

Simple LLM Prompting is State-of-the-Art for Robust and Multilingual Dialogue Evaluation

John Mendonça^{1,2,*}, Patrícia Pereira^{1,2},
Helena Moniz^{1,3}, João Paulo Carvalho^{1,2}, Alon Lavie^{4,5} and Isabel Trancoso^{1,2}

¹ INESC-ID, Lisbon

² Instituto Superior Técnico, University of Lisbon

³ Faculdade de Letras, University of Lisbon

⁴ Carnegie Mellon University, Pittsburgh

⁵ Phrase, Pittsburgh

john.mendonca@inesc-id.pt

Abstract

Despite significant research effort in the development of automatic dialogue evaluation metrics, little thought is given to evaluating dialogues other than in English. At the same time, ensuring metrics are invariant to semantically similar responses is also an overlooked topic. In order to achieve the desired properties of robustness and multilinguality for dialogue evaluation metrics, we propose a novel framework that takes advantage of the strengths of current evaluation models with the newly-established paradigm of prompting Large Language Models (LLMs). Empirical results show our framework achieves state of the art results in terms of mean Spearman correlation scores across several benchmarks and ranks **first place on both the Robust and Multilingual tasks** of the DSTC11 Track 4 “Automatic Evaluation Metrics for Open-Domain Dialogue Systems”, proving the evaluation capabilities of prompted LLMs.

1 Introduction

Automatic dialogue evaluation has largely been focused on evaluating select few languages. The main reason for this constraint is the lack of linguistic diversity in dialogue corpora, which leads to a lack of chatbots that cover other languages. As a result, the need for multilingual metrics has also been limited.

A possible solution to this issue is to leverage the latest batch of Large Language Models (LLMs) to synthetically generate multilingual dialogues. Some research has already been conducted to study the capabilities of these models (Guo et al., 2023; Bubeck et al., 2023) and the consensus appears to be that these models have achieved a proxy of a *formal* linguistic competence in the most studied

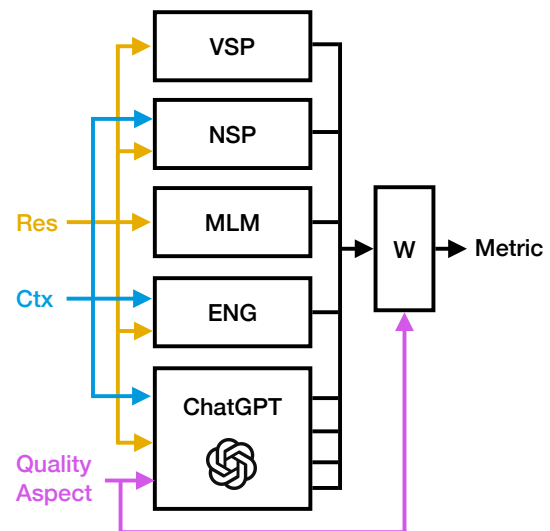


Figure 1: Proposed framework architecture. The **Response**, **Context** and **Quality Aspect** under evaluation are fed to the submetrics: **VSP** (Valid Sentence Prediction), **NSP** (Next Sentence Prediction), **MLM** (Masked Language Modelling), **ENG** (Engagement) and **ChatGPT**. Each submetric score is then weighted according to the aspect, yielding the final metric.

languages. That is, its responses follow linguistic conventions and are fluent and grammatical, but they might be inaccurate or even hallucinate (Guerreiro et al., 2023). More importantly, pertaining to dialogue, they also show signs of *functional* linguistic competence in its responses, i.e., discursive coherence, narrative structure and linguistic knowledge, even if not fully consistent (sometimes they do not consider context or situated information, and fail to adapt to users and domains).

Irrespective of these models’ limitations, it is clear their emergent capabilities allow for the development of chatbots with capabilities vastly beyond what earlier models were able to achieve. Yet,

* Work conducted as a visiting scholar at CMU.

an interesting research question lingers: *If these models are able to write responses that follow formal and functional linguistics rules, are they also capable of evaluating responses/dialogues in terms of these same rules?* Prior work has confirmed the language understanding capabilities of instruction-based LLMs for dialogue evaluation (Huynh et al., 2023). However, we are the first to study the evaluation capabilities of the newest batch of LLMs in terms of multilinguality and paraphrase robustness.

This paper presents our contribution to the DSTC11 track on Robust and Multilingual Automatic Evaluation Metrics for Open-Domain Dialogue Systems (Rodríguez-Cantelar et al., 2023), where we participated in both the Multilingual and Robustness tasks. This track is an excellent venue to benchmark the capabilities of these new LLMs for dialogue evaluation, as it evaluates properties that have been observed in these models. We propose a comprehensive framework, incorporating earlier encoder-based approaches and ChatGPT, as illustrated in Figure 1. By combining multiple models and submetrics through ensembling, our approach aims to improve the performance and robustness of dialogue evaluation, ultimately contributing to the advancement of dialogue system research and development.

