

# Conversational Document Prediction to Assist Customer Care Agents

Jatin Ganhotra,<sup>1</sup> Haggai Roitman,<sup>2\*</sup> Doron Cohen,<sup>1</sup> Nathaniel Mills,<sup>1</sup> Chulaka Gunasekara,<sup>1</sup> Yosi Mass,<sup>1</sup> Sachindra Joshi,<sup>1</sup> Luis Lastras,<sup>1</sup> David Konopnicki<sup>1</sup>

<sup>1</sup>IBM Research, <sup>2</sup>eBay Research

{*jatinganhotra, wnm3, lastrasl*}@us.ibm.com, *hroitman@ebay.com, chulaka.gunasekara@ibm.com,*  
{*doronc, yosimass, davidko*}@il.ibm.com, *jsachind@in.ibm.com*

## Abstract

A frequent pattern in customer care conversations is the agents responding with appropriate webpage URLs that address users' needs. We study the task of predicting the documents that customer care agents can use to facilitate users' needs. We also introduce a new public dataset<sup>1</sup> which supports the aforementioned problem. Using this dataset and two others, we investigate state-of-the-art deep learning (DL) and information retrieval (IR) models for the task. We also analyze the practicality of such systems in terms of inference time complexity. Our results show that a hybrid IR+DL approach provides the best of both worlds.

## 1 Introduction

Customer care (CC) agents play a crucial role as an organization's main representatives to the public. Our work is motivated by the observation that, in many conversations between CC agents and users, the former tend to provide links to documents that may help resolve user issues. This is a prevalent pattern that is found in around 5-9% of all customer care conversations in multiple domains that we have reviewed. To identify such documents, the agents manually extract the keywords from the conversation and search over their customer service knowledge base (Habibi and Popescu-Belis, 2015; Ferreira et al., 2019). Table 1 shows a conversation where the agent provides a URL<sup>2</sup> to the user.

Although responding with URLs is a common pattern, automating this process to aid the agents remains underexplored in the literature. This task of **Conversational Document Prediction (CDP)** can be viewed as a conversational search problem, where the entire conversation context or a

\* Work done when author was at IBM Research

<sup>1</sup>The Twitter dataset is available at:  
<https://github.com/IBM/twitter-customer-care-document-prediction>

<sup>2</sup>The terms URL and document are used interchangeably.

## Dialogue

U: My virtual keyboard seems to float in the screen. Not sure how to undo what I just did. Can you help me please?  
A: We're happy to help. To start, let us know which device you're working with, and the OS version installed on it.  
U: It is an iPad  
A: Ok, to check version, tap Settings; General ; About.  
U: It's iPad 4, 11 inch - model A1934  
A: Thank you. This article can help with how to merge a split keyboard and move the keyboard for an iPad: <https://support.apple.com/en-us/HT207521>. Let me know if this helps

Table 1: Sample dialog from Twitter where the Agent (A) utterance includes a URL to the User (U) query.

subset of it could be used as the query for retrieving matching documents. Compared to ad-hoc retrieval settings, using a conversational interface, the agent/system can ask clarification questions and interactively modify the search results as the conversation progresses (Zhang et al., 2018; Alian-nejadi et al., 2019).

The CDP task has been primarily addressed so far using "traditional" information retrieval (IR) techniques. Habibi and Popescu-Belis (2015) proposed a document recommender system by extracting keywords from a conversation using topic modeling techniques. Ferreira et al. (2019) have used a similar keyword extraction framework and reported their results on a proprietary dataset.

Many aspects of IR systems have undergone a revolution with the advent of powerful Deep Learning (DL) techniques in recent years (Mitra et al., 2018; Yang et al., 2019). Yet this superior performance comes with high demand in computational resources as well as longer inference times, which hinders their application in real-world IR systems. Thus, the attention has been focused on techniques that reduce the computation complexity at the runtime without hindering the performance (Reimers and Gurevych, 2019; Lu et al., 2020).

