

Positive & Negative Critiquing for VAE-based Recommenders

Diego Antognini

École Polytechnique Fédérale de Lausanne
Lausanne, Switzerland
diego.antognini@epfl.ch

Boi Faltings

École Polytechnique Fédérale de Lausanne
Lausanne, Switzerland
boi.faltings@epfl.ch

ABSTRACT

Providing explanations for recommended items allows users to refine the recommendations by *critiquing* parts of the explanations. As a result of revisiting critiquing from the perspective of multimodal generative models, recent work has proposed M&Ms-VAE, which achieves state-of-the-art performance in terms of recommendation, explanation, and critiquing. M&Ms-VAE and similar models allow users to negatively critique (i.e., explicitly disagree). However, they share a significant drawback: users cannot positively critique (i.e., highlight a desired feature). We address this deficiency with M&Ms-VAE⁺, an extension of M&Ms-VAE that enables positive and negative critiquing. In addition to modeling users’ interactions and keyphrase-usage preferences, we model their keyphrase-usage dislikes. Moreover, we design a novel critiquing module that is trained in a self-supervised fashion. Our experiments on two datasets show that M&Ms-VAE⁺ matches or exceeds M&Ms-VAE in recommendation and explanation performance. Furthermore, our results demonstrate that representing positive and negative critiques differently enables M&Ms-VAE⁺ to significantly outperform M&Ms-VAE and other models in positive *and* negative multi-step critiquing.

1 INTRODUCTION

Critiquing is a conversational recommendation method that incrementally adapts recommendations in response to user preferences [4]. Several studies have revisited critiquing with neural models [3, 9–11, 15]. They allow users to critique the recommendation by interacting with a set of attributes mined from user reviews. Recently, [2] proposed a multimodal variational-autoencoder, called M&Ms-VAE, that achieves state-of-the-art performance in terms of recommendation, explanation, and multi-step critiquing. Although M&Ms-VAE has made important contributions, it shares an important deficiency with other models: users can only express an explicit disagreement (i.e., negative critiquing); there is no means of highlighting a desired feature (i.e., positive critiquing) [12, 16].

One way to enable positive critiquing is to embed the positive critique identically to the negative one, with the advantage that the recommender remains the same. As we demonstrate empirically in Section 4.4, it underperforms in positive multi-step critiquing compared to negative critiquing. Another means of enabling positive critiquing consists of modeling the positive and negative critiques within the recommender. Although this approach improves critiquing performance, it affects recommendation and explanation performance, in a way that limits their practice performance.

This paper is published under the Creative Commons Attribution-NonCommercial-NoDerivs 4.0 International (CC-BY-NC-ND 4.0) license. Authors reserve their rights to disseminate the work on their personal web sites with the appropriate attribution.
arXiv, 2022, USA

In this study, we present M&Ms-VAE⁺, an extension of M&Ms-VAE, to enable positive critiquing. We take advantage of the multimodal modeling to introduce a third modality that represents users’ keyphrase-usage dislikes along with their interactions and keyphrase-usage preferences. Finally, we design a method to critique positively and negatively, which trained in a self-supervision manner.

We evaluate our method using two real-world, publicly available datasets. The results show that M&Ms-VAE⁺ (1) matches or exceeds M&Ms-VAE in recommendation and explanation performance and (2) outperforms M&Ms-VAE and other strong models by a large margin on positive and negative multi-step critiquing.

2 MIXTURE-OF-EXPERTS MULTIMODAL VAE

Notation. Before proceeding, we define the following notation:

- $U, I,$ and K : The user, item, and keyphrase sets, respectively.
- $\mathbf{R} \in \mathbb{R}^{|U| \times |I|}$: The binary user-by-item interaction matrix.
- $\mathbf{K}^+, \mathbf{K}^- \in \mathbb{R}^{|U| \times |K|}$: User-keyphrase matrices that reflect users’ keyphrase-usage preferences and dislikes in the reviews. \mathbf{k}_u^+ are the keywords mentioned by the user u , and \mathbf{k}_u^- are all the others. Similarly, \mathbf{K}^I is the item-keyphrase matrix.
- $\mathbf{c}_u^{+t}, \mathbf{c}_u^{-t} \in \mathbb{R}^{|K|}$: A one-hot vector whose only positive value indicates the index of the keyphrase to be critiqued positively and negatively by the user u at a given step t .
- $\mathbf{z}_u, \mathbf{z}_{c_u^{+t}}, \mathbf{z}_{c_u^{-t}} \in \mathbb{R}^{|H|}$: The latent representation of the user u , the positive critique \mathbf{c}_u^{+t} , and the negative critique \mathbf{c}_u^{-t} .

