

Overview of Robust and Multilingual Automatic Evaluation Metrics for Open-Domain Dialogue Systems at DSTC 11 Track 4

Mario Rodríguez-Cantelar¹, Chen Zhang², Chengguang Tang³, Ke Shi³, Sarik Ghazarian⁴, João Sedoc⁵, Luis Fernando D’Haro¹ and Alexander Rudnicky⁶

¹Speech Technology and Machine Learning Group - Universidad Politécnica de Madrid, Spain

²National University of Singapore, Singapore

³Tencent AI Lab, China

⁴University of Southern California, USA

⁵Department of Technology, Operations, and Statistics, New York University, USA

⁶Carnegie Mellon University, USA

Abstract

The advent and fast development of neural networks have revolutionized the research on dialogue systems and subsequently have triggered various challenges regarding their automatic evaluation. Automatic evaluation of open-domain dialogue systems as an open challenge has been the center of the attention of many researchers. Despite the consistent efforts to improve automatic metrics’ correlations with human evaluation, there have been very few attempts to assess their robustness over multiple domains and dimensions. Also, their focus is mainly on the English language. All of these challenges prompt the development of automatic evaluation metrics that are reliable in various domains, dimensions, and languages. This track in the 11th Dialogue System Technology Challenge (DSTC11) is part of the ongoing effort to promote robust and multilingual automatic evaluation metrics. This article describes the datasets and baselines provided to participants and discusses the submission and result details of the two proposed subtasks.

1 Introduction

Recent advances in large-scale neural language models (Devlin et al., 2019; Radford et al., 2019; Zhang et al., 2020) have led to significant attention in dialogue systems, especially in the open domain category. Significant research efforts are dedicated to boost the robustness of dialogue systems, that is, improving their capability to perform well across multiple domains, dimensions, and handling humans’ diverse expressions of the same ideas (e.g., paraphrasing or back-translation).

Automatic evaluation is an indispensable component for speeding up the development of robust

dialogue systems. Common metrics are based on word overlap, such as BLEU (Papineni et al., 2002) and ROUGE (Lin, 2004), which mainly focus on matching syntactic information with a set of golden references. Unfortunately, such metrics correlate poorly with human judgments (Liu et al., 2016) as in open-domain dialogue, there can be limitless feasible responses w.r.t. a dialogue context.

Alternatively, recently developed model-based metrics such as BERTscore (Sun et al., 2022), BLEURT (Sellam et al., 2020), FED (Mehri and Eskenazi, 2020a), and MDD-Eval (Zhang et al., 2022a), which take advantage of the strong semantic representation capability of pre-trained transformer language models, perform the evaluation at semantic and partially pragmatic levels. Some of them do not even need golden references as input. Regrettably, despite their improvement over the word-overlap metrics, these metrics are not perfect; that is, their correlation with human evaluation is still not strong. Moreover, most of them perform well only on a particular dimension (e.g., engagingness or coherence) (Zhang et al., 2022b), or specific to a single domain. In addition, their performance may be highly dependent on the datasets used for training and evaluation (Yeh et al., 2021).

Due to the lack of robust automatic evaluation metrics (Mehri and Eskenazi, 2020a), researchers have to resort to the time-consuming and cost-intensive human evaluation process to analyze the performance of their model and benchmark their proposed methods against baselines.

Furthermore, to the best of our knowledge, none of the existing metrics have been thoroughly tested in a multilingual setting. Metric generalization across different languages is highly desirable, as it

helps the transformation of state-of-the-art English-only dialogue systems into highly capable multilingual systems. Although multilingual pre-trained language models may exist and can be potentially used for training multilingual dialogue systems, human-annotations or high-quality dialogue datasets for languages other than English are very scarce or even nonexistent in the case of some low-resource languages. To address this problem, we take advantage of recent advances in neural machine translation and paraphrasing systems. Using existing high-quality services and models, it is possible to create new datasets for different languages and perform back-translation or paraphrasing to create additional data in the original language to improve and evaluate the robustness of existing metrics. To this end, we propose two subtasks in our track, and their details are listed as follows:

1.1 Track Details

This track consists of two tasks which are explained in more detail below.

Participants will develop effective open-ended and multilingual automatic dialogue evaluation metrics that perform similarly when evaluated in a new language. Participants will develop effective open-ended automatic dialogue evaluation metrics that perform robustly when evaluated over paraphrased/back-translated sentences in English. For both tasks, proposed metrics are expected to show the following two important properties, as indicated in (Deriu et al., 2021):

1. Correlated to human judgments - the metrics should produce evaluation scores that well correlate to human judgments (scores) across multiple languages or alternative responses (i.e., back-translated or paraphrased).
2. Explainable - the metrics should provide constructive and explicit feedback to the generative models in terms of the quality of their generated responses. For instance, if a generative model contradicts itself, the evaluation metrics should signal such behavior.

Participants can propose their own metrics or optionally improve the deep AM-FM (Zhang et al., 2021) baseline evaluation model provided by us. A leaderboard on the ChatEval platform¹ was provided to check the performance of their different

¹<https://chateval.org/dstc11>

proposed models compared to those submitted by other researchers.

For each evaluation task, Spearman’s correlation was used to compare the proposed evaluation metrics against human judgments. A final average score was calculated to rank the submitted metric models. Additional instructions to participants were provided through the Github repository² and by email on the main DSTC distribution list.

2 Task 1: Multilingual Automatic Metrics

In this task, the goal for participants is to propose effective automatic dialogue evaluation metrics that exhibit the properties mentioned above (Section 1.1) and perform well in a multilingual setup (English, Spanish, and Chinese). In concrete, participants were asked to propose a single multilingual model that could provide high correlations with human-annotations when evaluated in multilingual dialogues (development set in Section 2.1) and perform well in the hidden multilingual test set. Participants were required to use pre-trained multilingual models and train them to predict multidimensional quality metrics using self-supervised techniques and, optionally, fine-tune their system over a subset of the development data.

Finally, participants evaluated their models on the development and test sets, expecting to show similar performance in terms of correlations with human-annotations across three languages: English, Spanish, and Chinese. Only development and test sets have human-annotations, and only the test sets were manually translated or paraphrased/back-translated to guarantee the correlations with the original human-annotations on the English data.

2.1 Datasets

Datasets summary Table 1 shows the three clusters of datasets we used or created during the competition. The table shows information about the number of data used to train, develop, and test the proposed metrics. All these datasets clusters were available in English, Spanish, or Chinese and were back-translated into English. CHANEL and CDIAL include open-domain human-human conversations, while DSTC10 includes human-annotations on human-chatbot interactions. The type of annotations or metadata and how each clus-

²https://github.com/Mario-RC/dstc11_track4_robust_multilingual_metrics

ter was used (training/development/test) are indicated in the last three rows.

