

User-Inspired Posterior Network for Recommendation Reason Generation

Haolan Zhan, Hainan Zhang, Hongshen Chen, Lei Shen, Yanyan Lan, Zhuoye Ding, Dawei Yin

State Key Lab. of Computer Science, Institute of Software, Chinese Academy of Sciences

Institute of Computing Technology, Chinese Academy of Sciences

University of Chinese Academy of Sciences

Data Science Lab, JD.com, Baidu Inc.

zhanhl@ios.ac.cn, zhanghainan@jd.com, ac@chenhongshen.com

{shenlei17z, lanyanyan}@ict.ac.cn, dingzhuoye@jd.com, yindawei@acm.org

ABSTRACT

Recommendation reason generation, aiming at showing the selling points of products for customers, plays a vital role in attracting customers' attention as well as improving user experience. A simple and effective way is to extract keywords directly from the knowledge-base of products, i.e., attributes or title, as the recommendation reason. However, generating recommendation reason from product knowledge doesn't naturally respond to users' interests. Fortunately, on some E-commerce websites, there exists more and more user-generated content (user-content for short), i.e., product question-answering (QA) discussions, which reflect user-cared aspects. Therefore, in this paper, we consider generating the recommendation reason by taking into account not only the product attributes but also the customer-generated product QA discussions. In reality, adequate user-content is only possible for the most popular commodities, whereas large sums of long-tail products or new products cannot gather a sufficient number of user-content. To tackle this problem, we propose a user-inspired multi-source posterior transformer (MSPT), which induces the model reflecting the users' interests with a posterior multiple QA discussions module, and generating recommendation reasons containing the product attributes as well as the user-cared aspects. Experimental results show that our model is superior to traditional generative models. Additionally, the analysis also shows that our model can focus more on the user-cared aspects than baselines.

CCS CONCEPTS

• **Information systems** → **Recommendation systems**; *Computing methodologies*; • **Computer systems organization** → Robotics; • **Computing methodologies** → Natural Language Generation.

KEYWORDS

Recommendation Reason; Posterior Network; Self-attention

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SIGIR '20, July 25–30, 2020, Xi'an, China

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-XXXX-X/20/06...\$15.00

<https://doi.org/10.1145/1122445.1122456>

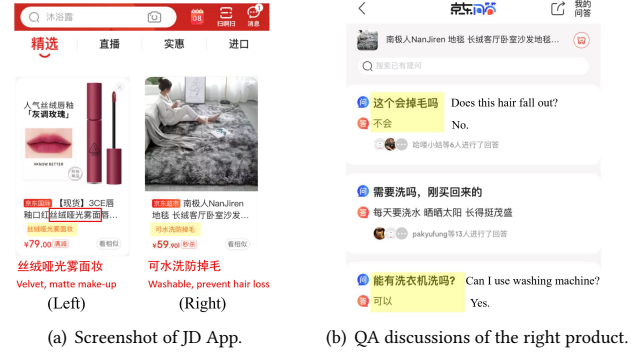


Figure 1: Two products with title and recommendation reason.

ACM Reference Format:

Haolan Zhan, Hainan Zhang, Hongshen Chen, Lei Shen, Yanyan Lan, Zhuoye Ding, Dawei Yin. 2020. User-Inspired Posterior Network for Recommendation Reason Generation. In *SIGIR '20: International ACM SIGIR Conference on Research and Development in Information Retrieval*, July 25–30, 2020, Xi'an, China. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/1122445.1122456>

1 INTRODUCTION

Nowadays, more and more customers are used to online-shopping based on the recommendation system rather than just search for a product directly. However, existing recommendation systems, just ranking the products, can no longer meet the requirements of customers. To save costs, most customers would like to browse products with the selling points in the limited mobile feeds. Thus, it is necessary to obviously display the selling points as the recommendation reasons to help customers learn the recommended products straightforwardly.

Furthermore, products with informative recommendation reasons can also attract customers' attention and receive more clicks. From screenshots in Figure 1(a), we can see that the red line, showing the informative characteristics of the product, i.e., “可水洗防掉毛(Washable, prevent hair loss)”, can arouse customers' interests and encourage them to click to see the details. Based on our online analysis, we find that products with recommendation reasons can receive more clicks than those without. Such effects of informative display have also been discussed in previous studies[7, 10], showing that when generated properly, the recommendation reasons will increase the profits of the recommendation system.

