

An Effective Domain Adaptive Post-Training Method for BERT in Response Selection

Taesun Whang^{1*} Dongyub Lee² Chanhee Lee³ Kisu Yang³ Dongsuk Oh³ Heuseok Lim³

¹Wisnut Inc.

²Kakao Corp.

³Korea University

taesunwhang@wisnut.co.kr, jude.lee@kakaocorp.com,
{chanhee0222, willow4, inow3555, limhseok}@korea.ac.kr

Abstract

We focus on multi-turn response selection in a retrieval-based dialog system. In this paper, we utilize the powerful pre-trained language model Bi-directional Encoder Representations from Transformer (BERT) for a multi-turn dialog system and propose a highly effective post-training method on domain-specific corpus. Although BERT is easily adopted to various NLP tasks and outperforms previous baselines of each task, it still has limitations if a task corpus is too focused on a certain domain. Post-training on domain-specific corpus (*e.g.*, Ubuntu Corpus) helps the model to train contextualized representations and words that do not appear in general corpus (*e.g.*, English Wikipedia). Experimental results show that our approach achieves new state-of-the-art on two response selection benchmarks (*i.e.*, Ubuntu Corpus V1, Advising Corpus) performance improvement by 5.9% and 6% on $R_{10}@1$.

Index Terms: Response selection, Human computer dialog system, Spoken language processing

1. Introduction

Human computer conversation system aims to have natural and consistent conversation. In general, dialog systems can be categorized into task-oriented dialog systems and non-task-oriented dialog systems (*i.e.*, chatbot systems). In this paper, we focus on one of non-task-oriented dialog systems, especially retrieval-based dialog systems, in which the task of predicting most likely next utterance from a set of candidates pool. While generative-based models provide a consistent but inaccurate response, retrieval-based models have a advantage of ensuring accurate answer.

Previous studies of response selection utilized recurrent neural networks [1] or convolutional neural networks [2] to represent dialog context and response, and predict their relevance score based on their representations. While these models represent dialog as a sequential sentence, utterance-response matching models build utterance-level encodings and utilize attention mechanism to catch relevant words between two sentences [3, 4]. Although attention-based utterance-response matching models are highly effective for predicting relevant response, self-attention-based matching model based on the work of Vaswani et al. [5] strengthens encoding sophisticated segment representations [6, 7, 8]. Also, a model based on contextualized language representations is also applied to response selection task in the work of Vig and Ramea [9] and showed effectiveness of the pre-trained contextual language models.

Recently, pre-trained contextualized language models, such as ELMo [10], BERT [11], and XLNet [12]) are proposed and showed great improvements on a wide range of NLP tasks, such as natural language inference, named entity recognition, and question answering. Despite their huge success, they still have limitations to represent contextual information in dialog corpus, specifically domain-specific corpus, since it is trained on general corpora (*e.g.*, English Wikipedia and Book Corpus). For example, Ubuntu Corpus, which is the most used corpus for evaluating retrieval-based dialog system, contains a number of terminologies and ubuntu manuals that do not usually appear in general corpora (*e.g.*, apt-get, lsmod, and grep). Since the corpus is biased towards a certain domain, pre-trained contextualized language models are not able to fully represent dialog contexts. In addition, conversation corpus, such as Twitter and Reddits, is mainly composed of colloquial expressions and abbreviation which are usually grammatically incorrect. In response selection task, one approach of training a certain domain knowledge embeddings is proposed. Chaudhuri et al. [13] proposed a method building domain understandable embeddings by incorporating external knowledge (*e.g.*, ubuntu manual description). In other NLP task, such as aspect extraction and review reading comprehension (RRC), there has been an attempt to learn domain specificity based on pre-trained model. Xu et al. [14] proposed BERT based post-training method for RRC to enhance domain-awareness. Since reviews and opinion-based texts have many differences compared to the original corpus of BERT, therefore, post-training of BERT with two powerful unsupervised objectives (*i.e.*, masked language model (MLM) and next sentence prediction (NSP)) on task-specific corpus enhance to produce task-awareness contextualized representations. Also, additional training on down-stream corpus shows effectiveness and performance improvements in various NLP tasks [15, 16].

