (2013) Characterization and computation of local nash equilibria in continuous games. In: IEEE TNN, pp 917–924
31. Rendle S (2010) Factorization machines. In: ICDM, pp 995–1000
32. Salimans T, Goodfellow I, Zaremba W, Cheung V, Radford A, Chen X (2016) Improved techniques for training gans. In: NIPS, pp 2234–2242
33. Sutskever I, Vinyals O, Le QV (2014) Sequence to sequence learning with neural networks. In: NIPS, pp 3104–3112
34. Tan Y, Zhang M, Liu Y, Ma S (2016) Rating-boosted latent topics: understanding users and items with ratings and reviews. In: IJCAI, pp 2640–2646
35. Tay Y, Luu AT, Hui SC (2018) Multi-pointer co-attention networks for recommendation. In: KDD, pp 2309–2318
36. Wang C, Blei DM (2011) Collaborative topic modeling for recommending scientific articles. In: KDD, pp 448–456
37. Wang Z, Zhang Y (2017) Opinion recommendation using a neural model. In: EMNLP, pp 1626–1637
38. Xu Y, Lam W, Lin T (2014) Collaborative filtering incorporating review text and co-clusters of hidden user communities and item groups. In: CIKM, pp 1661–1670
39. Yu LT, Zhang WN, Wang J, Yu Y (2016) Seqgan: sequence generative adversarial nets with policy gradient. In: AAAI
40. Zhang W, Wang J (2016) Integrating topic and latent factors for scalable personalized review-based rating prediction. TKDE 28(11):3013–3027
41. Zhang Y, Lai G, Zhang M, Zhang Y, Liu Y, Ma S (2014) Explicit factor models for explainable recommendation based on phrase-level sentiment analysis. In: SIGIR, pp 83–92
42. Zheng L, Noroozi V, Yu PS (2017) Joint deep modeling of users and items using reviews for recommendation. In: WSDM, pp 425–434