

# Efficient Task-Oriented Dialogue Systems with Response Selection as an Auxiliary Task

**Radostin Cholakov**  
High School of Mathematics  
Plovdiv, Bulgaria  
r.cholakov@obecto.com

**Todor Kolev**  
Obecto Ltd.  
Sofia, Bulgaria  
tkolev@obecto.com

## Abstract

The adoption of pre-trained language models in task-oriented dialogue systems has resulted in significant enhancements of their text generation abilities. However, these architectures are slow to use because of the large number of trainable parameters and can sometimes fail to generate diverse responses. To address these limitations, we propose two models with auxiliary tasks for response selection - (1) distinguishing distractors from ground truth responses and (2) distinguishing synthetic responses from ground truth labels. They achieve state-of-the-art results on the MultiWOZ 2.1 dataset with combined scores of 107.5 and 108.3 and outperform a baseline with three times more parameters. We publish reproducible code and checkpoints and discuss the effects of applying auxiliary tasks to T5-based architectures.

## 1 Introduction

Task-oriented dialogue (TOD) systems are developed to lead conversations with users and assist them with the completion of various tasks. Unlike traditional solutions which rely on natural language understanding, state tracking, language generation, and other modules, end-to-end systems utilize a single network for all required functionality (Young et al., 2013). The recent research in the field has concentrated on leveraging language models pre-trained on general-domain corpora (Devlin et al., 2018; Radford et al., 2019; Raffel et al., 2020) to produce more robust architectures fine-tuned specifically for TOD generation. This has bridged the gap between production-ready modularized pipelines and single-network models in terms of accuracy and human-sounding results. However, such architectures are big and computationally expensive; they are also prone to overfitting on the final task and "forgetting" useful capabilities from the pre-training phase (Greco et al., 2019; Kulhánek et al., 2021). Multiple studies (Section 2)

have demonstrated that learning related auxiliary tasks can improve the generation performance of a model while making it less affected by the overfitting issue.

In this paper, we study the effects of learning auxiliary response selection tasks together with an architecture based on the T5 (Raffel et al., 2020) text-to-text transformer. We use MTTOD (Lee, 2021), trained on the MultiWOZ 2.1 (Eric et al., 2019) dataset, as a baseline and evaluate two main approaches for response selection:

- Binary classifier to distinguish between encodings of ground truth responses and encodings of distractor sentences sampled from the dataset.
- Binary classifier to tell apart ground truth responses from decoder-generated sequences.

Reproducible code and model checkpoints are available at <https://github.com/radi-cho/RSTOD>.

## 2 Related Work

TOD sequence-to-sequence models usually generate a belief state based on the dialogue history and then use the belief state (in addition to the previous context) to generate a response (Lei et al., 2018).

**Pre-trained Language Models** such as BERT (Devlin et al., 2018), GPT-2 (Radford et al., 2019), and T5 (Raffel et al., 2020) significantly enhance dialogue systems when they are fine-tuned for sequence tasks. The first study to validate this on GPT-2 is (Budzianowski and Vulić, 2019). SOLOIST (Peng et al., 2020), UBAR (Yang et al., 2021), and SimpleTOD (Hosseini-Asl et al., 2020) further develop the end-to-end setting of the problem by considering database results and generated responses during training. MinTL (Lin et al., 2020) provides a minimalistic **transfer learning** dialogue system with multiple backbones. TOD-BERT (Wu

et al., 2020) utilizes a contrastive objective function to mimic a response selection task. (Yang et al., 2022) augments data by ignoring nonessential tokens and also adversarially filters “easy” samples to enhance model robustness.

**Auxiliary Learning** - training additional tasks which improve the performance of the primary text generation task - is increasingly applied in TOD systems. AuGPT (Kulhánek et al., 2021) demonstrates that response selection tasks are helpful on top of GPT-2. MTTOD (Lee, 2021) has a span selection auxiliary task. GALAXY (He et al., 2022) (with UniLM (Dong et al., 2019) as a backbone) optimizes four objectives, one of which is a selection between ground truth responses and randomly sampled responses. PPTOD (Su et al., 2022) is also trained for multiple tasks in a plug-and-play fashion (Dathathri et al., 2019).

