On the Generation of Medical Question-Answer Pairs

Sheng Shen,¹ Yaliang Li,² Nan Du,³ Xian Wu,³
Yusheng Xie,³ Shen Ge,³ Tao Yang,³ Kai Wang,³ Xingzheng Liang,³ Wei Fan³

¹University of California at Berkeley, ²Alibaba Group, ³Tencent

sheng.s@berkeley.edu, yaliang.li@alibaba-inc.com,

{ndu, kevinxwu, yushengxie, shenge, tytaoyang, ironswang, evelynliang, Davidwfan}@tencent.com

Abstract

Question answering (QA) has achieved promising progress recently. However, answering a question in real-world scenarios like the medical domain is still challenging, due to the requirement of external knowledge and the insufficient quantity of high-quality training data. In the light of these challenges, we study the task of generating medical QA pairs in this paper. With the insight that each medical question can be considered as a sample from the latent distribution of questions given answers, we propose an automated medical QA pair generation framework, consisting of an unsupervised key phrase detector that explores unstructured material for validity, and a generator that involves a multi-pass decoder to integrate structural knowledge for diversity. A series of experiments have been conducted on a real-world dataset collected from the National Medical Licensing Examination of China. Both automatic evaluation and human annotation demonstrate the effectiveness of the proposed method. Further investigation shows that, by incorporating the generated QA pairs for training, significant improvement in terms of accuracy can be achieved for the examination QA system.¹

Introduction

Due to the remarkable breakthrough of deep learning and natural language processing, question answering (QA) has gained increasing popularity in the past few years. Among QA's broad application domains, medical QA is one of the most appealing real-world application scenarios: People tend to consult others about health-related issues on online communities, which might be more affordable than visiting doctors in resource-limited areas.

Although QA systems with deep learning methods have achieved good performance, medical QA confronts particular difficulties against other domains. First, medical QA system requires highly accurate answers, and thus external and professional knowledge gathered from various sources are needed. Second, the size of available high-quality medical QA pairs is limited, as the labeling process by medical experts is time-consuming and expensive. Therefore, the performance of medical QA system is further constrained by the

paucity of high-quality QA pairs since it can hardly learn a good model from limited training data. Though (Roberts et al. 2017; Pampari et al. 2018) aim to enrich the dataset itself, but the efforts are still far from enough.

To tackle these difficulties, the generation of medical QA pairs plays an indispensable role. By automatic generation of high-quality medical QA pairs, external and professional knowledge can be incorporated, and the size of training data can be augmented. Therefore, we study this important task of medical QA pair generation in this paper. To be more specific, we assume that each medical answer corresponds to a distribution of valid questions, which should be constrained on external medical knowledge. Following this assumption, with more high-quality QA pairs generated based on the same knowledge as original QA pairs, the latent distribution of available medical QA pairs can be supplemented and thus medical QA system could learn unbiased model easily.

However, the generation of new medical QA pairs based on original ones is challenging: It is hard to simultaneously maintain the diversity and the validity of generated question-answer pairs. Existing question-answer pair generation methods (Yang et al. 2017; Song et al. 2018) either has external context to build upon or (Duan et al. 2017; Du and Cardie 2018; Yang et al. 2017) focused more on the word-level similarity, and it may generate lexically similar question-answer pairs to the original ones. These generated similar QA pairs are valid but of limited use for allowing the system to answer questions involving new knowledge. On the other hand, if more diversity in the discourse/sentence level is promoted, validity might not be guaranteed.

To ensure the validity of the generated medical QA pairs, we propose a retrieval and matching method to detect the key information of QA pairs in an unsupervised way using unstructured text materials such as patients' medical records, textbooks, and research articles.

To promote the diversity of the generated medical QA pairs while retaining validity, we propose two mechanisms to incorporate structured, unstructured knowledge for QA generation. We first explore global phrase level diversity and validity with a hierarchical Conditional Variational Autoencoder (CVAE) framework, which models phrase level relationship in original medical QA pairs, and generates the new pairs without breaking these relationships. We then propose a multi-pass decoder, in which all the local components

Figure 1: Overview of the proposed framework. Note that this question consists of N phrases and this figure shows the process where we are generating the k -th phrase.

(phrase type, entities in each phrase) are coupled together and are jointly optimized in an end-to-end fashion.

In order to demonstrate the effectiveness of the proposed generation method, we evaluate generated medical QA pairs through qualitative and quantitative measures, and the results confirm the high-quality of the generated medical QA pairs. Further, in an application of the proposed method to a medical certification exam, the experimental results show that the generated medical QA pairs improve the original QA system by six percent question-level accuracy.

Methodology

In this section, we introduce our framework for generating medical question-answer pairs based on existing pairs. For medical QA, we assume the same answer can be produced by multiple questions, for example, patients of *stiff neck(+)* with *pap test(+)* or *respiratory failure* can be diagnosed as the disease *Japanese encephalitis* due to the diversity of medical characteristics, while for a specific medical question, there is only one correct answer. Hence, we view the generating process of medical QA pairs as generating questions given a certain answer. Technically speaking, our framework for generating medical QA pairs can be considered as an approximation of the latent distribution of questions given answers and sampling new questions from the distribution. As shown in Fig 1, the whole framework involves a key phrase detector and an entity-guided CVAE based generator (eg-CVAE), which we describe in detail in the following subsections. Both the original QA pairs and the generated ones from our framework will be fed into the QA system as inputs for training.

