

Promise Relevance in Diversity-Promoting Dialogue System

Yifan Zhou

Computer Science and Engineering, UC San Diego, La Jolla

yfzhou@ucsd.edu

Abstract

In this article, we introduce an significant and interesting problem about how to improve relevance between message and response of diversity for a open-domain dialogue generation system. Inspired by SEQ2SEQ model with attention and context information, a MMI-context prototype model is proposed. After discussing related work about context aware decoding RNN model, we describe some results in preliminary experiments from both quantitative and qualitative aspects.

1 INTRODUCTION

An appealing response generation system should be expected to output grammatical, coherent responses that are diverse and interesting. In practice, however, neural conversation models tend to generate trivial or non-committal responses, often involving high-frequency phrases along the lines of I dont know or Im OK [1][2][3]. More importantly, however, how to enhance the relation between message and response turns out to attract little attention in recent research contribution. The underlying motivation comes from the fact that dialogue turns becomes difficult to continue once machine gives meaningless, non-causal responses, though of considerable diversity.

In neural language generation model, e.g. SEQ2SEQ, responses that seem more meaningful or specific can also be found in the N-best lists, but rank much lower [4]. In part at least, this behavior can be ascribed to the relative frequency of generic responses like I dont know in conversational datasets, in contrast with the relative sparsity of more contentful alternative responses. After changing objective, the number of dull responses decreases on the one hand, while

candidates are still densely similar to each other and lack relevance to message. Table 1 illustrates this phenomenon, showing top outputs from MMI-antiLM models [4]. All the top-ranked responses are not quite generic but has little relation with the given message.

0.22976082563400269	until i get you
0.09237494319677353	until i get my water
-0.1772318333387375	until i get my money
-1.1583690643310547	until i get my water or
-1.191847801208496	until i get my money and
-1.2212326526641846	until i get my money and i want to
-1.29261493632286133	until i get my money and i
-1.3613690137396316	until i get my money and i want something else
-1.3692959712982178	until i get my money and i want something else to do .
-1.3829419612884521	until i get my money and i want something else to know
-1.398069295883179	until i get my money and i want to go
-1.4009228944778442	until i get my water or something
-1.4635696411132812	until i get my money and i want something else to do
-1.7492773532867432	until i get my money and i want something else to know .
-1.8667898178100586	until i get my money and i want something else to do . .
-1.9897382259368896	until i get my water or something okay
-2.058776378631592	until i get my money and i want something else to know . .
Message:	i take my leave
Response:	until i get you

0.48085737228393555	boy just a moment
0.018881281837821007	boy i m gonna kill you
-0.2039129137992859	boy i m gonna take you out
-0.22754280269145966	boy i m gonna take you to another hospital
-0.2515827715396881	boy i m gonna take you to another
-0.2964247763156891	boy i m gonna take you to the hospital
-0.5100808143615723	boy i m gonna take you to another hospital tonight
-0.6366313099861145	boy i m gonna take you to the hospital i m gonna need
-0.7092328830566406	boy i m gonna take you to another hospital tonight .
-0.7412936091423035	boy i m gonna take you to the hospital i m sorry
-0.780777633190155	boy i m gonna take you to the hospital i m gonna hit you
-0.8495317101478577	boy i m gonna take you to the hospital i m gonna hit the
Message:	you better stay right there boy
Response:	boy just a moment

Figure 1: Top ranked responses generated by MMI-antiLM model (Li et al., 2015) using beam decode for message "I take my leave" and "you better stay right there boy"

More specifically, we find that it is common that in basic SEQ2SEQ model or even those diversity-promoting ones, e.g. MMI-antiLM, result of beam decoding, namely the N-best list, is still highly centralized to only a few ancestors or prefixes. Such phenomena might be owned to MLE objective[4], as well as short memory of RNN in practice. Meanwhile, pursuing diversity only does not aligned with human-like dialogue system. Because, many answer generated are not quite boring, though, they hardly convey relevant information regards to source. Intuitively, it seems desirable not only to take into account the dependency

of responses from previous tokens generated, but also replay the relation that connects message and response. In this article, we therefore propose potentially feasible solutions to address such problem.

