

MEDFILTER: Improving Extraction of Task-relevant Utterances through Integration of Discourse Structure and Ontological Knowledge

Sopan Khosla Shikhar Vashishth Jill Fain Lehman Carolyn Rose

Language Technologies Institute
Carnegie Mellon University, USA

{sopank, svashish, jfl, cprose}@cs.cmu.edu

Abstract

Information extraction from conversational data is particularly challenging because the task-centric nature of conversation allows for effective communication of implicit information by humans, but is challenging for machines. The challenges may differ between utterances depending on the role of the speaker within the conversation, especially when relevant expertise is distributed asymmetrically across roles. Further, the challenges may also increase over the conversation as more shared context is built up through information communicated implicitly earlier in the dialogue. In this paper, we propose the novel modeling approach MEDFILTER, which addresses these insights in order to increase performance at identifying and categorizing task-relevant utterances, and in so doing, positively impacts performance at a downstream information extraction task. We evaluate this approach on a corpus of nearly 7,000 doctor-patient conversations where MEDFILTER is used to identify medically relevant contributions to the discussion (achieving a 10% improvement over SOTA baselines in terms of area under the PR curve). Identifying task-relevant utterances benefits downstream medical processing, achieving improvements of 15%, 105%, and 23% respectively for the extraction of symptoms, medications, and complaints.

1 Introduction

In this paper, we propose a novel modeling approach that embodies insights regarding the organization of task-oriented conversations in order to improve performance at utterance classification over SOTA baseline approaches. Task-oriented conversations involve sharing task-relevant information that may be useful as the task ensues (Liu et al., 2019a; Kazi and Kahanda, 2019). Unfortunately, human-to-human conversations are less well structured than expository text, which is more often

the source material for information extraction and summarization. Expository text is typically structured top-down and organized around information flow. Task-oriented conversations, on the other hand, are typically organized around the task and knowledge of task structure provides an implicit scaffold for understanding. Thus speakers feel free to elide or imply important information rather than making it explicit. The challenges have been well documented (Waitzkin, 1989; Lacson et al., 2006). Prior work in utterance classification is a source of SOTA modeling approaches that perform relatively well despite these challenges while leaving much room for improvement.

Our evaluation in this paper specifically focuses on doctor-patient interactions. Doctor-patient interactions are task-oriented, expert-layperson interactions in which the concerns voiced by the layperson (e.g., symptoms), the underlying issue identified by the expert (e.g., complaint) and the prescribed solutions (e.g., medications) play a crucial part. Customer-service chats are another example of such dialogue. As in the general case, topic switching abounds: the doctor may jump from a question about a symptom to a statement providing initial assessment then back again, with or without waiting for a reply from the patient (which may, itself, be responsive or introduce a new concern). In addition, the participants make unequal contributions to different parts of the schema due to the inherent asymmetry between their roles in terms of knowledge and authority. Despite these challenges, humans are able to communicate very effectively in this way. Because of that, the issues increase as the conversation progresses and more shared context is built up, in part because of a certain amount of shared domain knowledge, despite differences in the extent and phrasing of it. In response to these insights, our proposed model, which we refer to as MEDFILTER, integrates elements of dis-

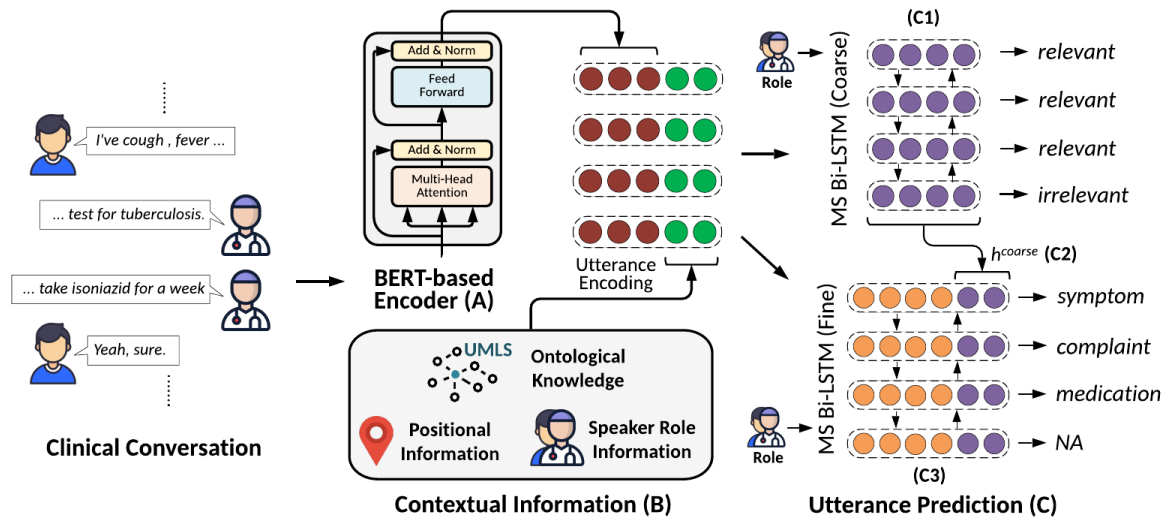


Figure 1: Overview of MEDFILTER. MEDFILTER first encodes each utterance of the given conversation using a BERT-based encoder (A). The obtained utterance embedding is concatenated with contextual information like speaker role, position of utterance in the conversation, and ontological knowledge (B). This is then fed to a MS-BiLSTM (C1) for medical relevance identification. MS-BiLSTM leverages speaker role information to learn speaker-specific context for each utterance. This contextual representation is concatenated with the utterance embedding (C2) and passed through another MS-BiLSTM (C3) which focuses on fine-grained categorization. Both tasks are jointly learned. Refer to Section 3 for more details.

course structure and ontological knowledge to improve utterance classification, the impact of which is also observed in a downstream extraction task. We evaluate the approach on a corpus of nearly 7,000 doctor-patient interactions as a case study.

Our proposed method, MEDFILTER¹, is illustrated in Figure 1 and described in detail in Section 3. Its architecture specifically reflects an awareness of the challenges above and begins to address them. In particular, the speaker’s role (i.e., doctor, patient, and other) and position within the interaction are both introduced as structuring variables. Insights from ontological knowledge are also made available through a domain ontology: specifically, the Unified Medical Language System (UMLS) (Bodenreider, 2004). From a more technical perspective, the architecture introduces a novel Multi-Speaker BiLSTM to learn role-specific context representations. MEDFILTER also benefits from the incorporation of a hierarchical loss that jointly learns the coarse-grained task of predicting medical relevance to improve fine-grained topic-based utterance classification. The ability to extract medically relevant utterances from doctor-patient conversations and categorize them into the medical topics/categories has a substantial practical impact in medical practice (Finley et al., 2018; Quiroz et al., 2019).

¹<https://github.com/sopankhosla/MedFilter>

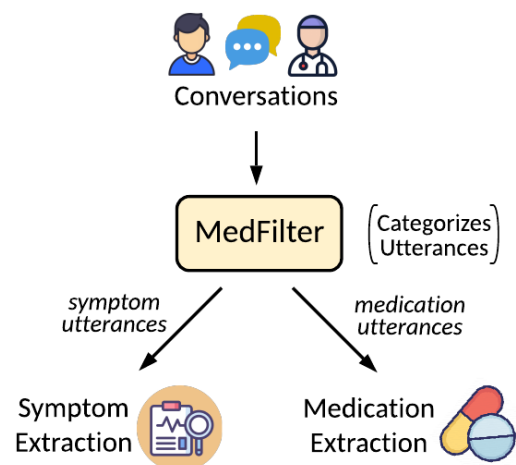


Figure 2: MEDFILTER as a part of extraction pipeline.

2 Related Work

Dialogue Summarization: In addition to the challenges noted earlier in the paper, other linguistic phenomena such as backchannels, false starts, and topic diffusion are prominent in human-to-human conversations. They add noise, which challenges the capabilities of otherwise effective summarization approaches such as pointer-generator networks (See et al., 2017; Liu et al., 2019b).

Some prior work has relied on an Information Extraction (IE) based approach to extract details about individual medical entities such as symptoms or medications (Du et al., 2019; Selvaraj

and Konam, 2019). However, recently, multiple studies (Lacson et al., 2006; Kocaballi et al., 2019; Liu et al., 2019a,b; Park et al., 2019) have shown the benefits of using the topical structure in goal-oriented dialogues to improve summarization. Within that scope, Liu et al. (2019a) introduce *key-point sequences* that describe the logical topic flow of the summary of customer-service chats. They propose a hierarchical transformer to predict these topics (key-points) for each utterance and use them as auxiliary labels to guide the summarization.