Overall, our contributions are the following:

- We show that ChatGPT is a *strong* evaluator of dialogues, outperforming typical encoder frameworks.
- We propose a new framework for dialogue evaluation that is multilingual and robust to paraphrases. In fact, our combined Encoder and ChatGPT framework ranks **1st place** on both the Multilingual and Robust metrics task.
- We discuss the outlook of Dialogue Evaluation in this new realm of LLMs.
- We open source the code and checkpoints of the submetrics at github.com/johndmendonca/DialEvalML.

2 Related work

2.1 Automatic Evaluation Metrics

Statistic-based metrics such as BLEU (Papineni et al., 2002), ROUGE (Lin, 2004), and METEOR (Banerjee and Lavie, 2005), are a popular choice to evaluate NLG (Natural Language Generation)

models as they are easy to employ. These metrics assume valid responses have significant word-overlap with the ground truth. However, this is not a valid assumption for dialogue: there are many equally good responses for a single utterance. As such, the correlation with Human Evaluation (HE) annotations is very low for these metrics (Liu et al., 2016), and they cannot be used to evaluate models whenever a gold-response is not available.

Earlier learned metrics such as ADEM (Lowe et al., 2017) and RUBER (Tao et al., 2018) explicitly predict HE annotations by initialising pre-trained Recurrent Neural Network response generators. Unlike ADEM, which is trained with HE-annotated data in a supervised manner, RUBER leverages negative samples. In both cases, a reference response is used to score the candidate response. As such, these metrics still suffer the same issues as word-overlap metrics.

The primary motivation for the negative sampling approach in RUBER was the need for extensive HE annotations in ADEM. Approaches similar to this are now the norm for training open-domain dialogue evaluation metrics. By using well-defined self-supervised tasks which correlate well with their corresponding aspects, the annotation limitations are mostly circumvented.

The most widely used self-supervised task is Next Sentence Prediction (NSP), as it is known to correlate well with HE that evaluate "*Context Awareness*". The typical approach is to finetune a pretrained encoder model with this automatically generated data (Mehri and Eskenazi, 2020b; Phy et al., 2020; Mendonca et al., 2022; Zhao et al., 2020; Zhang et al., 2022). More complex approaches leverage graph representations to model dialogue interactions explicitly (Huang et al., 2020; Zhang et al., 2021a). Another typically employed self-supervised task is Valid Sentence Prediction (VSP), which uses word-level noising techniques to generate negative samples and correlates well with HE that evaluate *Fluency* (Phy et al., 2020; Mendonca et al., 2022; Zhang et al., 2022).

Parallel to this trend, other annotation-free approaches in the literature have surfaced. For instance, qualities such as *Specificity* correlate reasonably well with metrics obtained directly from the MLM (Masked Language Modelling) loss calculated using pretrained encoder models (Mehri and Eskenazi, 2020b; Phy et al., 2020; Zhang et al., 2022).

Given the multifaceted nature of dialogue, dialogue quality metrics typically employ a combination of submetrics. Mehri and Eskenazi (2020a) leverage follow-up utterance from a pretrained decoder model to calculate 18 turn and dialogue-level submetrics, which are then used as inputs to a regression model for overall quality. In fact, Linear Regression is frequently used as a feature aggregation method in the literature (Jiang et al., 2022; Mehri and Eskenazi, 2020b). Alternatively, Phy et al. (2020) propose a hierarchical composition where they incorporate the quality aspects together in a way that aspects in the lower hierarchy need to be satisfied before aspects higher up are considered. Also worth mentioning is the work of Zhang et al. (2022), which proposes the so called Correlation Re-scaling method. Here, the contribution of each aspect is calculated from the individual correlations of the submetrics, obtained from a subset of HE.

2.2 Large Language Models

The widespread use of LLMs was established, practically speaking, with the work of Devlin et al. (2019), where a transformer architecture (Vaswani et al., 2017) is pretrained with substantial amounts of unlabelled text with a Masked Language Modelling (MLM) objective. With this architecture, a new paradigm in NLP surfaced, where the adaptation to downstream tasks was conducted by finetuning the pretrained model with supervised data. Later on, GPT-3 (Brown et al., 2020), which is trained with an autoregressive objective, showed competitive results by leveraging few-shot prompting. Nevertheless, given their training objective function, it was difficult for autoregressive LLMs to successfully perform downstream NLP tasks without substantial prompt engineering.

Ouyang et al. (2022) propose finetuning GPT-3 using a 3-step approach named Reinforcement Learning through Human Feedback (RLHF). In detail, the model is (1) initially finetuned using supervised data obtained from labelling prompts (SFT); (2) a reward model is trained using ranked responses given a prompt; (3) the policy is optimised against the reward model using the Proximal Policy Optimisation reinforcement learning algorithm (Schulman et al., 2017). As a testament to the power of this approach, ChatGPT took the world by storm in late 2022 thanks to its incredible human-like generation capabilities. This was achieved by including dialogues in all steps of RLHF.

