

A Systematic Evaluation of Response Selection for Open Domain Dialogue

Behnam Hedayatnia, Di Jin, Yang Liu, Dilek Hakkani-Tur

Amazon Alexa AI

{behnam, djinamzn, yangliud, hakkaniit}@amazon.com

Abstract

Recent progress on neural approaches for language processing has triggered a resurgence of interest on building intelligent open-domain chatbots. However, even the state-of-the-art neural chatbots cannot produce satisfying responses for every turn in a dialog. A practical solution is to generate multiple response candidates for the same context, and then perform response ranking/selection to determine which candidate is the best. Previous work in response selection typically trains response rankers using synthetic data that is formed from existing dialogs by using a ground truth response as the single appropriate response and constructing inappropriate responses via random selection or using adversarial methods. In this work, we curated a dataset where responses from multiple response generators produced for the same dialog context are manually annotated as appropriate (positive) and inappropriate (negative). We argue that such training data better matches the actual use case examples, enabling the models to learn to rank responses effectively. With this new dataset, we conduct a systematic evaluation of state-of-the-art methods for response selection, and demonstrate that both strategies of using multiple positive candidates and using manually verified hard negative candidates can bring in significant performance improvement in comparison to using the adversarial training data, e.g., increase of 3% and 13% in Recall@1 score, respectively.

1 Introduction

Building an open-domain dialog system to interact with users on a variety of topics can involve building multiple response generators (RG) with different functions (Paranjape et al., 2020). These RGs can be a mixture of generative, retrieval and template based methods. A response selector is then built to re-rank response candidates produced by different applicable RGs to determine the best response for a given turn. These response selectors

are based on either rule-based or model-based architectures (Papaioannou et al., 2017; Serban et al., 2017; Zhou et al., 2020; See and Manning, 2021a). Rule-based systems typically consist of manually-designed logic to rank hypotheses, whereas model-based approaches can either be conventional machine learning models or recent neural models that learn to rank candidates. As the number of RGs grows, a rule-based system can become cumbersome to maintain, whereas model-based methods can simplify the selection process as well as achieve better performance.

Latest work in model-based response selectors involves leveraging pretrained transformer models such as BERT (Devlin et al., 2019) and DialoGPT (Zhang et al., 2019). These selection models are often trained using existing dialog datasets that typically contain ground truth responses. Thus a focus of past response selection work is on the construction of inappropriate/negative responses, using methods such as random selection, utterance manipulation or leveraging user feedback (Whang et al., 2020; Han et al., 2021; Whang et al., 2021; Gu et al., 2020; Xu et al., 2020; Zhang and Zhao, 2021; Gao et al., 2020; See and Manning, 2021b; Gupta et al., 2021; Li et al., 2019). However, such synthesized datasets for response selection have the following known drawbacks. First of all, their claimed incorrect responses are not verified if they are actually incorrect. Second, these negative responses are easy to differentiate from positive ones since it is very likely that they will be on different topics from the context. Therefore, models trained on such easy negative responses will not be able to generalize to real-world settings, where multiple responses are generated given the same dialog context and many of them are strong candidates.

To resolve the aforementioned issues, we construct a new dataset (named RSD) for response selection by showing human annotators multiple response candidates produced by different RGs for

a given turn and dialog context, and asking them to annotate all responses that are appropriate for that specific dialog context. We leverage RSD to conduct a systematic evaluation of state-of-the-art methods for response selection, including existing trained models, DialogRPT (Gao et al., 2020) and BERT-FP (Han et al., 2021), and a BERT based ranker that we trained. Our experimental results show the following findings: (1) Models trained on RSD significantly outperform those trained on existing datasets, e.g., Reddit and Ubuntu, showing the benefit of bringing in human annotated data for this task; (2) Using manually verified hard negatives greatly outperforms using adversarial negatives; (3) Training on multiple positive candidates improves performance in comparison to a single positive candidate. Though these findings are most expected, this is the first empirical study that clearly shows that constructing a more realistic dataset benefits strongly over generating synthetic examples for response selection, and we hope such results can guide future research in this direction and deployment of open domain dialog systems.

2 Related Work

Previous work in response selection has been conducted in different domains, such as chatlogs (Lowe et al., 2015), e-commerce (Zhang et al., 2018b), and open-domain dialog (Wu et al., 2017; Zhang et al., 2018a; Smith et al., 2020; See and Manning, 2021b). Our work focuses on open-domain dialog, where current systems typically consist of multiple response generators, each of which is designed to deal with a certain domain. For example, in the Alexa Prize challenge (Ram et al., 2018; Gabriel et al.), most of the participating socialbots built by university teams consist of a variety of responders that are based on retrieval-based methods, template-based methods, or generative models (Konrád et al., 2021; Saha et al., 2021; Paranjape et al., 2020; Ram et al., 2018). In order to select the final response to present to users, both rule-based or model-based ranking models have been proposed (Ram et al., 2018; Papaioannou et al., 2017; Serban et al., 2017; Zhou et al., 2020; See and Manning, 2021a; Shalyminov et al., 2018). This approach is also common in other real-world systems such as XiaoIce that employs a manually-designed set of features to rank hypotheses (Zhou et al., 2020).

For training response selection models, typically

human-human dialogs are used, where positive examples are the ground truth responses and negative responses are often randomly selected or synthetically created since there are no labeled negative responses. Han et al. (2021) randomly selected responses from other dialogs or within the same dialog session. Whang et al. (2021) corrupted utterances by inserting, substituting and deleting random tokens. Xu et al. (2020) masked and shuffled utterances within a dialog. Li et al. (2019) selected negative responses from a batch based on their similarity scores from the positive response score. Gupta et al. (2021) used automatic methods such as replacing random tokens in a positive examples using a Mask-and-fill approach to create adversarial negative examples.

However, these sampling strategies do not ensure the selected negative responses are hard examples. In this work, rather than relying on approximation for negative responses, we perform turn level annotation of multiple response candidates for response appropriateness for a given dialog context. See and Manning (2021b); Gao et al. (2020) did construct hard negative examples by annotating responses from a single generative model for appropriateness; however, our work contains responses from a mixture of various RG methods.