This past work inspires our work in which we extend the approach and then apply it in the more challenging domain of doctor-patient interactions. We consider it more challenging both in terms of the number of utterances per conversation (avg. 225 vs 20) and topic switches (Kocaballi et al., 2019). To improve the key-point sequence utterance-level topic classification approach (Liu et al., 2019a), we propose MEDFILTER that models speaker-specific context augmented with ontological knowledge and a hierarchical loss function.

Intent Classification: The problem of classifying utterances into medical topics/categories has many similarities with the task of utterance-level intent classification (Zhang et al., 2019; Budzianowski et al., 2018b; Qu et al., 2019). In our case, medical categories act as coarse-grained intents that drive the content of the discussion. Much of the previous work in intent classification caters to creating better dialog agents that condition their responses on the intent of the previous utterance (Budzianowski et al., 2018a; Bocklisch et al., 2017). For instance, Chen et al. (2019); Kim et al. (2017) propose intent classification as a text classification task where each utterance is considered a complete, independent command. However, this is not true in our case as the discussion about a medical category might range over multiple utterances, each dependent on context. Hence, we tackle the classification problem as a sequence-labeling task.

Sequence Labeling in Dialogue: Most prior work that employs sequence labeling for utterance classification in dialogues (Raheja and Tetreault, 2019; Liu et al., 2017; Jiao et al., 2019a) evaluates their systems on dialogue-act classification (Shriberg et al., 2004, 1998) or emotion recognition datasets (Poria et al., 2019). In this paper, we adopt state-of-the-art modeling approaches from the emotion recognition task (Jiao et al., 2019a,b) to serve as baselines in our evaluation since our task has not

previously been benchmarked.

3 Proposed Method: MEDFILTER

The overall architecture of MEDFILTER is shown in Figure 1. The input to MEDFILTER is a transcribed clinical conversation \mathcal{C} of form $\{u_1, u_2, \dots, u_n\}$, where each u_i represents an utterance. Each utterance in the conversation is passed through a BERT-based encoder (Fig. 1A and Sec. 3.1) to get a fixed-dimensional representation. Contextual information such as speaker role, the utterance’s position in the conversation, and ontological knowledge (Fig. 1B and Sec. 3.2) is then appended to the BERT representation. The encoding is input to the coarse Multi-Speaker BiLSTM (MS-BiLSTM) model (Fig. 1C1) followed by a fully-connected layer to classify the relevance of utterances for topical classification. The representation created by MS-BiSLTM (Coarse) is then concatenated with the utterance encoding (Fig. 1C2) and the resulting vector is fed to the fine-grained MS-BiLSTM (Fig. 1C3) to classify utterances into different medical categories (Sec. 3.3). MEDFILTER is jointly optimized on both classification tasks.

3.1 BERT-based Encoder

Given the superior modeling capabilities of long-range dependencies in Transformer-based models (Vaswani et al., 2017), we use pre-trained BERT (Devlin et al., 2019) for encoding each utterance u_i . We first encode each token in the utterance using BERT, i.e., $[h_{i1}^{\text{BERT}}, h_{i2}^{\text{BERT}}, \dots, h_{im}^{\text{BERT}}]$, where h_{ij} represents BERT-encoding of j^{th} token of u_i . Now, following Reimers and Gurevych (2019), we use MEAN pooling for obtaining a representation for the entire utterance (h_i^{Text}). Since the original pre-trained BERT model is trained on a general web corpus such as Wikipedia, it might not generalize well to our corpus. Therefore, we further fine tune the BERT model in a supervised manner for the task of predicting the utterance type.

3.2 Contextual Information

In addition to encoding the text of an utterance, we also make use of the following types of contextual information.

1. Speaker Role Info: In conversations in general, speaker identity helps ground co-references like *I*, *You*. In doctor-patient conversations, each of the speakers play a specific role in the goals of the interaction. For example, the doctor is more

likely to discuss medications than the patient. To allow the representation to be sensitive to speaker information, we map the speaker roles, namely, doctor, patient, and other, to a d -dimensional embedding ($\mathbf{h}_i^{\text{speaker}}$) which is learned during training and given to the model along with the text-based representation.

2. Positional Info: Clinical conversations often follow a pattern where topics like symptoms and complaints are discussed earlier in the dialog and prescribed medications are narrated in the middle or toward the end. To include this signal in MEDFILTER, we partition all the utterances in a conversation into k equal parts based on their position. For instance, if the conversation has 40 utterances and $k = 4$ then the initial 10 belong to 1st partition and the next 10 belong to 2nd and so on. Similar to speaker role information, a trainable embedding is associated with each partition ($\mathbf{h}_i^{\text{position}}$).

3. Ontological Knowledge: UMLS (Unified Medical Language System) (Bodenreider, 2004) is a combination of a semantic network and a meta-thesaurus. The semantic network consists of a set of 127 broad subject categories, or semantic types, which provide a consistent categorization of all concepts represented in the meta-thesaurus. In MEDFILTER, we use Quick-UMLS (Soldaini and Goharian, 2016), which identifies clinical mentions in an utterance and retrieves the associated UMLS Concept Unique Identifiers (CUIs) and semantic type, to inform our model about the type of medical phrases present in the input. We believe that types such as *Pharmacologic Substance*, *Symptoms*, and *Diseases* can be helpful in correctly classifying the utterances. We assign a trainable embedding to each semantic type. However, since each utterance can contain multiple clinical mentions of varied semantic types, we average the semantic-type embeddings for each mention present in the utterance and pass it to the model ($\mathbf{h}_i^{\text{semantic}}$).

3.3 Utterance Prediction

The classifier takes in the extended representation for each utterance u_i in the conversation given as

$$\mathbf{h}_i = [\mathbf{h}_i^{\text{Text}}; \mathbf{h}_i^{\text{speaker}}; \mathbf{h}_i^{\text{position}}; \mathbf{h}_i^{\text{semantic}}].$$

To explicitly model the separate roles performed by each speaker (as discussed in Section 1), we propose a novel module **Multi-Speaker BiLSTM (MS-BiLSTM)** that includes speaker-level BiLSTMs to learn the context for each speaker type

separately. We note, for example, that when the doctor is prescribing medications to the patient, she is more likely to expand on her previous utterance in order to discuss different details about the medicine, whereas the patient is most likely to give simple acknowledgments or ask questions in her turn. Having separate speaker-level BiLSTMs allows MS-BiLSTM to model this difference in the use of context.

MS-BiLSTM takes in \mathbf{h}_i and s_i (utterance’s speaker) as input. \mathbf{h}_i is passed through a background BiLSTM (BiLSTM_{bg}) and different speaker-level BiLSTMs (BiLSTM_{s_j}). Thus, if there are 3 speaker roles in the conversation, then the extended representation for each utterance (\mathbf{h}_i) would be input to 4 BiLSTMs (1 background BiLSTM + 3 speaker BiLSTMs). The hidden representations from BiLSTM_{bg} and BiLSTM_{s_i} are combined using a sigmoid gate that is learned during training:

$$\begin{aligned} \hat{\mathbf{h}}_i^{\text{bg}} &= \text{BiLSTM}_{\text{bg}}(\mathbf{h}_i) \\ \hat{\mathbf{h}}_i^{\text{s}_j} &= \text{BiLSTM}_{\text{s}_j}(\mathbf{h}_i), \forall j \in \text{speakers} \\ g^s &= \sigma(w_g) \\ \mathbf{h}'_i &= g^{s_i} * \hat{\mathbf{h}}_i^{\text{s}_i} + (1 - g^{s_i}) * \hat{\mathbf{h}}_i^{\text{bg}}, \\ \mathbf{h}'_i &= \text{MS-BiLSTM}(\mathbf{h}_i, s_i). \end{aligned}$$

Each speaker-level BiLSTM (BiLSTM_{s_i}) only receives gradients for that speaker’s utterance (u_i) thus focusing on role-specific context. The gate between $\hat{\mathbf{h}}_i^{\text{s}_i}$ and $\hat{\mathbf{h}}_i^{\text{bg}}$ controls the relative importance of the role-specific and general-context representation learned by speaker-level and background BiLSTMs respectively.

In this paper, we focus on classifying an utterance into one or more out of three categories, namely *symptoms*, *complaints*, and *medications*. However, these categories can be combined together to create a coarse-grained task of predicting if the utterances are *medically relevant*. We leverage this coarse-grained supervision to create a hierarchical model with a joint-learning loss.

