

Adaptive Contrastive Learning on Multimodal Transformer for Review Helpfulness Predictions

Thong Nguyen^{1,2}, Xiaobao Wu³, Anh Tuan Luu^{3*},
Cong-Duy Nguyen³, Zhen Hai⁴, Lidong Bing⁴

¹National University of Singapore, Singapore

²VinAI Research, Vietnam

³Nanyang Technological University, Singapore

⁴DAMO Academy, Alibaba Group

e0998147@u.nus.edu, anhtuan.luu@ntu.edu.sg

Abstract

Modern Review Helpfulness Prediction systems are dependent upon multiple modalities, typically texts and images. Unfortunately, those contemporary approaches pay scarce attention to polish representations of cross-modal relations and tend to suffer from inferior optimization. This might cause harm to model’s predictions in numerous cases. To overcome the aforementioned issues, we propose Multimodal Contrastive Learning for Multimodal Review Helpfulness Prediction (MRHP) problem, concentrating on mutual information between input modalities to explicitly elaborate cross-modal relations. In addition, we introduce Adaptive Weighting scheme for our contrastive learning approach in order to increase flexibility in optimization. Lastly, we propose Multimodal Interaction module to address the unalignment nature of multimodal data, thereby assisting the model in producing more reasonable multimodal representations. Experimental results show that our method outperforms prior baselines and achieves state-of-the-art results on two publicly available benchmark datasets for MRHP problem.

1 Introduction

Current e-commerce sites such as Amazon, Ebay, etc., construct review platforms to collect user feedback concerning their products. These platforms play a fundamental role in online transactions since they help future consumers collect useful reviews which assist them in deciding whether to make the purchase or not. Unfortunately, nowadays the number of user-generated reviews is overwhelming, raising doubts related to the relevance and veracity of reviews. Therefore, there is a need to verify the quality of reviews before publishing them to prospective customers. As a result, this inspires a recent surge of interest targeting the Review Helpfulness Prediction (RHP) problem.

*Corresponding Author

Product Information

The Cooks Standard 6-Quart Stainless Steel Stockpot with Lid is made with 18/10 stainless steel with an aluminum disc layered in the bottom. The aluminum disc bottom provides even heat distribution and prevents hot spots. Tempered glass lid with steam hole vent makes viewing food easy. Stainless steel riveted handles offer durability. Induction compatible. Works on gas, electric, glass, ceramic, etc. Oven safe to 500F, glass lid to 350F. Dishwasher safe.



Review 1

I needed a stainless steel pot for canning my tomatoes. I learned the hard way that you have to use a non-reactive pot or else your end result will be inedible (I thought I was using stainless steel but quickly realized it wasn't) I headed to Amazon and came across this Cooks Standard SS Cookpot with cover and bought it after reading the reviews. I have had it for just under a year and it still looks just as good as the day I bought it. I couldn't be happier with my purchase! Oh, and by the way, this one actually is stainless steel unlike the other pot I bought that said it was and wasn't.

Review 2

I ordered it on May 21st. What a waste of time and money.



	Review 1	Review 2
Label score	4	1
MCR score	0.168	3.637
Our Model score	4.651	0.743

Table 1: Example of unreasonable predictions in the Multimodal Review Helpfulness Prediction task.

Two principal groups of early efforts focus on purely textual data. The first group follows feature engineering techniques, retrieving argument-based features (Liu et al., 2017), lexical features (Kr-

ishnamoorthy, 2015), and semantic features (Kim et al., 2006), as input to their classifier. Inherently, their methods are labor-intensive and vulnerable to the typical issues of conventional machine learning methods. Instead of relying on manual features, the second group leverages deep neural models, for instance, RNN (Alsmadi et al., 2020) and CNN (Chen et al., 2018), to learn rich features automatically. Nonetheless, their approach is ineffective because the helpfulness of a review is not only contingent upon textual information but also other modalities.

To cope with the above issues, recent works (Liu et al., 2021b; Han et al., 2022) proposed to utilize multi-modality via the Multi-perspective Coherent Reasoning (MCR) model. Hypothesizing that a review is helpful if it exhibits coherent text and images with the product information, those works take into account both textual and visual modality of the inputs, then estimate their coherence level to discern whether the reviews are *helpful* or *unhelpful*. However, the MCR model contains a detrimental drawback. Particularly, it aims to maximize the scores s_p of positive (helpful) product-review pairs while minimizing those s_n of negative (unhelpful) pairs. Hence, it was assumed that following the aforementioned manner would project features with similar semantics to stay close and those with disparate ones to be distant apart. Unfortunately, in multimodal learning, this was shown not to be the case, causing the model to learn ad-hoc representations (Zolfaghari et al., 2021). This is one reason leading to unreasonable predictions of MCR in Table 1. As it can be seen, even though Review 1 closely relates to the product of “6-Quart Stainless Steel Stockpot”, the model classifies it as *unhelpful*. In addition, the target of Review 2’s text content is vague because it does not specifically correspond to the “Stockpot”. In fact, it can be used for any product. Moreover, the image does not clearly show any hint of the “Stockpot” as well. Despite such vagueness, the output of MCR for Review 2 is still *helpful*.

As a remedy to this problem, we propose Cross-modal Contrastive Learning to mine the mutual information of cross-modal relations in the input to capture more sensible representations. Nonetheless, plainly applying symmetric gradient pattern, which is similar to MCR that they assign equivalent penalty to s_n and s_p , is inflexible. In cases that s_p is small and s_n is already negatively skewed, or both

s_p and s_n are positively skewed, it is irrational to assign equivalent penalties to both s_p and s_n . Last but not least, MCR directly leverages Coherent Reasoning, repeatedly enforcing alignment among modalities in the input. This ignores the unaligned nature of multimodal input, for example, images might only refer to a particular section in the text, hence do not completely align with the textual content. In consequence, strictly forming alignment can make the model learn inefficient multimodal representations (Tsai et al., 2019).