On the other hand, open-domain dialogs can have multiple appropriate responses for a given dialog context. Previous work has augmented dialog datasets with multiple positive examples (Mizukami et al., 2015; Khayrallah and Sedoc, 2020; Gupta et al., 2019; Sai et al., 2020; Zhang et al., 2020). Within open-domain dialogs, Gupta et al. (2019); Sai et al. (2020) augmented the DailyDialog dataset (Li et al., 2017) with multiple positive human written responses. In contrast, our dataset has multiple positive responses generated from models, which reduces the cost of human annotation significantly. The closest work to ours is (Sai et al., 2020) that constructed negative examples by asking annotators to copy information from the dialog context. We do not restrict the definition of negative examples to be copying information from the dialog context, since incorrect responses in open-domain dialog can have different issues, e.g., off-topic, contradicting or repetitive responses.

3 Datasets

As described earlier, most previous work in response selection has constructed test sets that typi-

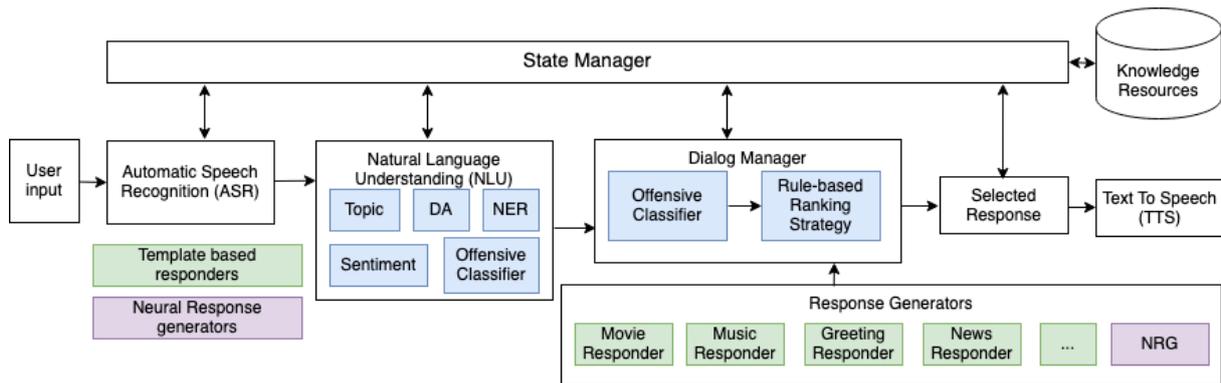


Figure 1: Architecture of our Open Domain Dialog System. NER = Named Entity Recognition, DA = Dialog Act

cally contain only one positive candidate and one or more synthetically created negative candidates. However, such negative responses may be easy for a model to detect. Additionally, in real-world open-domain dialogs there can be more than one positive response per turn. Therefore, in this work we constructed a more realistic dataset consisting of annotations for real response candidates. Our dataset consists of spoken interactions between a dialog system and real users.

3.1 Open Domain Dialog System

We first describe the open-domain dialog system used for data collection. The architecture of our dialog system is shown in Figure 1. Every user utterance in the dialog is sent into an ASR system whose output goes through a series of NLU modules that classifies topics, dialog acts, sentiment, extracts entities, and detects if user utterance is offensive. Our system then calls multiple response generators for the given dialog context and logs all the generated response candidates within the State Manager. The response presented to the user is selected by a rule-based ranker and then sent to the TTS module.

For popular topics in open domain dialogs, such as movies, music, recent news, we developed template-based response generators (highlighted in green in Figure 1) for the given dialog state. An example state and response for the movie domain is: when the user turn mentions a movie name (based on the NER result), we respond with information about the actor, the rating, or the plot of this certain movie. In addition to topic-specific template-based RGs, our system includes other template-based RGs for different dialog contexts, such as, greetings, topic switches, etc.

For every user turn, we also apply a neural

network-based response generation (NRG) model to produce a response, highlighted in purple in Figure 1. Our *NRG Responder* is a GPT2-XL (Radford et al., 2019) based model trained on real user conversation data described in Section 3.2. We discuss its training details in Appendix B.

The rule-based ranker uses predefined logic and the topic extracted from the user utterance to select domain specific template-based responders. If a template-based responder is not available it will use the NRG response as a fall back. Our system has just a few template-based RGs, and uses NRG responses for almost half of all turns.

3.2 Response Selection Data (RSD)

We deploy the dialog system described above within the Alexa Prize Socialbot framework (Ram et al., 2018) to interact with real users. A user initiates an interaction with our dialog system and consents to have their data being collected. These interactions end when the user requests to stop the conversation. At the end of each interaction, users are asked to leave a rating in the range of 1 to 5. We denote this dataset as real user interactions (RUI)¹. Our data consists of approximately 100k interactions and 2.5 million turns. For each user turn in RUI, we produced additional response candidates using variants of our *NRG Responder* to supplement the logged responses. These may be appropriate responses, or hard negative examples. The NRG variants we used include the following (Further model training details are in Appendix B).

- A GPT2-medium version of our *NRG Responder*.
- A GPT2-XL *NRG Responder* grounded on knowledge. When there is an entity in the user

¹All interactions are in English.

Data Split	# Dialogs	# positive responses	# negative responses	Avg. # responses at each turn	# Turns with no positive responses
RSD Train	1,501	17,778	78,273	5.67	8,871
RSD Test	142	2,995	6,298	5.36	309

Table 1: Dataset Statistics. For our experiments, we conduct 5 fold cross validation on all our training datasets and therefore do not have a dedicated development set.

turn, we search Wikipedia to find the article related to the entity, and perform knowledge selection and knowledge-grounded response generation.

- A GPT2-medium *NRG Responder* grounded on dialog acts (DA) (Hedayatnia et al., 2020).
- A GPT2-XL based sentiment controlled *NRG Responder*. When the user’s utterance shows some negative sentiment (e.g., when a person says “I’m depressed”), the NRG model generates a response conditioned on this emotion.

We worked with internal human annotators to set up an annotation pipeline. These internal annotators are not experts in the dialog domain; however, we worked closely with them to ensure they have a clear understanding of the task provided to them. In our annotation pipeline, for each turn in a dialog, we showed internal human annotators all the available responses produced by the template based generators and various NRG models², and asked them whether each response candidate is appropriate given the certain dialog context. An annotator can label multiple responses or none of them as appropriate. To determine if a response is appropriate we ask annotators to see if the response is relevant to the dialog context and that it does not contradict what was said in previous dialog system’s responses. For data annotation we randomly sampled a subset of RUI that contain dialogs with more than 5 turns and fewer than 30 turns. A snapshot of the interface for the annotation task can be found in Appendix C.