3 Problem Formulation

The main goal of this track was to develop and benchmark automatic open-ended dialogue evaluation metrics. Two tasks were proposed this year, Metrics for Multilingual Data and Robust metrics. For the Metrics for Multilingual Data task, participants were asked to construct quality metrics that perform well on a multilingual setup. For the Robust metrics task, the goal was to develop metrics that perform robustly when evaluated over back-translated/paraphrased sentences in English.

In both tasks, the proposed metrics were evaluated at the turn and dialogue level, without access to a reference. In a *turn-level* evaluation setting, the goal is, given prior dialogue history (frequently denoted as context) c of varying amount of turns, and a response r , to learn a scoring function (also known as metric) that assigns a score $f(c, r) \rightarrow s$. Conversely, in a *dialogue-level* evaluation setting, the goal is to evaluate the performance throughout the full dialogue.

Irrespective of the level of evaluation, the proposed metrics' outputs are typically compared against HE annotations that use a Likert scale, where the lowest value means lowest quality and highest value maximum quality. For this track, the performance of these metrics was evaluated by calculating the Pearson correlation between the calculated score and HE.

4 Methodology

Our framework, which we call DIALEVALML, can be viewed as a dual layered ensemble which are done at the **model** and **submetric** level, and that employ strong multilingual pretrained encoder and decoder models which were finetuned or prompted¹. In this section, we describe the step-by-step process of DIALEVALML, detailing the various components and methods employed.

4.1 Submetrics

Similar to other frameworks, including the best performing ones in last year's track (Zhang et al., 2022; Jiang et al., 2022) which take inspiration from the works of Phy et al. (2020); Sinha et al. (2020); Mehri and Eskenazi (2020b), we employ

¹We tried experimenting with metrics that use graph representations, but found implementing these metrics to be Multilingual and Robust, and including them in our framework, to be impractical, not to mention detrimental to performance in some instances.

several submetrics to evaluate dialogue responses – ranging from zero-shot prediction using pretrained LLMs to trained models using self-supervised and supervised methods – and weigh them according to the aspect we wish to predict.

4.1.1 VSP: Valid Sentence Prediction

Following [Sinha et al. \(2020\)](#), we train a regression model that is optimised to differentiate between positive samples and synthetic negative samples. **Positive** samples are perturbed by randomly applying one of the following: (1) no perturbation, (2) punctuation removal, (3) stop-word removal. **Negative** samples are generated by randomly applying one of the following rules: (1) word reorder (shuffling the ordering of the words); (2) word-drop; and (3) word-repeat (randomly repeating words).

4.1.2 NSP: Next Sentence Prediction

With the binary **NSP** (Next Sentence Prediction) task, the goal is to distinguish a positive example from a semantically negative one, given a context. We train a discriminative regression model using the following sampling strategy: **positive** responses are drawn directly from the dialog; **negative** responses are randomly selected and a token coverage test discards semantically similar sentences. All responses are processed using the positive-sample heuristic used by **VSP**.

For both tasks, the underlying goal is that paraphrased and/or translated responses should have the same coherence score as the original response, since they (in theory) convey the same message. In order to increase the robustness of our framework to paraphrased responses we propose a Siamese Neural Network. Simply put, we train an encoder model (denoted NSP-Siamese) to jointly optimise a Cosine Embedding Loss between the hidden states of the encoder model for the original and a paraphrase, and the individual errors between the predictions and the ground truth. We hypothesise this enables the model to compare the semantic coherence of the responses w.r.t the context, instead of more spurious features such as syntax.

A similar approach could’ve been employed for multilingual metrics, however, scaling to more languages is computationally expensive: one would either need a new model for each language, or a training procedure requiring a forward pass for each language, for each example.

4.1.3 MLM: Masked Language Modelling

Similar to [Mehri and Eskenazi \(2020b\)](#); [Phy et al. \(2020\)](#), we use a pretrained encoder model to calculate the MLM loss of all tokens of the response. The resulting **MLM** submetric is calculated as the sum of the individual losses.

4.1.4 ENG: Engagement

An important quality aspect of dialogue that is frequently overlooked is *Engagement*. Some work attempt to equate this aspect with *Specificity* and related metrics. However, we argue this is a reductive solution, as engagement is an abstract and multi-dimensional concept, thereby making a surface level evaluation of the response in terms of diversity insufficient.

As such, and following the methodology used for **VSP** and **NSP**, we train a discriminate model using RED (Reddit-based Engagement Dataset) ([Xu et al., 2022](#)) which we then use as a submetric denoted in our framework as **ENG**. This dataset is sourced from Reddit and is curated using a novel distant-supervision framework. This framework aggregates emotional, attentional, behavioural and reply engagement onto a single score denoted EN-DEX, which then has a hyperparameter threshold applied to it to cluster posts into positive and negative samples.

