

Dialogue-Contextualized Re-ranking for Medical History-Taking

Jian Zhu*

University of British Columbia

JIAN.ZHU@UBC.CA

Ilya Valmianski †

AuxHealth

ILYA@AUXHEALTH.IO

Anitha Kannan

Curai Health

ANITHA@CURAI.COM

Abstract

AI-driven medical history-taking is an important component in symptom checking, automated patient intake, triage, and other AI virtual care applications. As history-taking is extremely varied, machine learning models require a significant amount of data to train. To overcome this challenge, existing systems are developed using indirect data or expert knowledge. This leads to a training-inference gap as models are trained on different kinds of data than what they observe at inference time. In this work, we present a two-stage re-ranking approach that helps close the training-inference gap by re-ranking the first-stage question candidates using a dialogue-contextualized model. For this, we propose a new model, global re-ranker, which cross-encodes the dialogue with all questions simultaneously, and compare it with several existing neural baselines. We test both transformer and S4-based language model backbones. We find that relative to the expert system, the best performance is achieved by our proposed global re-ranker with a transformer backbone, resulting in a 30% higher normalized discount cumulative gain (nDCG) and a 77% higher mean average precision (mAP). As part of this work, we also release pre-trained checkpoints for bi-directional and autoregressive S4 models trained on Wikipedia and PubMed data.

1. Introduction

History taking is a critical component of a medical encounter ([Hampton et al., 1975](#)). It involves collecting relevant patient-reported information such as presenting symptoms, patient concerns as well as the past medical, psychological and social history. This information forms the basis of subsequent patient triage, diagnosis, and care planning. While history taking is an important component of the medical encounter, it is also one of the most time-consuming components ([Soltau et al., 2021](#); [Chen et al., 2020](#)) and when done incompletely can lead to triage, diagnostic, and treatment errors ([Hampton et al., 1975](#)).

Creating tools for automating portions of history taking has been an on-going effort for more than five decades ([Watson, 1965](#)). The simplest of such tools are static pre-visit questionnaires that are now used widely in US healthcare. However, static questionnaires tend to be long, ask not very relevant questions, and are not customized to patients' needs.

* Work done during an internship at Curai Health

† Work done while at Curai Health

More recently, there has been work on building intelligent systems that can adjust questions based on patient responses (see (Valmianski et al., 2021; Compton et al., 2021) and citations therein). However, developing the medical reasoning necessary for these systems is difficult. Existing approaches include using reinforcement learning with simulated patients (Tang et al., 2016; Kao et al., 2018), supervised learning on clinical notes (Valmianski et al., 2021), and expert systems (Compton et al., 2021).

In all of the previous works, the medical reasoning system was built on top of data that was a proxy for real doctor-patient interactions. This is because, on the one hand, there is little available data consisting of doctor-patient history-taking dialogue, on the other hand, the space of possible questions asked during history-taking is very large. Thus, training a history-taking model requires significant amounts of labeled interaction data. However, this data is more readily available from indirect sources such as medical notes, expert knowledge, or simulations. This creates a training-inference gap: the data that is used to train the model is not fully representative of the data that the model sees at inference time.

This training-inference gap has a significant impact on the quality of history taking, especially since only a few questions can be asked in a given encounter. This calls for an approach that reconciles the difficulty of supporting a large set of potential questions to a small set of pertinent questions attuned to the patient’s health issue, with only a small amount of direct training data.

In this paper, we start with an expert system and show how to use a relatively small amount of real doctor-patient dialogue data to close this training-inference gap. We take inspiration from the information retrieval literature where “retrieve and re-rank” is a popular paradigm for computationally efficient retrieval of documents from a large corpus (Nogueira and Cho, 2019; Lin et al., 2021). In our case, the “retrieve” part is performed by the expert systems which retrieves a list of possible questions to ask the patient, and a dialogue-trained re-ranker then “re-ranks” the possible questions. Because the re-ranking model takes the original expert system’s candidate questions, it does not need to predict over the space of all possible questions. Instead, it only needs to re-rank from a much smaller subset, which greatly simplifies the machine-learning task. Our model takes both the previous dialogue and the possible questions as free text entries, which means that the system can operate even if the underlying expert system is replaced with something else.

Our contributions are as follows:

1. We propose a two-step approach to history-taking question selection where we use an expert system to retrieve a list of candidate questions and then use a machine-learned re-ranker to get the top question to ask.
2. We propose a novel “global re-ranker” which embeds both the preceding dialogue and candidate questions into a single long string. We then train long context language models to predict the relevance of each question simultaneously.
3. We perform a careful study of other re-rankers for this task. This includes different architectures such as bi-encoder, cross-encoder, and autoregressive re-rankers. We examine different long context models including S4 (bi-directional and autoregressive), Nystromformer (bi-directional, variants with Nystrom attention and with full attention), and LongT5 (autoregressive). We examine the effect of different loss functions

from the pointwise, pairwise, and listwise families. Finally, we perform some ablation studies on the context length and the initial retrieval ordering.

4. We release checkpoints for S4 pre-trained both bidirectionally and autoregressively on the English subset of Wikipedia¹ and Pubmed PMC Open Access Subset² datasets.
5. We find that our global re-ranker approach performs better than other more traditional approaches. Furthermore, all re-rankers significantly improve the original expert system performance. We also find that S4-based, while worse than full-attention transformers, is competitive with the Nystrom-attention transformer.

