

MMAPS: End-to-End Multi-Grained Multi-Modal Attribute-Aware Product Summarization

Tao Chen¹, Ze Lin¹, Hui Li^{1*}, Jiayi Ji¹, Yiyi Zhou¹, Guanbin Li², Rongrong Ji¹

¹Key Laboratory of Multimedia Trusted Perception and Efficient Computing, Ministry of Education of China
Xiamen University

²School of Computer Science and Engineering, Sun Yat-sen University
{taochen, linze4014}@stu.xmu.edu.cn, {hui, zhouyiyi, rrji}@xmu.edu.cn
jjyxmu@gmail.com, liguanbin@mail.sysu.edu.cn

Abstract

Given the long textual product information and the product image, Multi-modal Product Summarization (MPS) aims to increase customers' desire to purchase by highlighting product characteristics with a short textual summary. Existing MPS methods can produce promising results. Nevertheless, they still 1) lack end-to-end product summarization, 2) lack multi-grained multi-modal modeling, and 3) lack multi-modal attribute modeling. To improve MPS, we propose an end-to-end multi-grained multi-modal attribute-aware product summarization method (MMAPS) for generating high-quality product summaries in e-commerce. MMAPS jointly models product attributes and generates product summaries. We design several multi-grained multi-modal tasks to better guide the multi-modal learning of MMAPS. Furthermore, we model product attributes based on both text and image modalities so that multi-modal product characteristics can be manifested in the generated summaries. Extensive experiments on a real large-scale Chinese e-commerce dataset demonstrate that our model outperforms state-of-the-art product summarization methods w.r.t. several summarization metrics. Our code is publicly available at: <https://github.com/KDEGroup/MMAPS>.

Keywords: product summarization, multi-modal learning

1. Introduction

With the development of the Internet, online shopping has become an integral part of people's daily life. Unlike brick-and-mortar stores, where customers can interact face-to-face with salespeople, people mainly learn about products through textual and pictorial descriptions in online stores. A product is typically described with a product title, one or a few product images and a long product description in the online store. For instance, Fig. 1 provides an example of a portable massage sticker sold in a Chinese online store.

However, long product descriptions increase the cognitive load and hurt the shopping experience. Hence, informative product summaries are critical for online stores to provide a better shopping experience and boost product sales. As new products emerge rapidly, manual summarization becomes cost prohibitive, leaving alone that it requires a certain level of expertise to write accurate and attractive product summaries. To overcome this problem, much effort has been devoted to designing product summarization methods (Chen et al., 2019; Li et al., 2020a; Song et al., 2022) that automatically summarizes product information to highlight product's characteristics and advantages.

Early methods (Wang et al., 2017; Khatri et al., 2018; Chen et al., 2019; Daultani et al., 2019)

mainly focus on leveraging product's textual information, such as descriptions and attributes. However, they ignore the product's visual modality which can provide rich information for product summarization. Recently, some works (Zhang et al., 2019; Li et al., 2020a; Song et al., 2022) consider multi-modal product information in product summarization, i.e., Multi-modal Product Summarization (MPS). As illustrated in Fig. 1, both the text modality and the visual modality are presented on product's web page and they convey cues regarding key product characteristics which are helpful information for generating eye-catching product summaries. Existing MPS methods have achieved promising performance. Nevertheless, they still suffer from several problems:

- **P1: Lack end-to-end product summarization.** State-of-the-art MPS approaches (Li et al., 2020a; Song et al., 2022) treat product attribute modeling and product summarization generation as two separate phases and train them independently, making it difficult to train and tune the complete process. Moreover, the error from the two phases is accumulated, negatively affecting the summarization.
- **P2: Lack multi-grained multi-modal modeling.** Existing MPS methods (Li et al., 2020a; Song et al., 2022) fuse the global information (e.g., sentence-level representation and image-level representation) from different modalities,

* Corresponding Author.



Figure 1: An example of a product in the CEPsum dataset (Li et al., 2020a). Product attributes are shown in red.

i.e., coarse-grained fusion. They neglect the importance of fine-grained multi-modal modeling (e.g., token-level and region-level representations alignment and fine-grained multi-modal product attribute modeling), which can improve the quality of generated product summaries.

- **P3: Lack multi-modal attribute modeling.** Product attributes describe key product characteristics (e.g., design and functionality) which help MPS models better understand products. However, existing works only model attributes from one modality, leading to inferior results. For instance, V2P (Song et al., 2022) models product attributes by extracting features from product images, while MMPG (Li et al., 2020a) considers attribute features from text.

To address the above issues, we propose an end-to-end Multi-Grained Multi-modal Attribute-aware Product Summarization method (MMAPS). Our contributions are summarized as follows:

- To deal with P1, we design an end-to-end multi-modal product summarization method MMAPS, which jointly models product attributes and summary generation. MMAPS can attend to the product characteristics from multiple modalities, helping generate coherent product summaries. The end-to-end learning process also reduces the difficulty of training and tuning.
- To remedy P2, we propose several multi-grained

multi-modal tasks to guide MMAPS. We design a coarse-grained dual-encoder contrastive learning task to align cross-modal coarse-grained information coarsely. Additionally, we design the fine-grained multi-modal alignment task and the fine-grained product attribute prediction task to endow MMAPS with the ability to capture fine-grained cross-modal product information.

- To handle P3, in the fine-grained product attribute prediction task, MMAPS models product attributes based on both text and image modality, which helps MMAPS understand multi-modal product characteristics and guides MMAPS to pay more attention to the significant features when generating product summaries.
- We conduct extensive experiments on a large-scale Chinese e-commerce dataset. Experimental results demonstrate that MMAPS outperforms state-of-the-art product summarization methods.

2. Related Work

2.1. Product Summarization

E-Commerce product summarization aims to generate text summaries that provide customers with the most valuable information about the product, increasing their desire to purchase the product.

Traditional works only take the long textual descriptions of the products as input. For example, Yuan et al. (Yuan et al., 2020) propose to use a dual-copy mechanism to generate faithful product summaries, which can selectively copy tokens from both product descriptions and attributes. EPCCG (Guo et al., 2022) utilizes Transformer (Vaswani et al., 2017) to summarize controllable product copywriting from product title, attributes, and OCR text in different aspects.

Despite their promising performance, earlier methods do not incorporate the products' visual signal in the text generation. The visual modality can help customers discriminate essential product characteristics and thus improve the quality of the generated summaries.