4.2 Exploiting Data Augmentation for Robust and Multilingual Evaluation

The main novelty of this year’s track is the release of training and development dialogue data that has been augmented with MT (Machine Translation) – for the Multilingual task – and Paraphrases – for the Robust task. These augmentations are subsequently scored to determine similarity against the original data: for MT, several COMET QE (Quality Estimation) scores ([Rei et al., 2020](#); [Zerva et al., 2021](#); [Rei et al., 2022](#)) were provided; for Paraphrases, the organisers provided cosine similarity scores of the sentence embeddings.

A naive approach to obtain competitive metrics in both tasks would be to simply introduce the full amount of augmented data during self-supervised and supervised training. However, [Mendonca et al. \(2023\)](#) showed that low quality augmentation affects the performance of models trained on MT augmented data, especially for **VSP**. Following this work, we select 5 and 75 % of the best translated data (ranked using COMET QE) for training of the **VSP** and **NSP** models respectively. For **ENG**, we

train different proportions of data and select the best performing ones.

4.2.1 ChatGPT

We briefly experimented with different prompts, and found the best performing prompt (irrespective of language) in a held-out internal set to be simply:

- Turn-level: *"Given the Context, evaluate from 1-5 the Response in terms of {aspect}. Provide a single score and nothing else."*
- Dialogue-level: *"Evaluate the following dialogue from 1-5 in terms of {aspect}. Provide a single score and nothing else."*

Unlike GPT-3, the API for ChatGPT does not output the log probabilities of the most likely tokens. As such, the measurement of quality is non-deterministic. We attempt to reduce output variability by reinforcing the desired output in the prompt (*"Provide a single score and nothing else."*) and by setting the temperature to 0. We report a mean absolute deviation of 0.0182 across 3 runs when querying *Appropriateness* on the provided en/dailydialog-grade dataset included in the development set. To facilitate ensembling in later stages, we normalise the predictions to [0,1].

The default processing step consists of searching for an integer in the response. However, there are some instances where ChatGPT fails to output the desired score: (1) When conducting dialogue level evaluation, the model sometimes outputs scores for each individual response. In these cases, we calculate the average score, similar to the dialogue-level encoder scores. (2) Less frequently, ChatGPT ignores the task and continues the conversation. Here, we prompt the model again until a score is provided.

4.3 Submetric Ensembling

Despite having a key role in NLG evaluation, HE has been performed while suffering from nontransparent and inconsistent annotation procedures. As such, annotations from different works one expects to report the same quality are frequently only nominal in nature. A good example is *Coherence*, with some definitions referring to it as (1) semantic relevance with respect to a previous sentence; (2) a theme/topic; or even (3) *Readability*, which is considered a different quality in other guidelines. [Howcroft et al. \(2020\)](#) provides an in-depth survey

of 165 NLG papers with human evaluations where these issues are highlighted.

Taking into account these facts, it is not clear we can successfully apply an empirical surjective mapping function from our submetrics to the quality aspects. Instead, we take a data-driven approach to generate this mapping, similar to the one proposed in [Zhang et al. \(2022\)](#). The main difference between the original Correlation Re-Scaling method and our approach is that, instead of zeroing the weights of submetrics that have a negative correlation with the given aspect, we take a probabilistic approach where we conduct a statistic significance test, i.e., we check if the p -value is higher than a given threshold. This ensures submetrics which are strongly and negatively correlated with the aspect (for example, **MLM** and *Fluency*) are still included in the ensembling².

4.4 Dialogue-level Evaluation

We obtain dialogue-level quality predictions from the encoder models – **NSP**, **VSP**, **MLM** and **ENG** – by averaging the individual turn-level predictions. These are combined with the dialogue-level predictions obtained by prompting ChatGPT with the full dialogue in the prompt.

5 Experiments

5.1 Datasets

For data preprocessing we used [spaCy](#). For the **VSP** and **NSP** models, we followed prior work and base the self-supervised data on DailyDialog ([Li et al., 2017](#)). For the language specific and multilingual models, we rank the translations using the provided WMT22 scores. Models using paraphrased responses are trained using the least similar responses (lowest score³).

The **ENG** model was trained using the RED dataset, more specifically on the 80k split with negative sampled data ([Xu et al., 2022](#)). Given it is an English dataset, we use MBART50⁴ ([Liu et al., 2020](#)) to augment the original dataset with Spanish and Chinese MT. Finally, we score it using the WMT20-COMET-QE-DA model ([Rei et al., 2020](#)).

²For some annotations, none of the metrics were statistically significant. In these cases, we resort to the original proposed approach.

³We also trained models using the highest scoring responses and report lower performance. This is in line with our intuition that lower scoring responses are more diverse, and as such more informative for training.

⁴We chose MBART50 as it is lightweight and open source.

For the paraphrase augmentation, we follow the organisers’ approach of using Parrot Paraphraser (Damodaran, 2021) and scoring the paraphrases with Cosine Similarity.