Key Phrase Detector

In order to approximate the unknown conditional distribution of medical questions given answer, we leverage external knowledge to exploit the intrinsic characteristics of medical questions that associate with the same answer. Specifically, every medical question Q consists of several phrases

$P_k, k \in [1, N]$, such as patient's symptoms, examination results. Each phrase is composed of several words. Among medical questions, there exist *key phrases* highly correlated with the answers (denoted as P'_k like *stiff neck(+)* in Fig 1). To detect the prior key phrases, we employ an unsupervised matching approach on unstructured medical text. Furthermore, to ensure the consistency of these key phrases in the generated new questions, we assign each phrase with a normalized significance score $s_k \in [0, 1]$, which is further used as the probability of replacing this phrase by the generated one or not in the generation process.

Rather than considering each phrase separately, we assume that the co-occurrence probability of a key phrase and answer indicates the significance of that phrase. To explore this co-occurrence information, we first use each medical QA pair as query to perform an Elasticsearch² (Gormley and Tong 2015) based retrieval over the medical materials. We also apply rules to ensure the presence of the answer in retrieved texts, denoted as $R_i, i \in [1, M]$ (M stands for the number of retrieved texts). An unsupervised matching strategy is proposed to model the relevance of a certain phrase P_k with the answer by matching P_k with all the R_i . Specifically, we divide each R_i into phrases P^{R_i} (each phrase contains multiple words), and represent each P^{R_i} and P_k into the same vector space. To produce that vector, we perform a hierarchical pooling over the word embedding $v_j, j \in [1, L]$ in that phrase following (Shen et al. 2018): first, average pooling over $v_{j,j+k-1}, j \in [1, L-k+1]$ within each sliding window (size is k); then, max pooling over the induced average-pooling vectors. We match every phrase $P_k, k \in [1, N]$ with the phrase splits $P^{R_i, i \in [1, M]}$ using cosine distance and store the highest score $s_k^{R_i}$. The unnormalized matching score for P_k with R is the mean value of $s_k^{R_i}, i \in [1, N]$. These scores for each phrase P_k in the QA pair will be normalized as the $s_k, s_k \in [0, 1]$ for final sampling decision with the Min-Max method. Specifically, in inference, we randomly samples

²<https://github.com/elastic/kibana>

Figure 2: Entity-guided CVAE based Generator. In this figure, we illustrate the detailed process to generate current phrase_k based on previous altered phrases_{1,...,k-1, k+1,N}.

$p'_k \in [0, 1]$. Then if $p'_k > s_k$ we will replace P_k with the generated phrase or retains P_k .

Entity-guided CVAE based Generator

A medical question has two levels of structures: one structure exists within a single phrase, which is dominated by local information of involved medical entities, and the other is a distinct across-phrase structure, which is characterized by aspects such as phrase types and the corresponding answer etc.. We thus explore the answer conditioned medical question generation in a two-level hierarchy: sequences of subsequences (iterative phrase generation process), and subsequences of words. Towards modeling the constraint over the whole question, we first use Conditional Variational Autoencoder. Moreover, towards modeling the internal structure within each phrase, we draw the idea from human's process to generate a complete question (start from a sketch and then details), and introduce a three-pass decoding process: first implicit type modeling, then explicit entities modeling, and finally phrase decoding.

Conditional Variational Autoencoder Motivated by (Serban et al. 2017), we adapt the original CVAE for dialog generation to our setting by considering question generation as an iterative phrase generation process in Figure 2. To this end, we represent each phrase generation procedure with three random variables: the phrase context c , the target phrase x , and a latent variable z that is used to capture the latent distribution over all valid phrases. For each phrase, c is composed of both the sequence of other phrases in the question and the corresponding answer. We then define the conditional distribution $P(x, z|c) = P(x|c, z) \cdot P(z|c)$ and set the learning target is to approximate $P(z|c)$ and $P(x|c, z)$ via deep neural networks (parametrized by θ). We refer to $P_\theta(z|c)$ as the prior network and $P_\theta(x|c, z)$ as the target phrase decoder. Then the generative process of x is summarized as first sampling a latent variable z from $P_\theta(z|c)$ (a parametrized Gaussian distribution.) and then generating x by $P_\theta(x|c, z)$.

The CVAE is trained to maximize the conditional log

likelihood of x given c , meanwhile minimizing the KL divergence between the posterior distribution $P(z|x, c)$ and a prior distribution $P(z|c)$. We assume that both z follow multivariate Gaussian distribution with a diagonal covariance matrix. Further, we introduce a recognition network $Q_\phi(z|x, c)$ to approximate the true posterior distribution $P(z|x, c)$. As proposed in (Sohn, Lee, and Yan 2015), CVAE can be efficiently trained with the *Stochastic Gradient Variational Bayes* (SGVB) framework (Kingma and Welling 2013) by maximizing the variational lower bound of the conditional log likelihood, which can be written as:

$$\mathcal{L}(\theta, \phi; x, c) = -KL(Q_\phi(z|x, c)||P_\theta(z|c)) + E_{Q_\phi(z|x, c)}[\log P_\theta(x|c, z)]. \quad (1)$$