To overcome the above problems, we propose an adaptive scheme to accomplish the flexibility in the optimization of our contrastive learning stage. Finally, we propose to adopt a multimodal attention module that reinforces one modality’s high-level features with low-level ones of other modalities. This not only relaxes the alignment assumption but also informs one modality of information of others, encouraging refined representation learning.

In sum, our contributions are three-fold:

- We propose an Adaptive Cross-modal Contrastive Learning for Review Helpfulness Prediction task by polishing cross-modal relation representations.
- We propose a Multimodal Interaction module which correlates modalities’ features without depending upon the alignment assumption.
- We conducted extensive experiments on two datasets for the RHP problem and found that our method outperforms other baselines which are both textual-only and multimodal, and obtains state-of-the-art results on those benchmarks.

2 Model Architecture

In this section we delineate the overall architecture of our MRHP model. Particular modules of our system are depicted in Figure 1.

2.1 Problem Definition

Given a product item p , which consists of a description T^p and images I^p , and a set of reviews $R = \{r_1, \dots, r_N\}$, where each review is composed of user-generated text T_i^r and images I_i^r , RHP model’s task is to generate the scores

$$s_i = f(p, r_i), \quad 1 \leq i \leq N \quad (1)$$

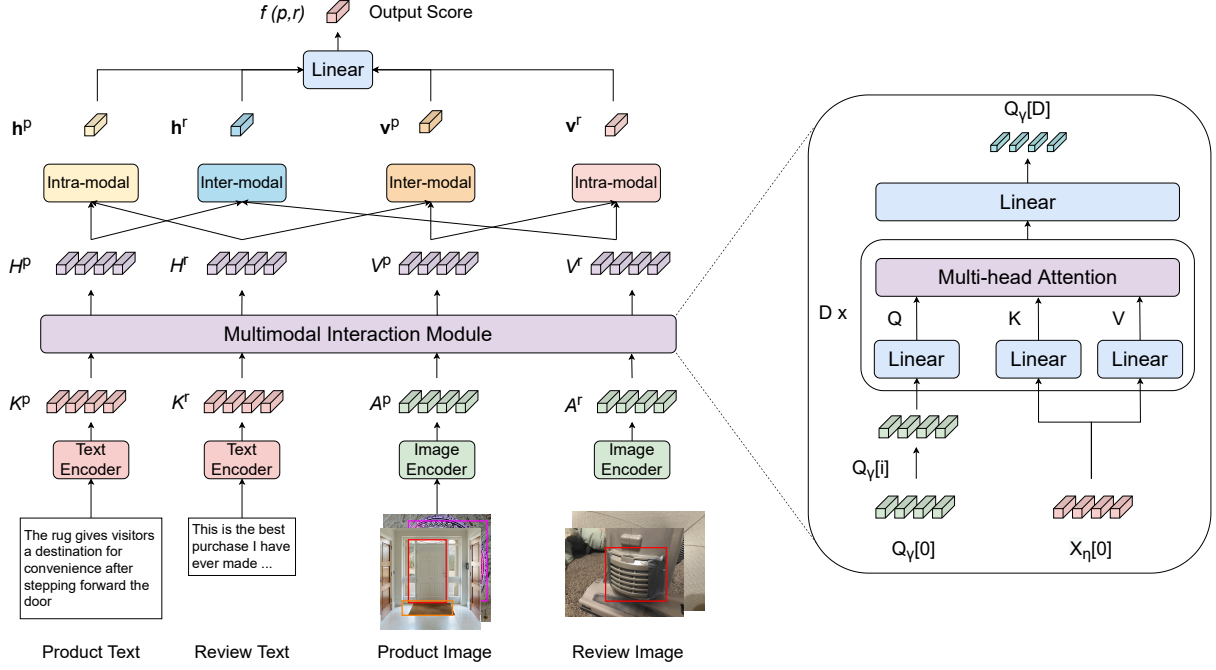


Figure 1: Diagram of our Multimodal Review Helpfulness Prediction model.

where N is the number of reviews for product p and f is the scoring function of the RHP model. Empirically, each score estimated by f indicates the helpfulness level of each review, and the ground-truth is the descending sort order of helpfulness scores.

2.2 Encoding Modules

Our model accepts product description T^p , product images I^p , review text T_i^r , and review images I_i^r as input. The encoding process of those elements is described as follows.

Text Encoding Product description and review text are sequences of words. Each sequence is indexed into the word embedding layer and then passed into the respective LSTM layer for product or review.

$$K^p = \text{LSTM}^p(\mathbf{W}_{\text{emb}}(T^p)) \quad (2)$$

$$K^r = \text{LSTM}^r(\mathbf{W}_{\text{emb}}(T^r)) \quad (3)$$

where $K^p \in \mathbb{R}^{l_p \times d}$, $K^r \in \mathbb{R}^{l_r \times d}$, l_p and l_r are the sequence lengths of product and review text respectively, and d is the hidden size.

Image Encoding We follow Anderson et al. (2018) to take detected objects as embeddings of the image. In particular, a pre-trained Faster R-CNN is applied to extract ROI features for m objects $\{\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_m\}$ from the product and review images. Subsequently, we encode extracted features using the self-attention module (SelfAttn) (Vaswani

et al., 2017)

$$A = \text{SelfAttn}(\{\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_m\}) \quad (4)$$

where $A \in \mathbb{R}^{m \times d}$ and d is the hidden size. Here we use A^p and A^r to indicate product and review image features, respectively.

