

Fast Prototyping a Dialogue Comprehension System for Nurse-Patient Conversations on Symptom Monitoring

Zhengyuan Liu^{*}, Hazel Lim^{*}, Nur Farah Ain Binte Suhaimi^{*}, Shao Chuen Tong[†],
Sharon Ong[‡], Angela Ng[‡], Sheldon Lee[†], Michael R. Macdonald[†],
Savitha Ramasamy^{*}, Pavitra Krishnaswamy^{*}, Wai Leng Chow[†] and Nancy F. Chen^{*}

^{*}Institute for Infocomm Research, A*STAR, Singapore

[†]Health Services Research, Changi General Hospital, Singapore

[‡]Health Management Unit, Changi General Hospital, Singapore

[†]Department of Cardiology, Changi General Hospital, Singapore

{liu.zhengyuan, nfychen}@i2r.a-star.edu.sg

Abstract

Data for human-human spoken dialogues for research and development are currently very limited in quantity, variety, and sources; such data are even scarcer in healthcare. In this work, we investigate fast prototyping of a dialogue comprehension system by leveraging on minimal nurse-to-patient conversations. We propose a framework inspired by nurse-initiated clinical symptom monitoring conversations to construct a simulated human-human dialogue dataset, embodying linguistic characteristics of spoken interactions like thinking aloud, self-contradiction, and topic drift. We then adopt an established bidirectional attention pointer network on this simulated dataset, achieving more than 80% F1 score on a held-out test set from real-world nurse-to-patient conversations. The ability to automatically comprehend conversations in the healthcare domain by exploiting only limited data has implications for improving clinical workflows through red flag symptom detection and triaging capabilities. We demonstrate the feasibility for efficient and effective extraction, retrieval and comprehension of symptom checking information discussed in multi-turn human-human spoken conversations.

1 Introduction

1.1 Problem Statement

Spoken conversations still remain the most natural and effortless means of human communication. Thus a lot of valuable information is conveyed and exchanged in such an unstructured form. In telehealth settings, nurses might call discharged patients who have returned home to continue to monitor their health status. Human language technology that can efficiently and effectively extract

key information from such conversations is clinically useful, as it can help streamline workflow processes and digitally document patient medical information to increase staff productivity. In this work, we design and prototype a dialogue comprehension system in the question-answering manner, which is able to comprehend spoken conversations between nurses and patients to extract clinical information¹.

1.2 Motivation of Approach

Machine comprehension of written passages has made tremendous progress recently. Large quantities of supervised training data for reading comprehension (e.g. SQuAD (Rajpurkar et al., 2016)), the wide adoption and intense experimentation of neural modeling (Seo et al., 2017; Wang et al., 2017), and the advancements in vector representations of word embeddings (Pennington et al., 2014; Devlin et al., 2018) all contribute significantly to the achievements obtained so far. The first factor, the availability of large scale datasets, empowers the latter two factors. To date, there is still very limited well-annotated large-scale data suitable for modeling human-human spoken dialogues. Therefore, it is not straightforward to directly port over the recent endeavors in reading comprehension to dialogue comprehension tasks.

In healthcare, conversation data is even scarcer due to privacy issues. Crowd-sourcing is an ef-

¹N.F.C., P.K., R.S. and C.W.L. conceptualized the overall research programme; N.F.C. and L.Z. proposed and developed the proposed approach; L.Z. and N.F.C. developed methods for data analysis; L.Z., L.J.H., N.F.S., and N.F.C. constructed the corpus; T.S.C, S.O., A.N. S.L.G, and M.R.M acquired, prepared, and validated clinical data; L.Z., N.F.C., P.K., S.L.G., M.R.M and C.W.L interpreted results; L.Z., N.F.C., L.J.H, N.F.S., P.K. and C.W.L wrote the paper.

efficient way to annotate large quantities of data, but less suitable for healthcare scenarios, where domain knowledge is required to guarantee data quality. To demonstrate the feasibility of a dialogue comprehension system used for extracting key clinical information from symptom monitoring conversations, we developed a framework to construct a simulated human-human dialogue dataset to bootstrap such a prototype. Similar efforts have been conducted for human-machine dialogues for restaurant or movie reservations (Shah et al., 2018). To the best of our knowledge, no one to date has done so for human-human conversations in healthcare.