2. Generalizable Insights about Machine Learning in the Context of Healthcare

One of the main challenges in using deep learning for healthcare is the lack of large annotated datasets. Obtaining large amounts of annotated data is costly and time-consuming because annotations need to be provided by trained healthcare professionals. Recent works have successfully leveraged the progress in the development of large language models that are trained on web-scale data. In many tasks, including medical history taking discussed in this paper, this approach introduces a training-inference gap: the data used to train the model does not fully represent the data that the model sees at inference time. In this context, our approach of retrieving a candidate set of answers (we use an expert system as the base model to provide candidates, but this can come from a large language model, too) and then using a learned reranker based on small amounts of labeled data is a promising alternative. As we show in this paper, such a reranker can be trained from public data sources and then finetuned to the task.

3. Related work

We study re-ranking history-taking recommendations based on doctor-patient dialogue. These dialogues tend to be long and exceed the typical token-length limits of transformer models. As such, there are two bodies of literature relevant to this work. § 3.1 discusses work on modern neural long-range language models that are able to encode the entire doctor-patient conversation. § 3.2 discusses work on re-ranking algorithms that can take the encoded dialogue and use it to re-rank history-taking questions.

3.1. Long-range transformers

Transformers (Vaswani et al., 2017) have become the mainstream architecture for natural language processing. With the self-attention mechanism, transformers can attend to all tokens in a sequence simultaneously, thereby being more powerful than classical architectures like convolutional (LeCun et al., 1989) or long short-term memory (LSTM) (Hochreiter and Schmidhuber, 1997) networks. However, due to its $O(n^2)$ complexity, the original transformer cannot process long sequences efficiently. Popular pre-trained language models such

1. <https://huggingface.co/datasets/wikipedia>

2. <https://www.ncbi.nlm.nih.gov/pmc/tools/textmining/>

as BERT (Devlin et al., 2019) and RoBERTa (Liu et al., 2019) can only process up to 512 tokens. Efforts have been made to reduce the computational complexity of transformers, as a variety of efficient transformers that can process long text sequences have been proposed, such as Reformer (Kitaev et al., 2019), Linformer (Wang et al., 2020), Longformer (Beltagy et al., 2020), BigBird (Zaheer et al., 2020), Performer (Choromanski et al., 2020), Nystromformer (Xiong et al., 2021), LongT5 (Guo et al., 2021), etc. In addition to transformers, alternative approaches have also been shown to be promising for processing very long sequences, notably state-space models (Gu et al., 2021, 2022; Mehta et al., 2022). The Structured State Space Sequence model (S4) (Gu et al., 2021) has significantly outperformed many long-range transformers in the Long Range Arena benchmark (Tay et al., 2020). In this paper we utilize the Nystromformer (Xiong et al., 2021) and the S4 model (Gu et al., 2021) as the possible non-autoregressive backbones and LongT5 as an autoregressive backbone.

3.2. Re-ranking

Modern information retrieval (IR) or question answering (QA) systems are usually divided into two stages (Nogueira and Cho, 2019; Lin et al., 2021). In the first stage, given a query, a large number of documents are retrieved using an efficient method. In the second stage, a computationally intensive but more accurate method is used to re-rank documents retrieved in the first stage. This is analogous to our problem statement where we use an expert system to retrieve a set of relevant history-taking questions, and then use a re-ranking algorithm as the second stage. Relevant related work addresses both the architectural choices that can be made for the re-rankers, as well as the loss functions used to train them.

Architectures for second-stage re-ranking. There are three main type of architectures (1) bi-encoders (Lin et al., 2021; Nogueira et al., 2019; Reimers and Gurevych, 2019; Thakur et al., 2021) (2) cross-encoders (Lin et al., 2021; Nogueira et al., 2019; Humeau et al., 2019) and (3) autoregressive re-rankers (Nogueira et al., 2020; Pradeep et al., 2021; Min et al., 2021). In the bi-encoder architecture, the query and the candidate document are encoded into vector representations by two separate encoders (though these two encoders can share the same weights (Reimers and Gurevych, 2019)). The relevance score is calculated as the distance between the two vector representations. In a cross-encoder, a query and a candidate document are concatenated together and fed into the cross-encoder in a single pass. In most use cases, cross-encoders outperform bi-encoders in document retrieval and ranking (Thakur et al., 2021; Humeau et al., 2019); however, bi-encoders are usually more efficient than cross-encoders, as all documents in a given corpus can be pre-computed and stored as dense embeddings for retrieval, thereby avoiding repeated computations (Reimers and Gurevych, 2019). Recent sequence-to-sequence models such as T5 (Raffel et al., 2020) have also been applied to autoregressive re-ranking. In this approach, the query and the documents are usually encoded by the encoder and the decoder either predicts whether the document is relevant (Nogueira et al., 2020) or directly generates the retrieved text in response to the query (Min et al., 2021).