Table 7 (Appendix A) provides a brief summary of all the statistics of the train, development, and test datasets. The datasets statistics including their number of utterances, avg. number of utterances in each conversation, avg. number of context/response words, type of annotations (turn or dialogue level), number of criteria, number of provided annotations, and type of dialogue systems used for generating responses are shown.

Train As training set, we used the data released during the CHANEL@JSALT2020^{3,4} (Rudnicky et al., 2020) workshop organized by Johns Hopkins University. This cluster consisted of a total of 18 well-known human-human dialogue datasets pre-processed and distributed in a standard format. The total number of dialogues was 393k (approximately 3M turns). An additional advantage of the data in this cluster is that they have been automatically translated back and forth using the same high-quality MS Azure translation service.⁵

Development As development set, the organizers provided data from two clusters of datasets: DSTC10 and CDIAL.

The first one was collected during DSTC10 Track 5 (Zhang et al., 2022c), consisting of more than 35k turn-level human-annotations, which were automatically translated into Spanish and Chinese, and then back-translated into English using MS Azure services.

Second, we used datasets provided by THU-COAI⁶ group (Conversational AI groups from Tsinghua University), naming this cluster of datasets CDIAL. It contains open-domain human-human dialogues. They are originally in Chinese and include 3,470 dialogues (approximately 130k turns). Furthermore, we provided Chinese to English translations through the SotA Tencent MT⁷ system.

Furthermore, Tencent AI manually annotated ~3k random H-H turns (~1k dialogues) of CDIAL in Chinese (at turn-and dialogue-level).

It is important to note that the development data is intended to help participants verify the multilin-

gualism and robustness capabilities of their models in terms of correlations with human-annotations.

Test Furthermore, in order to check the generalization capabilities of the proposed metrics from the participant, the test data included new English, Chinese, and Spanish data of human-chatbot interactions (Appendix B).

A new Human-Chatbot English dataset (HCEnglish) with ~2k turns (~60 dialogues) with three different SotA chatbots (ChatGPT (Radford et al., 2018), GPT-3.5 (Brown et al., 2020), and BlenderBot 3 (Shuster et al., 2022) (Giorgi et al., 2023)). This dataset was manually annotated (turn level and dialogue level) using Amazon Mechanical Turk (AMT), then translated from English to Chinese and Spanish using MS Azure.

In addition, a new Human-Chatbot Chinese dataset (HCChinese) consisting of ~5k turns (~500 dialogues) was generated with three different SotA chatbots (Chinese DialogPT, Microsoft’s Xiaoice (Zhou et al., 2020b) and Baidu’s Plato-XL (Bao et al., 2022)). This dataset was manually annotated (turn and dialogue level) by Tencent AI, and then translated from Chinese to English using the Tencent MT system.

Third, hidden data from the DSTC10 data was used for Spanish with a total of ~1500 turns (~700 dialogues). Existing turn-level annotations were used, as well as Spanish translations and English back-translations created using MS Azure, which were subsequently manually reviewed.

Table 2 shows the number of turns and dialogues for each test dataset for each language. The DSTC10 datasets did not have annotations at dialogue-level.

Metadata Since the quality of translated sentences can play an important role in the estimation of metric scores, quality annotations between the original sentence and its respective translation were delivered for each turn of all datasets. Machine translation Quality Estimation (QE) metric scores were given to participants using the QE COMET⁸ (Rei et al., 2020) system. In addition, for task 1, the cosine similarity between the original sentence and the translated sentence. Thanks to this information, participants could optionally discard dialogues or turns that potentially did not get a high translation quality estimation, therefore reducing potential noise but also allowing the creation of more robust metric systems.

³<https://github.com/CHANEL-JSALT-2020/datasets>

⁴<https://www.clsp.jhu.edu/chaval-chat-dialogue-modeling-and-evaluation/>

⁵<https://azure.microsoft.com/en-us/products/cognitive-services/translator/>

⁶<https://github.com/thu-coai>

⁷<https://www.tencentcloud.com/product/s/tmt>

⁸<https://github.com/Unbabel/COMET>

Dataset Name	CHANEL	DSTC10	CDIAL
#datasets	18	7	3
Language	English, Spanish/Chinese, and English back-translation	English, Spanish/Chinese, and English back-translation	English, Spanish/Chinese, and English back-translation
Dialogues Type	Human-Human Open-Domain	Human-Chatbot Open-Domain	Human-Human Open-Domain
#dialogues/utterances	+ 390,000 / + 3,000,000	+ 18,000 / + 55,000	+ 3,470 / +130,000
Annotations	Sentiment analysis and Toxicity	Sentiment analysis and Toxicity Turn /dialogue level human scores	Turn /dialogue level human scores
Task 1 Set	Public: Train	Public: Dev, Test Hidden: Automatic Translations	Public: Train, Dev
Task 2 Set	Public: Train	Public: Dev, Test Hidden: Manually back-translated/paraphrased	—

Table 1: Summary of train/development/test datasets.

Language Dataset	EN				ZH				ES			Global
	HCEnglish	HCChinese	DSTC10	Total	HCEnglish	HCChinese	DSTC10	Total	HCEnglish	DSTC10	Total	
Turns	1,700	478	114	2,292	364	1,672	123	2,159	55	333	388	4,839
Dialogues	59	40	-	99	15	160	-	175	3	-	3	277

Table 2: Summary statistics of the test dataset used for **task 1** at turn and dialogue level, and separated by language.

In addition, toxicity and sentiment analysis metadata were provided for the original turns in both the CHANEL and DSTC10 datasets for filtering and dialogue curation purposes, as well as to avoid potential biases. These metadata allowed participants to have a better reference of the dataset quality, being of great help for them to decide whether or not to use these original turn and their translations in the training of their evaluation models and, optionally, fine-tune multilingual pre-trained models allowing better performance on the proposed dialogue-oriented tasks.

Data Format All data given follow a unified data format to make the storage, handling, and retrieval easier for the participants. Detailed guidelines are available in the track repository.⁹

Dimensions For HCEnglish, Amazon Mechanical Turk (AMT) was used to collect annotations for each of the dimensions evaluated in the test data. Our annotations restricted the users to location US, >97% approval rate, >1000 HITs done, and a convenience pool of workers used for NLP evaluation tasks.¹⁰ This pool included workers from the pipeline paper (Zhang et al., 2022d) and cloudfiresearch. The average compensation was ~ 15\$/hr. We included text-based attention checks at the dialogue-level as well as and annotator agree-

ment (both with an expert as well as between crowd workers [some from the DSTC10 dataset]) time-based filters on the turn-level.

For the HCChinese data, we leveraged the power of Tencent MT¹¹ to perform the English-to-Chinese translation of the corpus, followed by training a team of six professional Chinese annotators to annotate the dialogues. The entire annotation process spanned a month and incurred costs of approximately 6,194 US dollars, which is in line with the expenses associated with other evaluation datasets. The average cost of annotating each dialogue was 2.36 US dollars. Finally, the average correlation coefficient for Adequacy scored by six annotators is 0.79, and 0.67 for Fluency.

