SCIENCE CHINA
Information Sciences

• LETTER •

April 2022, Vol. 65 149103:1–149103:2
https://doi.org/10.1007/s11432-020-3212-x

Example-guided stylized response generation in zero-shot setting

Guirong BAI[1,2], Shizhu HE[1,2*], Kang LIU[1,2] & Jun ZHAO[1,2]

[1] *National Laboratory of Pattern Recognition Institute of Automation, Chinese Academy of Sciences, Beijing* 100190*, China;*
[2] *University of Chinese Academy of Sciences, Beijing* 100049*, China*

Received 27 September 2020/Accepted 15 December 2020/Published online 14 December 2021

Citation Bai G R, He S Z, Liu K, et al. Example-guided stylized response generation in zero-shot setting. Sci China Inf Sci, 2022, 65(4): 149103, https://doi.org/10.1007/s11432-020-3212-x

Dear editor,

Stylized response generation is a valuable task for conversational agents. However, due to the lack of conversational data in various styles, learning to generate responses in the desired style is challenging. In this study, we propose an example-guided stylized response generation method in a zero-shot setting, which is able to fast generate responses in the desired style by utilizing only few non-parallel style examples as guidance. To achieve it, we train an example-guided generator only on a conversational dataset without any specific style and choose a suitable style example as the guidance for a generation when testing. In addition, to reduce the impact of bad training samples, we propose to reweight training samples with meta-learning. Experimental results show that our method significantly outperforms other models on stylizing responses in a low-resource setting and can generate style-specific responses even if no training data is available for the target style.

Definition. $D_{\rm conv} = [(x_1, y_1), \ldots, (x_n, y_n)]$ denotes a vanilla conversational dataset without any specific style, where x_i and y_i refer to the context and a corresponding response respectively. $D_{\rm style} = [s_1, s_2, \ldots, s_m]$ denotes a non-parallel and non-conversational dataset in a specific style, where s_i is one of the textual style examples. Examples in $D_{\rm style}$ do not have any aligned relation with samples in $D_{\rm conv}$ (non-parallel). Our goal is to train an example-guided model only on $D_{\rm conv}$ without $D_{\rm style}$ (zero-shot). When given $D_{\rm style}$ as style guidance for testing, the model is expected to generate context-relevant responses in the same style with examples in $D_{\rm style}$. In our low-resource setting, the number of style examples in $D_{\rm style}$ is limited.

Method. In our example-guided zero-shot stylized response generation, we propose to use style examples as the guidance to fast generate style-specific responses. To achieve it, we present an example-guided generator. It learns a pattern of how to transform the most similar example into the context-relevant response with minimum changes. As a result, the output response of the generator for testing is similar to the guidance example and thus is able to retain the linguistic style as much as possible.

However, the original training data $D_{\rm conv}$ only has context-response pairs, and does not provide such an example-guided generation format (context-example-response triples). Inspired by self-supervised learning [1], which learns a primary task where labeled data is not directly available but where the data itself provides a supervision signal for another auxiliary task which lets the network learn the primary one, we utilize original $D_{\rm conv}$ to construct such an example-guided conversational dataset $D_{\rm conv}^{\rm ref}$. Specifically, for each training sample in $D_{\rm conv}$, we obtain a guidance example by using the response y_i to retrieve another most similar response in $D_{\rm conv}$. In practice, we train an autoencoder on contexts and responses, and then use the Euclidean distance to find the nearest guidance example. Finally, training samples in $D_{\rm conv}^{\rm ref}$ consist of (x_i, r_i, y_i) triples, which denote contexts, guidance examples and responses respectively.

After obtaining training samples in an example-guided generation format, the generator fuses the contexts and guidance examples to output responses in a sequence-to-sequence framework [2]. Specifically, we use two different stacked gated-recurrent-unit (GRU) encoders for representations of context x_i ($z_{\rm S2S}(x)$) and guidance example r_i ($z_{\rm AE}(r)$) respectively. Then another GRU decoder augmented with context x_i is used to generate response y_i according to guidance example s_i.

$$h_t^{\rm S2S} = {\rm GRU}(h_{t-1}^{\rm S2S}, e(y_{t-1})), \tag{1}$$

$$P(y_t) = {\rm softmax}(W_{\rm S2S}[h_t^{\rm S2S}, z_{\rm S2S}(x)]), \tag{2}$$

where $h_t^{\rm S2S}$ is the decoder state at the t-th time step and $h_0^{\rm S2S} = z_{\rm AE}(r)$. $e(\cdot)$ denotes obtaining the word embedding. $[\cdot]$ denotes concatenation. Then, the example-guided generation objective is represented as $\mathcal{L}_y^{\rm S2S} = -\log p(y|z_{\rm AE}(r), z_{\rm S2S}(x)) = \sum_{t=1}^{T_y} P(y_t)$, where T_y denotes the length of the response.

