

Pattern-aware Data Augmentation for Query Rewriting in Voice Assistant Systems

Yunmo Chen^{1*} Sixing Lu² Fan Yang² Xiaojiang Huang² Xing Fan² Chenlei Guo²
¹Johns Hopkins University ²Amazon.com, Inc. USA

Abstract

Query rewriting (QR) systems are widely used to reduce the friction caused by errors in a spoken language understanding pipeline. However, the underlying supervised models require a large number of labeled pairs, and these pairs are hard and costly to be collected. Therefore, We propose an augmentation framework that learns patterns from existing training pairs and generates rewrite candidates from rewrite labels inversely to compensate for insufficient QR training data. The proposed framework casts the augmentation problem as a sequence-to-sequence generation task and enforces the optimization process with a policy gradient technique for controllable rewarding. This approach goes beyond the traditional heuristics or rule-based augmentation methods and is not constrained to generate predefined patterns of swapping/replacing words. Our experimental results show its effectiveness compared with a fully trained QR baseline and demonstrate its potential application in boosting the QR performance on low-resource domains or locales.

Introduction

Spoken language understanding (SLU) is widely used in voice assistant systems such as Amazon Alexa and Google Home, to extract the semantic information from an input voice query. Two main components in the SLU system, automatic speech recognition (ASR) and natural language understanding (NLU) might introduce errors from many resources such as the speaker’s accent or semantics ambiguity. These errors cascade and result in user dissatisfaction. Quite a few works are focusing on friction correction on the ASR component or the NLU component, including the query rewriting (QR) system (Chen et al. 2020) that leverages a deep learning architecture to handle longer context with little feature engineering. However, supervised QR learning requires a large quantity of paired utterance strings as training labels, which are usually collected through human annotation or the user’s own rephrases. Such pairs are limited for small domains or locales that lack data. Moreover, the annotation cost is significant given the high demand of annotations.

* The work was completed when the author was an intern at Amazon.

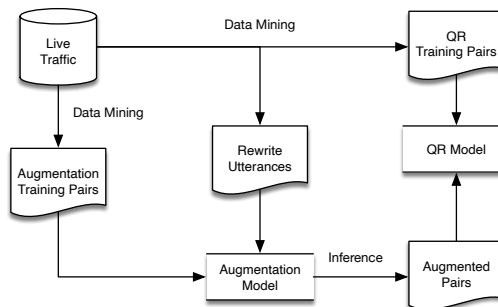


Figure 1: Augmentation System Design Overview.

To conquer the data scarcity problem, data augmentation has been studied in natural language processing (NLP). Approaches spanning from deterministic text editing to neural sequence-to-sequence modeling emerged in the past few years. (Wei and Zou 2019, *inter alia*) employed text editing techniques, where a series of candidate operations including replacement, insertion, swap, and deletion are performed to modify text sequences and introduce noise. But such pre-defined operations are limited to certain modifications and prone to generating unrealistic data examples. (Sennrich, Haddow, and Birch 2016, *inter alia*) showed the effectiveness of a back-translation method to augment neural machine translation with additional parallel data generated from monolingual corpora. (Culkin et al. 2020, *inter alia*) investigated paraphrastic augmentation which automatically expands the overall sizes and syntactic diversity via a paraphrasing model. Although their synthetic data is meaningful in semantics, they do not model the discrepancy between friction information thus cannot model specific error patterns in the QR scenario.

We cast the data augmentation task as a machine translation (MT) task, augmenting training pairs by translating from existing utterances with low friction (seen as rewrites) to the potential corrupted utterances (seen as requests). This framework is an inverse process of QR that denoises requests to rewrites. It follows the nature of voice assistant systems where all functionalities are usually limited and can be easily covered by existing correct utterances, while the

corrupted utterance varies due to the large variety of noise sources.

In addition, the designed framework not only boosts the quantity of training data but also extracts specific error patterns without underlying rules. For example, requests with ASR errors tend to be similar in phonemes but different in semantic meanings. To this end, we optimize the generation model by employing a reward that measures a underlying error pattern (e.g. phonetic dissimilarity for ASR errors). The policy gradient technique is adopted to optimize such discrete rewards along with the maximum likelihood estimation.