Recently, a few works resort to multi-modal product summarization which considers multi-modal information. For example, MMPG (Li et al., 2020a) introduce a multi-modal pointer-generator network to generate an aspect-aware textual summary for Chinese e-commerce products by integrating textual and visual product information.

V2P (Song et al., 2022) unifies the heterogeneous multi-modal data in the same embedding space by converting the vision modality into semantic attribute prompts.

2.2. Multi-Modal Self-Supervised Learning

Self-supervised learning has been widely used in pre-trained models (Liu et al., 2023). In multi-modal learning, many works (Li et al., 2021b; Sheng et al., 2021; Lin et al., 2020) have exploited contrastive learning, a type of self-supervised learning, to model the correspondence among different modalities. For instance, Oscar (Li et al., 2020b) learns cross-modal representations by predicting whether the image-text pair contains the original image representation or any polluted one. ALBEF (Li et al., 2021a) learns a similarity function such that parallel image-text pairs have higher similarity scores. Based on the learned function, ALBEF enforces the representations of an image and a text in a pair close to each other. AVTS (Korbar et al., 2018) leverage the consistency of videos and audios to train the deep encoder. MMV (Alayrac et al., 2020) further considers the consistency among video, audio and text.

3. Our Method MMAPS

In this section, we illustrate the details of our proposed MMAPS. Fig. 2 provides an overview of MMAPS. MMAPS takes the product information X , which contains the title and the long description of a product, and the product image I as input, and it generates a product summary that is much shorter than the product description.

3.1. Architecture Design

The architecture of MMAPS consists of four components. We directly use the design of BART decoder as the decoder in MMAPS to decode summaries. Text encoding and image encoding components are used for encoding textual features and visual features, respectively. The text-image fusion component is designed for fusing multi-modal representations that are prepared for the product summarization. In the following, we illustrate the details of text encoding, image encoding and text-image fusion component.

3.1.1. Text Encoding

The product information X is tokenized and a special token [CLS] is appended to the token sequence in order to represent the overall, sequence-level information after text encoding. Then, we use the pre-trained BART encoder as the text encoder in MMAPS and the token sequence is fed into it. The output encoded token representation sequence is denoted as $\mathbf{Z} = \{\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_L, \mathbf{z}_{\langle \text{cls} \rangle}\}$ where L is the length of the sequence, \mathbf{z}_i is the encoded representation of i -th token, and $\mathbf{z}_{\langle \text{cls} \rangle}$ is the

representation of [CLS]. $\mathbf{z}_{\langle \text{cls} \rangle}$ is inferred by all tokens in the sequence and it can be regarded as the sequence-level representation.

3.1.2. Image Encoding

We use the Faster R-CNN (Ren et al., 2015) pre-trained on Visual Genome (Krishna et al., 2017) to extract region representations. Specifically, we feed the input image I to the pre-trained Faster R-CNN and extract all the detected objects. We only retain m objects with the highest confidence. To keep the spatial information of the image, each image region i is encoded as the sum of three types of features (Cho et al., 2021):

1. RoI (Region of Interest) object feature \mathbf{v}_i ;
2. RoI bounding box coordinate feature $\mathbf{e}_i^{\text{box}}$, which is encoded with a linear layer;
3. Region id feature $\mathbf{e}_i^{\text{reg}}$, which is encoded by the embedding layer in the text encoder.

The final region representation \mathbf{o}_i of i can be obtained as follows:

$$\begin{aligned} \mathbf{v}_i, \mathbf{c}_i, \mathbf{r}_i &\in \text{Faster R-CNN}(I) \\ \mathbf{e}_i^{\text{box}} &= [\mathbf{c}_i, \mathbf{s}_i] \mathbf{W}_e + \mathbf{B}_e, \quad \mathbf{o}_i = \mathbf{v}_i + \mathbf{e}_i^{\text{box}} + \mathbf{e}_i^{\text{reg}} \end{aligned} \quad (1)$$

where \mathbf{c}_i is a 4-dimensional normalized position vector indicating the coordinates of the top-left and bottom-right corners, $\mathbf{s}_i \in \mathbb{R}^1$ indicates the corresponding area of \mathbf{c}_i , \mathbf{r}_i is the class distribution of region i which is later used in the Masked Region Modeling task introduced in Sec. 3.2.2, $[\cdot, \cdot]$ represents the concatenation operation, and \mathbf{W}_e and \mathbf{B}_e are learnable parameters. Faster R-CNN(I) outputs m triples $\{\mathbf{v}, \mathbf{c}, \mathbf{r}\}$ and each of them corresponds to an image region. If an image has less than m detected regions, we use vectors with all zeros as \mathbf{o} to pad region representations to have m region representations for that image.

\mathbf{o} is only position-based region representation. To further model the context of the region, we subsequently feed \mathbf{o} into a Transformer-based unimodal encoder. For an image, we construct a region representation sequence with its m region representations: $\mathbf{O} = \{\mathbf{o}_{\langle \text{cls} \rangle}, \mathbf{o}_1, \mathbf{o}_2, \dots, \mathbf{o}_m\}$. Similar to the [CLS] token used in text encoding, we prepend a learnable parameter $\mathbf{o}_{\langle \text{cls} \rangle}$ to \mathbf{O} to represent the global representation of the image. The image encoder consists of H stacked layers and each layer includes two sub-layers: 1) Multi-Head Attention (MHA) and 2) a position-wise Feed-Forward Network (FFN). Finally, we obtain the output visual features $\mathbf{G} = \{\mathbf{g}_{\langle \text{cls} \rangle}, \mathbf{g}_1, \mathbf{g}_2, \dots, \mathbf{g}_m\}$:

$$\mathbf{S}_V^h = \text{MHA}(\mathbf{U}_V^{h-1}) + \mathbf{U}_V^{h-1}, \quad \mathbf{U}_V^h = \text{FFN}(\mathbf{S}_V^h) + \mathbf{S}_V^h \quad (2)$$

where \mathbf{U}_V^0 is the input region representations \mathbf{O} .

