

SELF-Eval: Self-supervised Fine-grained Dialogue Evaluation

Longxuan Ma and Ziyu Zhuang and Weinan Zhang* and Mingda Li and Ting Liu
Research Center for Social Computing and Information Retrieval, Harbin Institute of Technology
lxma, zyzhuang, wnzhang, mdli, tliu@ir.hit.edu.cn

Abstract

This paper introduces a novel **Self-supervised Fine-grained Dialogue Evaluation** framework (**SELF-Eval**). The core idea is to model the correlation between turn quality and the entire dialogue quality. We first propose a novel automatic data construction method that can automatically assign fine-grained scores for arbitrarily dialogue data. Then we train **SELF-Eval** with a multi-level contrastive learning schema which helps to distinguish different score levels. Experimental results on multiple benchmarks show that SELF-Eval is highly consistent with human evaluations and better than the state-of-the-art models. We give a detailed analysis of the experiments in this paper. Our code is available on GitHub.

1 Introduction

Dialogue systems (DS) aim to satisfy human needs (Shum et al., 2018; Yan, 2018; Gao et al., 2019) such as information, communication, entertainment, etc. Appraising the quality of the DS responses reflects the system’s capability and provides insights into required further improvements (Finch and Choi, 2020; Deriu et al., 2021). Among the commonly used evaluation metrics, human evaluation is of high reliability but expensive to conduct, automatic metrics used in language generation (Perplexity (Bengio et al., 2000)) or machine translation (BLEU (Papineni et al., 2002), ROUGE (Lin, 2004), etc.) are easy to conduct but ineffective to reflect the dialogue quality (Liu et al., 2016; Novikova et al., 2017). Therefore, researchers have made great efforts to find more reliable automatic evaluation metrics that are highly correlated with human evaluation (Lowe et al., 2017; Tao et al., 2018; Mehri and Eskénazi, 2020a).

The current automatic dialogue evaluation metrics leverage semantic information (Huang et al., 2020; Mehri and Eskénazi, 2020b; Ye et al., 2021)

*Corresponding author

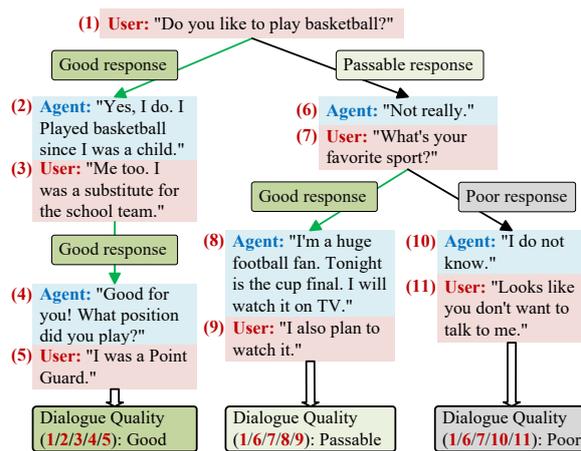


Figure 1: The fine-grained relationships between turn-level and dialogue-level quality.

to measure dialogue quality. For example, when evaluating the response quality, they either compute the semantic similarity between dialogue context and the generated response (Xu et al., 2018a; Tao et al., 2018; Dziri et al., 2019; Ghazarian et al., 2019) or measure the soft semantic overlap between ground-truth response and the model-generated one (Lowe et al., 2017; Xu et al., 2018b; Zhao et al., 2019; Zhang et al., 2020a; Yuan et al., 2021). When evaluating (Zhang et al., 2021) the overall dialogue quality, they either learn a dialogue-level representation for rating directly (Mesgar et al., 2020; Zhang et al., 2021) or calculate the score with the help of other indirect assists (Mehri and Eskénazi, 2020a). However, recent studies (Mehri and Eskénazi, 2020a; Sai et al., 2021; Yeh et al., 2021) show that current models can only work well for measuring the response or evaluating the entire dialogue. They could not perform well in both situations at the same time. It means that the dialogue representation they learned (Zhang et al., 2021) could not reflect both turn quality and the entire dialogue quality.

The dialogue quality is affected by all turns’ qualities in it (Gopalakrishnan et al., 2019) and this effect is accumulated in a multi-turn dialogue

(Li et al., 2021). Figure 1 shows these quality correlations between turns and dialogue. Each turn is marked with a serial number. Three dialogue examples are starting with the same user turn (1). The left example (1/2/3/4/5) shows that two good agent responses result in good overall quality. The middle (1/6/7/8/9) and right (1/6/7/10/11) examples show how lower-quality agent responses result in different dialogue qualities (passable or poor). The current open-domain dialogue evaluation methods fail to model the fine-grained correlations between turn quality and dialogue quality, which entail a poor dialogue representation for evaluation.

In this paper, we introduce an evaluation method that explicitly models the correlations between turn quality and dialogue quality. Specifically, we aim to learn a dialogue representation that can reflect each turn’s contribution, so that the evaluation score obtained by this representation aligns turn quality with the dialogue quality. To this end, we need to first obtain large amounts of dialogue data that reflects fine-grained correlations between turns quality and dialogue quality, then train an evaluation model to measure the fine-grained correlations. The contributions of this paper are:

- To the best of our knowledge, we are the first to explicitly model the fine-grained correlation between turns and the entire dialogue for open-domain dialogue evaluation.
- We introduce a simple but effective data construction method to align the turn-level quality with the overall dialogue quality. We design a **Self-supervised Fine-grained Dialogue Evaluation** model (Self-Eval) with a multi-level contrastive learning (MLCL) method. Our code and data are publicly available: <https://github.com/royny/Self-Eval>.
- Experiments on multiple benchmarks show that Self-Eval: 1) can evenly distinguish different replacement levels; 2) builds the correlations between turn qualities and dialogue qualities; 3) gets better correlation scores with human ratings than the state-of-the-art (SOTA) models.

2 Related Work

We first survey evaluation metrics in open-domain dialogue (sections 5.1 and 5.2), then compare related work in task-oriented dialogue (section 5.3).