We evaluate our proposed approach on English datasets. As the approach is agnostic to a specific language, it can be easily applied to other languages or even multilingual settings. To summarize, our contributions are as follows:

1. Design a learning-based data augmentation framework that boosts training data pairs without human annotations, reducing the cost for QR system.
2. Optimize the sequence generation through predefined rewards that measures distance between request and rewrite, Particularly phoneme and semantic distance in this paper, as they represent two main friction root causes in a spoken dialog system.

Approach

Problem formulation

For QR systems, each training example can be denoted as a pair of request utterance $U = (u_1, \dots, u_n)$ and rewrite utterance $R = (r_1, \dots, r_m)$, where $u_i \in \mathcal{V}$ and $r_j \in \mathcal{V}$. \mathcal{V} is a token vocabulary. As opposite to the QR that learns a function of $U \rightarrow R$, our data augmentation task is to learn a mapping $R \rightarrow U$. Under this mapping, a U is guaranteed to be paired with its corresponding input R , which can thus be used as a training example for QR. In addition, as both U and R are sequences, we leverage sequence-to-sequence technique to model the generation process.

Data augmentation as machine translation

We adopt the Transformer (Vaswani et al. 2017) architecture to model the conditional generation process and take a rewrite utterance R as input. We first encode the utterance using the Transformer encoder so that each token r_j has one vector representation as:

$$\mathbf{r}_j = \text{Emb}(r_j) \in \mathbb{R}^{d_{\text{tok}}} \quad (1)$$

Therefore, the rewrite utterance becomes a sequence of vectors $\mathbf{R} = (\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_m)$. A Transformer decoder takes the encoded rewrite sequence and autoregressively computes hidden states denoted as $\mathbf{h}_t \in \mathbb{R}^{d_{\text{hid}}}$. For each hidden state, we pass the vector \mathbf{h}_t to a linear layer with the output size being the size of the vocabulary \mathcal{V} . Softmax is applied to the output of size $|\mathcal{V}|$, yielding a distribution over the set of tokens:

$$p_\theta(u_t | u_1, \dots, u_{t-1}) = \frac{\exp \mathbf{w}_u^T \mathbf{h}_t}{\sum_{u' \in \mathcal{V}(e)} \exp \mathbf{w}_{u'}^T \mathbf{h}_t} \quad (2)$$

The log probability of the target sequence can thus be written as:

$$\begin{aligned} \log p_\theta(u_1, u_2, \dots, u_T) &= \log \prod_{t=1}^T p_\theta(u_t | u_1, \dots, u_{t-1}) \\ &= \sum_{t=1}^T \log p_\theta(u_t | u_1, \dots, u_{t-1}) \end{aligned} \quad (3)$$

Cross-entropy loss is used to minimize the negative log probability of the target sequence given the model parameters:

$$\begin{aligned} L(\theta) &= -\log p_\theta(\mathbf{U}) \\ &= -\log p_\theta(u_1, u_2, \dots, u_T) \end{aligned} \quad (4)$$

During the training, we incorporate the teacher forcing algorithm (Williams and Zipser 1989) to force the model to be exposed to the ground truth tokens. At the inference stage, we employ beam search to reduce search errors caused by the discrepancy between training and inference.

Pattern-aware sequence generation

To introduce pattern-aware objectives into the learning task, we employ policy gradient algorithm to approximate gradients with respect to non-differentiable metrics. Upon the end of the sequence generation process, encountering with [EOS] (end of sequence) token, the designated reward will be computed, denoted as r . We compute the reward in the sequence-level by comparing the pair of target sequence(s) and source sequence. Under this formulation, the optimization objective is to minimize the negative expected reward:

$$L(\theta) = -\mathbb{E}_{\mathbf{U} \sim p_\theta} [r(\mathbf{U}, \mathbf{R})] \quad (5)$$