We randomly split the annotated conversations into training and test sets. Table 1 shows the statistics of our annotated response selection data, denoted as RSD. Due to user privacy constraints, we cannot release this data. Note that we assume our response selector must always choose a response and therefore we drop turns where none of the responses are labeled as appropriate, and for each turn, we may have multiple positive and negative responses.

²Note that for the NRG models, we only use responses produced within a pre-defined timeout period.

3.3 RSD Training Variations

To show the importance of using hard negative and multiple positive candidates for response selection, we have also created five variations of the train set of RSD. In our experiments, for each variation we ran random sampling five times and report the average results.

- RSD Train with one positive candidate (denoted as “RSD 1 Pos.”). Based on the original RSD Train, we sample only one positive candidate for each turn from the multiple positive candidates, and keep all the annotated negative responses. This leads to 8,046 positive and 78,273 negative candidates.
- Synthetic Inter-Random. Based on the above-mentioned “RSD 1 Pos.” set, we further remove the human annotated negative candidates, and instead use five randomly selected responses from other dialogs and deem these as the new negative candidates. There are 8,046 positive and 40,230 negative candidates in this set. This approach to constructing negative candidates is commonly used in the literature. We experimented with different number of negative candidates and found sampling 5 negative candidates at each turn had the best results.
- Synthetic Intra-Random. Similar to the above set, we use one positive example and four randomly selected responses as negative, two drawn from a random different dialog and the other two from the same dialog as the candidate we are training on. This set contains 8,046 positive and 32,184 negative candidates. This approach to constructing negative candidates is proposed by (Han et al., 2021). We experimented with different number of negative candidates and found sampling 4 negative candidates at each turn had the best results.
- Synthetic Adversarial. Based on the above-mentioned “RSD 1 Pos.” set, we further create negative candidates using the Mask-and-Fill

approach from (Gupta et al., 2021). This approach uses the hierarchical masking function from (Donahue et al., 2020) to replace spans in a positive example with blank tokens that will be replaced with tokens predicted from an Infilling Language Model from (Donahue et al., 2020). For every turn, an average of 28.22 negative candidates were constructed using this approach. We experimented with different number of negative candidates and found sampling 10 negative candidates at each turn had the best results. In total, we have 8,046 positive and 76,307 negative candidates.

- **Synthetic Retrieval.** In this approach, we generate negative examples that are semantically similar to the positive example. This approach to constructing negative candidates is proposed by (Li et al., 2019). The motivation behind this approach is to create negative candidates that are somewhat similar to the positive candidate and use these as hard examples for the model to train on. Specifically, we use the *all-MiniLM-L6-v2* model (Wang et al., 2020) from HuggingFace³ and create a sentence embedding for each response in our dataset. At each turn we compute the cosine similarity between the positive candidate and all the other responses in the dataset. We then take responses that have a cosine similarity between 0.8 and 0.95 as a negative candidate. We experimented with different thresholds and found this had the best results. Using these thresholds we get an average of 2.2 negative candidates per turn. In total we have 8,046 positive and 17,778 negative candidates.

4 Response Selection Models

We have adopted two state-of-the-art methods for response selection and adapted them to our new dataset for a comprehensive empirical evaluation.

4.1 DialogRPT (Gao et al., 2020)⁴

DialogRPT is initialized with DialoGPT (Zhang et al., 2019) and trained using a contrastive loss function to predict a higher score for the positive response given the dialog context and a pair of one positive and one negative response. Trained on

³<https://huggingface.co/sentence-transformers/all-MiniLM-L6-v2>

⁴<https://github.com/golsun/DialogRPT/>

the Reddit dataset, five different ranker models are proposed by training DialogRPT on different synthesized labels (see the original paper for details).

4.2 BERT Models

We experiment with two different BERT model variants for response selection:

BERT-FP (Han et al., 2021)⁵: BERT-FP has achieved high scores on the Ubuntu Dialogue Corpus test set (Lowe et al., 2015). The authors post-train the Masked Language Model (MLM) head and Next Sentence Prediction (NSP) head of a BERT-base model (Devlin et al., 2019) on the Ubuntu corpus via unsupervised learning. Given a dialog context and a response, the NSP head is trained to predict whether a response is either: the ground truth, from a random dialog, or from a random turn in the same dialog. After post-training, the model is further fine-tuned on downstream data for response selection, where given a dialog context and a system response, the model classifies whether this is the correct response or not.

BERT-Ranker: We directly fine-tune a BERT-base (Devlin et al., 2019) model without the above-mentioned post-training step. We denote this model as BERT-Ranker.

Figure 2 illustrates the fine-tuning stage for both BERT models. To construct our input, we concatenate the dialog context with a system response and follow the same training procedure used by (Han et al., 2021), which uses the pooled output representation by the BERT model, passes it through a linear layer followed by a sigmoid function, and minimizes the binary cross-entropy function to predict whether the given system response is positive or negative.

5 Experiments

5.1 Experimental Setup

Following the previous work (Whang et al., 2020; Han et al., 2021; Whang et al., 2021; Gu et al., 2020; Xu et al., 2020; Zhang and Zhao, 2021), for evaluation metrics, we use MRR (mean reciprocal rank) and Recall at k (R@k), which is defined as the correct answer existing among the top-k candidates.