2.1 Calculation of Semantic Overlap

In this category, metrics are designed to measure the semantic similarity between the generated response and the dialogue context (Xu et al., 2018a; Tao et al., 2018; Ghazarian et al., 2019; Pang et al., 2020) or soft semantic overlap between the generated response and the reference response (Lowe et al., 2017; Xu et al., 2018b; Zhang et al., 2020a; Zhao et al., 2019; Yuan et al., 2021). Dziri et al. (2019) presented interpretable metrics for evaluating topic coherence by making use of distributed sentence representations. COMET (Rei et al., 2020) evaluated machine translation quality with a pre-trained model by minimizing the distance of the hypothesis with both reference and source text. The most similar work to ours is from Ye et al. (2021) that measures quantifiable coherence scores. *The differences between their work and ours are: 1) they focus on turn-level evaluations while we aim to evaluate both turn and dialogue-levels. Their method models the relationship between dialogue context and response while we model the fine-grained correlations between turns and the entire dialogue; 2) their method relies on the multi-level human annotations for dialogue quality while our method is free from these constrain. Self-Eval is trained in a self-supervised manner, using synthetic dialogue data and automatically annotated scores.*

2.2 Regression to a Reference Score

In this category, metrics learn to evaluate dialogue with scores that represent pre-defined dialogue attributes. BLEURT (Sellam et al., 2020) trained a BERT model with synthetic data and fine-tuned it on human ratings. GRADE (Huang et al., 2020) introduced dialogue topic transitions for coherence evaluation. USR (Mehri and Eskénazi, 2020b) leveraged RoBERTa and a regression model to approximate the specific scores rated by annotators. Mehri and Eskénazi (2020a) computed the log-likelihood of DialoGPT generating predefined positive or negative comments as the score. The most similar work to ours in this category is from Mesgar et al. (2020), they utilized dialogue act labels to help dialogue level representation learning and assist the performance of dialogue-level coherence evaluation. *The difference between their work and ours are: 1) they use dialogue act to assist the dialogue representation learning in a multi-task learning framework while we only use dialogue information; 2) their method only measures the co-*

herence of dialogue while ours measures multiple attributes of dialogue; 3) their method aims to distinguish good samples from bad ones while ours can assign fine-grained scores for each sample.

2.3 Related Work in Task-oriented DS

Besides the open-domain dialogue, there is also work in task-oriented dialogue similar to ours. They assume that users start a dialogue with a task-sensitive patience budget and the dialogue is finished by users when the task is completed or the budget runs out. Their model estimates user satisfaction at each turn and consumes some remaining budget. *The differences between their work and ours are: 1) they focus on task-oriented dialogue with explicitly dialogue purpose while ours are open-domain dialogue with no such goals; 2) they need expensive training and data collecting pipeline while ours do not need; 3) they set up an overall budget which is consumed during the dialogue while we learn a dialogue-level representation for evaluation.* When adopting our method to evaluate task-oriented dialogues, information such as intents and request types is required to determine task completion. Our model will require further improvements to utilize this information.

3 Our Proposed Method

3.1 Problem Statement

Given an n rounds ($2*n$ turns) dialogue $\mathbf{D} = [A_1, B_1, A_2, B_2, \dots, A_n, B_n]$ where A_i/B_i represents the i -th ($i \in \{1, 2, \dots, n\}$) turn from human-A/machine-B, respectively. The dialogue evaluation model takes \mathbf{D} as input and outputs a quality score (a scalar value) for it.

3.2 Data Construction

We need training data with quantitative annotation on both turn-level and dialogue-level (Figure 1). However, only very few dialogue data today have these kinds of labels (Gopalakrishnan et al., 2019) and the models trained with this kind of data are restricted by domain adaptability and generality. Inspired by previous work (Mesgar et al., 2020; Zhang et al., 2021), we adopt a replacement strategy that perturbs a dialogue at the semantic level. In this strategy, the easily accessible human-human dialogue is considered positive. The negative samples for this dialogue are constructed by replacing some turns with randomly selected turns from other

dialogues¹. These randomly selected turns bring multiple negative effects (topically in-congruent, semantic confusion, etc.) w.r.t the current dialogue context. However, different from previous works that replaced a fixed number of turns in a $n > 0$ rounds dialogue, we set multiple replacement strategies and randomly replace $i \in \{0, 1, \dots, n\}$ turns in it. One sample with more replacements is considered of worse overall quality (Gopalakrishnan et al., 2019; Li et al., 2021). Specifically, we assign a score 1 to the original dialogue and assign a score $(n - i)/n$ to the new dialogue that replaces i turns. By aligning the replaced turn numbers with a reference score, we get the required training data. Meanwhile, we avoid the quantity and domain limitation of human-annotated data and can easily obtain a large amount of fine-grained training data in different domains.

Notice that 1) we treat each round with equal weight in this paper, but there may be differences when replacing a turn in the first round (usually a greeting round) or the last one; 2) we hypothesize a linear relationship between the number of replacement turns and the overall dialogue quality, which is not necessarily true. For example, replacing 3 turns and more than 3 turns in the same 6-rounds dialogue may cause the same damage to the overall dialogue quality. We leave these problems for future work.

3.3 Training

During training, we want to minimize the distance between the predicted score and the reference score. This is a difficult regression task because **1)** unlike the coherence degree between dialogue context and response or semantic relationships in a Natural Language Inference, our automatic score as regression target lacks clear semantic meaning; **2)** dialogues with different replaced turns may have the same reference score. For example, the reference score of replacing 1 turn in a 2 rounds dialogue is equal to the score of replacing 2 turns in a 4 rounds dialogue; **3)** when the replacement level increases, it is hard for the model to distinguish the small differences. For example, the reference score of replacing 4 turns in an 8 rounds dialogue is close to the score of replacing 5 turns in an 8 rounds dialogue (0.5 and 0.375, respectively). To smooth the convergence process, we divide the training stage

¹To ensure the generality of our method, we did not use more complex sampling strategies.

