

Automatic Scene-based Topic Channel Construction System for E-Commerce

Peng Lin*, Yanyan Zou*, Lingfei Wu, Mian Ma, Zhuoye Ding, Bo Long

JD.com, Beijing, China

{linpeng47,zouyanyan6,lingfei.wu,mamian,dingzhuoye,bo.long}@jd.com

Abstract

Scene marketing that well demonstrates user interests within a certain scenario has proved effective for offline shopping. To conduct scene marketing for e-commerce platforms, this work presents a novel product form, scene-based topic channel which typically consists of a list of diverse products belonging to the same usage scenario and a topic title that describes the scenario with marketing words. As manual construction of channels is time-consuming due to billions of products as well as dynamic and diverse customers' interests, it is necessary to leverage AI techniques to automatically construct channels for certain usage scenarios and even discover novel topics. To be specific, we first frame the channel construction task as a two-step problem, i.e., scene-based topic generation and product clustering, and propose an E-commerce Scene-based Topic Channel construction system (i.e., ESTC) to achieve automated production, consisting of scene-based topic generation model for the e-commerce domain, product clustering on the basis of topic similarity, as well as quality control based on automatic model filtering and human screening. Extensive offline experiments and online A/B test validates the effectiveness of such a novel product form as well as the proposed system. In addition, we also introduce the experience of deploying the proposed system on a real-world e-commerce recommendation platform.

1 Introduction

Recently, e-commerce platforms have become an indispensable part of people's daily life. Different from brick-and-mortar stores where salespersons can hold face-to-face conversations to promote products and even recommend more products related to customers' interests, most recommendation systems of e-commerce platforms, such as Taobao¹, mainly display individual products in

which users might be interested (Zhou et al., 2018, 2019), as listed in Figure 1 (indicated as Recommendation Flow Page). Recently, scene marketing has become a new marketing mode for product promotion where particular application scenarios (i.e., scene) are created to demonstrate product functions and highlight features correspondingly (Zhao, 2020), which is also paramount for e-commerce platforms to improve user experience during online shopping (Kang et al., 2019; Fu et al., 2019). A practical usage scenario of products can help users better understand product functions and features, and also allow the platform to exhibit more products that hit customer's specific interests, so that the user experience and click rate might be improved. However, scenes do not always help. For example, displaying all related products belonging to the same scene in the recommendation flow page might harm the user experience, since they tend to be homogeneous.

To achieve scene marketing in e-commerce platforms, this work presents a novel product form, scene-based topic channel, which consists of a list of diverse products belonging to the same scenario, together with two short phrases (or sentences) as the topic title summarizing the scene. Exemplified by Figure 1, one primary product of a channel and the associated scene topic title (highlighted with red box) are displayed in the recommendation flow page. If a user is interested in the primary product and clicks on it, the user is then redirected to the topic channel page where diverse products belonging to the same usage scenario are displayed. Existing ways to constructing scene-based topic channel mainly rely on expert knowledge and past experience of business operators in grouping products into different functional categories with certain scene topics (Mansell, 2002; Cooke and Leydesdorff, 2006; Fernandez-Lopez and Corcho, 2010). However, such methods are highly expensive with low efficiency and even impractical since there

*The first two authors made equal contributions. Correspond to Yanyan Zou.

¹<https://www.taobao.com/>



Figure 1: A screenshot of a scene-based topic channel on an e-commerce platform, with only four products due to limited space. Text with underline in the right-side Translation column are used to connect the translated words with associated parts in the topic channel.

are billions of products in the e-commerce platforms. Therefore, in this work, we propose an E-commerce Scene-based Topic Channel construction system (i.e., ESTC) to automatically construct such scene-based topic channels, where the task is framed as a two-step problem, i.e., scene-based topic generation and product clustering. One intuitive solution to obtaining scene topics is to make use of topic models (Blei et al., 2003; Roberts et al., 2013; Grootendorst, 2022) or techniques from extractive summarization (Basave et al., 2014; Wan and Wang, 2016), which are, however, restricted to assigning topics within a predefined limited candidate set, while there are often emerging scenes in the e-commerce fields. Thus, like Alokaili et al. (2020), we propose to generate scene-based topic titles for products, which allows to create novel topics not featured in the training set.

Nevertheless, in practice, the limitation of labeled data for training (around 5000 instances) hinders the generation quality of the model. On the other hand, we observe that generated topic titles, describing the same scenario, might be slightly different in formulation. Simply grouping products based on exact string match of generated topic titles results in channels with rare products. To address above issues, we first develop a pre-trained model in the e-commerce field to improve generation quality. Then, a semantic similarity based clustering method is designed to conduct product clustering to form the channel. Finally, to ensure the user

experience online, we further design a quality control module to strictly filter out undesired channels, such as inconsistent topic titles, or channels with irrelevant topic-product pairs. Our contributions are summarized as follows:

- A topic generation model in e-commerce field is proposed to generate scene-based topic titles for products, which is flexible to produce topics for emerging products and allows the system to discover novel scene topics.
- A semantic similarity based clustering method is designed to aggregate products with similar topic titles and form scene-based channels, which is able to improve the product diversity.
- A quality control module is designed to ensure the quality of the artificially constructed channels before they are released online.
- We introduce the overall architecture of deployed system where the ESTC has been successfully implemented into a real-world e-commerce platform.
- To the best of our knowledge, this is the first work on automatically constructing scene-based topic channel for scene marketing in e-commerce platforms.