Given the product information, such as title, attributes and relevant QAs, RRG task aims at generating selling points of the product. The explicit approach can be inspired by the Short Title Generation (STG) task [7] for E-commerce, which automatically generates short titles by extracting key information from item-content, i.e., product title and its attributes. The researchers propose to use the pointer-network to copy the brand name and the key information from the product title and attributes as the short title. Similarly, RRG can be viewed as a modified form of the STG, with the emphasis on creating a natural reason from the product item-content by taking the key information as selling points. As shown in Figure 1(a), the example “丝绒哑光面妆 (Velvet, matte makeup)” is an item-content based recommendation reason. It is generated by extracting the selling points directly from title and attributes, which can be seen as an important and explicit method for the RRG task.

However, this item-content based approach is not concerned with the user’s experience and the most-cared questions, which may hurt the effectiveness of the generated reason. Obviously, the RRG should take more respect to the user-cared aspects rather than the item-content. Therefore, how to incorporate user-cared aspects into the recommendation reason is critical for RRG. To tackle this problem, we propose to use the user-content (QA discussions) as the most-cared aspects from the user’s experience. Taking the right product in Figure 1(b) as an example, some customers post their questions about the product and others answer these questions. The recommendation reason can be summarized as “可水洗不掉毛 (Washable, prevent hair loss)”. Therefore, the most-cared aspect of the product can be mined from the user-content.

In reality, adequate user-content is only possible for the most popular commodities, whereas large sums of long-tail products or new products can not gather sufficient amount of product user-content. In this paper, we propose a novel model named Multi-Source Posterior Transformer (MSPT) to explicitly model these two aspects and take into account the user-content with an adaptive posterior module. Specifically, for the item-content, we first use the Bi-LSTM encodes the title and attributes, and then use the self-attention [9] to obtain the title representation and the attribute representation. Besides, for the user-content, we utilize the product-guided attention module to encode the QA sentences into dense representation with a hierarchical method. Thirdly, we train the posterior network with feature fusion between QA pairs and products through minimizing the KL divergence of prior and posterior distributions. Finally, we use the posterior representation as input for decoding phase.

For experiments, we construct a large-scale JDRRG dataset containing 122,405 product items from JD App. We compare our model with several state-of-the-art generation baselines using both automatic and human evaluations. The results demonstrate that our model performs significantly better than existed methods. The main contributions of this paper include: (1) We introduce a new RRG task for recommendation. (2) We propose a novel multi-source posterior transformer network to model both the item-content and user-content. (3) We build a dataset and conduct rigorous experiments to demonstrate the effectiveness of our proposed model.

2 PROBLEM STATEMENT

Task Definition: Given a product title, attribute words and its relative question answering pairs, the RRG task aims to utilize inherent product information to identify the privilege of products and users’ preference, as well as to generate the coherent and appealing recommendation reasons. Formally, given a product title T composed of a sequence of words $\{t_1, \dots, t_N\}$, a set of attribute words $W = \{w_1, \dots, w_M\}$ and a relative question-answering set defined as $\{\{Q_1, A_1\}, \dots, \{Q_n, A_n\}\}$, the RRG task is to learn the generative model $G(\cdot)$. Each question in Q is defined as $Q_n = \{q_n^1, \dots, q_n^L\}$, where q_n^i is the i -th word in the sentence of Q_n and L is the max length of sentence. Each answer in A is defined as $A_n = \{a_n^1, \dots, a_n^J\}$, where a_n^j is the j -th word in the answer sentence. The corresponding recommendation reason is defined as $Y = \{y^1, \dots, y^S\}$, where y^i is the i -th word and S is the max length.

Dataset Collection: As there is no appropriate corpus for our RRG task currently, we build a new Chinese dataset, JDRRG. We collect our dataset from JD App, a Chinese E-commerce website. E-commerce is a popular way of shopping, where customers can post questions about products in public and other customers can give reply answers. JDRRG dataset consists of 122,405 products and relative 3.79 million pairs of QA discussions from real customers and 122,405 annotated recommendation reasons from our invited annotators (all professional writers). Each product has 30.9 different QA pairs on average. The average length of each question sentence, answer sentence, the product title, attributes and recommendation reason are 18.05, 18.59, 23.2, 16.03 and 10.5 respectively.