Hierarchical Modeling: In this architecture, the extended representation (\mathbf{h}_i) and the corresponding speaker role (s_i) are first passed through a coarse-grained MS-BiLSTM and a fully-connected layer followed by softmax to be classified into one of the two categories $\{\textit{Medically Relevant}, \textit{Irrelevant}\}$. The representation $\mathbf{h}_i^{\text{coarse}}$ learned by this MS-BiLSTM would model the differences between medically relevant and irrelevant text which

can also benefit fine-grained classification. Hence, $\mathbf{h}_i^{\text{coarse}}$ is concatenated with \mathbf{h}_i and sent to the fine-grained MS-BiLSTM which focuses on the multi-label classification into the three categories discussed earlier:

$$\begin{aligned} \mathbf{h}_i^{\text{coarse}} &= \text{MS-BiLSTM}_{\text{coarse}}(\mathbf{h}_i, s_i), \\ \mathbf{h}_i'' &= [\mathbf{h}_i; \mathbf{h}_i^{\text{coarse}}], \\ \mathbf{h}_i^{\text{fine}} &= \text{MS-BiLSTM}_{\text{fine}}(\mathbf{h}_i'', s_i), \\ \mathbf{p}_{\text{coarse}} &= \text{softmax}(\mathbf{W}_c \mathbf{h}_i^{\text{coarse}} + \mathbf{b}), \\ \mathbf{p}_{\text{fine}} &= \sigma(\mathbf{W}_f \mathbf{h}_i^{\text{fine}} + \mathbf{b}). \end{aligned}$$

Both tasks are jointly optimized and the hyperparameter β controls the relative strength of the medical-relevance classification loss (L_{coarse}):

$$L = L_{\text{fine}} + \beta L_{\text{coarse}}.$$

Such a loss function could also be used in other utterance classification tasks where classes follow a hierarchical structure. For instance, in emotion classification (Poria et al., 2019), the fine-grained categories (e.g, happiness, anger, etc.) can be combined to create an *emotive* class, and a coarse-grained classifier could be used to learn features that differentiate between *emotive* and *neutral* utterances.

4 Experimental Setup

4.1 Corpus Description

Our data set comprises 6,862 annotated transcripts of real and de-identified doctor-patient conversations with an average of 225 utterances per conversation, primarily from the doctor and patient but occasionally including contributions from nurses, caregivers, and other attendees as well. The annotation guidelines were developed by a team of professional medical scribes and NLP experts. Annotators were trained to identify the medically-relevant utterances in a given conversation and assign one or more (out of 15 possible) tags to each utterance. Each of these tags represents a medical category like *symptom*, *previous medical history*, *diagnosis* etc. Most conversations contain some informal, social interactions with utterances that are irrelevant to the downstream clinical tasks.²

In this work, we leverage the labels to train MEDFILTER on the task of utterance classification and

²An example dialogue is included in Appendix (Sec. A.2).

focus on three categories, namely, *symptoms*, *complaints*, and *medications*, where *medications* include past/current medications taken by the patient and prescriptions given by the doctor.³ We choose the above-mentioned categories as they are found in every office visit, and most closely generalize to other domains like customer-service chats. However, our approach can be easily generalized for capturing other aspects such as *previous medical history*, *diagnosis*, and *assessments* as well. We set aside a random sample of 627 and 592 conversations for validation and testing respectively.

4.2 Baselines

Since sequence-labelling models haven't been applied to utterance classification in doctor-patient conversations previously, we compare our proposed method, MEDFILTER, against baseline methods that give SOTA results on utterance-level emotion recognition data sets. **HiGRU-sf** (Jiao et al., 2019b) is a hierarchical gated recurrent unit (HiGRU) framework with an utterance-level GRU and a conversation-level GRU. **BiF-AGRU** (Jiao et al., 2019a) denotes a two-level BiGRU fusion model with uni-directional AGRU for attentive context representation. **UniF-BiAGRU** is similar to BiF-AGRU, but uses a uni-directional GRU for contextual utterance representation and a bi-directional AGRU for attentive context. For implementation, we use the official code provided by the authors.^{4,5}

Evaluation Metric: We use the mean area under the PR curve (AUC), a widely used metric in multi-label classification setting (Riedel et al., 2013; Mintz et al., 2009), as our evaluation metric. It is also used for early stopping and hyperparameter tuning.⁶

5 Utterance Classification Results

MEDFILTER performs better than any of the baseline approaches in assigning utterances in doctor-patient conversations to medically relevant categories. Table 1 presents the AUC scores for different utterance-labeling models on our test set. Each result is the mean of 5 independent runs with different seeds.

A BERT-based classifier that passes the mean of token-level embeddings through an FC layer

³Additional statistics are included in Appendix (Sec. A.1).

⁴<https://github.com/wxjiao/HiGRUs>

⁵<https://github.com/wxjiao/AGHMN>

⁶Hyperparameters are included in Appendix (Table A4)

	Methods	AUC (x100)
Baselines	UniF-BiAGRU	40.9 (0.51)
	BiF-AGRU	40.9 (0.37)
	HiGRU-sf	43.1 (0.45)
BERT variants	BERT	33.5 (0.08)
	Clinical BioBERT-FT	36.1 (0.11)
	BERT-FT	36.2 (0.08)
With Context	BERT BiLSTM FT	44.5 (0.22)
	BERT-FT BiLSTM	45.8 (0.16)
Our Method	MEDFILTER	47.2 (0.26)

Table 1: Utterance classification results on the test-set (**Avg.** (**std. dev.**)). Results on valid-set are shown in the Appendix. The improvements are statistically significant ($p < 0.01$).

gives a low score of 33.5 AUC. When the BERT encoder is fine-tuned along with the classification layer (BERT-FT), the performance jumps to 36.2 underlining the benefits of fine-tuning BERT (Devlin et al., 2019). We also find that using Clinical BioBERT-FT (fine-tuned) does not beat BERT-FT. This is partly because the former is further pre-trained on MIMIC notes (Alsentzer et al., 2019) which are much more formal than medical conversations and thus the additional knowledge does not transfer well to our corpus.

Adding context to BERT-based models, using, e.g. BiLSTM, gives substantial boosts. End-to-end fine-tuned BERT BiLSTM (BERT BiLSTM FT) performs worse than BERT-FT BiLSTM that passes fine-tuned BERT embeddings through a BiLSTM as non-learnable features. MEDFILTER, which further includes contextual information, uses MS-BiLSTM in place of BiLSTM, and optimizes a hierarchical loss, significantly outperforms all baselines and obtains 1.4 absolute AUC points over BERT-FT BiLSTM (2nd best). It also surpasses emotion recognition SOTA methods like HiGRU-sf by 4.1 AUC points.

Ablation Results: To understand the importance of each module in MEDFILTER, we perform a cumulative ablation study (Figure 3). We find that removing individual modules results in notably reduced performance. The model that does not incorporate hierarchical modeling, shows a dip of 0.4 AUC points. This suggests that the information learned in the medical-relevance prediction layer aids the final classification task. Further, replacing MS-BiLSTM with a simple BiLSTM leads to a drop of an additional 0.6 AUC points, revealing the importance of modeling speaker-specific context. Without contextual information, we see a reduc-

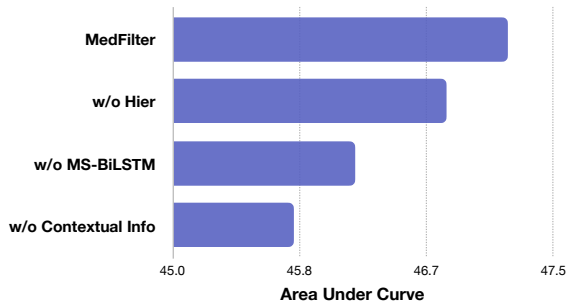


Figure 3: Cumulative Ablation Results

tion of 0.4 AUC points. This shows that features like speaker role, position, and semantic types are essential for our task.

6 Impact of Utterance Classification on Downstream Medical Extraction

The results in the previous section portray the effectiveness of MEDFILTER at sorting important utterances in clinical conversations into medically relevant categories. Such filtering, when included in the pipeline (for example, as a pre-processing step), can assist downstream medical processing methods to focus on utterances that contain information pertinent for their tasks (Figure 2), by improving the signal-to-noise ratio in the input. In this section, we evaluate whether the use of MEDFILTER to prune irrelevant utterances is advantageous for symptom, medication, and complaint extraction.

6.1 Task Setup

The extractor takes the conversation as input and outputs the discussed symptoms/medications/complaints within.

Conversation-level labels for all three extraction tasks are taken from a predefined set provided by the corpus annotators. For **symptoms**, they include 14 coarse-grained classes to represent different body systems (e.g., cardiovascular) and 178 fine-grained ones for the corresponding issues (e.g., palpitations). Given the small size of the training data, we use the coarse-grained body-systems for symptom extraction. We then manually curate a list of different symptoms corresponding to each body-system using UMLS and use their UMLS CUIs as labels.⁷

For **medications**, we manually link medication labels to their corresponding UMLS (Bodenreider, 2004) concepts and group them using hierarchies

⁷Refer to Table A10 for the final list of symptom labels.