1.3 Human-human Spoken Conversations

Human-human spoken conversations are a dynamic and interactive flow of information exchange. While developing technology to comprehend such spoken conversations presents similar technical challenges as machine comprehension of written passages (Rajpurkar et al., 2018), the challenges are further complicated by the interactive nature of human-human spoken conversations:

(1) Zero anaphora is more common: Coreference resolution of spoken utterances from multiple speakers is needed. For example, in Figure 1(a) *headaches*, *the pain*, *it*, *head bulging* all refer to the patient’s headache symptom, but they were uttered by different speakers and across multiple utterances and turns. In addition, anaphors are more likely to be omitted (see Figure 1(a) A4) as this does not affect the human listeners understanding, but it might be challenging for computational models.

(2) Thinking aloud more commonly occurs: Since it is more effortless to speak than to type, one is more likely to reveal her running thoughts when talking. In addition, one cannot retract what has been uttered, while in text communications, one is more likely to confirm the accuracy of the information in a written response and revise if necessary before sending it out. Thinking aloud can lead to self-contradiction, requiring more context to fully understand the dialogue; e.g., in A6 in Figure 1(a), the patient at first says he has none of the symptoms asked, but later revises his response saying that he does get dizzy after running.

(3) Topic drift is more common and harder to detect in spoken conversations: An example is shown in Figure 1(a) in A3, where *No* is actu-

Q1: Hi Mr.[name], you were discharged on [date]. There are some questions i'd like to check with you.
A1: Ok, Ok.

Q2: Is your left leg still swollen? You said so the last time I called you?
A2: Yes, only a bit when I drink too much water.

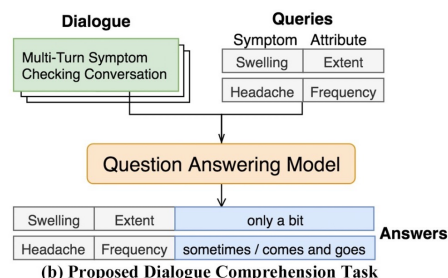
Q3: Well, I see, how about your cough?
A3: Cough? No... I get headaches sometimes.

Q4: How is the pain?
A4: Quite serious, comes and goes, you know.

Q5: When did it first start?
A5: Ah... I first felt my head bulging maybe since two weeks ago.

Q6: Headache, ok, got it. Any other symptoms? No cough, no dizziness, no breathlessness?
A6: No no, I think I'm ok... Ah, I get dizzy after running.

(a) Multi-Turn Spoken Dialogue Example



(b) Proposed Dialogue Comprehension Task

Symptom	Activity	Extent	Time	Frequency	Location
Swelling	when drink too much water	only a bit			left leg
Headache		quite serious	since two weeks ago	sometimes / comes and goes	
Dizziness	after running				
No Symptom:					
Cough, Breathlessness.					

(c) Clinical Use Case: Extractive Summary of Symptom Checking Dialogues

Figure 1: Dialogue comprehension of symptom checking conversations.

ally referring to *cough* in the previous question, and then the topic is shifted to *headache*. In spoken conversations, utterances are often incomplete sentences so traditional linguistic features used in written passages such as punctuation marks indicating syntactic boundaries or conjunction words suggesting discourse relations might no longer exist.

1.4 Dialogue Comprehension Task

Figure 1(b) illustrates the proposed dialogue comprehension task using a question answering (QA) model. The input are a multi-turn symptom checking dialogue D and a query Q specifying a *symptom* with one of its *attributes*; the output is the extracted answer A from the given dialogue. A training or test sample is defined as $S = \{D, Q, A\}$. Five attributes, specifying certain details of clinical significance, are defined to characterize the an-

swer types of *A*: (1) the *time* the patient has been experiencing the symptom, (2) *activities* that trigger the symptom (to occur or worsen), (3) the *extent* of seriousness, (4) the *frequency* occurrence of the symptom, and (5) the *location* of symptom. For each symptom/attribute, it can take on different linguistic expressions, defined as *entities*. Note that if the queried symptom or attribute is not mentioned in the dialogue, the groundtruth output is “No Answer”, as in (Rajpurkar et al., 2018).