### 3 Method

#### 3.1 Dialogue System Baseline

As a baseline, we use the end-to-end system setting introduced in (Lee, 2021) (Figure 1) with T5 encoder-decoder backbone. The encoder input consists of a dialogue history concatenated with a user utterance. A *belief state* decoder generates a sequence of a domain name, slot names, and slot values. There is an option for querying a domain-specific database based on the belief state to generate a *DB state* which is then used to condition a final *response* decoder. The response output contains a *system action* state - a sequence of a domain name, action types and slots - and a *system response*. Since the decoder works autoregressively<sup>1</sup>, response generation is automatically conditioned on the system action.

As shown in figure 1, MTTOD utilizes a classifier as an auxiliary task for span matching, inspired by recent dialogue state tracking approaches. Labels for this task are the extractive informable slots defined in (Gao et al., 2020).

The loss to be jointly minimized is

$$\mathcal{L} = \alpha \mathcal{L}_{belief} + \beta \mathcal{L}_{resp} + \gamma \mathcal{L}_{span} \quad (1)$$

where  $\mathcal{L}_{belief}$  and  $\mathcal{L}_{resp}$  are negative log-likelihood language modeling losses for the two decoders and  $\mathcal{L}_{span}$  is a cross-entropy loss for the span task. For compatibility with (Lee, 2021) the

<sup>1</sup>An autoregressive decoder uses information from previous time steps to generate the value at the current time step.

coefficients  $\alpha$ ,  $\beta$  and  $\gamma$  are set to 1.0, 1.0 and 0.5 respectively. Refer to section 5 for baseline benchmarks.

#### 3.2 Response Selection as an Auxiliary Task

Our study aims to evaluate the effects of using response selection as an additional auxiliary task for the presented T5-based dialogue system. We propose two variants for such a task (Figure 2) and modify the full objective to

$$\mathcal{L} = \alpha \mathcal{L}_{belief} + \beta \mathcal{L}_{resp} + \gamma \mathcal{L}_{span} + \delta \mathcal{L}_{select} \quad (2)$$

In our experiments  $\delta$  is also set to 0.5.

##### 3.2.1 Distinguishing distractor encodings

The first proposal for a response selection task in our system is a binary classifier head - a linear layer or a simple multilayer perceptron - distinguishing randomly sampled *distractor* responses from ground truth responses. During training, the dialogue context  $C_t$  at time step  $t$  (consisting of history  $H_t$  and user utterance  $U_t$ ) is concatenated with both the ground truth labels  $T_t$  - forming a sequence  $(C_t, T_t)$  - and a distractor response  $D_t$  sampled from the dataset - forming a sequence  $(C_t, D_t)$ . Encodings for both sequences are generated by the already implemented T5 encoder and are then fed to the response selection head. The class label is 0 for  $(C_t, D_t)$  and 1 for  $(C_t, T_t)$ . The binary cross entropy loss to be minimized is defined as

$$\begin{aligned} \mathcal{L}_{select} = & -\log p(l = 1 | C_t, T_t) \\ & -\log p(l = 0 | C_t, D_t) \end{aligned} \quad (3)$$

$$p(l = 1 | C_t, T_t) = \text{sigmoid}(\phi_a(\phi_E(C_t, T_t))) \in \mathbb{R}^1 \quad (4)$$

$$p(l = 0 | C_t, D_t) = 1 - \text{sigmoid}(\phi_a(\phi_E(C_t, D_t))) \in \mathbb{R}^1$$

where  $\phi_E$  denotes the encoder and  $\phi_a$  - the final classifier.

Optimizing the auxiliary response selection task affects the gradients of the encoder parameters. We empirically prove that this is beneficial for the overall score improvements on multiple metrics.