At timestamp k of the whole generation process to produce a question phrase, the phrase encoder is a bidirectional recurrent neural network (Schuster and Paliwal 1997) with a gated recurrent unit (GRU (Chung et al. 2014)) to encode each phrase P_k into a fixed-size vector by concatenating the last hidden states of the forward and backward RNN as $[hv_k, hv_k]$. This basic phrase context encoder is a one-layer GRU network that encodes the $N - 1$ context phrases (in training, the context phrases are from the original question; in testing, the preceding $k - 1$ phrases are from the generated question) as $hv_{1:k-1}$ with $hv_{k+1:N}$. The last hidden state hv^c of the phrase context encoder is concatenated with the corresponding answer embedding a and $c = [hv^c, a]$. As we assume z follows an isotropic Gaussian distribution, the recognition network $Q_\phi(z|x, c) \sim N(\mu, \sigma^2 I)$, the prior network $P_\theta(z|c) \sim N(\mu', \sigma'^2 I)$, and then we get:

$$\begin{bmatrix} \mu \\ \log(\sigma^2) \end{bmatrix} = W_r \begin{bmatrix} x \\ c \end{bmatrix} + b_r, \begin{bmatrix} \mu' \\ \log(\sigma'^2) \end{bmatrix} = MLP_p(c). \quad (2)$$

The reparameterization trick (Kingma and Welling 2013) that uses formed parameter to treat z as deterministic node is adopted to get samples from $N(z; \mu, \sigma^2 I)$ in training (recognition network) and from $N(z; \mu', \sigma'^2 I)$ in testing (prior network). The final phrase decoder at timestamp k is a one-layer GRU network with initial state set as $W_k[z, c] + b_k$.

The words will be predicted sequentially by the phrase decoder.

Phrase-type Augmented Encoder Inspired by (Parvez et al. 2018)’s insights to facilitate text generation with entity type, we similarly introduce phrase type in the medical domain as a similar source of structural information (the intuition behind specific phrases such as lab examination and physical characteristics employed by doctors). Rather than focusing on word level, we assume each phrase information involves two levels of characteristics: 1) global characteristic as the surrounding or context phrases’ type information; 2) local characteristic as entity type knowledge within each phrase. Moreover, to address the difficulty of acquiring labeled data from experts, we propose to directly utilize a structured entity dictionary and model the phrase type in a contextualized way following (Peters et al. 2018).

To this end, we design a sequence labeling task for pre-training, whose learning goal is to predict each word’s type (for those words not in the entity dictionary, the type is considered as “other”) over the whole question.

A Bi-LSTM-CRF model, which takes each word’s embedding in the question as input and their types as output, is applied in the pre-training task. We use Bi-LSTM layer to encode word-level local features, and CRF layer to capture sentence-level type information. As the pre-training task’s accuracy can achieve 97.08%, we assume that the hidden states of Bi-LSTM for each word k as $h_k[\vec{h}_k, \overleftarrow{h}_k]$ can encode the contextualized type information. Considering that each phrase can be split into multiple words, the phrase type information is introduced by performing max-pooling over each word’s h_k . We then concatenate contextualized type vector t_k at timestamp k to generate phrase type vector $hv'_k = [hv_k, t_k]$ for P_k (clustering as 6 T_* in Figure 2). t_k is pre-trained through the sequence labeling task, and different for each timestamp of the whole generation procedure. The new $x' = hv'_k$ will be then applied for the recognition network.

Entity-guided Decoder Other than only conditioning on the corresponding answer, we introduce extra constraints on latent z to keep it meaningful during decoding process. Drawn the insights from the process of human generating a complete question (start from a sketch and then details) in (Xia et al. 2017), we propose a multiple pass decoding procedure to incorporate inter-phrase level and intra-phrase level information as constraints. We thus model the contextualized type t , which is imposed by the entity dictionary, at the first pass to ensure the consistency of type information across phrases. We then conjecture entities to be the skeleton within each phrase, and explicitly model entities e at the second pass. We promote diversity in our generation process by adding entity-level variation during inference, allowing the production of phrases with similar semantics towards the same answer but containing diverse entities.

We assume that the generation of phrase P_k as x depends on c, z, t and e ; e relies on c, z, t ; and t relies on c, z . During training, the initial state of the final decoder is $d_k = W_k[z, c, t, e] + b_k$ and the input is $[w_{1:n^k}, t, e_k]$

where $w_{1:n^k}$ is the word embedding of words in x and e_k is average pooling embedding of the entire entity embedding in x . In the first type-prediction pass, there is an MLP to predict $t' = MLP_t(z, c)$ based on z and c . In the second entity-prediction pass, another MLP is used to predict $e_{softmax'} = MLP_e(z, c, t)$ based on z, c and t . Then $e_{softmax'}$ is multiplied with the whole entity embedding matrix for the aggregation of the e'_k . In the testing stage, the predicted t' and e'_k are used in the final phrase decoder.

Training Objective

To induce meaningful latent variable z , we explicitly model the generation of x as a multi-pass process, which might relieve the posterior collapse problem (He et al. 2019) motivated by (Zhao, Zhao, and Eskenazi 2017) in enriching the information in posterior distribution of z with dialog actions.

Specifically, by introducing phrase-type information in the first pass, we suppose that the generation of x is based on c, z and t , where t is based on c . Then the modified variational lower bound for eg-CVAE without entity modeling:

$$\begin{aligned} \mathcal{L}(\theta, \phi; x, c, t) = & -KL(Q_\phi(z|x, c, t)||P_\theta(z|c)) \\ & + E_{Q_\phi(z|x, c, t)}[\log P_\theta(t|c, z)] \\ & + E_{Q_\phi(z|x, c, t)}[\log P_\theta(x|c, z, t)]. \end{aligned} \quad (3)$$

To refine phrase-type information into detailed entities in the second pass, we model e explicitly based on the assumption that the produce of x is divided into two phases: exploiting phrase-type to generate e ; and using e, t, c and z to generate x . Thus the final eg-CVAE model is to maximize:

$$\begin{aligned} \mathcal{L}(\theta, \phi; x, c, t, e) = & -KL(Q_\phi(z|x, c, t, e)||P_\theta(z|c)) \\ & + E_{Q_\phi(z|x, c, t, e)}[\log P_\theta(t|c, z)] \\ & + E_{Q_\phi(z|x, c, t, e)}[\log P_\theta(e|c, z, t)] \\ & + E_{Q_\phi(z|x, c, t, e)}[\log P_\theta(x|c, z, t, e)]. \end{aligned} \quad (4)$$

Furthermore, the KL annealing (Serban et al. 2016b) technique as gradually increasing the weight of the KL term from 0 to 1 during training and auxiliary bag-of-words loss of x as in (Zhao, Zhao, and Eskenazi 2017) are also adopted.