2.2 Dimensions Evaluated

Since open-domain dialogue systems have multifacet nature, the evaluation can be accomplished from different perspectives. Since this is the case in both development and test data of task 1 (multilingual) and task 2 (robust), we include the following dimensions at turn-level and dialogue-level annotations (Mehri et al., 2022):

– Turn-level dimensions:

Appropriateness - The response is appropriate given the preceding dialogue.

Content Richness - The response is informative, with long sentences including multiple entities and conceptual or emotional words.

⁹https://github.com/Mario-RC/dstc11_track4_robust_multilingual_metrics/blob/main/dstc11_track4-datasets-format.md

¹⁰Without the convenience pool our annotator agreement was near random.

¹¹<https://cloud.tencent.com/product/tmt>

Grammatical Correctness - Responses are free of grammatical and semantic errors.

Relevance - Responses are on-topic with the immediate dialogue history.

– **Dialogue-level dimensions:**

Coherence - Throughout the dialogue, is the system maintaining a good conversation flow.

Engageness/Likeability - Throughout the dialogue, the system displays a likeable personality.

Informativeness - Throughout the dialogue, the system provides unique and non-generic information.

Overall - The overall quality of and satisfaction with the dialogue.

Furthermore, when choosing the test dimensions, the annotations available in the train and development data were taken into account to keep them balanced and homogeneous.

The dimensions chosen at the turn level show how much the responses are appropriate, informative including multiple entities and conceptual or emotional words, free of grammatical and semantic errors, and on-topic with the immediate dialogue history. The dimensions chosen at the dialogue level show how much the system maintains a good conversation flow, engages well with the user, provides unique and non-generic information, and the overall quality of the system.

Table 3 summarizes the dimensions for each test data set. As can be seen, the DSTC10 set only has human turn-level annotations.

2.3 Baseline

We provide a multilingual variant of deep AM-FM (Zhang et al., 2021)(used previously during Track5 at DSTC10) as the baseline model. The formulation of both AM and FM remains unchanged except that we switch their original English-based pre-trained language models to multilingual models. For the adequacy metric (AM), we use XLM-R¹² (Conneau et al., 2020) to extract sentence-level embeddings of both the response and the last sentence in the corresponding dialogue context. Then, the cosine similarity of the two embeddings is the AM score assigned to the corresponding response. For the fluency metric (FM), we adopt the multi-

¹²<https://huggingface.co/sentence-transformers/xlm-r-100langs-bert-base-nli-sts-b-mean-tokens>

Sets	Dimensions			
<u>DSTC10</u>				
DSTC10-turn	A	CR	GC	R
ChatEval-turn	A			
JSALT-turn	A			
<u>HCChinese</u>				
HCChinese-dial	C	EL	I	O
HCChinese-turn		CR	GC	R
<u>HCEnglish</u>				
HCEnglish-dial	C	EL	I	O
HCEnglish-turn	A	CR	GC	R
<u>Test data</u>				
Test-dial	C	EL	I	O
Test-turn	A	CR	GC	R

C: Coherence
I: Informativeness
A: Appropriateness
GC: Grammatical Correctness
EL: Engageness/Likeability
O: Overall
CR: Content Richness
R: Relevance

Table 3: Summary of the dimensions (human-annotations) available for each dataset used in the test data, both at the turn and dialogue level.

lingual GPT-2¹³ as the backbone language model. The conditional probability of the response w.r.t. the context given by the multilingual GPT-2 model serves as the FM score of the response. The final AM-FM score is the arithmetic mean of both metric scores. All information related to the baseline model, such as code and data, can be found in this GitHub repository.¹⁴

2.4 Participants

In Task 1, 4 teams participated, which provided a total of 16 submissions. Participants were asked to provide a brief description of the system for their proposals. The two system descriptions provided by the participants are shown below:

Team 4 Their approach utilizes two submetric groups, XLM-R and ChatGPT, for evaluating dialogue responses. The XLM-R group employs the XLM-Roberta-Large encoder model, consisting of NSP (Next Sentence Prediction), VSP (Valid Sentence Prediction), MLM (Masked Language Modeling), and ENG (Engagement) submetrics. The NSP submetric ensembles three models trained on English and multilingual data, while the VSP model

¹³<https://huggingface.co/ai-forever/mGPT>

¹⁴<https://github.com/karthik19967829/DSTC11-Benchmark>

combines different models. The ENG submetric uses an ensemble of encoder models trained on the ENDEX engagement dataset (Xu et al., 2022). The MLM submetric utilizes the pre-trained XLM-R-large model with a Language Modeling head. The ChatGPT group prompts gpt-3.5-turbo to evaluate responses based on the dimensions of the DSTC11 test, with submetrics for dialogue and turn level. Weighted sums of the submetrics are calculated, with the weights learned from a subset of the dev dataset. For the test set, four variations were submitted, including weighted sums of XLM-R and ChatGPT, direct mapping of ChatGPT, and a weighted sum of all models.

In addition, Team 4 used the metadata provided. During their tests performed for task 1 they discovered that increasing the machine translated data affected the performance of the trained models. Therefore, they made use of the quality estimations computed with the COMET MTQE model to use only the best translated dialogues.

For task 2, they trained their models using the least similar, and separately the most similar ones, based on cosine similarity and Levenshtein distance. They found that there was a good correlation between using the paraphrase score and their model performance, with lower scores bringing higher performance and vice versa. They deduced that the lower-scored responses were more diverse and therefore more informative for training.

Team 7 Their Parallel Corpus Alignment Framework enhances model evaluation on parallel corpora, focusing on Robust and Multilingual Automatic Evaluation Metrics for Open-Domain Dialogue systems. By utilizing xlm-roberta-large and bert-base as baseline models, they leverage representations from different languages, paraphrases, and translations to align parallel corpora in the semantic space. Through contrastive learning and multi-dataset distillation, they strengthen the model’s scoring robustness and evaluation capability across various data domains.

2.5 Results

Table 4 shows the results on test data for task 1 at turn and dialogue level. To calculate the scores in the table, the following procedure was followed: 1. Data are separated in each language (English, Chinese, and Spanish); 2. Then, for each language separately, Spearman’s correlation coefficients are calculated for each dimension independently; 3.

Next, we calculate the mean of the correlations of the dimensions in each language (columns EN, ZH, and ES); 4. Finally, we calculate the final mean (Global column) of the language columns.