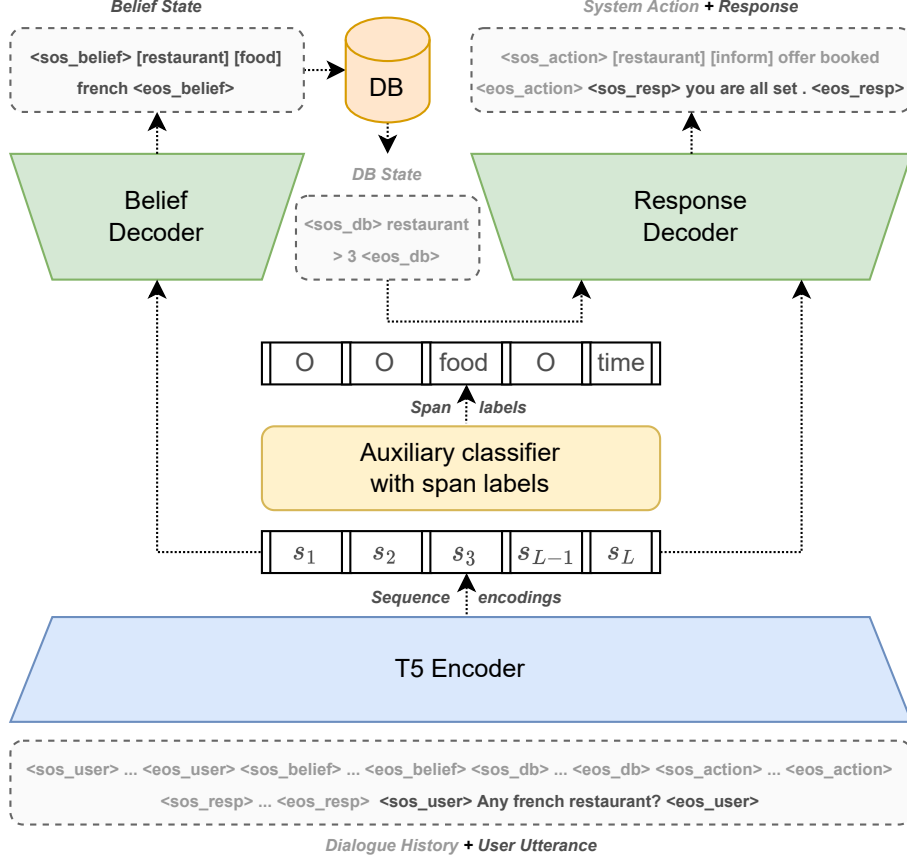


Figure 1: Dialogue generation architecture.

### 3.2.2 Distinguishing generated sequences

We also propose another independent auxiliary task for response selection inspired by Generative Adversarial Networks (Goodfellow et al., 2014). Its goal is to distinguish between responses from the transformer  $R_t$  and ground truth sequences  $T_t$ .

The baseline response decoder generates a sequence of token logits  $\pi_1, \pi_2, \dots, \pi_k$ , where  $\pi_i$  is a vector of unnormalized class outputs over a vocabulary with size  $v$ . To obtain token ids we usually apply

$$\arg \max_j [\log \pi_{ij}], \quad j \in [1, v - 1] \quad (5)$$

for every  $\pi_i$ . However, such a step is not differentiable, and when subsequent layers are optimized, transformer gradients won't be affected, making the auxiliary task useless. One way to overcome the limitation is to re-encode the sequences as previously described in 3.2.1 and thus backpropagate knowledge to the T5 encoder. Instead, we propose a classifier that works with differentially sampled token representations and backpropagates knowledge to the whole architecture during training.

We sample vocabulary class probabilities  $y_{i1}, y_{i2}, \dots, y_{iv}$  for every token representation  $\pi_i$  from a Gumbel-Softmax distribution (Jang et al., 2016; Maddison et al., 2016; Gumbel, 1954):

$$y_{ij} = \frac{\exp((\log(\pi_{ij}) + g_j)/\tau)}{\sum_{k=1}^v \exp((\log(\pi_{ik}) + g_k)/\tau)} \quad (6)$$

where  $\tau$  is a temperature, treated as a hyperparameter, and  $g_j$  is a noise sample from a Gumbel(0, 1) distribution which can be computed by drawing a  $u \sim \text{Uniform}(0, 1)$  and applying

$$g = -\log(-\log(u)) \quad (7)$$

For consistency with ground truth response sequences which are represented with  $v$ -dimensional one-hot vectors  $\hat{y}_i$ , we programmatically<sup>2</sup> convert the probabilities  $y_i$  to one-hot vectors but compute gradients with their continuous values.