Model Overview. An important feature of M&Ms-VAE [2] is learning the joint distribution $p(\mathbf{r}_u, \mathbf{k}_u^+)$ under partial observations: we aim to recommend and generate keyphrase explanation jointly and independently from each observed variable (i.e., modality). Thanks to the multimodal assumption, \mathbf{r}_u and \mathbf{k}_u^+ are conditionally independent given the common latent variable \mathbf{z}_u ; an unobserved variable can be safely ignored when evaluating the marginal likelihood. As Fig. 1 (left) shows, we write the joint log-likelihood as follows:

$$\log p(\mathbf{r}_u, \mathbf{k}_u^+) \geq \mathbb{E}_{q_\Phi(\mathbf{z}_u | \mathbf{r}_u, \mathbf{k}_u^+)} [\log p_{\Theta_r}(\mathbf{r}_u | \mathbf{z}_u) + \log p_{\Theta_{k^+}}(\mathbf{k}_u^+ | \mathbf{z}_u)] - \beta \text{D}_{\text{KL}}[q_\Phi(\mathbf{z}_u | \mathbf{r}_u, \mathbf{k}_u^+) || p(\mathbf{z}_u)], \quad (1)$$

where the prior $p(\mathbf{z}_u)$ is a normal distribution and β is a hyperparameter that controls the strength of the regularization relative.

Currently, learning $q_\Phi(\mathbf{z}_u | \mathbf{r}_u, \mathbf{k}_u^+)$ requires \mathbf{r}_u and \mathbf{k}_u^+ to be fully observed. [2] remedied this problem by using a mixture of experts (MoE): $q_\Phi(\mathbf{z}_u | \mathbf{r}_u, \mathbf{k}_u^+) = \zeta(q_{\Phi_r}(\mathbf{z}_u | \mathbf{r}_u), q_{\Phi_{k^+}}(\mathbf{z}_u | \mathbf{k}_u^+))$. The MoE $\zeta(\cdot)$ takes the average of the experts when \mathbf{r}_u and \mathbf{k}_u^+ are both observed or the unimodal expert when only one modality is observed.

Finally, the proposed training scheme mimics weakly supervised learning to train the individual inference networks q_{Φ_r} and $q_{\Phi_{k^+}}$:

$$\mathcal{L}_{M\&Ms}(\mathbf{R}, \mathbf{K}^+) = \sum_u \text{ELBO}(\mathbf{r}_u, \mathbf{k}_u^+) + \text{ELBO}(\mathbf{r}_u) + \text{ELBO}(\mathbf{k}_u^+) \quad (2)$$

Algorithm 1 Creation of Synthetic Critiquing Datasets

```

1: function CREATE_CRITIQUING_DATASET( $K, \mathbf{R}^{\text{val}}, \mathbf{K}^{\text{I}}, \hat{\mathbf{K}}^{\text{U}}$ )
2:   Synthetic datasets  $D^+ \leftarrow \{\}, D^- \leftarrow \{\}$ 
3:   for each user  $u$  and target item  $i$ , where  $\mathbf{r}_{u,i}^{\text{val}} = \mathbf{1}$  do
4:     Randomly sample a critique  $\mathbf{c}^- \in K \setminus \mathbf{k}_i^{\text{I}}$  // Neg. critique
5:     Compute the (un-)affected item sets  $I_{\mathbf{c}^-}^+$  and  $I_{\mathbf{c}^-}^-$  by  $\mathbf{c}^-$ 
6:     Update  $D^- \leftarrow D^- \cup \{(u, i, \mathbf{c}^-, I_{\mathbf{c}^-}^+, I_{\mathbf{c}^-}^-)\}$ 
7:     Randomly sample a critique  $\mathbf{c}^+ \in \mathbf{k}_i^{\text{I}} \setminus \hat{\mathbf{k}}_i^{\text{U}}$  // Pos. critique
8:     Compute the (un-)affected item sets  $I_{\mathbf{c}^+}^+$  and  $I_{\mathbf{c}^+}^-$  by  $\mathbf{c}^+$ 
9:     Update  $D^+ \leftarrow D^+ \cup \{(u, i, \mathbf{c}^+, I_{\mathbf{c}^+}^+, I_{\mathbf{c}^+}^-)\}$ 
10:  return Synthetic datasets  $D^-$  and  $D^+$ 

```

2.1 Negative Keyphrase-based Critiquing

Given the predicted explanation $\hat{\mathbf{k}}_u^+$ and the recommendation $\hat{\mathbf{r}}_u$, the user u can accept or refine the recommendation. In negative critiquing, the user iteratively provides a keyphrase to critique \mathbf{c}_u^{-t} (i.e., disagreement) and obtains a new recommendation $\hat{\mathbf{r}}_u^t$ until he is satisfied. The critique representation $\mathbf{z}_{\mathbf{c}_u^{-t}}$ is encoded via $q_{\Phi_{k^+}}(\mathbf{z}_u | \mathbf{k}_u^+)$.