Team	Turn-level			Global	Rank
	EN	ZH	ES		
Baseline	0.2940	0.0753	0.1826	0.1840	4
Team 2	0.1469	0.1054	0.0808	0.1110	5
Team 4	0.4818	0.3936	0.5890	0.4881	1
Team 5	0.3702	0.0701	0.1983	<i>0.2129</i>	3
Team 7	0.2214	0.3112	0.5644	<u>0.3657</u>	<u>2</u>

Team	Dialogue-level			Global	Rank
	EN	ZH	ES		
Baseline	0.2414	0.4648	0.8080	<u>0.5047</u>	2
Team 4	0.5342	0.7133	0.8080	0.6852	1
Team 5	0.1865	0.1356	0.6830	<i>0.3350</i>	3

Table 4: Spearman’s correlations of the baseline and average correlations of each team’s metrics on **turn-level** and **dialogue-level** test sets for **task 1**. The first position is shown in bold, the second in underline and the third in italics.

To rank each team, the best submission was used according to the calculated global score. Teams 4, 7 and 5 were the best performers at turn-level. Regarding dialogue-level, team 4 was the best performer, followed by the baseline model and then by team 3. In particular, the performance of team 4 is outstanding in all languages.

This shows that team’s 4 model is very effective not only at the global multilingual level, but also in each language separately, showing a very high performance in Spanish, followed by English and then Chinese. This highlights the need of a multilingual metric capable of performing in Chinese to match the results obtained in Spanish or English.

At dialogue-level, team 4 also demonstrated a very high correlation. Having a wide margin of advantage over team 5 and the base model. It should be noted that for Spanish at the dialogue-level, the amount of data was scarce, then producing non-statistical significant results and making difficult to analyze the reason for so high correlation results.

3 Task 2: Robust Evaluation Metrics

In this task, the goal of the participants was to propose robust metrics for automatic evaluation of English dialogues that exhibit previously mentioned properties (subsection 1.1) while being robust when dealing with paraphrased/back-translated English sentences. Here, the expected performance for the

proposed metrics was that they could be on par with the correlations with human-annotations obtained over the original sentences. As robustness criteria proposed, paraphrased/back-translated sentences should have the same semantic meaning as the original sentence but different wording. Task 2 was only evaluated for the English language.

Participants had the opportunity to evaluate their models with developmental data composed of paraphrased/back-translated sentences and their respective human annotations.

3.1 Datasets

Train, development, and test For task 2, the same task 1 datasets were used. However, to evaluate robustness, paraphrases and back-translated data were used. Thus, for task 2, the original datasets data was provided, in addition to the back-translations and paraphrases of the original sentences, but not the translations to other languages. Table 5 shows the number of turns and dialogues for each test data set sent to the participants. The DSTC10 datasets did not have annotations at dialogue-level.

Dataset	HCEnglish	DSTC10	Total
Turns	1,701	404	2,105
Dialogues	59	-	59

Table 5: Summary statistics of **task 2** test datasets at the turn and dialogue level.

For creating semantically similar sentences, we relied on two options: back-translations and a paraphraser model. For back-translations we used MS Azure MT services or Tencent MT system. Then for the paraphraser model, we used PARROT¹⁵ (Damodaran, 2021). Multiple paraphrases were generated for all the original English sentences in each dataset.

For this task, paraphrases were preferable to back-translations. The reason is that current translation systems have a very high quality, so back-translations are often too similar to the original sentence, or even identical, not meeting in this case the robustness criterion proposed in task 2.

Metadata For this specific task, participants received as metadata the Levenshtein¹⁶ distance

¹⁵https://github.com/jsedoc/Parrot_Paraphraser

¹⁶The Levenshtein distance is a numerical measure indicating the similarity between two strings. A higher Levenshtein

calculated for all paraphrases generated from the original sentences, in all datasets. For task 1, QE annotations were given using the same COMET model. In this case, they were calculated between the original sentence and its respective paraphrases separately, and between the original sentence and respective back-translation.

Moreover, participants were given the Cosine Similarity calculation between the original sentence and its respective paraphrases, as well as between the original sentence and its back-translation. Finally, participants were notified of the provided metadata, as well as the toxicity and sentiment analysis annotations, for them to filter potentially biased or noised sentences.

Dimensions Human-annotations for development and test data were the same as for task 1.

3.2 Dimensions Evaluated

As the data for task 2 are the same as those in task 1, the nature of the data is common in both tasks. Therefore, the dimensions used to evaluate the models, both at the turn and dialogue level, are shared between the two tasks.

3.3 Baseline

The same baseline was used for task 2, as for task 12.3. In this case, paraphrases were used instead of multilingual sentences to evaluate robustness.

3.4 Participants

For this task, a total of 5 teams participated and sent a total of 21 submissions. Participants were asked to provide a brief description of their systems. Team 4 and 7 used the same models as for task 1, therefore they are described in Section 2.4. Below, we provide detailed description for teams 3 and 6.

Team 3 To address the variability of metrics in evaluating different dimensions and mitigate overfitting on scarce human-annotated data, they propose IDEL. This approach combines multiple metrics to achieve a higher correlation with human judgment across all dimensions. To avoid overfitting, they employed a list-wise learning-to-rank objective, leveraging the relative positions of examples rather than absolute coordinates. Furthermore, they utilized the LLaMa 65B dataset and the in-context-learning method for direct evaluation of examples, considering their context.

Team 6 Their approach focused on predicting turn-level qualities. They utilized pre-trained Large distance signifies a greater difference between the two strings.

Language Models (LLMs) with manually designed prompts and two selected dialogues as few-shot examples to adapt the LLM output. Additionally, they built a feed-forward neural network (FNN) using frozen LLM representations as features to predict the desired metrics. Another submission employed the ChatGPT API with optimized prompts and dynamically obtained dialogue examples. Hyperparameters were selected based on manual annotations of 157 testing examples. However, for grammaticality metric scores, randomly generated scores were submitted due to uninformative constant scores predicted by the LLM.

3.5 Results

Team results for turn and dialogue levels on the test data for task 2 are provided in Table 6. To calculate the scores and provide the ranking, the following procedure was followed: a. Calculate the Spearman’s correlation coefficients for each dimension independently; b. Calculate the mean of the correlations of the dimensions; c. Calculate the mean (Global column) of the language columns.

Team	Turn-level		Team	Dialogue-level	
	Global	Rank		Global	Rank
Baseline	0.3387	4	Baseline	0.4800	1
Team 1	0.1537	6	Team 1	0.1111	4
Team 3	0.2697	5	Team 3	<i>0.2196</i>	3
Team 4	0.4890	1	Team 4	<u>0.3031</u>	<u>2</u>
Team 6	<u>0.4190</u>	<u>2</u>			
Team 7	<i>0.3833</i>	3			

Table 6: Spearman’s correlations of the baseline and average correlations of each team’s metrics on **turn-level** and **dialogue-level** test sets for **task 2**. The first position is shown in bold, the second in underline and the third in italics.

The best presentation according to the overall score calculated was used to rank each team. Teams 4, 6 and 7 were the best performers at the turn-level. At dialogue-level, the baseline model provided the best performance, followed by team 4 and team 3.