Medication Extraction	Macro F1	Micro F1
QuickUMLS (All Text)	25.4	33.5
QuickUMLS (MR BERT-FT BiLSTM)	32.6	61.9
QuickUMLS (MU BERT-FT BiLSTM)	34.0	67.3
QuickUMLS (MR MEDFILTER)	34.2	62.8
QuickUMLS (MU MEDFILTER)	35.9	68.9

Table 2: Results(%) for Medication Extraction. **MR**=Medically Relevant (Symptom + Complaints + Medications) Utterances, **MU**=Medication Utterances.

from NCI Thesaurus (Sioutos et al., 2007).⁸ We pass each medication name through QuickUMLS to get a list of possible CUIs for the term in UMLS. We take the candidate CUI with a similarity of 1 and find its NCI hierarchy in the UMLS metathesaurus. The four topmost nodes in the hierarchy are extracted, which act as the pseudo-label for that CUI. In order to reduce the class-imbalance, some of these hierarchies are combined to form a coarser label. This reduces the number of labels to 31. Finally, *Others* label is added, which inhabits medicine names (in the test-set) that do not correspond to any of the previous 31 labels. This reduces the label count to 32 for medications.⁹

Complaints in our corpus range from follow-up visits to disease names to vaccine requests. Similar to medication extraction, we leverage SNOMED-CT hierarchies¹⁰ to constraint the tag list to 11, where the first 10 represent diseases of different body systems and *Others* encompasses complaints like follow-up, vaccine requests, medication refill requests, etc. (Table A11).

We use the same train/val/test split as defined for the utterance classification experiments in Section 4.1. The performance of the extraction pipeline is evaluated on Micro and Macro-F1 scores.

6.2 Extractor Details

All three extraction tasks are modeled as multi-label classification. We leverage a state-of-the-art medical entity-linking tool, QuickUMLS (Soldaini and Goharian, 2016)¹¹, that takes in a conversation and outputs UMLS CUIs corresponding to all identified candidate concepts. Concepts with a similarity measure of 1 are chosen as predictions. For symptom extraction, the predictions are compared against a manually created list of CUIs (presented

⁸<https://ncit.nci.nih.gov>

⁹Refer to Table A13 for the final list of medication labels.

¹⁰<https://www.nlm.nih.gov/healthit/snomedct/index.html>

¹¹<https://github.com/Georgetown-IR-Lab/QuickUMLS>

Symptom Extraction	Macro F1	Micro F1
QuickUMLS (All Text)	33.9	42.7
QuickUMLS (MR BERT-FT BiLSTM)	36.4	47.4
QuickUMLS (SU BERT-FT BiLSTM)	35.9	49.2
QuickUMLS (MR MEDFILTER)	35.2	47.4
QuickUMLS (SU MEDFILTER)	36.1	49.3

Table 3: Results(%) for Symptom Extraction. **MR**=Medically Relevant, **SU**=Symptom Utterances.

in Appendix Table A12) for symptoms associated with each of the 14 Body Systems. The presence of a symptom of body-system b is determined by the presence of the predicted CUIs in the target list for that body system. We compare the NCI and the SNOMED-CT hierarchies of the predicted concepts against the label hierarchies for medications and complaints, respectively. Concepts that do not fit into one of the specific categories are grouped under the label *Others*. In the next section, we report the results for the best performing filtering thresholds.¹²

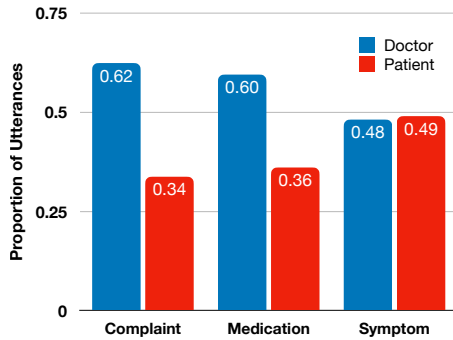
6.3 Results

We find that the performance of the baseline medication and symptom extractor QuickUMLS (All Text) is substantially boosted by filtering out irrelevant utterances (Tables 2 and 3). Pruning medically irrelevant utterances using MEDFILTER (MR MEDFILTER) improves Micro F1 by 29.3 and 4.7 points for medication and symptom extraction, respectively. If only the medication/symptomatic utterances (MU/SU) are input to the extractors, the results improve further.

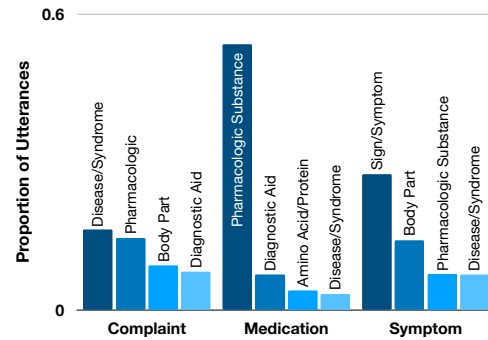
Results for complaint extraction are shown in Table 4. We find that the QuickUMLS extractor does not perform well on complaint extraction. However, consistent with the other two categories' trends, pruning irrelevant utterances before sending the conversation through the extractor improves performance. Micro-F1 score increases from 35.6 for All Text to 43.7 for CU MEDFILTER.

Pruning done using MEDFILTER seems to be more beneficial than BERT-FT BiLSTM (2nd best utterance classifier in Table 1) for medication and complaint extraction, however they perform equally well for symptom extraction. This suggests that the benefits from the inclusion of discourse structure, domain knowledge, and a hierarchical loss function, do not transfer well to symptom ex-

¹²Micro F1 vs filtering threshold graphs are presented in the Appendix (Figures A5 and A4).



(a) Proportion of utterances, in each medical category, spoken by different speaker roles.



(b) Most frequent semantic types in terms of the proportion of utterances they occur in.

Figure 4: Contextual Information: Different speaker roles contribute asymmetrically towards different medical topics/categories in the dialogue (Figure 4a). Furthermore, phrases with UMLS semantic types *Pharmacologic Substance*, *Sign/Symptom*, and *Disease/Syndrome* occur quite frequently in medical, symptom, and complaint utterances respectively (Figure 4b).

Complaint Extraction	Macro F1	Micro F1
QuickUMLS (All Text)	10.0	35.6
QuickUMLS (MR BERT-FT BiLSTM)	10.9	40.3
QuickUMLS (CU BERT-FT BiLSTM)	11.1	43.0
QuickUMLS (MR MEDFILTER)	11.1	40.7
QuickUMLS (CU MEDFILTER)	11.1	43.7

Table 4: Results(%) for Complaint Extraction (CE). **MR**=Medically Relevant, **CU**=Complaint Utterances.

traction. In Section 7, we investigate the kinds of utterance classification errors MEDFILTER makes, that need to be addressed to further improve the symptom extraction pipeline.

7 Discussion

Why does contextual information help? Ablation results (Figure 3) show that incorporating speaker role information and UMLS semantic-type information provides significant improvements in AUC scores for utterance classification. In Figure 4a, we plot the proportion of utterances from different medical categories against their speakers. While both parties contribute equally to symptom discussions, there is a clear asymmetry in the number of *medication* and *complaint* utterances spoken by the doctor and the patient, explaining the contribution of speaker role information in differentiating *medication/complaint* utterances from others.

We also plot the distribution of the four most frequent UMLS semantic types present in the utterances of different medical categories (Figure 4b). For medications, we find that UMLS entities with semantic type *Pharmacologic Substance* are present in more than 55% of the medication utterances indicating that its detection

is a knowledge-dependent task. Similarly, and supporting our hypothesis, *Disease/Syndrome* and *Sign/Symptom* are the most frequent semantic types in complaint and symptom utterances, respectively.

Error Analysis: In this section, we present a deeper analysis of some of the systematic knowledge-extraction errors made by MEDFILTER that limit its performance in recognizing medically-relevant utterances.

1. Informal Language: The model sometimes overlooks informal references to symptoms. For instance, utterances such as *PT: I feel something unusual in my leg* or *PT: My heart beats funny!* discuss musculoskeletal and cardiovascular symptoms but do not use medical terms to refer to them. These patterns seem to be more frequent in patient utterances, likely because they are less familiar with medical terminology. Off-the-shelf entity-linkers, like QuickUMLS (Soldaini and Goharian, 2016), do not transfer well to spoken medical conversations. They are unable to recognise the correct UMLS concepts (and semantic types) corresponding to the colloquial symptomatic phrases which reduces their effectiveness as features.¹³ For instance, for the utterance *PT: My heart is racing.*, QuickUMLS outputs 1A rather than 1B:

Input: PT: My heart is racing
(A): PT: My [heart]_{body_part} is racing. (1)
(B): PT: My [[heart]_{body_part} is racing]_{symptom}.