We incorporate the self-critic sequence training (SCST) (Rennie et al. 2017) to approximate the gradient $\nabla_\theta L(\theta)$. We also set a reward baseline that is independent from the generated sequence to stabilize the policy gradient approximation. The reward obtained by the current model used at test time is set as a baseline $r(\hat{\mathbf{U}})$. The gradient is then approximated as the relative reward value against the baseline:

$$\nabla_\theta L(\theta) = -\mathbb{E}_{\mathbf{U} \sim p_\theta} [(r(\mathbf{U}, \mathbf{R}) - r(\hat{\mathbf{U}}, \mathbf{R})) \nabla_\theta \log p_\theta(\mathbf{U})] \quad (6)$$

Reward design

We design the following three reward functions. All r are normalized within $r \in [0, 1]$.

Phonetic reward This reward is to learn the pattern of error cascaded from the ASR system. To compute the reward r_p , we use an out-of-box grapheme-to-phoneme (G2P) model (Park 2019) to convert the decoded target sequence \mathbf{U} and the source sequence \mathbf{R} to their corresponding phoneme sequences, \mathbf{U}_p and \mathbf{R}_p respectively. We compute the normalized Levenshtein distance of the phoneme sequences as:

$$\begin{aligned} r_p(\mathbf{U}, \mathbf{R}) &= \text{LevDist}(\text{G2P}(\mathbf{U}), \text{G2P}(\mathbf{R})) \\ &= \text{LevDist}(\mathbf{U}_p, \mathbf{R}_p) \end{aligned} \quad (7)$$

Dataset	Split	# of examples
Augmentation	train	3,471,063
	dev	433,377
	test	434,233
InfRewrite*		3,844,430

Table 1: Augmentation dataset statistics. * denotes the inference data that will be used to generate requests to be paired as training data for QR system training.

Utterance dissimilarity reward We adopt a pre-trained semantic encoder that outputs a probability of the similarity between two utterances.¹ This encoder embeds the request utterance and rewrite utterance as two vectors, $\mathbf{u}_d, \mathbf{r}_d \in \mathbb{R}^d$ respectively and uses cosine distance to measure the similarity between these two vectors. The dissimilarity reward is thus denoted as:

$$r_d(\mathbf{U}, \mathbf{R}) = 1 - \cos(\mathbf{u}_d, \mathbf{r}_d) \quad (8)$$

Combined reward We also investigate a linear combination of the phonetic reward and utterance dissimilarity reward. This design is motivated by the error patterns that are similar in phonemes but different in semantics. The combination is controlled by a factor $\alpha \in [0, 1]$. The combined reward is:

$$r_c(\mathbf{U}, \mathbf{R}) = \alpha r_p + (1 - \alpha) r_d \quad (9)$$

Experiments and Discussions

Training data for augmentation model

The training set for the augmentation model are anonymized rewrite-request pairs. These pairs are constructed under a data mining process using ASR nth best transcripts from hundreds of domains. In such pairs, the best ASR string from ASR model serves as the “rewrite”; “requests” are corrupted wrong ASR output strings of the same user utterance. The statistics of this dataset are shown in Table 1 named as *Augmentation* dataset. This dataset is in the form of $(\mathbf{U}^*, \mathbf{R})$ tuple, where the \mathbf{U}^* is treated as a ground truth sequence that the model learns from.

At the inference and augmentation stage, the input rewrites are utterances that successfully processed by the Amazon Alexa and confirmed with positive user feedbacks. Those utterances are considered to be “golden” rewrite utterances from the history. Note that the “golden” inputs do not have the scarcity problem so that the augmentation model can produce as many output pairs as needed. For fair comparison, we only adopt rewrites in the evaluation QR system training set as the “golden” inputs. The statistics of the data are also shown in Table 1 named as *InfRewrite* dataset.

The test data in the paper includes two sets: one is sampled from friction traffics to specifically measure ASR error recovery, and the other adopts users own rephrases² as la-

¹Note that any other semantic models/functions can substitute it.

²The pairs of one user’s first utterance and following rephrased utterances are also annotated by human judges.

