

ECoh: Turn-level Coherence Evaluation for Multilingual Dialogues

John Mendonça^{1,2}, Isabel Trancoso^{1,2} and Alon Lavie^{3,4}

¹ INESC-ID, Lisbon

² Instituto Superior Técnico, University of Lisbon

³ Carnegie Mellon University, Pittsburgh

⁴ Phrase, Pittsburgh

{john.mendonca, isabel.trancoso}@inesc-id.pt, alavie@cs.cmu.edu

Abstract

Despite being heralded as the new standard for dialogue evaluation, the closed-source nature of OpenAI’s GPT-4 model poses challenges for the research community. Motivated by the need for lightweight, open source, and multilingual automated dialogue evaluators, this paper introduces GENRESCO (Generated Responses targeting Coherence). GENRESCO is a novel LLM-generated dataset comprising over 130k negative and positive responses and accompanying explanations seeded from XDailyDialog and XPersona covering English, French, German, Italian, and Chinese. Leveraging GENRESCO, we propose ECOH¹ (Evaluation of Coherence), a family of evaluators trained to assess response coherence across multiple languages. Experimental results demonstrate that ECOH achieves multilingual coherence detection capabilities superior to the teacher model (GPT-3.5-Turbo) on GENRESCO, despite being based on a much smaller architecture. Furthermore, the explanations provided by ECOH closely align in terms of quality with those generated by the teacher model.

1 Introduction

With LLMs showcasing impressive reasoning and dialogue understanding capabilities vastly superior to any prior NLP technologies, human evaluation has more recently been complemented with automatic evaluations using GPT-4 (OpenAI, 2024). However, GPT-4 as an automated evaluator has its downsides. Perhaps the main downside is it being a closed source model hidden behind a paid API, making accessibility difficult for those outside the coverage area and lacking extensive financial resources, while also lacking transparency in its development. In contrast, and to the best of our knowledge, the study of open source alternatives to GPT-4 based dialogue evaluation is mostly limited

¹Pronounced "Echo".

Context:

A: Hello. I bought a China dress in your shop this morning.

B: Yes?

A: I bought it one size up by mistake.

B: Oh, did you?

A: I wonder if you can change it to one size down.

Reference Response. B: Yes, of course. Will you come with the receipt ?

"The response acknowledges the request and offers a solution to accommodate the customer’s needs. The answer is Yes."

Random Negative Sample. B: I’d like a book about law.

"The response is completely unrelated to the situation discussed. The answer is No."

GENRESCO Positive Sample. B: Absolutely, if you bring the dress back to the shop with the receipt, we can exchange it for a smaller size for you.

"The response acknowledges the request and offers a solution to the problem. The answer is Yes."

GENRESCO Negative Sample. B: Oh, that’s great to hear! I hope it fits perfectly.

"The response does not acknowledge the request for a size change and instead expresses an unrelated sentiment. The answer is No."

Table 1: Example of automatically generated negative samples obtained with random response selection, and obtained from GENRESCO (§2), our proposed dataset. The explanations are generated using one of our proposed models, ECOH-4B-ML (§3).

to the benchmarking of open source and open access LLMs or finetuning with dialogue data (Huynh et al., 2023; Zhang et al., 2023, 2024). These works suggest that LLMs struggle to outperform older encoder-based metrics trained using negative sampling approaches for relevance (e.g. random response selection). However, it is important to point out that these benchmarks have several limitations.

First and foremost, the high performance of these encoder-based models can be explained by the fact that the benchmarks themselves are based on old generative models that exhibit relevance issues that are easy to detect. For instance, in Table 1, metrics trained using random negative sampling strategies for relevance will output a positive score to all

responses except the random negative one. As such, these metrics struggle to evaluate contemporary chatbots, since these typically output fluent and semantically relevant responses.

Furthermore, only a select few benchmarks are multilingual. Whilst there is work that attempts to evaluate the multilingual capabilities of dialogue evaluation metrics (Mendonca et al., 2023; Zhang et al., 2023), they use translated benchmarks. This assumes that critical errors typically produced by these older models (e.g. irrelevance), are not influenced by language. However, more complex quality aspects such as coherence may have nuances that make them unique to certain cultures. Depending on the context, some culture specific details may or may not be implicitly inferred (Hall, 1959).