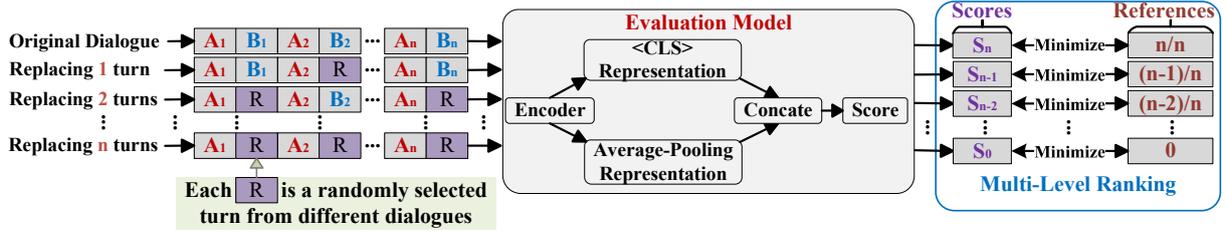


Figure 2: The training procedure of Self-Eval. A_i and B_i is the i -th turn from speaker A and B, respectively. R is the randomly replaced utterances. The replacement positions can be arbitrary rounds.

into **coarse** and **fine-grained** and name the training stages Multi-level Contrastive learning (MLCL) schema. MLCL is model agnostic and can be used for any similar tasks.

Figure 2 shows the training process of our model. We choose RoBERTa (Liu et al., 2019) as encoder. The input for a dialogue sample \mathbf{D} is $[\langle \text{CLS} \rangle, A_1, B_1, \dots, A_n, B_n]$, where " $\langle \text{CLS} \rangle$ " is a special token and n is different for different \mathbf{D} . We first obtain two kinds of dialogue representations. The $\langle \text{CLS} \rangle$ representation is the output vector E_{cls} of the first token $\langle \text{CLS} \rangle$. The *Pooling* representation $E_{pooling}$ is obtained by average-pooling all token representations of $[A_1, B_1, A_2, B_2, \dots, A_n, B_n]$. The final dialogue representation h_D is $[E_{cls}; E_{pooling}]$, where $[\cdot]$ is the concatenation operation. The h_D is passed through a Multi-layer Perceptron (MLP) to get the predicted quality score $S^{\mathbf{D}}$ for \mathbf{D} :

$$S^{\mathbf{D}} = \sigma(W_2 \cdot \mu(W_1 \cdot h_D + b_1) + b_2), \quad (1)$$

where $W_{1,2}$ and $b_{1,2}$ are training parameters; σ/μ is the sigmoid/tanh function, respectively.

3.3.1 The Coarse Training Stage

Formally, given a training corpus $\mathbf{C} = \{\mathbf{D}_m\}_{m=1}^M$ where \mathbf{D}_m is the m -th dialogue with n_m -rounds. For each \mathbf{D}_m , we can replace $i \in \{0, 1, \dots, n_m\}$ turns and get a replaced version of \mathbf{D}_m , named \mathbf{D}_m^i . Each replacement level has its own reference score $(n_m - i)/n_m$. In the first training stage, we combine a separation loss and a compactness loss as the multi-level ranking (mlr) loss. The mlr loss helps the model learn a coarse granularity ranking ability for multi-levels.

The **separation loss** aims to separate the features of different replacement examples by distinguishing their scores. For each replacement level $i \in \{0, 1, \dots, n_m\}$, we first calculate a centroid score $S^{\mathbf{D}_m^i} = \frac{1}{K_i} \sum_{k=1}^{K_i} S_k^{\mathbf{D}_m^i}$ where $S_k^{\mathbf{D}_m^i}$ is the quality score of a dialogue example with i turns replaced, K_i is the number of contrastive samples for

\mathbf{D}_m in this replacement level². The separation loss between different replacement levels is:

$$l_m^{sep} = \sum_{j=0}^{n_m-1} \sum_{l=j+1}^{n_m} \max(0, \omega * \lambda + S^{\mathbf{D}_m^j} - S^{\mathbf{D}_m^l}), \quad (2)$$

where $\lambda=1/(n_m - 1)$ is the lower bound for the distance between two centroid scores³, $\omega = l - j$ is the weight used for amplifying the lower bound according to the quality-level gap.

The **compactness loss** aims to compact the examples within the same level, which served as a regularization role to avoid outlier exceptions for each level. Specifically, the dialogue quality score $S_k^{\mathbf{D}_m^i}$ for $k \in \{1, 2, \dots, K_i\}$ is forced to be closer to the corresponding centroid $S^{\mathbf{D}_m^i}$ as follows:

$$l_m^{com} = \sum_{i=0}^{n_m} \sum_{k=1}^{K_i} \max(0, |S_k^{\mathbf{D}_m^i} - S^{\mathbf{D}_m^i} - \mu|), \quad (3)$$

where μ is the upper bound for the distance between the centroid of a certain replacing level and the score within this level⁴. The mlr loss is:

$$L^{mlr} = \sum_{m=1}^M (l_m^{sep} + l_m^{com}). \quad (4)$$

The original multi-level ranking method is proposed by Ye et al. (2021) and has three secondary losses: separation loss, compactness loss, and ordering loss. The difference between the multi-level ranking methods we used and what they used is that we remove the ordering loss and compute the difference instead of the L1 distance between different centroid scores so that the ordering loss is covered by the separation loss. Our method can save training time and keep equal performance.

²For example, when replacing i turns in a n rounds dialogue, we have total $K_i=n!/((i!(n-i)!)$ contrastive samples.

³For example, $\lambda=0.5$ when there are 3 reference score levels. The expecting 3 centroids are around 1, 0.5, and 0.

⁴For example, if we set λ to 0.3 and μ to 0.1 when there are 3 reference levels. The expecting 3 level ranges are around $[0.9, 1]$, $[0.4, 0.6]$, and $[0, 0.1]$.

Dataset	Ds	turns	words	turns/D
Empathetic	24,846	107,208	1.7M	4.3
ConvAI-2	18,878	278,192	3.3M	14.7
DailyDialog	12,096	100,360	1.4M	8.3
DailyDialog++	19,071	215,625	1.2M	4.3
GRADE	1,200	2,400	61K	2.0
FED	500	5,603	49K	11.2
DSTC-9	2,200	59,840	533K	27.2

Table 1: Statistics of datasets. "D" means dialogue.