2 Related Work

2.1 Reading Comprehension

Large-scale reading comprehension tasks like SQuAD (Rajpurkar et al., 2016) and MARCO (Nguyen et al., 2016) provide question-answer pairs from a vast range of written passages, covering different kinds of factual answers involving entities such as location and numerical values. Furthermore, HotpotQA (Yang et al., 2018) requires multi-step inference and provides numerous answer types. CoQA (Reddy et al., 2018) and QuAC (Choi et al., 2018) are designed to mimic multi-turn information-seeking discussions of the given material. In these tasks, contextual reasoning like coreference resolution is necessary to grasp rich linguistic patterns, encouraging semantic modeling beyond naive lexical matching. Neural networks contribute to impressive progress in semantic modeling: distributional semantic word embeddings (Pennington et al., 2014), contextual sequence encoding (Sutskever et al., 2014; Gehring et al., 2017) and the attention mechanism (Luong et al., 2015; Vaswani et al., 2017) are widely adopted in state-of-the-art comprehension models (Seo et al., 2017; Wang et al., 2017; Devlin et al., 2018).

While language understanding tasks in dialogue such as domain identification (Ravuri and Stolcke, 2015), slot filling (Kurata et al., 2016) and user intent detection (Wen et al., 2016) have attracted much research interest, work in dialogue comprehension is still limited, if any. It is labor-intensive and time-consuming to obtain a critical mass of annotated conversation data for computational modeling. Some propose to collect text data from human-machine or machine-machine dialogues (Li et al., 2016; Shah et al., 2018). In such cases, as human speakers are aware of current limitations of dialogue systems or due to predefined assumptions of user simulators, there are

fewer cases of zero anaphora, thinking aloud, and topic drift, which occur more often in human-human spoken interactions.

2.2 NLP for Healthcare

There is emerging interest in research and development activities at the intersection of machine learning and healthcare^{2 3}, of which much of the NLP related work are centered around social media or online forums (e.g., (Wallace et al., 2014; Lyles et al., 2013)), partially due to the world wide web as a readily available source of information. Other work in this area uses public data sources such as MIMIC⁴ in electronic health records: text classification approaches have been applied to analyze unstructured clinical notes for ICD code assignment (Baumel et al., 2017) and automatic intensive emergency prediction (Grnarova et al., 2016). Sequence-to-sequence textual generation has been used for readable notes based on medical and demographic recordings (Liu, 2018). For mental health, there has been more focus on analyzing dialogues. For example, sequential modeling of audio and text have helped detect depression from human-machine interviews (Al Hanai et al., 2018). However, few studies have examined human-human spoken conversations in healthcare settings.

3 Real-World Data Analysis

3.1 Data Preparation

We used recordings of nurse-initiated telephone conversations for congestive heart failure patients undergoing telemonitoring, post-discharge from the hospital. The clinical data was acquired by the Health Management Unit at Changi General Hospital. This research study was approved by the SingHealth Centralised Institutional Review Board (Protocol 1556561515). The patients were recruited during 2014-2016 as part of their routine care delivery, and enrolled into the telemonitoring health management program with consent for use of anonymized versions of their data for research.

The dataset comprises a total of 353 conversations from 40 speakers (11 nurses, 16 patients, and 13 caregivers) with consent to the use of anonymized data for research. The speakers are

²ML4H: Machine Learning for Health, Workshop at NeurIPS 2018 <https://ml4health.github.io/2018/>

³2018 Workshop on Health Intelligence (W3PHIAI 2018) <http://w3phiai2018.w3phi.com/>

⁴<https://mimic.physionet.org/>

Utterance Type	%
Open-ended Inquiry	31.8
Detailed Inquiry	33.0
Multi-Intent Inquiry	15.5
Reconfirmation Inquiry	21.3
Inquiry with Transitional Clauses	8.5
Yes/No Response	52.1
Detailed Response	29.4
Response with Revision	5.1
Response with Topic Drift	11.1
Response with Transitional Clauses	9.5
Sampled Turn Number	1200

Table 1: Linguistic characterization of inquiry-response types and their occurrence frequency from the seed data in Section 3.2.