These key observations motivate our work. In order to move towards the development of metrics that evaluate dialogue coherence and are multilingual, we propose GENRES COH (**Generated Responses targeting Coherence**), a collection of positive and negative responses focused on coherence. Our dataset, generated using strong LLMs, contains over 130k responses in different languages (English, French, German, Italian, and Chinese), together with their corresponding explanations (in English). By prompting an LLM, we are able to (1) obtain positive samples that are in distribution (LLMs frequently output more verbose responses than their human counterparts); (2) obtain negative samples that remain semantically relevant but contain coherence and logical consistency issues, which may be more informative during training, and that are more representative of current limitations of LLMs.

With this dataset, we train a family of evaluators we call ECOH (**E**valuation of **C**oherence)². Our results demonstrate that distilling Coherence knowledge from a strong LLM allows us to obtain multilingual coherence detection performance of **.945** F1 score using a 0.5B model, which is superior to both the teacher models’ (GPT-3.5-Turbo) **.910** and a much larger model of the same family (QWEN1.5-7B-CHAT - **0.825**). Furthermore, the explanations provided by ECOH are of higher quality than QWEN1.5-7B-CHAT, scoring an average of over **4** out of **5** on most instances, as reported by GPT-4 evaluations.

²github.com/johndmendonca/Ecoh

2 GENRES COH responses dataset

This section introduces GENRES COH, a multilingual, large-scale response collection that targets coherence, seeded from well established dialogue datasets (§2.1), and generated using LLMs (§2.2). Table 1 presents an example from this dataset. For additional examples of this dataset in other languages, see Appendix A.2.

2.1 Dataset Sources

Our work leverages two distinct dataset sources: XDailyDialog (Liu et al., 2023b) and XPersona (Lin et al., 2021). For training, development and testing, we use XDailyDialog, a multilingual extension of DailyDialog with human translations covering German-DE, Italian-IT and Chinese-ZH. XDailyDialog includes 13K parallel dialogues, amounting to 52K dialogues and 410K utterances. During our pre-processing step we noted a substantial overlap of dialogues between the provided test and training/validation sets of XDailyDialog. As a result, we excluded these dialogues (amounting to 20%) from the test set.

In order to gauge the extensibility to other dialogue datasets and languages, we additionally include XPersona data in our GENRES COH test set. XPersona is a multilingual extension of the PersonaChat dataset (Zhang et al., 2018) with human revised machine translations for six languages. Besides English, we include Italian-IT, Chinese-ZH, and an additional unseen language, French-FR, in our experiments. For each language, we extract 1K contexts from the test set for response generation.

For contrastive comparison, we also use DailyDialog++ (Sai et al., 2020), a similar curation effort which uses the original DailyDialog dataset, and where annotators were asked to create five additional relevant responses and five adversarial irrelevant responses for each context.

2.2 Generation

Development set We leverage GPT-3.5-turbo³ (Ouyang et al., 2022) as the strong LLM to generate, given prior dialogue context, a positive and a negative response, paired with a brief explanation of the issue (or lack thereof). Each response pair is generated given a context of at least 2 turns up until the length of the dialogue except the last turn (this ensures the response is generated from a still

³`gpt-3.5-turbo-0125` and `gpt-4-1106-preview` accessed via OpenAI’s API in early April.

Dataset	Size (# contexts)	Response Avg. length	Explanation Avg. length	Response MTLD
DailyDialog++ (2020)				
Random	9,259/1,028/1,142	9.40	-	169.94
Adversarial		10.70	-	186.42
GENRESCO-H-DEV				
DailyDialog-LATIN	51,873/5,080	14.74	15.03	105.03
DailyDialog-ZH		23.06	14.54	54.38
GENRESCO-H-TEST				
DailyDialog-LATIN	4,770	14.82	26.27	155.28
DailyDialog-ZH		24.79	26.04	69.03
PersonaChat-LATIN	1,000	15.37	27.89	204.61
PersonaChat-ZH		28.81	27.78	76.66