3.3.2 The Fine-grained Training Stage

After the coarse training stage, the model has learned to rank multi-level scores, which can be seen as an approximate fitting to the reference labels. To make the training more smooth and more efficient, we add an R-drop loss (Liang et al., 2021) aside from the mlr loss to obtain a more robust representation for each dialogue. The robust representation will help the convergence of the model. Specifically, one input dialogue will go through the model twice and obtain two scores $S_{k,first}^{\mathbf{D}_m^j}$ and $S_{k,second}^{\mathbf{D}_m^j}$, then the model will minimize the distance between the two scores as follows:

$$L^{drop} = \sum_{m=1}^M \sum_{j=0}^{n_m} \sum_{k=1}^{K_i} (S_{k,first}^{\mathbf{D}_m^j} - S_{k,second}^{\mathbf{D}_m^j})^2, \quad (5)$$

The overall Loss of the fine-grained training stage L^{final} is computed as follows:

$$L^{final} = L^{mlr} + L^{drop}. \quad (6)$$

4 Experimental Settings

4.1 Datasets

The datasets used in this paper are shown in Table 1. **Empathetic Dialogue** dataset (Rashkin et al., 2019) simulates real life dialogue in which the interlocutor needs to identify and recognize the feelings of others. **ConvAI-2** (Zhang et al., 2018; Dinan et al., 2019) mimics the scene where each interlocutor tries to understand each other by incorporating persona information. **DailyDialog** (Li et al., 2017) reflects our daily communication and covers different topics such as interpersonal relationships and health. **DailyDialog++** (Sai et al., 2020) is a multi-reference open-domain dialogue dataset with 3 groups (relevant, irrelevant, and adversarial) of responses for each context, each group has 5 different responses. **GRADE** dataset (Huang et al., 2020) contains 300 dialogue examples from Empathetic Dialogue and DailyDialog, and 600 dialogue examples from ConvAI2. Each example has 2 turns with human-annotated relevance

scores. **FED** (Mehri and Eskénazi, 2020a) is a set of human-machine and human-human conversations with eighteen fine-grained quality scores in both turn and dialogue levels. **DSTC-9** (Gunasekara et al., 2020) was collected on the DialPort platform through direct interaction between real users and open-domain chat systems.

4.2 Baselines

We choose the following SOTA models: **GPT-2** (Pang et al., 2020) computes the log-likelihood of the response conditional on the dialogue context normalized by the length of the response; **QuantidCE** (Ye et al., 2021) uses BERT (Devlin et al., 2019) to get dialogue-level representations and proposes a multi-level ranking method to train a quantifiable turn-level coherence metric; **FED** (Mehri and Eskénazi, 2020a) computes the log-likelihood of DialoGPT (Zhang et al., 2020b) generating predefined positive or negative comments as the quality score. It can measure both turn and dialogue-level qualities; **DynaEval** (Zhang et al., 2021) integrates turn representations from RoBERTa into dialogue-level representation with a graph convolutional network, then adopts contrastive learning to distinguish positive and negative samples.

We also test with different settings of Self-Eval. The model shown in Figure 2 is named **Self-Eval(full)**, in which the training dialogue can be any rounds. The first different setting is that we use fixed rounds of dialogue data for training. We set all dialogues to 2 rounds and have 3 different replacement strategies: the original dialogue and replacing 1 or 2 turns. This setting is named **Self-Eval(simple)**. Besides, we have the following settings for the ablation study. **Self-Eval(-mlr)** and **Self-Eval(-drop)** means we remove the multi-level ranking loss and D-drop loss, respectively. When removing the mlr loss, we use a binary cross-entropy (BCE) loss instead. It means the model makes a binary decision between original dialogue and dialogue with replacements. We use BCE loss to show our multi-level ranking method is better than a two-level loss when learning a dialogue representation for evaluation.

4.3 Implementation Details

The setting of the baseline models follows the papers that proposed them. The pre-trained models (BERT, RoBERTa, DialoGPT, GPT-2) are based on the public Pytorch implementation (<https://github.com/huggingface/transformers>).

model	Rep-0	Rep-1	Rep-2	overall
QuantiDCE	0.688	0.486	0.654	0.609
DynaEval	0.812	0.595	0.699	0.702
Self-Eval(simple)	0.962	0.891	0.904	0.919
Self-Eval(full)	0.973	0.898	0.914	0.928

Table 2: Accuracy of predicting replacement levels.

The hyper-parameters which are not introduced in this section follow the original implementation in the link. During fine-tuning, we truncate the input dialogue length to 512 tokens. Among the 7 datasets we used, only DSTC-9 has dialogue examples that exceed 512 tokens and the percentage is 13.6%. We set the max contrastive sample number to 8. All models are learned with Adam optimizer with $\beta_1 = 0.9$ and $\beta_2 = 0.999$. We use a single Tesla A100s GPU with 40GB memory, the batch size is 15. The average training time for each epoch is around 4 hours (2 hours for the first training stage and 2 hours for the second training stage). The initial learning rate is set to 0.005 and decays to 0.002 in the second stage. A dropout of 0.5 is also applied. When training Self-Eval(full), the GPU memory occupation is 39GB. μ is set to 0.1.⁵

4.4 Evaluation Metrics

Following previous works (Mehri and Eskénazi, 2020a; Zhang et al., 2021), we choose two metrics to correlated with manual evaluations. **Pearson Correlation** (Freedman et al., 2007) measures the linear correlation between two sets of data. **Spearman Correlation** (Zar, 2005) assesses the monotonic relationships between two variables. Besides, we use **Accuracy** measures the percentage of correct ranking for multi-level replacement.

5 Experimental Results and Analysis

We aim to answer the following questions about Self-Eval: **(Q1)** *can it assign reasonable scores for multiple replacement levels?* (See section 5.1) **(Q2)** *does it outperform state-of-the-art methods and truly model the correlations between turns/dialogue?* (See section 5.2 and 5.3) **(Q3)** *how do the different components contribute to its performance?* (See section 5.4) **(Q4)** *what can we learn from case study?* (See section 5.5)

⁵When evaluating dialogue-level qualities with turn-level metrics, we measure all context-response pairs in a dialogue and use their average as the final score. When evaluating turn-level qualities with dialogue-level metrics, we treat the context-response pair as an entire dialogue.

model	Pearson	Spearman	average
DynaEval	0.093	0.101	0.097
FED	0.128	0.120	0.124
Self-Eval(simple)	0.158	0.165	0.162
Self-Eval(full)	0.163	0.173	0.168

Table 3: Evaluation on DSTC-9 dialogue-level quality.