Considering team 4 results, both in task 1 and task 2, it can be considered their model as the overall best in the competition, being good at multilingual level as well as in robustness. However, the performance of the baseline model at dialogue level is far superior to that of team 4, showing there is still room for improvement.

4 Conclusions and Future Work

This paper presents a comprehensive overview of Track 4 on "Robust and Multilingual Automatic Evaluation Metrics for Open-Domain Dialogue Systems" organized as part of the 11th Dialogue System Technology Challenge (DSTC11). The track was divided into two subtasks addressing an important problems in Dialogue Systems: the design of automatic evaluation metrics for multilingual dialogues and dialogue robustness when dealing with paraphrases or back-translations.

Each task was divided at turn and dialogue level. At the turn level, 5 teams actively participated and at the dialogue level 3 teams participated. Having some of the teams participated at both levels. Some of the teams obtained interesting results that effectively contribute to the state-of-the-art of automatic evaluation models for multilingual dialogues. However, the results at the language level show a disparate performance among the different languages, giving room for improvement in the evaluation of other languages. The overall performance of the participants was satisfactory, with some teams outperforming the baseline model both in language and globally, as well as at the turn and dialogue levels. However, we can see that the automatic evaluation is still an open problem as correlation scores are still below 0.7 in the best of the cases.

The second task was also subdivided at turn and dialogue level. At the turn level, 5 teams actively participated and at dialogue-level 3 teams, with some of the teams having participated at both levels. At the turn level, several teams outperformed the baseline model. However, no team was able to outperform the baseline model at dialogue-level, showing that there is still room for improvement.

As future work, we plan to increase the number of databases, as well as to provide better baseline models. We also want to include a larger number of dimensions so that the evaluations performed are more complete, covering more different aspects of the dialogue. For task 1, it is planned to extend the number of available languages, to create multilingual models with a wider spectrum, thus widening the scope of the competition and attracting more participants who are fluent in other languages. For task 2 we want to propose higher quality paraphrases, such as those generated with models like GPT-4 (OpenAI, 2023).

Acknowledgements

This work is supported by project BEWORD (PID2021-126061OB-C43) funded by MCIN/AEI/10.13039/501100011033 and, as appropriate, by “ERDF A way of making Europe”, by the “European Union”, and by the European Commission through Project ASTOUND (101071191 — HORIZON-EIC-2021-PATHFINDERCHALLENGES-01). We gratefully acknowledge valuable efforts from Tencent AI Lab who supports Chinese translation and annotation of datasets by funding and infrastructure. Thanks to THU-CoAI (Conversational AI groups from Tsinghua University) for providing their Chinese datasets as part of the challenge data. Thanks to Unbabel for providing the COMET MTQE scores annotations as part of the challenge data. This contribution was supported by national funds through Fundação para a Ciência e a Tecnologia (FCT) with references PRT/BD/152198/2021 and UIDB/50021/2020, and by the P2020 program MAIA led by Unbabel (LISBOA-01-0247-FEDER-045909). We also give thanks to MS Azure services (especially to Irving Kwong) for their sponsorship to continue processing new datasets for the research community. This research project is supported by the NYU ChatEval Team led by João Sedoc. This research project is supported in part by a grant from Amazon to Alexander Rudnicky, Carnegie Mellon University. Thanks to Karthik Ganesan, Sarik Ghazarian, James Hagerty, Zhang Chen and Alex Rudnicky for developing the baseline model as part of the challenge tasks.

References

- Daniel Adiwardana, Minh-Thang Luong, David R. So, Jamie Hall, Noah Fiedel, Romal Thoppilan, Zi Yang, Apoorv Kulshreshtha, Gaurav Nemade, Yifeng Lu, and Quoc V. Le. 2020. [Towards a human-like open-domain chatbot](#).
- Rafael E. Banchs. 2012. [Movie-DiC: a movie dialogue corpus for research and development](#). In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 203–207, Jeju Island, Korea. Association for Computational Linguistics.
- Siqi Bao, Huang He, Fan Wang, Hua Wu, Haifeng Wang, Wenquan Wu, Zhihua Wu, Zhen Guo, Hua Lu, Xinxian Huang, Xin Tian, Xinchao Xu, Yingzhan Lin, and Zheng-Yu Niu. 2022. [PLATO-XL: Exploring the large-scale pre-training of dialogue generation](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2022*, pages 107–118, Online only. Association for Computational Linguistics.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Alessandra Cervone and Giuseppe Riccardi. 2020. [Is this dialogue coherent? learning from dialogue acts and entities](#). In *Proceedings of the 21th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 162–174, 1st virtual meeting. Association for Computational Linguistics.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Prithviraj Damodaran. 2021. Parrot: Paraphrase generation for nlu.
- Cristian Danescu-Niculescu-Mizil and Lillian Lee. 2011. Chameleons in imagined conversations: A new approach to understanding coordination of linguistic style in dialogs. In *Proceedings of the Workshop on Cognitive Modeling and Computational Linguistics, ACL 2011*.
- Jan Deriu, Alvaro Rodrigo, Arantxa Otegi, Guillermo Echevoyen, Sophie Rosset, Eneko Agirre, and Mark Cieliebak. 2021. Survey on evaluation methods for dialogue systems. *Artificial Intelligence Review*, 54:755–810.
- 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, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Emily Dinan, Stephen Roller, Kurt Shuster, Angela Fan, Michael Auli, and Jason Weston. 2019. [Wizard of wikipedia: Knowledge-powered conversational agents](#).
- Michel Galley, Chris Brockett, Xiang Gao, Jianfeng Gao, and Bill Dolan. 2019. Grounded response generation task at dstc7. In *AAAI Dialog System Technology Challenges Workshop*.