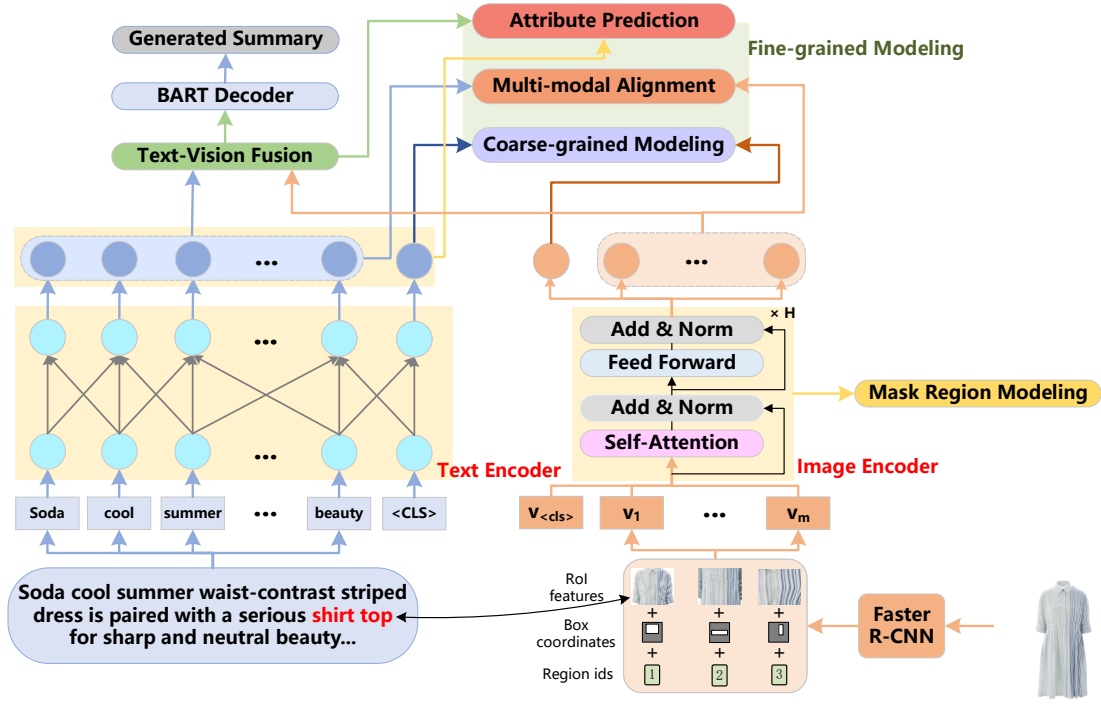


Figure 2: Overview of MMAPS. Chinese product information has been translated into English.

3.1.3. Text-Image Fusion

MMAPS contains a multi-modal fusion component. It receives the encoded text features \mathbf{Z} and image features \mathbf{G} from the text and image encoding component and outputs fused, multi-modal features containing related information derived from textual and visual modality, aiming to complement the information of textual modality while potentially removing noisy information.

The process of the multi-modal fusion is shown as follows:

$$\begin{aligned} \mathbf{Q} &= \mathbf{Z}\mathbf{W}_q, \quad \mathbf{K} = \mathbf{G}\mathbf{W}_k, \quad \mathbf{V} = \mathbf{G}\mathbf{W}_v \\ \mathbf{C} &= \text{CMA}(\mathbf{Q}, \mathbf{K}, \mathbf{V}), \quad \mathbf{F} = \sigma([\mathbf{Z}, \mathbf{C}]\mathbf{W}_f + \mathbf{B}_f) \\ \mathbf{Z}' &= [\mathbf{Z}, \mathbf{F} \otimes \mathbf{C}]\mathbf{W}_{z'} + \mathbf{B}_{z'} \end{aligned} \quad (3)$$

where $\text{CMA}(\cdot)$ indicate the cross-modality attention layer using the Multi-Head Attention mechanism (i.e., $\text{MHA}(\cdot)$ in Eq. 2 using different input as query, key and value), $\sigma(\cdot)$ refers to the Sigmoid function, \mathbf{W}_* and \mathbf{B}_* are learnable parameters, and \otimes is element-wise vector multiplication. Specifically, the textual features \mathbf{Z} are linearly projected into queries \mathbf{Q} and the visual features \mathbf{G} are linearly mapped to key-value pairs \mathbf{K} and \mathbf{V} . Next, to retain the pre-trained text features from BART and overcome the noise brought by the visual modality, we apply a forget gate \mathbf{F} . Finally, we concatenate the textual features \mathbf{Z} and the result of $\mathbf{F} \otimes \mathbf{C}$ to generate the multi-modal features \mathbf{Z}' .

3.2. Multi-Modal Multi-task Learning

We design several tasks to guide the multi-modal learning of MMAPS:

3.2.1. Product Summarization

To supervise the product summarization task, given input text X and image I , we fulfill the output-level supervision by minimizing the negative log-likelihood:

$$\mathcal{L}_{\text{PS}} = - \sum_{y \in Y} \sum_{t=1}^{|y|} \log(p(y_t | I, X, y_1, \dots, y_{t-1})), \quad (4)$$

where y_t is the t -th token in the ground truth summary y , $|y|$ is the number of tokens in y and Y is the ground truth summary set.

3.2.2. Masked Region Modeling

We adopt the Masked Region Modeling task to improve the performance of the image encoder. For input images, MMAPS samples S image regions with a probability of 15% and masks the sampled regions. Then, we train the model to predict the class of the masked regions. For the s -th masked region, we use \mathbf{g}_s and \mathbf{r}_s to denote the visual feature output by the image encoder and the class distribution detected by Faster R-CNN (i.e., \mathbf{r} in Eq. 1), respectively. MMAPS minimizes the KL divergence of the predicted class distribution and \mathbf{r}_s :

$$\mathcal{L}_{\text{MRM}} = \sum_{s=1}^S D_{\text{KL}}(\mathbf{r}_s \| \text{MLP}(\mathbf{g}_s)), \quad (5)$$

where $\text{MLP}(\cdot)$ is a two-layer perception for classification.

3.2.3. Multi-grained Multi-modal Modeling

As product summarization involves two correlated modalities, modeling the correlation between the text modality and the image modality will help MMAPS generate expressive text that corresponds to the product image. Moreover, both coarse-grained information (sentence-level representation and image-level representation) and fine-grained information (token-level representation and region-level representation) from the two modalities are related to depicting the product and their cross-modal correlation affects the quality of the generated product summaries. Consequently, we design coarse-grained modeling and fine-grained modeling to help MMAPS better capture the cross-modal correlation.