When testing in a zero-shot setting via the example-guided generator, we need an encoder to obtain reliable representations of unseen style examples. To solve the prob-

* Corresponding author (email: shizhu.he@nlpr.ia.ac.cn)

Table 1 Results of different models

	arXiv-like					Holmes-like				
	Style	Relevancy		Diversity		Style	Relevancy		Diversity	
	Neural	BLEU1	BLEU2	Distinct1	Distinct2	Neural	BLEU1	BLEU2	Distinct1	Distinct2
Rand	0.99	12.1	1.7	0.13	0.56	0.60	13.1	1.9	0.15	0.62
Retrieval	0.84	15.5	2.3	0.06	0.19	0.20	10.7	1.7	0.04	0.15
Human	0.43	29.0	16.3	0.31	0.81	0.46	26.5	13.7	0.16	0.60
S2S+LM	0.03	15.0	1.5	**0.10**	**0.43**	0.02	15.5	2.4	**0.06**	**0.35**
MTask	0.01	14.5	2.3	0.06	0.21	0.01	17.0	3.8	0.04	0.16
StyleFusion	0.01	14.8	2.6	0.04	0.16	0.02	18.6	**3.9**	0.03	0.14
FewGuide	**0.33**	**15.0**	**2.8**	0.07	0.23	**0.48**	**19.0**	3.6	0.05	0.19

lem, we pretrain an autoencoder (AE) model to reconstruct responses in $D_{\rm conv}$. Specifically, we obtain response representations $z_{\rm AE}(y)$ using a GRU encoder of the AE model. Then we employ a decoder of the AE model to reconstruct responses. The objective is formulated as follows:

$$h_t^{\rm AE} = {\rm GRU}(h_{t-1}^{\rm AE}, e(y_{t-1})), \tag{3}$$

$$P(y_t) = {\rm softmax}(W_{\rm AE} h_t^{\rm AE}), \tag{4}$$

$$\mathcal{L}_y^{\rm AE} = -\log p(y|z_{\rm AE}(y)) = \sum_{t=1}^{T_y} P(y_t), \tag{5}$$

where the encoder for $z_{\rm AE}(y)$ is shared with the one for $z_{\rm AE}(r)$ in the example-guided generator. Moreover, the cell of GRU decoder for $h_t^{\rm AE}$ also shares parameters with the one for $h_t^{\rm S2S}$.

Similarly, we also reconstruct the guidance example with the same encoder and decoder. The object is as follows:

$$h_t^{\rm AE} = {\rm GRU}(h_{t-1}^{\rm AE}, e(r_{t-1})), \tag{6}$$

$$P(r_t) = {\rm softmax}(W_{\rm AE} h_t^{\rm AE}), \tag{7}$$

$$\mathcal{L}_r^{\rm AE} = -\log p(y|z_{\rm AE}(r)) = \sum_{t=1}^{T} P(r_t), \tag{8}$$

where T_r denotes the length of the guidance example. Both generating responses and reconstructing guidance examples use $z_{\rm AE}(r)$ as the initial decoder state, and they also share the same GRU cell for the decoder. The difference between them is whether there is a context to augment the decoder when predicting words.

In addition, we encourage representations to be distributed in a homogeneous manner scattered. Formally, the regularization term is represented as $\mathcal{L}_{\rm scatter} = -\sum_{i,j\in{\rm batch},i\neq j} \frac{d(z_{\rm AE}(y_i), z_{\rm AE}(y_j))}{N^2-N}$, where N denotes the batch size and $d(a,b)$ denotes the Euclidean distance.

For training, we optimize parameters at the same time by minimizing the total loss function, i.e., $\mathcal{L} = \mathcal{L}_y^{\rm S2S} + \mathcal{L}_y^{\rm AE} + \mathcal{L}_r^{\rm AE} + \mathcal{L}_{\rm scatter}$, where a (x_i, r_i, y_i) triple is a training sample in the mini-batch.

For inference, we choose a suitable style example from limited candidates in $D_{\rm style}$ as the guidance example s for the context x, and then send the chosen guidance example s into the example-guided generator to fast output style-specific responses y similar to s without extra training. To obtain the suitable style example s, we pretrain a specific autoencoder M on contexts and responses like the process of reconstruction object. Then we use the encoder of M to obtain representations of s and x, and compute the Euclidean distance between them to choose the nearest s as the suitable style example.

To obtain diverse responses for a given context, we sample different latent vectors. Specifically, we add a random vector σ of a given length $|\sigma|$ to $z_{\rm S2S}(x)$, i.e., $P(y) = -\log p(y|z_{\rm AE}(s), z_{\rm S2S}(x) + \sigma)$.

In addition, we employ a trade-off ρ to balance the degree of dependency on guidance examples, i.e., $P(y_t) = \rho P_{\rm AE}(y_t|s) + (1-\rho) P_{\rm S2S}(y_t|s,x)$. When predicting words, we fuse the word distribution of reconstructing guidance examples and generating responses. As ρ increasing, generated responses are more dependent on guidance examples. In practice, we set $\rho = 0.5$.

Experiments. Following [3], we compare our method (FewGuide) with different models. The results are presented in Table 1. We can see that our FewGuide has a high style intensity with comparable performance on relevancy and diversity. In contrast, other methods are unable to stylize responses in a low-resource setting and thus their style intensity is very low. It demonstrates that our method can generate stylized responses with only a few non-parallel style examples. Especially for style metric, MTask and StyleFusion can hardly work when style examples are limited.

Conclusion. We propose an example-guided zero-shot stylized response generation method. The experimental results show that our method significantly outperforms competitive models.

Acknowledgements This work was supported by National Key Research and Development Program of China (Grant No. 2020AAA0106400), National Natural Science Foundation of China (Grant Nos. 61533018, U1936207, 61976211, 61702512), and Independent Research Project of National Laboratory of Pattern Recognition (Grant No. Z-2018013). This research work was also supported by Youth Innovation Promotion Association CAS.

References

1 Raina R, Battle A, Lee H, et al. Self-taught learning: transfer learning from unlabeled data. In: Proceedings of the 24th International Conference on Machine Learning, 2007. 759–766

2 Sutskever I, Vinyals O, Le Q V. Sequence to sequence learning with neural networks. In: Proceedings of the 27th International Conference on Neural Information Processing Systems, 2014. 3104–3112

3 Gao X, Zhang Y Z, Lee S, et al. Structuring latent spaces for stylized response generation. In: Proceedings of Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), 2019. 1814–1823