2 Proposed Method

The development of the proposed ESTC system consists of three main parts, including scene-based topic generation for each products, scene-based

product clustering to aggregate products with similar topic titles, as well as the quality control module to ensure the quality of AI-generated channels. We also include a simple data augmentation module to discover weakly supervised data in order to improve the diversity of generated topic titles.

2.1 Scene-based Topic Generation

In this work, we propose to generate the scene-based topic titles for each product. To be specific, given input information $X = (x_1, x_2, \dots, x_{|X|})$ of a product P , including product’s title T , a set of attributes A and side information O obtained through optical character recognition techniques, paired with scene-based topic title $Y = (y_1, y_2, \dots, y_{|Y|})$, we aim to learn model parameters θ and estimate the conditional probability:

$$P(Y|X; \theta) = \prod_{t=1}^{|Y|} p(y_t|y_{<t}; X; \theta)$$

where $y_{<t}$ stands for all tokens in a scene title before position t (i.e., $y_{<t} = (y_1, y_2, \dots, y_{t-1})$).

Pretraining with E-commerce Corpus Pre-trained models (Radford et al., 2019; Devlin et al., 2019; Lewis et al., 2020; Raffel et al., 2020; Zou et al., 2020; Xue et al., 2021) have proved effective in many downstream tasks, however, most of which are developed on English corpora from general domains, such as news articles, books, stories and web text. In our scenario, we aim to produce topic titles in Chinese that summarize certain usage scenarios of products. Therefore, a model is required to understand the products through its associated information (such as title, semi-structured attributes) and generate scene-based topic titles, where we argue that the model should learn knowledge from e-commerce fields and thus propose to further pre-train models in domain (Gururangan et al., 2020). Specifically, besides the product title, attribute set as well as side information, we also collect the corresponding advertising copywriting of products from e-commerce platforms for the second phase of pre-training. We adopt the UniLM (Dong et al., 2019) with BERT initialization as backbone structure.

Recall that the product attributes A is a set without fixed order. We observe that input containing same attributes yet in different orders might results in different outputs. On the other hand, UniLM is an encoder-decoder shared architecture. To reinforce both the understanding and generation ability

of no-order input information, in addition to the original pre-training objectives of UniLM, we also propose two objectives to adapt the target domain:

- **Consistency Classification:** Given a product title-attributes pair, this task aims to classify if the two refer to the same product. For the positive example, the attributes and the title describe the same product and attributes are randomly concatenated as a sequence to introduce disorder noises. For the negative example, we randomly select attributes from a different product.
- **Sentence Reordering:** We split the product copywriting into pieces according to marks (such as comma and period). Such pieces are then shuffled and concatenated as a new text sequence. The model takes the shuffled sequence as input and learns to generate the original copywriting.

After the second phase of pre-training in the target e-commerce domain, we fine-tune the pre-trained model on the scene-based topic generation dataset.

2.2 Scene-based Product Clustering

One intuitive solution to constructing a scene-based topic channel is to group products with exactly the same generated topic titles. However, we observe there exists channels with similar topic titles, each of which merely contains several products, while we expect one channel has diverse products to ensure user experience. Therefore, we design a clustering module to aggregate products with semantically similar topic titles.

Topic Encoding To better learn scene-based topic representations and distinguish different topic titles, we take all topic titles from training set as input and employ the SimCSE (Gao et al., 2021) to further fine-tune the e-commerce pre-trained UniLM model in an unsupervised fashion. The embeddings of the last layer are used as the initialization for product clustering.

Product Clustering This module aims to group products with semantically similar topic titles into a cluster, in other words, a product list for a channel. Since we do not have prior knowledge of how many topic clusters the topic generation model would produce, we adopt the hierarchical clustering (Sahoo et al., 2006) where the number of kernels is

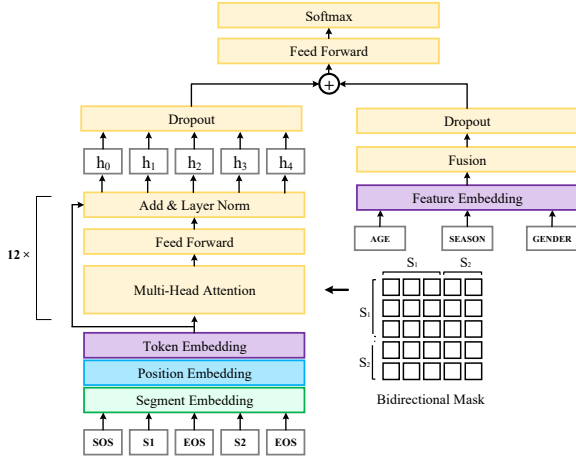


Figure 2: Correlation scoring model structure.

not required. To be specific, we adopt the bottom-up version, namely Agglomerative Nesting, which treats each sample as a leaf node and uses an iterative method for aggregation. In each iteration, two nodes with the highest similarity score are merged to form a new parent node. The iterative process stops when the shortest distance among all nodes is greater than a preset threshold. It is worthy noting each cluster might align with multiple topic titles and a list of products. The display order of products within a channel is decided by recommendation strategies, which is not focus of this work.