Experiments

Dataset

To validate the effectiveness of the proposed method, we collect real-world medical QA pairs from the National Medical Licensing Examination of China (denoted as NM-LEC_QA). The collected NMLEC_QA dataset contains 18,798 QA pairs, and we generate new QA pairs based on these original ones. We adopt NMLEC_2017 as the test set to evaluate the QA system, which will not be used in QA pair generation. The medical entity dictionary is extracted from medical Wikipedia-style pages³, and the constructed dictionary covers 19 types of medical entities. The unstructured medical materials consists of 2,130,128 published paper in medical domain and 518 professional medical textbooks.

³<http://www.xywy.com/>

Table 1: Performance comparison under automatic evaluation metrics.

Method	BLEU			BOW Embedding			intra-dist		inter-dist	
	Precision	Recall	F1	Average	Extreme	Greedy	dist-1	dist-2	dist-1	dist-2
HRED	0.435	0.737	0.547	0.753	0.705	0.809	0.837	0.912	0.205	0.255
VHRED	0.454	0.705	0.533	0.863	0.872	0.887	0.803	0.991	0.562	0.538
type-CVAE	0.507	0.748	0.572	0.872	0.852	0.892	0.831	0.997	0.555	0.581
entity-CVAE	0.541	0.781	0.613	0.891	0.903	0.874	0.840	0.996	0.533	0.554
eg-CVAE	0.450	0.611	0.494	0.802	0.793	0.819	0.867	0.994	0.637	0.589

Baselines

We compare the performance of the proposed method **eg-CVAE** with two recently-proposed text generation methods: *HRED* (Serban et al. 2016a), a sequence-to-sequence model with a hierarchical RNN encoder, and *VHRED* (Serban et al. 2017), a hierarchical conditional VAE model. We also test the contribution of the multiple steps of our decoder of type modeling or entity modeling process: *type-CVAE* with type decoding as the only-pass, and *entity-CVAE* with entity decoding as the only-pass.

Evaluation based on Automatic Metric

Automatically evaluating the quality of generated text remains challenging (Liu et al. 2016), and thus we design automatic evaluation metrics for our specific scenario. As mentioned above, we assume that each QA pair can be considered as a question sampled from a latent answer-conditioned distribution. Based on each original question-answer pair, we generate N new questions by iteratively sampling candidate phrases determined by each s_i and choosing phrases using beam search (Sutskever, Vinyals, and Le 2014). As the generation procedure is at the phrase-level, we evaluate each generated question by comparing the generated phrases with the original and averaging evaluation results over all the phrases in the questions.

We adopt the following three standard metrics to measure the quality of the generated questions from lexical, semantic and diversity perspectives.

- *Smoothed Sentence-level BLEU* (Papineni et al. 2002; Chen and Cherry 2014): BLEU is a popular metric to measure the geometric mean of modified n -gram precision with a length penalty. As N new questions are generated, we define the n -gram precision and n -gram recall as the average and the maximum value of N n -gram BLEU scores respectively. We use 3-gram with smoothing technique, and BLEU scores are normalized to $[0, 1]$.
- *Cosine similarity of Bag-of-words (BOW) embeddings*: a metric matches phrase embeddings through the average, extreme or greedy strategy over all the word embeddings in the phrases (Forgues et al. 2014; Rus and Lintean 2012). The score is the cosine distance between the two produced vectors. We used pretrained embeddings ⁴

⁴Implementation details are in supplementary material. *Average*: cosine similarity between the averaged word embeddings; *Extrema* (Forgues et al. 2014): cosine similarity between the biggest extreme values among the word embeddings of the two phrases;

and denote the three metrics as “Average”, “Extreme” and “Greedy”.

- *Distinct* (Gu et al. 2018): a metric computes the diversity of the generated phrases. The ratio of unique n -grams over all n -grams in the generated phrases is denoted as *distinct- n* . We further define *intra-dist* as the average of distinct values within each sampled phrase and *inter-dist* as the distinct value among all sampled phrases.

We compare the proposed method eg-CVAE with the aforementioned baselines on the collected real-word NM-LEC_QA dataset, and report the experiment results in Table 1. The highest score in each column is in bold for clarity. In the following, we discuss the results in details.

First, we examine the results in terms of similarity using BLEU and BOW metrics. Our proposed method eg-CVAE is designed to promote diversity, and thus the semantic similarity score is not that high. The vanilla CVAE-based VHRED does not involve any constraint on the latent distribution of z , and the HRED (Serban et al. 2016a) models the decoding process in a definite way without further manipulation on the hidden context, so their semantic similarity scores are medium. A variant of the proposed method type-CVAE models prior type information, and another variant entity-CVAE models entity explicitly. These constraints facilitate models to generate more similar QA pairs to the original.