Metric	Twitter	Mac-Support	Telco-Support
# of dialogs	13500	83436	1634562
# of dialogs with an URL in agent utterance	13500	10470	99733
# of dialogs with an in-domain URL in agent utterance	13500	7531	84126
# of dialogs with an in-domain valid URL in agent utterance	11025	4611	48565
Valid/Total # of unique URLs	2004 / 3585	522 / 1130	318 / 1203
Avg./Max. # of turns per dialogue	1.2 / 8.0	8.2 / 80.5	9.9 / 75.0
Avg./Max. dialog length (in tokens)	40.5 / 503	145.8 / 1481	334.5 / 5390
Avg./Max. URL content length (in tokens)	537.4 / 20492	311.6 / 7765	877.2 / 7889
vocabulary size	11646	10454	34099
train/dev/test split size	10000 / 525 / 500	3677 / 467 / 467	38850 / 4857 / 4858

Table 2: The overall statistics of the three datasets.

In this work, we formulate the CDP task to support CC agents. We further release a new public dataset which enables research on the aforementioned task and investigate the performance of state-of-the-art DL and IR models side-by-side on a number of datasets. We also analyze the runtime complexity of such systems, and propose a hybrid solution which is applicable in real-life systems.

## 2 Data

We explore the CDP task using three datasets which contain human-to-human conversations between users and CC agents. Two of these datasets are internal: one from an internal customer support service on Mac devices (*Mac-Support*) and another from an external client in the telecommunication domain (*Telco-Support*). We also release a new *Twitter* dataset, containing conversations between users and CC agents in 25 organizations on the Twitter platform<sup>3</sup>. We summarize the statistics of the three datasets in Table 2.

For our internal datasets, we filter out dialogs where: a) the agent doesn't provide a URL to the user, b) the URL is not in-domain (e.g. Google searches, Microsoft forums, etc.), and focus on URLs from internal customer service knowledge base, and c) the URL is either no longer valid or has no content (e.g., login page). For Twitter dataset, we used the user\_timeline API to collect the tweets from agents containing in-domain URLs. The dialogs were constructed starting from these tweets and identifying the previous user and agent tweets to these tweets. If a dialog contains multiple URLs, we only use the dialog till the first agent utterance containing a URL. The details for document content extraction are in Appendix.

From Table 2, we observe that, around 5-9% dialogs include a URL document provided by the

<sup>3</sup>The Twitter dataset is available at: <https://github.com/IBM/twitter-customer-care-document-prediction>

agent. We also note that the website content for organizations gets updated frequently as many URLs return 404 errors. The average number of turns in a dialog and dialog length (in tokens) is much smaller for *Twitter* in comparison to the *Mac-Support* and *Telco-Support* datasets. Our experiments results in Section 4, particularly BM25 and IRC in Table 4, demonstrate the importance of dialog context for the CDP task, even when that context is not very rich, as is the case for the short dialogs of Twitter.

## 3 Approaches

We now formally introduce the CDP task and notations below. We then describe two alternative approaches (IR and DL) and their hybrid that we evaluate for this task.

### 3.1 Task Definition

We regard the CDP task as a dialogue-based document classification task, similar to next utterance classification (Lowe et al., 2015). This is achieved by processing the data as described in Section 3.3, without requiring any human labels.

Formally, let  $d = \{s_1:t_1, s_2:t_2, \dots, s_n:t_n\}$  denote an  $n$ -turn dialog, where  $s_i$  represents the speaker (*user* -  $U$  or *agent* -  $A$ ), and  $t_i$  represents the  $i^{th}$  utterance. The dialog history is concatenated together to form dialog context of length  $m$ , represented as  $d = (d_1, d_2, \dots, d_i, \dots, d_m)$ , where  $d_i$  is the  $i^{th}$  word in context. Let  $Y$  denote the set of all documents which can be recommended to the user. Similar to dialogs, each document  $y \in Y$  is represented as  $y = (y_1, y_2, \dots, y_j, \dots, y_n)$ , where  $y_j$  is the  $j^{th}$  word in the document. Given dialog query  $d$ , the goal of the CDP task is to recommend  $k$  documents in  $Y$  to the agent. For evaluation, we use Recall@ $k$  and Mean Reciprocal Rank, where the model is asked to select the  $k$  most likely documents, and it is correct if the correct URL document is among these  $k$  documents.

### 3.2 Information Retrieval approaches

Following previous works, the first approach we evaluate for this task is based on IR models. We use an Apache Lucene index, employed with English language analyzer and default BM25 similarity (Robertson and Zaragoza, 2009). Documents in the index are represented using two fields. The first field contains the actual document content. The second field augments the document’s representation with the text of all dialogs that link to it in the train-set (Amitay et al., 2005).