2.3 Quality Control

Although our method can generate good-quality channels most of the time, there is still possibility that the generated channels might not be accurate: 1) the generated topic title is semantically incoherent; 2) the topic title and associated products are not related according to the product usage scenarios. Thus, in order to alleviate above issues and ensure a reasonably good experience online, we design two modules, sentence coherence and correlation scoring models, to remove unexpected samples.

Topic Coherence Model We empirically observe that the generated topic titles might suffer coherent issues, like repetition and incompleteness. Thus, we design a topic coherence model to classify if a generated topic title is coherent. To be specific, the model is the e-commerce UniLM model with a softmax layer for classification. During training, we treat the online published topic titles as positive examples. The negative ones are synthesized:

- Samples with repetition: For a positive example of topic title, each unigram and bigram

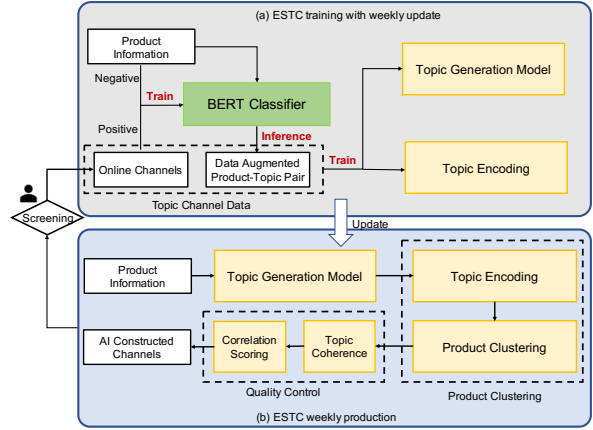


Figure 3: ESTC deployment on e-commerce platform.

is selected and repeated for one or two times with equal probability.

- Incomplete samples: We randomly remove the last two bigram or unigram tokens of a positive topic title.

We randomly select above synthesized samples to make the number of negative examples equal to the size of positive examples. Recall that a cluster might have multiple topic title candidates, the one with highest coherence score by the topic coherence model is used as final topic title. If all title candidates are classified as incoherent, then we simply remove such a cluster. After this module, each cluster is a scene-based topic channel with a list of products belonging to the same scene as well as a topic title summarizing the scene.

Correlation Scoring Model We design another binary classification model, i.e., correlation scoring model, to identify if the topic title and products are scene-based related. As illustrated in Figure 2, the e-commerce UniLM model takes as input the product information of a single product X as well as the generated topic title Y and determine whether they are related by the relevant scene. For better learning the product usage scenario, we also take into account the product profile information, such as age, season, and gender profiles, and employ a feed-forward layer to encode such features.

Likewise, product-topic title pairs from online published topic channels are considered as positive examples. The negative samples are obtained by randomly selecting mismatched product-topic title pairs. As a result, the number of negative examples is the same as the positive ones.

For each constructed channel, we use this model to check each product-topic title pair and remove

Model	SacreBLEU	ROUGE-1	ROUGE-2	ROUGE-L	BLEU	METEOR	DR(%)
BART	1.92	7.50	1.02	7.01	3.20	8.63	1.09
UniLM-BERT	2.05	7.87	1.11	7.42	3.45	8.70	0.87
E-commerce UniLM	2.08	8.01	1.12	7.56	3.47	8.78	0.88
E-commerce UniLM + DA	2.17	7.68	1.21	7.36	3.68	8.70	12.07

Table 1: The results of different topic generation model.

products that are unrelated to the topic.

2.4 Data Augmentation

Initially, the online existing (i.e., human-created) topic channels are quite limited which might hinder the model performance. Moreover, we would like to construct novel channels. Thus, we propose a UniLM-based binary classification model to discover more and diverse product-topic title candidate pairs. To be specific, the existing online product-topic title pairs are considered as positive examples. Similar to Zhang et al. (2022), a product with its the side information O from product detail images are considered as negative examples. After the classification model is trained, we use such a model to further extract more data for training. Negative examples with high probability scores are augmented into the training set.

3 Deployment

We have successfully deployed the proposed ESTC system on a real-world e-commerce platform. Figure 3 demonstrates the workflow of the deployed system with weekly update. Firstly, the data augmentation module is utilized to augment existing online channels. The augmented data is then used to train the topic generation and encoding models. Since there are thousands of millions products online, we weekly update the model and re-construct the channels to discover novel scene-based topic channels. To ensure a proper user experience, human screening is necessary before publishing channels online.

4 Experiment

4.1 Topic Generation Results

Data Collection The data for developing scene-based topic generation model consists of two sources: existing online channels (including human-created) and augmented samples, collected from a publicly available online e-commerce platform, JD.com². For the scene-based topic channels, we collected online channels from the product

²<https://www.jd.com/>

Dataset	#PT	#T	IL	OL
Human	177,412	5,186	69.34	13.44
Mined	111,572	82,834	74.21	12.54

Table 2: The statistics of topic generation dataset. #PT denotes the number of product-topic pairs, #T denotes the number of topic titles, IL denotes the average length of input product information sequence and OL denotes the average length of topic titles.

form ‘‘Goods List’’ from the platform, which were reviewed by human.