Coarse-grained Multi-modal Modeling. We design a coarse-grained dual-encoder contrastive learning task to jointly optimize the image encoder and the text encoder by contrasting the text-image pairs against others in the same batch. Specifically, for each image-text pair (positive pair) in a batch of B image-text pairs, we use its image and all the text in the remaining $B - 1$ pairs to form $B - 1$ negative image-text pairs. Then, the image encoder and the text encoder are trained to maximize the similarity between the image-level representation ($\mathbf{g}_{\langle \text{cls} \rangle}$) and the sentence-level representation ($\mathbf{z}_{\langle \text{cls} \rangle}$) for a positive pair and minimize the similarity of representations corresponding to the $B(B-1)$ negative pairs. The loss of the contrastive learning is shown as follows:

$$\begin{aligned} \mathbf{g}_{\text{norm}}^i &= \mathbf{W}_g \mathbf{g}_{\langle \text{cls} \rangle}^i / \left\| \mathbf{W}_g \mathbf{g}_{\langle \text{cls} \rangle}^i \right\|_2 \\ \mathbf{z}_{\text{norm}}^j &= \mathbf{W}_z \mathbf{z}_{\langle \text{cls} \rangle}^j / \left\| \mathbf{W}_z \mathbf{z}_{\langle \text{cls} \rangle}^j \right\|_2 \\ \mathcal{L}_{i2t} &= -\frac{1}{B} \sum_i \log \frac{\exp(\mathbf{g}_{\text{norm}}^{i\top} \mathbf{z}_{\text{norm}}^i / \tau)}{\sum_{j=1}^B \exp(\mathbf{g}_{\text{norm}}^{i\top} \mathbf{z}_{\text{norm}}^j / \tau)} \quad (6) \\ \mathcal{L}_{t2i} &= -\frac{1}{B} \sum_i \log \frac{\exp(\mathbf{z}_{\text{norm}}^{i\top} \mathbf{g}_{\text{norm}}^i / \tau)}{\sum_{j=1}^B \exp(\mathbf{z}_{\text{norm}}^{i\top} \mathbf{g}_{\text{norm}}^j / \tau)} \\ \mathcal{L}_{\text{CMM}} &= \mathcal{L}_{i2t} + \mathcal{L}_{t2i} \end{aligned}$$

where $\mathbf{g}_{\text{norm}}^i$ and $\mathbf{z}_{\text{norm}}^j$ are normalized representations of the image in the i -th pair and the text in the j -th pair, respectively. \mathbf{W}_g and \mathbf{W}_z are weights and τ is a learnable temperature parameter.

Fine-grained Multi-modal Modeling. The coarse-grained multi-modal modeling improves the quality of sequence-level representations and image-level representations via aligning global semantic information. Besides, product attributes, which depict the key product characteristics that customers are most concerned with, also require fine-grained multi-modal modeling to better model

attributes from two modalities. Therefore, we further propose two fine-grained multi-modal modeling tasks to enhance MMAPS.

1. **Fine-grained Multi-modal Alignment.** We design a fine-grained multi-modal alignment task to train MMAPS to model the semantic correlation between the text modality and the image modality by inspecting individual regions and tokens. This way, MMAPS can capture the salient information that appears in both modalities. To do this, we use Hausdorff distance, which can measure the similarity between two sequences of different features, to align textual and visual features:

$$\begin{aligned} d(\mathbf{G}, \mathbf{Z}) &= \max_i \min_j \left\| \frac{\text{MLP}(\mathbf{g}_i)}{\|\text{MLP}(\mathbf{g}_i)\|_2} - \frac{\mathbf{z}_j}{\|\mathbf{z}_j\|_2} \right\|_2 \\ d(\mathbf{Z}, \mathbf{G}) &= \max_j \min_i \left\| \frac{\text{MLP}(\mathbf{g}_i)}{\|\text{MLP}(\mathbf{g}_i)\|_2} - \frac{\mathbf{z}_j}{\|\mathbf{z}_j\|_2} \right\|_2 \quad (7) \\ d_H(\mathbf{G}, \mathbf{Z}) &= \max \{d(\mathbf{G}, \mathbf{Z}), d(\mathbf{Z}, \mathbf{G})\} \\ \mathcal{L}_{\text{HD}} &= d_H^2 \end{aligned}$$

In Eq. 7, we use a two-layer perception $\text{MLP}(\cdot)$ to map the encoded visual features into the text representation space. The above Hausdorff Loss can capture the local boundary information and enforce multi-modal alignment, making features in different modalities that correspond to the same characteristics getting closer in the common semantic space.

2. **Fine-grained Multi-Modal Product Attribute Prediction.** Each product in e-commerce platform is typically described by some product attributes. For instance, Fig. 2 shows that the product attribute ‘‘shirt top’’ is manifested in both the text and the image. Including the descriptions of product attributes in the generated text improves the quality of product summaries and makes the summaries eye-catching. Moreover, they help customers quickly understand and distinguish the bright spots of different products. Therefore, we design a multi-modal product attribute prediction task, which trains MMAPS to predict product attributes, to endow MMAPS with the ability of understanding multi-modal product characteristics and guide MMAPS to pay more attention to the significant features when generating product summaries. For a product in the dataset, we assign a multi-hot vector as the ground-truth attribute vector, denoted as $\mathbf{y}^a = (\mathbf{y}_1^a, \dots, \mathbf{y}_N^a)$, where $\mathbf{y}_l^a = 1$ denotes the product has the l -th attribute, $\mathbf{y}_l^a = 0$ otherwise. N is the total number of the attributes. To predict product attributes, we feed the text representation \mathbf{Z} and the multi-modal representation \mathbf{Z}' into a feed-forward layer to

Category	Home Appliances	Clothing	Cases & Bags
#Train Sample	437,646	790,297	97,510
#Valid Sample	10,000	10,000	5,000
#Test Sample	10,000	10,000	5,000
Avg. Length of Input	335	286	299
Avg. Length of Output	79	78	79

Table 1: Statistics of data. The unit of the average length of input/output is one Chinese character.

predict the attribute vector:

$$\hat{\mathbf{y}}^a = \sigma(\mathbf{W}_y^{(1)}(\mathbf{W}_y^{(2)} \sum_{i=1}^L \mathbf{z}'_i + \mathbf{W}_y^{(3)} \sum_{j=1}^L \mathbf{z}_j + \mathbf{W}_y^{(4)} \mathbf{z}_{\langle cls \rangle}) + \mathbf{B}_y), \quad (8)$$

where \mathbf{W} and \mathbf{B} are trainable weights, and $\hat{\mathbf{y}}^a$ indicates the predicted attributes. We use a binary cross entropy loss for the multi-modal product attribute prediction task:

$$\mathcal{L}_{\text{ATT}} = \min_{\Theta_a} - \sum^B (\mathbf{y}^a \ln(\hat{\mathbf{y}}^a) + (\mathbf{1} - \mathbf{y}^a) \ln(\mathbf{1} - \hat{\mathbf{y}}^a)), \quad (9)$$

where $\mathbf{1} \in \mathbb{R}^N$ is a vector with all elements being one, and Θ_a stands for corresponding to-be-learned parameters.