In this work, we propose an effective post-training method for a multi-turn conversational system. To the best of our knowledge, it is the first attempt to adopt BERT_{base} on the most popular response selection benchmark, Ubuntu Corpus V1. We demonstrate that NSP is especially considered as an important task for response selection, since classifying whether given two sentences are *IsNext* or *NotNext* is the ultimate objective of response selection. Also, we append [EOT] to the end of utterance that model can learn relationships among utterances during the period of post-training. Furthermore, our approach outperforms previous state-of-the-art performance by 5.9% on $R_{10}@1$. We also evaluate on the recently released data set in the dialog system technology challenges 7 (DSTC 7) [17] and outperforms the 1st place of the challenges.

* Work performed while at Korea University

2. Related Work

Lowe et al. [1] proposed a new response selection benchmark dataset, Ubuntu Corpus V1, with a dual encoder baseline model. They utilized a Ubuntu based models, specifically a vanilla RNN, a long short-term memory (LSTM), and a bi-directional LSTM. The sequential representations of a dialog context and a response are encoded by the RNN-based encoder and are utilized to predict the probability score of a given response being the next utterance in a given dialogue context. Kadlec et al. [2] conducted experiments that apply convolutional neural networks to a dual encoder architecture that offers better performance than that of the vanilla RNN model. Even though each turn of the dialog should be considered in a multi-turn conversation, a given dialog context was regarded as a sequence in the previous baselines. To alleviate this, Zhou et al. [18] proposed the MultiView model that encodes both word-level and utterance-level representations. However, it does not have the ability to fully capture the relevance of the dialog context and the response.

Therefore, more recently, the utterance-level dialog-response matching models were proposed to enable the model to readily catch the relevance of the dialog and the response [3, 4]. Wu et al. [3] proposed a sequential matching network, utilizing both word embeddings and their sequential representations encoded by gated recurrent unit (GRU) encoder to build matching matrices between the dialogue context and the response. Zhang et al. [4] proposed a highly effective turns-aware aggregation methodology and exercised self-matching attention to fuse the representation of each utterance. Another approach to the dialogue-response matching model proposed by Dong and Huang (2018), is designed based on enhanced sequential inference model (ESIM) [19] model. ESIM achieved a good performance on a natural language inference task and is a proper model for discriminating if given two sentences are relevant or not. As a self-attention-based model (Vaswani et al., 2017) achieved significant performance improvements on various NLP tasks; it was also adapted for the response selection task [6, 7, 8]. Zhou et al. [6] used self-attention and cross-attention simultaneously so as to capture token-level dependencies as well as relevant segment pairs from a dialog and a response. One of most recent models, interaction-over-interaction model alleviates shallow interaction between dialog-response matching by stacking multiple interaction blocks [20]. Yuan et al. [8] pointed out previous dialog-response matching models in terms of excessive use of dialog context information, therefore they proposed multi-hop selector to filter only relevant utterances to response. They achieved state-of-the-art results on three response selection benchmark datasets (*i.e.*, Ubuntu Corpus V1, Douban Corpus, and E-commerce).

3. Our Approach

3.1. Domain Post-Training

BERT is designed to be easily applied to other nlp tasks with a fine-tuning manner. Since it is pre-trained on general corpus (*e.g.*, Wikipedia Corpora), it is insufficient to have enough supervision of task-specific words and phrases during the period of fine-tuning. To alleviate this issue, we post-train BERT on our task-specific corpora that helps model understand certain domain. The model is trained with two objectives, MLM and NSP, which are highly effective to learn contextual representations from the corpora. One example (Ubuntu Corpus) of domain post-training of BERT for response selection is described in Table 1.