Scoring functions for ranking. There are also several scoring paradigms possible for ranking: (1) “pointwise” scores where the relevance of a query is computed on a per-document basis (similar to cross-encoders) (Nogueira et al., 2019), (2) “pairwise” scores

where documents are ranked relative to each other in pairs (Nogueira et al., 2019), and (3) “listwise” scores where a list of candidates are ranked simultaneously (Liu et al., 2009). Prior studies show that pairwise and listwise approaches tend to outperform pointwise approaches (Nogueira et al., 2019; Pradeep et al., 2021; Min et al., 2021; Zerveas et al., 2021; Chen et al., 2021). In this study, we compare all three (‘pointwise’, ‘pairwise’, and ‘listwise’) approaches for re-ranking history-taking questions but mainly focus on listwise approaches.




Figure 1: (a) Bi-encoder, (b) the cross-encoder and (c) the proposed global reranker.

Connection to our proposed global re-ranker. In this paper we propose a novel re-ranker approach we call “global re-ranker.” In this approach, all candidate documents are concatenated into a single input that is then processed by a long context language model. For schematic comparison between bi-encoder, cross-encoder, and global re-ranker please see Figure 1. For a more detailed description of the method see § 5. Concatenating pairs of documents into a single string has been previously done both in bi-directional (Nogueira et al., 2019) and autoregressive (Pradeep et al., 2021) paradigms. These models only investigate pairwise scoring due to length constraints imposed by the pre-trained transformer. For listwise ranking with more than two documents, previous approaches focused on ranking only the extracted embeddings (Chen et al., 2021; Ai et al., 2019), which doesn’t model the deep semantic relationships between candidate documents.

4. Closing the train-inference gap with re-ranking

An overview of our approach to closing the train-inference gap in an existing history-taking system can be seen in Figure 2. We first use an expert system to suggest relevant history-taking questions and then use a deep neural network contextualized by the entire doctor-patient dialogue to re-rank expert system suggestions.

The goal of re-ranking is, given the prior dialogue context \mathbf{d} and a list of n candidate history-taking questions $Q = [q_1, q_2, \dots, q_n]$, to generate a new list Q' which consists of (possibly reordered) questions from Q such that the higher relevance questions appear earlier in the sequence. In our case, the candidate questions are generated using an in-house Expert System, and the ground truth labels $\mathbf{y} = [y_1, y_2, \dots, y_n], y_i \in \{0, 1\}$ represent whether a doctor asked a given recommended question (1 if the question was asked, 0 if the question was not asked). A doctor may ask multiple questions at the same time, thus multiple elements of \mathbf{y} can have a value of 1, see § 6.1 for more details on how the ground truth is produced. Finally, in all of the models studied in this work, the re-ranking is achieved by assigning scores $\mathbf{s} = [s_1, s_2, \dots, s_n]$ to each question in Q , and then constructing Q' by reordering using scores in \mathbf{s} .

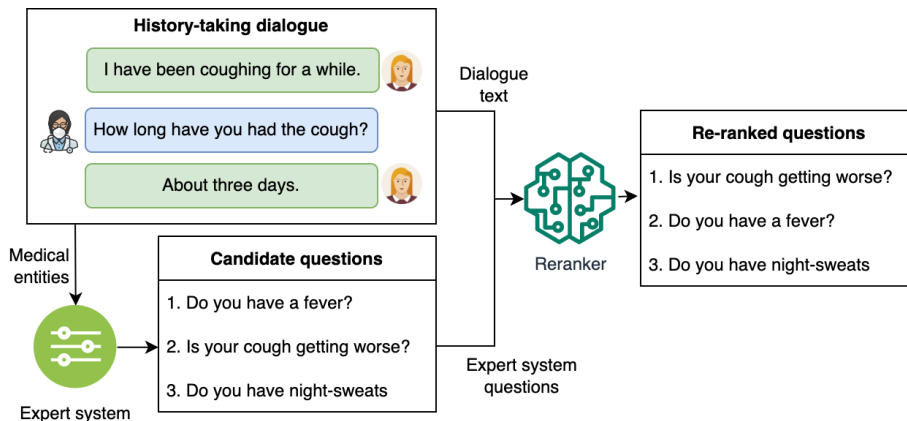


Figure 2: Overview of the proposed two-stage history-taking workflow. An expert system suggests candidate questions based on relevant entities extracted from the dialogue. A machine-learned deep neural network re-ranker then re-ranks the candidate questions based on the dialogue text.

5. Global re-ranker

We propose the global re-ranker, an accurate and efficient listwise re-ranking method. In this approach (see Figure 1(c) for a schematic), the history-taking dialogue and all candidate history-taking questions are concatenated into a single text input, using which the model then assigns the ranking scores to all questions simultaneously. The global re-ranker directly encodes all texts through the language model, thereby ensuring deep semantic interactions not only between the dialogue and the candidate questions but also between all candidate questions.

The input text to the global re-ranker is the concatenation of both the dialogue context and all the candidate questions: $[\text{CLS}] \mathbf{d} [\text{SEP}] \mathbf{q}_1 [\text{MASK}_1] [\text{SEP}] \mathbf{q}_2 [\text{MASK}_2] [\text{SEP}] \dots \mathbf{q}_n [\text{MASK}_n] [\text{SEP}]$, where the $[\text{SEP}]$ token is used to mark the boundaries of candidate questions. The $[\text{MASK}_i]$ token is the pooling token for the preceding question \mathbf{q}_i . For each pooling token $[\text{MASK}_i]$, the global reranker predicts a score s_i , which represents the relevance for \mathbf{q}_i . We also added type embeddings to every input tokens to indicate whether it belongs to the dialogue or the candidate questions. The actual number of candidate questions provided by the expert system ranged from 3 to 40.