5.1 Ranking Capability (Q1)

This experiment tests whether an evaluation model assigns higher scores for dialogues with less replacement. Table 2 shows the accuracy results of QuantiDCE, DynaEval, and Self-Eval(simple/full), all models 1) are trained with DailyDialog++ and test with the DailyDialog++ test set; 2) use base-sized pre-trained models as backbones. QuantiDCE is chosen because 1) it is trained for classification and fits perfectly for this experiment; 2) it represents the SOTA turn-level metric. DynaEval is chosen because it is the SOTA dialogue-level metric. We define 3 replacement levels: the original dialogue (Rep-0), the dialogue with 1 replacing turn (Rep-1), and more than 1 replacing turn (Rep-2). Each replacement level has 5010 samples.

We can see that Self-Eval(full) gets the highest performance on all replacement levels. Between the multi-level ranking models, Self-Eval(full) outperforms QuantiDCE by 52.4%. Between the dialogue-level ranking models, Self-Eval(full) surpasses DynaEval by 32.2%. Notably, the accuracy gaps between Rep-(0, 1, and 2) of Self-Eval(full) are 0.075/-0.016, which are much smaller than the gaps of QuantiDCE (0.202/-0.168) and DynaEval (0.217/-0.104). The results show that 1) Self-Eval can evenly distinguish the 3 replacement levels; 2) the MLCL method we used shows advantages over multi-level learning in QuantiDCE and contrastive learning in DynaEval.

5.2 Experiments on DSTC-9 (Q2)

Table 3 shows the experimental results on DSTC-9 data. The DSTC-9 dataset is difficult to evaluate because of two reasons: 1) it contains direct interaction between real users and multiple open-domain chit-chat systems. These chit-chat systems are trained with dialogue data in different domains compared with ours. In another word, the DSTC-9 dialogue data is out-of-domain for our model and can be used to test the generality of our method; 2) the average turns in a dialogue is around 27.2, which is the longest among all datasets we used and also much longer than the training data we used. It is difficult for evaluation models to give a score for

Dialogue Aspects	GPT-2	Q-DCE	FED	D-Eval	S-E(s)	S-E(f)	(-drop)	(-mlr)	(-drop,-mlr)	Human
Dialogue-level (11 quality aspects) Spearman Correlation										
Coherence	0.122	0.191	0.251	0.424	0.423	0.436	0.332	0.340	0.137	0.809
Error Recovery	<i>0.097</i>	0.109	0.165	0.351	0.363	0.393	0.252	0.269	0.135	0.840
Consistency	0.093	0.332	<i>0.116</i>	0.326	0.246	0.347	0.233	0.318	0.124	0.562
Diversity	0.145	<i>-0.014</i>	0.420	0.342	0.283	0.263	0.197	0.116	<i>0.022</i>	0.789
Topic Depth	<i>0.094</i>	<i>-0.054</i>	0.476	0.375	0.316	0.327	0.204	0.177	<i>0.004</i>	0.833
Likability	0.178	0.098	0.262	0.357	0.345	0.390	0.285	0.275	<i>0.074</i>	0.838
Understanding	<i>0.073</i>	0.210	0.306	0.373	0.364	0.406	0.329	0.306	0.108	0.809
Flexibility	0.135	0.093	0.293	0.361	0.307	0.317	0.233	0.184	<i>0.082</i>	0.816
Informativeness	0.119	<i>0.063</i>	0.288	0.302	0.311	0.318	0.184	0.194	<i>0.019</i>	0.806
Inquisitiveness	<i>0.070</i>	0.115	0.163	0.294	0.401	0.421	0.309	0.267	<i>0.085</i>	0.769
Overall	0.121	0.140	0.443	0.428	0.428	0.435	0.252	0.272	<i>0.055</i>	0.830
Turn-level (9 quality aspects) Spearman Correlation										
Interestingness	-0.097	-0.163	0.408	0.197	0.146	0.183	0.143	0.142	0.063	0.819
Engagement	<i>-0.096</i>	-0.138	0.318	0.119	0.149	0.206	0.126	0.128	0.047	0.798
Specificity	-0.114	-0.171	0.267	0.161	0.097	0.169	0.139	0.112	<i>0.069</i>	0.790
Relevance	-0.103	<i>0.085</i>	0.152	0.171	0.263	0.282	0.174	0.175	0.101	0.753
Correctness	<i>0.041</i>	<i>0.080</i>	0.133	0.165	0.233	0.291	0.165	0.163	<i>0.083</i>	0.780
S.Appropriateness	<i>-0.081</i>	0.112	0.155	0.112	0.208	0.247	0.143	0.132	<i>0.068</i>	0.682
Understandable	<i>-0.076</i>	0.195	0.111	0.116	0.136	0.173	0.110	0.111	<i>0.046</i>	0.522
Fluency	-0.154	<i>0.071</i>	0.224	<i>0.016</i>	0.095	<i>0.038</i>	<i>0.034</i>	<i>0.036</i>	<i>-0.047</i>	0.714
Overall	<i>-0.090</i>	<i>0.014</i>	0.209	0.207	0.255	0.292	0.146	0.140	<i>0.067</i>	0.820

Table 4: Comparison of both dialogue and turn level Spearman correlations on the FED evaluation dataset. Q-DCE/D-Eval/S-E(s)/S-E(f)/"S." is short for QuantiDCE/DynaEval/Self-Eval(simple)/Self-Eval(full)/Semantically, respectively. Scores with p-values larger than 0.01 are italicized (indicating statistical insignificance).

such a long conversation. In this experiment, all models are fine-tuned on DailyDialog++ and the RoBERTa-based models are all first pre-trained on Empathetic Dialogue, ConvAI-2, and DailyDialog.