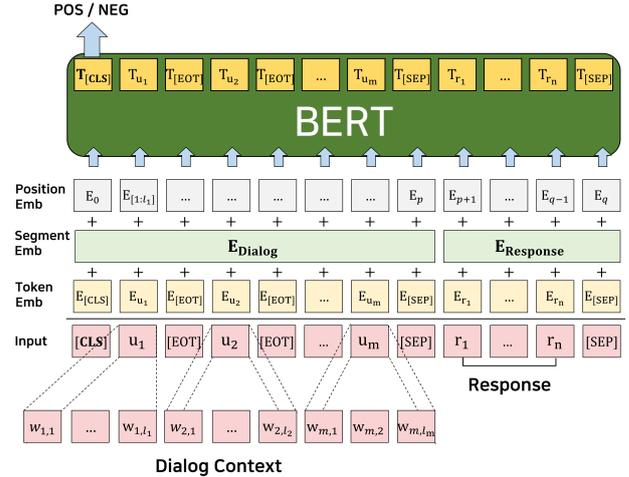


Figure 1: BERT for Response Selection

In the example of Masked LM, model can learn that `sudo` command is needed when trying `apt` install in Ubuntu system, which is not generally showed from universal corpora. Unlike general sentence, multi-turn dialog system is composed of a set of utterances. We append “End Of Turn” token [EOT] to the end of each turn to make the model catch each utterance is finished at the point. By conducting NSP during the post-training, model also can train given two sentences are sequential and relevant, which is the common ultimate goal of response selection. To optimize the model domain post-training (DPT) loss is calculated adding mean MLM likelihood and that of NSP, formulated as

$$\mathcal{L}_{DPT} = \mathcal{L}_{MLM} + \mathcal{L}_{NSP}. \quad (1)$$

[Masked Language Model]

Input : `sudo ##o apt - get install cuda`

\Rightarrow [MASK] ##o apt - get install cuda

[Next Sentence Prediction]

Input : `how do I [MASK] cuda library
in terminal ? [EOT] [SEP] (Text A)`

`You can search with apt - cache
search [MASK] [EOT] [SEP] (Text B)`

Label : IsNext

Table 1: An example of domain post-training in Ubuntu Corpus

3.2. Fine-tuning BERT for Response Selection

Our overall approach is described in Figure 1. We transform the task of ranking responses from the candidates pool into binary classification problem by approaching a pointwise method. We denote a training set as a triples $\mathcal{D} = \{(c_i, r_i, y_i)\}_{i=1}^N$, where $c = \{u_1, u_2, \dots, u_m\}$ is a dialog context consists of a set of utterances. An utterance $u_i = \{w_{i,1}, w_{i,2}, \dots, w_{i,l_i}\}$ is composed of a set of word tokens $w_{i,k}$, where $1 \leq k \leq l_i$ and l_i is the length of i -th utterance. Response is denoted as $r_i = \{r_1, r_2, \dots, r_n\}$ (n is the number of tokens in the response), and ground truth $y_i \in \{0, 1\}$. We define BERT input $\mathbf{x} = ([CLS], u_1, [EOT], u_2, [EOT], \dots, u_m, [SEP], r_1, \dots, r_n, [SEP])$. Maximum sequence length of the dialog context

and response is denoted as q . Position, segment, and token embeddings are added and fed into the BERT layers. The BERT contextual representations of [CLS] token, $\mathbf{T}_{[CLS]}$, is utilized to classify whether a given dialog context and response is `IsNextUtterance` or not. We feed $\mathbf{T}_{[CLS]}$ to single-layer perceptron to compute the model prediction score,

$$g(c, r) = \sigma(\mathbf{W}_{task}^\top \mathbf{T}_{[CLS]} + b), \quad (2)$$

where \mathbf{W}_{task} is a task-specific trainable parameter. We use cross entropy loss as the objective function to optimize our model, formulated as

$$\mathcal{L}_{loss} = - \sum_{(c,r,y) \in \mathcal{D}} y \log(g(c, r)) + (1 - y) \log(1 - g(c, r)). \quad (3)$$