While self-attention itself does not assume any inherent order of the input sequence, pretrained transformer models usually encode the text sequentially due to the presence of positional embeddings. In the current task, it is expected that a language model learns the sequential relations between words within \mathbf{d} and \mathbf{q}_i . From our ablation experiments (see § 7.2), we found that the best performance is achieved when the model is agnostic to the order of input questions $[\mathbf{q}_1, \mathbf{q}_2, \dots, \mathbf{q}_n]$. In order to remove the positional bias, we reset the positional embedding when each new question starts.

We selected three different neural architectures to implement the global reranker, all of which can process long textual sequences. The first two approaches are based on the Nystromformer (Xiong et al., 2021), which was originally proposed to be an efficient trans-

	Train	Dev	Test
Num. Encounters	12105	311	655
Num. Samples	26106	626	1361
Avg. Length of Dialog.	287.2	374.7	288.8
Num. Selected Questions	4.0	3.8	4.1
Num. Candidate Questions	27.9	26.6	28.1
Avg. Length of Questions	8.0	8.1	8.0

Table 1: Statistics of different data splits. Text lengths were calculated based on words.

former. We experiment with Nystromformer with both Nystrom attention turned on and turned off (in which case it uses full attention). We use Nystromformer as the base of our “full attention” transformer because this enables us to leverage the pretrained Nystromformer checkpoints that had been trained on long texts and retain the good performance of full attention. We learned from pilot experiments that other efficient transformers such as Longformer (Beltagy et al., 2020) failed to converge. The third neural architecture is a state-space model, S4, which has been shown to process long sequences more effectively than many transformers (Gu et al., 2021).

To train the global re-ranker, we compared a variety of loss functions across point-wise, pair-wise and listwise approaches in the learning-to-rank framework (Liu et al., 2009). The point-wise baseline was trained with binary cross-entropy. For pairwise loss functions, we tested the RankNet (Burges et al., 2005) and LambdaRank (Burges et al., 2006). The listwise loss functions we used were ListNet (Cao et al., 2007), ListMLE (Xia et al., 2008), ApproxNDCG (Qin et al., 2010) and NeuralNDCG (Pobrotyn and Bialobrzeski, 2021), the latter two of which directly optimized the Normalized Discounted Cumulative Gain (NDCG) metrics.

6. Experiments

6.1. Data

The medical dialogue data was collected from a portion of real doctor-patient interactions collected on our text-based medical service platform, Curai. In a typical interaction, the physician asks a series of history-taking questions that can be entered either as free text or selected from a list of recommendations. These recommendations are made using the Expert System that forms the first stage in our proposed workflow. At each dialogue turn where recommended questions are asked, the doctor selected questions are marked as relevant and the not-selected questions are marked as irrelevant. This forms a natural dataset of doctor annotated selections on which we train our re-rankers.

The dataset consists of 13071 encounters. We filtered non-history-taking dialogue turns using in-house dialogue segmentation model, similar to Wang et al. (2022). The detailed statistics of our data are displayed in Table 1.

6.2. Metrics

For evaluation, we adopted two common ranking metrics, normalized discounted cumulative gain (nDCG)(Järvelin and Kekäläinen, 2000) and mean average precision (mAP) (Lin et al., 2021). The mAP assumes binary relevance whereas nDCG can work with both binary and continuous relevance. Specifically for global re-rankers, the average metrics over 5 repeated runs of evaluations were reported. In each run, the order of candidate questions fed to the global re-ranker was randomly reshuffled to mitigate positional biases.

6.3. Baseline approaches

In addition to the global ranker, we also implement three widely adopted baseline ranking approaches: bi-encoder, cross-encoder, and autoregressive re-ranker.

Bi-encoder. In the bi-encoder architecture (see Figure 1(a)), the dialogue query and the candidate questions are encoded by two separate encoders f_D and f_Q , and the relevance score between the two resulting vector representations are computed with cosine similarity. The bi-encoder learns an embedding space where the dialogue representation is close to the most relevant questions while being distant from less relevant questions. The training objective is to minimize the InfoNCE loss function (Oord et al., 2018) through contrastive learning with 7 negatives randomly sampled from the list of recommended candidate questions by the Expert System. The temperature parameter of the InfoNCE loss was set to 0.05 throughout the training (Gao et al., 2021).

Cross-encoder. In the cross-encoder architecture (see Figure 1(b)) the prior dialogue is concatenated with a candidate question. The cross-encoder f_C assigns a relevance score to this candidate question using a classification head on top of the contextual representation of the dialogue and the query. We consider transformers and S4-based models. For transformers, the [CLS] token is treated as the contextual representation. For the bi-directional S4 re-rankers, we use average pooling of the last layer to obtain the contextual representations. All cross-encoder variants are trained with the binary cross-entropy loss.

Autoregressive re-ranker. We also consider autoregressive re-rankers (Nogueira et al., 2020; Pradeep et al., 2021). For a transformer baseline, we use a pre-trained LongT5 (Guo et al., 2021). The query and the document are concatenated together to form the input sequence: `Query: d Document: q_i Relevant:`, which is fed into the encoder. The decoder then predicts `true` for relevant documents or `false` for irrelevant documents. During inference, a softmax function is applied to the logits of the `true` and the `false` tokens to normalize the results across multiple queries.