- Salvatore Giorgi, Shreya Havaladar, Farhan Ahmed, Zuhaib Akhtar, Shalaka Vaidya, Gary Pan, Lyle H. Ungar, H. Andrew Schwartz, and Joao Sedoc. 2023. [Human-centered metrics for dialog system evaluation](#).
- Karthik Gopalakrishnan, Behnam Hedayatnia, Qinlang Chen, Anna Gottardi, Sanjeev Kwatra, Anu Venkatesh, Raefer Gabriel, and Dilek Hakkani-Tür. 2019. [Topical-Chat: Towards Knowledge-Grounded Open-Domain Conversations](#). In *Proc. Interspeech 2019*, pages 1891–1895.
- Prakhar Gupta, Shikib Mehri, Tiancheng Zhao, Amy Pavel, Maxine Eskenazi, and Jeffrey P. Bigham. 2019. [Investigating evaluation of open-domain dialogue systems with human generated multiple references](#).
- Carmen Haro, Oriol A Rangel-Zúñiga, Juan F Alcalá-Díaz, Francisco Gómez-Delgado, Pablo Pérez-Martínez, Javier Delgado-Lista, Gracia M Quintana-Navarro, Blanca B Landa, Juan A Navas-Cortés, Manuel Tena-Sempere, et al. 2016. Intestinal microbiota is influenced by gender and body mass index. *PLoS one*, 11(5):e0154090.
- Ryuichiro Higashinaka, Kotaro Funakoshi, Yuka Kobayashi, and Michimasa Inaba. 2016. [The dialogue breakdown detection challenge: Task description, datasets, and evaluation metrics](#). In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 3146–3150, Portorož, Slovenia. European Language Resources Association (ELRA).
- Chao-Chun Hsu, Sheng-Yeh Chen, Chuan-Chun Kuo, Ting-Hao Huang, and Lun-Wei Ku. 2018. [Emotion-Lines: An emotion corpus of multi-party conversations](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- 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)*, pages 9230–9240, Online. Association for Computational Linguistics.
- S Lee, H Schulz, A Atkinson, J Gao, K Suleman, L El Asri, M Adada, M Huang, S Sharma, W Tay, et al. 2019. Multi-domain task-completion dialog challenge. *Dialog system technology challenges*, 8(9).
- Yanran Li, Hui Su, Xiaoyu Shen, Wenjie Li, Ziqiang Cao, and Shuzi Niu. 2017. [DailyDialog: A manually labelled multi-turn dialogue dataset](#). In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 986–995, Taipei, Taiwan. Asian Federation of Natural Language Processing.
- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Chia-Wei Liu, Ryan Lowe, Iulian Serban, Mike 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 *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2122–2132, Austin, Texas. Association for Computational Linguistics.
- Shikib Mehri, Jinho Choi, Luis Fernando D’Haro, Jan Deriu, Maxine Eskenazi, Milica Gasic, Kallirroi Georgila, Dilek Hakkani-Tur, Zekang Li, Verena Rieser, Samira Shaikh, David Traum, Yi-Ting Yeh, Zhou Yu, Yizhe Zhang, and Chen Zhang. 2022. [Report from the nsf future directions workshop on automatic evaluation of dialog: Research directions and challenges](#).
- Shikib Mehri and Maxine Eskenazi. 2020a. [Unsupervised evaluation of interactive dialog with DialogGPT](#). In *Proceedings of the 21th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 225–235, 1st virtual meeting. Association for Computational Linguistics.
- Shikib Mehri and Maxine Eskenazi. 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*, pages 681–707, Online. Association for Computational Linguistics.
- Erinc Merdivan, Deepika Singh, Sten Hanke, Johannes Kropf, Andreas Holzinger, and Matthieu Geist. 2020. Human annotated dialogues dataset for natural conversational agents. *Applied Sciences*, 10(3):762.
- Nikita Moghe, Siddhartha Arora, Suman Banerjee, and Mitesh M. Khapra. 2018. [Towards exploiting background knowledge for building conversation systems](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2322–2332, Brussels, Belgium. Association for Computational Linguistics.
- Yixin Nie, Mary Williamson, Mohit Bansal, Douwe Kiela, and Jason Weston. 2020. I like fish, especially dolphins: Addressing contradictions in dialogue modeling. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020*.
- OpenAI. 2023. [Gpt-4 technical report](#).
- 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.

- Soujanya Poria, Devamanyu Hazarika, Navonil Majumder, Gautam Naik, Erik Cambria, and Rada Mihalcea. 2019. [MELD: A multimodal multi-party dataset for emotion recognition in conversations](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 527–536, Florence, Italy. Association for Computational Linguistics.
- Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. 2018. Improving language understanding by generative pre-training.
- 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.
- 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 Annual Meeting of the Association for Computational Linguistics*, pages 5370–5381, Florence, Italy. 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)*, pages 2685–2702, Online. Association for Computational Linguistics.
- Alexander Rudnicky, Rafael Banchs, Luis F. D’Haro, João Sedoc, Zhang Chen, Mario Rodríguez-Cantelar, Andrew Koh Jin Jie, et al. 2020. [Chanel-metrics: Chat/dialogue modeling and evaluation report](#).
- João Sedoc, Daphne Ippolito, Arun Kirubarajan, Jai Thirani, Lyle Ungar, and Chris Callison-Burch. 2019. [ChatEval: A tool for chatbot evaluation](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 60–65, Minneapolis, Minnesota. Association for Computational Linguistics.
- Abigail See, Stephen Roller, Douwe Kiela, and Jason Weston. 2019. [What makes a good conversation? how controllable attributes affect human judgments](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1702–1723, Minneapolis, Minnesota. Association for Computational Linguistics.
- Thibault Sellam, Dipanjan Das, and Ankur Parikh. 2020. [BLEURT: Learning robust metrics for text generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7881–7892, Online. Association for Computational Linguistics.
- Lifeng Shang, Zhengdong Lu, and Hang Li. 2015. [Neural responding machine for short-text conversation](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1577–1586, Beijing, China. Association for Computational Linguistics.
- Kurt Shuster, Jing Xu, Mojtaba Komeili, Da Ju, Eric Michael Smith, Stephen Roller, Megan Ung, Moya Chen, Kushal Arora, Joshua Lane, Morteza Behrooz, William Ngan, Spencer Poff, Naman Goyal, Arthur Szlam, Y-Lan Boureau, Melanie Kambadur, and Jason Weston. 2022. [Blenderbot 3: a deployed conversational agent that continually learns to responsibly engage](#).
- Tianxiang Sun, Junliang He, Xipeng Qiu, and Xuanjing Huang. 2022. [BERTScore is unfair: On social bias in language model-based metrics for text generation](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 3726–3739, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Bibek Upadhyay and Vahid Behzadan. 2020. Sentimental liar: Extended corpus and deep learning models for fake claim classification. In *2020 IEEE International Conference on Intelligence and Security Informatics (ISI)*, pages 1–6. IEEE.
- Yida Wang, Pei Ke, Yinhe Zheng, Kaili Huang, Yong Jiang, Xiaoyan Zhu, and Minlie Huang. 2020a. A large-scale chinese short-text conversation dataset. In *Natural Language Processing and Chinese Computing: 9th CCF International Conference, NLPCC 2020, Zhengzhou, China, October 14–18, 2020, Proceedings, Part I 9*, pages 91–103. Springer.
- Yida Wang, Pei Ke, Yinhe Zheng, Kaili Huang, Yong Jiang, Xiaoyan Zhu, and Minlie Huang. 2020b. A large-scale chinese short-text conversation dataset. In *Natural Language Processing and Chinese Computing*, pages 91–103, Cham. Springer International Publishing.
- Yu Wu, Wei Wu, Chen Xing, Ming Zhou, and Zhoujun Li. 2017. [Sequential matching network: A new architecture for multi-turn response selection in retrieval-based chatbots](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 496–505, Vancouver, Canada. Association for Computational Linguistics.
- Guangxuan Xu, Ruibo Liu, Fabrice Harel-Canada, Nischal Reddy Chandra, and Nanyun Peng. 2022. [EnDex: Evaluation of dialogue engagingness at scale](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 4884–4893, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Yi-Ting Yeh, Maxine Eskenazi, and Shikib Mehri. 2021. [A comprehensive assessment of dialog evaluation](#)