To blend the user representation \mathbf{z}_u with the t^{th} critique representation $\mathbf{z}_{\mathbf{c}_u^{-t}}$, [2] introduced a blending function that treats each critique as independent and uses gated recurrent units [5]: $\xi(\mathbf{z}_u, \mathbf{z}_{\mathbf{c}_u^{-t}}) = \text{GRU}([\mathbf{z}_u; \mathbf{z}_{\mathbf{c}_u^{-t}}])$, where $;$ denotes the concatenation operation. The blending module $\xi(\cdot)$ is optimized in a self-supervised fashion. First, a synthetic dataset D^- is created following Alg. 1 (Lines 1-6 and Line 10). For each negative critique \mathbf{c}^- inconsistent with the target item, we compute the item set $I_{\mathbf{c}^-}^+$ (symmetrically $I_{\mathbf{c}^-}^-$), which contains (symmetrically does not contain) the critique. Second, a max-margin ranking-based objective encourages $\xi(\cdot)$ to learn how to re-rank the items according to the critique \mathbf{c}^- :

$$\begin{aligned} \mathcal{L}_{M\&Ms}^{\text{neg}}(\hat{\mathbf{R}}^0, \hat{\mathbf{R}}^1, u, \mathbf{c}^-, I_{\mathbf{c}^-}^+, I_{\mathbf{c}^-}^-) &= \sum_{i^+ \in I_{\mathbf{c}^-}^+} \max\{0, h - (\hat{\mathbf{r}}_{u,i^+}^0 - \hat{\mathbf{r}}_{u,i^+}^1)\} \\ &+ \sum_{i^- \in I_{\mathbf{c}^-}^-} \max\{0, h - (\hat{\mathbf{r}}_{u,i^-}^1 - \hat{\mathbf{r}}_{u,i^-}^0)\}. \end{aligned} \quad (3)$$

3 ENABLING POSITIVE CRITIQUING

A drawback of M&Ms-VAE and other models ([9–11, 15]) is that users can express a disagreement but cannot highlight a feature. To enable positive critiquing, we propose three incremental solutions.

① *Extending the blending module.* A suboptimal way to enable positive critiquing is to adjust *only* the blending module and the self-supervised task. The advantage of this approach is that there is no need to retrain the generative and inference networks: the recommendation and explanation performance remain the same. The representation $\mathbf{z}_{\mathbf{c}_u^{+t}}$ of the positive critique \mathbf{c}_u^{+t} is computed identically to $\mathbf{z}_{\mathbf{c}_u^{-t}}$, using the inference model $q_{\Phi_{k^+}}(\mathbf{z}_u | \mathbf{k}_u^+)$. Then, we decompose the blending module into $\xi^+(\mathbf{z}_u, \mathbf{z}_{\mathbf{c}_u^{+t}}) = \text{GRU}([\mathbf{e}^+; \mathbf{z}_u; \mathbf{z}_{\mathbf{c}_u^{+t}}])$ and $\xi^-(\mathbf{z}_u, \mathbf{z}_{\mathbf{c}_u^{-t}}) = \text{GRU}([\mathbf{e}^-; \mathbf{z}_u; \mathbf{z}_{\mathbf{c}_u^{-t}}])$. $\xi^+(\cdot)$ and $\xi^-(\cdot)$ share the same weights. Both \mathbf{e}^+ and \mathbf{e}^- are trainable parameters that condition the gated mechanism to compute a representation that reflects an explicit agreement with \mathbf{c}_u^{+t} and a disagreement with \mathbf{c}_u^{-t} .

To include positive critiquing examples, we create a second synthetic dataset D^+ (see Alg. 1). Echoing the approach used in

Section 2.1, we sample a critique \mathbf{c}^+ that is not part of the predicted keyphrase explanation $\hat{\mathbf{k}}_u^+$ but is consistent with the target item’s features and compute the item sets $I_{\mathbf{c}^+}^+$ and $I_{\mathbf{c}^+}^-$.

$$\begin{aligned} \mathcal{L}_{M\&Ms}^{\text{pos,neg}}(\cdot) &= \mathcal{L}_{M\&Ms}^{\text{neg}}(\cdot) + \sum_{i^+ \in I_{\mathbf{c}^+}^+} \max\{0, h - (\hat{\mathbf{r}}_{u,i^+}^1 - \hat{\mathbf{r}}_{u,i^+}^0)\} \\ &+ \sum_{i^- \in I_{\mathbf{c}^+}^-} \max\{0, h - (\hat{\mathbf{r}}_{u,i^-}^0 - \hat{\mathbf{r}}_{u,i^-}^1)\}. \end{aligned} \quad (4)$$