5.2 Training and Hyperparameters

We used XLM-RoBERTa-large (Conneau et al., 2020) as the encoder model for the experiments. This model is the multilingual version of RoBERTa, pretrained on CommonCrawl data containing 100 languages. We used a single Quadro RTX 6000 24GB GPU for the encoder experiments, and accessed ChatGPT (gpt-3.5-turbo) in late March using the OpenAI API.

For the **VSP**, **NSP** and **ENG** metrics, a token representing the speaker was added for each turn, and a maximum history length of 3 turns was used during training. For predictions in the development and test sets we include the full conversational context whenever possible. If it surpasses input size limitations, we iteratively remove turns from the context, starting from the oldest one. We applied a regression head consisting of a 2-layer MLP with a hidden size of 1024 and a hyperbolic tangent function as activation for prediction. All parameters were trained/finetuned using Adam optimiser (Kingma and Ba, 2015). The fully finetuned models used a learning rate of 3e-6 and were trained for 3 epochs using a batch size of 16. Evaluation was conducted every 10,000 steps. The best performing model on the evaluation set was selected for testing. For the **MLM** metric, we used the existing LM head available in the Transformers library (Wolf et al., 2020).

With respect to the model-level ensembling, we conduct simple unweighted averaging of the predictions of the models. For the submetric-level ensembling, we define the mask threshold as $p > 0.05$ and square the correlations following Zhang et al. (2022). For testing, we define a mapping from the development quality aspects to the test-set aspects and obtain the final weights by averaging the weights obtained on the test set.

5.3 Model ensembling

In order to determine the best combination of models to include in our model ensemble, all encoder based models that require training were trained using different subsets of data. This includes the original (EN) English data, the corresponding augmentations in Chinese (ZH), Spanish (ES) and Para-

Submetric	Model	Language				
		EN	ES	ZH	PA	ALL
VSP	EN	0.195	0.173	0.161	0.067	0.149
	ES	0.156	0.183	0.158	0.012	0.127
	ZH	0.179	0.111	0.102	0.086	0.119
	PA	0.212	0.193	0.198	0.062	0.166
	ML5	0.195	0.168	0.157	0.040	0.140
NSP	EN	0.279	0.256	0.286	0.267	0.272
	ES	0.266	0.257	0.282	0.251	0.264
	ZH	0.246	0.238	0.298	0.232	0.254
	PA	0.307	0.279	0.286	0.279	0.288
	ML75	0.300	0.284	0.311	0.272	0.292
ENG	EN	0.319	0.275	0.251	0.260	0.276
	ML5	0.310	0.268	0.214	0.275	0.267
	ML10	0.334	0.296	0.243	0.279	0.288
	ML20	0.379	0.324	0.274	0.316	0.324
	ML50	0.340	0.263	0.258	0.289	0.287
	PA	0.265	0.245	0.213	0.265	0.247

Table 1: Spearman Correlation scores of our trained model variants on all Language benchmarks on the full development set. The best score for each submetric and language is highlighted in **bold**. Models included in the final ensemble are in **bold**, except for NSP, which also includes NSP-Siamese.

phrases (PA) and the QE-ranked multilingual augmentation (MLXX)⁵.

Spearman correlation results are presented in Table 1. For the **VSP** submetric, we note that the inclusion of translations is detrimental to performance. In fact, the best performing models are PA, followed by EN. This contrasts with **NSP**, where we observe that the inclusion of more translated data improves performance. For **ENG**, the best performance is obtained with 20% of translated data. We include the 10 and 50% models in our framework to take advantage of ensembling.

5.4 Track Results

For the track we submitted 4 different systems, exploring the contribution of the different components of our framework:

- **System 1 (DIALEVALML)**: Submetric ensembling of ChatGPT + XLM-R.
- **System 2**: Submetric ensembling of XLM-R.
- **System 3**: Submetric ensembling of ChatGPT.
- **System 4**: Direct mapping of ChatGPT submetrics.

Table 2 identifies the turn-level weights calculated for testing for System 1.

⁵We only include the best performing ML models.

Aspect	VSP	NSP	MLM	ENG	cGPT-A	cGPT-R	cGPT-C	cGPT-G
Appropriateness	0.039	0.176	0.017	0.0511	0.165	0.185	0.181	0.185
Relevance	0.014	0.214	0.003	0.023	0.188	0.210	0.160	0.190
Content Richness	0.176	0.085	0.181	0.238	0.039	0.022	0.210	0.048
Grammatical Correctness	0.021	0.084	-0.06	0.061	0.238	0.242	0.155	0.258

Table 2: Calculated submetric weights of System 1 for test set quality aspects. Highest weight per aspect in **bold**.