- metrics. In *The First Workshop on Evaluations and Assessments of Neural Conversation Systems*, pages 15–33, Online. Association for Computational Linguistics.
- Chen Zhang, Luis Fernando D’Haro, Thomas Friedrichs, and Haizhou Li. 2022a. MDD-Eval: Self-training on augmented data for multi-domain dialogue evaluation. In *AAAI 2022*.
- Chen Zhang, Luis Fernando D’Haro, Qiquan Zhang, Thomas Friedrichs, and Haizhou Li. 2022b. **FineD-eval: Fine-grained automatic dialogue-level evaluation**. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 3336–3355, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Chen Zhang, Luis Fernando D’Haro, Rafael E Banchs, Thomas Friedrichs, and Haizhou Li. 2021. Deep amfm: Toolkit for automatic dialogue evaluation. *Conversational Dialogue Systems for the Next Decade*, pages 53–69.
- Chen Zhang, João Sedoc, Luis Fernando D’Haro, Rafael Banchs, and Alexander Rudnicky. 2022c. Automatic evaluation and moderation of open-domain dialogue systems. In *AAAI 2022*.
- Lining Zhang, João Sedoc, Simon Mille, Yufang Hou, Sebastian Gehrmann, Daniel Deutsch, Elizabeth Clark, Yixin Liu, Miruna Clinciu, Saad Mahamood, and Khyathi Chandu. 2022d. **Needle in a haystack: An analysis of finding qualified workers on mturk for summarization**.
- Saizheng Zhang, Emily Dinan, Jack Urbanek, Arthur Szlam, Douwe Kiela, and Jason Weston. 2018a. **Personalizing dialogue agents: I have a dog, do you have pets too?** In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2204–2213, Melbourne, Australia. Association for Computational Linguistics.
- Yizhe Zhang, Siqi Sun, Michel Galley, Yen-Chun Chen, Chris Brockett, Xiang Gao, Jianfeng Gao, Jingjing Liu, and Bill Dolan. 2020. **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*, pages 270–278, Online. Association for Computational Linguistics.
- Zhuosheng Zhang, Jiangtong Li, Pengfei Zhu, Hai Zhao, and Gongshen Liu. 2018b. **Modeling multi-turn conversation with deep utterance aggregation**. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 3740–3752, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Tianyu Zhao, Divesh Lala, and Tatsuya Kawahara. 2020. Designing precise and robust dialogue response evaluators. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 26–33. Association for Computational Linguistics.
- Hao Zhou, Minlie Huang, Tianyang Zhang, Xiaoyan Zhu, and Bing Liu. 2018a. Emotional chatting machine: Emotional conversation generation with internal and external memory. In *AAAI 2018*.
- Hao Zhou, Minlie Huang, Tianyang Zhang, Xiaoyan Zhu, and Bing Liu. 2018b. **Emotional chatting machine: Emotional conversation generation with internal and external memory**. *Proceedings of the AAAI Conference on Artificial Intelligence*, 32(1).
- Hao Zhou, Chujie Zheng, Kaili Huang, Minlie Huang, and Xiaoyan Zhu. 2020a. **KdConv: A Chinese multi-domain dialogue dataset towards multi-turn knowledge-driven conversation**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7098–7108, Online. Association for Computational Linguistics.
- Kangyan Zhou, Shrimai Prabhumoye, and Alan W Black. 2018c. **A dataset for document grounded conversations**. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 708–713, Brussels, Belgium. Association for Computational Linguistics.
- Li Zhou, Jianfeng Gao, Di Li, and Heung-Yeung Shum. 2020b. **The design and implementation of XiaoIce, an empathetic social chatbot**. *Computational Linguistics*, 46(1):53–93.

A Appendix: Datasets statistics

Table 7 shows all the data sets that make up the training, development, and test sets. In detail, it shows the number of dialogues, turns per dialogue, average number of turns per dialogue, average number of words per turn, as well as the granularity of the annotations (at turn and/or dialogue level), the original language of the dataset and into which languages it is translated.

B Appendix: Existing Benchmark Datasets

Descriptions of the datasets that constitute the DSTC10 benchmark can be found at [Zhang et al. \(2022c\)](#). Details of the remaining evaluation datasets are as follows:

ECM-Eval - The test instances in ECM-Eval test set are sampled from the Emotional Short-Text Conversation (ESTC) dialogue corpus ([Zhou et al., 2018b](#)), which is built on top of the Short-Text Conversation (STC) dataset ([Shang et al., 2015](#)). ESTC is designed to build Chinese empathetic dialogue systems. The dialogues are crawled from