3 PROPOSED MODEL

As shown in Figure 2, our proposed MSPT model consists of an item-content self-attention module, a user-content QA attention module, an adaptive posterior module followed with decoding phase.

3.1 Item-content Self-Attention Module

3.1.1 Word-level Encoder. Given a product title $T = \{t_1, \dots, t_N\}$ as the input, a standard Bi-LSTM first maps the input sequence to a fixed-dimension vector $\mathbf{h}_T = [h_L, h_R]$. At the same time, we also use another Bi-LSTM to encode the attributes W as \mathbf{h}_W .

3.1.2 Title and Attributes Self-Attention. Given the word-level representation \mathbf{h}_T and \mathbf{h}_W of title and attributes, the initial input of the self-attention [9] is the sum of the word-level representation and position encoding (PE), denoting as $PE(\mathbf{h}_T)$ and $PE(\mathbf{h}_W)$, respectively. $\mathbf{E}_T^{(0)}$ and $\mathbf{E}_W^{(0)}$ are the initial input representations at the first layer. Self-attention is a special attention mechanism in a normal Transformer architectures [9]. At the l -th layer, the output representation is defined as below:

$$\mathbf{E}_T^{(l)} = \text{FFN}(\text{MHA}(\mathbf{E}_T^{(l-1)}, \mathbf{E}_T^{(l-1)}, \mathbf{E}_T^{(l-1)})), \quad (1)$$

$$\mathbf{E}_W^{(l)} = \text{FFN}(\text{MHA}(\mathbf{E}_W^{(l-1)}, \mathbf{E}_W^{(l-1)}, \mathbf{E}_W^{(l-1)})), \quad (2)$$

where $\mathbf{E}_T^{(l)}$, $\mathbf{E}_W^{(l)}$ denote the output representations after the l -th layer. The sub-layer FFN is a position-wise fully connected feed-forward network and $\text{MHA}(\mathbf{Q}, \mathbf{K}, \mathbf{V})$ [9] is a multi-head attention function. The final product representation is defined as:

$$\mathbf{H}_{item} = \lambda_1 \mathbf{E}_T^{(N)} + (1 - \lambda_1) \mathbf{E}_W^{(N)}, \quad (3)$$

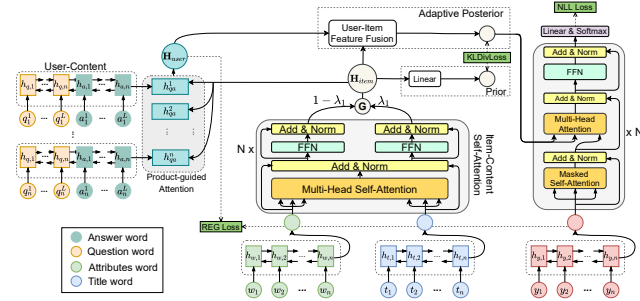


Figure 2: The framework of our proposed Multi-Source Posterior Transformer Network (MSPT).

where λ_1 is a parameter, and $E_T^{(N)}$ and $E_W^{(N)}$ is the final representation of title and attributes that output from the N -th layer.

3.2 User-content QA Attention Module

The user-content QA attention module consists of the word-level encoder for each QA sentence and context-level encoder for multiple QA pairs with the product-guided attention mechanism. The product-guided attention (PGA) mechanism examines all the relationships to highlight the most relevant QA pairs to the product. This type of attention is guided by the current product representation which is enhanced with corresponding information from the QA discussions. For each question sentence Q_n and answer sentence A_n in product a given product, we concatenate both sentences as input, and use the Bi-LSTM encoder to obtain its low-dimension representation h_{QA}^n , the same in section 3.1.1. Given the QA representation $h_{QA} = \{h_{QA}^1, h_{QA}^2, \dots, h_{QA}^n\}$, the context-level representation for multiple QA pairs is defined as a weighted sum of these h_{QA}^k :

$$H_{user} = \sum_{k=1}^n \delta_k^{qa} h_{QA}^k. \quad (4)$$

where weight δ_k^{qa} of each representation h_{QA}^k is defined as:

$$\delta_k^{qa} = \text{softmax}((W_1 H_{item} + b_1)^T (W_2 h_{QA}^k + b_2)). \quad (5)$$

where W_1, b_1, W_2 and b_2 are learned parameters. H_{item} is the output of the item-content self-attention module in section 3.1.2. we obtain the final representation of the multiple QA discussion H_{user} .