4. Experiments

4.1. Datasets and Training Setup

We evaluate our model on two multi-turn dyadic data sets, Ubuntu IRC (Internet Relay Chat) Corpus V1 [1] and Advising Corpus¹ [17]. For the Ubuntu Corpus, training set is composed of 0.5M dialog context containing positive and negative response with the ratio of 1:1. Each validation and test set contains 50k dialog context with one positive response and nine negative responses. Advising Corpus consists of 100k dialogs for training set and 500 for validation and test set. All sets contain one positive response and 99 negative responses. We only use one negative sample for training to make same conditions with Ubuntu Corpus. For an evaluation metric, we use $R_n@k$, evaluating if the ground truth exists in top k from n candidates [1, 3]. We also use another evaluation metric mean reciprocal rank (MRR) [21].

4.2. Implementation Details

The models are implemented using the Tensorflow library [22]. We use the uncased BERT_{base} model² as a base code for our experiments. The batch size is set to 32 and the maximum sequence length is set to 320, specifically 280 for a dialog context and 40 for a response. We post-train the model more on Ubuntu Corpus V1 and Advising Corpus, 200,000 steps and 100,000 steps, respectively. The model is optimized using Adam weight decay optimizer with learning rate of 3e-5.

4.3. BERT Models

BERT_{base} [11] is a vanilla BERT_{base} model for response selection task. Uncased BERT_{base} model checkpoint is utilized as initial weights of the model, and single-layer perceptron are trained during the fine-tuning phase.

BERT-DPT is a BERT_{base} model that is post-trained on domain-specific corpus, Ubuntu Corpus V1 and Advising Corpus. The unsupervised objectives of the BERT model, MLM and NSP tasks are conducted during the DPT phase.

BERT-VFT is a model that selects the number of top layers to fine-tuning in BERT-DPT. We evaluate the model with varying $T = \{0, 2, 4, 6, 8, 10, 12\}$, where T is the number of layers which are tuned during training time. When $T = 12$, all the BERT-DPT layers are fine-tuned and only embedding layers are fine-tuned when $T = 0$. We fill in the best score based on MRR

¹<https://github.com/IBM/dstc-noesis>

²<https://github.com/google-research/bert>

in Table 4.

BERT-VFT(DA) performs data augmentation technique by increasing the number of negative samples. The number of negative samples are decided by several experiments, when the number is 4 shows the best results. Negative responses are randomly resampled for every epoch from response candidates pool.

Model	$R_{10}@1$	$R_{10}@2$	$R_{10}@5$
DualEncoder _{rnn}	0.403	0.547	0.819
DualEncoder _{cnn}	0.549	0.684	0.896
DualEncoder _{lstm}	0.638	0.784	0.949
DualEncoder _{bilstm}	0.630	0.780	0.944
MultiView	0.662	0.801	0.951
SMN	0.726	0.847	0.961
AK-DE-biGRU	0.747	0.868	0.972
DUA	0.752	0.868	0.962
DAM	0.767	0.874	0.969
MRFN	0.786	0.886	0.976
IoI	0.796	0.894	0.974
MSN	<u>0.800</u>	<u>0.899</u>	<u>0.978</u>
BERT _{base}	0.817	0.904	0.977
BERT-DPT	0.851	0.924	0.984
BERT-VFT	0.855	0.928	0.985
BERT-VFT(DA)	0.858	0.931	0.985

Table 2: Model comparison on Ubuntu Corpus V1.

4.4. Performance on Ubuntu Corpus V1

For Ubuntu Corpus V1, we compared our model with the following baseline methods and the results are given in Table 2.

Dual Encoder is simple dialogue-response matching model based on RNN, CNN, LSTM, and BiLSTM [2]. Each dialog context and response is represented by each encoder and utilized to obtain relevance score of given two sentences.

MultiView model is based on both token-level and utterance-level representations to catch utterance-level information in a dialog. CNN encoder is firstly used to encode each utterance representation, and Gated recurrent unit (GRU) encoder is used to encode both token-level representations and utterance-level representations [18].

SMN [3] proposed utterance-response matching methods, specifically attention matrices created by word embeddings and sequential representations are encoded by CNN and they are fed into GRU encoder to obtain a probability score.