We also leveraged the optical character recognition (OCR) and classification techniques to extract key information about the product from product detail images. Firstly, texts are extracted from the images. Then, the extracted texts are ranked in descending order of their importance and relevance using the classification model. Finally, highly ranked texts are selected and merged as the final OCR input of products for topic generation.

In the end, we have 5,186 topic titles created by human and 82,834 topic title candidates from product side information. We further split the whole dataset into training, validation and test set with a ratio of 80%:10%:10%. The online channels are considered as ground-truth. Details are listed in Table 2. Moreover, we constructed product-OCR text (i.e., side information) pairs for the data augmentation module.

Comparison We use SacreBLEU (Post, 2018), ROUGE (Lin, 2004), BLEU (Papineni et al., 2002), and METEOR (Lavie et al., 2004) to measure quality of outputs by different generation models. We also design a new metric, difference rate (i.e., DR), to measure the novelty of generated topics, which is the ratio of the number of novel topics (i.e., not appearing in the training set) and the total number of generated ones. We consider publicly available models pre-trained on Chinese corpus as baselines, including BART (Shao et al., 2021) and UniLM with BERT initialization. As listed in Table 1, our E-commerce UniLM model achieves best performance for most evaluation metrics. With augmented data (denoted as +DA), the performance

Model	Silhouette	Recall	Precision	F1
B.O.W	0.264	90.8	88.7	88.0
Word2Vec	0.220	88.3	86.5	85.3
BERT	0.200	75.7	74.9	73.7
+SimCSE	0.283	88.9	89.0	88.1
E-commerce UniLM	0.248	72.3	68.7	68.0
+SimCSE	0.283	96.4	96.0	95.7

Table 3: The performance of different topic encoding models for clustering.

of our model is further improved with more novel topic titles produced, which shows the effectiveness of the data augmentation module.

4.2 Product Clustering Results

Dataset The clustering module works in an unsupervised fashion, while labeled data is still required for model evaluation. We manually create a data set for clustering evaluation, containing 65 different topic title samples, belonging to 18 groups.

Metrics We adopt the distance-based Silhouette Coefficient (Rousseeuw, 1987) to evaluate the performance of topic clustering. To investigate how well a clustering matches reference partitions of the test data, we further design two metrics.

For each topic sample i from cluster j , the precision score is calculated as:

$$P_{i,j} = \frac{TP_{i,j}}{N_j} \quad (1)$$

where $TP_{i,j}$ denotes the number of correctly grouped topic i in cluster j , N_j is the number of samples in cluster j . Similarly, the recall score is calculated as:

$$R_{i,j} = \frac{TP_{i,j}}{T_i} \quad (2)$$

where T_i is the total number of topic i found across all clusters.

The F1-measure score is computed as the harmonic mean of precision and recall.

Comparison We compare different sentence embedding-based clustering methods, including bag of words (i.e., B.O.W), Word2Vec (Mikolov et al., 2013), BERT as well as our E-commerce UniLM model. As listed in Table 3, models with SimCSE achieve better clustering performance. It is worthy noting that, the Silhouette score is not consistent with our designed metric scores. We practically observed that higher F1 scores indicate better clustering results for topics.

4.3 Quality Control Results

We also conducted human evaluation to investigate the effectiveness of each module for quality control, where for each setting, 1000 constructed channels are presented for human screening and report overall acceptance rate that is the ratio of the validated channels and the all candidates. As listed in Table 5, considering both topic coherence and correlation scoring modules results in the highest acceptance rate, demonstrating strengths of quality control module.

4.4 Experiments of Different Clustering Methods

As listed in Table 4, we compare two clustering method, K-means and hierarchical clustering methods, where the initial embedding are taken from different models. The hierarchical clustering with SimCSE enhanced e-commerce UniLM model achieves best performance.

4.5 Online A/B Test

To demonstrate the payoff generated by ESTC system, a standard A/B testing is conducted to evaluate the benefit of deploying scene-based topic channels on an e-commerce mobile app. After launching such a new product form, the Click-Through Rate (CTR) is improved by 3.20%, compared to the one without scene-based topic channels, which shows the values of AI-generated scene-based topic channels. We note that the comparison between human-created and AI-generated channels is difficult to fairly determine, since there are many factors mattering the online performance, such as recommendation strategies of products within channels.

More details about generated samples are included in Appendix.

5 Lessons Learned During Deployment

Several lessons we have learned during model deployment could be beneficial for other like-minded practitioners who wish to deploy cutting-edge AI technologies into real-world applications, such as the importance of real-world data quality and business understandings.