2. Physical Manifestations of Symptoms: Internal symptoms often manifest themselves physically as a digression from the natural ability to

¹³*Sign/Symptom* entities are identified in less than 30% of total symptomatic utterances (Figure 4b).

perform typical activities. For instance, when the patient says *I can't do anything after I'm back from office* or *I can only walk up one flight of stairs*, she might be implicitly mentioning a cardiovascular symptom. A sizeable subset of such examples includes usage of duration or frequency to convey the implicit deviation, like

Duration: PT: I only sleep for **4 hours**,

Frequency: PT: I go to the bathroom **10 times** at night,

Quantity: DR: I see you are up to **300 pounds** now.

8 Future Work

For a system to correctly classify the samples of the above two categories, it needs both to generalize to patient-generated language, and to have a semantic understanding of whether the description strays from normal. Incorporating data from online self-disclosure sites like medical subreddits and discussion forums (Basaldella and Collier, 2019) during training might prove beneficial for learning better representations for such vocabulary. Concept normalization data sets (Miftahutdinov and Tutubalina, 2019; Lee et al., 2017) could also be leveraged in this regard. Our approach of training the BERT encoder separately from the context encoder would allow MEDFILTER to learn from such non-dialogue resources.

Extraction tasks (Section 6) mostly evaluate the ability of MEDFILTER to recognize utterances that contain the most information about the name or type of the medication, symptom, or complaint. However, to quantify the context-level benefits of MEDFILTER, especially the speaker-specific context modeling (MS-BiLSTM), on downstream processing, we need to evaluate the system on problems like regimen extraction (Selvaraj and Konam, 2019) or symptom summarization (Liu et al., 2019b). Such tasks require utterance classification models to correctly identify utterances that discuss fine-grained details about the topic and would therefore evaluate a model's ability to solve multiple challenges like coreference resolution, speaker-specific context detection, thread identification, etc. Such an evaluation is a part of future work.

9 Conclusion

In this paper, we have proposed a novel text classification approach that specifically leverages insights into the organization of task-oriented conversations in order to improve performance at topic-based utterance classification over SOTA baseline

approaches. In particular, we have demonstrated that our utterance classification model, MEDFILTER, benefits from discourse information, domain knowledge, speaker-specific context modeling, and a hierarchical loss to reach a new state-of-the-art performance on a doctor-patient interactions dataset. We find that using topic-based utterance classification in general, and MEDFILTER in particular, as a pre-processing step before medical extraction tasks, significantly improves the extraction scores. We believe that the contributions made in this work would also generalize to other kinds of expert-lay dialogue like customer-service chats.

Acknowledgements

We thank the anonymous EMNLP reviewers for their insightful comments. We are also grateful to the members of the TELEDIA group at LTI, CMU for the invaluable feedback. This work was funded in part by Dow Chemical, University of Pittsburgh Medical Center, and Microsoft.

References

- Emily Alsentzer, John Murphy, William Boag, Wei-Hung Weng, Di Jin, Tristan Naumann, and Matthew McDermott. 2019. [Publicly available clinical BERT embeddings](#). In *Proceedings of the 2nd Clinical Natural Language Processing Workshop*, pages 72–78, Minneapolis, Minnesota, USA. Association for Computational Linguistics.
- Marco Basaldella and Nigel Collier. 2019. Bioredit: Word embeddings for user-generated biomedical nlp. In *Proceedings of the Tenth International Workshop on Health Text Mining and Information Analysis (LOUHI 2019)*, pages 34–38.
- Tom Bocklisch, Joey Faulkner, Nick Pawlowski, and Alan Nichol. 2017. Rasa: Open source language understanding and dialogue management. *arXiv preprint arXiv:1712.05181*.
- Olivier Bodenreider. 2004. The unified medical language system (umls): integrating biomedical terminology. *Nucleic acids research*, 32 Database issue:D267–70.
- Paweł Budzianowski, Tsung-Hsien Wen, Bo-Hsiang Tseng, Iñigo Casanueva, Ultes Stefan, Ramadan Osman, and Milica Gašić. 2018a. Multiwoz - a large-scale multi-domain wizard-of-oz dataset for task-oriented dialogue modelling. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Paweł Budzianowski, Tsung-Hsien Wen, Bo-Hsiang Tseng, Iñigo Casanueva, Stefan Ultes, Osman Ramadan, and Milica Gasic. 2018b. Multiwoz-a large-

- scale multi-domain wizard-of-oz dataset for task-oriented dialogue modelling. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 5016–5026.
- Cen Chen, Chilin Fu, Xu Hu, Xiaolu Zhang, Jun Zhou, Xiaolong Li, and Forrest Sheng Bao. 2019. Reinforcement learning for user intent prediction in customer service bots. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 1265–1268.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [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, Minneapolis, Minnesota. Association for Computational Linguistics.
- Nan Du, Kai Chen, Anjuli Kannan, Linh Tran, Yuhui Chen, and Izhak Shafran. 2019. [Extracting symptoms and their status from clinical conversations](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 915–925, Florence, Italy. Association for Computational Linguistics.
- Gregory Finley, Erik Edwards, Amanda Robinson, Michael Brenndoerfer, Najmeh Sadoughi, James Fone, Nico Axtmann, Mark Miller, and David Suendermann-Oeft. 2018. [An automated medical scribe for documenting clinical encounters](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Demonstrations*, pages 11–15, New Orleans, Louisiana. Association for Computational Linguistics.
- Wenxiang Jiao, Michael R Lyu, and Irwin King. 2019a. Real-time emotion recognition via attention gated hierarchical memory network. *arXiv preprint arXiv:1911.09075*.
- Wenxiang Jiao, Haiqin Yang, Irwin King, and Michael R. Lyu. 2019b. [HiGRU: Hierarchical gated recurrent units for utterance-level emotion recognition](#). 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 397–406, Minneapolis, Minnesota. Association for Computational Linguistics.
- Nazmul Kazi and Indika Kahanda. 2019. Automatically generating psychiatric case notes from digital transcripts of doctor-patient conversations. In *Proceedings of the 2nd Clinical Natural Language Processing Workshop*, pages 140–148.
- Young-Bum Kim, Sungjin Lee, and Karl Stratos. 2017. Onenet: Joint domain, intent, slot prediction for spoken language understanding. In *2017 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, pages 547–553. IEEE.
- Ahmet Baki Kocaballi, Enrico Coiera, Huong Ly Tong, Sarah J White, Juan C Quiroz, Fahimeh Rezazadegan, Simon Willcock, and Liliana Laranjo. 2019. A network model of activities in primary care consultations. *Journal of the American Medical Informatics Association*, 26(10):1074–1082.
- Ronilda C Lacson, Regina Barzilay, and William J Long. 2006. Automatic analysis of medical dialogue in the home hemodialysis domain: structure induction and summarization. *Journal of biomedical informatics*, 39(5):541–555.
- Kathy Lee, Sadid A Hasan, Oladimeji Farri, Alok Choudhary, and Ankit Agrawal. 2017. Medical concept normalization for online user-generated texts. In *2017 IEEE International Conference on Healthcare Informatics (ICHI)*, pages 462–469. IEEE.
- Chunyi Liu, Peng Wang, Jiang Xu, Zang Li, and Jieping Ye. 2019a. Automatic dialogue summary generation for customer service. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1957–1965.
- Yang Liu, Kun Han, Zhao Tan, and Yun Lei. 2017. [Using context information for dialog act classification in DNN framework](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2170–2178, Copenhagen, Denmark. Association for Computational Linguistics.
- Zhengyuan Liu, Angela Ng, Sheldon Lee Shao Guang, Ai Ti Aw, and Nancy F. Chen. 2019b. [Topic-aware pointer-generator networks for summarizing spoken conversations](#). In *IEEE Automatic Speech Recognition and Understanding Workshop, ASRU 2019, Singapore, December 14-18, 2019*, pages 814–821. IEEE.
- Zulfat Miftahutdinov and Elena Tutubalina. 2019. Deep neural models for medical concept normalization in user-generated texts. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: Student Research Workshop*, pages 393–399.
- Mike Mintz, Steven Bills, Rion Snow, and Dan Jurafsky. 2009. Distant supervision for relation extraction without labeled data. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP: Volume 2-Volume 2*, pages 1003–1011. Association for Computational Linguistics.
- Jihyun Park, Dimitrios Kotzias, Patty Kuo, Robert L Logan IV, Kritzia Merced, Sameer Singh, Michael Tanana, Efi Karra Taniskidou, Jennifer Elston Lafata,