Reward	Rewrite (input)	Request (output)
Phonetic	<i>what is your favorite clothes</i>	<i>what is your favorite close</i>
	<i>play love by wale</i>	<i>play law by wale</i>
	<i>what does a dog sound like</i>	<i>what does a dug sound like</i>
Semantic	<i>when does summer end in this year</i>	<i>when does summer this year</i>
	<i>one hundred seventeen times two</i>	<i>one hundred and seventy times two</i>
	<i>air conditioning in the summer</i>	<i>conditioning in the summer</i>
Combined	<i>is today the last day of ramadan</i>	<i>is today the last of ramadan</i>
	<i>turn your own volume down</i>	<i>turn on your own volume down</i>
	<i>i do not appreciate you</i>	<i>i don't not appreciate you</i>

Table 2: Sample inputs and outputs from the augmentation models trained with different rewards.

els. Both of them are very challenging testsets with around 10K samples for the evaluation of QR performance.

Evaluation of QR retrieval system

The underlying end-to-end evaluation system is a dense retrieval model powered by faiss(Johnson, Douze, and Jégou 2017)³. The search system indexes among user satisfied utterances from the live traffic in one month. We augment the training set with synthetic data generated by our augmentation model. We then train a deep neural network on it. The training objective is to minimize the distance of the input pair and maximize the distances to its neighbor pairs. We use the precision within top rank K position (P@K) as the metric.

Baselines The baseline uses the *train* split in the *Augmentation* dataset only, without any synthetic data generated by the augmentation model. The baseline indicates how the end-to-end QR model performs with limited existing training pairs.

Training We trained four augmentation models: 1) a fine-tuned BART translation model (Lewis et al. 2020) without using any rewards; 2) a Policy Gradient based (PG-based) sequence-to-sequence model using the phonetic reward; 3) a PG-based sequence-to-sequence model using the utterance dissimilarity reward; 4) a PG-based sequence-to-sequence model using the combined reward.

We use a maximum number of tokens at 1024 to construct a mini-batch and set the maximum utterance length to be 25 tokens. The learning rate is set to 3×10^{-5} and warm up for 200 steps. The total training step is set to 20,000 with early stopping. For *combined reward* experiment, the weight between two rewards is set to be $\alpha = 0.5$. We use 7 AWS EC2 instances with a total of 56 Tesla V100 GPUs for training. All experiments are done using fairseq(Ott et al. 2019).

Results and analysis

Table 2 shows sample inputs and outputs from the trained augmentation models. The experimental evaluation results are summarized on Table 3 and Table 4⁴. There are five cat-

³The retrieval model is only one of many models used by the whole QR system and does not reflect any actual production model or performance of the Alexa product.

⁴Due to business reasons, relative improvements are indicated in both tables. All numbers are statistically significant. The baseline number for Table 3 and 4 is in lower 50s.

Augment Model	QR training input	Data Size	p@1	p@5
Baseline	Training data only	3.45M	0%	0% (range of 50%)
Translation model	Training+Synthetic	7.28M	3.7%	2.1%
phonetic reward	Training+Synthetic	7.35M	6.5%	2.4%
semantic reward	Training+Synthetic	7.34M	5.0%	2.5%
combined reward	Training+Synthetic	7.35M	3.7%	0.8%

Table 3: Experimental results on the friction testset. precision is relative number. “QR training input” indicates what is the training data for the downstream QR model: *Training only* only uses the *train* split of the *Augmentation* dataset; *Training+Synthetic* combines the synthetic data and the *train* split. “Size” indicates the resulted numbers of training examples.

Augment Model	QR training input	Data Size	p@1	p@5
Baseline	Training only	3.45M	0%	0% (range of 50%)
phonetic reward	Training+Synthetic	7.35M	0%	0.1%
semantic reward	Training+Synthetic	7.34M	3.3%	1.3%
combined reward	Training+Synthetic	7.35M	2.5%	3.4%

Table 4: Experimental results on the user rephrase testset. precision is relative number.

egories of training data for QR evaluation system: 1) baseline QR precision trained with no synthetic data; 2) QR precision trained with the merged synthetic data from the BART augmentation model; 3) QR precision trained with the merged data from the PG-based phonetic reward augmentation model; 4) QR precision trained with the merged data from PG-based utterance dissimilarity reward augmentation model; 5) QR precision trained with the merged data from the PG-based combined reward augmentation model.