- Data quality matters model performance. Besides the model capacity, the quality of training data is of paramount importance. The cleaning procedures of raw data (e.g., removing poor samples from training set and speci-

Model	Kmeans				HC			
	Silhouette	Recall	Precision	F1	Silhouette	Recall	Precision	F1
B.O.W	0.221	81.9	74.3	72.4	0.264	90.8	88.7	88.0
Word2Vec	0.191	82.3	80.3	78.7	0.220	88.3	86.5	85.3
BERT	0.150	68.3	64.8	62.7	0.200	75.7	74.9	73.7
+SimCSE	0.262	90.0	86.2	86.6	0.283	88.9	89.0	88.1
E-commerce UniLM	0.210	73.5	69.0	68.4	0.248	72.3	68.7	68.0
+SimCSE	0.248	84.5	83.1	81.9	0.283	96.4	96.0	95.7

Table 4: The performance of different topic encoding models and different clustering models.

fying important attributes) plays a critical role in model development.

- Business understandings and logics advance AI model launching. The AI constructed scene-based topic channels are not fool proof. Thus, in order to ensure a reasonably good user experience, post-processing, based on insightful business understandings and logics, of AI constructed channels in the production platform is necessary to filter out any inconsistent or low-quality contents.

6 Related Work

Previous studies (Lau et al., 2011; Bhatia et al., 2016; Mei et al., 2007) on topic mining mainly first retrieve candidate topic labels from reference corpora and then conduct topic ranking to select the best topic label. Lau et al. (2010) simply take a word from a top-N terms as the topic label. Knowledge bases are also adopted to retrieve topic labels by matching candidate topic words to knowledge concepts (Magatti et al., 2009; Hulpus et al., 2013). Techniques from extractive summarization have also been used for topic extraction (Basave et al., 2014; Wan and Wang, 2016), which typically extract summary sentences from the input text related to topics. Recent years have witnessed neural networks are successfully leveraged to improve performance of topic modeling techniques, such as incorporating neural embeddings into existing LDA-like models (Bianchi et al., 2021; Thompson and Mimno, 2020), as well as the clustering embedding based approaches (Sia et al., 2020; Angelov, 2020; Grootendorst, 2022). A potential limitation of such methods is that the topic labels are within a predefined limited candidate set, while there are often emerging scenes in the e-commerce fields. Therefore, similar to Alokaili et al. (2020), we design a pre-trained model in e-commerce domain to generate scene-based topic titles, which allows to generate novel topics not featured in training set.

Architecture	Acceptance Rate (%)
ESTC w/o Quality Control	51.6
+ Topic Coherence	65.6
+ Correlation Scoring	60.6
+ Both	75.0

Table 5: Human evaluation for quality control.

Natural language processing techniques have been widely used in e-commerce fields to improve user experience, including automatic product copywriting generation (Zhang et al., 2022; Wang et al., 2022), online product review generation (Fan et al., 2019; Liu et al., 2021) and question generation (Gao et al., 2020; Deng et al., 2020). Differently, we propose to leverage natural language generation and clustering techniques to automatically construct scene-based topic channels, which, to the best of our knowledge, is novel.

7 Conclusion

This work aims to automatically construct scene-based topic channels. According to the understanding of business requirements, we propose to first generate topic titles for each product following by conducting product clustering to form a channel and design a novel framework, consisting of topic generation, product clustering and post-processing modules. The extensive offline experiments and online A/B test have demonstrated the effectiveness of the proposed approach. Incorporating user behaviors (e.g., click preference) into channel construction processes is worthy investigating in future. For example, the generated title and the clustered product are personalized.

8 Ethical Considerations

The data used in this work are collected from a publicly available online e-commerce platform, where the collection process is consistent with the terms of use, intellectual property and privacy rights of the platform. The annotated data for clustering evaluation are constructed by authors, where the

process is fair for all models. Please note that no private user data was used during data collection process.

The proposed ESTC system can be deployed on various e-commerce platforms where the scene marketing is required. On the other hand, the display style (or the product form) can be changed according to the practical needs, where the proposed system can provide products that belong to same usage scenarios.

Moreover, the AI constructed channels are not fool proof. Thus, as we discussed in the paper, in order to ensure the users can have a reasonably good experience, quality control and human screening of AI generated channels in the production platform is necessary to filter out any inconsistent or low-quality content.

9 Acknowledgements

We thank all the anonymous reviewers for their constructive comments.