For DialogRPT, we run their five different rankers out of the box over RSD Test in a zero-shot fashion and find that the human vs random ranker scores the highest for both MRR and Recall,

⁵https://github.com/hanjanghoon/BERT_FP

Model	Train Data	MRR	R@1	R@2	R@3	R@4	R@5
DialogRPT	Reddit	0.681	0.481	0.730	0.868	0.939	0.988
BERT-FP	Ubuntu	0.684	0.486	0.742	0.864	0.930	0.979
DialogRPT	RSD Train	0.787	0.647	0.834	0.910	0.979	0.992
BERT-FP	RSD Train	0.795	0.657	0.841	0.931	0.973	0.994
BERT-R	RSD Train	0.796	0.659*	0.843*	0.936*	0.980*	0.995*
BERT-R	RSD 1 Pos.	0.762(0.06)	0.628(0.05)	0.806(0.06)	0.894(0.07)	0.941(0.07)	0.958(0.07)
BERT-R	Synthetic IE	0.688(0.01)	0.488(0.01)	0.741(0.00)	0.880(0.00)	0.949(0.00)	0.984(0.00)
BERT-R	Synthetic IA	0.698 (0.00)	0.506 (0.00)	0.750 (0.00)	0.879 (0.00)	0.948 (0.00)	0.983 (0.00)
BERT-R	Synthetic Adv	0.712 (0.00)	0.532 (0.00)	0.753 (0.00)	0.884 (0.00)	0.950 (0.00)	0.987 (0.00)
BERT-R	Synthetic Ret	0.718 (0.00)	0.533 (0.01)	0.776 (0.00)	0.902 (0.00)	0.961 (0.00)	0.990 (0.00)

Table 2: Model results on RSD Test. Results for BERT-R (Ranker) using Synthetic datasets are computed by sampling candidates with five different seeds and averaging the model prediction results across those runs. Standard deviations are in parentheses. IE (Inter-Random), IA (Intra-Random), Adv (Adversarial), Ret (Retrieval) are the four different ways of creating negative examples described in Section 3.3. Recall numbers marked with * mean that the improvement is statistically significant compared with Synthetic Ret (mcnemar with p-value < 0.05).

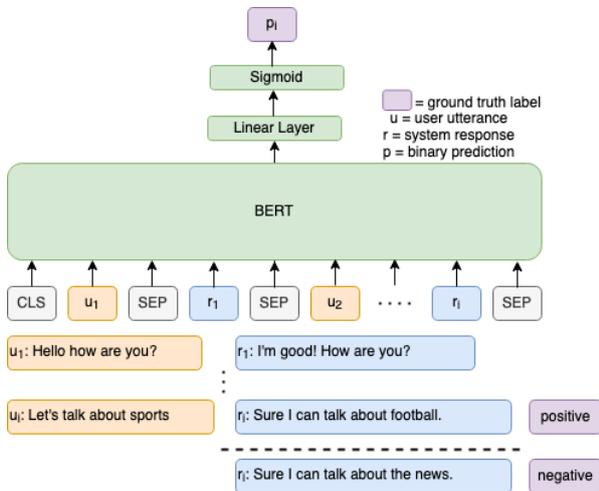


Figure 2: Model architecture of BERT-Ranker and BERT-FP.

therefore we fine-tune this model on RSD Train following the same training approach in the original paper. Since we have p positive and n negative candidates for each turn, we can obtain $p \times n$ example pairs. For our BERT models, we finetune both BERT-FP and BERT-Ranker on RSD Train. To evaluate the effect of positive and negative examples, we finetune the BERT-Ranker using different RSD training variations described in Section 3.3.

We also implemented model ensembling for all the methods. We first divide the training set into five folds, and each time we choose four of them for model training and the remaining one for validation. In this way, we obtain five trained models, and then average their prediction probability outputs on the test set to get the final prediction scores. Further training details are provided in Appendix A.

5.2 Results

Table 2 shows the results on RSD Test using different models and training configurations. From the table, we have the following findings:

- We observe that there is no performance improvement when training BERT-FP on RSD Train versus BERT-Ranker on RSD Train. Therefore the post-training process via optimizing the MLM and NSP objectives proposed in the BERT-FP model does not bring an extra advantage.
- By comparing DialogRPT trained on both Reddit and RSD Train as well as BERT-FP trained on Ubuntu and RSD Train, we can see that the same models trained on our labeled data lead to much better performance, because of the matched training and testing setup.
- We observe that training BERT-Ranker on adversarially created negatives (Synthetic Adv.) and (Synthetic Ret.) outperforms using random negatives within the same dialog (Synthetic IA), achieving Recall@1 scores 0.532 and 0.506, respectively. However, training on adversarial examples (Synthetic Adv.) and (Synthetic Ret.) still significantly underperforms training on human-verified hard negatives (RSD Train), which achieved a Recall@1 score of 0.659.
- We see the benefit of leveraging multiple positive responses, i.e., BERT-Ranker (RSD Train) outperforms BERT-Ranker (RSD 1 Pos.) with Recall@1 scores of 0.659 and 0.628, respectively.

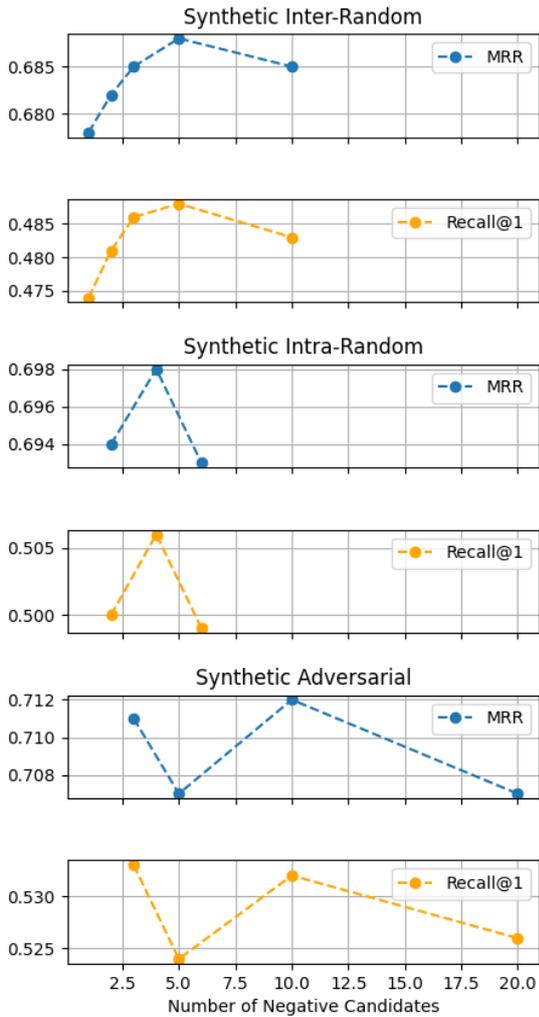


Figure 3: Results when increasing the number of negative candidates in the Synthetic Datasets Training set. See Recall@2-5 results in Appendix D

Example 1

Dialog History:

System: If you don't mind me asking, what type of music are you into ?