- David C Atkins, et al. 2019. Detecting conversation topics in primary care office visits from transcripts of patient-provider interactions. *Journal of the American Medical Informatics Association*, 26(12):1493–1504.
- Soujanya Poria, Devamanyu Hazarika, Navonil Majumder, Gautam Naik, Erik Cambria, and Rada Mihalcea. 2019. MELD: A multimodal multi-party dataset for emotion recognition in conversations. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 527–536, Florence, Italy. Association for Computational Linguistics.
- Chen Qu, Liu Yang, W Bruce Croft, Yongfeng Zhang, Johanne R Trippas, and Minghui Qiu. 2019. User intent prediction in information-seeking conversations. In *Proceedings of the 2019 Conference on Human Information Interaction and Retrieval*, pages 25–33.
- Juan C. Quiroz, Liliana Laranjo, Ahmet Baki Kocaballi, Shlomo Berkovsky, Dana Rezazadegan, and Enrico Coiera. 2019. Challenges of developing a digital scribe to reduce clinical documentation burden. *npj Digital Medicine*, 2(1):114.
- Vipul Raheja and Joel Tetreault. 2019. Dialogue Act Classification with Context-Aware Self-Attention. 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 3727–3733, Minneapolis, Minnesota. Association for Computational Linguistics.
- Nils Reimers and Iryna Gurevych. 2019. Sentencebert: Sentence embeddings using siamese bert-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Sebastian Riedel, Limin Yao, Andrew McCallum, and Benjamin M Marlin. 2013. Relation extraction with matrix factorization and universal schemas. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 74–84.
- Abigail See, Peter J Liu, and Christopher D Manning. 2017. Get to the point: Summarization with pointer-generator networks. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1073–1083.
- Sai P Selvaraj and Sandeep Konam. 2019. Medication regimen extraction from medical conversations. *arXiv*, pages arXiv–1912.
- Elizabeth Shriberg, Rebecca Bates, Paul Taylor, Andreas Stolcke, Daniel Jurafsky, Klaus Ries, Noah Coccaro, Rachel Martin, Marie Meteer, and Carol Van Ess-Dykema. 1998. Can prosody aid the automatic classification of dialog acts in conversational speech? *Language and Speech*, 41(3–4):439–487.
- Elizabeth Shriberg, Raj Dhillon, Sonali Bhagat, Jeremy Ang, and Hannah Carvey. 2004. The ICSI meeting recorder dialog act (MRDA) corpus. In *Proceedings of the 5th SIGdial Workshop on Discourse and Dialogue at HLT-NAACL 2004*, pages 97–100, Cambridge, Massachusetts, USA. Association for Computational Linguistics.
- Nicholas Sioutos, Sherri de Coronado, Margaret W. Haber, Frank W. Hartel, Wen-Ling Shaiu, and Lawrence W. Wright. 2007. Nci thesaurus: A semantic model integrating cancer-related clinical and molecular information. *J. of Biomedical Informatics*, 40(1):30–43.
- Luca Soldaini and Nazli Goharian. 2016. Quickumls: a fast, unsupervised approach for medical concept extraction. In *MedIR workshop, sigir*, pages 1–4.
- 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 I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 5998–6008. Curran Associates, Inc.
- Howard Waitzkin. 1989. A critical theory of medical discourse: Ideology, social control, and the processing of social context in medical encounters. *Journal of Health and Social Behavior*, pages 220–239.
- Chenwei Zhang, Yaliang Li, Nan Du, Wei Fan, and S Yu Philip. 2019. Joint slot filling and intent detection via capsule neural networks. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5259–5267.

Appendix

A Dataset Details

A.1 Dataset Statistics

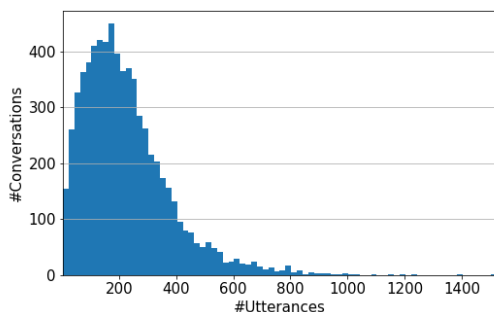


Figure A1: Distribution of the number of utterances in each conversation for the entire dataset.

The de-identified doctor-patient dialogue corpus used in this work was made available by University of Pittsburgh Medical Center (UPMC) and Abridge AI Inc.. Most of the conversations in this corpus are follow-up encounters between cardiovascular/general medicine doctors and patients. Figure A1 shows the distribution of the number of utterances in each conversation. The number ranges from as low as 3 to as high as 1521 with a mean of 225. The proportion of medically relevant utterances in a conversation is quite low (Table A1). As shown, utterances that belong to the three categories combined make up less than 10% of the conversation portraying the amount of noise present in doctor-patient conversations with regards to further medical processing.

Category	#MR-Utt	#MR-Utt/#Utt (%)
Complaints	6.17 (3.40)	4.34 (6.06)
Symptoms	3.56 (4.00)	1.98 (2.37)
Medications	4.79 (3.70)	3.10 (5.49)

Table A1: Avg. (std. dev.) medically relevant utterances (MR-Utt) in each medically relevant category.

In Table A2, we show the average position in the doctor-patient conversation where the speakers start discussing different medical topics. Several of the the encounters are follow-up discussions about a pre-existing complaint. Therefore, patient’s current condition with respect to the complaint is often discussed earlier in the conversation. This is generally followed by a discussion about different body systems and associated symptoms that may

be bothering the patient, which allows the doctor to prescribe suitable medications.

Category	Relative Position
Complaints	0.133 (0.043)
Symptoms	0.321 (0.057)
Medications	0.524 (0.069)

Table A2: Avg. (std. dev.) relative position in the conversation where speakers start discussing different medical topics.

The above-mentioned flow is merely an ideal depiction of the logical path that could be followed in the dialogue. However, real conversations in the corpus contain multiple topic-switches. For example, discussion of a symptom could be followed by medication which could then lead into a discourse about another symptom and so on.

Utterance	Labels
1 Check if conversation can be added	
1 DR: Good Morning.	
2 PT: Good Morning.	
3 DR: I’m here with, [PATIENT NAME].	
4 DR: Last time I saw you, you were getting pains in your left leg. Is it still the case?	C,S
5 PT: Yes, I do.	S
6 DR: Okay, and generally, what are you doing when you get the pains?	
7 PT: Um, usually just a heating pad or, you know, ice.	
8 DR: Right, but what causes the pains, is what I was getting at?	S
9 PT: Uh, I think just the strain of, like, walking, or, or exercise.	S
10 DR: All right.	
11 DR: I think I am going to ask you to try some Baclofen.	M
12 DR: This is a patch you put on the foot when it’s bothering you.	M
13 DR: Try one patch.	M
14 DR: It’ll last up to 6 hours.	
15 PT: Okay.	
16 DR: If you like it, let me know.	
17 DR: We’ll get you a prescription.	
18 PT: Okay.	
19 DR: The difference is, you can put this exactly where you need it on the foot and since it’s going through the skin, it’s not rough on your stomach like, let’s say, Ibuprofen or Aspirin or any of the over the counter stuff would be.	
20 PT: Okay.	

Table A3: A constructed example conversation (S = Symptoms, C = Complaints, and M = Medications). Because conversations in the corpus cannot be published or distributed without agreement, the example here is based on a corpus conversation but with the details changed.

A.2 Example Conversation

An example conversation (details modified) from our corpus is shown in Table A3. Utterances 4,5,8,9

in the conversation discuss a symptom, with the Patient’s reply (A3:5) acting as an important information about confirmation of the presence of that symptom. Furthermore, A3:8 and A3:9 provide additional details about the physical activities that cause the symptom. Although, the symptom name is discussed only in A3:4, information presented in the other utterances plays an important role in the clinical note. Similarly, utterances A3:11,12,13 discuss the medication *Baclofen*. Doctor prescribes the medication in A3:11. She provides further information like frequency of usage in A3:12, and dosage in A3:13, all of which is extremely important for regimen extraction. A3:19 contains names of two medications however it is not a medication utterance. This is case because the utterance does not discuss any medication that the patient is currently taking or being prescribed. The doctor is merely comparing the benefits of her prescribed medication against two popular pain pills.

A.3 Symptom and Medication Extraction Labels

In addition to identifying the type of each utterance, corpus annotators also provide a class label to the symptoms and medications from a predefined set. For symptoms, guidelines include 178 classes of the form $\langle \text{Body System} \rangle: \langle \text{symptom} \rangle$ (e.g. Cardiovascular: Palpitations). Given the small size of the training data, instead of predicting given symptom classes, we predict the body system with which a symptom is associated (Table A10). Table A12 contains the list of target UMLS CUIs for each body system that are used as labels for Symptom Extraction. Please note that the list is manually curated and therefore is not exhaustive. For medications, we manually link each medication label in our training-set to its corresponding UMLS (Bodenreider, 2004) concept and group them using hierarchies from NCI Thesaurus (Sioutos et al., 2007) (Table A13).