AK-DE-biGRU [13] proposed a method incorporating domain knowledge (*i.e.*, Ubuntu manual description). Domain knowledge embeddings are created based on manual description, BiGRU encoder’s forward and backward last hidden states are concatenated to build the embeddings. Pre-trained word embeddings and created domain knowledge embeddings are added and fed into an attention module.

DUA [4] is a method of modeling utterance aggregation. By conducting turns-aware aggregation and gated self attention, the model can give weights to more relevant utterance with response.

DAM [6] is based on transformer [5] encoder and apply both self-attention and cross-attention to obtain matching scores. Each attention matching score is aggregated by 3D matching image.

MRFN [20] proposed multi representation fusion network and highlighted the effectiveness of fusing strategy. They proposed a methodology to fuse multiple representations of words, con-

texts, and attention.

IoI [7] build interaction blocks to help model conduct deep interactions between utterances and response. All representations extracted from the blocks are aggregated to obtain a relevance score.

MSN [8] is a previous state-of-the-art model, which proposed multi-hop selector network to control which dialog context information is reflected for matching response.

In Ubuntu Corpus V1, it is observed that BERT-VFT achieves new state-of-the-art performance and it obtains performance improvement by 5.5%, 2.9%, 0.7% in terms of $R_{10}@k$, where $k=\{1, 2, 5\}$ respectively, compared to the previous state-of-the-art method (*i.e.*, MSN). Focusing on $R_{10}@1$ metric, the performance of a vanilla BERT_{base} is 0.817 and comparing our main approach, which is BERT-DPT, shows better result (improvement by 3.4%).

In addition, we especially point out comparing BERT-VFT with other self-attention-based models, such as DAM, IoI, and MSN, since all models are built with similar model architecture. Especially, domain-specific optimized BERT-VPT model shows performance improvement by 8.8% in terms of $R_{10}@1$ compared to the general transformer based model (*i.e.*, DAM). In the case of MSN, even though it uses word2vec [23] as pre-trained word embeddings, it shows comparable results with BERT_{base} model. However, BERT-VFT that understands domain specific corpus shows 5.5% higher results in terms of $R_{10}@1$, which has significant gap between two models.

4.5. Performance on Advising Corpus

As shown in Table 3, we compare our approach with two existing baselines on Advising Corpus, proposed by Vig and Remea [9] and Chen et al. [24] in DSTC 7. The former baseline evaluate BERT_{base} model on the Advising Corpus, but there is substantial performance difference from what we obtain. We believe that different implementation frameworks and hyperparameters, such as learning rates, are the main reason why performance differences exist between our work and that of Vig and Remea [9]. The first place of the challenge was the work conducted by Chen et al. [24], BERT-VFT outperforms 6% in terms of $R_{10}@1$. DPT shows its effectiveness on Advising Corpus, leading 3.4% higher results in terms of $R_{10}@1$ compared to the BERT baseline.

Model	$R_{100}@1$	$R_{100}@10$	$R_{100}@50$	MRR
Vig and Remea [9]	0.186	0.580	0.942	0.312
Chen et al. [24]	0.214	0.630	0.948	0.339
BERT _{base}	0.236	0.656	0.946	0.359
BERT-DPT	0.270	0.668	0.942	0.395
BERT-VFT	0.274	0.654	0.932	0.400
BERT-VFT(DA)	0.274	0.664	0.942	0.399

Table 3: Evaluation results on the Advising Corpus.

4.6. Variable Fine-tuning

Houlsby et al. [25] proposed efficient fine-tuning approach that only updates a few top layers of BERT during fine-tuning period. They suggested that small data sets may be sub-optimal, when fine-tuning the whole BERT layers. Inspired by this assumption, we conduct an experiment (Table 4) varying T layers which is fine-tuned during the period of training ($T=\{0, 2, 4, 6, 8, 10, 12\}$). The models achieve the best performance when $T=4$ and $T=6$, for Ubuntu and Advising, respectively. We demonstrate that utilizing this application is

effective on not only small sets but also domain-specific sets. Variable fine-tuning also helps for reducing computational resources (*e.g.*, GPU memory) as well as showing performance improvements.