User: i like rock my favorite band is the arctic monkeys

BERT Ranker (RSD Train): Thanks for sharing! If you don't mind me asking, what is your favorite rock song ?

BERT Ranker (Synthetic Inter-Random): i see, you like rock your favorite band is the arctic monkeys. i like rock too. who is your favorite artist?

Example 2

Dialog History:

System: i'm happy to talk about movies. so, have you seen a good movie lately?

User: oh yeah

BERT Ranker (RSD Train): nice. what movie did you see?

BERT Ranker (Synthetic Adversarial): tom hanks is a great actor. are you a fan?

Example 3

Dialog History:

System: so do you have any pets at home?

User: yeah i've got two

System: cool what is it?

User: it's a dog

BERT Ranker (RSD Train): wow! that's great! i like dogs! what is its breed?

BERT Ranker (Synthetic Retrieval): cool! what kind of pet?

Figure 4: Example predictions of BERT Ranker (RSD Train) outperforming BERT Ranker trained on Synthetic Datasets. Due to privacy concerns, these example dialogs are from an internal author.

Example 1

Dialog History:

System: How about we talk about movies, news or music. So, what topic would you like to talk next ?

User: tell me about

Model selected response: ok. i can fetch ratings from rotten tomatoes, check for fresh movies in the theatres, tell you movie plot and reviews, and more.

A correct response: i didnt hear you clearly, could you please repeat?

Issue: User utterance 'tell me about' was incomplete due to ASR error. The model did not recognize the error and focused on the dialog topic instead which was movies.

Example 2

Dialog History:

System: so, tell me, what was the last movie you liked?

User: i don't remember

Model selected response: i don't know this movie. space odyssey 2001 is great, in my opinion. it is a long-standing classic, one of the first films about artificial intelligence. tell me, what movie genres do you prefer?

A correct response: i have found some popular new movies: uncut gems and the rest of us. do you want to know details about uncut gems?

Issue: The model assumes the user fulfilled the system's question by providing a movie even though the user didn't.

Figure 5: Example predictions of BERT Ranker (RSD Train). Due to privacy concerns, these example dialogs are from an internal author.

5.3 Analysis

The advantage of creating negative examples via random or synthetic approaches is the ability to automatically increase the number of training examples. To further evaluate this, we vary the number of negative candidates in Synthetic Inter-Random, Intra-Random, and Adversarial, and report the corresponding MRR and Recall@1 scores, in Figure 3. We see that for our Synthetic Datasets increasing the number of negative candidates to a certain point improves performance for both MRR and Recall@1, after which the performance will degrade.

Increasing the number of negative candidates for (Synthetic Inter-Random) and (Synthetic Intra-Random) increases the likelihood of retrieving a candidate that is a false negative. This can bring noise and confusion to the model during training time. Increasing the size of the corpus could mitigate this issue; however, it can be expensive to collect a large enough dataset to see its benefits.⁶ The advantage of (Synthetic Adv.) is the ability to create a large number of negative candidates without collecting more data; however, as seen in Figure 3 the decrease in MRR and Recall@1 when sampling

⁶Large datasets such as Reddit are known to be noisy and could degrade performance.

more candidates may be due to false negatives and therefore still need to be manually verified.

5.4 Qualitative Examples

We provide examples of our BERT-Ranker models in Figure 4. In Example 1 both responses selected by the models acknowledge the user's artist preference; however, BERT-Ranker (Synthetic Inter-Random) chooses a response that repeats the question already answered by the user while BERT-Ranker (RSD Train) does not. In Example 2, BERT-Ranker (RSD Train) provides a more coherent response versus BERT-Ranker (Synthetic Adversarial) which has an abrupt topic change. In Example 3, BERT Ranker (Synthetic Retrieval) repeats the same question asked in the dialog history.

Figure 5 shows two typical erroneous examples. In the examples we also provide an explanation for the errors. It is worth pointing out that incorrect ASR output (word errors or end point detection errors such as the first example) is a source of errors to confuse our models. Gopalakrishnan et al. (2020) has observed similar issues for the task of response generation in speech-based dialog systems. Future work such as training on synthetic/actual ASR errors is needed to improve the robustness of models for such ASR issues.

5.5 Limitations

Our evaluation is done on a dialog dataset that contains a limited number of responders and only GPT2 is used as a neural response generation model. Synthetically created examples may perform better on datasets with a wider variety of neural response generation models. Future work would involve collecting response selection data annotated with a wider variety of responders.

6 Conclusion

In this work, we have curated a new dataset for response selection, which contains multiple positive responses and human verified hard negatives. We conducted a comprehensive evaluation of SOTA response selection models and various techniques to construct negative candidates to demonstrate the benefit of the dataset. Even though RSD requires manual annotation we see that training on our dataset greatly outperforms methods that use only one positive example and generate adversarial negative candidates.

7 Ethics and Broader Impact

Our work involves re-ranking responses from a dialog system. We acknowledge that we are using data from real users who have not been paid for these interactions. We also acknowledge there may be biases in the demographics of the user population.