For autoregressive S4, when we followed the Long-T5 method, we found it to highly unstable and failed to converge, similar to what was found in the literature (Nogueira et al., 2020) on its dependency certain keywords e.g., true/false. Therefore, we followed the same setting as in the cross-encoder, except that the underlying model is autoregressive rather than bi-directional. Here, the concatenated dialogue and a candidate question are fed into the S4 re-ranker and the average pooling of the last layer is classified as either relevant or irrelevant through a classification head.

6.4. Implementation

S4 model pretraining. The S4 model was based on the original implementation of S4 language model (Gu et al., 2021), in which the S4 layers were used as a drop-in replacement for the self-attention layers in a typical transformer. We implemented a 12-layer bidirectional and autoregressive S4 models. We set the hidden dimensions to 768 in order to match the parameter count of mainstream pretrained transformers (such as BERT-base (Devlin et al., 2019)), and the number of state-space machines (SSM) to 128 with 64 states for each SSM.

Both bidirectional S4 and autoregressive S4 models were pretrained on large-scale texts. The autoregressive S4 was pretrained with the casual language modeling task on the whole English subset of Wikipedia. The second iteration of pretraining, initialized with the pretrained Wikipedia checkpoint, was on the whole Pubmed PMC Open Access Subset. The bidirectional S4 models were pretrained on the same datasets but with the mask language modeling task using the same masking settings as in BERT (Devlin et al., 2019). The maximum sequence length for pretraining was set to 8192 and the effective batch size was 256.

All models were optimized with AdamW optimizer with a learning rate of 1e-4 and the learning rate was dynamically adjusted using the Cosine Scheduler with a warm-up step of 1000. The pretraining took place on 8×RTX 3090 GPU with 24GB of memory. The training was stopped when the evaluation loss stopped to decrease (~12k steps for all models). The autoregressive and bi-direction checkpoints pre-trained on these datasets will be released together with this paper.

Transformer implementation. Transformer models were all implemented through the `Transformers` package (Wolf et al., 2020) with default dimensions. The autoregressive model was LongT5 (Guo et al., 2021) initialized by the `long-t5-tglobal-base` checkpoint. Other transformers were based on the Nystromformer (Xiong et al., 2021) with initialization from the public checkpoint `uw-madison/nystromformer-4096`.

Re-ranker training. For global re-rankers, the maximum input length was set to 4096 with an effective batch size of 32. For other models, the effective batch size was 64 and the maximum length was 2048, as this length was enough to cover almost all of the data samples. Models were trained with a maximum of 5 epochs and only the model with the best validation performance was kept. All models were trained using the AdamW optimizer (Loshchilov and Hutter, 2018) with a learning rate of 5e-5. We used a cosine scheduler with a warm-up step of 1000 to automatically adjust the learning rate during training. All ranking models were trained on a single V100 GPU with 16GB of memory.

7. Results

7.1. Main results

Our main results are summarized in Table 2. All neural re-ranking models outperform the baseline Expert System in both metrics, suggesting that re-ranking does up-rank the more relevant history-taking questions. Among the neural baselines, the transformer-based cross-encoder outperforms the bi-encoder, which is consistent with previous findings (Nogueira

Model	Dev		Test	
	nDCG	mAP	nDCG	mAP
Expert System (Baseline)	0.592	0.383	0.570	0.349
Bi-encoder (Transformer)	0.690	0.548	0.677	0.531
Cross-encoder				
Transformer	0.718	0.584	0.706	0.566
Nystromformer	0.653	0.496	0.654	0.497
Bidirectional S4 (Wiki Pretraining)	0.648	0.490	0.641	0.481
Bidirectional S4 (Pubmed Pretraining)	0.643	0.483	0.630	0.464
Autoregressive Re-ranker				
LongT5-base	0.690	0.546	0.678	0.529
Autoregressive S4 (Wiki Pretraining)	0.658	0.502	0.648	0.490
Autoregressive S4 (Pubmed Pretraining)	0.654	0.498	0.642	0.484
Global Re-ranker				
Transformer				
+ Pointwise loss: BCE	0.744	0.618	0.743	0.618
+ Pairwise loss: RankNet	0.739	0.612	0.735	0.603
+ Pairwise loss: LambdaLoss	0.739	0.616	0.739	0.612
+ Listwise loss: ListNet	0.737	0.609	0.740	0.610
+ Listwise loss: ListMLE	0.727	0.597	0.721	0.587
+ Listwise loss: ApproxNDCG	0.701	0.555	0.697	0.550
+ Listwise loss: NeuralNDCG	0.742	0.617	0.741	0.612
Nystromformer				
+ Pointwise loss: BCE	0.684	0.537	0.678	0.530
Bidirectional S4 (Wiki Pretraining)				
+ Pointwise loss: BCE	0.667	0.516	0.663	0.510
Bidirectional S4 (PubMed Pretraining)				
+ Pointwise loss: BCE	0.697	0.556	0.670	0.518

Table 2: Results of reranking experiments.

et al., 2019). Surprisingly, the LongT5 autoregressive re-ranker, despite having more parameters (220M parameters), also performs worse than the cross-encoder (\sim 110M parameters).