38 to 88 years old, equally distributed across gender, and comprise a range of ethnic groups (55% Chinese, 17% Malay, 14% Indian, 3% Eurasian, and 11% unspecified). The conversations cover 11 topics (e.g., medication compliance, symptom checking, education, greeting) and 9 symptoms (e.g., chest pain, cough) and amount to 41 hours.

Data preprocessing and anonymization were performed by a data preparation team, separate from the data analysis team to maintain data confidentiality. The data preparation team followed standard speech recognition transcription guidelines, where words are transcribed verbatim to include false starts, disfluencies, mispronunciations, and private self-talk. Confidential information were marked and clipped off from the audio and transcribed with predefined tags in the annotation. Conversation topics and clinical symptoms were also annotated and clinically validated by certified telehealth nurses.

3.2 Linguistic Characterization on Seed Data

To analyze the linguistic structure of the inquiry-response pairs in the entire 41-hour dataset, we randomly sampled a seed dataset consisting of 1,200 turns and manually categorized them to different types, which are summarized in Table 1 along with the corresponding occurrence frequency statistics. Note that each given utterance could be categorized to more than one type. We elaborate on each utterance type below.

Open-ended Inquiry: Inquiries about general well-being or a particular symptom; e.g., “*How are you feeling?*” and “*Do you cough?*”

Detailed Inquiry: Inquiries with specific details that prompt yes/no answers or clarifications; e.g., “*Do you cough at night?*”

Multi-Intent Inquiry: Inquiring more than one symptom in a question; e.g., “*Any cough, chest pain, or headache?*”

Reconfirmation Inquiry: The nurse reconfirms particular details; e.g., “*Really? At night?*” and “*Serious or mild?*”. This case is usually related to explicit or implicit coreferencing.

Inquiry with Transitional Clauses: During spoken conversations, one might repeat what the other party said, but it is unrelated to the main clause of the question. This is usually due to private self-talk while thinking aloud, and such utterances form a transitional clause before the speaker starts a new topic; e.g., “*Chest pain... no chest pain, I see... any cough?*”.

Yes/No Response: Yes/No responses seem straightforward, but sometimes lead to misunderstanding if one does not interpret the context appropriately. One case is tag questions: A: “*You don’t cough at night, do you?*” B: “*Yes, yes*” A: “*cough at night?*” B: “*No, no cough*”. Usually when the answer is unclear, clarifying inquiries will be asked for reconfirmation purposes.

Detailed Response: Responses that contain specific information of one symptom, like “*I felt tightness in my chest*”.

Response with Revision: Revision is infrequent but can affect comprehension significantly. One cause is thinking aloud so a later response overrules the previous one; e.g., “*No dizziness, oh wait... last week I felt a bit dizzy when biking*”.

Response with Topic Drift: When a symptom/topic like headache is inquired, the response might be: “*Only some chest pain at night*”, not referring to the original symptom (headache) at all.

Response with Transitional Clauses: Repeating some of the previous content, but often unrelated to critical clinical information and usually followed by topic drift. For example, “*Swelling... swelling... I don’t cough at night*”.

4 Simulating Symptom Monitoring Dataset for Training

We divide the construction of data simulation into two stages. In Section 4.1, we build templates and expression pools using linguistic analysis followed by manual verification. In Section 4.2, we present our proposed framework for generating simulated training data. The templates and framework are verified for logical correctness and clinical soundness.