Pearson and Spearman correlations between the model-generated scores and the corresponding human evaluation scores are computed in Table 3. FED and DynaEval are chosen because they are the SOTA dialogue-level evaluation models. We can see that both Self-Eval(simple) and Self-Eval(full) largely outperform SOTA baselines even though all models are affected by the out-of-domain and long conversation problems. The results show that Self-Eval is capable of learning a dialogue representation for evaluating dialogue-level quality even in a difficult dataset such as DSTC-9. To further verify the generality of this evaluation ability, we test with other out-of-domain datasets in the following sections.

5.3 Experiments on FED (Q2)

Table 4 shows the experimental results on FED data. In both dialogue and turn-level evaluations, Spearman correlations between the model-generated scores and the corresponding human evaluation scores are computed. Models are trained in the same setting as experiments on DSTC-9.

5.3.1 Dialogue-level Evaluation

There are 11 different aspects of the FED dialogue-level evaluation. GPT-2 and QuantiDCE are SOTA

turn-level evaluation metrics. They evaluate a dialogue based on the aggregation of scores of all the context-response pairs within the dialogue. We can observe that most of their correlation scores (21 out of 22) on dialogue aspects are lower than those of FED and DynaEval. The results are consistent with the conclusion of previous studies (Yeh et al., 2021) that turn-level quality evaluation may be insufficient to assess the dialogue-level performance.

FED has the highest scores on Diversity, Topic Depth, and Overall. These results may indicate that the DialoGPT-based evaluation model (FED) is better at measuring these three attributes than the RoBERTa-based models (DynaEval and Self-Eval). The reason is that DialoGPT uses a large amount of Reddit data for training. The diverse topics and variation expressions in Reddit data provide DialoGPT with more insights on these attributes, especially the dialogue-level Overall attribute. In contrast, DynaEval and Self-Eval are trained with fewer dialogue data (fewer topics and variation expressions). The DynaEval focuses on modeling the dependency between pairs of utterances and the Self-Eval focuses on modeling the correlations between turns and the entire dialogue. They are more useful for evaluating Coherence, Error Recovery, and Consistency aspects which reflect the interaction between turns. Specifically, Self-Eval owns the highest correlation scores in 7 out of 11 dialogue aspects (Coherence, Error Recovery, Con-

sistency, likability, Understanding, informativeness, and Inquisitiveness) and the second-highest correlation scores on Flexibility and Overall. Self-Eval successfully learns to measure these attributes with our replacement strategies. The MLCL training method captures the various dialogue attributes and entails good dialogue-level representations. The dialogue-level evaluation tasks are benefiting from this representation. One way to improve the Diversity and Topic Depth scores of RoBERTa-based models is to pre-train them with dialogue data that contains more topics and domains. We can also notice that Self-Eval(full) is better than Self-Eval(simple) in most aspects. It means the simplified training method used by Self-Eval(simple) is not as strong as the original method introduced in Figure 2.

5.3.2 Turn-level Evaluation

There are 9 different aspects of the FED turn-level evaluation. The turn-level metrics (GPT-2 and QuantiDCE) only get better correlations on 6 out of 18 aspects than the dialogue-level metrics (FED and DynaEval). The results indicate that the generality of these two turn-level evaluation models is not strong. They work well only in constrained environments or on specific datasets. The FED model achieves the highest correlation on Interestingness, Engagement, Specificity, and Fluency. The reason is that the DialoGPT used by FED is trained with an auto-regressive mode and models language generation word by word. DialoGPT focuses more on the token-level correlations and is effective for evaluating the naturalness of an utterance. In contrast, all the RoBERTa-based models (DynaEval and Self-Eval) perform poorly for token-level aspects. This is because they focus on the correlations in the turn level and do not pay enough attention to the token level. One way to strengthen the fluency and Specificity aspects of Self-Eval is to introduce token-level perturbation strategies in training data, such as word drop and addition (Sai et al., 2021). These strategies provide negative samples with semantical or grammatical mistakes which may also be used for setting multi-level turn qualities for training. We consider this token-level perturbation as future work. What’s more, we have a similar finding to Zhang et al. (2021) that Self-Eval(s) and FED complement each other at turn-level. It means that they both perform well in aspects that the other one is not good at. Self-Eval achieves the highest correlation in Relevance, Correctness, Se-

U1: My partner left me the other day.
U2: That’s rough, I’m sorry to hear that.
R: Being a punching bag in a relationship is no good. It’s a 2 way street. Is your partner doing their part?
Scores(Human / Self-Eval / DynaEval): 0.77 / 0.84 / 0.50

U1: i was so stressed when i found out that i did not get accepted in my dream college.
U2: Oh no. Did you have a good backup plan?
R: thats cool, i hope you have a good time.
Scores(Human / Self-Eval / DynaEval): 0.25 / 0.09 / 0.46

Table 5: Case study on GRADE. U1/U2 are the dialogue context and R is the response to be evaluated.

mantically Appropriateness, Understandable, and Overall. The Self-Eval(simple)/(full) outperforms the best baseline 39.7%/22.0% on turn-level Overall, respectively. It also has the second-highest scores on Engagement and Specificity. The results are consistent with the dialogue-level evaluation where Self-Eval has good results on aspects that reflect the interaction between turns. As in the dialogue-level, Self-Eval(full) is still better than Self-Eval(simple) on most aspects in turn-level.

To sum up the experiments on the FED dataset, Self-Eval performs well for both dialogue-level and turn-level evaluations, especially the latter. The reason is that the training process of Self-Eval not only models the correlation between turns and the entire dialogue but also models the inner connection between context and response. Our method successfully aligns the semantic information shared by turns and dialogue and shows good domain adaptability (on both DSTC-9 and FED).

5.4 Ablation Study (Q3)

Table 4 also shows the ablation study of the Self-Eval(full). Removing R-drop loss (-drop) in the fine-grained training stage causes more declines than (-mlr) in dialogue-level evaluations. This is because the R-drop loss helps Self-Eval to learn more robust dialogue representations. Replacing the mlr loss with BCE loss (-mlr) causes more declines than (-drop) in turn-level evaluations. This indicates that the mlr loss helps to distinguish the turn replacement levels and the semantic inconsistency information caused by the replacements. When removing both mlr and R-drop losses, the performance declines significantly and the results become statistical insignificance.