<sup>2</sup>Refer to the `hard` flag in [http://pytorch.org/docs/stable/generated/torch.nn.functional.gumbel\\_softmax](http://pytorch.org/docs/stable/generated/torch.nn.functional.gumbel_softmax)

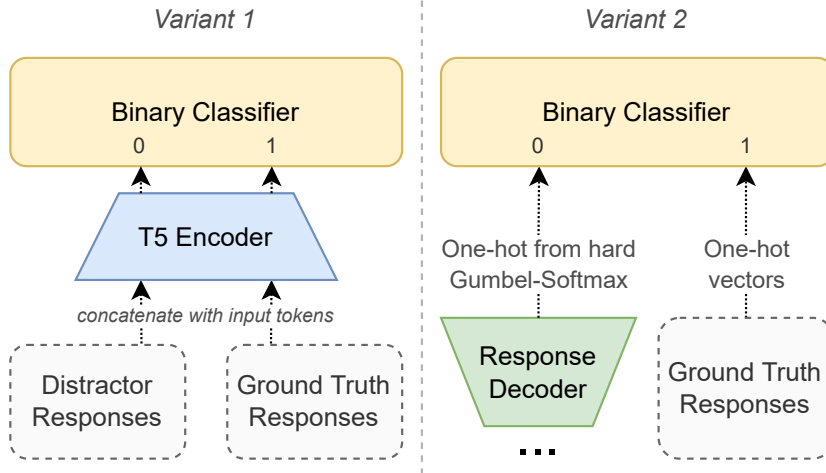


Figure 2: Binary classification response generation tasks.

Finally, both  $y$  and  $\hat{y}$  are fed to the binary classifier  $\phi_b$  and the loss to be minimized is computed as

$$\mathcal{L}_{select} = -\log p(l = 1 | \hat{y}) - \log p(l = 0 | y) \quad (8)$$

$$\begin{aligned} p(l = 1 | \hat{y}) &= \text{sigmoid}(\phi_b(\hat{y})) \in \mathbb{R}^1 \\ p(l = 0 | y) &= 1 - \text{sigmoid}(\phi_b(y)) \in \mathbb{R}^1 \end{aligned} \quad (9)$$

## 4 Experiments

### 4.1 Dataset

In our workflow, we use the large-scale TOD dataset MultiWOZ 2.1 (Eric et al., 2019) for benchmarks and comparisons with baselines. We follow the preprocessing techniques from (Zhang et al., 2020; Lee, 2021) to replace the specific slot values with placeholders. Table 1 presents more in-depth details and statistics on the contents of the dataset.

### 4.2 Training procedure

Train/development/test sets are generated with 80%/10%/10% of the samples. We optimize the objectives from section 3 for 15 epochs and report the results from the best performing checkpoint on the development set. In our experiments, we tested different learning rate schedule strategies and found the best results to be achieved with a constant learning rate initialized as  $5 \times 10^{-4}$  with liner warmup for the first 10% of the samples.

For variant 2 of our architecture, a scheduler is used to linearly decrease the Gumbel-Softmax temperature  $\tau$  with each training iteration. The optimal initial value for  $\tau$  used to derive the results in Table 2 is 4 and is gradually decreased to 0.8.

### 4.3 Evaluation Metrics

During inference, the response selection head is not used and the model performs the same way in terms of speed as the T5-small baseline. We compute *BLEU* (Papineni et al., 2002), *Inform* and *Success* metrics for both architecture variants. *Inform* validates whether system entities are correct and *Success* checks whether relevant information is given for all user inquiries. A combined score is derived as  $0.5 \times (\text{Inform} + \text{Success}) + \text{BLEU}$  which is consistent with previous studies.

## 5 Results

### 5.1 MultiWOZ Benchmarks

Table 2 compares the calculated benchmarks for the two proposed variants of our auxiliary task. As a baseline, we present the results of MTTOD with *T5 base* backbone, which has more than 360 million trainable parameters. In contrast, our models, which use *T5 small* as a backbone, have 102.2 and 105.5 million parameters but achieve higher overall results with total scores of 107.4 and 108.3, respectively.

## 6 Discussion

Response selection tasks similar to variant 1 of our architecture have been previously applied in models for chit-chatting and question answering (Wolf et al., 2019). For TOD systems such tasks are used in architectures with GPT-2 (AuGPT) and UniLM (GALAXY) backbones resulting in responses with higher text-generation metrics. Our study is the first to provide an in-depth analysis of whether a T5-based model in a task-oriented setting would