3.3 Adaptive Posterior Module

Given the user-content representation H_{user} and item-content representation H_{item} , the goal of adaptive posterior module is to enhance our model to learn the user-cared aspect from QA discussions. The adaptive posterior module consists of two sub-modules: a prior module and a posterior module. In the prior module, the product representation is updated by a linear transform module:

$$\tilde{H}_{prior} = \sigma(W_3 H_{item} + b_3). \quad (6)$$

In the posterior module, we define the posterior representation by adopting the user-item feature fusion mechanism. We concatenate the item-content representation H_{item} and the user-content representation H_{user} at first. And then, we adopt a linear transform to get the final representation, named hard fusion:

$$\tilde{H}_{post} = \sigma(W_4 [H_{item}, H_{user}] + b_4). \quad (7)$$

Model	RG-1	RG-2	RG-L	BU-1	BU-2	MTR
TextRank	8.04	2.19	7.32	7.89	1.74	2.03
SEQ2SEQ	10.60	2.86	8.43	8.38	1.91	2.35
PG-BiLSTM	12.32	3.59	9.35	10.54	2.35	3.25
MS-Ptr	13.45	4.41	10.65	11.24	2.69	3.98
Transformer	11.21	3.46	9.68	9.89	2.01	3.47
CopyTrans	12.37	4.39	10.50	10.95	2.24	4.04
HierTrans	11.94	3.80	9.96	10.58	2.36	3.94
Hi-MAP	13.85	4.65	11.48	11.03	2.85	4.31
MSPT-QA	13.54	4.23	10.59	10.94	2.47	3.93
MSPT-PGA	15.03	4.89	12.24	12.06	2.84	4.53
MSPT (hardFS)	15.97	5.62	12.76	13.13	3.26	4.97
MSPT (softFS)	16.52	5.83	13.38	15.69	3.52	5.16

Table 1: Metric-based evaluation on JDRRG dataset (%).

Besides, we also provide a soft fusion choice while employing the gate mechanism to fuse H_{item} and H_{user} .

$$\tilde{H}_{post} = \sigma(W_5 (\lambda_2 H_{item} + (1 - \lambda_2) H_{user}) + b_5). \quad (8)$$

Clearly, the discrepancy between prior and posterior distributions introduces great challenges in training the model. To tackle this problem, we introduce the KL divergence loss, to measure the proximity between the prior representation and the posterior representation. The KL divergence is defined as follows:

$$\mathcal{L}_{KL}(\theta) = D_{KL}(p(y_t | \tilde{H}_{post}) || p(y_t | \tilde{H}_{prior}; \theta)), \quad (9)$$

where θ denotes the model parameters.

3.4 Training Objectives

We employ three different loss functions: KL divergence, NLL and Regular loss. All loss functions are also elaborated in Figure 2.

NLL Loss. The objective of NLL loss is to quantify the difference between the ground truth reason and the generated reason by our model. It minimizes Negative Log-Likelihood (NLL):

$$\mathcal{L}_{NLL}(\theta) = - \sum_{t=1}^S \log p(y_t | y_{<t}, T, W, \{Q_i, A_i\}_{i=1}^n; \theta), \quad (10)$$

where θ denotes the model parameters and $y_{<t}$ denotes the previously generated words.