② *Introducing a third modality.* Above, we used the *same* model $q_{\Phi_{k^+}}(\cdot)$ to infer $\mathbf{z}_{\mathbf{c}_u^{+t}}$ and $\mathbf{z}_{\mathbf{c}_u^{-t}}$. However, neither \mathbf{c}^+ and \mathbf{c}^- embed the same meaning, whereas $q_{\Phi_{k^+}}(\cdot)$ considers at training only users’ keyphrase-usage preferences. Therefore, we introduce a new modality \mathbf{k}_u^- that represents users’ keyphrase-usage dislikes: the keyphrases that are not part of the user’s profile. As Fig. 1 (M&Ms-VAE³) shows, we update Eq. 1 as follows:

$$\begin{aligned} \log p(\mathbf{r}_u, \mathbf{k}_u^+, \mathbf{k}_u^-) &\geq \mathbb{E}_{q_{\Phi}(\mathbf{z}_u | \mathbf{r}_u, \mathbf{k}_u^+, \mathbf{k}_u^-)} [\log p_{\Theta_r}(\mathbf{r}_u | \mathbf{z}_u) + \log p_{\Theta_{k^+}}(\mathbf{k}_u^+ | \mathbf{z}_u) \\ &+ \log p_{\Theta_{k^-}}(\mathbf{k}_u^- | \mathbf{z}_u)] - \beta \text{D}_{\text{KL}}[q_{\Phi}(\mathbf{z}_u | \mathbf{r}_u, \mathbf{k}_u^+, \mathbf{k}_u^-) || p(\mathbf{z}_u)]. \end{aligned} \quad (5)$$

We augment the mixture of experts with the extra unimodal posterior: $q_{\Phi}(\mathbf{z}_u | \mathbf{r}_u, \mathbf{k}_u^+, \mathbf{k}_u^-) = \zeta(q_{\Phi_r}(\mathbf{z}_u | \mathbf{r}_u), q_{\Phi_{k^+}}(\mathbf{z}_u | \mathbf{k}_u^+), q_{\Phi_{k^-}}(\mathbf{z}_u | \mathbf{k}_u^-))$. We also update the training strategy as $\mathcal{L}_{M\&Ms^3}(\mathbf{R}, \mathbf{K}^+, \mathbf{K}^-)$:

$$\sum_u \text{ELBO}(\mathbf{r}_u, \mathbf{k}_u^+, \mathbf{k}_u^-) + \text{ELBO}(\mathbf{r}_u) + \text{ELBO}(\mathbf{k}_u^+) + \text{ELBO}(\mathbf{k}_u^-). \quad (6)$$

Finally, we compute the critique representations $\mathbf{z}_{\mathbf{c}_u^{+t}}$ and $\mathbf{z}_{\mathbf{c}_u^{-t}}$ using $q_{\Phi_{k^+}}(\mathbf{z}_u | \mathbf{k}_u^+)$ and $q_{\Phi_{k^-}}(\mathbf{z}_u | \mathbf{k}_u^-)$. However, the joint likelihood now includes the generative model $p_{\Theta_{k^-}}(\mathbf{k}_u^- | \mathbf{z}_u)$, which leads to two problems: (1) it is unclear what value $\hat{\mathbf{k}}_u^-$ offers the user in addition to the explanation $\hat{\mathbf{k}}_u^+$ and the recommendation $\hat{\mathbf{r}}_u$, and (2) $\hat{\mathbf{k}}_u^-$ is redundant, because it is equivalent to $K \setminus \hat{\mathbf{k}}_u^+$. It might affect the latent space and the recommendation and explanation performance.

③ *Removing $p_{\Theta_{k^-}}(\mathbf{k}_u^- | \mathbf{z}_u)$.* Because we aim to develop a generative model of the form $p_{\Theta}(\mathbf{r}_u, \mathbf{k}_u^+, \mathbf{z}_u)$, and thanks to the multimodal factorization, we can rewrite Eq. 5 by safely removing $p_{\Theta_{k^-}}(\mathbf{k}_u^- | \mathbf{z}_u)$:

$$\begin{aligned} \log p(\mathbf{r}_u, \mathbf{k}_u^+) &\geq \mathbb{E}_{q_{\Phi}(\mathbf{z}_u | \mathbf{r}_u, \mathbf{k}_u^+, \mathbf{k}_u^-)} [\log p_{\Theta}(\mathbf{r}_u, \mathbf{z}_u) + \log p_{\Theta_{k^+}}(\mathbf{k}_u^+ | \mathbf{z}_u)] \\ &- \beta \text{D}_{\text{KL}}[q_{\Phi}(\mathbf{z}_u | \mathbf{r}_u, \mathbf{k}_u^+, \mathbf{k}_u^-) || p(\mathbf{z}_u)]. \end{aligned} \quad (7)$$