B Hyperparameters

All our experiments are performed on a single Nvidia GeForce GTX 1080 Ti GPU. For MEDFILTER and other BERT-based baselines, we divide the conversations into windows of 128 utterances to ensure fair comparison against BERT-BiLSTM FT, which cannot process more than 128 utterances at a time due to GPU constraints. Other hyperparameters are presented in Table A4. We perform man-

Hyper-parameter	Search Range	Best
GRU hidden size in baselines	[100, 300, 512]	300
Max. utterance length	[64]	64
BERT embedding size	[768]	768
#Speakers	[3]	3
Speaker embedding size	[3, 4, 8, 16]	8
Number of bins (k)	[4]	4
Position embedding size	[4]	4
Semantic Type embedding size	[8, 16]	8
BiLSTM hidden size	[512, 1024, 2048]	1024
Weight of L_{coarse} (β)	[0, 0.25, 0.5, 1, 5]	1
Learning-rate	[0.0005, 0.001, 0.01]	0.0005
Batch-size	[8, 16, 32]	16

Table A4: Hyper-parameters. We search over the entire Cartesian product of the different hyper-parameters mentioned here. Best values are chosen using mean AUC of PR curve metric.

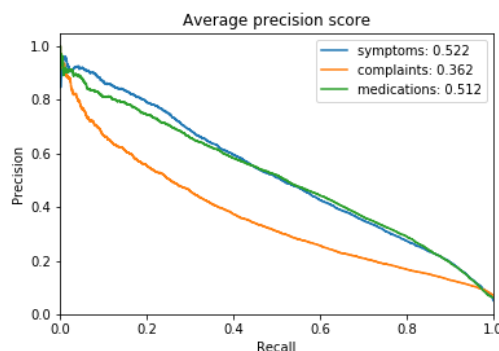


Figure A2: Category-wise PR curves for BERT-FT BiLSTM

ual tuning on the entire range of hyper-parameters. AUC under the PR curve metric was chosen to select the best configuration. Results were not very sensitive to different non-zero values of β .

C Utterance Classification

C.1 PR Curves

Figures A2 and A3 show the precision-recall curves for each category separately. MEDFILTER improves utterance classification for all three categories. For symptom classification, the AUC scores improve from 52.2 to 53.5. However, symptom extraction results (Section 6 in the main paper) suggest that most of this improvement is on identifying utterances that discuss fine-grained details about symptom discussion and not on recognizing the utterance that contains the actual symptom name.

C.2 Performance on Validation Set

Table A5 shows the performance of different utterance classification models on the validation set.

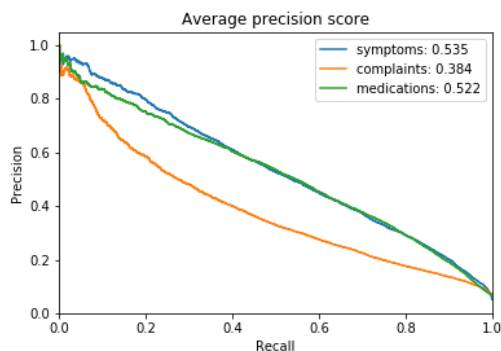


Figure A3: Category-wise PR curves for MEDFILTER

Similar to the trend shown on test-set, MEDFILTER beats all of the baselines reaching a score of 50.5 AUC points.

Methods	Val AUC	#Param	Time (hrs)
UniF-BiAGRU	42.7	1.3M	1
BiF-AGRU	42.9	1.3M	1
HiGRU-sf	45.0	2.6M	0.45
BERT	35.9	110M	-
Clinical BioBERT-FT	38.5	110M	10
BERT-FT	38.5	110M	10
BERT BiLSTM FT	47.9	125M	12
BERT-FT BiLSTM	49.6	125M	10 + 0.1
MEDFILTER	50.5	169M	10 + 1

Table A5: Results on val-set and the number of trainable parameters corresponding to each utterance classification model. The time taken by models that use BERT-FT is shown as a sum of two numbers as fine-tuning BERT is only done once, which is then used for both BERT-FT BiLSTM and MEDFILTER.

D Downstream Medical Extraction

D.1 Micro-F1 vs Threshold

Figure A4 and A5 show how the performance of medication (ME) and symptom extraction (SE) varies against different utterance topic prediction probability thresholds. We plot the results for BERT-FT BiLSTM and MEDFILTER for brevity. Micro F1 scores for ME increase monotonically when the threshold is increased from 0 to ~ 0.75 (Figure A4). This suggests that QuickUMLS medication extractor has low precision that is substantially improved when we prune irrelevant utterances. However, the graph for SE (Figure A5) shows that the extractor’s performance is dominated by its recall. Pruning helps with improving the precision however does not help with the low recall. This explains the lower gains as compared to

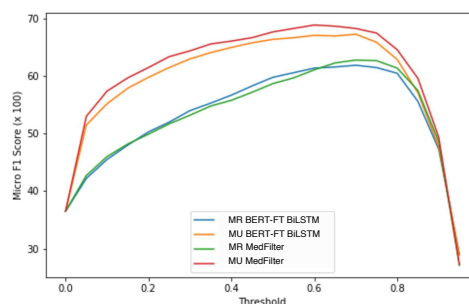


Figure A4: Medication Extraction: Micro-F1 vs Threshold

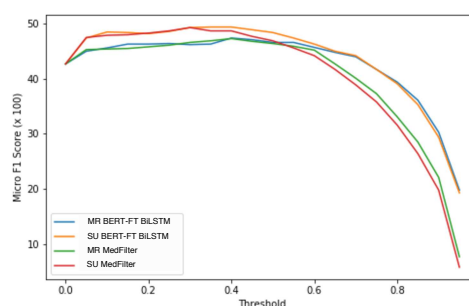


Figure A5: Symptom Extraction: Micro-F1 vs Threshold

ME when using topic-based utterance classification in the SE pipeline (Table 3 in the main paper).

D.2 Oracle Results

Table A6 contains results for medication extraction when medically relevant (MR) or medication (MU) utterances are chosen using an oracle (MR/MU Oracle). Similarly, oracle results for symptom and complaint extraction are shown in Tables A7 and A8, respectively.

We find that there is still a substantial room for improvement in the symptom extraction pipeline. By just improving the topic-based utterance classifier, one can observe a potential jump of 5 Micro-F1 points in symptom extraction. However, we do not observe this trend for medication extraction where the topic-classification done by MEDFILTER performs much better than the Oracle.

D.2.1 Why does MEDFILTER perform better than Oracle on Medication Extraction?

Extraction experiments (like medication extraction or symptom extraction) evaluate the performance at the conversation-level. So, where the medication

Medication Extraction	Macro F1	Micro F1
QuickUMLS (All Text)	25.4	33.5
QuickUMLS (MR Oracle)	30.0	41.3
QuickUMLS (MU Oracle)	37.6	58.3
QuickUMLS (MR MEDFILTER)	34.2	62.8
QuickUMLS (MU MEDFILTER)	35.9	68.9

Table A6: Results(%) for Medication Extraction (ME). **MR**=Medically Relevant (Symptom + Complaints + Medications) Utterances, **MU**=Medication Utterances.

Symptom Extraction	Macro F1	Micro F1
QuickUMLS (All Text)	33.9	42.7
QuickUMLS (MR Oracle)	36.9	47.2
QuickUMLS (SU Oracle)	41.9	54.5
QuickUMLS (MR MEDFILTER)	35.2	47.4
QuickUMLS (SU MEDFILTER)	36.1	49.3

Table A7: Results(%) for Symptom Extraction (SE). **MR**=Medically Relevant, **SU**=Symptom Utterances.

name gets extracted from within the conversation is irrelevant to the task.

Oracle picks utterances that would be sufficient for a human to identify the medications discussed in the dialogue. However, they might not be adequate for an automatic string-matching based extractor like QuickUMLS. Since QuickUMLS uses non-contextual surface-level features to identify medication names, it would look for phrases (in the input given to it) that match the surface requirements. So, it is possible for the Oracle utterances not to contain the proper surface-level forms that QuickUMLS could leverage for extracting medications. Furthermore, the utterances categorized as medication utterances by MedFilter on the other hand, even though incorrect, might contain the medication names in the form QuickUMLS expects, thus improving the score over the Oracle. One should note, however, that a perfect downstream extractor would not suffer from these side-effects.

Complaint Extraction	Macro F1	Micro F1
QuickUMLS (All Text)	10.0	35.6
QuickUMLS (MR Oracle)	10.6	38.8
QuickUMLS (SU Oracle)	13.4	44.3
QuickUMLS (MR MEDFILTER)	11.1	40.7
QuickUMLS (CU MEDFILTER)	11.1	43.7

Table A8: Results(%) for Complaint Extraction (CE). **MR**=Medically Relevant, **CU**=Complaint Utterances.

