Visual Question Rewriting for Increasing Response Rate

Jiayi Wei

jwei37@ucsc.edu

University of California, Santa Cruz

Yi Zhang

yiz@ucsc.edu

University of California, Santa Cruz

Xilian Li

xli237@ucsc.edu

University of California, Santa Cruz

Xin Eric Wang

xwang366@ucsc.edu

University of California, Santa Cruz

ABSTRACT

When a human asks questions online, or when a conversational virtual agent asks a human questions, questions triggering emotions or with details might more likely to get responses or answers. We explore how to automatically rewrite natural language questions to improve the response rate from people. In particular, a new task of Visual Question Rewriting (VQR) task is introduced to explore how visual information can be used to improve the new question(s). A data set containing ~4K bland&attractive question-images triples is collected. We developed some baseline sequence to sequence models and more advanced transformer-based models, which take a bland question and a related image as input, and output a rewritten question that's expected to be more attractive. Offline experiments and mechanical Turk based evaluations show that it's possible to rewrite bland questions in a more detailed and attractive way to increase response rate, and images can be helpful.

CCS CONCEPTS

• **Computing methodologies** → **Natural language processing**; **Computer vision**.

KEYWORDS

Natural language processing, Computer vision, Question rewriting, Sequence2sequence, Transformer, Response Rate

ACM Reference Format:

Jiayi Wei, Xilian Li, Yi Zhang, and Xin Eric Wang. 2021. Visual Question Rewriting for Increasing Response Rate. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '21)*, July 11–15, 2021, Virtual Event, Canada. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3404835.3463114>

1 INTRODUCTION

When buying complex products such as real estates, furniture, automobiles, buyers often shop locally and ask local store staff various product related questions before making a purchasing decision. However, with the advance of online shopping or virtual reality technologies, and with the impact of covid-19, people are starting to

make complex purchases and asking questions online. Some questions are answered by unpaid specialists, sellers or other buyers, however, many questions are not answered.

After analyzing online QA web sites or discussion forums, we found more attractive questions, such as those with details or emotional words, are more likely to be answered. However, majority normal users are not good at writing attractive questions, or may not take the effort to do so, when shopping around and browsing product images. Thus we propose to automatically rewrite a simple question into a more attractive one, and we expect the visual image associated with the original question might help.

This paper proposes a new Visual Question Rewriting (VQR) problem that takes an image and a simple natural language question related to the image as input, and output a more attractive question. We created a data set containing many <image, question, more attractive question> triplets based on Houzz.com, an online community with many questions about published home designs. Some questions on Houzz received answers from the designer who shared the design image. Some examples are shown in Fig 1. A baseline Seq2Seq [20] model and a transformer-based model are implemented to explore whether the proposed research direction is promising and what results look like.

Figure 1: Example of image, simple question and detailed question triples.

2 RELATED WORK

This work is mainly related to and motivated by a large amount of prior research on natural language processing, computer vision and the interplay between the two areas.

Natural Language Processing has been studied since the late 1940s. NLP field has several important sub-tasks. Summarizing, text generation, story writing, text simplification, rewriting questions

This work is licensed under a Creative Commons Attribution International 4.0 License.

SIGIR '21, July 11–15, 2021, Virtual Event, Canada.

© 2021 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8037-9/21/07.

<https://doi.org/10.1145/3404835.3463114>

in context of a dialog systems, and grammar error correction are existing NLP tasks related our problem. The performances of those tasks have improved significantly in the last several years due to the advance of deep learning technologies and the availability of large scale training data. The models we proposed in this paper are motivated by these prior research work, especially the recent advances [8, 10, 17, 18, 23]. However, the goal of our new rewriting task is different: we are focusing on extending and enriching a user's input so that it's more attractive.

Computer vision(CV) systems have improved significantly over the last decade and perform better than human in some standard classification or object detection tasks [12, 19]. Over the last several years, two main tasks combining natural language and computer vision have emerged: Visual Question Answering [1, 2, 11] and Image Captioning [14]. Existing VQA tasks mainly focus on selecting the best answer from candidates or produce a short answer for questions based on given images, while our new task focuses on how to generate a more attractive question to increase response rate, based on visual information and the original text question. Usually, the answer of our questions can not be derived from the images and needs to be answered by a human. These differences make our problem novel and more difficult to evaluate. That's why we start with a specific domain in this paper.