To summarize, we model both keyphrase-usage preferences and dislikes of users, while predicting only the keyphrase explanation alongside the recommended items, as Fig. 1 (M&Ms-VAE⁺) shows. The inference networks are identical to those of M&Ms-VAE³ and similar to the generative networks of M&Ms-VAE. A crucial question remains: how do we train $q_{\Phi_{k^-}}(\mathbf{z}_u | \mathbf{k}_u^-)$? Once again, we take advantage of the multimodal modeling and incorporate the partial-paired observations $(\mathbf{r}_u, \mathbf{k}_u^-)$ and $(\mathbf{k}_u^+, \mathbf{k}_u^-)$ into the training scheme:

$$\begin{aligned} \mathcal{L}_{M\&Ms^+}(\mathbf{R}, \mathbf{K}^+, \mathbf{K}^-) &= \sum_{u \in \mathcal{U}} \text{ELBO}(\mathbf{r}_u, \mathbf{k}_u^+, \mathbf{k}_u^-) \\ &+ \text{ELBO}(\mathbf{r}_u) + \text{ELBO}(\mathbf{k}^+) + \text{ELBO}(\mathbf{r}_u, \mathbf{k}_u^-) + \text{ELBO}(\mathbf{k}^+, \mathbf{k}_u^-). \end{aligned} \quad (8)$$

4 EXPERIMENTS

4.1 Datasets

We run experiments on two real-world datasets: Yelp [17] and HotelRec [1]. Each contains over 100k reviews with five-star ratings. As in [2, 9, 11], we extract the positive keyphrases from user reviews for the explanations. We consider keyphrases negative when they are not mentioned by users in their reviews. Each dataset contains




Figure 1: Probabilistic-graphical-model views of M&Ms-VAE [2] (left), M&Ms-VAE³ (middle), and M&Ms-VAE⁺ (right). In all models, both the implicit feedback r_u and the keyphrase-usage preferences k_u^+ are generated from user u 's latent representation z_u . Regarding the keyphrase-usage dislikes k_u^- , M&Ms-VAE³ considers k_u^- a third modality whereas M&Ms-VAE⁺ treats it only as an input variable. Solid lines denote the generative model, whereas dashed lines denote the variational approximation.

Table 1: Descriptive statistics of the datasets.

Dataset	#Users	#Items	#Interactions	Sparsity	#Keyphrases
Yelp	9,801	4,706	140,496	99.70%	234
Hotel	7,044	4,874	143,612	99.58%	141

complete observations and is split into 60% 20%/20% for the train, dev, and test sets. We binarize the ratings with the thresholds $t > 4.5$ and $t > 3.5$ for restaurants and hotels, respectively.

4.2 Experimental Settings

We treat the prior and the likelihood as normal and multinomial distributions. Each inference and generative network is composed of a two-layer neural network with a tanh activation function. We use dropout [14] and the Adam optimizer [8] with a learning rate of $5 \cdot 10^{-5}$. We anneal linearly the regularization parameter β of the KL terms. We tune each model on the NDCG on the validation set with a random search and a maximum of 50 trials. For critiquing, we tune the M&Ms-VAE-based models on the synthetic datasets.

4.3 RQ 1: Recommendation and Explanation Performance Comparison

4.3.1 Baselines. We compare M&Ms-VAE⁺ with the state-of-the-art M&Ms-VAE model, M&Ms-VAE³ (see Section 3), and the following baselines. **PLRec** [13] is a linear recommender that solves the scalability problem; it projects the user preferences into smaller space prior to a linear regression. **PureSVD** [6] builds a similarity matrix through SVD decomposition of the implicit rating matrix. **CE-VAE** [11] is a major improvement of the neural collaborative filtering model [7] with an explanation and a critiquing module¹.

4.3.2 Top-N Recommendation Performance. We report the following five metrics: R-Precision and NDCG; MAP, Precision, and Recall at 10. We present the main results in Table 2 (left columns).

Interestingly, M&Ms-VAE³ and M&Ms-VAE⁺ perform similarly to M&Ms-VAE, and the three models obtain the best results. This validates that their modifications do not affect recommendation performance (M&Ms-VAE⁺ even performs slightly better, particularly on the Hotel dataset). On both datasets, M&Ms-VAE-based models significantly outperform CE-VAE by an average factor of 1.7. These results indicate that the multimodal modeling of M&Ms-VAE-based

¹We remove the prior model CE-VNCF [15] because it consistently underperformed.

models and their training strategies are more robust than those of CE-VAE, which learns a bidirectional mapping between the latent space and the keyphrases, which perturbs the training process.