References

- 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.
- Chris Donahue, Mina Lee, and Percy Liang. 2020. Enabling language models to fill in the blanks. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2492–2501.
- Raefer Gabriel, Yang Liu, Anna Gottardi, Mihail Eric, Anju Khatri, Anjali Chadha, Qinlang Chen, Behnam Hedayatnia, Pankaj Rajan, Ali Binici, et al. Further advances in open domain dialog systems in the third alexa prize socialbot grand challenge.
- Xiang Gao, Yizhe Zhang, Michel Galley, Chris Brockett, and William B Dolan. 2020. Dialogue response ranking training with large-scale human feedback data. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 386–395.
- Karthik Gopalakrishnan, Behnam Hedayatnia, Longshaokan Wang, Yang Liu, and Dilek Hakkani-Tür. 2020. Are neural open-domain dialog systems robust to speech recognition errors in the dialog history? an empirical study.
- Jia-Chen Gu, Tianda Li, Quan Liu, Zhen-Hua Ling, Zhiming Su, Si Wei, and Xiaodan Zhu. 2020. Speaker-aware bert for multi-turn response selection in retrieval-based chatbots. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, pages 2041–2044.
- Prakhar Gupta, Shikib Mehri, Tiancheng Zhao, Amy Pavel, Maxine Eskenazi, and Jeffrey P Bigham. 2019. Investigating evaluation of open-domain dialogue systems with human generated multiple references. *arXiv preprint arXiv:1907.10568*.
- Prakhar Gupta, Yulia Tsvetkov, and Jeffrey P Bigham. 2021. Synthesizing adversarial negative responses for robust response ranking and evaluation. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 3867–3883.
- Janghoon Han, Taesuk Hong, Byoungjae Kim, Youngjoong Ko, and Jungyun Seo. 2021. [Fine-grained post-training for improving retrieval-based dialogue systems](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1549–1558, Online. Association for Computational Linguistics.
- Behnam Hedayatnia, Karthik Gopalakrishnan, Seokhwan Kim, Yang Liu, Mihail Eric, and Dilek Hakkani-Tur. 2020. Policy-driven neural response generation for knowledge-grounded dialog systems. In *Proceedings of the 13th International Conference on Natural Language Generation*, pages 412–421.
- Huda Khayrallah and João Sedoc. 2020. Smrter chatbots: Improving non-task-oriented dialog with simulated multi-reference training. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings*, pages 4489–4505.
- Gary King and Langche Zeng. 2001. Logistic regression in rare events data. *Political analysis*, 9(2):137–163.
- Jakub Konrád, Jan Pichl, Petr Marek, Petr Lorenc, Van Duy Ta, Ondřej Kobza, Lenka Hýlová, and Jan Šedivý. 2021. Alquist 4.0: Towards social intelligence using generative models and dialogue personalization. *arXiv preprint arXiv:2109.07968*.
- Jia Li, Chongyang Tao, Wei Wu, Yansong Feng, Dongyan Zhao, and Rui Yan. 2019. [Sampling matters! an empirical study of negative sampling strategies for learning of matching models in retrieval-based dialogue systems](#). 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 1291–1296, Hong Kong, China. Association for Computational Linguistics.
- Yanran Li, Hui Su, Xiaoyu Shen, Wenjie Li, Ziqiang Cao, and Shuzi Niu. 2017. Dailydialog: A manually labelled multi-turn dialogue dataset. In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 986–995.
- Ryan Lowe, Nissan Pow, Iulian Vlad 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.
- Stefano Mezza, Alessandra Cervone, Evgeny Stepanov, Giuliano Tortoreto, and Giuseppe Riccardi. 2018. Iso-standard domain-independent dialogue act tagging for conversational agents. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 3539–3551.
- Masahiro Mizukami, Hideaki Kizuki, Toshio Nomura, Graham Neubig, Koichiro Yoshino, Sakriani Sakti,

- Tomoki Toda, and Satoshi Nakamura. 2015. Adaptive selection from multiple response candidates in example-based dialogue. In *2015 IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU)*, pages 784–790. IEEE.
- Ioannis Papaioannou, Amanda Cercas Curry, Jose L Part, Igor Shalyminov, Xinnuo Xu, Yanchao Yu, Ondrej Dušek, Verena Rieser, and Oliver Lemon. 2017. Alana: Social dialogue using an ensemble model and a ranker trained on user feedback.
- Ashwin Paranjape, Abigail See, Kathleen Kenealy, Haojun Li, Amelia Hardy, Peng Qi, Kaushik Ram Sadagopan, Nguyet Minh Phu, Dilara Soylu, and Christopher D Manning. 2020. Neural generation meets real people: Towards emotionally engaging mixed-initiative conversations. *arXiv preprint arXiv:2008.12348*.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.
- Ashwin Ram, Rohit Prasad, Chandra Khatri, Anu Venkatesh, Raefer Gabriel, Qing Liu, Jeff Nunn, Behnam Hedayatnia, Ming Cheng, Ashish Nagar, et al. 2018. Conversational ai: The science behind the alexa prize. *arXiv preprint arXiv:1801.03604*.
- Sougata Saha, Souvik Das, Elizabeth Soper, Erin Pacquetet, and Rohini K Srihari. 2021. Proto: A neural cocktail for generating appealing conversations. *arXiv preprint arXiv:2109.02513*.
- Ananya B Sai, Akash Kumar Mohankumar, Siddhartha Arora, and Mitesh M Khapra. 2020. Improving dialog evaluation with a multi-reference adversarial dataset and large scale pretraining. *Transactions of the Association for Computational Linguistics*, 8:810–827.
- Abigail See and Christopher Manning. 2021a. [Understanding and predicting user dissatisfaction in a neural generative chatbot](#). In *Proceedings of the 22nd Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 1–12, Singapore and Online. Association for Computational Linguistics.
- Abigail See and Christopher D Manning. 2021b. Understanding and predicting user dissatisfaction in a neural generative chatbot. In *Proceedings of the 22nd Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 1–12.
- Iulian V Serban, Chinnadhurai Sankar, Mathieu Germain, Saizheng Zhang, Zhouhan Lin, Sandeep Subramanian, Taesup Kim, Michael Pieper, Sarath Chandar, Nan Rosemary Ke, et al. 2017. A deep reinforcement learning chatbot. *arXiv preprint arXiv:1709.02349*.
- Igor Shalyminov, Ondřej Dušek, and Oliver Lemon. 2018. Neural response ranking for social conversation: A data-efficient approach. In *Proceedings of the 2018 EMNLP Workshop SCAI: The 2nd International Workshop on Search-Oriented Conversational AI*, pages 1–8.
- Eric Michael Smith, Mary Williamson, Kurt Shuster, Jason Weston, and Y-Lan Boureau. 2020. Can you put it all together: Evaluating conversational agents’ ability to blend skills. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2021–2030.
- Wenhui Wang, Furu Wei, Li Dong, Hangbo Bao, Nan Yang, and Ming Zhou. 2020. Minilm: Deep self-attention distillation for task-agnostic compression of pre-trained transformers. *Advances in Neural Information Processing Systems*, 33:5776–5788.
- Taesun Whang, Dongyub Lee, Chanhee Lee, Kisu Yang, Dongsuk Oh, and Heuseok Lim. 2020. An effective domain adaptive post-training method for bert in response selection. In *Interspeech*.
- Taesun Whang, Dongyub Lee, Dongsuk Oh, Chanhee Lee, Kijong Han, Dong-hun Lee, and Saebyeok Lee. 2021. Do response selection models really know what’s next? utterance manipulation strategies for multi-turn response selection.
- Thomas Wolf, Victor Sanh, Julien Chaumond, and Clement Delangue. 2019. Transfertransfo: A transfer learning approach for neural network based conversational agents. *arXiv preprint arXiv:1901.08149*.
- Yu Wu, Wei Wu, Chen Xing, Ming Zhou, and Zhoujun Li. 2017. Sequential matching network: A new architecture for multi-turn response selection in retrieval-based chatbots. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 496–505.
- Ruijian Xu, Chongyang Tao, Daxin Jiang, Xueliang Zhao, Dongyan Zhao, and Rui Yan. 2020. Learning an effective context-response matching model with self-supervised tasks for retrieval-based dialogues. *arXiv preprint arXiv:2009.06265*.
- Saizheng Zhang, Emily Dinan, Jack Urbanek, Arthur Szlam, Douwe Kiela, and Jason Weston. 2018a. Personalizing dialogue agents: I have a dog, do you have pets too? In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2204–2213.
- Yichi Zhang, Zhijian Ou, and Zhou Yu. 2020. Task-oriented dialog systems that consider multiple appropriate responses under the same context. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 9604–9611.
- Yizhe Zhang, Siqi Sun, Michel Galley, Yen-Chun Chen, Chris Brockett, Xiang Gao, Jianfeng Gao, Jingjing Liu, and Bill Dolan. 2019. DialoGPT: Large-scale generative pre-training for conversational response generation. *arXiv preprint arXiv:1911.00536*.