Variable Fine-Tuning	Ubuntu	Advising
All layers	0.907	0.395
Top 10 layers	0.908	0.393
Top 8 layers	0.909	0.389
Top 6 layers	0.909	0.400
Top 4 layers	0.910	0.392
Top 2 layers	0.907	0.387
Only Embedding Layers	0.904	0.386

Table 4: Variable fine-tuning in BERT model. All models are evaluated on MRR metric.

4.7. MLM and NSP in BERT

Table 5 shows the effectiveness of two unsupervised objectives of BERT, MLM and NSP in Ubuntu Corpus V1. Also, it shows that appending special token [EOT] at the end of each utterance is highly effective for response selection task. During the DPT phase, MLM helps the model learn contextual representations from the target domain corpus and it is trained sequential consistency of the given two sentences by conducting NSP.

In the case of NSP, it does not influence much in improvement of performance because what NSP do during the DPT phase is the same task as the model do during the fine-tuning phase. While conducting only NSP is not helpful for corpus understanding, MLM leads significant improvement in performance regardless of using the special token.

Post-Training	Special Token	$R_{10}@1$	$R_{10}@2$	$R_{10}@5$	MRR
NSP	without	0.816	0.903	0.977	0.884
MLM	without	0.834	0.916	0.981	0.896
MLM + NSP	EOT	0.839	0.920	0.982	0.900
NSP	with	0.819	0.906	0.978	0.886
MLM	with	0.838	0.918	0.982	0.899
MLM + NSP	EOT	0.851	0.924	0.984	0.907

Table 5: Comparison of MLM and NSP on Ubuntu Corpus V1. Experiments are conducted depending on the use of [EOT].

5. Conclusion

In this paper, a highly effective post-training method for a multi-turn response selection is proposed and evaluated. Our approach achieved new state-of-the-art results for two response selection benchmark data sets, Ubuntu Corpus V1 and Advising Corpus. For future work, we will utilize external domain knowledge, such as ubuntu manual description in Ubuntu Corpus or curriculum information in Advising Corpus, for enhancing domain understandings.

6. Acknowledgements

This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIP) (No.NRF-2017M3C4A7068189). This research was supported by the MSIT (Ministry of Science and ICT), Korea, under the ITRC (Information Technology Research Center) support program (IITP-2020-2018-0-01405) supervised by the IITP (Institute for Information Communications Technology Planning Evaluation).