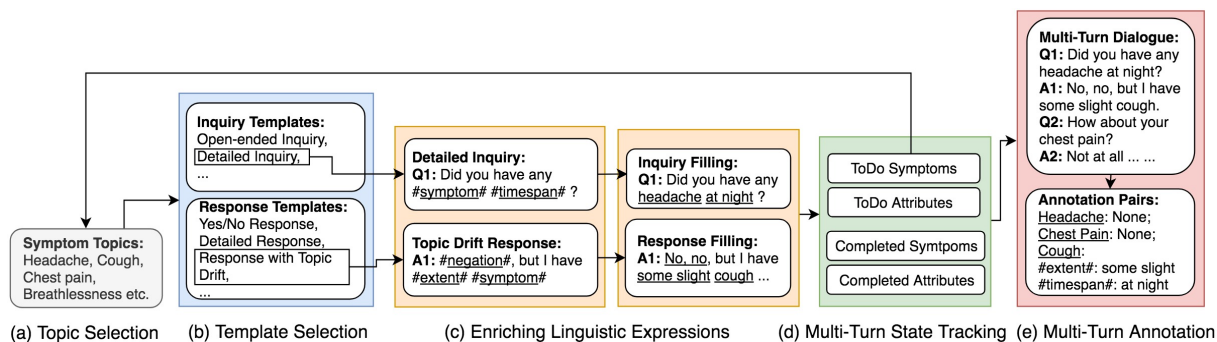


Figure 2: Simulated data generation framework.

4.1 Template Construction

4.1.1 Linguistically-Inspired Templates

Each utterance in the seed data is categorized according to Table 1 and then abstracted into templates by replacing entity phrases like *cough* and *often* with respective placeholders “*#symptom#*” and “*#frequency#*”. The templates are refined through verifying logical correctness and injecting expression diversity by linguistically trained researchers. As these replacements do not alter the syntactic structure, we interchange such placeholders with various verbal expressions to enlarge the simulated training set in Section 4.2. Clinical validation was also conducted by certified telehealth nurses.

4.1.2 Topic Expansion & Symptom Customization

For the 9 symptoms (e.g. chest pain, cough) and 5 attributes (e.g., extent, frequency), we collect various expressions from the seed data, and expand them through synonym replacement. Some attributes are unique to a particular symptom; e.g., “*left leg*” in *#location#* is only suitable to describe the symptom *swelling*, but not the symptom *headache*. Therefore, we only reuse general expressions like “*slight*” in *#extent#* across different symptoms to diversify linguistic expressions.

4.1.3 Expression Pool for Linguistic Diversity

Two linguistically trained researchers constructed expression pools for each symptom and each attribute to account for different types of paraphrasing and descriptions. These expression pools are used in Section 4.2 (c).

4.2 Simulated Data Generation Framework

Figure 2 shows the five steps we use to generate multi-turn symptom monitoring dialogue samples.

(a) Topic Selection: While nurses might prefer to inquire the symptoms in different orders depending on the patient’s history, our preliminary analysis shows that modeling results do not differ noticeably if topics are of equal prior probabilities. Thus we adopt this assumption for simplicity.

(b) Template Selection: For each selected topic, one inquiry template and one response template are randomly chosen to compose a turn. To minimize adverse effects of underfitting, we redistributed the frequency distribution in Table 1: For utterance types that are below 15%, we boosted them to 15%, and the overall relative distribution ranking is balanced and consistent with Table 1.

(c) Enriching Linguistic Expressions: The placeholders in the selected templates are substituted with diverse expressions from the expression pools in Section 4.1.3 to characterize the symptoms and their corresponding attributes.

(d) Multi-Turn Dialogue State Tracking: A greedy algorithm is applied to complete conversations. A “completed symptoms” list and a “to-do symptoms” list are used for symptom topic tracking. We also track the “completed attributes” and “to-do attributes”. For each symptom, all related attributes are iterated. A dialogue ends only when all possible entities are exhausted, generating a multi-turn dialogue sample, which encourages the model to learn from the entire discussion flow rather than a single turn to comprehend contextual dependency. The average length of a simulated dialogue is 184 words, which happens to be twice as long as an average dialogue from the real-world evaluation set. Moreover, to model the roles of the respondents, we set the ratio between patients and caregivers to 2:1; this statistic is inspired by the real scenarios in the seed dataset. For both the caregivers and patients, we assume equal probability of both genders. The corresponding pronouns in the conversations are thus determined by

the role and gender of these settings.

(e) Multi-Turn Sample Annotation: For each multi-turn dialogue, a query is specified by a symptom and an attribute. The groundtruth output of the QA system is automatically labeled based on the template generation rules, but also manually verified to ensure annotation quality. Moreover, we adopt the unanswerable design in (Rajpurkar et al., 2018): when the patient does not mention a particular symptom, the answer is defined as “No Answer”. This process is repeated until all logical permutations of symptoms and attributes are exhausted.