2.3 Multimodal Interaction Module

We consider two components γ , η with their inputs X_γ , X_η , where η is the concatenation of input elements apart from the one in γ . For instance, if $\gamma = K^p$, then $\eta = [K^r, A^p, A^r]$, where $[\cdot, \cdot]$ indicates the concatenation operation. We define each cross-modal attention block to have three components Q , K , and V :

$$Q_\gamma = X_\gamma \cdot W_{Q_\gamma} \quad (5)$$

$$K_\eta = X_\eta \cdot W_{K_\eta} \quad (6)$$

$$V_\eta = X_\eta \cdot W_{V_\eta} \quad (7)$$

where $W_{Q_\gamma} \in \mathbb{R}^{d_\gamma \times d_k}$, $W_{K_\eta} \in \mathbb{R}^{d_\eta \times d_k}$, and $W_{V_\eta} \in \mathbb{R}^{d_\eta \times d_v}$ are weight matrices. The interaction between γ and η is computed in the cross-attention manner

$$Z_\gamma = \text{CM}_\gamma(X_\gamma, X_\eta) = \text{softmax} \left(\frac{Q_\gamma \cdot K_\eta^T}{\sqrt{d_k}} \right) \cdot V_\eta \quad (8)$$

Our full module comprises D layers of the above-mentioned attention block, as indicated in the right

part of Figure 1. Theoretically, the computation is carried out as follows

$$Q_\gamma[0] = X_\gamma \quad (9)$$

$$T[i] = \text{CM}_\gamma[i](\text{LN}(Q_\gamma[i-1]), \text{LN}(X_\eta)) \quad (10)$$

$$U_\gamma[i] = T[i] + Q_\gamma[i-1] \quad (11)$$

$$Q_\gamma[i] = \text{GeLU}(\text{Linear}(U_\gamma[i])) \quad (12)$$

where LN denotes layer normalization operator. We iteratively estimate cross-modal features for product text, product images, review text, and review images with a view to obtaining H^p , V^p , H^r , and V^r .

$$H^p = Q_k^p[D], \quad V^p = Q_a^p[D] \quad (13)$$

$$H^r = Q_k^r[D], \quad V^r = Q_a^r[D] \quad (14)$$

After our cross-modal interaction module, we proceed to pass features to undertake relation fusion in three paths: intra-modal, inter-modal, and intra-review.

Intra-modal Fusion The intra-modal alignment is calculated for two relation kinds: (1) product text - review text and (2) product image - review image. Firstly, we learn alignment among intra-modal features via self-attention modules

$$H^{\text{intraM}} = \text{SelfAttn}([H^p, H^r]) \quad (15)$$

$$V^{\text{intraM}} = \text{SelfAttn}([V^p, V^r]) \quad (16)$$

Then intra-modal hidden representations are fed to a CNN, and continuously a max-pooling layer to attain salient entries

$$\mathbf{z}^{\text{intraM}} = \text{MaxPool}(\text{CNN}([H^{\text{intraM}}, V^{\text{intraM}}])) \quad (17)$$

Inter-modal Fusion Similar to intra-modal alignment, inter-modal one is calculated for two types of relations as well: (1) product text - review image and (2) product image - review text. The first step is also to relate feature components using self-attention modules

$$H^{\text{prd_txt} - \text{rvw_img}} = \text{SelfAttn}([H^p, V^r]) \quad (18)$$

$$H^{\text{prd_img} - \text{rvw_txt}} = \text{SelfAttn}([V^p, H^r]) \quad (19)$$

We adopt a mean-pool layer to aggregate inter-modal features and then concatenate the pooled vectors to construct the final inter-modal represen-

tation

$$I^{\text{prd_txt} - \text{rvw_img}} = \text{MeanPool}(H^{\text{prd_txt} - \text{rvw_img}}) \quad (20)$$

$$I^{\text{prd_img} - \text{rvw_txt}} = \text{MeanPool}(H^{\text{prd_img} - \text{rvw_txt}}) \quad (21)$$

$$\mathbf{z}^{\text{interM}} = [I^{\text{prd_txt} - \text{rvw_img}}, I^{\text{prd_img} - \text{rvw_txt}}] \quad (22)$$

Intra-review Fusion The estimation of intra-review module completely mimics the inter-modal manner. The only discrimination is that the estimation is taken upon two different relations: (1) product text - product image and (2) review text - review image.

$$H^{\text{prd_txt} - \text{prd_img}} = \text{SelfAttn}([H^p, V^p]) \quad (23)$$

$$H^{\text{rvw_txt} - \text{rvw_img}} = \text{SelfAttn}([H^r, V^r]) \quad (24)$$

$$G^{\text{prd_txt} - \text{prd_img}} = \text{MeanPool}(H^{\text{prd_txt} - \text{prd_img}}) \quad (25)$$

$$G^{\text{rvw_txt} - \text{rvw_img}} = \text{MeanPool}(H^{\text{rvw_txt} - \text{rvw_img}}) \quad (26)$$

$$\mathbf{z}^{\text{intraR}} = [G^{\text{prd_txt} - \text{prd_img}}, G^{\text{rvw_txt} - \text{rvw_img}}] \quad (27)$$

Finally, we concatenate intra-modal, inter-modal, and intra-review output, and then feed the concatenated vector to the linear layer to obtain the ranking score:

$$\mathbf{z}^{\text{final}} = [\mathbf{z}^{\text{intraM}}, \mathbf{z}^{\text{interM}}, \mathbf{z}^{\text{intraR}}] \quad (28)$$

$$f(p, r_i) = \text{Linear}(\mathbf{z}^{\text{final}}) \quad (29)$$

3 Training Strategies

3.1 Adaptive Cross-modal Contrastive Learning

In this section, we explain the formulation and adaptive pattern along with its derivation of our Cross-modal Contrastive Learning.