For a given (dialog) query  $d$ , matching documents are retrieved using four different ranking steps, which are combined using a cascade approach (Wang et al., 2011). Following (Van Gysel et al., 2016), we obtain an initial pool of candidate documents using a lexical query aggregation approach. To this end, each utterance  $t_i \in d$  is represented as a separate weighted query-clause, having its weight assigned relatively to its sequence position in the dialog (Van Gysel et al., 2016). Various sub-queries are then combined using a single disjunctive query. The second ranker evaluates each document  $y$  obtained by the first ranker against an expanded query (applying relevance model (Lavrenko and Croft, 2001)). The third ranker applies a manifold-ranking approach (Xu et al., 2011), aiming to score content-similar documents (measured by Bhattacharyya language-model based similarity) with similar scores.

The last ranker in the cascade treats the dialog query  $d$  as a verbose query and applies the *Fixed-Point* (FP) method (Paik and Oard, 2014) for weighting its words. Yet, compared to “traditional” verbose queries, dialogs are further segmented into distinct utterances. Using this observation, we implement an *utterance-biased* extension for enhanced word-weighting. To this end, we first score the various utterances based on the initial FP weights of words they contain and their relative position. We then propagate utterance scores back to their associated words. The IR model is denoted as *IRC*, short for *IR-Cascade* in Table 4.

### 3.3 Neural approaches

The second type of approaches we evaluate are neural models. We process the datasets to construct triples of  $\langle$ dialog context ( $d$ ), URL document content ( $y$ ), label (1/0) $\rangle$  from each dialog. For each  $d$ , we create a set of  $k + 1$  triples: one triple containing the correct URL provided by the agent (label

- 1), and  $k$  triples containing incorrect URLs randomly sampled from  $Y$  (label - 0). We explore different values for  $k$  and share additional results in Appendix. During evaluation, we evaluate a given dialog context against the set of all documents ( $Y$ ).

We evaluate the CDP task using three state-of-the-art neural models: *Enhanced Sequential Inference Model* (ESIM) proposed by Chen et al. (2017) which performs well on Natural Language Inference (NLI) and next utterance prediction tasks (Dong and Huang, 2018), *BertForSequenceClassification* model (Wolf et al., 2019) and *SBERT*. We next briefly describe the details for these models.

#### 3.3.1 ESIM

The ESIM model takes two input sequences: dialog context ( $d$ ) and document content ( $y$ ), and feeds them through BiLSTM to generate local context-aware word representations denoted by  $\vec{d}$  and  $\vec{y}$ . A co-attention matrix  $E$ , where  $E_{ij} = \vec{d}_i^T \vec{y}_j$ , computes the similarity between  $d$  and  $y$ . The attended dialog context and document content vectors denoted by  $\tilde{d}$  and  $\tilde{y}$  are computed using  $E$ , which represent the most relevant word in  $y$ ’s content for each word in  $d$ ’s context and vice-versa.

This local inference information is enhanced by computing the difference and the element-wise product for the tuple  $\langle \vec{d}, \tilde{d} \rangle$  as well as for  $\langle \vec{y}, \tilde{y} \rangle$ . The difference and element-wise product are then concatenated with the original vectors,  $\vec{d}$  and  $\tilde{d}$  or  $\vec{y}$  and  $\tilde{y}$  respectively. The concatenated vectors are then fed to another set of BiLSTMs to compose the overall inference between the two sequences. Finally, the result vectors are converted to a fixed-length vector by max pooling and fed to a final classifier.

#### 3.3.2 BERT

We use pre-trained BERT (Devlin et al., 2019) in two settings: a) fine-tuned on the training set, and b) an additional pre-training step on unlabeled data (dialogs in the training set and all documents) followed by fine-tuning on the training set (denoted as BERT\* in Table 4). In both settings, evaluation is done on the test set.

We utilize the binary classifier (*BertForSequenceClassification*) of BERT, commonly used for GLUE tasks (Wang et al., 2018) as follows. A dialog context  $d$  and a document  $y$  are fed together to BERT as a sequence ([CLS]  $d$  [SEP]  $y$  [SEP]). To adapt to BERT’s limitation of maximum sequence length, we use 512 tokens and feed BERT

Model	# param.	Inf. time (sec.)
BM25	2	0.02
IRC	12	0.03
ESIM	3.7M	2.37
BERT(*)	110M	0.95
SBERT(*)	110M	0.04

Table 3: Inference time for a single query from Twitter test set on a V100-PCI-E-32GB GPU

with the 256 tokens each from  $d$  and  $y$ , decided by a heuristic explained in Appendix. We use the hidden state of the [CLS] token as the representation of the pair. Training is done using positive and negative examples (Sec. 3.3) with cross-entropy loss. Re-ranking of candidate documents  $\{y \in Y\}$  for a given context  $d$  is done through the confidence score of each pair  $(d, y)$  which belongs to the positive class.