On the other hand, from the view of diversity, the proposed method eg-CVAE has the highest score over distinct metrics. This is because that we hierarchically generate new questions based on the latent answer-conditioned distribution, rather than a definite decoding process. As pointed out in (Serban et al. 2017), this hierarchical strategy can prevent diversity being injected at the low level.

Human Evaluations ⁵

Following (Li et al. 2018), we further conduct human evaluation on 10% samples from NMLEC_QA training dataset and the corresponding generated QA pairs by our methods and baselines. Three experts (real doctors) were asked to assess each QA pair from three perspectives: 1) Consistent: How consistent the generated QA is compared with the original one? 2) Informative: How informative the generated QA

Greedy (Rus and Lintean 2012): matching words in two phrases greedily based on their embeddings’ cosine similarity and averaging the obtained scores.

⁵We also propose a reusable method for evaluation using human annotation of key phrases in supplementary material.

Table 2: Human evaluation results. * indicates the difference between eg-CVAE and other baselines are statistically significant ($p < 0.01$) by two-tailed t-test.

Method	Consist.	Informat.	Fluency
HRED	3.68*	3.38*	3.93*
VHRED	2.79*	3.52	3.79*
type-CVAE	3.53*	3.42	4.03*
entity-CVAE	3.68*	3.38*	4.08*
eg-CVAE	4.09	3.62	4.43

is against the original one? 3) Fluent: How fluent the phrases of a generated question are? Each perspective is assessed with a score from 1 (worst) to 5 (best). The average results are presented in Table 2.

The results show that our model consistently outperforms the seq2seq-baseline model (HRED) and the vanilla CVAE-based method (VHRED). The type-level and entity-level modelings of medical questions make the key information consistent. The prior information from these two levels of modeling also ensures the good ability of our model to generate informative and fluent questions.

Moreover, the implicit type-level modeling via aggregated embedding introduces more variance but less consistence against explicit entity-level modeling via concrete entities, which inspires us to combine them together in the eg-CVAE.

Qualitative Analysis

To further qualitatively analyze the proposed method through real cases, we compare the generated QA pairs from different models in Figure 3. Each example consists of an original valid QA pair and three generated questions, which are sampled based on the raw one through beam-search. We can clearly see our eg-CVAE retains both one-to-many diversity property and validity of each phrase’s generation.

We compared three models here including HRED, CVAE and eg-CVAE.⁶ For HRED, we can observe that the generated questions’ diversity is limited since the model tends to repeat the seed phrases (e.g., the meaningless repetition of “RBC” and “anxiety”) and the important information describing topographic shape (e.g., “lower than” in “HB is lower than normal”) is lost. On the contrary, CVAE explores the discourse-level diversity but ambiguous phrases like “wbc $3.45 \times 10^{12}/l$ ” in $Q1$, which indicates potential inflammation rather than anemia, are often generated in a key place. Similarly, in $Q3$ from CVAE “sudden fever after menstruation, discomfort” in most cases indicates endocrine disorders rather than anemia.

For eg-CVAE, we can see it explores discourse-level diversity by generating symptoms like “whitish complexion” in $Q1$ that are not existing in the Q . In terms of the validity, the generated imperative semantics of the non-key phrases are consistent with the implicit semantics of the original questions of anemia. For example, although the semantics

⁶We include detailed case comparison between eg-CVAE, type-CVAE and entity-CVAE in supplementary material.

Table 3: Usefulness of the generated QA pairs. * indicates difference between the original setting and the new setting is statistically significant ($p < 0.01$).⁷

Dataset	Accuracy
Original	61.97
+ HRED QA	58.78
+ VHRED QA	62.28
+ type-CVAE QA	65.27*
+ entity-CVAE QA	64.67*
+ eg-CVAE QA	67.96*

of “the poor face”, “anxious” and “whitish complexion” in Q and $Q3$, $Q1$ are different, they does not influence on the overall diagnosis of “anemia”. The generated “the normal systolic blood pressure” and “normal liver” do not affect the judgment of “anemia” as they are normal body signal, too.

Evaluation on a QA System

To further study the usefulness of the generated medical QA pairs, we integrate such generated pairs into a QA system, which is an attention-based model (Cui et al. 2017) for NM-LEC_QA dataset. The results are summarized in Table 3.

For baseline methods, integrating the generated QA pairs from HRED hurts the accuracy without augmented data. As pointed out in (Serban et al. 2017), HRED is very likely to favor short-term predictions instead of long-term predictions. As shown in Figure 3, rather than globally considering context phrases to generate a meaningful phrase for the current slot, HRED tends to repeat the predicted correct word. The lack of diversity and repeat of common words lead to the discrepancy in the generated questions’ distribution and the original one, which may cause the degradation and introduce noise to the original dataset. From the results of vanilla CVAE-based VHRED, we can see that the improvement exists but is marginal. we presumes that is because the lack of constraint on the latent distribution leads to weak guidance from the corresponding answer and the unlabeled textbook for generated questions from VHRED.

Two variants of the proposed method, entity-CVAE and type-CVAE, generated QA pairs that boost the original QA system with 3-4% accuracy. Each of them introduces external constraints on the latent variable in the decoding phase, which may help to diversify the generated questions while keeping linguistic and structural relationships within original questions. Furthermore, type-CVAE generates QA pairs that seem to be more helpful to the QA system. This benefit may come from the modeling of type information, which allows the generated questions to be relatively more diverse and thus introduces more useful knowledge. The proposed method eg-CVAE combines the advantages of entity-CVAE and type-CVAE, building a three-pass decoding process, and thus improves the QA system to achieve the highest accuracy. These observations further demonstrate the usefulness of the generated medical QA pairs by eg-CVAE.