Table 1: MultiWOZ dataset statistics

| Domain                         | Train dialogues | Dev dialogues | Test dialogues |
|--------------------------------|-----------------|---------------|----------------|
| Police                         | 245             | 0             | 0              |
| Hospital                       | 287             | 0             | 0              |
| Attraction                     | 127             | 11            | 12             |
| Taxi                           | 326             | 57            | 52             |
| Train                          | 282             | 30            | 33             |
| Hotel                          | 513             | 56            | 67             |
| Restaurant                     | 1199            | 50            | 62             |
| Train + Attraction             | 883             | 148           | 163            |
| Hotel + Attraction             | 437             | 55            | 50             |
| Restaurant + Attraction        | 396             | 78            | 70             |
| Restaurant + Train             | 875             | 157           | 155            |
| Restaurant + Hotel             | 462             | 59            | 49             |
| Hotel + Train                  | 1077            | 149           | 144            |
| Restaurant + Hotel + Taxi      | 454             | 41            | 42             |
| Restaurant + Attraction + Taxi | 431             | 53            | 59             |
| Hotel + Attraction + Taxi      | 444             | 56            | 42             |
| Total                          | 8438            | 1000          | 1000           |

Table 2: Benchmark results on MultiWOZ 2.1

| Model        | Backbone | Selection task | Parameters | Inform       | Success      | BLEU         | Score         |
|--------------|----------|----------------|------------|--------------|--------------|--------------|---------------|
| MTTOD*       | T5 base  | None           | 360.9 M    | 92.30        | 84.00        | 19.41        | 107.56        |
| MTTOD*       | T5 small | None           | 102.2 M    | 89.20        | 80.50        | 19.14        | 103.99        |
| RSTOD (ours) | T5 small | After encoder  | 102.2 M    | 92.10        | 83.30        | <b>19.69</b> | 107.39        |
| RSTOD (ours) | T5 small | Differentiable | 105.5 M    | <b>93.50</b> | <b>84.70</b> | 19.24        | <b>108.34</b> |

\* MTTOD benchmarks are reproduced using its public source code. A slight deviation from the results in (Lee, 2021) is caused by a correction in the evaluation scripts as acknowledged on <https://github.com/bepoetree/MTTOD>.

benefit from selection tasks. The results we present are consistent with related literature since we also observe an increase in generation performance.

Most of the solutions relying on pre-trained language models have big amounts of trainable parameters making them slow to train. In our study, we use a modification of the baseline with T5-small instead of T5-base, reducing the parameters more than 3 times. In variant 1 the shared encoder is responsible for processing more sequences than the baseline - it is slower to train but identical in terms of inference speed and amount of storage space required for its weights. Variant 2 is comparable in terms of train-time and inference-time speed to the baseline but is able to achieve a higher overall score. It employs techniques for overcoming backpropagation issues with the discrete token representations of a generated response sequence<sup>3</sup>.

<sup>3</sup>Usually text is generated by picking the most likely tokens

## 7 Future Work

Directions for further research on the topic of TOD systems include testing our proposals on bigger backbone models to evaluate their effectiveness against overfitting, experimenting with additional auxiliary tasks for the current baseline, and introducing data augmentations. Also, whether our classifier heads could be used during inference to perform real-time response selection should be explored.

As a long-term development in the field, we consider various possibilities for building production-ready end-to-end dialogue systems by employing reinforcement learning or semi-supervised learning methods. More experimentally, a generative adversarial network for creative text generation could

from a probability distribution over the token space. This is not a differentiable operation and prevents gradient computations.



also be tested.

## 8 Conclusion

In this paper, we propose two independent auxiliary tasks for response selection on top of a TOD system transformer baseline. Both tasks demonstrate state-of-the-art results on multiple text generation metrics despite having 3+ times less trainable parameters. The first variant involves a classifier, distinguishing between distractor and ground truth responses, which affects the transformer encoder during training and achieves results consistent with related literature. The second variant applies a novel technique for the TOD problem and involves a classifier, distinguishing between synthetic and ground truth responses. We publish reproducible code implementations of our proposals and present potential directions for future research.