Team	EN		ZH		ES		ML-AVG		Rank	
	Turn	Dial	Turn	Dial	Turn	Dial	Turn	Dial	Turn	Dial
Baseline (AM-FM)	0.2940	0.2414	0.0753	0.4648	0.1826	0.8080	0.1840	0.5047	4	2
Team 2	0.1469	-	0.1054	-	0.0808	-	0.1110	-	5	-
Team 4 (us)									1	1
- S1 (DIALEVALML)	0.4818	0.5342	0.3936	0.7133	0.5890	0.8080	0.4881	0.6852		
- S2	0.2625	0.3295	0.3096	0.7030	0.5056	0.2500	0.3592	0.4275		
- S3	0.4795	0.5251	0.3656	0.6701	0.5409	0.8080	0.4620	0.6677		
- S4	0.4586	0.5039	0.3618	0.5859	0.5412	0.5915	0.4539	0.5604		
Team 5	0.3702	0.1865	0.0701	0.1356	0.1983	0.6830	0.2129	0.3350	3	3
Team 7	0.2214	-	0.3112	-	0.5644	-	0.3657	-	2	-

Table 3: Average Spearman correlation across the 4 dimensions evaluated for the baseline Deep AM-FM (Zhang et al., 2021b) and all participating teams on the Task 1 (Multilingual metrics) test set. **Bold** denotes the best result for the corresponding Language, *italic* denotes our best submission.

Task 1: Multilingual Metrics The results for each team for Task 1 are presented in Table 3, together with all of our submissions. In all languages at both the dialogue and turn level, our submissions vastly outperform others, with the exception of S2, which has comparable results with other participants. This clearly demonstrates the conversational understanding ChatGPT possesses. As expected, the best submission is S1, which conducts submetric ensembling with the XLM-R submetrics. This is followed by S3 and S4, which are exclusive ChatGPT submissions with and without ensembling, respectively.

Task 2: Robust Metrics The results for each team for Task 2 are presented in Table 4. Similar to Task 1, in Task 2, our ChatGPT submissions outperform other teams. However, at the dialogue level, the best performing model is AM-FM.

5.5 Example predictions

Given the widely publicised emergent capabilities of current LLMs, it is worthwhile exploring where their quality predictions diverge from the annotators. To do so, we checked all instances where ChatGPT (System 4) diverges from the Human Evaluation (HE) annotations by more than 3 points. In all of the detected examples, we noted ChatGPT consistently underestimated the quality of the response when compared to HE.

We present in Table 5 two representative examples. In the first example, we see that ChatGPT

Team	Turn (rank)	Dial (rank)
Baseline (AM-FM)	0.3387 (4)	0.4800 (1)
Team 1	0.1537 (6)	0.1111 (4)
Team 3	0.2697 (5)	0.2196 (3)
Team 4 (us)		
- S1 (DIALEVALML)	0.4890 (1)	<i>0.3031</i> (2)
- S2	0.3320	0.2335
- S3	0.4756	0.2979
- S4	0.4427	0.2492
Team 6	0.4190 (2)	-
Team 7	0.3833 (3)	-

Table 4: Average Spearman correlation and corresponding rank across the 4 dimensions evaluated for the baseline Deep AM-FM and all participating teams on the Task 2 (Robust metrics) test set. **Bold** denotes the best result for the corresponding Language, *italic* denotes our best submission.

erroneously underestimates quality due to the inclusion of "weed" in the response. We posit this is due to the RLHF finetuning, which conditions the model to avoid inappropriate or divisive topics. In the second example, we see ChatGPT has trouble understanding the conversation. Although one could argue the HE scores for *Correctness* and *Appropriateness* are too high, it seems clear the response is undeserving of a minimum score for all aspects. In fact, if one prompts the model to provide an explanation for *Content Richness*, it replies the following: "The response attempts to provide some content related to the topic of adolescent sadness, but it is vague and lacks depth. The mention of "Qibing" without any explanation or context

Chatbot:	Tennis racquets used to be made from cloth strips stuffed with feathers haha! Then eventually evolved into today's technology			
Human:	oh i never knew that			
Chatbot:	Indeed, there was also aluminium construction improvements allowed for the introduction of oversized rackets by 1975 which was manufactured by weed			
Label:	A: 4.7	C: 4.7	G: 4.3	R: 5.0
ChatGPT:	A: 1.0	C: 2.0	G: 2.0	R: 1.0
	<i>"The mention of "weed" as the manufacturer may cause confusion or be perceived as inappropriate without further context."</i>			
Best:	A: 2.2	C: 2.6	G: 1.9	R: 2.2
Human:	Don't cry for breakfast.			
Chatbot:	Cry t for comfort			
Human:	In the end, young people are always inexplicably sad in adolescence. Which one is Qibing?			
Label:	A: -	C: 5.0	G: 5.0	R: 3.0
ChatGPT:	A: 1.0	C: 1.0	G: 1.0	R: 1.0
	<i>"The response does not directly relate to the context or provide a meaningful answer. It seems unrelated and out of place. The mention of "Qibing" without any explanation further adds to the confusion."</i>			
Best:	A: 1.3	C: 1.9	G: 1.1	R: 1.2

Table 5: Example turn-level predictions for *Appropriateness*, *Content Richness*, *Grammatical Correctness* and *Relevance*. We include the ChatGPT explanation for *Appropriateness*.

leaves the reader confused. The response could benefit from more specific and informative details about the topic to increase its content richness.". However, if anything, the inclusion of the last sentence increases the richness of the response. Yet, it seems ChatGPT is conflating *Content Richness* with *Relevance*. We observe the same behaviour in all other instances we studied, and is in line with the submetric weights (Table 2).