The complete QR system constitutes of many more precision layers to serve customers. We only select one of the retrieval models to serve the purpose of proof-of-concept, hence relatively low baseline numbers for precision.

Comparing the “Translation model” with the baseline in Table 3, there is 3.7% improvement in P@1 and 2.1% improvement in P@5, which demonstrate the power of the synthetic data added. Noting that the augmentation model input is the same as the original QR training pairs, it only boosts the original label’s corrupted patterns. In practice, any “golden” utterances can be used as inputs and they will largely increase the potential pairs out of limited scopes.

Among the three PG-based augmentation models, the phonetic reward augmentation model performs the best. It improves 6.5% P@1 and 2.4% P@5 using the merged data. The main reason is that the first testset emphasizes on evaluating ASR friction cases and maximizing phonetic distance represents the ASR error best. Minor difference on the phonetic string may not generate corrupted utterances that have a large intention gap.

The user rephrase testset also demonstrates the power of synthetic supplements as shown in Table 4. The rephrase set emphasizes rewrites that are more likely being semantic rephrasing and intention switching. The semantic reward augmentation model achieves better performance on this testset. The merged training data increases the precision P@1 by 3.3% and P@5 by 1.3%. One may notice that the

precision improvement is not huge. This is because our exploration for such augmentation models is on English whose original training dataset might reach the capacity of the simple evaluation model. However, even for such almost fully trained QR models, additional synthetic training data still proves to catch additional friction patterns under different and challenging evaluation sets. We believe this augmentation framework will benefit more for locales and domains with less resource, and we also would like to address such issues in the future work.

Conclusion and Future work

This work proposes a data augmentation framework targeting at boosting high-quality training data for query rewriting systems. The experimental results show the effectiveness of our approach. In the future, we will explore different problem formulations such as text refinement in the augmentation process, bringing in more utterance variations. Different reward designs are another promising direction to explore. Moreover, low-resource settings in domains or locales are also worth investigation.

References

- Chen, Z.; Fan, X.; Ling, Y.; Mathias, L.; and Guo, C. 2020. Pre-training for query rewriting in a spoken language understanding system. *ICASSP 2020*.
- Culkin, R.; Hu, J. E.; Stengel-Eskin, E.; Qin, G.; and Durme, B. V. 2020. Iterative paraphrastic augmentation with discriminative span alignment. *CoRR* abs/2007.00320.
- Johnson, J.; Douze, M.; and Jégou, H. 2017. Billion-scale similarity search with gpus. *arXiv:1702.08734*.
- Lewis, M.; Liu, Y.; Goyal, N.; Ghazvininejad, M.; Mohamed, A.; Levy, O.; Stoyanov, V.; and Zettlemoyer, L. 2020. BART: denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In *ACL 2020*.
- Ott, M.; Edunov, S.; Baevski, A.; Fan, A.; Gross, S.; Ng, N.; Grangier, D.; and Auli, M. 2019. fairseq: A fast, extensible toolkit for sequence modeling. In *NAACL-HLT 2019*.
- Park, Kyubyong Kim, J. 2019. g2pe. github.com/Kyubyong/g2p.
- Rennie, S. J.; Marcheret, E.; Mroueh, Y.; Ross, J.; and Goel, V. 2017. Self-critical sequence training for image captioning. In *CVPR 2017*.
- Sennrich, R.; Haddow, B.; and Birch, A. 2016. Improving neural machine translation models with monolingual data. In *ACL 2016*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; and Polosukhin, I. 2017. Attention is all you need. In *NuerIPS 2017*.
- Wei, J. W., and Zou, K. 2019. EDA: easy data augmentation techniques for boosting performance on text classification tasks. In *EMNLP-IJCNLP 2019*.
- Williams, R. J., and Zipser, D. 1989. A learning algorithm for continually running fully recurrent neural networks. *Neural Comput.*