4.3.3 Top-K Keyphrase-explanation Performance. We evaluate the models in terms of keyphrase-explanation performance. PLRec and PureSVD do not contain an explanation module. In Table 2 (right columns), we report NDCG, MAP, Precision, and Recall at 10.

Overall, M&Ms-VAE and M&Ms-VAE⁺ obtain the best results. Nevertheless, M&Ms-VAE³ underperforms compared to (respectively outperforms) CE-VAE on the Hotel (respectively Yelp) dataset. It highlights the trade-off between recommendation and explanation that results from the additional training objective or modality.

4.4 RQ 2: Multi-step Negative and Positive Critiquing Performance Comparison

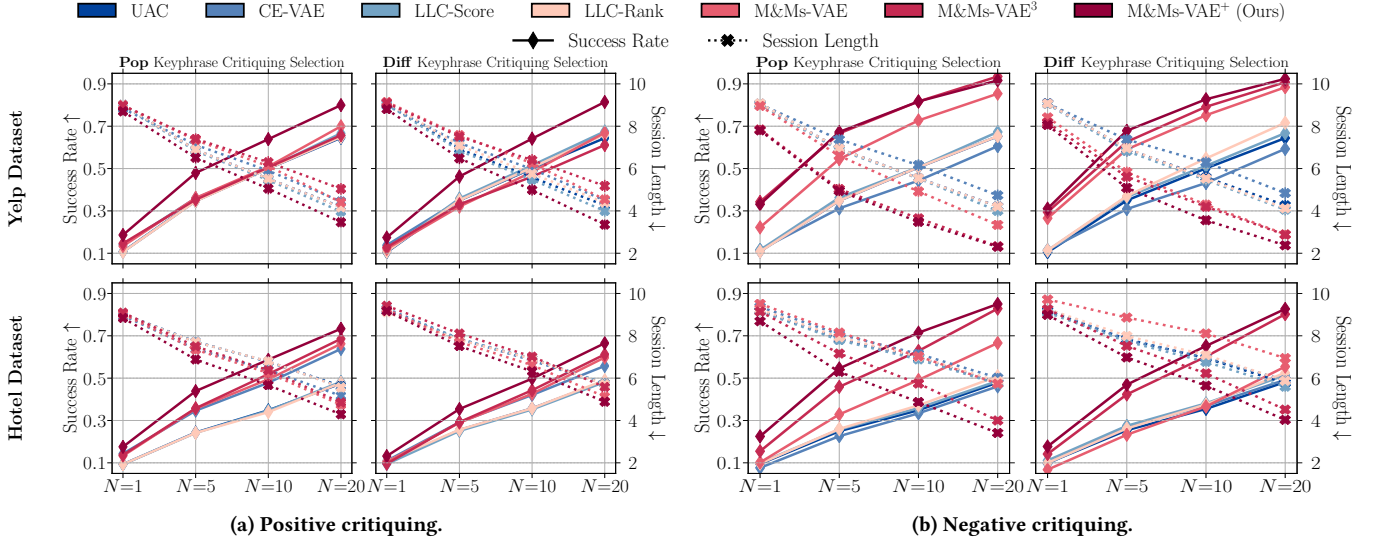
4.4.1 Baselines. We use **CE-VAE**, which learns an inverse mapping between a critique and the latent space. During critiquing, it averages the user embedding with the critique embedding. We also employ uniform average critiquing (**UAC**) [10], in which the user and critique embeddings are averaged. Finally, we include **LLC-Score** [10] and **LLC-Rank** [9], which extend PLRec [13] to co-embed keyphrases into the user embedding. Then, they compute a weighted average between the user's and each critique's embedding. Weights are optimized via linear programming; LLC-Score uses a margin-based objective, and LLC-Rank uses a ranking-based one. All models were designed for negative critique. For positive critiquing scenarios, we adapt their inner workings accordingly.

4.4.2 User Simulation. Following prior work [2, 9, 10], we run a user simulation to assess the models in a multi-step conversational recommendation scenario. We limit the critiquing iterations to a maximum of 10 turns. The conversations stop if the target item appears within the top-N recommendations on that iteration. The simulation includes all users and follows Alg. 1, with the following differences: (1) we consider only target items from the **test** set, and (2) on Line 7, \hat{k}_u^+ is replaced by k_u^+ to guarantee the same sequence of critiques and thus a fair comparison. We assess positive and negative critiquing performance in two separate simulations.