2 RELATED WORK

Maximum Mutual Information is adopted in recent diversity-promoting dialogue generation system. [4] Inspired by optimization objective in speech recognition domain, MMI considers not only the conditional probability which is given by traditional Maximum Likelihood Estimation, and also the probability of responses. From another point of view, the latter one could be seen as penalty to those phrases and sentences with high frequency. As a result, it might prefer those not much common which enhances diversity. Penalizing siblings in beam decode is another approach to improve diversity. By assigning different penalty level for intra-sibling ranking, the model would favor choosing hypotheses from diverse parents. [5] There are also plenty of other prior work addressing diversity problem during decoding like producing multiple outputs that are mutually diverse, either non-redundant summary sentences or N-best lists. [6][7] Nevertheless, most of them are evaluation driven which emphasizes much on a single aspect.

3 METHODOLOGY

One reason that generation tends to produce similar, dull and noncommittal response regardless of context information lies in limited influential time of LSTM or GRU unit in sequence to sequence model. For example, faced with two different questions as input, first two words of generated responses might be the same, like I dont. However, the following phrases could be and expected to be distinct and meaningful like I dont need more coffee and I dont care about who wins, instead of I dont know what you are talking about almost all the time. In other words, decoding process after the first few words seems to take account into little, if any, information from source. Consequently, the response tends to show less relevance and diversity.

A lot of prior work contributes to how to improve long-term memory. Recently, adding attention mechanism appears to be one of the most popular methods. In many situation like machine trans-

lation [8][10] and document summarization [9], it works well in terms of focusing on particular positions in source input. Nevertheless, from either point of view of experimental results or intuition behind, we argue that attention is not quite suitable for this open-domain dialogue system.

After applying attention mechanism on basic SEQ2SEQ architecture[10], the training time is far longer and performance drops considerably instead. By sampling test data and visualizing attention weights, it is often the case that one single token in source sentence gather majority of attention for each token decoded.

It is not difficult to explain that in dialogue generation scenario there are relatively much fewer corresponding tokens or phrases in source and target, compared to other application like machine translation. For instance, with seriously one of the best shows!!! fed in, the MMI-antiLM model probably generates love it ! cant wait for the next one... But its hard to clarify that each word in response should be associated with a single word in the context sentence.

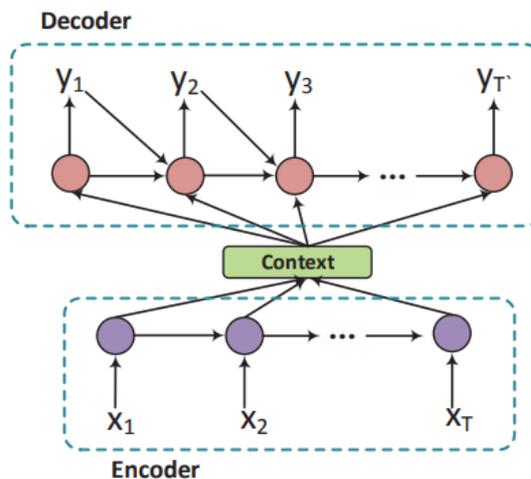


Figure 2: SEQ2SEQ model using context information for each decoding step

Considering those factors discussed above, a decoding process that combines last hidden state and also fixed context information from source should be promising, for both diversity and relevance. Such architecture is broadly applied in domains like summarization and document classification [11], which helps system capture historical records efficiently. In addition, in some recent work of neural dialogue generation paper considering context sensitivity[1], although not using

sequence to sequence model, context vector becomes a component of inputs for recurrent decoding process. What's more, in hierarchical architecture which takes as source input more than one utterances also adopts context vector for each token generation.

Static model Context vector has fixed length but varies as for different choice to compute. In this SEQ2SEQ scenario, for example, one way is to use simply linear or non-linear combination of hidden state in encoder, naturally like average or max. As to incorporating method, either concatenated form or weighted sum form is reasonable. It's important to assign proportion to context and decoder hidden state, in terms of dimension, since to some degree context vector may interface coherent generation than purely using prefix.

Dynamic model For every position during decoded utterance, the way it combines prefix and context information may vary a lot. Thus, we claim it both reasonable and feasible to add one control module that determine which kind and degree of incorporating should take for particular token. For the sum form, it essentially assign two positive number to context vector and hidden vector, which can be jointly learned with the whole network. For concatenated or other form, the answer is simply yes, no, or both. Such decision space is discrete which brings reinforcement learning to train this module. Specifically, policy gradient method [12] would show more power in both performance and ease of training, due to limited action space.

4 PRELIMINARY EXPERIMENT

Dataset The corpus used for this chatbot is open subtitles movie dataset [13]. This raw corpus is around 2.5 million lines, where odd lines are messages and even lines correspond to response. Before split into training, development and test set, it is preprocessing by lowercasing and filtering out extremely short or long utterances.