- Zhuosheng Zhang, Jiangtong Li, Pengfei Zhu, and Hai Zhao. 2018b. Modeling multi-turn conversation with deep utterance aggregation. In *Proceedings of the 27th International Conference on Computational Linguistics (COLING 2018)*, pages 3740–3752.
- Zhuosheng Zhang and Hai Zhao. 2021. Structural pre-training for dialogue comprehension. *arXiv preprint arXiv:2105.10956*.
- Li Zhou, Jianfeng Gao, Di Li, and Heung-Yeung Shum. 2020. The design and implementation of xiaoice, an empathetic social chatbot. *Computational Linguistics*, 46(1):53–93.
- Naitian Zhou and David Jurgens. 2020. Condolences and empathy in online communities. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 609–626.

A Response Selection Model Training Details

All our BERT-base (Devlin et al., 2019) models are trained with a batch size of 32 on 1 NVIDIA V100 GPU with 16GB memory. We use the Adam optimizer with a learning rate of $1e-5$ and the model is trained for 2 epochs. We use a sequence length of 256 tokens. To deal with the label imbalance, we compute a weighted loss where the loss for a positive candidate is up-weighted by a factor of α and the loss for a negative candidate is down-weighted by a factor of β . We follow (King and Zeng, 2001) and compute α by taking the sum of the number of positive and negative candidates and divide by the number of labels times the number of positive candidates. The same is done for β but we divide by the number of negative candidates instead. In our experiments $\alpha = 5.35$ and $\beta = 0.55$.

For the DialogRPT-human vs ranker model, we train with a batch size of 4 on 8 NVIDIA V100 GPUs with 16GB memory each. We use the Adam optimizer with a learning rate of $3e-5$. We use a sequence length of 50 tokens and the model is trained for 3 epochs.

B NRG Training Details

We train all our NRG models on the RUI dataset described in Section 3.2. This dataset is split into a 90/10/10 train, valid, test split. All of our models are initialized with GPT2 (Radford et al., 2019) based models and were trained with a batch size of 2 on 8 NVIDIA A100 GPUS with 32GB memory each. We use the Adam optimizer and a learning rate of $6.25e-5$. Each model is trained for 3 epochs and we finetune both the Language Modeling Head and Multiple Choice Head of GPT2 in a Transfer-Transfo fashion (Wolf et al., 2019). The Multiple Choice Head is finetuned with 1 randomly selected negative candidate. We leverage the HuggingFace’s transformers library for all our models.¹ Detailed descriptions of our NRG variants are provided as below.

NRG Responder: Is a GPT2-XL model where the input is the dialog context which is truncated to 64 tokens.

NRG Responder GPT2-medium: Is a GPT2-medium model where the input is the dialog context which is truncated to 64 tokens.

NRG Responder grounded on knowledge: Is a

GPT2-XL model where the dialog context is truncated to 256 tokens and a single knowledge sentence is truncated to 32 tokens. The dialog context and knowledge sentence are concatenated together to be used as input into the model.

NRG Responder grounded on dialog acts (DA): Is a GPT2-XL model where the dialog context is truncated to 64 tokens and each dialog act has its own embedding that is randomly initialized and updated during finetuning. The dialog context and DA are concatenated together to be used as input into the model. When training this model we automatically label the RUI dataset with a dialog act tagger² and use those DAs as the ground truth. The DA labels used are from (Mezza et al., 2018) e.g. Feedback, Yes-No question, Statement.

During inference, a sequence of dialog acts are determined using a rule-based dialog policy which are used as input into the model to control the generated response. For example, a Yes-No question dialog act will cause the model response to generate a question (Hedayatnia et al., 2020).

NRG Responder grounded on sentiment: Is a GPT2-XL model where the dialog context is truncated to 64 tokens. There is an embedding representing negative sentiment that is randomly initialized and updated during finetuning. The dialog context and negative sentiment are concatenated together to be used as input into the model to control the generated response. This controllability allows the model is able to generate a sympathetic response when the user expresses negative sentiment. When training such a model, we automatically label the RUI dataset with an off the shelf sentiment classifier (Zhou and Jurgens, 2020) and use those sentiment tags as the ground truth.

C Response Selection Annotation Details

Our annotation framework is shown in Figure C.1. A human annotator is shown a dialog context and a set of response candidates are shown below. The annotator can then check off however many responses they deem as appropriate with respect to the dialog context. All responses not selected are considered inappropriate.