3.3. Putting All Together

In summary, training MMAPS involves optimizing four parts and the overall objective is defined as follows:

$$\mathcal{L} = \mathcal{L}_{\text{PS}} + \lambda_1 \mathcal{L}_{\text{MRM}} + \lambda_2 \mathcal{L}_{\text{CMM}} + \lambda_3 \mathcal{L}_{\text{FMM}}, \quad (10)$$

where $\mathcal{L}_{\text{FMM}} = \mathcal{L}_{\text{HD}} + \mathcal{L}_{\text{ATT}}$ and λ_* are pre-defined loss weights.

4. Experiment

In this section, we report and analyze the experimental results in order to answer the following research questions:

- **RQ1.** Does MMAPS outperform state-of-the-art product summarization methods w.r.t. to different summarization metrics?
- **RQ2.** Does each component in MMAPS contribute to its overall performance?
- **RQ3.** Is MMAPS sensitive to the setting of task weights in Eq. 10?
- **RQ4.** Can MMAPS generate more coherent and descriptive summaries than baselines?

4.1. Experimental Settings

Dataset. We use the CEPsum dataset¹ (Li et al., 2020a), which is collected from an e-commerce

¹<https://github.com/hrlinlp/cepsum>

platform in China. It includes around 1.4 million products covering three categories: Home Appliances, Clothing, and Cases & Bags. Each product in CEPsum contains a long product description, a product title, a product image, and a high-quality summary written by humans. Tab. 1 provides the data statistics.

For the multi-modal product attribute prediction task, we construct the pre-defined attribute vocabulary following V2P (Song et al., 2022). Specifically, we use Jieba² to tokenize the dataset and then perform the part-of-speech tagging. The attribute vocabulary for each category consists of adjectives and nouns with more than one Chinese character and have appeared in more than a pre-defined number of products' summaries. According to the scale of different product categories, the pre-defined number threshold for Home Appliance, Clothing, and Cases & Bags categories are set to 5,000, 10,000, and 1,000, respectively.

Implementation Details. We use the pre-trained bart-base-chinese³ to construct the text encoder, which has six layers. For the image encoder, we stack 4 layers (i.e., H in Sec. 3.1.2) with 8 attention heads and the hidden dimensionality is 2,048. The parameters of the image encoder are initialized randomly. The batch size B of Eq. 6 is 16. The dropout rate is 0.1. The maximum length L of the text sequence is set to 400, and the retained number M of regions per image is set to 36. We conduct a grid search for task weights λ_1 , λ_2 and λ_3 in Eq. 10. The default λ_1 , λ_2 and λ_3 are set to 0.8, 0.05 and 0.3, respectively. We also report the impact of performance when using different task weights in Sec. 4.4. The models are optimized using Adam (Kingma and Ba, 2015) and a cosine learning rate schedule with the initial learning rate of $3e^{-5}$. Notably, we report the average of three runs with different random seeds on the testing set as experimental results in this paper.

Evaluation Metric. We use ROUGE (1, 2, L) (Lin, 2004), BLEU (1, 2, 3, 4) (Papineni et al., 2002), S-BLEU (NLTK, 2023), METEOR (Banerjee and Lavie, 2005) and BERTScore (Zhang et al., 2020) as the evaluation metrics. BLEU measures the micro-average precision. We also adopt S-BLEU that averages the sentence-level BLEU-4 scores (i.e., macro-average precision). Instead of exact matching, BERTScore measures the relevance between generated summaries and ground truth in the semantic space of BERT (Devlin et al., 2019). Thus, it correlates with human judgments.

Baselines. We adopt the following methods, which are prevalently used in the experiments of

²<https://github.com/fxsjy/jieba>

³<https://huggingface.co/fnlpl/bart-base-chinese>

Category	Method	R-1	R-2	R-L	B-1	B-2	B-3	B-4	S-B	M	BS
Home Appliances	Lead	22.15	9.46	19.59	26.08	19.30	14.49	11.17	10.58	35.95	63.34
	Seq2seq	26.88	7.98	20.91	39.32	25.53	17.01	11.48	9.66	21.44	65.26
	PG	27.67	9.24	23.67	41.58	28.04	20.01	14.76	13.19	25.44	64.99
	MMPG	32.88	11.88	21.96	36.20	19.07	10.40	5.54	2.90	16.47	65.13
	VG-BART (Multi-head)	32.10	11.72	25.15	48.34	34.25	24.92	18.59	16.83	29.58	67.73
	VG-BART (Dot-product)	32.07	11.53	<u>25.46</u>	48.25	33.87	24.45	<u>18.13</u>	16.47	29.52	67.72
	V2P	34.47	12.63	25.09	<u>52.23</u>	36.49	25.30	17.81	15.24	30.31	68.76
	MMAPS	34.91	13.20	26.44	54.22	38.65	28.26	21.22	19.66	<u>32.32</u>	69.76
	Improvement	1.28%	4.51%	3.85%	3.28%	11.31%	16.72%	15.76%	16.82%	-10.10%	1.45%
Clothing	Lead	19.84	7.13	17.77	16.78	10.76	7.18	5.10	4.01	28.41	60.61
	Seq2seq	31.33	9.88	23.02	35.71	21.84	14.29	9.54	7.43	23.46	68.22
	PG	30.92	9.66	24.31	36.28	21.90	14.08	9.22	7.39	26.07	67.96
	MMPG	30.73	10.29	21.25	28.32	11.69	5.15	0.57	2.24	14.47	66.07
	VG-BART (Multi-head)	31.32	10.80	<u>23.49</u>	36.74	23.29	15.58	<u>10.81</u>	8.69	27.23	67.79
	VG-BART (Dot-product)	30.18	10.40	22.82	35.30	22.37	15.09	10.59	8.36	26.32	66.81
	V2P	35.05	11.98	22.62	42.41	<u>25.04</u>	<u>15.70</u>	9.16	6.55	27.96	<u>68.40</u>
	MMAPS	34.05	<u>11.52</u>	25.09	<u>41.43</u>	26.05	17.17	11.62	9.26	<u>28.30</u>	69.57
	Improvement	-2.85%	-3.84%	10.92%	-2.31%	4.03%	9.36%	7.49%	6.56%	-0.39%	1.71%
Cases & Bags	Lead	20.15	7.32	18.26	17.93	11.71	7.89	5.60	4.56	<u>29.23</u>	61.05
	Seq2seq	28.59	8.07	20.60	23.64	14.11	8.85	5.74	4.38	17.79	68.16
	PG	32.18	9.73	24.73	40.06	23.97	15.00	9.58	7.70	26.43	68.14
	MMPG	32.69	11.78	22.27	30.39	12.46	5.26	2.14	0.49	14.93	65.44
	VG-BART (Multi-head)	30.76	10.46	24.88	38.16	24.31	16.22	11.30	9.39	25.81	67.11
	VG-BART (Dot-product)	31.21	10.74	<u>25.13</u>	38.61	24.71	16.27	11.09	9.10	26.39	67.81
	V2P	<u>34.65</u>	<u>11.89</u>	<u>24.53</u>	<u>43.88</u>	<u>27.32</u>	16.64	10.17	7.40	28.25	67.74
	MMAPS	35.08	12.08	25.60	44.35	28.20	18.45	12.45	10.57	29.37	68.95
	Improvement	1.24%	1.60%	1.87%	1.07%	3.22%	10.88%	10.18%	12.57%	0.48%	1.16%