Table 2: Comparison of statistics for different negative sample datasets. DD denotes XDailyDialog, PC XPersona. Dataset size denotes the number of unique contexts from which responses were obtained for training/validation/test subsets. MTLD denotes the Measure of Textual Lexical Diversity (McCarthy, 2005) of the responses. We report statistics for Latin script languages (denoted LATIN and covering EN,DE,FR,IT), separated from Chinese-ZH. For Average length, LATIN is calculated using words, whereas ZH uses characters.

ongoing conversation). We set the temperature to 0.7, the top- p to 1, and the maximum number of tokens to 300, thereby enforcing smaller explanations which in turn should reduce inference costs. Despite sharing the same contexts, the responses and corresponding explanations are not necessarily translations of the English subset. This allows the model to freely generate responses that are more likely to occur (for the positive samples) or more representative of coherence issues in that particular language, instead of being a translation from English. The prompt used for this generation is included in Appendix A.1.

Test set For testing, we employ GPT-4 (OpenAI, 2024) to ensure higher quality outputs and reduce in-distribution biases from the training set. GPT-4 has been shown to match human annotations on quality, from general NLP tasks to highly specialised fields (West et al., 2022; Raunak et al., 2023; Savelka et al., 2023).

Human validation In order to verify the outputs of GPT-4, we additionally conduct a human validation step involving one expert linguist from each language. We randomly sample 100 examples from the XDailyDialog test set, and report an appropriateness rate that exceeds 97%, thus validating the response and explanation generation process using GPT-4. Details regarding human validation are provided in Appendix A.3.

2.3 Statistical Analysis

We present relevant statistics for our dataset, together with DailyDialog++ in Table 2. Since the test set for our dataset is generated by GPT-4, we opt to present the statistics separately.

Firstly, despite GENRESCO-H boasting a much larger context set, amounting to 51k/5k for training/validation, each context only has a single positive and negative response, whereas DailyDialog++ contains 5 positive responses and an additional 5 adversarial negative responses.

When comparing the average length of responses, we note that GENRESCO-H responses are longer than the human curated responses of DailyDialog++. This verbosity is a known behaviour by LLMs, since they are conditioned to output longer responses due to the Reinforcement learning from human feedback (RLHF) step, at least when compared to humans (Kamalloo et al., 2023). Additionally, we note that the response lengths remain similar across the development and test sets, whereas the explanations are much longer in the test set.

For a more fine grained analysis of the responses, we measure their lexical diversity using the Measure of Textual Lexical Diversity (MTLD) (McCarthy, 2005).⁴ Since DailyDialog++ contains 5 responses per context, we calculate the average diversity when considering the responses individually. We observe that the diversity of human responses for DailyDialog++ is larger than the ones generated by GPT-3.5-Turbo for the development set,

⁴Calculated using lexical-diversity Python package.

but similar to the ones generated by GPT-4 for the test set. This disparity is to be expected, given the performance differences between the two models in creative writing tasks.⁵

It is important to note that the adversarial responses from DailyDialog++ exhibit greater diversity compared to those from GENRESCO. This is because the tasks are slightly different: in DailyDialog++, annotators were asked to generate new irrelevant responses by incorporating certain words from the context directly or indirectly into their responses. This stands in contrast to our approach, which prioritises coherence while preserving relevance. As such, the introduction of diverse words into the response is constrained by the fact relevance must be upheld.

3 ECOH

This section presents ECOH, our proposed family of response coherence evaluators. We initially present the method of formulating the task of coherence evaluation as explainable QA (Question Answering) (§3.1). Then, we describe in detail how our evaluator is trained (§3.2) and evaluate its performance on different settings (§3.3).

3.1 Problem Formulation

Turn-level dialogue coherence evaluation consists of the assessment of a response hypothesis h given a dialogue history (frequently denoted as context) c of varying amount of turns, and optionally one or more references r and/or external knowledge k . The goal is to learn a scoring function that assigns a score $f(c, k, r, h) \rightarrow s$ for each individual quality aspect. This scoring function is compared against human judgements, which annotate the same context-response pairs. These responses are evaluated by humans using, for instance, a binary (0, 1) judgement or a [1, 5] Likert scale, where the lowest value means lowest quality and highest value maximum quality.