Regular Loss. It aims to enforce the relevancy between the QA discussion representation and the true reason. Then, the Regular loss is defined to minimize:

$$\mathcal{L}_{REG}(\theta) = - \sum_{t=1}^S \log p(y_t | H_{user}; \theta). \quad (11)$$

our final loss is a combination of these three loss functions:

$$\mathcal{L}(\theta) = \mathcal{L}_{KL}(\theta) + \mathcal{L}_{NLL}(\theta) + \mathcal{L}_{REG}(\theta). \quad (12)$$

4 EXPERIMENTS

4.1 Experimental Settings

Baseline Approaches: We choose the following baseline approaches to compare with our model: **TextRank** is an extraction method using a graph-based ranking algorithm [5]. **SEQ2SEQ** is a basic end-to-end model based on standard LSTM units and the attention mechanism [8]. **PG-BiLSTM** is the bi-directional LSTM with pointer generator mechanism [6]. **MS-Ptr** represents the multi-source pointer network for short product title generation [7]. **CopyTrans** is the

Model	MSPT vs.			kappa
	Win	Loss	Tie	
PG-BiLSTM	57%	14%	29%	0.475
MS-Ptr	43%	21%	36%	0.491
CopyTrans	46%	28%	26%	0.427
HierTrans	51%	19%	30%	0.528
Hi-MAP	40%	32%	28%	0.443

Table 2: Human evaluation of MSPT on JDRRG dataset.

Transformer model with the copy mechanism [3]. **HierTrans** represents the hierarchical transformer for multi-document summarization (MDS) tasks [4]. **Hi-MAP** is the hierarchical network with MMR-self-attention mechanism for MDS task [1].

Model Variants: To evaluate the product-guided attention module, User-Item feature fusion module and posterior module, we also employ some degraded MSPT models to investigate the effect of our proposed mechanisms: **MSPT-QA** removes the QA discussions and relevant product-guided attention and posterior module during training process. **MSPT-PGA** is MSPT model without the product-guided attention mechanism in user-content QA attention module. **MSPT (hardFS)** is the MSPT model with hard feature fusion. **MSPT (softFS)** is the MSPT model with soft feature fusion.

Evaluation Measures: We employ both automatic and human evaluation to measure the quality of generated recommendation reason. In automatic metrics, we report ROUGE-1/2/L (RG for short), BLEU-1/2 (BU for short) and METEOR (MTR for short). For human evaluation, we take the win-loss-tie evaluation method.

4.2 Results and Analysis

Metric-based Evaluation. The metric-based evaluation results are shown in Table 1. Our proposed MSPT model outperforms the best. Taking the ROUGE metrics as an example, the ROUGE-1 value of the MSPT is 16.52, which is significantly better than MS-Ptr, CopyTrans and Hi-MAP, i.e., 13.45, 12.37 and 13.85. The METEOR metrics of our model is also higher than other baseline models, indicating that our model can generate more informative and fluency recommendation reason. We also conducted a significant test, showing that the improvement is significant, i.e., $p\text{-value} < 0.01$.

Human Evaluation. The results of the human evaluation are shown in Table 2. The percentage of win, loss, and tie, as compared with the baseline models, is shown to evaluate the relevance and the informativeness of the generated reasons by MSPT. As showed in 2, the percentage of win is greater than the loss, indicating that our MSPT model is better than all baseline methods. MSPT achieves a preference gain (i.e., the win ratio minus the loss ratio) of 43%, 21%, 18%, 22% and 18% respectively, compared to PG-BiLSTM, MS-Ptr, CopyTrans, HierTrans and Hi-MAP, indicating that MSPT model can generate not only relevant but also informative reasons than baseline models. The Fleiss' kappa [2] value demonstrates the consistency of different annotators.

Ablation Study. As shown in Table 1, we also conduct an ablation study with model variants. Above all, to validate the enhancement of QA discussions, we remove this module and get the MSPT-QA model. The metric evaluation shows that the QA discussions perform very importantly. Besides, we also remove the product-guided attention module to get MSPT-PGA model. It also lost in the competition with MSPT(soft). Therefore, we can see that the product