7. References

- [1] R. Lowe, N. Pow, I. Serban, and J. Pineau, "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*, 2015, pp. 285–294.
- [2] R. Kadlec, M. Schmid, and J. Kleindienst, "Improved deep learning baselines for ubuntu corpus dialogs," *Computing Research Repository*, vol. arXiv:1510.03753, 2015.
- [3] Y. Wu, W. Wu, C. Xing, M. Zhou, and Z. Li, "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)*, 2017, pp. 496–505.
- [4] Z. Zhang, J. Li, P. Zhu, H. Zhao, and G. Liu, "Modeling multi-turn conversation with deep utterance aggregation," in *Proceedings of the 27th International Conference on Computational Linguistics*. Association for Computational Linguistics, 2018, pp. 3740–3752.
- [5] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," in *Proceedings of the Advances in Neural Information Processing Systems*, 2017, pp. 5998–6008.
- [6] X. Zhou, L. Li, D. Dong, Y. Liu, Y. Chen, W. X. Zhao, D. Yu, and H. Wu, "Multi-turn response selection for chatbots with deep attention matching network," in *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2018, pp. 1118–1127.
- [7] C. Tao, W. Wu, C. Xu, W. Hu, D. Zhao, and R. Yan, "One time of interaction may not be enough: Go deep with an interaction-over-interaction network for response selection in dialogues," in *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 2019, pp. 1–11.
- [8] C. Yuan, W. Zhou, M. Li, S. Lv, F. Zhu, J. Han, and S. Hu, "Multi-hop selector network for multi-turn response selection in retrieval-based chatbots," 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)*, 2019, pp. 111–120.
- [9] J. Vig and K. Ramea, "Comparison of transfer-learning approaches for response selection in multi-turn conversations," in *7th Edition of the Dialog System Technology Challenges at AAAI 2019*, 2019.
- [10] M. Peters, M. Neumann, M. Iyyer, M. Gardner, C. Clark, K. Lee, and L. Zettlemoyer, "Deep contextualized word representations," in *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, 2018, pp. 2227–2237.
- [11] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 2019, pp. 4171–4186.
- [12] Z. Yang, Z. Dai, Y. Yang, J. Carbonell, R. R. Salakhutdinov, and Q. V. Le, "Xlnet: Generalized autoregressive pretraining for language understanding," in *Proceedings of the Advances in Neural Information Processing Systems*, 2019, pp. 5753–5763.
- [13] D. Chaudhuri, A. Kristiadi, J. Lehmann, and A. Fischer, "Improving response selection in multi-turn dialogue systems by incorporating domain knowledge," in *Proceedings of the 22nd Conference on Computational Natural Language Learning*. Association for Computational Linguistics, 2018, pp. 497–507.
- [14] H. Xu, B. Liu, L. Shu, and P. S. Yu, "Bert post-training for review reading comprehension and aspect-based sentiment analysis," *Computing Research Repository*, vol. arXiv:1904.02232, 2019.
- [15] J. Phang, T. Févry, and S. R. Bowman, "Sentence encoders on stilts: Supplementary training on intermediate labeled-data tasks," *Computing Research Repository*, vol. arXiv:1811.01088, 2018.
- [16] A. Wang, Y. Pruksachatkun, N. Nangia, A. Singh, J. Michael, F. Hill, O. Levy, and S. Bowman, "Superglue: A stickier benchmark for general-purpose language understanding systems," in *Proceedings of the Advances in Neural Information Processing Systems*, 2019, pp. 3266–3280.
- [17] C. Gunasekara, J. K. Kummerfeld, L. Polymenakos, and W. Lasecki, "DSTC7 task 1: Noetic end-to-end response selection," in *Proceedings of the First Workshop on NLP for Conversational AI at ACL 2019*, 2019, pp. 60–67.
- [18] X. Zhou, D. Dong, H. Wu, S. Zhao, D. Yu, H. Tian, X. Liu, and R. Yan, "Multi-view response selection for human-computer conversation," in *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, 2016, pp. 372–381.
- [19] Q. Chen, X. Zhu, Z.-H. Ling, S. Wei, H. Jiang, and D. Inkpen, "Enhanced LSTM for natural language inference," in *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2017, pp. 1657–1668.
- [20] C. Tao, W. Wu, C. Xu, W. Hu, D. Zhao, and R. Yan, "Multi-representation fusion network for multi-turn response selection in retrieval-based chatbots," in *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining*. ACM, 2019, pp. 267–275.
- [21] E. M. Voorhees *et al.*, "The trec-8 question answering track report," in *Trec*, vol. 99, 1999, pp. 77–82.
- [22] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard *et al.*, "Tensorflow: A system for large-scale machine learning," in *Proceedings of the 12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)*, 2016, pp. 265–283.
- [23] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean, "Distributed representations of words and phrases and their compositionality," in *Proceedings of the Advances in Neural Information Processing Systems*, 2013, pp. 3111–3119.
- [24] Q. Chen and W. Wang, "Sequential attention-based network for noetic end-to-end response selection," in *7th Edition of the Dialog System Technology Challenges at AAAI 2019*, 2019.
- [25] N. Houlsby, A. Giurgiu, S. Jastrzebski, B. Morrone, Q. de Larousilhe, A. Gesmundo, M. Attariyan, and S. Gelly, "Parameter-efficient transfer learning for nlp," in *Proceedings of the 36th International Conference on Machine Learning*, 2019.