Cross-modal Contrastive Learning First of all, we extract hidden states of helpful product-review pairs. Second of all, hidden features are max-pooled to extract meaningful entries.

$$\mathbf{h}^p = \text{MaxPool}(H^p), \quad \mathbf{h}^r = \text{MaxPool}(H^r) \quad (30)$$

$$\mathbf{v}^p = \text{MaxPool}(V^p), \quad \mathbf{v}^r = \text{MaxPool}(V^r) \quad (31)$$

We formulate our contrastive learning framework taking positive and negative pairs from the above-mentioned cross-modal features. In our framework, we hypothesize that pairs established by modalities

of the same sample are positive, whereas those formed by modalities of distinct ones are negative.

$$\mathcal{L}_{\text{CE}} = - \sum_{i=1}^B \text{sim}(\mathbf{t}_i^1, \mathbf{t}_i^2) + \sum_{j=1, k=1, j \neq k}^B \text{sim}(\mathbf{t}_j^1, \mathbf{t}_k^2) \quad (32)$$

where $\mathbf{t}^1, \mathbf{t}^2 \in \{\mathbf{h}^p, \mathbf{h}^r, \mathbf{v}^p, \mathbf{v}^r\}$, and B denotes the batch size in the training process.

Adaptive Weighting The standard contrastive objective suffers from inflexible optimization due to irrational gradient assignment to positive and negative pairs. As a result, to tackle the problem, we propose the Adaptive Weighting Strategy for our contrastive framework. Initially, we introduce weights ϵ^p and ϵ^n to represent distances from the optimum, then integrate them into positive and negative terms of our loss.

$$\begin{aligned} \mathcal{L}_{\text{AdaptiveCE}} = & - \sum_{i=1}^B \epsilon_i^p \cdot \text{sim}(\mathbf{t}_i^1, \mathbf{t}_i^2) \\ & + \sum_{j=1, k=1, j \neq k}^B \epsilon_{j,k}^n \cdot \text{sim}(\mathbf{t}_j^1, \mathbf{t}_k^2) \end{aligned} \quad (33)$$

where $\epsilon_i^p = [o^p - \text{sim}(\mathbf{t}_i^1, \mathbf{t}_i^2)]_+$ and $\epsilon_{j,k}^n = [\text{sim}(\mathbf{t}_j^1, \mathbf{t}_k^2) - o^n]_+$. Investigating the intuition to determine the values for o^p and o^n , we continue to conduct derivation and arrive in the following theorem

Theorem 1 *Adaptive Contrastive Loss (33) has the hyperspherical form:*

$$\begin{aligned} \mathcal{L}_{\text{AdaptiveCE}} = & \sum_{i=1}^B \left(\text{sim}(\mathbf{t}_i^1, \mathbf{t}_i^2) - \frac{o^p}{2} \right)^2 \\ & + \sum_{j=1, k=1, j \neq k}^B \left(\text{sim}(\mathbf{t}_j^1, \mathbf{t}_k^2) - \frac{o^n}{2} \right)^2 - C, \end{aligned}$$

where $C > 0$

We provide the proof for Theorem (1) in the Appendix section. As a consequence, theoretically the contrastive objective arrives in the optimum when $\text{sim}(\mathbf{t}_i^1, \mathbf{t}_i^2) = \frac{o^p}{2}$ and $\text{sim}(\mathbf{t}_j^1, \mathbf{t}_k^2) = \frac{o^n}{2}$. Based upon this observation, in our experiments we set $o^p = 2$ and $o^n = 0$.

3.2 Training Objective

For the Review Helpfulness Prediction problem, the model’s parameters are updated according to

the pairwise ranking loss as follows

$$\mathcal{L}_{\text{ranking}} = \sum_i \max(0, \beta - f(p_i, r^+) + f(p_i, r^-)) \quad (34)$$

where r^+ and r^- are random reviews in which r^+ possesses a higher helpfulness level than r^- . We jointly combine the contrastive goal with the ranking objective of the Review Helpfulness Prediction problem to train our model

$$\mathcal{L} = \mathcal{L}_{\text{AdaptiveCE}} + \mathcal{L}_{\text{ranking}} \quad (35)$$

4 Experiments

Dataset	Split	Category (Product / Review)		
		Clothing	Electronics.	Home
Lazada	Train & Dev	8K/130K	5K/52K	4K/16K
	Test	2K/32K	1K/13K	1K/13K
Amazon	Train & Dev	16K/349K	13K/325K	18K/462K
	Test	4K/87K	3K/80K	5K/111K

Table 2: Statistics of MRHP datasets.

4.1 Datasets

We evaluate our methods on two publicly available benchmark datasets for MRHP task: Lazada-MRHP and Amazon-MRHP.

Lazada-MRHP (Liu et al., 2021b) consists of product items and artificial reviews on Lazada.com, an e-commerce platform in Southeast Asia. All of the texts in the dataset are expressed in Indonesian.

Amazon-MRHP (Liu et al., 2021b) is collected from Amazon.com, the large-scale international e-commerce platform. Product information and associated reviews are in English and extracted between 2016 and 2018.

Both datasets comprise 3 categories: (i) Clothing, Shoes & Jewelry (Clothing), (ii) Electronics (Electronics), and (iii) Home & Kitchen (Home). We present the statistics of them in Table 2.