### 3.3.3 Sentence-BERT (SBERT)

We also explore SBERT (Reimers and Gurevych, 2019), which uses a Siamese network structure to fine-tune the pre-trained BERT network and derive semantically meaningful sentence embeddings. The sentence embeddings for  $d$  and  $y$  are derived by adding a pooling operation (default: mean) on the BERT outputs and then can be compared using cosine-similarity to achieve low inference time. We fine-tune SBERT in the same two settings as BERT mentioned above. The input handling and evaluation is same as BERT above. The fine-tuning and hyperparameter details are available in Appendix.

### 3.4 An hybrid approach

To investigate the real-world use of our approaches, we compare (in Table 3) the number of parameters of each model and inference time for a single query from the Twitter test set. The IRC model is much faster in comparison to the neural models. For incorporating the additional performance gain from neural models (in Table 4), we introduce an hybrid approach by a two-stage pipeline where we utilize the IRC model to generate a ranking of the document pool  $Y$ . The top- $k$  documents ( $k : 20$ ) are then re-ranked through ESIM and recommended to the CC agent. This hybrid approach (IRC+ESIM) combines the best of both worlds.

## 4 Results and Analysis

Results are presented in Table 4. We provide training setting and hyperparameter details for all neural models in Appendix. We observe that the ESIM model performs best across all datasets and the IRC

Model	R@1	R@2	R@5	R@10
<i>Mac-Support</i>				
BM25	0.199	0.278	0.394	0.479
IRC	0.411	0.567	0.734	0.809
ESIM	0.419	0.602	0.758	0.848
BERT	0.319	0.441	0.655	0.809
BERT*	0.315	0.447	0.698	0.818
SBERT	0.096	0.177	0.299	0.460
SBERT*	0.128	0.203	0.319	0.496
IRC+ESIM	<b>0.496</b>	<b>0.684</b>	<b>0.872</b>	<b>0.985</b>
<i>Telco-Support</i>				
BM25	0.032	0.068	0.182	0.313
IRC	0.405	0.551	0.735	0.867
ESIM	0.676	0.806	0.911	0.951
BERT	0.523	0.699	0.866	0.918
BERT*	0.569	0.748	0.891	0.927
SBERT	0.250	0.391	0.612	0.758
SBERT*	0.360	0.506	0.711	0.826
IRC+ESIM	<b>0.721</b>	<b>0.863</b>	<b>0.942</b>	<b>0.964</b>
<i>Twitter</i>				
BM25	0.088	0.150	0.224	0.306
IRC	0.420	0.554	0.728	0.802
ESIM	0.474	0.590	0.680	0.772
BERT	0.400	0.418	0.424	0.428
BERT*	0.370	0.382	0.386	0.386
SBERT	0.182	0.246	0.354	0.442
SBERT*	0.224	0.308	0.484	0.644
IRC+ESIM	<b>0.559</b>	<b>0.684</b>	<b>0.819</b>	<b>0.902</b>

Table 4: Performance of models on the test set of three datasets. R@ $k$  refers to Recall at position  $k$ . MRR and corresponding validation results are in Appendix.

model performs comparably to the ESIM model except for Telco-Support dataset. We observe a significant performance reduction with BERT models in comparison to both IRC and ESIM models. The BERT\* model (additional pre-training) improves performance for Telco-Support dataset, but is still inferior to ESIM model. The SBERT models provide the benefit of low inference time, but reduce performance further. We conclude that for CDP task, explicit cross-attention between dialog context  $d$  and document  $y$  present in ESIM is crucial. The BERT models try to incorporate cross-attention through self-attention on the concatenated  $\langle d, y \rangle$  pair sequence, but still lag behind.

Finally, the hybrid approach (IRC+ESIM) provides a significant boost in performance (e.g., between +7%-20% in R@1), and reduces the inference time of ESIM. This demonstrates the benefit and importance of combining IR models that are based on exact matching, with neural models that further allow semantic inference in the domain for real-world applications.