For the critique selection, we assume the user selects a keyphrase to critique according to two strategies: (1) **Pop**: the most popular keyphrase and (2) **Diff**: the keyphrase that deviates the most from the target item keywords. We compare the top recommended items'

Table 2: Recommendation and keyphrase-explanation results. Bold and underline denote the best and second-best results.

Model	Top-N Recommendation					Top-K Keyphrase Explanation				
	R-Precision	NDCG	MAP@10	Prec.@10	Rec.@10	NDCG@10	MAP@10	Prec.@10	Rec.@10	
Yelp	PLRec	0.0191	0.0703	0.0189	0.0166	0.0513	-	-	-	-
	PureSVD	0.0253	0.0825	0.0249	0.0206	0.0597	-	-	-	-
	CE-VAE	0.0136	0.0533	0.0132	0.0119	0.0367	0.2763	0.6356	0.5876	0.1819
	M&Ms-VAE	<u>0.0264</u>	0.0909	<u>0.0261</u>	0.0223	<u>0.0682</u>	<u>0.2787</u>	0.6428	0.5935	0.1834
	M&Ms-VAE ³	0.0263	0.0904	0.0259	<u>0.0220</u>	0.0685	0.2786	0.6405	0.5919	<u>0.1839</u>
	M&Ms-VAE ⁺ (Ours)	0.0270	<u>0.0905</u>	0.0256	0.0216	0.0671	0.2797	<u>0.6413</u>	<u>0.5931</u>	0.1842
Hotel	PLRec	0.0242	0.0664	0.0234	0.0190	0.0466	-	-	-	-
	PureSVD	0.0179	0.0541	0.0173	0.0145	0.0357	-	-	-	-
	CE-VAE	0.0147	0.0538	0.0146	0.0137	0.0334	0.3521	0.9113	0.8629	0.2105
	M&Ms-VAE	<u>0.0272</u>	0.0804	<u>0.0265</u>	<u>0.0227</u>	<u>0.0555</u>	0.3595	0.9393	0.8776	<u>0.2131</u>
	M&Ms-VAE ³	0.0262	0.0779	0.0261	0.0217	0.0529	0.3396	0.8668	0.8506	0.2075
	M&Ms-VAE ⁺ (Ours)	0.0279	<u>0.0797</u>	0.0266	0.0228	0.0561	<u>0.3591</u>	<u>0.9361</u>	0.8776	0.2133

**Figure 2: Multi-step critiquing performance after 10 turns of conversation. We simulate (a) positive and (b) negative critiquing. We report the average success rate (left y-axis) and session length (right y-axis) at different Top-N with 95% confidence interval.**

keyphrase frequency to the target item’s keyphrases and select the keyphrase with the largest frequency differential.

4.4.3 Multi-Step Critiquing Performance. We assess the models with the average success rate and session length at different Top-N. For each user and target item, we sample 299 unseen items, as [2].

Fig. 2a shows the results for positive critiquing. Impressively, M&Ms-VAE⁺ significantly outperforms all methods on both metrics and datasets. It confirms that positive critiques are effectively embedded. On the Yelp dataset, M&Ms-VAE³ and M&Ms-VAE obtain success rates similar to those of the baselines but the worst session lengths. However, UAC and LLCs clearly underperform on the Hotel dataset. These results validate our intuition that representing negative and positive critiques via the same encoder is suboptimal.

Regarding the depiction of negative critiquing in Fig. 2b, there is a large margin between M&Ms-VAE-based models and baselines. Remarkably, the new positive critiquing terms in the loss function

of Eq. 4 dramatically improve the success rate of M&Ms-VAE⁺ and M&Ms-VAE³ compared to M&Ms-VAE. Also, M&Ms-VAE⁺ obtains session lengths significantly shorter than those of M&Ms-VAE³.

These observations highlight the importance of inferring the representations of positive and negative critiques differently. Adding users’ keyphrase-usage dislikes as a third modality hinders the positive critiquing performance compared to M&Ms-VAE due to the generative model $p_{\Theta_k^-}$, that perturbs its training. Finally, M&Ms-VAE⁺ remedies this problem and produces the best results overall.

5 CONCLUSION

We present M&Ms-VAE⁺, an extension of M&Ms-VAE that enables positive and negative multi-step critiquing. The novelty relies on modeling users’ keyphrase-usage dislikes as a third modality while dropping its generative model, an approach made possible by the multimodal factorization. This enables the model to learn powerful positive and negative critique representation. The key results show

that M&Ms-VAE⁺ matches or exceeds M&Ms-VAE in recommendation and explanation and is the first model to obtain substantially better positive and negative multi-step critiquing performance.