4.2 Implementation Details

We use a 1-layer LSTM with hidden dimension size of 128. We initialize our word embedding with fastText embedding (Bojanowski et al., 2017) for Lazada-MRHP dataset and 300-dimensional GloVe pretrained word vectors (Pennington et al., 2014) for Amazon-MRHP dataset. We set our multimodal attention module to have $D = 5$ attention layers. For the visual modality, we extract 2048-dimensional ROI features from each image and encode them into 128-dimensional vectors. Our

Type	Method	Clothing			Electronics			Home		
		MAP	N@3	N@5	MAP	N@3	N@5	MAP	N@3	N@5
Text-only	BiMPM	60.0	52.4	57.7	74.4	67.3	72.2	70.6	64.7	69.1
	EG-CNN	60.4	51.7	57.5	73.5	66.3	70.8	70.7	63.4	68.5
	Conv-KNRM	62.1	54.3	59.9	74.1	67.1	71.9	71.4	65.7	70.5
	PRH-Net	62.1	54.9	59.9	74.3	67.0	72.2	71.6	65.2	70.0
Multimodal	SSE-Cross	66.1	59.7	64.8	76.0	68.9	73.8	72.2	66.0	71.0
	DR-Net	66.5	60.7	65.3	76.1	69.2	74.0	72.4	66.3	71.4
	MCR	68.8	62.3	67.0	76.8	70.7	75.0	73.8	67.0	72.2
	Our Model	70.3	64.7	69.0	78.2	72.4	76.5	75.2	68.8	73.7

Table 3: Helpfulness Prediction results on Lazada-MRHP dataset.

Type	Method	Clothing			Electronics			Home		
		MAP	N@3	N@5	MAP	N@3	N@5	MAP	N@3	N@5
Text-only	BiMPM	57.7	41.8	46.0	52.3	40.5	44.1	56.6	43.6	47.6
	EG-CNN	56.4	40.6	44.7	51.5	39.4	42.1	55.3	42.4	46.7
	Conv-KNRM	57.2	41.2	45.6	52.6	40.5	44.2	57.4	44.5	48.4
	PRH-Net	58.3	42.2	46.5	52.4	40.1	43.9	57.1	44.3	48.1
Multimodal	SSE-Cross	65.0	56.0	59.1	53.7	43.8	47.2	60.8	51.0	54.0
	DR-Net	65.2	56.1	59.2	53.9	44.2	47.5	61.2	51.8	54.6
	MCR	66.4	57.3	60.2	54.4	45.0	48.1	62.6	53.5	56.6
	Our Model	67.4	58.6	61.6	56.5	47.6	50.8	63.5	54.6	57.8

Table 4: Helpfulness Prediction results on Amazon-MRHP dataset.

entire model is trained end-to-end with Adam optimizer (Kingma and Ba, 2014) and batch size of 32. For the training objective, we set the value of the margin in the ranking loss to be 1.

4.3 Baselines

We compare our proposed architecture against the following baselines:

- **BiMPM** (Wang et al., 2017): a ranking model which encodes input sentences in two directions to ascertain the matching result.
- **Conv-KNRM** (Dai et al., 2018): a CNN-based model which encodes n-gram of multiple lengths and uses kernel pooling to generate the final ranking score.
- **EG-CNN** (Chen et al., 2018): a CNN-based model targeting data scarcity and OOV problem in RHP task via taking advantage of character-based representations and domain discriminators.
- **PRH-Net** (Fan et al., 2019): a baseline to predict helpfulness of a review by taking into

consideration both product text and product metadata.

- **DR-Net** (Xu et al., 2020): a cross-modality approach that models contrast in associated contexts by leveraging decomposition and relation modules.
- **SSE-Cross** (Abavisani et al., 2020): multi-modal model to fuse different modalities with stochastic shared embeddings.
- **MCR** (Liu et al., 2021b): a baseline model focusing on coherent reasoning.

4.4 Automatic Evaluation

In Table 3 and 4, we follow previous work (Liu et al., 2021b) to report Mean Average Precision (MAP), Normalized Discounted Cumulative Gain (NDCG@N) (Järvelin and Kekäläinen, 2017) where $N = 3$ and $N = 5$. As it can be seen, multi-modal approaches achieve better performance than text-only ones.

For Lazada-MRHP dataset, we achieve an absolute improvement of NDCG@3 of 2.4 points in

Dataset	Clothing			Electronics			Home		
	MAP	N@3	N@5	MAP	N@3	N@5	MAP	N@3	N@5
Lazada	$4.48 \cdot 10^{-2}$	$1.55 \cdot 10^{-2}$	$3.93 \cdot 10^{-2}$	$4.54 \cdot 10^{-3}$	$1.05 \cdot 10^{-4}$	$2.63 \cdot 10^{-3}$	$1.09 \cdot 10^{-3}$	$3.40 \cdot 10^{-2}$	$3.68 \cdot 10^{-3}$
Amazon	$3.45 \cdot 10^{-2}$	$4.22 \cdot 10^{-2}$	$1.86 \cdot 10^{-2}$	$4.37 \cdot 10^{-3}$	$2.81 \cdot 10^{-2}$	$3.04 \cdot 10^{-2}$	$2.04 \cdot 10^{-3}$	$3.30 \cdot 10^{-3}$	$6.50 \cdot 10^{-3}$

Table 5: Significance test of the results of our model against MCR model.

Clothing, NDCG@5 of 1.5 points in Electronics, and MAP of 1.4 points in Home over the previous best method, which is MCR. In addition, our model also obtains better results than the best text-only RHP model, which is PRH-Net, with a gain of NDCG@3 of 9.8 points in Clothing, NDCG@5 of 4.3 points in Electronics, and MAP of 3.6 points in Home. Those results prove that our method can produce reasonable rankings for associated reviews.

For Amazon dataset, which is written in English, our model outperforms MCR on all 3 categories, by NDCG@5 of 1.4 points in Clothing, 2.7 points in Electronics, and 1.2 points in Home, respectively. These results have verified that our interaction module and optimization approach can come up with more useful multimodal fusion than previous state-of-the-art baselines, not only in English context but other language one as well.

We also perform significance tests to evaluate the statistical significance of our improvement on two datasets Amazon-MRHP and Lazada-MRHP, and note p-values in Table 5. As shown in the table, all of the p-values are smaller than 0.05, verifying the statistical significance in the enhancement of our method against prior best MRHP model, MCR (Liu et al., 2021b).