⁷We calculate statistical significance based on the bootstrap test in (Noreen 1989) with 10k samples.

Original QA	Q	男,1岁,脸色苍白1个月,易疲倦,时而焦虑,胃纳差,体检:肝脏位于肋下3cm,质中,血红蛋白(Hb)低于正常, 红细胞数(RBC)低于正常,平均红细胞体积(MCV)低于正常 , mch 20pg, mchc 0.26,最可能的诊断是 Male, 1 year-old, face been pale for 1 month, often tired, sometimes anxious, poor appetite, examination result: liver is 3 cm below the ribs, between hard and soft, hb is lower than normal, rbc is lower than normal, mcv is lower than normal, mch 20pg, mchc 0.26, the most likely diagnosis is
	A	贫血 (anemia)
HRED	Q1	女,1岁,脸色苍白1个月,焦虑,体检:肝脏肋下未触及,Hb正常, RBC RBC 1.02 正常, MCV正常,mch 10pg,mchc 10.3cm,最可能的诊断是 Female, 1 year-old, face been pale for 1 month, anxiety, anxiety, examination result: untouched liver under the ribs, hb normal, rbc count rbc count 1.02 normal, mcv normal, mch 10pg, mchc 10.3cm, the most likely diagnosis is
	Q2	男,1岁,脸色 3个月,焦虑,体检:肝脏肋下未触及,Hb 正常, RBC 1.02 正常, MCV 正常,mch 0.1pg,mchc 0.26,最可能的诊断是 Male, 1 year-old, face 3 months, anxiety, anxiety, anxiety, examination result: liver under the ribs, hb normal, rbc count 1.02 normal, mcv normal, mch 0.1pg, mchc 0.26, The most likely diagnosis is
	Q3	女,1岁,脸色苍白1个月,焦虑,体检:肝脏肋下未触及,体温39.3oC,Hb高于正常, RBC 1.02正常, MCV正常, mch 10pg, mchc 0.26,最可能的诊断是 Female, 1 year-old, face has been pale for 1 month, anxiety, anxiety, examination result: liver under the ribs, body temperature 39.3oC, hb higher than normal, rbc count 1.02 normal, mcv normal, mch 10pg, mchc 0.26, the most likely diagnosis is
CVAE	Q1	男,58岁,头发少1个月,易疲倦,全身无力,体检脸色发黄,肝肋下2.5cm,胃纳差,尿蛋白(+),wbc $3.45 \times 10^{12}/L$, MCV低于正常, mch 24 pg, mchc 0.26, 最可能的诊断是 Male, 58 year-old, hair deficiency for 1 month, easy to get tired, general weakness, examination result: yellow complexion, liver is 2.5cm under the ribs, poor storage of stomach, urine protein (+), wbc $3.45 \times 10^{12}/L$, mcv lower than normal, mch 24pg, mchc 0.26, the most likely diagnosis is
	Q2	女,4岁,发热,伴面色污秽,腹痛1天,查体,双腿呈非凹陷性水肿,大便WBC高于正常,尿中WBC高于正常,MCV低于正常, mch 20pg, mchc 0.26,最可能的诊断是 Female, 4 year-old, fever, with facial stains, diarrhea, abdominal pain for 1 day, examination result: legs with non-sink edema, wbc in stool is higher than normal, the number of wbc in urine is higher than normal, mcv below normal, mch 20pg, mchc 0.26, the most likely diagnosis is
	Q3	女,21岁,脸色苍白1个月,经后4天,偶有轻微不适,月经后突发发热,不适,偶有下腹痛,质硬,体检脸色发黄,Hb低于正常,RBC低于正常,MCV低于正常, mch 20pg, mchc 0.26,最可能的诊断是 Female, 21 year-old, face been pale for 1 month, 4 days after menstruation, mild discomfort occasionally, sudden fever after menstruation, discomfort, occasional lower abdominal pain, hard, examination result: yellow complexion, hb below normal, rbc count lower than normal, mcv lower than normal, mch 20pg, mchc 0.26, the most likely diagnosis is
Eg-CVAE	Q1	患儿,7岁,脸色苍白时有吞咽不下,精神萎靡伴焦虑,一周加重,查体,身高83cm,肝肋外观质硬,Hb低,RBC低于正常,MCV低于正常, mch 20pg, mchc 0.26,最可能的诊断是 Child patient, 7 year-old, pale complexion, unable to swallow occasionally, mental wilting with anxiety, a week of aggravation, examination result: height is 83cm, liver hard, low hb, rbc count is below normal, mcv is lower than normal, mch 20pg, mchc 0.26, the most likely diagnosis is
	Q2	患者,男,10岁,脸色发黄1个月,时有腹痛,伴焦虑,有时加重,体检:肝脏位于肋下2cm,周身淋巴结增生,Hb低于正常,RBC低于正常,MCV低于正常, mch 20pg, mchc 0.26,最可能的诊断是 Patient, male, 10 year-old, Face has been yellow for 1 month, diarrhea occasionally, with anxiety, sometimes aggravates, examination result: the liver is 2cm below the ribs, surrounding tissue proliferate, the hb is lower than normal, rbc count lower than normal, mcv lower than normal, mch 20pg, mchc 0.26, the most likely diagnosis is
	Q3	男,8岁,脸色差时有疲倦, 焦虑,查体:收缩压正常,舒张压正常,脾脏正常,外观毛玻璃样,Hb60K/L,rbc $3.3 \times 10^{12}/L$, MCV低于正常, mch 20pg, mchc 0.26,最可能的诊断是 Male, 8 year-old, with poor face, feel tired occasionally, anxious, examination result: normal systolic blood pressure, normal diastolic blood pressure, normal spleen, appearance of ground glass, hb 60K/L, rbc $3.3 \times 10^{12}/L$, mcv is lower than normal, mch 20pg, mchc 0.26, the most likely diagnosis is