References

- Areej Alokaili, Nikolaos Aletras, and Mark Stevenson. 2020. [Automatic generation of topic labels](#). In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval, SIGIR 2020, Virtual Event, China, July 25-30, 2020*, pages 1965–1968. ACM.
- Dimo Angelov. 2020. Top2vec: Distributed representations of topics. *arXiv preprint arXiv:2008.09470*.
- Amparo Elizabeth Cano Basave, Yulan He, and Ruifeng Xu. 2014. Automatic labelling of topic models learned from twitter by summarisation. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 618–624.
- Shraey Bhatia, Jey Han Lau, and Timothy Baldwin. 2016. Automatic labelling of topics with neural embeddings. In *International Conference on Computational Linguistics*. Association for Computational Linguistics, ACL Anthology.
- Federico Bianchi, Silvia Terragni, Dirk Hovy, Debora Nozza, and Elisabetta Fersini. 2021. Cross-lingual contextualized topic models with zero-shot learning. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1676–1683.
- David M. Blei, Andrew Y. Ng, and Michael I. Jordan. 2003. [Latent dirichlet allocation](#). *J. Mach. Learn. Res.*, 3:993–1022.
- Philip Cooke and Loet Leydesdorff. 2006. Regional development in the knowledge-based economy: The construction of advantage. *The Journal of Technology Transfer*, 31:5–15.
- Yang Deng, Wenxuan Zhang, and Wai Lam. 2020. Opinion-aware answer generation for review-driven question answering in e-commerce. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, pages 255–264.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4171–4186. Association for Computational Linguistics.
- Li Dong, Nan Yang, Wenhui Wang, Furu Wei, Xiaodong Liu, Yu Wang, Jianfeng Gao, Ming Zhou, and Hsiao-Wuen Hon. 2019. Unified language model pre-training for natural language understanding and generation. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, pages 13063–13075.
- Miao Fan, Chao Feng, Lin Guo, Mingming Sun, and Ping Li. 2019. Product-aware helpfulness prediction of online reviews. In *2019 World Wide Web Conference, WWW 2019*, pages 2715–2721. Association for Computing Machinery, Inc.
- Mariano Fernandez-Lopez and Oscar Corcho. 2010. *Ontological Engineering: with examples from the areas of Knowledge Management, e-Commerce and the Semantic Web*. Springer Publishing Company, Incorporated.
- Min Fu, Qiang Chen, Wei Lin, Pei Wang, and Wei Zhang. 2019. Constructing a scene-based knowledge system for e-commerce industries: Business analysis and challenges. *Data Intelligence*, 1(3):224–237.
- Shen Gao, Xiuying Chen, Chang Liu, Li Liu, Dongyan Zhao, and Rui Yan. 2020. Learning to respond with stickers: A framework of unifying multi-modality in multi-turn dialog. In *Proceedings of the Web Conference 2020*, pages 1138–1148.
- 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, EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 7-11 November, 2021*, pages 6894–6910. Association for Computational Linguistics.
- Maarten Grootendorst. 2022. Bertopic: Neural topic modeling with a class-based tf-idf procedure. *arXiv preprint arXiv:2203.05794*.

- Suchin Gururangan, Ana Marasovic, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A. Smith. 2020. [Don't stop pretraining: Adapt language models to domains and tasks](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 8342–8360. Association for Computational Linguistics.
- Ioana Hulpus, Conor Hayes, Marcel Karnstedt, and Derek Greene. 2013. Unsupervised graph-based topic labelling using dbpedia. In *Proceedings of the sixth ACM international conference on Web search and data mining*, pages 465–474.
- Wang-Cheng Kang, Eric Kim, Jure Leskovec, Charles Rosenberg, and Julian McAuley. 2019. Complete the look: Scene-based complementary product recommendation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Jey Han Lau, Karl Grieser, David Newman, and Timothy Baldwin. 2011. Automatic labelling of topic models. In *Proceedings of the 49th annual meeting of the association for computational linguistics: human language technologies*, pages 1536–1545.
- Jey Han Lau, David Newman, Sarvnaz Karimi, and Timothy Baldwin. 2010. Best topic word selection for topic labelling. In *Coling 2010: Posters*, pages 605–613.
- Alon Lavie, Kenji Sagae, and Shyamsundar Jayaraman. 2004. The significance of recall in automatic metrics for mt evaluation. In *AMTA*, pages 134–143.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. [BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81.
- Junhao Liu, Zhen Hai, Min Yang, and Lidong Bing. 2021. 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.
- Davide Magatti, Silvia Calegari, Davide Ciucci, and Fabio Stella. 2009. Automatic labeling of topics. In *Proceedings of the 2009 Ninth International Conference on Intelligent Systems Design and Applications*, pages 1227–1232.
- Robin Mansell. 2002. Constructing the knowledge base for knowledge-driven development. *Journal of Knowledge Management*, 6(4):317–329.
- Qiaozhu Mei, Xuehua Shen, and ChengXiang Zhai. 2007. Automatic labeling of multinomial topic models. In *Proceedings of the 13th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 490–499.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.
- 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.
- Matt Post. 2018. A call for clarity in reporting bleu scores. *arXiv preprint arXiv:1804.08771*.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *Journal of Machine Learning Research*, 21(140):1–67.
- Margaret Roberts, Brandon Stewart, Dustin Tingley, and Edoardo Airoldi. 2013. The structural topic model and applied social science. *Neural Information Processing Society*.
- Peter J Rousseeuw. 1987. Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *Journal of computational and applied mathematics*, 20:53–65.
- Nachiketa Sahoo, Jamie Callan, Ramayya Krishnan, George Duncan, and Rema Padman. 2006. Incremental hierarchical clustering of text documents. In *Proceedings of the 15th ACM international conference on Information and knowledge management*, pages 357–366.
- Yunfan Shao, Zhichao Geng, Yitao Liu, Junqi Dai, Fei Yang, Li Zhe, Hujun Bao, and Xipeng Qiu. 2021. Cpt: A pre-trained unbalanced transformer for both chinese language understanding and generation. *arXiv preprint arXiv:2109.05729*.
- Suzanna Sia, Ayush Dalmia, and Sabrina J Mielke. 2020. Tired of topic models? clusters of pretrained word embeddings make for fast and good topics too! In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1728–1736.
- Laure Thompson and David Mimno. 2020. Topic modeling with contextualized word representation clusters. *arXiv preprint arXiv:2010.12626*.