The best performance is achieved by the global re-ranker for both transformer and S4 architectures, regardless of the loss functions chosen. Among the various loss functions, the pointwise binary cross-entropy (BCE) performs the best. Our hypothesis is that since our ground truth relevance scores are binary rather than continuous, the current task does not make full use of the listwise loss functions.

The effectiveness of the global re-ranker lies in the fact that it attends to the semantic interactions not only between the dialogue and the candidate questions but also between the

candidate questions themselves. This allows the model to exploit the dependencies between history-taking questions, such as co-occurrence statistics, to improve ranking outcomes.

It is also worth noting that, despite its outstanding performance in some long sequence processing benchmark (Gu et al., 2021), S4 still lags behind transformers in the current task. One reason could be that the S4 model here has only been pre-trained on a comparatively small amount of texts, while transformers have been pre-trained on huge amounts of texts. Furthermore, the text sequences in our task range from a few hundred to about three thousand words, which might not be long enough for S4 to reveal its full potential.

7.2. Ablation analysis

We conducted ablation analyses on the global re-ranker to assess the impact of dialogue context length, the effect of type embeddings, and the effect of shuffling candidate question order. The results are displayed in Table 3.

Context length ablations. When ablating on context length, only the *last* N tokens of the dialogue were considered (full model is 4096 tokens, ablation are 3072, 2048, and 1024 tokens). While most of the text sequences were shorter than 1000 tokens, truncating texts still decreases test performance on some text sequences that are particularly long (longer than 1024), as some important information could be removed. In general, the global re-ranker benefits from getting more dialogue contexts, though this benefit seems to diminish after expanding to more than 2048 tokens.

Effect of position and type embeddings. We find that the removal of type embeddings (which are learned embeddings that differentiate whether the token is from dialogue or a candidate question) has almost no impact on the test performance. We reset the positional embeddings for each candidate questions in the input sequence, as this might help the model learn to be agnostic to the order of questions. We trained a model that used sequential positional embeddings for the input sequence. It turned out that positional embeddings played a minor role in training the global re-ranker.

Effect of shuffling. We tested the importance of permutation invariance with regard to the order of input candidate questions. The list of candidate questions $[\mathbf{q}_1, \mathbf{q}_2, \dots, \mathbf{q}_n]$ were concatenated with the prior dialogue as an input to the model. We found that while the expert system should produce questions in order or relevance, performance was significantly higher when the model was trained with shuffled order. We believe that this forces the model to learn to re-rank the questions without falling back to the original order of the candidate questions.

8. Discussion

In this work, we address an important problem of closing the training-inference gap for automated medical history-taking. Our approach, inspired by modern neural information retrieval systems, has two stages: (1) we use an expert system to suggest a list of candidate questions (out of possible thousands), (2) we train a machine-learned re-ranking model to re-rank expert system-suggested questions based on the free text of the doctor-patient dialogue.

Ablation	Dev		Test	
	nDCG	mAP	nDCG	mAP
- Maximum length: 4096 (Full)	0.744	0.618	0.743	0.618
- Maximum length: 3072	0.741	0.614	0.739	0.611
- Maximum length: 2048	0.746	0.622	0.747	0.622
- Maximum length: 1024	0.747	0.623	0.737	0.609
- No type embedding	0.733	0.610	0.732	0.607
- Sequential position embedding	0.749	0.625	0.741	0.613
- No random shuffling of questions	0.515	0.313	0.523	0.319

Table 3: Results of ablation studies on global re-ranker

To perform re-ranking (stage 2), we introduce a new approach which we call “global re-ranker”, and compare it to existing neural baselines. We also explore several language model back-bones including various transformers and structure-state-space (S4) models³. We find that while all neural re-ranking models outperform the original expert system, the global re-ranker with a full-attention transformer backbone performs the best with a 30% increase in nDCG and 77% increase in mAP over the first-stage recommendations.

While our results directly show the effectiveness of training a re-ranking model on top of an expert system for history taking, we believe this approach can also be applied to other decision support systems. The conditions under which this approach is beneficial are the following: (1) There exists a scoring system that has a training-inference gap (2) The space of possible predictions is very large, and as such would require a lot of data to machine-learn from scratch. One example beyond history-taking where we believe these conditions are satisfied is medical diagnosis prediction. There are many expert-system-derived diagnosis models, and training a diagnosis model from scratch can be difficult as the space of possible diagnosis is very large. Re-ranking could be used to close the gap between an off-the-shelf diagnostic expert system and the practice’s actual patient population outcomes.

Limitations This work is still limited in several ways. While our proposed global re-ranker had exhibited best overall performance over other ranking models, it is still computationally inefficient as the full attention transformers have quadratic computational complexity in processing long sequence. This will become a more serious bottleneck as the dialogue gets longer or the number of candidate questions increases. Secondly, the global re-ranker only learns the association between history-taking questions and the dialogue contexts from languages but it does not have the underlying medical knowledge. It will be paramount to augment such models with real medical knowledge such that it makes more informed decisions and does not biased against low frequency long-tail history-taking questions. In the future, we plan to investigate more effective approaches to encode long textual contexts and inject knowledge into language models.