Training The naive context model is constructed based on SEQ2SEQ recurrent neural network model. The input of decoder is simply concatenated form of hidden state and context vector with same dimension, plus the token predicted last step. Slightly different with that, two GRU layer architecture is adopted, where each layer contains 1000 hidden neurons. Word embeddings of size 1000 are learned jointly with the whole network to

project vocabulary of size 20000. GRU parameter are initialized from a uniform distribution in $[-0.1, 0.1]$. Stochastic gradient descent use Adam optimizer, with initial learning rate of 0.0005, reduced by three quarter after each 10 epochs. Batch size is set to 48. This implementation was completed on a single NVIDIA GTX 1080 Ti GPU process, lasting for about 12 hours basic model training.

Evaluation As for evaluation, we choose three measurements as evaluation metric: BLEU-1 score, unique uni-grams and unique bi-grams. The reason why BLEU-1 is used is because it is not necessary to force predicted sentence to be identical to the single reference in high order.

Model	BLEU	distinct-I	distinct-II
SEQ2SEQ	.301	.103	.412
MMI-antiLM	.275	.186	.467
MMI-context	.289	.178	.478

Table 1: Performance on test dataset of baseline as well as two MMI models

As seen in table 1, both MMI-antiLM and MMI-context have little bit lower BLEU-I scores than SEQ2SEQ baseline. What's more, MMI-antiLM and MMI-context achieve better performance on both distinct uni-grams and bi-grams measurements. However, information relevance could hardly be illustrated through automatic metrics, thus we turn to qualitative analysis.

Message	i take my leave
SEQ2SEQ	my life is quite clear
MMI-antiLM	until i get you
MMI-context	and i will take care of it
Message	you better stay right there boy
SEQ2SEQ	i am not a monster but i am a animal
MMI-antiLM	boy just a moment
MMI-context	all right...

Table 2: Responses generated from three models

After collecting samples, we find that replies from MMI models are more informative, and also MMI context model shows more relation between message and response.

5 DISCUSSION

Given that it is hard to measure automatically the performance of open-domain dialogue system, most of currently used metric evaluation can

only represents very limited aspects. From machine's perspective, it is worth-while to find trade-off between response coherency, diversity and relevance. This work only attempts to walk in the very beginning of this interesting field, remaining lots of proposed extensions and unknown architecture to explore, including end-to-end system or even more complicated ones using reinforcement learning etc.

Reference

[1] Sordoni, Alessandro, et al. "A neural network approach to context-sensitive generation of conversational responses." arXiv preprint arXiv:1506.06714 (2015).

[2] Vinyals, Oriol, and Quoc Le. "A neural conversational model." arXiv preprint arXiv:1506.05869 (2015).

[3] Serban, Iulian Vlad, et al. "Building End-To-End Dialogue Systems Using Generative Hierarchical Neural Network Models." AAAI. Vol. 16. 2016.

[4] Li, Jiwei, et al. "A diversity-promoting objective function for neural conversation models." arXiv preprint arXiv:1510.03055 (2015).

[5] Li, Jiwei, Will Monroe, and Dan Jurafsky. "A simple, fast diverse decoding algorithm for neural generation." arXiv preprint arXiv:1611.08562 (2016).

[6] Carbonell, Jaime, and Jade Goldstein. "The use of MMR, diversity-based reranking for re-ordering documents and producing summaries." Proceedings of the 21st annual international ACM SIGIR conference on Research and development in information retrieval. ACM, 1998.

[7] Gimpel, Kevin, et al. "A systematic exploration of diversity in machine translation." Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing. 2013.

[8] Bahdanau, Dzmitry, Kyunghyun Cho, and Yoshua Bengio. "Neural machine translation by jointly learning to align and translate." arXiv preprint arXiv:1409.0473 (2014).

[9] Tang, Duyu, Bing Qin, and Ting Liu. "Document modeling with gated recurrent neural network for sentiment classification." Proceedings of the 2015 conference on empirical methods in natural language processing. 2015.

[10] Luong, Minh-Thang, Hieu Pham, and Christopher D. Manning. "Effective approaches to attention-based neural machine translation." arXiv

preprint arXiv:1508.04025 (2015).

[11] Yang, Zichao, et al. "Hierarchical attention networks for document classification." Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies. 2016.

[12] Sutton, Richard S., et al. "Policy gradient methods for reinforcement learning with function approximation." Advances in neural information processing systems. 2000.

[13] Tiedemann, Jrg. "News from OPUS-A collection of multilingual parallel corpora with tools and interfaces." Recent advances in natural language processing. Vol. 5. 2009.