3 DATASET COLLECTION

To support the new research, we crawled images and their related questions from Houzz.com. ~3K images with ~18K questions are collected. To reconstruct a data set with <image, attractive question, bland question> triplets, questions that have some response(s) and contain more than 1 sentence are used, assuming those questions are attractive. For each attractive question, the key question sentence is selected as the corresponding bland question. As a result, the new VQR dataset contains ~4K image-attractive question-bland question triples. We split it into 4 : 1 for training and testing.

4 METHODS

As a preliminary research, we evaluated 4 different straightforward model architectures. These models are motivated by existing research in NLP[22] and image captioning[3, 9]. They either take the original question or (original question, associated image) as input, and output a new question.

4.1 Baseline Models

We implemented a basic Seq2Seq architecture for natural language generation as the baseline solution. In order to explore the impact of the image input features, we did experiments with or without the visual input, as illustrated in Fig 2.

Firstly tokenizing all text information (input and output questions), provided us a vocabulary of ~5K words. Then all input and targeting sentences were mapped into sequences with length of 30 and 50, respectively. The baseline model has 2 layers of Gated recurrent units(GRUs) [4] for both encoder and decoder. With the output from the encoder, the attention mechanism was applied to generate the attention weights which are used to pass weighted sum of encoder output to the decoder part. The attention equations [15] used in

Figure 2: Baseline model illustration. The architecture without the green dashed box is the pure language model. The part in the green dashed box is for the visual part.

our work are as below, where h_s denotes encoder output and h_t denotes the current hidden state vector, v , W are model parameters, α_{ts} are attention weights, c_t is a context vector for the decoder:

$$score(h_t, h_s) = v_{\alpha}^T \tanh(W_{\alpha} [h_t; h_s]) \quad (1)$$

$$\alpha_{ts} = \text{softmax}(score(h_t, h_s)) \quad (2)$$

$$c_t = \sum_s (\alpha_{ts} h_s) \quad (3)$$

Then, the embedding vector from the previous position and the acquired *vector* are concatenated together and passed into the decoder. In the experiments, we set the length of the embedding vector and hidden status vector 256 and 768 respectively.

Next, visual information is added to the model to see whether it can boost the performance. All images were resized as $224 * 224$ and feed to a pre-trained ResNet-50 [12]. The feature map before the fully-connected layer was used as the input for our model (Fig 2). After the forward pass, we received a $7 * 7 * 2048$ feature set. We flattened the feature map to $49 * 2048$. We implemented the similar attention method on the flattened image feature (Fig 2), because the $7 * 7$ grid possess the position information on this patch. The language generation part remained the same here.

4.2 Transformer Based Model

The self-attention [21] model is originally designed for natural language understanding (NLU) tasks. BERT [6, 7, 21] drives it to a higher level, which updates the state-of-the-art metrics in most NLU tasks. Transformer-based models allow the encoder and decoder to see the entire input sequence all at once and to directly model these dependencies using attentions. More conveniently, BERT model pre-trained on a large corpus is publicly available, which enables researchers to start with a good model and fine-tune on their specific tasks.

Much work has been attempted to combine BERT with Seq2Seq models. UniLM [7] is one of the most successful and decent one. The original BERT model sets mask on words randomly to force

Table 1: Model performance on our VQR data.

	baseline	baseline+vis	transformer	transformer+vis
BLEU	0.194	Nan	0.528	0.547
RG-1	0.198	Nan	0.753	0.765
RG-2	0.167	Nan	0.614	0.651
RG-L	0.307	Nan	0.870	0.885

the model better understand the generic context. Similarly, UniLM treats Seq2Seq as language completion. In this way, the attention on the input and output are bi-directional and uni-directional respectively, which can satisfy the requirement for language generation. Besides, with this specific mask, the language generation task is completed with a single BERT model and powerful pre-trained BERT weights.

Figure 3: Illustration of attention mask in the transformer based language generation model.