Xiaojun Wan and Tianming Wang. 2016. Automatic labeling of topic models using text summaries. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2297–2305.

Zeming Wang, Yanyan Zou, Yuejian Fang, Hongshen Chen, Mian Ma, Zhuoye Ding, and BO Long. 2022. Interactive latent knowledge selection for e-commerce product copywriting generation. In *Proceedings of The Fifth Workshop on e-Commerce and NLP (ECNLP 5)*, pages 8–19.

Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. 2021. **mT5: A massively multilingual pre-trained text-to-text transformer**. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 483–498, Online. Association for Computational Linguistics.

Xueying Zhang, Yanyan Zou, Hainan Zhang, Jing Zhou, Shiliang Diao, Jijia Chen, Zhuoye Ding, Zhen He, Xueqi He, Yun Xiao, et al. 2022. Automatic product copywriting for e-commerce. In *Proceedings of the AAAI Conference on Artificial Intelligence*.

Mingxiao Zhao. 2020. Data-driven scene marketing based on consumer insight. In *2020 International Conference on Big Data Economy and Information Management (BDEIM)*, pages 61–65. IEEE.

Guorui Zhou, Na Mou, Ying Fan, Qi Pi, Weijie Bian, Chang Zhou, Xiaoqiang Zhu, and Kun Gai. 2019. Deep interest evolution network for click-through rate prediction. In *Proceedings of the AAAI conference on artificial intelligence*.

Guorui Zhou, Xiaoqiang Zhu, Chenru Song, Ying Fan, Han Zhu, Xiao Ma, Yanghui Yan, Junqi Jin, Han Li, and Kun Gai. 2018. Deep interest network for click-through rate prediction. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*, pages 1059–1068.

Yanyan Zou, Xingxing Zhang, Wei Lu, Furu Wei, and Ming Zhou. 2020. Pre-training for abstractive document summarization by reinstating source text. In *Proceedings of the Empirical Methods in Natural Language Processing*.

A Experiment Settings

A.1 The Hyper-parameters of Scene-based Topic Generation Model

In this section, we introduce the detailed setting of proposed Scene-based Topic Generation. To generate the scene-based topic titles for each product, we design a topic generation model based on UniLM, the input sequence and output sequence are encoded by the same attention module with different attention masks. The model is a 12-layer

transformer with multi-head attentions. During training, the learning rate is 0.00007, the warmup proportion is set to 0.2 and the batch size is 1024. The detailed hyper-parameters are listed in Table 6. The rest of the parameters are set by default.

hyper-parameters	value
learning_rate	0.00007
warmup_proportion	0.2
batch_size	1024
max_input_length	120
max_output_length	20
beam_size	4
embedding_size	768
hidden_dropout_prob	0.1
hidden_size	768
layer_norm_eps	1e-12
max_position_embeddings	250
num_attention_heads	3
num_hidden_layers	12
activation	"gelu"
vocab_size	21128

Table 6: The detailed hyper-paramters of architecture of E-commerce UniLM.

A.2 The Hyper-parameters of Topic Encoding Model

The topic encoding model encodes topic texts into vector features of specified dimensions, which facilitates clustering by common numerical clustering models. In this section, we introduce the detailed setting of proposed theme encoding model. We first obtain 700k topic data by the inference results of the E-commerce UniLM model. Then we employ SimCSE³ to fine-tune the second pre-trained UniLM model in the e-commerce domain. The backbone model is a 12-layer transformer. The learning rate is 0.00003 and the batch size is 64. The rest of the parameters are set by default.

hyper-parameters	value
num_train_epochs	4
max_len	32
train_batch_size	64
learning_rate	3e-5
max_seq_length	32
evaluation_strategy	steps
pooler_type	cls
temp	0.05

Table 7: The detailed hyper-paramters of architecture of topic encoding.

³<https://github.com/princeton-nlp/SimCSE>

A.3 The Hyper-parameters of Topic Coherence Model

Topic Coherence Model is 12-layer transformer with a feed-forward network and a softmax layer to distinguish whether the input topic is coherent. The learning rate is 0.00005 and the batch size is 2048. The rest of the parameters are set by default.

hyper-parameters	value
learning_rate	0.00005
warmup_proportion	0.1
batch_size	2048
max_len	32
embedding_size	768
hidden_dropout_prob	0.1
hidden_size	768
layer_norm_eps	1e-12
max_position_embeddings	250
num_attention_heads	3
num_hidden_layers	12
activation	"gelu"
vocab_size	21128

Table 8: The detailed hyper-parameters of architecture of topic coherence model.