## 5 Conclusion and Future Work

We introduced the Conversational Document Prediction (CDP) task and investigated the perfor-



mance of state-of-the-art DL and IR models. We also release a new public Twitter dataset on the CDP task. In this work, we considered only URL documents with content. Other potential document types that could be considered are PDFs, doc etc. and URLs without content (e.g. login, tracking). We plan to address these challenges in future work.

## References

- Mohammad Aliannejadi, Hamed Zamani, Fabio Crestani, and W Bruce Croft. 2019. Asking clarifying questions in open-domain information-seeking conversations. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 475–484.
- Einat Amitay, Adam Darlow, David Konopnicki, and Uri Weiss. 2005. Queries as anchors: selection by association. In *Proceedings of the 16th ACM Conference on Hypertext and Hypermedia*, pages 193–201.
- Qian Chen, Xiaodan Zhu, Zhen-Hua Ling, Si Wei, Hui Jiang, and Diana Inkpen. 2017. Enhanced lstm for natural language inference. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 1657–1668.
- 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.
- Jianxiong Dong and Jim Huang. 2018. Enhance word representation for out-of-vocabulary on ubuntu dialogue corpus. *arXiv preprint arXiv:1802.02614*.
- Diogo Ferreira, Mário Antunes, Diogo Gomes, and Rui L Aguiar. 2019. Classy: A conversational aware suggestion system. In *Multidisciplinary Digital Publishing Institute Proceedings*, volume 31, page 40.
- Maryam Habibi and Andrei Popescu-Belis. 2015. Keyword extraction and clustering for document recommendation in conversations. *IEEE/ACM Transactions on audio, speech, and language processing*, 23(4):746–759.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Victor Lavrenko and W. Bruce Croft. 2001. [Relevance based language models](#). In *Proceedings of the 24th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '01*, page 120–127, New York, NY, USA. Association for Computing Machinery.
- Ryan Lowe, Nissan Pow, Iulian V Serban, and Joelle Pineau. 2015. The ubuntu dialogue corpus: A large dataset for research in unstructured multi-turn dialogue systems. In *Proceedings of the 16th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 285–294.
- Wenhao Lu, Jian Jiao, and Ruofei Zhang. 2020. Twinbert: Distilling knowledge to twin-structured bert models for efficient retrieval. *arXiv preprint arXiv:2002.06275*.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.
- Bhaskar Mitra, Nick Craswell, et al. 2018. An introduction to neural information retrieval. *Foundations and Trends® in Information Retrieval*, 13(1):1–126.
- Jiaul H. Paik and Douglas W. Oard. 2014. [A fixed-point method for weighting terms in verbose informational queries](#). In *Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management, CIKM '14*, page 131–140, New York, NY, USA. Association for Computing Machinery.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543.
- 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 and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3973–3983.
- Stephen Robertson and Hugo Zaragoza. 2009. [The probabilistic relevance framework: Bm25 and beyond](#). *Found. Trends Inf. Retr.*, 3(4):333–389.
- Christophe Van Gysel, Evangelos Kanoulas, and Maarten de Rijke. 2016. [Lexical query modeling in session search](#). In *Proceedings of the 2016 ACM International Conference on the Theory of Information Retrieval, ICTIR '16*, page 69–72, New York, NY, USA. Association for Computing Machinery.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2018. Glue: A multi-task benchmark and analysis platform for natural language understanding. In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 353–355.

Lidan Wang, Jimmy Lin, and Donald Metzler. 2011. [A cascade ranking model for efficient ranked retrieval](#). In *Proceedings of the 34th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '11, page 105–114, New York, NY, USA. Association for Computing Machinery.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, R’emi Louf, Morgan Funtowicz, and Jamie Brew. 2019. Huggingface’s transformers: State-of-the-art natural language processing. *ArXiv*, abs/1910.03771.

Bin Xu, Jiajun Bu, Chun Chen, Deng Cai, Xiaofei He, Wei Liu, and Jiebo Luo. 2011. [Efficient manifold ranking for image retrieval](#). In *Proceedings of the 34th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '11, page 525–534, New York, NY, USA. Association for Computing Machinery.

Wei Yang, Haotian Zhang, and Jimmy Lin. 2019. Simple applications of bert for ad hoc document retrieval. *arXiv preprint arXiv:1903.10972*.

Yongfeng Zhang, Xu Chen, Qingyao Ai, Liu Yang, and W Bruce Croft. 2018. Towards conversational search and recommendation: System ask, user respond. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*, pages 177–186.