We firstly evaluated the transformer-based language model without image input in Figure 3. Then, we did experiments by adding image-based inputs, again using ResNet-50 as the feature extractor. We found the attention mechanism on image doesn't work on the baseline method, thus we used the generic image feature instead of the grid image feature, which means the feature after the dense layer is used as our visual input to be combined with the language feature. Visual information is included through conditional normalization [5] (Figure 4):

$$\hat{x} = \gamma(c) \times \frac{x - \mu}{\sigma + \epsilon} + \beta(c) \quad (4)$$

$$\mu = \frac{1}{N} \sum x, \quad \sigma^2 = \frac{1}{N} \sum (x - \mu)^2$$

where x and c are language feature and conditional input (image feature) respectively, μ and σ^2 are mean and variance, and ϵ is a small constant number. This conditional normalization are firstly proposed and widely utilized as conditional batch normalization [5, 13, 16]. In our work, we modified the pre-trained BERT to generate language sequence and applied layer normalization as the main normalization method. β and γ variables in BERT's normalization layers are changed to control the sequence output. In traditional normalization layers, β and γ are learned through parameter updates. However, with conditional normalization, β and γ are generated from input, thus the normalization layer output is conditional on its side input. Because β and γ in the pre-trained

model have fixed length, a dense layer is used to convert the image features into the required length.

Figure 4: The illustration of conditional layer normalization.

5 EXPERIMENTS AND ANALYSIS

This section presents the experiment details and results. The RNN based Seq2Seq models are called baseline and baseline+vis, and the transformer based models are called transformer and transformer+vis.

The results are shown in Table 1. The baseline models performs poorly compared with the transformer models. The baseline+vis model didn't converge at the training stage, probably due to the lack of pre-trained models and the number of training data in our task is not enough for leaning embedding. Transformer models perform much better, probably because BERT weights used are pre-trained on large amount of training data sets across multiple language tasks jointly, which makes BERT more general and not over-fit on any single task. Besides, with attention mask specified for language generation, the attention of the input part is bidirectional and the attention of the output part is unidirectional. Therefore, the transformer model performs better than the baseline model, which is common for other NLP tasks.

With the help of conditional layer normalization, the related image features really contribute to language generation and improve performance. Rather than the bigger feature map, the image feature vector before *softmax* contains more generic context information. Besides, using the generic context as input, there is no need to learn attentions based on VQR data, which seems to be a difficult task and the main reason for baseline+vis not converging. Adding image feature through layer normalization effectively controls its influence, especially with zero initialization, as this accelerates the training process and avoids the harm of warming up.

Further analysis was done by examining some representative examples. As shown in Figure 5, the baseline model does not know when and where to stop, a common problem shared by many standard RNN methods. It also fails to acquire the most obvious pattern in the training data: copy the original question. As we mentioned before, the baseline+vis model doesn't converge during training, thus produce weird results, so we won't show it here. The output of the transformer models seems more attractive, with emotional parts added at the beginning or end of the original bland sentence. High level details "type, color, size, etc." could be added occasionally (middle example) by the transformer model, this is due to the trained language model, instead of the influence of the associated

Input	what manufacturer created the parts in this lighting set up?	can you tell me a little bit more about the beams?	is this a prefinished red oak hardwood or finished on site?
Target	love the warehouse look! the lighting is very creatively done. how did you do the lighting. it looks like it is attached or going through a tube. what manufacturer created the parts in this lighting set up? susan	type of beams? love this! can you tell me a little bit more about the beams? what kind of wood are they?	floor? is this a prefinished red oak hardwood or finished on site? and what color stain?
Baseline	what size do you have your time .	are the hood ? is it s beautiful ! can you tell me a little info on the hood ? is it s beautiful ! can you tell me a little info on the hood ? is it s beautiful ! can you tell me a little info	i m having brick on floor . is it difficult to maintain ? i want to decide on these wood floor plan of wood floor . is it difficult to maintain ? i want to decide on these wood floor plan of wood floor . is it difficult
Transformer	we love the lighting set up! what manufacturer created the parts in this lighting set up? thank you!	can you tell me a little bit more about the beams? type, color, size, etc. thank you. beautiful work!	kitchen flooring. i love the flooring! is this a prefinished red oak hardwood or finished on site? thank you!
Transformer+Vis	what manufacturer created the parts in this lighting set up? i'm in love with this. it's a beautiful feel to it.	beams. can you tell me a little bit more about the beams? what type of wood? what type?	flooring. is this a prefinished red oak hardwood or finished on site? what is the stain?

Figure 5: Examples of result on the VQR test data.

image. On the other hand, the transformer+vis model is able to add new information inferred from the image, such as "what type of wood?", "what is the stain". This suggests image helps adding extra details to the bland question.

6 HUMAN EVALUATIONS

We hired several Amazon Mechanical Turks to evaluate the generated results on test data. For each image, we present the question generated by transformer+vis model and the original bland question, let 3 different workers choose which question is more attractive independently, then use the majority vote as the final judge for each image. Out of 767 images, 652 rewriting questions (85.8%) are judged to be better than simple bland questions. This demonstrates that the rewriting model successfully generates detailed and emotional questions and makes the question more attractive.