Figure 3: Case study for generated QA pairs of different methods (the key phrases in original QA pair are in bold)

Related Work

Question Generation (Heilman and Smith 2010) has attracted increasing attention in recent years. However, most existing work only focuses on the similarity of generated questions with the original ones, but ignores the usefulness in training a QA system of generated questions given answers. Earlier work in question generation employed rule-based approaches to transform input texts into corresponding questions, usually requiring some well-designed general rules (Mitkov and others 2003), templates (Labutov, Basu, and Vanderwende 2015) or syntactic transformation heuristics (Ali, Chali, and Hasan 2010). Recent studies leveraged neural networks to generate questions in an end-to-end fashion. (Du, Shao, and Cardie 2017) applied the attention-based sequence-to-sequence model to generate questions in the context of reading comprehension. In medical QA, (Roberts et al. 2017; Pampari et al. 2018) targets the same problem as us from the dataset angle. (Walonoski et al. 2017) is similar to us, but they focus on the state transition of patient records.

Other existing work, which tackles the usefulness and models the question-answer pair generation directly, still sets the diversity of questions for the corresponding answer aside and requires related context in prior. (Serban et al. 2016b) applied the encoder-decoder framework to generate question-answering pairs from built knowledge base triples. (Subramanian et al. 2018) formulated the question-answer pair generation in reading comprehension, where each pair will be given one high-quality context and the answer is a text span of the context, separately with the answer detection and question generation problem. (Wang et al. 2017) leveraged policy gradient techniques to further improve the generation quality. Coreference knowledge is also introduced for question-answer pair generation from Wikipedia articles with the context in (Du and Cardie 2018). (Duan et al. 2017) investigated integrating generated questions from given con-

text to the question-answering system on sentence selection tasks, which leveraged both rule-based features and neural networks to approximate the semantics of generated questions with original ones. (Yang et al. 2017; Song et al. 2018) also leveraged the generated QA for QA system. But they all have the external context in SQuAD (Rajpurkar et al. 2016) to build upon, which does not exist in our medical setting.

Compared to existing work, our work introduces structure information of QA pairs generation in medical domain, which does not involve any prior context. To ensure the validity of generated QA pairs, we proposed an unsupervised detector to automatically explore external materials. We also proposed to model the question-answer pair generation problem directly as approximating the latent distribution of medical questions with the corresponding answer.

Conclusions

In this paper, we introduced a novel framework, consisting of an unsupervised key phrase detector and an Entity-guided CVAE-based generator, for automated question-answer pair generation in the medical domain. Different from existing seq2seq models that involve a definite encoding-decoding procedure to restrict the generation scope, or traditional CVAE models that directly approximate the posterior distribution over the latent variables to a simple prior, the proposed method models the generation process as a multi-pass procedure (type, entity and phrase as constraints over the latent distribution) to ensure both validity and diversity. Experiments on a real-world dataset from the National Medical Licensing Examination of China demonstrate that the proposed method outperforms existing methods and can generate more diverse, informative and valid medical QA pairs that further benefit the examination QA system. We will investigate more on the generalizability of proposed method on standard dataset like SQuAD (Rajpurkar et al. 2016) and

its integration with popular pretrained model (Devlin et al. 2019) in the future work.