¹<https://github.com/huggingface/transformers>

²We annotate a subset of the RUI dataset for dialog acts and train an RNN model on these annotations

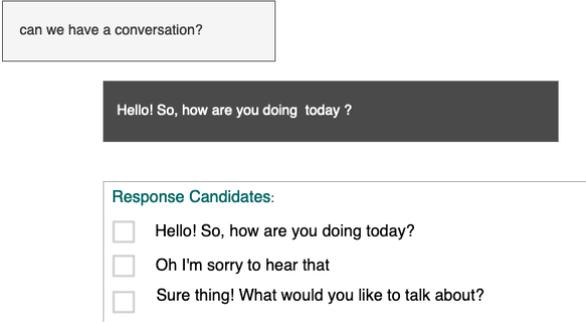


Figure C.1: Annotation framework to collect RSD Train/Test. Due to privacy concerns, this example dialog is from an internal author.

D Evaluation Metrics for Synthetic Datasets

In Figures D.2, D.3 and D.4 we show all the metrics for our BERT Ranker model trained on each of our Synthetic Datasets with different number of sampled negative candidates. We see for all metrics as the number of negative candidates increase results either degrade or taper off.

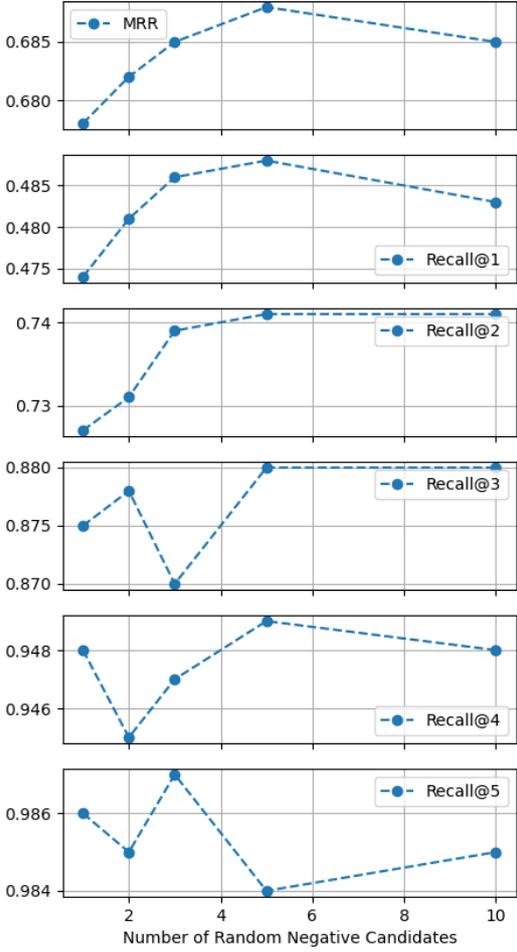


Figure D.2: Increasing the number of randomly sampled negative candidates in the Synthetic Inter-Random Training set.

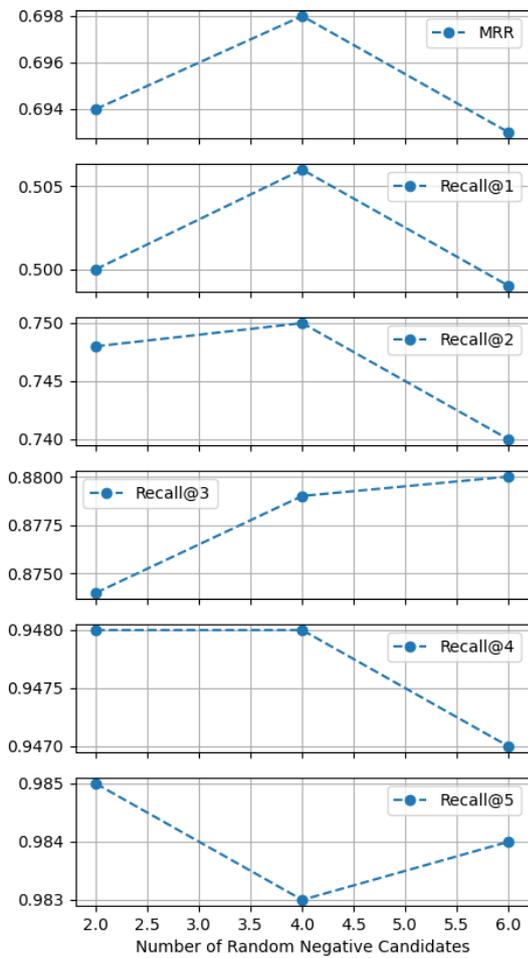


Figure D.3: Increasing the number of randomly sampled negative candidates in the Synthetic Intra-Random Training set.

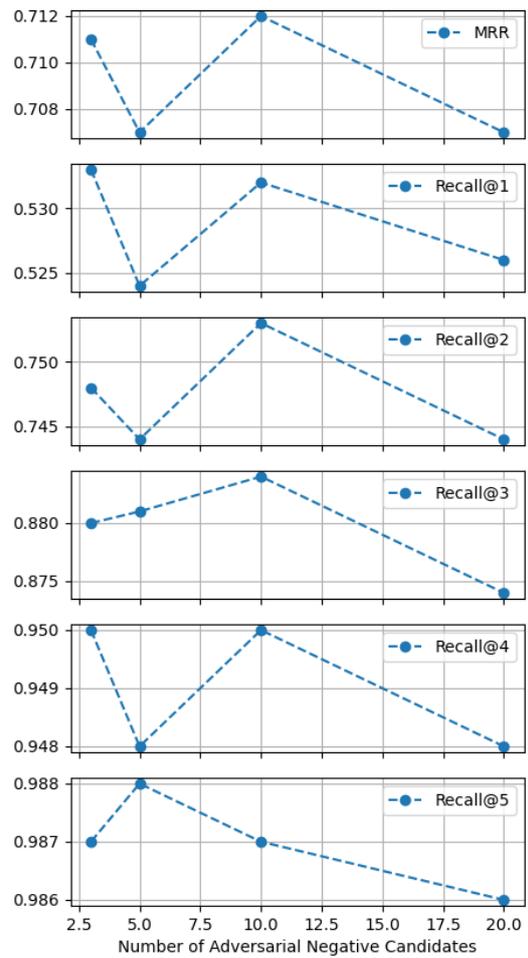


Figure D.4: Increasing the number of randomly sampled negative candidates in the Synthetic Adversarial Training set.