Table 2: Performance of all methods on three product categories. The best results are shown in bold and the second-best results are underlined. The percentages of the improvement are obtained by comparing MMAPS with the best baseline. ROUGE, BLEU, SENTENCE-BLEU, METEOR and BERTScore are denoted by R, B, S-B, M and BS, respectively.

existing product summarization works (Li et al., 2020a; Song et al., 2022), as baselines in our experiments:

- **Lead** (Li et al., 2020a; Song et al., 2022) directly extracts the first 80 characters of the long product descriptions as product summaries.
- **Seq2seq** (Khatri et al., 2018) is the standard neural architecture used for text generation. It takes the long product descriptions as input and outputs corresponding product summaries.
- **Pointer-Generator** (See et al., 2017) is a hybrid method consisting of the pointer network and the Seq2seq architecture. The pointer network helps reproduce input in the generated summaries by copying words from the long product descriptions.
- **MMPG**⁴ (Li et al., 2020a) is a multi-modal dual-encoder pointer-generator network for product summarization, where the convolutional neural networks are used to encode product images.
- **VG-BART (Dot-product) and VG-BART (Multi-head)**⁵ (Yu et al., 2021) adopt BART as the backbone for summarization, and use dot-product based fusion and multi-head based fusion to inject visual information, respectively.
- **V2P**⁶ (Song et al., 2022) adopts BART as the backbone. It enhances product summarization

by predicting product attributes and using attribute prompts extracted from product images.

4.2. Overall Performance (RQ1)

We report the results in Tab. 2. From the results, we can observe that:

1. MMAPS achieves the best results in most cases, especially outperforming other methods by a large margin on BLEU metrics. VG-BART, V2P and MMAPS all use the same PLM (BART) as the backbone, but MMAPS performs much better than others in terms of most metrics. For example, our method exceeds the best baseline by 15.76% and 16.82% for BLEU-4 and S-BLEU on Home Appliances, respectively. The results indicate that MMAPS consistently generates high-quality product summaries that are closer to human-written summaries.
2. In a few cases, baselines outperform MMAPS. However, they do not show robust performance as MMAPS. For example, Lead shows a superior performance on the METEOR metric, which considers the number of chunks. Moreover, as explained by the authors of V2P (Song et al., 2022), V2P obtains better scores on Rouge-1 and Rouge-2 on the Clothing category (the performance of MMAPS is 2.85% and 3.84% worse than V2P on Rouge-1 and Rouge-2) since human-written summaries in the Clothing category are more likely to contain the vision-related attributes. However, compared to MMAPS, both Lead and V2P do not show

⁴<https://github.com/hrlinlp/cepsum>

⁵<https://github.com/HLTCHKUST/VG-GPLMs>

⁶https://xuemengsong.github.io/V2P_

Code.rar


Product Detail Information:	Product Summary
<p>播凉夏汽水糖 2019夏新品女装 纯棉连衣裙 女撞色条纹 抽绳收腰显瘦 中长款仙女裙 鸢尾蓝条, 不对称, 撞色, 五分袖, 简约, 中长裙, 落肩袖, 文艺, 套头, 棉, 衬衫裙, 方领, 宽松, 品名: 播凉夏汽水糖, 背面较软, 袖肥, 袖长适中, 不同的条纹配色, 新颖亮眼. 宽松撞色条纹可收腰连衣裙, 严肃感的衬衫上衣, 利落的中性美...</p> <p>Soda cool summer soda sugar 2019 summer new women's cotton dress women's contrasting color stripes drawstring waist slimming mid-length fairy dress iris blue stripes, asymmetrical, contrasting colors, five-quarter sleeves, simple, mid-length skirt, dropped shoulder sleeves, art, pullover, cotton, shirt skirt, square collar, loose, product name: Boliangxia soda candy, soft back, fat sleeves, moderate sleeve length, different stripes and colors, novel and eye-catching. Loose contrasting striped waist dress, serious shirt upper body, neat neutral beauty...</p>	<p>GT Summary: 这款宽松撞色条纹可收腰连衣裙, 严肃感的衬衫上衣, 带出了利落的中性美。不对称的撞色线条跳跃活力, 打破视觉冲击, 具有感染力。中腰两边抽拉式抽绳增加趣味, 张弛有度间表达心境。 This loose, waist-contrast striped dress is paired with a serious shirt top for sharp and neutral beauty. The asymmetrical contrasting color lines are energetic, break the visual impact, and are infectious. The drawstring design on both sides of the middle waist add interest, and express the state of mind with a certain degree of relaxation.</p> <p>V2P: 宽松柔软舒适的面料, 不对称的撞色线条, 时尚简约的版型, 刺绣的T恤图案, 吸汗透气, 印花的圆领设计, 白色的飘带点缀在胸前, 增添了几分浪漫气息, 让人眼前一亮。适合日常穿着! Loose, soft and comfortable fabric, asymmetric contrasting color lines, stylish and simple version, embroidered T-shirt pattern, sweat-absorbing and breathable, printed round neck design, white streamers dotted on the chest, adding a bit of romantic atmosphere, let people's eyes light up. Perfect for everyday wear! (ROUGE-I: 20.15%, C-BLEU: 16.35%, METEOR: 23.62%)</p> <p>MMAPS: 这款连衣裙采用了撞色的条纹元素, 新颖亮眼。不对称的撞色线条跳跃活力。中腰两边抽拉式抽绳的设计, 可以自由调节松紧, 穿着舒适, 宽松的衬衫上身, 利落的中性美, 具有感染力。 This dress uses color-contrast striped elements, which are novel and eye-catching. The asymmetrical contrasting color lines are energetic. The drawstring design on both sides of the middle waist can be adjusted freely. It is comfortable to wear. The loose shirt top with sharp and neutral beauty is contagious. (ROUGE-I: 51.61%, C-BLEU: 49.11%, METEOR: 61.05%)</p> <p>MMAPS-w/o-Image: 播撞色条纹连衣裙, 臻选优质舒软面料材质, 赋予整款连衣裙更好的质感与格调, 穿着更显气质与档次。条纹设计优雅大气, 倍显女士的知性。条纹收腰设计, 倍添女士性感。 The contrasting color striped dress is made of high-quality and soft fabric materials, which endows the whole dress with better texture and style, and shows more temperament and class when worn. The striped design is elegant and majestic, showing the intellectuality of women. The striped waist design adds feminine sexiness. (ROUGE-I: 23.15%, C-BLEU: 8.88%, METEOR: 19.79%)</p>
<p>Product Image:</p> 	