References

- Ali, H.; Chali, Y.; and Hasan, S. A. 2010. Automation of question generation from sentences. In *Proceedings of QG Workshop*.
- Chen, B., and Cherry, C. 2014. A systematic comparison of smoothing techniques for sentence-level BLEU. In *Proceedings of ACL Workshop*.
- Chung, J.; Gulcehre, C.; Cho, K.; and Bengio, Y. 2014. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*.
- Cui, Y.; Chen, Z.; Wei, S.; Wang, S.; Liu, T.; and Hu, G. 2017. Attention-over-attention neural networks for reading comprehension. In *Proceedings of the ACL*.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of NAACL*.
- Du, X., and Cardie, C. 2018. Harvesting paragraph-level question-answer pairs from wikipedia. In *Proceedings of ACL*.
- Du, X.; Shao, J.; and Cardie, C. 2017. Learning to ask: neural question generation for reading comprehension. In *Proceedings of ACL*.
- Duan, N.; Tang, D.; Chen, P.; and Zhou, M. 2017. Question generation for question answering. In *Proceedings of EMNLP*.
- Forgues, G.; Pineau, J.; Larchevêque, J.-M.; and Tremblay, R. 2014. Bootstrapping dialog systems with word embeddings. In *Proceedings of NIPS Workshop*.
- Gormley, C., and Tong, Z. 2015. *Elasticsearch: The definitive guide: A distributed real-time search/analytics engine*.
- Gu, X.; Cho, K.; Ha, J.-W.; and Kim, S. 2018. Dialogwae: multimodal response generation with conditional wasserstein auto-encoder. *CoRR abs/1805.12352*.
- He, J.; Spokoyny, D.; Neubig, G.; and Berg-Kirkpatrick, T. 2019. Lagging inference networks: posterior collapse in variational autoencoders. *arXiv preprint arXiv:1901.05534*.
- Heilman, M., and Smith, N. A. 2010. Good question! statistical ranking for question generation. In *Proceedings of NAACL-HLT*.
- Kingma, D. P., and Welling, M. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- Labutov, I.; Basu, S.; and Vanderwende, L. 2015. Deep questions without deep understanding. In *Proceedings of ACL*.
- Li, W.; Xiao, X.; Lyu, Y.; and Wang, Y. 2018. Improving neural abstractive document summarization with explicit information selection modeling. In *Proceedings of EMNLP*.
- Liu, C.-W.; Lowe, R.; Serban, I. V.; Noseworthy, M.; Charlin, L.; and Pineau, J. 2016. How not to evaluate your dialogue system: an empirical study of unsupervised evaluation metrics for dialogue response generation. In *Proceedings of EMNLP*.
- Mitkov, R., et al. 2003. Computer-aided generation of multiple-choice tests. In *Proceedings of NAACL-HLT Workshop*, 17–22.
- Noreen, E. W. 1989. *Computer-Intensive Methods for Testing Hypotheses: An Introduction*.
- Pampari, A.; Raghavan, P.; Liang, J.; and Peng, J. 2018. emrqa: A large corpus for question answering on electronic medical records. In *Proceedings of EMNLP*.
- Papineni, K.; Roukos, S.; Ward, T.; and Zhu, W.-J. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of ACL*.
- Parvez, M. R.; Chakraborty, S.; Ray, B.; and Chang, K.-W. 2018. Building language models for text with named entities. In *Proceedings of ACL*.
- Peters, M. E.; Neumann, M.; Iyyer, M.; Gardner, M.; Clark, C.; Lee, K.; and Zettlemoyer, L. 2018. Deep contextualized word representations. In *Proceedings of NAACL-HLT*.
- Rajpurkar, P.; Zhang, J.; Lopyrev, K.; and Liang, P. 2016. Squad: 100,000+ questions for machine comprehension of text. In *Proceedings of EMNLP*.
- Roberts, K.; Demner-Fushman, D.; Voorhees, E. M.; Hersh, W. R.; Bedrick, S.; Lazar, A. J.; and Pant, S. 2017. Overview of the trec 2017 precision medicine track. In *TREC*.
- Rus, V., and Lintean, M. 2012. A comparison of greedy; optimal assessment of natural language student input using word-to-word similarity metrics. In *Proceedings of NAACL-HLT Workshop*.
- Schuster, M., and Paliwal, K. K. 1997. Bidirectional recurrent neural networks. *IEEE TSP*.
- Serban, I. V.; Sordoni, A.; Bengio, Y.; Courville, A.; and Pineau, J. 2016a. Building end-to-end dialogue systems using generative hierarchical neural network models. In *Proceedings of AAAI*.
- Serban, I. V.; García-Durán, A.; Gulcehre, C.; Ahn, S.; Chandar, S.; Courville, A.; and Bengio, Y. 2016b. Generating factoid questions with recurrent neural networks: the 30m factoid question-answer corpus. In *Proceedings of ACL*.
- Serban, I. V.; Sordoni, A.; Lowe, R.; Charlin, L.; Pineau, J.; Courville, A.; and Bengio, Y. 2017. A hierarchical latent variable encoder-decoder model for generating dialogues. In *Proceedings of AAAI*.
- Shen, D.; Wang, G.; Wang, W.; Min, M. R.; Su, Q.; Zhang, Y.; Li, C.; Henao, R.; and Carin, L. 2018. Baseline needs more love: on simple word-embedding-based models; associated pooling mechanisms. In *Proceedings of ACL*.
- Sohn, K.; Lee, H.; and Yan, X. 2015. Learning structured output representation using deep conditional generative models. In *Proceedings of NIPS*.
- Song, L.; Wang, Z.; Hamza, W.; Zhang, Y.; and Gildea, D. 2018. Leveraging context information for natural question generation. In *Proceedings of the NAACL*.
- Subramanian, S.; Wang, T.; Yuan, X.; Zhang, S.; Bengio, Y.; and Trischler, A. 2018. Neural models for key phrase extraction; question generation. In *Proceedings of ACL Workshop*.
- Sutskever, I.; Vinyals, O.; and Le, Q. 2014. Sequence to sequence learning with neural networks. In *Proceedings of NIPS*.
- Walonoski, J.; Kramer, M.; Nichols, J.; Quina, A.; Moesel, C.; Hall, D.; Duffett, C.; Dube, K.; Gallagher, T.; and McLachlan, S. 2017. Synthea: An approach, method, software mechanism for generating synthetic patients records. *Journal of the AMIA*.
- Wang, X. Y. T.; Gülcehre, Ç.; Sordoni, A.; Bachman, P.; Zhang, S.; Subramanian, S.; and Trischler, A. 2017. Machine comprehension by text-to-text neural question generation. In *Proceedings of ACL Workshop*.
- Xia, Y.; Tian, F.; Wu, L.; Lin, J.; Qin, T.; Yu, N.; and Liu, T.-Y. 2017. Deliberation networks: Sequence generation beyond one-pass decoding. In *Proceedings of NIPS*.
- Yang, Z.; Hu, J.; Salakhutdinov, R.; and Cohen, W. W. 2017. Semi-supervised qa with generative domain-adaptive nets. In *Proceedings of ACL*.
- Zhao, T.; Zhao, R.; and Eskenazi, M. 2017. Learning discourse-level diversity for neural dialog models using conditional variational autoencoders. In *Proceedings of ACL*.