D.3 Supervised Extractor

For symptom extraction (SE), we also show the benefits of using topic-based utterance classification on a supervised-classification based SE approach that leverages a BiLSTM with attention (BiLSTM-Attn) for the problem of predicting the symptoms present in a conversation.

D.3.1 BiLSTM-Attn

Each utterance in the conversation is passed through the embedding layer and a BiLSTM layer to obtain a contextualized representation.

$$h_i = \text{BiLSTM}(e(s_i), h_{i-1})$$

$$\mathbf{H}_i = \{h_1, h_2, \dots, h_n\}$$

where $e(\cdot)$ is the embedding function. The final state of the BiLSTM is re-weighted using attention calculated as shown in Equation A1.

$$h_{final} = [\vec{h}_0; \vec{h}_{n-1}]$$

$$\mathbf{S} = \mathbf{H}_i h_{final}$$

$$\mathbf{A} = \text{softmax}(\mathbf{S})$$

$$h'_{final} = \mathbf{H}_i^T \mathbf{A}$$
(2)

This allows our model to pay attention to important utterances in the conversation to extract symptom information. We pass h'_{final} through a linear classifier and a sigmoid layer to get logits for each possible symptom label (Table A10).

D.3.2 Experimental Setup

Similar to the QuickUMLS based extractor, we use Micro and Macro F1 scores to evaluate the performance of the supervised extraction pipeline. BiLSTM-Attn (All Text) model takes in the entire conversation as input, whereas the other variants are given only a subset of utterances. MR Oracle/MEDFILTER models are trained on the medically relevant utterances as output by the oracle. Similarly, SU Oracle/MEDFILTER models are trained on the Oracle symptom utterances in each conversation in the training-set. Therefore, topic-based classification is used as a pre-processing step in the pipeline.

D.3.3 Results

We present the results for symptom extraction (SE) using a BiLSTM-Attn model in Table A9. We find that using topic-based utterance classification to remove irrelevant utterances before passing the conversation through the BiLSTM-Attn improves

the SE performance of the pipeline (4 point jump in Micro F1). The results are further improved when the Oracle symptom utterances (SU Oracle) are input to the BiLSTM-Attn.

Symptom Extraction	Macro F1	Micro F1
BiLSTM-Attn (All Text)	28.1	57.7
BiLSTM-Attn (MR Oracle)	29.3	59.4
BiLSTM-Attn (SU Oracle)	31.5	66.6
BiLSTM-Attn (MR MEDFILTER)	28.7	58.4
BiLSTM-Attn (SU MEDFILTER)	29.8	61.7

Table A9: BiLSTM-Attn results(%) for Symptom Extraction (SE). **MR**=Medically Relevant, **SU**=Symptom Utterances.

Symptom Labels
Cardiovascular
General
Musculoskeletal
Respiratory
Endocrine
Ear Nose Throat
Eyes
Gastrointestinal
Genital
Head
Neurological
Psychiatric
Skin
Urinary

Table A10: Symptom Extraction Labels (Body Systems)

Chief Complaint Labels
General
Disorder of hematopoietic structure
Disorder of integument, immune system, endocrine
Disorder of musculoskeletal system
Disorder of digestive system
Disorder of the genitourinary system
Disorder of respiratory system
Disorder of breast
Disorder of nervous system
Disorder of cardiovascular system
Others

Table A11: Complaint labels in the dataset. The label names represent the children of the SNOMED-CT hierarchy: SNOMED CT Concept/Clinical Finding/Finding by site/ Disorder by body site.

Label	Target CUI List
General	C0036572, C0015672, C0424653, C0015967, C3714552
Skin Head	C0234233, C0178298, C0015230, C0151908, C0362076, C0042571, C0018681, C0220870, C0012833
Eyes	C0235267, C0015397, C0007222, C1705500, C0012634, C2107992, C0017178, C0085635, C0848332, C0521707, C0152227, C0151827, C0017601, C0015230
Ent	C0027424, C2926602, C0699744, C0030193, C0031350, C0013456, C0018621, C0009443, C0851354, C0018021, C0017672, C0024117, C0036572, C2012701, C0041912, C0042571, C0019825, C0242429, C0427008, C0497156, C1135208, C0151908
Genital	C0567522, C3539891, C3539893, C0149741, C0020624, C2127567, C3539020, C0030193, C0424849, C3539896, C0567523, C0577573, C0007947, C0282005, C0017412, C2129032, C0023533, C4029890, C3539892, C0850758, C0438692, C0567526, C0039591, C0036918, C0036917, C3539890, C0232861, C1657982, C0036916, C3539022, C0877338, C1658964, C1868932, C0423610, C4552766, C0024902, C0234233, C3539023, C0030794, C2075679, C0156398, C1391387, C2030274, C0567519, C0017411, C2032395, C2126231, C0236078, C3539889, C3539895, C0849787, C2032396, C0019693
Respiratory	C0857427, C0013404, C0149514, C1396850, C0041312, C0206526, C0019079, C0006277, C0041296, C0010200, C0024115, C0034067, C0030524, C0152874, C0004096, C0041322, C0043144, C0275904
Cardiovascular	C0013404, C0795691, C0235710, C0008031, C0035436, C0002871, C0020538, C0018799, C0497234
Gastrointestinal	C0011991, C0019196, C0019112, C2032722, C0030193, C0178298, C0854495, C4748517, C0019158, C0018834, C0237938, C0854496, C0849766, C0239549, C0149696, C3553270, C0014724, C0814152, C0000737, C1321898, C0596601, C0085293, C2697368, C0016977, C0949135, C0011226, C0018932, C0017178, C0019159, C0027497, C0687713, C0341286, C0009806, C4728126, C1258215, C0920703, C0019163
Urinary	C0392525, C0262655, C0018965, C0239725, C0042029, C0455880, C4087409, C0152032, C0022650, C0021167, C0030193, C0558489
Musculoskeletal	C0030193, C0040822, C0858888, C0026857, C1405877, C0158026, C0003864, C0003123, C0030554, C0085593, C0427086, C0426579, C3714552, C0231528, C0036572, C0003873, C0028643, C0424653, C0003862, C0423572, C0427008, C0007859, C0541786, C0522057, C0018099, C2242996, C0015967, C1328469, C0263776, C0015230
Psychiatric	C0542476, C1579931, C0235108, C0497307, C0027769
Neurologic	C0036572, C0233407, C0042571, C0018681, C0039070, C1660797, C0312422, C0012833, C1135208
Endocrine	C0024117, C0020175, C0085602, C0041912, C0848390, C0009443, C0020615, C0221500

Table A12: CUI Target List for Symptom Extraction

Medication Labels
DFCBM/Chemical Modifier/Toxin
DFCBM/Dietary Supplement
DFCBM/Drug or Chemical by Structure
DFCBM/Food or Food Product
DFCBM/Industrial Aid
DFCBM/Natural Product
DFCBM/Pharmacologic Substance/Adjuvant
DFCBM/Pharmacologic Substance/AA Blood or Body Fluid
DFCBM/Pharmacologic Substance/AA Cardiovascular System
DFCBM/Pharmacologic Substance/AA Digestive System or Metabolism
DFCBM/Pharmacologic Substance/AA Integumentary System
DFCBM/Pharmacologic Substance/AA Musculoskeletal System
DFCBM/Pharmacologic Substance/AA Nervous System
DFCBM/Pharmacologic Substance/AA Organs of Special Senses
DFCBM/Pharmacologic Substance/AA Respiratory System
DFCBM/Pharmacologic Substance/Anti-Infective Agent
DFCBM/Pharmacologic Substance/Antineoplastic Agent
DFCBM/Pharmacologic Substance/Biological Agent
DFCBM/Pharmacologic Substance/Cation Channel Blocker
DFCBM/Pharmacologic Substance/Chemopreventive Agent
DFCBM/Pharmacologic Substance/Combination Medication
DFCBM/Pharmacologic Substance/Endothelin Receptor Antagonist
DFCBM/Pharmacologic Substance/Enzyme Inhibitor
DFCBM/Pharmacologic Substance/Hormone Therapy Agent
DFCBM/Pharmacologic Substance/Immunotherapeutic Agent
DFCBM/Pharmacologic Substance/Prostaglandin Analogue
DFCBM/Pharmacologic Substance/Protective Agent
DFCBM/Pharmacologic Substance/Protein Synthesis Inhibitor
DFCBM/Physiology-Regulatory Factor
Activity/Clinical or Research Activity/Intervention or Procedure
Manufactured Object/Diagnostic, Therapeutic, or Research Equipment
Others

Table A13: Medication Extraction Labels (DFCBM = Drug, Food, Chemical or Biomedical Material, AA = Agent Affecting).