## A Appendix: Additional results

The results for corresponding validation performance for Information Retrieval and Neural approaches, as well as Mean Reciprocal Rank (MRR) metric for both validation and test sets on all datasets are available in Table 7.

### A.1 Negative samples for neural approaches

For creating training data for our neural approaches, we create  $k$  triples containing incorrect URLs sampled randomly from set of all documents. We experimented with different values for the hyperparameter *num\_negative\_samples* used for generating the training data. The results for ESIM model for the Mac-Support dataset are presented in Table 5. We observe that increasing the number of negative samples doesn’t improve the ESIM model performance significantly and *num\_negative\_samples* - 4 provides us the best of both worlds, i.e. good performance and lower training time, in comparison

to using a higher negative sample ratio. We use the same value for all neural models for all datasets.

### A.2 Input handling for BERT models

To handle the BERT model input limitation of 512 tokens max sequence length, we feed BERT with 256 tokens each from dialog context  $d$  and document content  $y$ . We observe that the initial sentences in a URL document always capture the core gist of the document, so we always use the first 256 tokens from the document content. For dialog context, we observe that as the conversation progresses over multiple turns and the user query gets more complex, the conversation shifts from the original query to another problem in many dialogs. We explore two input approaches for deciding which tokens to consider if dialog context sequence length  $|d| > 256$ :

1. Input-A: Truncate the dialog context  $d$  to consider only the first 256 tokens from the dialog context.
2. Input-B: Ignore tokens in the middle of dialog context sequence to reduce the  $|d|$  to 256.

The results for both approaches for BERT model on the Telco-Support dataset are in Table 6. We use the same heuristic for all neural models for all datasets.

negative samples	R@1	R@2	R@5	R@10	MRR
4	0.417	0.535	0.676	0.745	0.534
7	0.400	0.507	0.633	0.728	0.510
10	0.419	0.509	0.683	0.779	0.534
14	0.419	0.518	0.678	0.747	0.531

Table 5: Performance of ESIM model on the validation set of Mac-Support dataset for different values of *num\_negative\_samples*.

dialog input	R@1	R@2	R@5	R@10	MRR
Input-A	41.19	58.72	80.03	88.8	57.26
Input-B	52.38	70.89	87.13	92.4	67.04

Table 6: Performance of BERT model on the validation set of Telco-Support dataset for dialog context sequence input handling.

Model	Validation					Test				
	R@1	R@2	R@5	R@10	MRR	R@1	R@2	R@5	R@10	MRR
<i>Mac-Support</i>										
BM25	0.169	0.236	0.358	0.437	0.246	0.199	0.278	0.394	0.479	0.283
IRC	0.407	0.548	0.713	0.805	0.537	0.411	0.567	0.734	0.809	0.546
ESIM	0.417	0.569	0.741	0.831	0.561	0.419	0.602	0.758	0.848	0.573
BERT	0.302	0.447	0.625	0.769	0.444	0.319	0.441	0.655	0.809	0.470
BERT*	0.332	0.501	0.711	0.824	0.497	0.315	0.447	0.698	0.818	0.471
SBERT	0.079	0.137	0.267	0.434	0.185	0.096	0.177	0.299	0.460	0.207
SBERT*	0.100	0.149	0.312	0.494	0.214	0.128	0.203	0.319	0.496	0.238
<i>Telco-Support</i>										
BM25	0.039	0.069	0.178	0.313	0.100	0.032	0.068	0.182	0.313	0.097
IRC	0.409	0.549	0.737	0.859	0.547	0.405	0.551	0.735	0.867	0.546
ESIM	0.683	0.803	0.913	0.953	0.782	0.676	0.806	0.911	0.951	0.779
BERT	0.524	0.709	0.871	0.924	0.670	0.523	0.699	0.866	0.918	0.667
BERT*	0.568	0.745	0.899	0.936	0.706	0.569	0.748	0.891	0.927	0.702
SBERT	0.266	0.403	0.626	0.760	0.423	0.250	0.391	0.612	0.758	0.410
SBERT*	0.365	0.518	0.724	0.834	0.521	0.360	0.506	0.711	0.826	0.512
<i>Twitter</i>										
BM25	0.111	0.156	0.265	0.375	0.177	0.088	0.150	0.224	0.306	0.148
IRC	0.499	0.625	0.777	0.819	0.611	0.420	0.554	0.728	0.802	0.549
ESIM	0.548	0.642	0.747	0.806	0.642	0.474	0.590	0.680	0.772	0.579
BERT	0.474	0.489	0.489	0.489	0.482	0.400	0.418	0.424	0.428	0.411
BERT*	0.474	0.484	0.484	0.484	0.479	0.370	0.382	0.386	0.386	0.377
SBERT	0.226	0.295	0.373	0.447	0.311	0.182	0.246	0.354	0.442	0.276
SBERT*	0.321	0.417	0.573	0.683	0.439	0.224	0.308	0.484	0.644	0.349