Figure 3: A comparison between the product summaries generated by MMAPS and V2P. GT indicates ground truth. The English texts are translated from the corresponding Chinese texts. The same or semantically similar descriptions are highlighted in red. The product appearances manifested in the image are underlined.

- consistently good performance on all metrics across the three categories.
- We also find that V2P obtains better results than other baselines on unigram and bigram metrics (R-1, R-2, B-1 and B-2) while it is inferior to VG-BART and MMAPS on trigram, 4-gram and sentence-level metrics (R-L, B-3, B-4 and S-B). The results indicate that, compared to VG-BART and MMAPS, V2P cannot generate summaries containing adjectives and nouns with relatively more words (e.g., 4-gram text), making the generated summaries of V2P less coherent than those generated by MMAPS. This is also observed in the case study reported in Sec. 4.5.
- ### 4.3. Ablation Study (RQ2)
- To investigate the contribution of each component in MMAPS, we design the following variants of MMAPS for the ablation study:
- MMAPS-w/o-MRM:** To verify the importance of the masked region modeling task, we remove it by setting $\lambda_1 = 0$ in Eq. 10.
 - MMAPS-w/o-CMM:** To show the effect of coarse-grained multi-modal modeling, we take this task out by setting $\lambda_2 = 0$ in Eq. 10.
- MMAPS-w/o-FMM:** To show the necessity of fine-grained multi-modal modeling, we exclude it by setting λ_3 to zero in Eq. 10. Due to page limit, we only show the ablation study results of the above methods over the Cases & Bags in Tab. 3. We have the following findings:
 - The mask region modeling task improves the quality of the generated summaries on all metrics, demonstrating that predicting the class of the masked regions is helpful to the image encoder and it enhances product summarization.
 - Introducing the coarse-grained multi-modal modeling exhibits notable benefits, showing that capturing the semantic correlation between two modalities at the sequence level can facilitate generating high-quality product summaries.
 - The two fine-grained multi-modal modeling tasks show a positive impact on the model performance, suggesting that refining the fine-grained token-level and region-level representations and capturing their mutual influence are beneficial to product summary generation.
- ### 4.4. Impact of Task Weights (RQ3)
- We inspect the sensitivity of MMAPS to task weights in Eq. 10 over the category Cases & Bags.




Figure 4: Sensitivity analysis of task weights on Cases & Bags.

Model	R-1	R-2	R-L	B-1	B-2	B-3	B-4	S-B	M	BS
MMAPS	35.08	12.08	25.60	44.35	28.20	18.45	12.45	10.57	29.37	69.76
MMAPS-w/o-MRM	34.44	11.35	24.51	44.20	27.58	17.44	11.19	8.79	28.34	69.74
MMAPS-w/o-CMM	33.88	10.32	24.65	43.10	25.89	15.95	10.11	7.96	27.25	68.63
MMAPS-w/o-FMM	34.05	11.76	24.81	44.55	27.88	18.17	12.30	9.84	27.13	69.40

Table 3: The results of ablation studies on Cases & Bags. The best results are shown in bold.

We vary the values from 0 to 1 at the step of 0.1. From Fig. 4 we can see that MMAPS performs worst when setting one of the three task weights to zero. This confirms the importance of the masked region modeling task, the coarse-grained multi-modal modeling and the fine-grained multi-modal modeling to the overall performance of MMAPS. Besides, we observe that MMAPS performs relatively stably when the three task weights are non-zero. This implies that MMAPS is not sensitive to different task weights as long as they are non-zero.

4.5. Case Study (RQ4)

Fig. 3 provides a case study for better understanding the difference between MMAPS and V2P, which is the best baseline in most cases. From Fig. 3, we have the following observations:

1. V2P cannot generate grammatically corrected summaries and it produces many short phrases containing only product attributes. This phenomenon explains why V2P shows good performance on BLEU-1 in Tab. 2, but performs poorly when tested using n-gram based metrics with a large n (e.g., BLEU-3, BLEU-4 and S-BLEU). Differently, MMAPS can generate more coherent and descriptive product summaries containing rich product attribute information described by n-grams with a large n . For example, the generated summary of MMAPS contains several n-grams with a relatively large n , e.g., “不对称的撞色线条跳跃活力” (“The asymmetrical contrasting color lines are energetic”) and “中腰两边抽拉式抽绳” (“the drawstring design on both sides of the middle waist”). And these n-grams are also contained in the human-written summary. This observation shows that MMAPS can capture the prod-

uct attributes which are better described using n-grams with large n values. Therefore, MMAPS can generate more coherent and readable summaries as discussed in Sec. 4.2.

2. From the generated summaries, we can see that MMAPS is able to model the product appearances manifested in the image (i.e., the underlined parts: “asymmetrical”, “shirt top” and “The drawstring design on both sides of the middle waist”), while MMAPS-w/o-Image cannot. Hence, the image modality indeeds enhances the product summarization by providing additional signal to guide the model to generate the attractive summaries.