Name	#Turns	#Dialogues	Average Turn/Dial	Average Words/Turn	Annotation Granularity	Original Language	Translation
Train							
DBDC (Higashinaka et al., 2016)	8,509	415	20.5	7.31	Turn	En	Zh/Es
CMU_DoG (Zhou et al., 2018c)	95,305	4,221	22.58	17.93	Turn	En	Zh/Es
Cornell Movie-Dialogs (Danescu-Niculescu-Mizil and Lee, 2011)	304,713	83,097	3.67	13.72	Turn	En	Zh/Es
DailyDialog (Li et al., 2017)	102,960	13,116	7.85	13.96	Turn	En	Zh/Es
DECODE (Nie et al., 2020)	296,105	35,426	8.36	15.05	Turn	En	Zh/Es
EmotionLines (Hsu et al., 2018)	14,503	1,000	14.50	10.53	Turn	En	Zh/Es
EmpathicDialogues (Rashkin et al., 2019)	107,220	24,850	4.31	15.88	Turn	En	Zh/Es
Holl-E (Moghe et al., 2018)	91,452	9,071	10.08	17.74	Turn	En	Zh/Es
MEENA (Adiwardana et al., 2020)	3,675	193	19.04	9.14	Turn	En	Zh/Es
MELD (Poria et al., 2019)	23,197	1,592	14.57	10.98	Turn	En	Zh/Es
MetalWOz (Lee et al., 2019)	432,036	37,884	11.40	8.47	Turn	En	Zh/Es
Movie-DiC (Banchs, 2012)	512,582	65,215	7.86	13.82	Turn	En	Zh/Es
PersonaChat (Zhang et al., 2018a)	162,064	10,907	14.86	11.72	Turn	En	Zh/Es
SentimentLIAR (Upadhayay and Behzadan, 2020)	12,781	12,781	1.00	20.16	Turn	En	Zh/Es
Switchboard Coherence (Cervone and Riccardi, 2020)	12,059	1,000	12.06	20.55	Turn	En	Zh/Es
Topical-Chat (Gopalakrishnan et al., 2019)	235,281	10,784	21.82	23.23	Turn	En	Zh/Es
Wizard of Wikipedia (Dinan et al., 2019)	201,999	22,311	9.05	18.83	Turn	En	Zh/Es
Wochat (Haro et al., 2016)	19,881	607	32.75	6.75	Turn	En	Zh/Es
Total	2,636,322	334,470	236.26	255.77			
Development							
ConvAI2-GRADE (Huang et al., 2020)	1,800	600	3.0	12.07	Turn	En	Zh/Es
DailyDialog-GRADE (Huang et al., 2020)	900	300	3.0	12.60	Turn	En	Zh/Es
DailyDialog-GUPTA (Gupta et al., 2019)	2,460	500	4.92	12.37	Turn	En	Zh/Es
DailyDialog-ZHAO (Zhao et al., 2020)	4,248	900	4.72	12.41	Turn	En	Zh/Es
DSTC7 (Galley et al., 2019)	34,650	9,990	3.47	15.39	Turn	En	Zh/Es
Empathetic-GRADE (Huang et al., 2020)	900	300	3.0	16.65	Turn	En	Zh/Es
FED-Dial (Mehri and Eskenazi, 2020a)	1,715	125	13.72	11.1	Dial	En	Zh/Es
FED-Turn (Mehri and Eskenazi, 2020a)	3,888	375	10.37	10.78	Turn	En	Zh/Es
HUMOD (Merdivan et al., 2020)	37,468	9,499	3.94	7.97	Turn	En	Zh/Es
Persona-SEE (See et al., 2019)	39,792	3,316	12.0	9.0	Dial	En	Zh/Es
PersonaChat-USR (Mehri and Eskenazi, 2020b)	2,790	300	9.3	12.08	Turn	En	Zh/Es
PersonaChat-ZHAO (Zhao et al., 2020)	4,614	900	5.13	12.06	Turn	En	Zh/Es
TOPICAL-USR (Mehri and Eskenazi, 2020b)	4,032	360	11.2	23.16	Turn	En	Zh/Es
ECM-Eval (Zhou et al., 2018a)	3,004	1,502	2.0	13.13	Turn	Zh	En
KdConv-Eval (Zhou et al., 2020a)	3,499	354	9.88	21.11	Turn	Zh	En
LCCC-Eval (Wang et al., 2020a)	3,009	589	5.11	11.72	Turn	Zh	En
Total	148,769	29,910	104.76	212.64			
Test							
BlenderBot3 (Giorgi et al., 2023; Shuster et al., 2022)	679	21	32.33	16.96	Turn/Dial	En	Zh/Es
ChatGPT (Giorgi et al., 2023; Radford et al., 2018)	462	21	22	91.07	Turn/Dial	En	Zh/Es
GPT-3.5 (Giorgi et al., 2023; Brown et al., 2020)	560	17	32.94	23.73	Turn/Dial	En	Zh/Es
HCCChinese	2,017	187	10.79	8.08	Turn/Dial	Zh	En
ChatEval (Sedoc et al., 2019)	400	200	2	8.13	Turn	En	Zh/Es
DSTC10 (Zhang et al., 2022c)	112	28	4	14	Turn	En	Zh/Es
JSALT (Rudnicky et al., 2020)	46	13	3.54	17.26	Turn	En	Zh/Es
Total	4,276	487	107.60	179.23			

Table 7: Summary of the train, development, and test datasets. Some information comes from Yeh et al. (2021).

Weibo and post-processing, such as the removal of trivial responses and filtering out potential advertisements, has been conducted by Shang et al. (2015). The dialogues are automatically annotated by pre-trained emotion classifiers along six different emotion categories, such as angry, happy, sad, etc. The dialogues in ESTC are much shorter. Most contain only a single post-response pair.

LCCC-Eval - Data in LCCC-Eval are sampled from the Large-scale Cleaned Chinese Conversation dialogue corpus (LCCC) (Wang et al., 2020b). The LCCC corpus is designed for pretraining the Chinese dialogue model. The dialogues are mainly collected from Weibo, a Chinese microblogging

website¹⁷ and other open-source Chinese dialogue corpora, such as the Douban Conversation (Wu et al., 2017) and the E-Commerce Conversation Corpus (Zhang et al., 2018b). All the dialogues belong to the general domain and a rigorous cleaning process, which is based on a series of heuristic rules and several classifiers, is conducted to filter out dialogues with noise, such as dirty words, special characters, facial expressions, ungrammatical sentences, etc. Both ESTC and LCCC are released by the THU-COAI group for research purposes at <https://www.luge.ai/#/>

KdConv-Eval - KdConv-Eval is constructed based

¹⁷https://en.wikipedia.org/wiki/Sina_Weibo

on the KdConv corpus (Zhou et al., 2020a), a multi-domain Chinese dialogue dataset towards multi-turn knowledge-driven conversation. The corpus links the subjects of multi-turn discussions to knowledge graphs. It encompasses conversations from three categories (movies, music, and travel). These conversations involve detailed exchanges about relevant subjects and seamlessly move between a variety of topics. We sampled 354 dialogues from the original corpus to form the KdConv-Eval test dataset.

HCChinese - Dialogues in HCChinese are collected by interacting with three state-of-the-art Chinese chatbots, Baidu Plato-XL (Bao et al., 2022), Microsoft XiaoIce (Zhou et al., 2020b), and a Chinese DialoGPT model that is trained in a similar manner to DialoGPT (Zhang et al., 2020). We chat with the chatbots on a diverse set of topics, such as entertainment, relationship, arts, travel, food, etc. The discussion of sensitive topics, such as politics and race, was avoided. A manual check is performed on each dialogue, and those containing inappropriate responses were filtered out. In the end, we collected 531 human-chatbot multi-turn conversations with 207 from Plato-XL, 224 from XiaoIce, and 100 dialogues from the Chinese DialoGPT.

TBD-Q1-2023 Three Bot Dialog Evaluation Corpus (TBD-Q1-2023 OR TBD; Quarter 1 of 2023) from Giorgi et al. (2023) consists of dialogues with three chatbots: ChatGPT (Radford et al., 2018), GPT-3 (Brown et al., 2020), and BlenderBot3 (Shuster et al., 2022). Student participants were told to have a long conversation with the chatbots on a range of topics of their choosing. All conversations were collected between November 2022 and March 2023. They collected 21 dialogues with an average of 14.6 turns per dialogue. Conversations for **BlenderBot3** were directly from the website <https://blenderbot.ai/>. The **ChatGPT** conversations were taken directly from the website <https://chat.openai.com/>. Finally, **GPT3** used the text-davinci-003 model with the prompt *Hal is a chatbot that attempts to answer questions with useful responses:*. The GPT-3 model parameters were temperature of 0.5, max tokens of 289, top-p of 0.3, frequency penalty of 0.5, and presence penalty of 0.