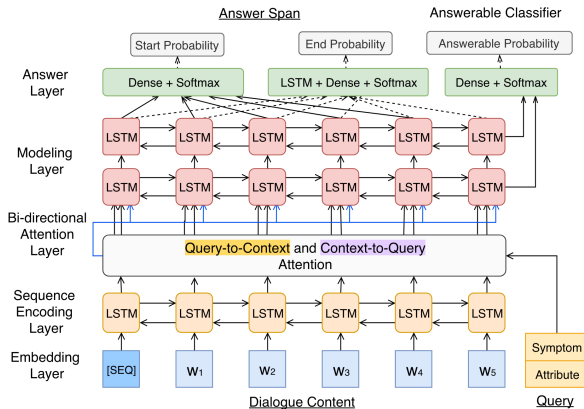


Figure 3: Bi-directional attention pointer network with an answerable classifier for dialogue comprehension.

5 Experiments

5.1 Model Design

We implemented an established model in reading comprehension, a bi-directional attention pointer network (Seo et al., 2017), and equipped it with an answerable classifier, as depicted in Figure 3. First, tokens in the given dialogue D and query Q are converted into embedding vectors. Then the dialogue embeddings are fed to a bi-directional LSTM encoding layer, generating a sequence of contextual hidden states. Next, the hidden states and query embeddings are processed by a bi-directional attention layer, fusing attention information in both context-to-query and query-to-context directions. The following two bi-directional LSTM modeling layers read the contextual sequence with attention. Finally, two respective linear layers with softmax functions are used to estimate token i ’s p_i^{start} and p_i^{end} probability of the answer span A .

In addition, we add a special tag “[SEQ]” at the head of D to account for the case of “No answer” (Devlin et al., 2018) and adopt an answer-

Training Samples	10k	50k	100k	150k
Base Evaluation Set:				
EM Score	15.41	80.33	89.68	91.45
F1 Score	50.63	89.18	92.27	94.17
Augmented Evaluation Set:				
EM Score	11.59	57.22	78.29	72.12
F1 Score	49.36	74.53	85.69	82.75
Real-World Evaluation Set:				
EM Score	38.81	42.93	78.23	75.41
F1 Score	46.29	52.68	80.18	78.09

Table 2: QA model evaluation results. Each sample is a simulated multi-turn conversation.

able classifier as in (Liu et al., 2018). More specifically, when the queried symptom or attribute is not mentioned in the dialogue, the answer span should point to the tag “[SEQ]” and answerable probability should be predicted as 0.

5.2 Implementation Details

The model was trained via gradient backpropagation with the cross-entropy loss function of answer span prediction and answerable classification, optimized by Adam algorithm (Kingma and Ba, 2015) with initial learning rate of 0.001. Pre-trained GloVe (Pennington et al., 2014) embeddings (size = 200) were used. We re-shuffled training samples at each epoch (batch size = 16). Out-of-vocabulary words (< 0.05%) were replaced with a fixed random vector. L2 regularization and dropout (rate = 0.2) were used to alleviate overfitting (Srivastava et al., 2014).

5.3 Evaluation Setup

To evaluate the effectiveness of our linguistically-inspired simulation approach, the model is trained on the simulated data (see Section 4.2). We designed 3 evaluation sets: (1) **Base Set** (1,264 samples) held out from the simulated data. (2) **Augmented Set** (1,280 samples) built by adding two out-of-distribution symptoms, with corresponding dialogue contents and queries, to the Base Set (“bleeding” and “cold”, which never appeared in training data). (3) **Real-World Set** (944 samples) manually delineated from the the symptom checking portions (approximately 4 hours) of real-world dialogues, and annotated as evaluation samples.

5.4 Results

Evaluation results are in Table 2 with exact match (EM) and F1 score in (Rajpurkar et al., 2016) metrics. To distinguish the correct answer span from the plausible ones which contain the same words, we measure the scores on the position indices of

	EM	F1
Augmented Evaluation Set:		
Best-trained Model	78.29	85.69
w/o Bi-Attention	72.08	78.57
w/o Pre-trained Embedding	56.98	72.31
Real-World Evaluation Set:		
Best-trained Model	78.23	80.18
w/o Bi-Attention	70.52	74.09
w/o Pre-trained Embedding	60.88	66.47

Table 3: Ablation experiments on 100K training size.

tokens. Our results show that both EM and F1 score increase with training sample size growing and the optimal size in our setting is 100k. The best-trained model performs well on both the Base Set and the Augmented Set, indicating that out-of-distribution symptoms do not affect the comprehension of existing symptoms and outputs reasonable answers for both in- and out-of-distribution symptoms. On the Real-World Set, we obtained 78.23 EM score and 80.18 F1 score respectively.