6 Discussions

The results from our work on both tasks (Section 5.4) reveals that ChatGPT vastly outperforms typical encoder approaches that are trained to discriminate positive samples from artificially generated negative ones. It is important to note that, compared to the months worth of research dedicated to optimise our encoder models (including curation, training and selection), we were able to easily outperform all other teams and our own encoder models with a day's worth of prompt engineering. This is, in our opinion, a turning point in the paradigm of dialogue evaluation.

In any case, we do find instances where ChatGPT fails to accurately evaluate aspects of quality, as identified in Section 5.5. Future research directions may attempt to tackle the issues of score calibration

by providing prompts that include examples and/or explicitly provide guidelines for scoring.

However, given the current landscape on dialogue generation, and as our submission suggests, dialogue evaluation, it is important to reflect on the value of current quality estimation frameworks. One might argue performing HE or developing metrics that evaluate responses and/or dialogues in terms of linguistic competence (e.g. *Grammatical Correctness* or *Coherence*) is no longer informative for the current and future crop of LLMs. Besides becoming ever so clear that these models no longer output responses that are incoherent or incorrect, we are reaching the point where these models are better evaluators than humans themselves (Gilardi et al., 2023). As such, developing metrics that correlate well with HE is becoming increasingly questionable.

One of the main contention points w.r.t the deployment of these models to the public pertain to their "safety" and "trustworthiness". But while "trustworthiness" can be evaluated by connecting the outputs to external and verifiable sources, the notion of "safety" is much more ambiguous. Kempt et al. (2023) suggests considering Positionality, Acceptability, and Value Alignment (PAVA) as features chatbots should have to fulfil appropriateness requirements. However, automatically evaluating if a chatbot has these features using current dialogue evaluation protocols seems implausible. Instead, the development of challenge sets for validation (such as the ones proposed in Valmeekam et al. 2023) appears to be the logical next step for evaluation of future chatbots⁶.

7 Conclusion

This paper presents a novel open-domain and reference-free dialogue evaluation framework that leverages strong pretrained LLMs using finetuning and zero-shot prompting. These models, combined with effective ensembling strategies, substantially outperform the previous automatic evaluation paradigm of only training LMs with semisupervised training objectives. In fact, DIALEVALML ranks 1st on both the Robust (1st turn-level, 2nd dialogue level) and Multilingual (1st on both levels) tasks of Track 4 at DSTC11.

⁶See [OpenAI Evals](#) for recent collaborative research efforts in this direction.

Acknowledgements

This research was supported by the Portuguese Recovery and Resilience Plan through project C645008882-00000055 (Responsible.AI), and 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 (LISBOA-01-0247-FEDER-045909).

References

- Satanjeev Banerjee and Alon Lavie. 2005. [METEOR: An automatic metric for MT evaluation with improved correlation with human judgments](#). In *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, pages 65–72, Ann Arbor, Michigan. 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.
- Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrike, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, Harsha Nori, Hamid Palangi, Marco Tulio Ribeiro, and Yi Zhang. 2023. [Sparks of Artificial General Intelligence: Early experiments with GPT-4](#).
- 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.
- 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.
- Fabrizio Gilardi, Meysam Alizadeh, and Maël Kubli. 2023. [ChatGPT Outperforms Crowd-Workers for Text-Annotation Tasks](#).
- Nuno M. Guerreiro, Duarte Alves, Jonas Waldendorf, Barry Haddow, Alexandra Birch, Pierre Colombo, and André F. T. Martins. 2023. [Hallucinations in large multilingual translation models](#).
- Biyang Guo, Xin Zhang, Ziyuan Wang, Minqi Jiang, Jinran Nie, Yuxuan Ding, Jianwei Yue, and Yupeng Wu. 2023. [How Close is ChatGPT to Human Experts? Comparison Corpus, Evaluation, and Detection](#).
- David M. Howcroft, Anya Belz, Miruna-Adriana Clinciu, Dimitra Gkatzia, Sadid A. Hasan, Saad Mahamood, Simon Mille, Emiel van Miltenburg, Sashank Santhanam, and Verena Rieser. 2020. [Twenty years of confusion in human evaluation: NLG needs evaluation sheets and standardised definitions](#). In *Proceedings of the 13th International Conference on Natural Language Generation*, pages 169–182, Dublin, Ireland. Association for Computational Linguistics.
- 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.
- Jessica Huynh, Cathy Jiao, Prakhar Gupta, Shikib Mehri, Payal Bajaj, Vishrav Chaudhary, and Maxine Eskenazi. 2023. [Understanding the effectiveness of very large language models on dialog evaluation](#).
- Zhijia Jiang, Guanghui Ye, Dongning Rao, Di Wang, and Xin Miao. 2022. [IM²: an interpretable and multi-category integrated metric framework for automatic dialogue evaluation](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 11091–11103, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Hendrik Kempt, Alon Lavie, and Saskia K. Nagel. 2023. [Appropriateness is all you need!](#)
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- 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.