Table 7: Performance of models on the validation and test sets for the three datasets.  $R@k$  refers to Recall at position  $k$  in all documents, denoted as  $R@1$ ,  $R@2$ ,  $R@5$  and  $R@10$ . MRR refers to the Mean Reciprocal Rank.

## B Appendix: Model Training and Hyperparameter Details

### B.1 ESIM model

We used 300-dimensional Glove pre-trained vectors (Pennington et al., 2014), 100-dimensional word2vec vectors (Mikolov et al., 2013) and 80-dimensional character embedding vectors for generating the word representation. For training word2vec vectors, we use the gensim API with the following hyper-parameters: size=100, window=10, min\_count=1 and epochs=20. We also incorporate character embeddings to our ESIM implementation (Dong and Huang, 2018). The final prediction layer is a 2-layer fully-connected feed-forward neural network with ReLU activation. We use sigmoid function and minimize binary cross-entropy loss for training and updating the model. We used Adam (Kingma and Ba, 2014) with a learning rate of 0.001 and exponential decay with a decay rate of 0.96 decayed every 5000 steps. The

number of hidden units for BiLSTMs was 256. For the prediction layers, we used 256 hidden units with ReLU activation.

### B.2 Additional pretraining for BERT model

We use the BERT-Base, Uncased model from [BERT-Base-Uncased](#) - 12-layer, 768-hidden, 12-heads, 110M parameters - as the base model for our experiments. For convenience, we refer to BERT-Base-Uncased as \$BERT below. We use the code from [Google-Research Bert Github](#) for creating pretraining data as well as to run additional pretraining on our domain data. We only use the training dialogs and contents from all documents for creating pretraining data. The hyperparameters used for creating pretraining data are:

```

vocab_file=$BERT/vocab.txt
do_lower_case=True
max_seq_length=512
max_predictions_per_seq=20
masked_lm_prob=0.15

```

```
random_seed=12345
dupe_factor=10
```

The hyperparameters used to run LM-pretraining are:

```
train_batch_size=16
max_seq_length=512
max_predictions_per_seq=20
num_train_steps=100000
num_warmup_steps=10000
save_checkpoints_steps=20000
learning_rate=5e-5
```

### B.3 Fine-tuning BERT model

The hyperparameters used for further fine-tuning BERT model are:

```
do_lower_case=True
max_seq_length=512
per_gpu_eval_batch_size=24
per_gpu_train_batch_size=24
learning_rate=2e-5
num_train_epochs=5
```

The model is periodically evaluated on the validation set after  $n$  steps, which is decided based on the training dataset size.

### B.4 Fine-tuning SBERT model

The hyperparameters used for fine-tuning SBERT model are:

```
do_lower_case=True
max_seq_length=256
batch_size=16
```

```
learning_rate=2e-5
num_train_epochs=5
optimizer=Adam
```

We use a linear learning rate warm-up over 10% of the training data. We fine-tune SBERT with a 3-way softmax-classifier objective function and the default pooling strategy is MEAN. The `max_seq_length` is 256 each for dialog context  $d$  and document content  $y$ . For SBERT\*, we use the same additional pre-trained BERT\* model from before. The model is periodically evaluated on the validation set after  $n$  steps, which is decided based on the training dataset size.

## C Appendix: Extracting content from URL documents

For the internal Mac-Support dataset, the document content for each URL was obtained by API calls to the customer service knowledge base. For the Telco-Support and Twitter datasets, we capture the HTML content using a Selenium Chrome webdriver, which renders the URL document by loading all CSS styling and Javascript. The extracted HTML was cleaned through a Markdown generation pipeline, where we manually identify and filter the DOM tags (using CSS id and/or class) which correspond to header(s), footer, navigation bars etc. This process is repeated for each URL domain in both datasets. The tools for data preprocessing are available here: <https://github.com/IBM/MDfromHTML>.