We also hired Turks to compare questions generated by transformer vs transformer+visual. 410 questions (53.4%) generated by the transformer+vis model are preferred by workers, which suggest extra information brought by the image input might make the rewriting more attractive.

7 CONCLUSIONS

We propose a new task of Visual Question Rewriting (VQR) to convert a natural language question into a more attractive one, so that it's more likely to be answered. A new data set was collected and will be shared with the research community. As a first step, we implemented two straightforward approaches: a standard sequence to sequence model and a transformer-based model, and found the transformer based model performs much better. We modified the original BERT model, and transferred image features as conditional layer normalization input to influence the output sentences, and found this transformer+visual model can rewrite a question with more details. User evaluation demonstrated visual question rewriting is more attractive, thus a promising direction to increase question response rate.

Although conversational IR is not the focus of this paper, how to make conversational agents, such as conversational recommendation or QA systems more attractive, is a major challenge in practice. This work can also be adapted to re-write questions made by a virtual assistant to increase user engagements with the system.

REFERENCES

- [1] Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. 2018. Bottom-up and top-down attention for image captioning and visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 6077–6086.
- [2] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. 2015. Vqa: Visual question answering. In *Proceedings of the IEEE international conference on computer vision*. 2425–2433.
- [3] Xinlei Chen and C Lawrence Zitnick. 2015. Mind’s eye: A recurrent visual representation for image caption generation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2422–2431.
- [4] Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning Phrase Representations using RNN Encoder-Decoder for Statistical Machine Translation. *arXiv:1406.1078* [cs.CL]
- [5] Harm de Vries, Florian Strub, Jérémie Mary, Hugo Larochelle, Olivier Pietquin, and Aaron Courville. 2017. Modulating early visual processing by language. *arXiv:1707.00683* [cs.CV]
- [6] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).
- [7] 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 *Advances in Neural Information Processing Systems*. 13063–13075.
- [8] Ahmed Elgohary, Denis Peskov, and Jordan Boyd-Graber. 2019. Can you unpack that? learning to rewrite questions-in-context. *Can You Unpack That? Learning to Rewrite Questions-in-Context* (2019).
- [9] Hao Fang, Saurabh Gupta, Forrest Iandola, Rupesh K Srivastava, Li Deng, Piotr Dollár, Jianfeng Gao, Xiaodong He, Margaret Mitchell, John C Platt, et al. 2015. From captions to visual concepts and back. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 1473–1482.
- [10] Seraphina Goldfarb-Tarrant, Haining Feng, and Nanyun Peng. 2019. Plan, Write, and Revise: an Interactive System for Open-Domain Story Generation. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*. 89–97.
- [11] Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. 2017. Making the V in VQA matter: Elevating the role of image understanding in Visual Question Answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 6904–6913.
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 770–778.
- [13] Xun Huang and Serge Belongie. 2017. Arbitrary style transfer in real-time with adaptive instance normalization. In *Proceedings of the IEEE International Conference on Computer Vision*. 1501–1510.
- [14] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. 2014. Microsoft coco: Common objects in context. In *European conference on computer vision*. Springer, 740–755.
- [15] Minh-Thang Luong, Hieu Pham, and Christopher D. Manning. 2015. Effective Approaches to Attention-based Neural Machine Translation. *arXiv:1508.04025* [cs.CL]
- [16] Takeru Miyato and Masanori Koyama. 2018. cGANs with projection discriminator. *arXiv preprint arXiv:1802.05637* (2018).
- [17] Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding by generative pre-training.
- [18] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI blog* 1, 8 (2019), 9.
- [19] Karen Simonyan and Andrew Zisserman. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014).
- [20] Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*. 3104–3112.
- [21] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems* 30 (2017), 5998–6008.
- [22] Thomas Wolf, Julien Chaumond, Lysandre Debut, Victor Sanh, Clement Delangue, Anthony Moi, Pierric Cistac, Morgan Funtowicz, Joe Davison, Sam Shleifer, et al. 2020. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*. 38–45.
- [23] Qionghai Xu, Chenchen Xu, and Lizhen Qu. 2019. ALTER: Auxiliary Text Rewriting Tool for Natural Language Generation. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP): System Demonstrations*. 13–18.