Error analysis suggests the performance drop from the simulated test sets is due to the following: 1) sparsity issues resulting from the expression pools excluding various valid but sporadic expressions. 2) nurses and patients occasionally chit-chat in the Real-World Set, which is not simulated in the training set. At times, these chit-chats make the conversations overly lengthy, causing the information density to be lower. These issues could potentially distract and confuse the comprehension model. 3) an interesting type of infrequent error source, caused by patients elaborating on possible causal relations of two symptoms. For example, a patient might say “*My giddiness may be due to all this cough*”. We are currently investigating how to close this performance gap efficiently.

5.5 Ablation Analysis

To assess the effectiveness of bi-directional attention, we bypassed the bi-attention layer by directly feeding the contextual hidden states and query embeddings to the modeling layer. To evaluate the pre-trained GloVe embeddings, we randomly initialized and trained the embeddings from scratch. These two procedures lead to 10% and 18% performance degradation on the Augmented Set and Real-World Set, respectively (see Table 3).

6 Conclusion

We formulated a dialogue comprehension task motivated by the need in telehealth settings to extract key clinical information from spoken conver-

sations between nurses and patients. We analyzed linguistic characteristics of real-world human-human symptom checking dialogues, constructed a simulated dataset based on linguistically inspired and clinically validated templates, and prototyped a QA system. The model works effectively on a simulated test set using symptoms excluded during training and on real-world conversations between nurses and patients. We are currently improving the model’s dialogue comprehension capability in complex reasoning and context understanding and also applying the QA model to summarization and virtual nurse applications.

Acknowledgements

Research efforts were supported by funding for Digital Health and Deep Learning I2R (DL2 SSF Project No: A1718g0045) and the Science and Engineering Research Council (SERC Project No: A1818g0044), A*STAR, Singapore. In addition, this work was conducted using resources and infrastructure provided by the Human Language Technology unit at I2R. The telehealth data acquisition was funded by the Economic Development Board (EDB), Singapore Living Lab Fund and Philips Electronics Hospital to Home Pilot Project (EDB grant reference number: S14-1035-RF-LLF H and W).

We acknowledge valuable support and assistance from Yulong Lin, Yuanxin Xiang, and Ai Ti Aw at the Institute for Infocomm Research (I2R); Weiliang Huang at the Changi General Hospital (CGH) Department of Cardiology, Hong Choon Oh at CGH Health Services Research, and Edris Atikah Bte Ahmad, Chiu Yan Ang, and Mei Foon Yap of the CGH Integrated Care Department.

We also thank Eduard Hovy and Bonnie Webber for insightful discussions and the anonymous reviewers for their precious feedback to help improve and extend this piece of work.

References

- Tuka Al Hanai, Mohammad Ghassemi, and James Glass. 2018. Detecting depression with audio/text sequence modeling of interviews. In *Proc. of INTERSPEECH*, pages 1716–1720.
- Tal Baumel, Jumana Nassour-Kassis, Raphael Cohen, Michael Elhadad, and Noemie Elhadad. 2017. Multi-label classification of patient notes a case study on ICD code assignment. In *Workshop on Health Intelligence, AAAI*.