5.5 Case Study (Q4)

We randomly select 2 examples from GRADE for the case study (Table 5). The task is to evaluate the response when giving dialogue context. We com-

pare the human rating (the relevance scores) with the scores given by DynaEval and Self-Eval(full). They are both based on RoBERTa and could provide more insight into our model. In both cases, the scores given by Self-Eval are closer to the human rating score than DynaEval. This is consistent with the experimental results in the turn-level Relevance of the FED dataset. In both cases, scores from Self-Eval are more polarized than the human evaluations. This indicates that humans may be reluctant to give extreme scores and Self-Eval could improve its performance by penalizing scores that are too extreme. However, whether this penalizing works for dialogue attributes other than relevance requires further study.

6 Conclusion

We propose to measure dialogue quality by modeling the fine-grained correlations between turns and the entire dialogue. We introduce our data construction method and Self-Eval model. Experiments show that Self-Eval builds fine-grained correlations and gets better correlation scores with human ratings than SOTA models. We think our method may have two potential applications: 1) It can be used alone in the evaluation of dialogue tasks after training the Self-Eval model with a large amount of in-domain data; 2) It can be combined with other evaluation models (such as FED) to evaluate the dialogue task by integrating the advantages of different evaluation models. In the future, we would like to improve our method by 1) employing multi-granularity turn-level scores; 2) modeling the nonlinear relationships between replacement numbers and dialogue quality.

7 Acknowledgments

This paper is supported by the Science and Technology Innovation 2030 Major Project of China (No. 2021ZD0113302) and National Natural Science Foundation of China (No. 62076081, No. 61772153 and No. 61936010).

References

Yoshua Bengio, Réjean Ducharme, and Pascal Vincent. 2000. A neural probabilistic language model. In *NIPS*, pages 932–938. MIT Press.

Jan Deriu, Álvaro Rodrigo, Arantxa Otegi, Guillermo Echegoyen, Sophie Rosset, Eneko Agirre, and Mark Cieliebak. 2021. Survey on evaluation methods for dialogue systems. *Artif. Intell. Rev.*, 54(1):755–810.

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, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4171–4186. Association for Computational Linguistics.

Emily Dinan, Varvara Logacheva, Valentin Malykh, Alexander H. Miller, Kurt Shuster, Jack Urbanek, Douwe Kiela, Arthur Szlam, Iulian Serban, Ryan Lowe, Shrimai Prabhunoye, Alan W. Black, Alexander I. Rudnicky, Jason Williams, Joelle Pineau, Mikhail S. Burtsev, and Jason Weston. 2019. The second conversational intelligence challenge (convai2). *CoRR*, abs/1902.00098.

Nouha Dziri, Ehsan Kamaloo, Kory Wallace Mathewson, and Osmar R. Zaiane. 2019. Evaluating coherence in dialogue systems using entailment. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 3806–3812. Association for Computational Linguistics.

Sarah E. Finch and Jinho D. Choi. 2020. Towards unified dialogue system evaluation: A comprehensive analysis of current evaluation protocols. In *Proceedings of the 21th Annual Meeting of the Special Interest Group on Discourse and Dialogue, SIGdial 2020, 1st virtual meeting, July 1-3, 2020*, pages 236–245. Association for Computational Linguistics.

David Freedman, Robert Pisani, Roger Purves, and Ani Adhikari. 2007. *Statistics*.

Jianfeng Gao, Michel Galley, and Lihong Li. 2019. Neural approaches to conversational AI. *Foundations and Trends in Information Retrieval*, 13(2-3):127–298.

Sarik Ghazarian, Johnny Tian-Zheng Wei, Aram Galstyan, and Nanyun Peng. 2019. Better automatic evaluation of open-domain dialogue systems with contextualized embeddings. *CoRR*, abs/1904.10635.

Karthik Gopalakrishnan, Behnam Hedayatnia, Qinglang Chen, Anna Gottardi, Sanjeev Kwatra, Anu Venkatesh, Raefer Gabriel, and Dilek Hakkani-Tür. 2019. Topical-chat: Towards knowledge-grounded open-domain conversations. pages 1891–1895. ISCA.

R. Chulaka Gunasekara, Seokhwan Kim, Luis Fernando D’Haro, Abhinav Rastogi, Yun-Nung Chen, Mihail Eric, Behnam Hedayatnia, Karthik Gopalakrishnan, Yang Liu, Chao-Wei Huang, Dilek Hakkani-Tür, Jinchao Li, Qi Zhu, Lingxiao Luo, Lars Liden, Kaili Huang, Shahin Shayandeh, Runze Liang, Baolin Peng, Zheng Zhang, Swadheen Shukla, Minlie Huang, Jianfeng Gao, Shikib Mehri, Yulan Feng,