Product Title	雅邦荔枝冻金箔果冻口红不易掉色唇膏女士口红 Alobon lychee gold foil jelly lipstick, not-easy-to-fade lady lipstick
Attributes	金箔, 布丁, 口红, 保湿 (gold foil, jelly, lipstick, moisturizing)
QA Pair 1	Q: 保湿怎样? 容易掉色吗? → A1: 保湿效果不错, 我朋友也买了。/// A2: 不掉色, 不沾杯。/// A3: 可以, 挺滋润的 Q: What about preserving moisture? Is the lipstick easy to fade? → A1: Pretty good, good moisturizing effect, my friends also bought it. /// A2: Don't fade, not sticky. /// A3: Great, it's moist.
QA Pair 2	Q: 这款口红有什么味道的? → A1: 淡淡的水果味, 挺好闻。/// A2: 樱桃味, 总体来说还是不错的。 Q: What kinds of smells does this lipstick have? → A1: Light fruit flavor, it smells good. /// A2: Cherry flavor, overall good.
Ground Truth	保湿不易掉色, 果冻质地更丝滑 Moisturizing, not easy to fade, silkier with jelly texture.
MS-Ptr	口感很清脆, 精心礼盒装 Very crisp taste, well-crafted gift pack
HierTran	持久不掉色, 清爽不油腻 Lasting, refreshing and non-greasy
Hi-MAP	保湿滋润不掉色, 包装不错易保湿 Moisturizing, no fading, good packaging, easy to moisturize
MSPT	保湿润唇不掉色, 水果口味质量佳 Moisturizing for lips, does not fade, fruit scent, and good quality

Table 3: Case study of MSPT and baselines on JDRRG.

title and attributes have a good effect on selecting the relevant information in QA discussions. Moreover, the different feature fusion mechanism also makes a difference. The MSPT model with soft fusion performs better than the model with hard fusion.

Case Study. To facilitate a better understanding of our model, we present some examples in Table 3. We select three baseline models: MS-Ptr, HierTrans and Hi-MAP, since these models outperform other baselines in metric-based evaluation. As shown in 3, our proposed MSPT model generates more aspects of product with considering QA discussions. For example, the recommendation reason generated by HierTrans and MS-Ptr, such as “精心礼盒装 (well-crafted gift pack)” and “清爽不油腻 (Refreshing and non-greasy)”, contains some of the incorrect information. The Hi-MAP can generate relevant and appropriate recommendation reason such as “保湿滋润不掉色 (Moisturizing, no fading)”. However, this reason sentence contains repetitive information, such as “保湿 (Moisturize)”. Our proposed MSPT model contains more characteristics, i.e., “水果口味 (fruit scent)” and “不掉色 (does not fade)”, since it utilizes the QA discussions and enhances the user-cared aspects while generating recommendation reasons.

5 CONCLUSION AND FUTURE WORK

In this paper, we proposed a challenging recommendation reason generation task in E-commerce. To tackle the problem, we proposed a novel multiple source posterior transformer to incorporate user-cared aspects into the recommendation reason. Besides, we constructed a new benchmark dataset JDRRG to study and evaluate this task. The experimental results demonstrated that our model could outperform baselines. In future work, we plan to further investigate the proposed model with emotion classification.

REFERENCES

- [1] Alexander R Fabbri and Irene et al. Li. 2019. Multi-News: A Large-Scale Multi-Document Summarization Dataset and Abstractive Hierarchical Model. In *ACL*.
- [2] J. L. Fleiss. 1971. Measuring nominal scale agreement among many raters. *Psychological bulletin* (1971).
- [3] S. Gehrmann, Y. Deng, and A. M. Rush. 2018. Bottom-up abstractive summarization. In *CoRR*.
- [4] Y. Liu and M. Lapata. 2019. Hierarchical Transformers for Multi-Document Summarization. In *ACL*.
- [5] R. Mihalcea and P. Tarau. 2004. Textrank: Bringing order into text. In *EMNLP*.
- [6] A. See, P. J. Liu, and C. D. Manning. 2017. Get To The Point: Summarization with Pointer-Generator Networks. In *ACL*.
- [7] F. Sun, P. Jiang, H. Sun, C. Pei, W. Ou, and X. Wang. 2018. Multi-source pointer network for product title summarization. In *CIKM*.
- [8] M. Tan, C. D. Santos, B. Xiang, and B. Zhou. 2015. Lstm-based deep learning models for non-factoid answer selection. In *CoRR*.
- [9] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin. 2017. Attention is all you need. In *NIPS*.
- [10] J. Zhang, P. Zou, Z. Li, Y. Wan, X. Pan, Y. Gong, and S. Y. Philip. 2019. Multi-Modal Generative Adversarial Network for Short Product Title Generation in Mobile E-Commerce. In *NAACL-HLT*.