- Eunsol Choi, He He, Mohit Iyyer, Mark Yatskar, Wentaoh Yih, Yejin Choi, Percy Liang, and Luke Zettlemoyer. 2018. Quac: Question answering in context. In *Proc. of EMNLP*, pages 2174–2184.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. BERT: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann N Dauphin. 2017. Convolutional sequence to sequence learning. *arXiv preprint arXiv:1705.03122*.
- Paulina Grnarova, Florian Schmidt, Stephanie L Hyland, and Carsten Eickhoff. 2016. Neural document embeddings for intensive care patient mortality prediction. *arXiv preprint arXiv:1612.00467*.
- Diederik P Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *Proc. of ICLR*.
- Gakuto Kurata, Bing Xiang, Bowen Zhou, and Mo Yu. 2016. Leveraging Sentence-level Information with Encoder LSTM for Semantic Slot Filling. In *Proc. of EMNLP*, pages 2077–2083.
- Xiujun Li, Zachary C Lipton, Bhuwan Dhingra, Lihong Li, Jianfeng Gao, and Yun-Nung Chen. 2016. A user simulator for task-completion dialogues. *arXiv preprint arXiv:1612.05688*.
- Peter J Liu. 2018. Learning to write notes in electronic health records. *arXiv preprint arXiv:1808.02622*.
- Xiaodong Liu, Wei Li, Yuwei Fang, Aerin Kim, Kevin Duh, and Jianfeng Gao. 2018. Stochastic answer networks for squad 2.0. *arXiv preprint arXiv:1809.09194*.
- Minh-Thang Luong, Hieu Pham, and Christopher D Manning. 2015. Effective approaches to attention-based neural machine translation. In *Proc. of EMNLP*, pages 1412–1421.
- Courtney R Lyles, Andrea López, Rena Pasick, and Urmimalla Sarkar. 2013. 5 mins of uncomfyness is better than dealing with cancer 4 a lifetime: an exploratory qualitative analysis of cervical and breast cancer screening dialogue on twitter. *Journal of Cancer Education*, 28(1):127–133.
- Tri Nguyen, Mir Rosenberg, Xia Song, Jianfeng Gao, Saurabh Tiwary, Rangan Majumder, and Li Deng. 2016. MS MARCO: A human generated machine reading comprehension dataset. *arXiv preprint arXiv:1611.09268*.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. GloVe: Global vectors for word representation. In *Proc. of EMNLP*, pages 1532–1543.
- Pranav Rajpurkar, Robin Jia, and Percy Liang. 2018. Know What You Don’t Know: Unanswerable Questions for SQuAD. In *Proc. of ACL*, pages 784–789.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. SQuAD: 100,000+ questions for machine comprehension of text. In *Proc. of EMNLP*, pages 2383–2392.
- Suman Ravuri and Andreas Stolcke. 2015. Recurrent neural network and lstm models for lexical utterance classification. In *Proc. of INTERSPEECH*.
- Siva Reddy, Danqi Chen, and Christopher D Manning. 2018. CoQA: A conversational question answering challenge. *arXiv preprint arXiv:1808.07042*.
- Minjoon Seo, Aniruddha Kembhavi, Ali Farhadi, and Hannaneh Hajishirzi. 2017. Bidirectional attention flow for machine comprehension. In *Proc. of ICLR*.
- Pararth Shah, Dilek Hakkani-Tür, Gokhan Tür, Abhinav Rastogi, Ankur Bapna, Neha Nayak, and Larry Heck. 2018. Building a conversational agent overnight with dialogue self-play. *arXiv preprint arXiv:1801.04871*.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In *NIPS*, pages 3104–3112.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *NIPS*, pages 5998–6008.
- Byron C Wallace, Michael J Paul, Urmimalla Sarkar, Thomas A Trikalinos, and Mark Dredze. 2014. A large-scale quantitative analysis of latent factors and sentiment in online doctor reviews. *Journal of the American Medical Informatics Association*, 21(6):1098–1103.
- Wenhui Wang, Nan Yang, Furu Wei, Baobao Chang, and Ming Zhou. 2017. Gated self-matching networks for reading comprehension and question answering. In *Proc. of ACL*, volume 1, pages 189–198.
- Tsung-Hsien Wen, David Vandyke, Nikola Mrksic, Milica Gasic, Lina M Rojas-Barahona, Pei-Hao Su, Stefan Ultes, and Steve Young. 2016. A network-based end-to-end trainable task-oriented dialogue system. In *Proc. of ACL*, pages 438–449.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William W Cohen, Ruslan Salakhutdinov, and Christopher D Manning. 2018. HotpotQA: A Dataset for Diverse, Explainable Multi-hop Question Answering. In *Proc. of EMNLP*, pages 2369–2380.