- 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.
- Yinhan Liu, Jiatao Gu, Naman Goyal, Xian Li, Sergey Edunov, Marjan Ghazvininejad, Mike Lewis, and Luke Zettlemoyer. 2020. [Multilingual denoising pre-training for neural machine translation](#). *Transactions of the Association for Computational Linguistics*, 8:726–742.
- 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 *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1116–1126.
- Shikib Mehri and Maxine Eskenazi. 2020a. [Unsupervised evaluation of interactive dialog with DialoGPT](#). 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.
- John Mendonca, Alon Lavie, and Isabel Trancoso. 2022. [QualityAdapt: an automatic dialogue quality estimation framework](#). In *Proceedings of the 23rd Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 83–90, Edinburgh, UK. Association for Computational Linguistics.
- John Mendonca, Alon Lavie, and Isabel Trancoso. 2023. Towards multilingual automatic open-domain dialogue evaluation. In *Proceedings of the 24th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, Prague, Czechia. Association for Computational Linguistics.
- Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. 2022. [Training language models to follow instructions with human feedback](#).
- 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.
- Vitou Phy, Yang Zhao, and Akiko Aizawa. 2020. [Deconstruct to reconstruct a configurable evaluation metric for open-domain dialogue systems](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 4164–4178, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Ricardo Rei, Craig Stewart, Ana C Farinha, and Alon Lavie. 2020. [Unbabel’s participation in the WMT20 metrics shared task](#). In *Proceedings of the Fifth Conference on Machine Translation*, pages 911–920, Online. Association for Computational Linguistics.
- Ricardo Rei, Marcos Treviso, Nuno M. Guerreiro, Chrysoula Zerva, Ana C Farinha, Christine Maroti, José G. C. de Souza, Taisiya Glushkova, Duarte Alves, Luisa Coheur, Alon Lavie, and André F. T. Martins. 2022. [CometKiwi: IST-unbabel 2022 submission for the quality estimation shared task](#). In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 634–645, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Mario Rodríguez-Cantelar, Chen Zhang, Chengguang Tang, Ke Shi, Sarik Ghazarian, João Sedoc, Luis Fernando D’Haro, and Alexander Rudnicky. 2023. Overview of robust and multilingual automatic evaluation metrics for open-domain dialogue systems at dstc 11 track 4. In *DSTC11: The Eleventh Dialog System Technology Challenge*, 24th Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL), Prague, Czechia.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. [Proximal Policy Optimization Algorithms](#).
- Koustuv Sinha, Prasanna Parthasarathi, Jasmine Wang, Ryan Lowe, William L. Hamilton, and Joelle Pineau. 2020. [Learning an unreferenced metric for online dialogue evaluation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2430–2441, Online. Association for Computational Linguistics.
- Chongyang Tao, Lili Mou, Dongyan Zhao, and Rui Yan. 2018. [Ruber: An unsupervised method for automatic evaluation of open-domain dialog systems](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.
- Karthik Valmeekam, Sarath Sreedharan, Matthew Marquez, Alberto Olmo, and Subbarao Kambhampati. 2023. [On the Planning Abilities of Large Language Models \(A Critical Investigation with a Proposed Benchmark\)](#).
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all](#)

- [you need](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. 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.
- Chrysoula Zerva, Daan van Stigt, Ricardo Rei, Ana C Farinha, Pedro Ramos, José G. C. de Souza, Taisiya Glushkova, Miguel Vera, Fabio Kepler, and André F. T. Martins. 2021. [IST-unbabel 2021 submission for the quality estimation shared task](#). In *Proceedings of the Sixth Conference on Machine Translation*, pages 961–972, Online. Association for Computational Linguistics.
- Chen Zhang, Yiming Chen, Luis Fernando D’Haro, Yan Zhang, Thomas Friedrichs, Grandee Lee, and Haizhou Li. 2021a. [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 (Volume 1: Long Papers)*, pages 5676–5689, Online. Association for Computational Linguistics.
- Chen Zhang, Luis Fernando D’Haro, Rafael E. Banchs, Thomas Friedrichs, and Haizhou Li. 2021b. [Deep AM-FM: Toolkit for Automatic Dialogue Evaluation](#), pages 53–69. Springer Singapore, Singapore.
- Pengfei Zhang, Xiaohui Hu, Kaidong Yu, Jian Wang, Song Han, Cao Liu, and Chunyang Yuan. 2022. [MME-CRS: Multi-Metric Evaluation Based on Correlation Re-Scaling for Evaluating Open-Domain Dialogue](#). *arXiv preprint arXiv:2206.09403*.
- 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, Online. Association for Computational Linguistics.