3. As part of this publication, we release bi-directional and autoregressive S4 checkpoints pre-trained on the English Wikipedia and Pubmed PMC Open Access Subset. The code and checkpoints are available at: <https://github.com/curai/curai-research/tree/main/dialogue-contextualized-re-ranking>

8.1. Ethics

This work was done as part of a quality improvement activity as defined in 45CFR §46.104 (d)(4)(iii) – secondary research for which consent is not required for the purposes of “health care operations.”

References

- Qingyao Ai, Xuanhui Wang, Sebastian Bruch, Nadav Golbandi, Michael Bendersky, and Marc Najork. Learning groupwise multivariate scoring functions using deep neural networks. In *Proceedings of the 2019 ACM SIGIR international conference on theory of information retrieval*, pages 85–92, 2019.
- Iz Beltagy, Matthew E Peters, and Arman Cohan. Longformer: The long-document transformer. *arXiv preprint arXiv:2004.05150*, 2020.
- Chris Burges, Tal Shaked, Erin Renshaw, Ari Lazier, Matt Deeds, Nicole Hamilton, and Greg Hullender. Learning to rank using gradient descent. In *Proceedings of the 22nd international conference on Machine learning*, pages 89–96, 2005.
- Christopher Burges, Robert Ragno, and Quoc Le. Learning to rank with nonsmooth cost functions. *Advances in neural information processing systems*, 19, 2006.
- Zhe Cao, Tao Qin, Tie-Yan Liu, Ming-Feng Tsai, and Hang Li. Learning to rank: from pairwise approach to listwise approach. In *Proceedings of the 24th international conference on Machine learning*, pages 129–136, 2007.
- Shu Chen, Zeqian Ju, Xiangyu Dong, Hongchao Fang, Sicheng Wang, Yue Yang, Jiaqi Zeng, Ruisi Zhang, Ruoyu Zhang, Meng Zhou, Penghui Zhu, and Pengtao Xie. Meddialog: A large-scale medical dialogue dataset. *CoRR*, 2020.
- Xiaoyang Chen, Kai Hui, Ben He, Xianpei Han, Le Sun, and Zheng Ye. Co-bert: A context-aware bert retrieval model incorporating local and query-specific context. *arXiv preprint arXiv:2104.08523*, 2021.
- Krzysztof Marcin Choromanski, Valerii Likhoshesterov, David Dohan, Xingyou Song, Andreea Gane, Tamas Sarlos, Peter Hawkins, Jared Quincy Davis, Afroz Mohiuddin, Lukasz Kaiser, et al. Rethinking attention with performers. In *International Conference on Learning Representations*, 2020.
- Rhys Compton, Ilya Valmianski, Li Deng, Costa Huang, Namit Katariya, Xavier Amatriain, and Anitha Kannan. Medcod: A medically-accurate, emotive, diverse, and controllable dialog system. In *Machine Learning for Health*, pages 110–129. PMLR, 2021.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, June 2019.

- Tianyu Gao, Xingcheng Yao, and Danqi Chen. SimCSE: Simple contrastive learning of sentence embeddings. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6894–6910, November 2021.
- Albert Gu, Karan Goel, and Christopher Re. Efficiently modeling long sequences with structured state spaces. In *International Conference on Learning Representations*, 2021.
- Albert Gu, Ankit Gupta, Karan Goel, and Christopher Ré. On the parameterization and initialization of diagonal state space models. *arXiv preprint arXiv:2206.11893*, 2022.
- Mandy Guo, Joshua Ainslie, David Uthus, Santiago Ontanon, Jianmo Ni, Yun-Hsuan Sung, and Yinfei Yang. Longt5: Efficient text-to-text transformer for long sequences. *arXiv preprint arXiv:2112.07916*, 2021.
- John R Hampton, MJ Harrison, John R Mitchell, Jane S Prichard, and Carol Seymour. Relative contributions of history-taking, physical examination, and laboratory investigation to diagnosis and management of medical outpatients. *Br Med J*, 2(5969):486–489, 1975.
- Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- Samuel Humeau, Kurt Shuster, Marie-Anne Lachaux, and Jason Weston. Poly-encoders: Architectures and pre-training strategies for fast and accurate multi-sentence scoring. In *International Conference on Learning Representations*, 2019.
- Kalervo Järvelin and Jaana Kekäläinen. Ir evaluation methods for retrieving highly relevant documents. SIGIR '00. Association for Computing Machinery, 2000. ISBN 1581132263.
- Hao-Cheng Kao, Kai-Fu Tang, and Edward Chang. Context-aware symptom checking for disease diagnosis using hierarchical reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.
- Nikita Kitaev, Lukasz Kaiser, and Anselm Levskaya. Reformer: The efficient transformer. In *International Conference on Learning Representations*, 2019.
- Y. LeCun, B. Boser, J. S. Denker, D. Henderson, R. E. Howard, W. Hubbard, and L. D. Jackel. Backpropagation applied to handwritten zip code recognition. *Neural Computation*, 1(4):541–551, 1989. doi: 10.1162/neco.1989.1.4.541.
- Jimmy Lin, Rodrigo Nogueira, and Andrew Yates. Pretrained transformers for text ranking: Bert and beyond. *Synthesis Lectures on Human Language Technologies*, 14(4):1–325, 2021.
- Tie-Yan Liu et al. Learning to rank for information retrieval. *Foundations and Trends® in Information Retrieval*, 3(3):225–331, 2009.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019.

- Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2018.
- Harsh Mehta, Ankit Gupta, Ashok Cutkosky, and Behnam Neyshabur. Long range language modeling via gated state spaces. *arXiv preprint arXiv:2206.13947*, 2022.
- Sewon Min, Kenton Lee, Ming-Wei Chang, Kristina Toutanova, and Hannaneh Hajishirzi. Joint passage ranking for diverse multi-answer retrieval. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6997–7008, November 2021.
- Rodrigo Nogueira and Kyunghyun Cho. Passage re-ranking with bert. *arXiv preprint arXiv:1901.04085*, 2019.
- Rodrigo Nogueira, Wei Yang, Kyunghyun Cho, and Jimmy Lin. Multi-stage document ranking with bert. *arXiv preprint arXiv:1910.14424*, 2019.
- Rodrigo Nogueira, Zhiying Jiang, Ronak Pradeep, and Jimmy Lin. Document ranking with a pretrained sequence-to-sequence model. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 708–718, November 2020.
- Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.
- Przemyslaw Pobrotyn and Radoslaw Bialobrzeski. Neuralndcg: Direct optimisation of a ranking metric via differentiable relaxation of sorting. *ArXiv*, abs/2102.07831, 2021.
- Ronak Pradeep, Rodrigo Nogueira, and Jimmy Lin. The expando-mono-duo design pattern for text ranking with pretrained sequence-to-sequence models. *arXiv preprint arXiv:2101.05667*, 2021.
- Tao Qin, Tie-Yan Liu, and Hang Li. A general approximation framework for direct optimization of information retrieval measures. *Information retrieval*, 13(4):375–397, 2010.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, Peter J Liu, et al. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21(140):1–67, 2020.
- Nils Reimers and Iryna Gurevych. Sentence-BERT: Sentence embeddings using Siamese BERT-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992, Hong Kong, China, November 2019. Association for Computational Linguistics. doi: 10.18653/v1/D19-1410.
- Hagen Soltau, Mingqiu Wang, Izhak Shafran, and Laurent El Shafey. Understanding medical conversations: Rich transcription, confidence scores & information extraction. *CoRR*, 2021.
- Kai-Fu Tang, Hao-Cheng Kao, Chun-Nan Chou, and Edward Y Chang. Inquire and diagnose: Neural symptom checking ensemble using deep reinforcement learning. In *NIPS Workshop on Deep Reinforcement Learning*, 2016.

- Yi Tay, Mostafa Dehghani, Samira Abnar, Yikang Shen, Dara Bahri, Philip Pham, Jinfeng Rao, Liu Yang, Sebastian Ruder, and Donald Metzler. Long range arena: A benchmark for efficient transformers. In *International Conference on Learning Representations*, 2020.
- Nandan Thakur, Nils Reimers, Johannes Daxenberger, and Iryna Gurevych. Augmented sbert: Data augmentation method for improving bi-encoders for pairwise sentence scoring tasks. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 296–310, 2021.
- Ilya Valmianski, Nave Frost, Navdeep Sood, Yang Wang, Baodong Liu, James J. Zhu, Sunil Karumuri, Ian M. Finn, and Daniel S. Zisook. Smarttriage: A system for personalized patient data capture, documentation generation, and decision support. In *Proceedings of Machine Learning for Health*, volume 158 of *Proceedings of Machine Learning Research*, pages 75–96. PMLR, 04 Dec 2021.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- Mengqian Wang, Ilya Valmianski, Xavier Amatriain, and Anitha Kannan. Learning functional sections in medical conversations: iterative pseudo-labeling and human-in-the-loop approach, 2022.
- Sinong Wang, Belinda Z Li, Madian Khabisa, Han Fang, and Hao Ma. Linformer: Self-attention with linear complexity. *arXiv preprint arXiv:2006.04768*, 2020.
- DL Watson. Health interview responses compared with medical records. *Vital and health statistics. Ser. 1, Programs and collection procedures*, (3):1–74, 1965.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 conference on empirical methods in natural language processing: system demonstrations*, pages 38–45, 2020.
- Fen Xia, Tie-Yan Liu, Jue Wang, Wensheng Zhang, and Hang Li. Listwise approach to learning to rank: theory and algorithm. In *Proceedings of the 25th international conference on Machine learning*, pages 1192–1199, 2008.
- Yunyang Xiong, Zhanpeng Zeng, Rudrasis Chakraborty, Mingxing Tan, Glenn Fung, Yin Li, and Vikas Singh. Nyströmformer: A nyström-based algorithm for approximating self-attention. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 14138–14148, 2021.
- Manzil Zaheer, Guru Guruganesh, Kumar Avinava Dubey, Joshua Ainslie, Chris Alberti, Santiago Ontanon, Philip Pham, Anirudh Ravula, Qifan Wang, Li Yang, et al. Big bird: Transformers for longer sequences. *Advances in Neural Information Processing Systems*, 33:17283–17297, 2020.

George Zerveas, Navid Rekasaz, Daniel Cohen, and Carsten Eickhoff. Coder: An efficient framework for improving retrieval through contextualized document embedding reranking. *arXiv preprint arXiv:2112.08766*, 2021.