- Carla Gordon, Seyed Hossein Alavi, David R. Traum, Maxine Eskénazi, Ahmad Beirami, Eunjoon Cho, Paul A. Crook, Ankita De, Alborz Geramifard, Satwik Kottur, Seungwhan Moon, Shivani Poddar, and Rajen Subba. 2020. Overview of the ninth dialog system technology challenge: DSTC9. *CoRR*, abs/2011.06486.
- Lishan Huang, Zheng Ye, Jinghui Qin, Liang Lin, and Xiaodan Liang. 2020. GRADE: automatic graph-enhanced coherence metric for evaluating open-domain dialogue systems. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 9230–9240. 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, IJCNLP 2017, Taipei, Taiwan, November 27 - December 1, 2017 - Volume 1: Long Papers*, pages 986–995. Asian Federation of Natural Language Processing.
- Ziming Li, Dookun Park, Julia Kiseleva, Young-Bum Kim, and Sungjin Lee. 2021. A data-driven approach to estimate user satisfaction in multi-turn dialogues. *CoRR*, abs/2103.01287.
- Xiaobo Liang, Lijun Wu, Juntao Li, Yue Wang, Qi Meng, Tao Qin, Wei Chen, Min Zhang, and Tie-Yan Liu. 2021. R-drop: Regularized dropout for neural networks. In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 10890–10905.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81.
- Chia-Wei Liu, Ryan Lowe, Iulian Serban, Michael Noseworthy, Laurent Charlin, and Joelle Pineau. 2016. How NOT to evaluate your dialogue system: An empirical study of unsupervised evaluation metrics for dialogue response generation. In *EMNLP*, pages 2122–2132. The Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized BERT pretraining approach. *CoRR*, abs/1907.11692.
- Ryan Lowe, Michael Noseworthy, Iulian Vlad Serban, Nicolas Angelard-Gontier, Yoshua Bengio, and Joelle Pineau. 2017. Towards an automatic turing test: Learning to evaluate dialogue responses. In *ACL (1)*, pages 1116–1126. Association for Computational Linguistics.
- Shikib Mehri and Maxine Eskénazi. 2020a. Unsupervised evaluation of interactive dialog with dialogpt. In *Proceedings of the 21th Annual Meeting of the Special Interest Group on Discourse and Dialogue, SIGdial 2020, 1st virtual meeting, July 1-3, 2020*, pages 225–235. Association for Computational Linguistics.
- Shikib Mehri and Maxine Eskénazi. 2020b. USR: an unsupervised and reference free evaluation metric for dialog generation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 681–707. Association for Computational Linguistics.
- Mohsen Mesgar, Sebastian Bückner, and Iryna Gurevych. 2020. Dialogue coherence assessment without explicit dialogue act labels. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 1439–1450. Association for Computational Linguistics.
- Jekaterina Novikova, Ondrej Dusek, Amanda Cercas Curry, and Verena Rieser. 2017. Why we need new evaluation metrics for NLG. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, EMNLP 2017, Copenhagen, Denmark, September 9-11, 2017*, pages 2241–2252. Association for Computational Linguistics.
- Bo Pang, Erik Nijkamp, Wenjuan Han, Linqi Zhou, Yixian Liu, and Kewei Tu. 2020. Towards holistic and automatic evaluation of open-domain dialogue generation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 3619–3629. Association for Computational Linguistics.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *ACL*, pages 311–318. ACL.
- Hannah Rashkin, Eric Michael Smith, Margaret Li, and Y-Lan Boureau. 2019. Towards empathetic open-domain conversation models: A new benchmark and dataset. In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 5370–5381. Association for Computational Linguistics.
- Ricardo Rei, Craig Stewart, Ana C. Farinha, and Alon Lavie. 2020. COMET: A neural framework for MT evaluation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 2685–2702. Association for Computational Linguistics.
- Ananya B. Sai, Tanay Dixit, Dev Yashpal Sheth, Sreyas Mohan, and Mitesh M. Khapra. 2021. Perturbation checklists for evaluating NLG evaluation metrics. In *Proceedings of the 2021 Conference on Empirical*

- Methods in Natural Language Processing, EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 7-11 November, 2021*, pages 7219–7234. Association for Computational Linguistics.
- 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. *Trans. Assoc. Comput. Linguistics*, 8:810–827.
- Thibault Sellam, Dipanjan Das, and Ankur P. Parikh. 2020. BLEURT: learning robust metrics for text generation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 7881–7892. Association for Computational Linguistics.
- Heung-Yeung Shum, Xiaodong He, and Di Li. 2018. From eliza to xiaoice: challenges and opportunities with social chatbots. *Frontiers of IT & EE*, 19(1):10–26.
- Chongyang Tao, Lili Mou, Dongyan Zhao, and Rui Yan. 2018. RUBER: an unsupervised method for automatic evaluation of open-domain dialog systems. In *AAAI*, pages 722–729. AAAI Press.
- Xinnuo Xu, Ondrej Dusek, Ioannis Konstas, and Verena Rieser. 2018a. Better conversations by modeling, filtering, and optimizing for coherence and diversity. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 3981–3991. Association for Computational Linguistics.
- Zhen Xu, Nan Jiang, Bingquan Liu, Wenge Rong, Bowen Wu, Baoxun Wang, Zhuoran Wang, and Xiaolong Wang. 2018b. LSDSCC: a large scale domain-specific conversational corpus for response generation with diversity oriented evaluation metrics. In *NAACL-HLT*, pages 2070–2080. Association for Computational Linguistics.
- Rui Yan. 2018. "chitty-chitty-chat bot": Deep learning for conversational AI. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden*, pages 5520–5526. ijcai.org.
- Zheng Ye, Liucun Lu, Lishan Huang, Liang Lin, and Xiaodan Liang. 2021. Towards quantifiable dialogue coherence evaluation. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021*, pages 2718–2729. Association for Computational Linguistics.
- Yi-Ting Yeh, Maxine Eskénazi, and Shikib Mehri. 2021. A comprehensive assessment of dialog evaluation metrics. *CoRR*, abs/2106.03706.
- Weizhe Yuan, Graham Neubig, and Pengfei Liu. 2021. Bartscore: Evaluating generated text as text generation. In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 27263–27277.
- Jerrold H Zar. 2005. Spearman rank correlation. *Encyclopedia of biostatistics*, 7.
- Chen Zhang, Yiming Chen, Luis Fernando D’Haro, Yan Zhang, Thomas Friedrichs, Grandee Lee, and Haizhou Li. 2021. Dynaeval: Unifying turn and dialogue level evaluation. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021*, pages 5676–5689. Association for Computational Linguistics.
- Saizheng Zhang, Emily Dinan, Jack Urbanek, Arthur Szlam, Douwe Kiela, and Jason Weston. 2018. Personalizing dialogue agents: I have a dog, do you have pets too? In *ACL (1)*, pages 2204–2213. Association for Computational Linguistics.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020a. Bertscore: Evaluating text generation with BERT. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Yizhe Zhang, Siqi Sun, Michel Galley, Yen-Chun Chen, Chris Brockett, Xiang Gao, Jianfeng Gao, Jingjing Liu, and Bill Dolan. 2020b. DIALOGPT : Large-scale generative pre-training for conversational response generation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations, ACL 2020, Online, July 5-10, 2020*, pages 270–278. Association for Computational Linguistics.
- Wei Zhao, Maxime Peyrard, Fei Liu, Yang Gao, Christian M. Meyer, and Steffen Eger. 2019. Moverscore: Text generation evaluating with contextualized embeddings and earth mover distance. 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, Hong Kong, China, November 3-7, 2019*, pages 563–578. Association for Computational Linguistics.