4.5 Case Study

In Table 1, we introduce an example of one product item and two reviews extracted from Electronics category of Amazon-MRHP dataset. Whereas MCR fails to predict relevant helpfulness scores, our model successfully produces sensible rankings for both of them. We hypothesize that our Multimodal Interaction module learns more meaningful representations and Adaptive Contrastive Learning framework acquires more logical hidden states of relations among input elements. Thus, our model is able to generate more rational outcomes.

4.6 Ablation Study

In this section, we proceed to study the impact of (1) Adaptive Contrastive Learning framework and (2) Cross-modal Interaction module.

Adaptive Contrastive Learning It is worth noting from Table 6 that plainly integrating contrastive learning brings less enhancement to the performance, with the improvement of NDCG@3 dropping 0.53 points in Lazada-MRHP dataset, NDCG@5 waning 0.84 points in Amazon-MRHP dataset. Furthermore, completely removing contrastive objective hurts performance, as NDCG@3 score decreasing 0.77 points in Lazada-MRHP, and MAP score declining 1.06 points in Amazon-MRHP. We hypothesize that the model loses the ability to learn efficient representations for cross-modal relations.

Cross-modal Interaction In this ablation, we eliminate the cross-modal interaction module. As shown in Table 6, without the module, the improvement is downgraded, for instance, N@3 drops 1.89 points in Lazada-MRHP dataset, MAP shrinks 1.39 points in Amazon-MRHP dataset. It is hypothesized that without the module, the model is rigidly dependent upon the alignment nature among multimodal input elements, which brings about insensible modeling because in most cases, cross-modal elements are irrelevant to be bijectively mapped together.

Dataset	Model	MAP	N@3	N@5
Lazada	Our Model	78.15	72.43	76.49
	- w/o Adaptive Weighting	77.90	71.90	75.97
	- w/o Contrastive Objective	77.69	71.66	75.85
	- w/o Cross-modal Module	77.32	70.54	74.86
Amazon	Our Model	56.49	47.62	50.79
	- w/o Adaptive Weighting	56.03	46.98	49.95
	- w/o Contrastive Objective	55.43	46.30	49.02
	- w/o Cross-modal Module	55.10	45.67	48.50

Table 6: Ablation study in Electronics category of Lazada-MRHP and Amazon-MRHP datasets.

4.7 Impact of Contrastive Learning on Cross-modal Relations

Despite improved performances, it remains a quandary that whether the enhancement stems from more meaningful representations of input samples, which we hypothesize as a significant benefit of our contrastive learning framework. For deeper investigation, we decide to statistically measure

Label	Model	Intra-modal		Inter-modal		Intra-review	
		CS	L2	CS	L2	CS	L2
1	MCR	0.785 ± 0.002	3.852 ± 0.067	0.843 ± 0.002	11.719 ± 0.001	0.845 ± 0.002	14.631 ± 0.001
	Our Model	0.875 ± 0.002	6.545 ± 0.007	0.957 ± 0.002	13.934 ± 0.027	0.953 ± 0.002	15.160 ± 0.036
4	MCR	0.533 ± 0.004	1.014 ± 0.051	0.712 ± 0.010	9.476 ± 0.001	0.617 ± 0.001	8.519 ± 0.001
	Our Model	0.433 ± 0.001	0.981 ± 0.005	0.564 ± 0.001	4.179 ± 0.017	0.538 ± 0.001	3.827 ± 0.020

Table 7: Intra-modal, Inter-modal, and Intra-review distances in Home category of Lazada-MRHP dataset.

Label	Model	Intra-modal		Inter-modal		Intra-review	
		CS	L2	CS	L2	CS	L2
1	MCR	0.785 ± 0.006	8.532 ± 0.292	0.686 ± 0.001	9.696 ± 0.300	0.880 ± 0.002	9.620 ± 0.217
	Our Model	0.971 ± 0.001	10.663 ± 0.770	0.976 ± 0.001	13.234 ± 0.493	0.970 ± 0.001	12.222 ± 0.431
4	MCR	0.697 ± 0.009	3.045 ± 0.139	0.624 ± 0.001	3.179 ± 0.830	0.781 ± 0.001	5.098 ± 0.636
	Our Model	0.571 ± 0.001	1.572 ± 0.037	0.488 ± 0.001	1.460 ± 0.008	0.487 ± 0.001	3.555 ± 0.001

Table 8: Intra-modal, Inter-modal, and Intra-review distances in Home category of Amazon-MRHP dataset.

distances among input samples using standard distance functions. Table 7 and 8 reveal the results of our experiment. In particular, we estimate the cosine distance (CS) and L2 distance (L2) between tokens of (1) product text - review text and product image - review image (intra-modal), (2) product text - review image and product image - review text (inter-modal), and (3) product text - product image and review text - review image (intra-review), then calculate the mean value of all samples. As it can be seen, our frameworks are more efficient in attracting elements of helpful pairs and repelling those of unhelpful pairs.

5 Related Work

5.1 Review Helpfulness Prediction

Past works that pursue Review Helpfulness Prediction (RHP) dilemma follow text-only approaches. In general, they extract salient information, for instance lexical (Krishnamoorthy, 2015), argument (Liu et al., 2017), and emotional features (Martin and Pu, 2014) from reviews. Subsequently, these features are fed to a standard classifier such as Random Forest (Louppe, 2014) in order to produce the output score. Inspired by the meteoric development of computation resources, contemporary approaches seek to take advantage of deep learning techniques to tackle the RHP problem. For instance, Wang et al. (2017) propose multi-perspective matching between review and product information via applying attention mechanism. Furthermore, Chen et al. (2018); Dai et al. (2018) adapt CNN models to learn textual representations in various views.

In reality, review content are not only determined by texts but also other modalities. As a consequence, Fan et al. (2019) integrate metadata information of the target product into the prediction model. Abavisani et al. (2020) filter out uninformative signals before fusing various modalities. Moreover, Liu et al. (2021b) perform coherent reasoning to ascertain the matching level between product and numerous review items.