5. Conclusion

In this paper, we propose MMAPS for multi-modal product summarization and it is able to model product attributes and produce coherent and descriptive summaries simultaneously. Our experiments show that MMAPS exceeds state-of-the-art product summarization methods. In the future, we plan to incorporate cross-grained contrast learning, i.e., the contrast between coarse-grained representations and fine-grained representations. We will also consider generating personalized product summarization based on user preferences and affinity towards different product characteristics.

Acknowledgments

This work was partially supported by National Science and Technology Major Project (No. 2022ZD0118201) and National Natural Science Foundation of China (No. 62002303, 42171456).

References

- Jean-Baptiste Alayrac, Adrià Recasens, Rosalia Schneider, Relja Arandjelovic, Jason Ramapuram, Jeffrey De Fauw, Lucas Smaira, Sander Dieleman, and Andrew Zisserman. 2020. Self-supervised multimodal versatile networks. In *NeurIPS*.
- Satanjeev Banerjee and Alon Lavie. 2005. ME-TEOR: an automatic metric for MT evaluation with improved correlation with human judgments. In *IEEevaluation@ACL*, pages 65–72.
- Qibin Chen, Junyang Lin, Yichang Zhang, Hongxia Yang, Jingren Zhou, and Jie Tang. 2019. Towards knowledge-based personalized product description generation in e-commerce. In *KDD*, pages 3040–3050.
- Jaemin Cho, Jie Lei, Hao Tan, and Mohit Bansal. 2021. Unifying vision-and-language tasks via text generation. In *ICML*, volume 139 of *Proceedings of Machine Learning Research*, pages 1931–1942.
- Vijay Daultani, Lasguido Nio, and Young-joo Chung. 2019. Unsupervised extractive summarization for product description using coverage maximization with attribute concept. In *ICSC*, pages 114–117.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*, volume 1, pages 4171–4186.
- Xiaojie Guo, Qingkai Zeng, Meng Jiang, Yun Xiao, Bo Long, and Lingfei Wu. 2022. Automatic controllable product copywriting for e-commerce. In *KDD*, pages 2946–2956.
- Chandra Khatri, Gyanit Singh, and Nish Parikh. 2018. [Abstractive and extractive text summarization using document context vector and recurrent neural networks](#). *arXiv Preprint*.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *ICLR (Poster)*.
- Bruno Korbar, Du Tran, and Lorenzo Torresani. 2018. Cooperative learning of audio and video models from self-supervised synchronization. In *NeurIPS*, pages 7774–7785.
- Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannic Kalantidis, Li-Jia Li, David A. Shamma, Michael S. Bernstein, and Li Fei-Fei. 2017. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *Int. J. Comput. Vis.*, 123(1):32–73.
- Haoran Li, Peng Yuan, Song Xu, Youzheng Wu, Xiaodong He, and Bowen Zhou. 2020a. Aspect-aware multimodal summarization for chinese e-commerce products. In *AAAI*, pages 8188–8195.
- Junnan Li, Ramprasaath R. Selvaraju, Akhilesh Gotmare, Shafiq R. Joty, Caiming Xiong, and Steven Chu-Hong Hoi. 2021a. Align before fuse: Vision and language representation learning with momentum distillation. In *NeurIPS*, pages 9694–9705.
- Xiaoni Li, Yu Zhou, Yifei Zhang, Aoting Zhang, Wei Wang, Ning Jiang, Haiying Wu, and Weiping Wang. 2021b. Dense semantic contrast for self-supervised visual representation learning. In *ACM Multimedia*, pages 1368–1376.
- Xiujun Li, Xi Yin, Chunyuan Li, Pengchuan Zhang, Xiaowei Hu, Lei Zhang, Lijuan Wang, Houdong Hu, Li Dong, Furu Wei, Yejin Choi, and Jianfeng Gao. 2020b. Oscar: Object-semantics aligned pre-training for vision-language tasks. In *ECCV*, volume 12375, pages 121–137.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81.
- Lilang Lin, Sijie Song, Wenhan Yang, and Jiaying Liu. 2020. MS2L: multi-task self-supervised learning for skeleton based action recognition. In *ACM Multimedia*, pages 2490–2498.
- Xiao Liu, Fanjin Zhang, Zhenyu Hou, Li Mian, Zhaoyu Wang, Jing Zhang, and Jie Tang. 2023. Self-supervised learning: Generative or contrastive. *IEEE Trans. Knowl. Data Eng.*, 35(1):857–876.
- NLTK. 2023. nltk.translate.bleu_score module. https://www.nltk.org/api/nltk.translate.bleu_score.html. Accessed: 2023-05-01.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *ACL*, pages 311–318.
- Shaoqing Ren, Kaiming He, Ross B. Girshick, and Jian Sun. 2015. Faster R-CNN: towards real-time object detection with region proposal networks. In *NIPS*, pages 91–99.
- Abigail See, Peter J. Liu, and Christopher D. Manning. 2017. Get to the point: Summarization with pointer-generator networks. In *ACL (1)*, pages 1073–1083.

- Changchong Sheng, Matti Pietikäinen, Qi Tian, and Li Liu. 2021. Cross-modal self-supervised learning for lip reading: When contrastive learning meets adversarial training. In *ACM Multimedia*, pages 2456–2464.
- Xuemeng Song, Liqiang Jing, Dengtian Lin, Zhongzhou Zhao, Haiqing Chen, and Liqiang Nie. 2022. V2P: vision-to-prompt based multi-modal product summary generation. In *SIGIR*, pages 992–1001.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *NIPS*, pages 5998–6008.
- Jinpeng Wang, Yutai Hou, Jing Liu, Yunbo Cao, and Chin-Yew Lin. 2017. A statistical framework for product description generation. In *IJCNLP*, volume 2, pages 187–192.
- Tiezheng Yu, Wenliang Dai, Zihan Liu, and Pascale Fung. 2021. Vision guided generative pre-trained language models for multimodal abstractive summarization. In *EMNLP*, pages 3995–4007.
- Peng Yuan, Haoran Li, Song Xu, Youzheng Wu, Xiaodong He, and Bowen Zhou. 2020. On the faithfulness for e-commerce product summarization. In *COLING*, pages 5712–5717.
- Jianguo Zhang, Pengcheng Zou, Zhao Li, Yao Wan, Xiuming Pan, Yu Gong, and Philip S. Yu. 2019. Multi-modal generative adversarial network for short product title generation in mobile e-commerce. In *NAACL-HLT (2)*, pages 64–72.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. Bertscore: Evaluating text generation with BERT. In *ICLR*.