A.4 The Hyper-parameters of Correlation Scoring Model

To filter out bad cases where topic title is not suitable for product usage scenarios, we design another binary classification model, i.e., correlation scoring model, to identify if the topic title and products are scene-based related. We concatenate the product description information X and topic title Y as the input of UniLM. For better learning the product usage scenario, we also take into account the product profile information, such as age, season, and gender profiles, and employ an embedding layer and a feed-forward layer to encode such features. The detailed hyper-parameters of correlation scoring model are listed in Table 9. The rest of the parameters are set by default.

B The Generated Scene-based Topic

In this subsection, we show some example topics generated by topic generation models, as well as examples of topics generated by the entire system, as shown in Table 10 and Table 11, respectively.

C Examples of Scene Marketing

In recent years, many companies have begun to use scene marketing to promote products, as shown in

hyper-parameters	value
learning_rate	0.00005
warmup_proportion	0.1
batch_size	2048
max_len	158
feature_embedding_size	300
feature_fusion_size	300
feature_vocab_size	13
embedding_size	768
hidden_dropout_prob	0.1
hidden_size	768
layer_norm_eps	1e-12
max_position_embeddings	250
num_attention_heads	3
num_hidden_layers	12
activation	"gelu"
vocab_size	21128

Table 9: The detailed hyper-parameters of architecture of correlation scoring model.

Figure 4, which are scene marketing for IKEA⁴ and Amazon⁵. IKEA combines different types of furniture in a scene room to highlight key attributes of furniture, such as storage and simple shape. Amazon also exhibits the functional scenarios for products, like babysitting and party games. Such scene marketing can help consumers understand the functions and features of products, which may improve user experience and product conversion rates.

⁴<https://www.ikea.com/>

⁵<https://www.amazon.com/>


Input	Generated Topic
创意烟灰缸生日送礼送男朋友 Creative ashtray birthday gift for boyfriend	送男友好物 @ 为爱精挑细选 Gifts For Boyfriends @ Carefully Selected For Love
夏季宽松休闲翻领男上衣 Summer Loose Casual Lapel Men's Top	精选t恤 @ 夏日清凉出行 Selection Of T-shirts @ Summer Cool Outting
网红款渔夫帽 Web celebrity's fisherman hat	防晒合集 @ 清凉防晒一夏 Sun Protection Collection @ All Summer Cool Sun Protection
龙井2022新茶绿茶茶礼盒装 Longjing 2022 new tea green tea gift box	女婿必买 @ 教你一招搞定老丈人 Son-in-law Must Buy @ Teach You A Trick To Get Father-in-law

Table 10: The generated scene-based topic titles by the topic generation model. We use @ to separate two phrases of the topic title.


Generated Topic	Product List
初生好礼 @ 虎娃新生儿礼物 Newborn Gift @ Tiger Baby Newborn Gift	尿裤尿不湿学步裤吸湿透气 Moisture absorbent breathable diaper toddler pants 初生宝宝幼儿浴巾被子防惊跳睡袋 Newborn baby bath towel quilt anti-startle sleeping bag 婴儿记忆棉乳胶枕头枕芯 Baby memory foam latex pillow 婴儿配方奶粉2段850克 Infant formula milk powder 2 stage 850g
快乐露营 @ 露营运动欢乐时光 Happy Camping @ Camping Sports Happy Hour	露营灯强光手电筒帐篷灯 Camping lights glare flashlight tent lights 户外折叠桌椅便携式野外可折叠野餐桌子 Portable outdoor folding picnic table 登山露营保暖加宽双人户外棉睡袋 Mountaineering camping warm widening double outdoor cotton sleeping bag 大空间防风3-4人三秒速开全自动速搭帐篷 Large space windproof 3-4 people three seconds to open fully automatic tent
尽情挥洒汗水 @ 是兄弟一起上球场 Sports Sweat @ On The Court With Your Brother	高帮板鞋男子经典运动休闲鞋篮球文化鞋 High-top sneakers men's classic sneaker, basketball culture shoes 男装梭织运动长裤运动服男 Men's woven sports trousers sportswear 简约经典训练系列男子圆领套头休闲百搭卫衣 Simple and classic training series casual all-match sweatshirt 针织五分裤男透气舒适夏季短裤男运动裤子 Knitted 1/2 pants men's breathable and comfortable summer running workout joggers

Table 11: The generated scene-based topic channel of ESTC system. We use @ to separate two phrases of the topic title.








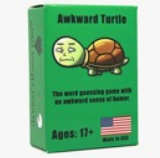
Examples of Scene Marketing for Different Companies


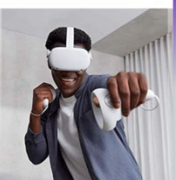
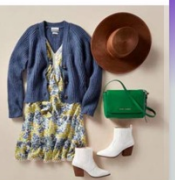
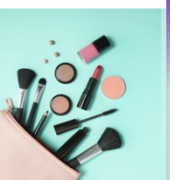


PAX bedroom storage furniture collection
A flexible closet collection that offers clothing solutions to fit different sized spaces. Keeping all your clothes neatly organized and in their place, while still showing off your individual style.



MALM series
Simple shape, versatile with various home styles.

Must have Baby products		Adult party games	
			
			

Shop Laptops & Tablets 	Oculus 	Women's Fashion 	Beauty Picks 
---	--	--	---






Figure 4: Examples of scene marketing for different companies.