5.2 Contrastive Estimation

Different from architectural techniques such as Knowledge Distillation (Hinton et al., 2015; Hahn and Choi, 2019; Nguyen and Luu, 2022) or Variational AutoEncoder (Zhao et al., 2020; Nguyen et al., 2021; Nguyen and Luu, 2021; Wang et al., 2019), Contrastive Learning has been introduced as a representation-based but universal mechanism to enhance natural language processing performance. Proposed by Chopra et al. (2005), Contrastive Learning has been widely adopted in myriad problems of Natural Language Processing (NLP).

As an approach to polish text representations, Gao et al. (2021); Zhang et al. (2021); Liu et al. (2021a); Nguyen and Luu (2021) employ contrastive loss to advance sentence embeddings and topic representations. For downstream tasks, Cao and Wang (2021) propose negative sampling strategies to generate noisy output so that the model can learn to distinguish correct summaries from incorrect ones in Document Summarization. For Spoken Question Answering (SQA), You et al. (2021) introduce augmentation algorithms in their contrastive learning stage so as to capture noisy-invariant rep-

representations of utterances. Additionally, Ke et al. (2021) inherit the formulation of the contrastive objective to construct distillation loss which transfers knowledge of the previous task to the current one. Their proposals are to improve tasks in the Aspect Sentiment Classification domain. Unfortunately, despite the surge of interest in exercising contrastive learning for NLP, research works to adapt the method to the MRHP task have been scant.

6 Conclusion

In this paper, we propose methods to polish representation learning for the Multimodal Review Helpfulness Prediction task. In particular, we aim to advance cross-modal relation representations by learning mutual information through contrastive learning. In order to further enhance our framework, we propose an adaptive weighting strategy to encourage flexibility in optimization. Moreover, we integrate a cross-modal interaction module to loose the model’s reliance on unalignment nature among modalities, continuing to refine multimodal representations. Our framework is able to outperform prior baselines and achieve state-of-the-art results on the MRHP problem.

7 Limitations

Despite the novelty and benefits of our method for Multimodal Review Helpfulness Prediction (MRHP) problem, it does include some drawbacks. Firstly, even though empirical results demonstrate that our approach not only works in English contexts, we have not conducted the verification in multilingual circumstances, in which product or review texts are written in different languages. If a model is corroborated to work efficiently in such contexts, it is capable of providing myriad benefits for practical implementation, for example, e-commerce applications can leverage such one single model for multiple cross-lingual scenarios. Furthermore, our work can also be extended to other domains. For instance, in movie assessment, we need to determine whether the review suits the material in the film, or visual scenes in the comment are consistent with the textual content. These would form our prospective future directions.

Secondly, in the MRHP problem, there are several relationships that contrastive learning could exploit to burnish the performance. In particular, performing contrastive discrimination between two

sets of reviews is able to furnish the model with useful set-based representations, which consolidate general knowledge for better helpfulness prediction. Similar insights are applicable for two sets of product information. At the moment, we leave such promising perspectives for future work.

8 Acknowledgement

This work was supported by Alibaba Innovative Research (AIR) programme with research grant AN-GC-2021-005.

References

- Mahdi Abavisani, Liwei Wu, Shengli Hu, Joel Tetreault, and Alejandro Jaimes. 2020. Multimodal categorization of crisis events in social media. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14679–14689.
- Abdalraheem Alsmadi, Shadi AlZu’bi, Mahmoud Al-Ayyoub, and Yaser Jararweh. 2020. Predicting helpfulness of online reviews. *arXiv preprint arXiv:2008.10129*.
- Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. 2018. Bottom-up and top-down attention for image captioning and visual question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6077–6086.
- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. Enriching word vectors with subword information. *Transactions of the association for computational linguistics*, 5:135–146.
- Shuyang Cao and Lu Wang. 2021. Cliff: Contrastive learning for improving faithfulness and factuality in abstractive summarization. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6633–6649.
- Cen Chen, Yinfei Yang, Jun Zhou, Xiaolong Li, and Forrest Bao. 2018. Cross-domain review helpfulness prediction based on convolutional neural networks with auxiliary domain discriminators. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 602–607.
- Sumit Chopra, Raia Hadsell, and Yann LeCun. 2005. Learning a similarity metric discriminatively, with application to face verification. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’05)*, volume 1, pages 539–546. IEEE.
- Zhuyun Dai, Chenyan Xiong, Jamie Callan, and Zhiyuan Liu. 2018. Convolutional neural networks