REFERENCES

- [1] Diego Antognini and Boi Faltings. 2020. HotelRec: a Novel Very Large-Scale Hotel Recommendation Dataset. In *Proceedings of the 12th Language Resources and Evaluation Conference*. European Language Resources Association, Marseille, France, 4917–4923. <https://aclanthology.org/2020.lrec-1.605>
- [2] Diego Antognini and Boi Faltings. 2021. Fast Multi-Step Critiquing for VAE-Based Recommender Systems. In *Fifteenth ACM Conference on Recommender Systems* (Amsterdam, Netherlands) (*RecSys '21*). Association for Computing Machinery, New York, NY, USA, 209–219. <https://doi.org/10.1145/3460231.3474249>
- [3] Diego Antognini, Claudiu Musat, and Boi Faltings. 2021. Interacting with Explanations through Critiquing. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, Zhi-Hua Zhou (Ed.). International Joint Conferences on Artificial Intelligence Organization, 515–521. <https://doi.org/10.24963/ijcai.2021/72> Main Track.
- [4] Li Chen and Pearl Pu. 2012. Critiquing-based recommenders: survey and emerging trends. *User Modeling and User-Adapted Interaction* 22, 1-2 (04 2012). <https://doi.org/10.1007/s11257-011-9108-6>
- [5] Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. 2014. Empirical Evaluation of Gated Recurrent Neural Networks on Sequence Modeling. *arXiv preprint arXiv:1412.3555* (2014). <http://arxiv.org/abs/1412.3555>
- [6] Paolo Cremonesi, Yehuda Koren, and Roberto Turrin. 2010. Performance of Recommender Algorithms on Top-n Recommendation Tasks. In *Proceedings of the Fourth ACM Conference on Recommender Systems* (Barcelona, Spain) (*RecSys '10*). Association for Computing Machinery, New York, NY, USA, 39–46. <https://doi.org/10.1145/1864708.1864721>
- [7] Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. 2017. Neural Collaborative Filtering. In *Proceedings of the 26th International Conference on World Wide Web* (Perth, Australia) (*WWW '17*). International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 173–182. <https://doi.org/10.1145/3038912.3052569>
- [8] Diederik P. Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. In *3rd International Conference on Learning Representations, ICLR 2015*. <http://arxiv.org/abs/1412.6980>
- [9] Hanze Li, Scott Sanner, Kai Luo, and Ga Wu. 2020. A Ranking Optimization Approach to Latent Linear Critiquing for Conversational Recommender Systems. In *Fourteenth ACM Conference on Recommender Systems* (Virtual Event, Brazil) (*RecSys '20*). Association for Computing Machinery, New York, NY, USA, 13–22. <https://doi.org/10.1145/3383313.3412240>
- [10] Kai Luo, Scott Sanner, Ga Wu, Hanze Li, and Hojin Yang. 2020. Latent Linear Critiquing for Conversational Recommender Systems. In *Proceedings of The Web Conference 2020* (Taipei, Taiwan) (*WWW '20*). Association for Computing Machinery, New York, NY, USA, 2535–2541.
- [11] Kai Luo, Hojin Yang, Ga Wu, and Scott Sanner. 2020. Deep Critiquing for VAE-Based Recommender Systems. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval* (Virtual Event, China) (*SIGIR '20*). Association for Computing Machinery, New York, NY, USA, 1269–1278. <https://doi.org/10.1145/3397271.3401091>
- [12] Diana Andreea Petrescu, Diego Antognini, and Boi Faltings. 2021. Multi-Step Critiquing User Interface for Recommender Systems. In *Fifteenth ACM Conference on Recommender Systems* (Amsterdam, Netherlands) (*RecSys '21*). Association for Computing Machinery, New York, NY, USA, 760–763. <https://doi.org/10.1145/3460231.3478886>
- [13] Suvash Sedhain, Hung Bui, Jaya Kawale, Nikos Vlassis, Branislav Kveton, Aditya Krishna Menon, Trung Bui, and Scott Sanner. 2016. Practical linear models for large-scale one-class collaborative filtering. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence*. 3854–3860.
- [14] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research* 15, 1 (2014), 1929–1958.
- [15] Ga Wu, Kai Luo, Scott Sanner, and Harold Soh. 2019. Deep Language-Based Critiquing for Recommender Systems. In *Proceedings of the 13th ACM Conference on Recommender Systems* (Copenhagen, Denmark) (*RecSys '19*). Association for Computing Machinery, New York, NY, USA, 137–145. <https://doi.org/10.1145/3298689.3347009>
- [16] Hojin Yang, Tianshu Shen, and Scott Sanner. 2021. Bayesian Critiquing with Keyphrase Activation Vectors for VAE-based Recommender Systems. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval* (*SIGIR-21*). Online.
- [17] Yelp. 2016. Yelp Open Dataset 2016 - <https://www.yelp.com/dataset>. <https://www.yelp.com/dataset>