## References

- Paweł Budzianowski and Ivan Vulić. 2019. Hello, it’s gpt-2—how can i help you? towards the use of pre-trained language models for task-oriented dialogue systems. *arXiv preprint arXiv:1907.05774*.
- Sumanth Dathathri, Andrea Madotto, Janice Lan, Jane Hung, Eric Frank, Piero Molino, Jason Yosinski, and Rosanne Liu. 2019. Plug and play language models: A simple approach to controlled text generation. *arXiv preprint arXiv:1912.02164*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Li Dong, Nan Yang, Wenhui Wang, Furu Wei, Xiaodong Liu, Yu Wang, Jianfeng Gao, Ming Zhou, and Hsiao-Wuen Hon. 2019. Unified language model pre-training for natural language understanding and generation. *Advances in Neural Information Processing Systems*, 32.
- Mihail Eric, Rahul Goel, Shachi Paul, Abhishek Sethi, Sanchit Agarwal, Shuyang Gao, and Dilek Hakkani-Tur. 2019. Multiwoz 2.1: Multi-domain dialogue state corrections and state tracking baselines. *arXiv preprint arXiv:1907.01669*.
- Shuyang Gao, Sanchit Agarwal, Tagyoung Chung, Di Jin, and Dilek Hakkani-Tur. 2020. From machine reading comprehension to dialogue state tracking: Bridging the gap. *arXiv preprint arXiv:2004.05827*.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. *Advances in neural information processing systems*, 27.
- Claudio Greco, Barbara Plank, Raquel Fernández, and Raffaella Bernardi. 2019. Psycholinguistics meets continual learning: Measuring catastrophic forgetting in visual question answering. *arXiv preprint arXiv:1906.04229*.
- Emil Julius Gumbel. 1954. *Statistical theory of extreme values and some practical applications: a series of lectures*, volume 33. US Government Printing Office.
- Wanwei He, Yinpei Dai, Yinhe Zheng, Yuchuan Wu, Zheng Cao, Dermot Liu, Peng Jiang, Min Yang, Fei Huang, Luo Si, et al. 2022. Galaxy: A generative pre-trained model for task-oriented dialog with semi-supervised learning and explicit policy injection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 10749–10757.
- Ehsan Hosseini-Asl, Bryan McCann, Chien-Sheng Wu, Semih Yavuz, and Richard Socher. 2020. A simple language model for task-oriented dialogue. *Advances in Neural Information Processing Systems*, 33:20179–20191.
- Eric Jang, Shixiang Gu, and Ben Poole. 2016. Categorical reparameterization with gumbel-softmax. *arXiv preprint arXiv:1611.01144*.
- Jonáš Kulhánek, Vojtěch Hudeček, Tomáš Nekvinda, and Ondřej Dušek. 2021. Augpt: Auxiliary tasks and data augmentation for end-to-end dialogue with pre-trained language models. *arXiv preprint arXiv:2102.05126*.
- Yohan Lee. 2021. [Improving end-to-end task-oriented dialog system with a simple auxiliary task](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 1296–1303, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Wenqiang Lei, Xisen Jin, Min-Yen Kan, Zhaochun Ren, Xiangnan He, and Dawei Yin. 2018. Sequicity: Simplifying task-oriented dialogue systems with single sequence-to-sequence architectures. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1437–1447.
- Zhaojiang Lin, Andrea Madotto, Genta Indra Winata, and Pascale Fung. 2020. Mintl: Minimalist transfer learning for task-oriented dialogue systems. *arXiv preprint arXiv:2009.12005*.
- Chris J Maddison, Andriy Mnih, and Yee Whye Teh. 2016. The concrete distribution: A continuous relaxation of discrete random variables. *arXiv preprint arXiv:1611.00712*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.

- Baolin Peng, Chunyuan Li, Jinchao Li, Shahin Shayan-deh, Lars Liden, and Jianfeng Gao. 2020. Soloist: Few-shot task-oriented dialog with a single pre-trained auto-regressive model. *arXiv preprint arXiv:2005.05298*.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, Peter J Liu, et al. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21(140):1–67.
- Yixuan Su, Lei Shu, Elman Mansimov, Arshit Gupta, Deng Cai, Yi-An Lai, and Yi Zhang. 2022. [Multi-task pre-training for plug-and-play task-oriented dialogue system](#).
- 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*.
- Chien-Sheng Wu, Steven C.H. Hoi, Richard Socher, and Caiming Xiong. 2020. [TOD-BERT: Pre-trained natural language understanding for task-oriented dialogue](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 917–929, Online. Association for Computational Linguistics.
- Shiquan Yang, Xinting Huang, Jey Han Lau, and Sarah Erfani. 2022. Robust task-oriented dialogue generation with contrastive pre-training and adversarial filtering. *arXiv preprint arXiv:2205.10363*.
- Yunyi Yang, Yunhao Li, and Xiaojun Quan. 2021. Ubar: Towards fully end-to-end task-oriented dialog system with gpt-2. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 14230–14238.
- Steve Young, Milica Gašić, Blaise Thomson, and Jason D. Williams. 2013. [Pomdp-based statistical spoken dialog systems: A review](#). *Proceedings of the IEEE*, 101(5):1160–1179.
- 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.