- for soft-matching n-grams in ad-hoc search. In *Proceedings of the eleventh ACM international conference on web search and data mining*, pages 126–134.
- Miao Fan, Chao Feng, Lin Guo, Mingming Sun, and Ping Li. 2019. Product-aware helpfulness prediction of online reviews. In *The World Wide Web Conference*, pages 2715–2721.
- Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021. Simcse: Simple contrastive learning of sentence embeddings. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6894–6910.
- Sangchul Hahn and Heeyoul Choi. 2019. Self-knowledge distillation in natural language processing. In *Proceedings of the International Conference on Recent Advances in Natural Language Processing (RANLP 2019)*, pages 423–430.
- Wei Han, Hui Chen, Zhen Hai, Soujanya Poria, and Lidong Bing. 2022. Sancl: Multimodal review helpfulness prediction with selective attention and natural contrastive learning. *arXiv preprint arXiv:2209.05040*.
- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. Distilling the knowledge in a neural network (2015). *arXiv preprint arXiv:1503.02531*, 2.
- Kalervo Järvelin and Jaana Kekäläinen. 2017. Ir evaluation methods for retrieving highly relevant documents. In *ACM SIGIR Forum*, volume 51, pages 243–250. ACM New York, NY, USA.
- Zixuan Ke, Bing Liu, Hu Xu, and Lei Shu. 2021. Classic: Continual and contrastive learning of aspect sentiment classification tasks. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6871–6883.
- Soo-Min Kim, Patrick Pantel, Timothy Chklovski, and Marco Pennacchiotti. 2006. Automatically assessing review helpfulness. In *Proceedings of the 2006 Conference on empirical methods in natural language processing*, pages 423–430.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Srikumar Krishnamoorthy. 2015. Linguistic features for review helpfulness prediction. *Expert Systems with Applications*, 42(7):3751–3759.
- Che Liu, Rui Wang, Jinghua Liu, Jian Sun, Fei Huang, and Luo Si. 2021a. Dialoguesec: Dialogue-based contrastive learning of sentence embeddings. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 2396–2406.
- Haijing Liu, Yang Gao, Pin Lv, Mengxue Li, Shiqiang Geng, Minglan Li, and Hao Wang. 2017. Using argument-based features to predict and analyse review helpfulness. *arXiv preprint arXiv:1707.07279*.
- Junhao Liu, Zhen Hai, Min Yang, and Lidong Bing. 2021b. Multi-perspective coherent reasoning for helpfulness prediction of multimodal reviews. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 5927–5936.
- Gilles Louppe. 2014. Understanding random forests: From theory to practice. *arXiv preprint arXiv:1407.7502*.
- Lionel Martin and Pearl Pu. 2014. Prediction of helpful reviews using emotions extraction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 28.
- Thong Nguyen and Anh Tuan Luu. 2021. Contrastive learning for neural topic model. *Advances in Neural Information Processing Systems*, 34:11974–11986.
- Thong Nguyen, Anh Tuan Luu, Truc Lu, and Tho Quan. 2021. Enriching and controlling global semantics for text summarization. *arXiv preprint arXiv:2109.10616*.
- Thong Thanh Nguyen and Anh Tuan Luu. 2022. Improving neural cross-lingual abstractive summarization via employing optimal transport distance for knowledge distillation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 11103–11111.
- Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543.
- Yao-Hung Hubert Tsai, Shaojie Bai, Paul Pu Liang, J Zico Kolter, Louis-Philippe Morency, and Ruslan Salakhutdinov. 2019. Multimodal transformer for unaligned multimodal language sequences. In *Proceedings of the conference. Association for Computational Linguistics. Meeting*, volume 2019, page 6558. NIH Public Access.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008.
- Yue Wang, Jing Li, Hou Pong Chan, Irwin King, Michael R Lyu, and Shuming Shi. 2019. Topic-aware neural keyphrase generation for social media language. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2516–2526.
- Zhiguo Wang, Wael Hamza, and Radu Florian. 2017. Bilateral multi-perspective matching for natural language sentences. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, pages 4144–4150.

- Nan Xu, Zhixiong Zeng, and Wenji Mao. 2020. Reasoning with multimodal sarcastic tweets via modeling cross-modality contrast and semantic association. In *Proceedings of the 58th annual meeting of the association for computational linguistics*, pages 3777–3786.
- Chenyu You, Nuo Chen, and Yuexian Zou. 2021. Self-supervised contrastive cross-modality representation learning for spoken question answering. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 28–39.
- Dejiao Zhang, Shang-Wen Li, Wei Xiao, Henghui Zhu, Ramesh Nallapati, Andrew O Arnold, and Bing Xiang. 2021. Pairwise supervised contrastive learning of sentence representations. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 5786–5798.
- He Zhao, Dinh Phung, Viet Huynh, Trung Le, and Wray Buntine. 2020. Neural topic model via optimal transport. In *International Conference on Learning Representations*.
- Mohammadreza Zolfaghari, Yi Zhu, Peter Gehler, and Thomas Brox. 2021. Crossclr: Cross-modal contrastive learning for multi-modal video representations. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1450–1459.

A Hyperspherical Form of Adaptive Contrastive Loss

We have the initial formulation of the adaptive contrastive loss

$$\mathcal{L}_{\text{AdaptiveCE}} = - \sum_{i=1}^B \epsilon_i^p \cdot \text{sim}(\mathbf{t}_i^1, \mathbf{t}_i^2) + \sum_{j=1, k=1, j \neq k}^B \epsilon_{j,k}^n \cdot \text{sim}(\mathbf{t}_j^1, \mathbf{t}_k^2) \quad (36)$$

We first substitute $\epsilon_i^p = [o^p - \text{sim}(\mathbf{t}_i^1, \mathbf{t}_i^2)]_+$ and $\epsilon_{j,k}^n = [\text{sim}(\mathbf{t}_j^1, \mathbf{t}_k^2) - o^n]_+$ into the above equation,

$$\mathcal{L}_{\text{AdaptiveCE}} = \sum_{i=1}^B \text{sim}(\mathbf{t}_i^1, \mathbf{t}_i^2)^2 - o^p \cdot \text{sim}(\mathbf{t}_i^1, \mathbf{t}_i^2) + \sum_{j=1, k=1, j \neq k}^B \text{sim}(\mathbf{t}_j^1, \mathbf{t}_k^2)^2 - o^n \cdot \text{sim}(\mathbf{t}_j^1, \mathbf{t}_k^2) \quad (37)$$

$$= \sum_{i=1}^B \left(\text{sim}(\mathbf{t}_i^1, \mathbf{t}_i^2) - \frac{o^p}{2} \right)^2 + \sum_{j=1, k=1, j \neq k}^B \left(\text{sim}(\mathbf{t}_j^1, \mathbf{t}_k^2) - \frac{o^n}{2} \right)^2 - C \quad (38)$$

where $C = \left(\frac{o^p}{2}\right)^2 + \left(\frac{o^n}{2}\right)^2$. Now we obtain the spherical form of our contrastive loss.