In our work, we consider Coherence as being a binary quality aspect. Despite being frequently annotated in the literature on a Likert Scale, what can be considered a response that is neither coherent or incoherent is mostly left to the interpretation of the annotator. Given that we are leveraging an LLM for generation, we find it unfeasible to generate a balanced dataset that contains intermediate levels

⁵It is important to point out, that a higher temperature value would likely result in higher diversity, with a possible trade off in performance.

of coherence. Instead, we generate a positive and a negative response in terms of coherence and label it accordingly. This contrastive sampling strategy for coherent responses is also followed in most metric development work for Relevance or Sensibleness, where models are typically trained using self-supervised learning strategies that sample negative responses by random selection (Mehri and Eskenazi, 2020b; Yeh et al., 2021; Mendonça et al., 2023). Lacking any external knowledge with respect to each dialogue, we then further simplify the reference-free evaluation of coherence as a Question Answering (QA) task ($f(c, h) \rightarrow s \in (0, 1)$), with model responses being either coherent ("Yes") or incoherent ("No").

3.2 Experimental Setup

Model Specification We employ the QWEN1.5-CHAT family of LLMs (Bai et al., 2023) for our models. QWEN1.5 contains LLMs of various sizes, ranging from 0.5B up to 72B and support all the languages of XDailyDialog. We limit our finetuning experiments up to 4B due to the tradeoff between performances and compute. We feed the dialogue context to the model and ask it to provide a "Yes"/"No" answer to the question "*Given the context, is the response Coherent?*". The model is trained to also output a succinct explanation to the answer. We opted with asking for the explanation first, before answering the question, in order to leverage the autoregressive nature of the model. In theory, this should guarantee that final answer be informed by the explanation.⁶ Additional training details are available in B.

Baselines We contrastively compare our proposed approach against several models. We begin by including models trained using random negative responses from DailyDialog: a ROBERTA-LARGE model (Liu et al., 2019) (which we train ourselves – see Appendix B for details); and UNIEVAL (Zhong et al., 2022) (which uses T5 as base model). Since these models output a probability score, we assume the model outputs the positive class when the $p > 0.5$. Additionally, we conduct zero and one shot (with English and language specific examples) inference using QWEN1.5-CHAT to determine if finetuning on GENRESCO adds improvements to the performance of the base model. We also com-

⁶Chiang and Lee (2023) has shown that dialogue evaluation performance is not *always* better when requesting the explanation first. We leave this analysis for future work.

pare against GPT-3.5-Turbo (Ouyang et al., 2022), the teacher model which was used to generate the development set of GENRESCOH, and which is weaker than our expert (GPT-4).

3.3 Main Results

Since the coherence labels are binary, we report detection results using F1-score and Point Biserial Correlation. Additionally, we compute the BLEU-4 score of the generated short explanation using the GPT-4 explanation as a reference. Since BLEU compares overlap in tokens instead of comparing meaning, we also employ GPT-4 as a drop-in replacement for human annotators, and ask it to assess the explanations of 200 random responses from the models that output an explanation.

Model	ρ_{pb}	F1	BLEU	GPT-4
1 (always positive)	NaN	.333	-	-
NSP-ROBERTA	.1651	.430	-	-
UNIEVAL	<u>.3272</u>	<u>.500</u>	-	-
QWEN1.5-CHAT				
0.5B	.2226	.600	3.80	1.84 \pm 1.12
1.8B	.5212	.740	2.58	2.39 \pm 1.29
4B	.5850	.783	<u>8.16</u>	3.18 \pm 1.60
7B	<u>.7918</u>	<u>.890</u>	4.63	<u>3.95</u> \pm 1.48
GPT-3.5-Turbo	.8256	.910	5.25	4.55 \pm 1.08
ECOH-EN				
0.5B	.7756	.878	16.02	3.80 \pm 1.43
1.8B	.8242	.908	17.30	4.13 \pm 1.29
4B	.9185	.960	17.92	<u>4.45</u> \pm 0.96
ECOH-ML				
0.5B	.8882	.945	17.00	3.99 \pm 1.36
1.8B	.9019	.953	17.28	4.24 \pm 1.28
4B	<u>.9491</u>	<u>.975</u>	<u>18.05</u>	4.29 \pm 1.12

Table 3: Reported results on GENRESCOH-DD-TEST, averaged across all languages. ρ_{pb} denotes Point Biserial Correlation. ECOH-EN and ECOH-ML denote the finetuned models using English data and all multilingual data, respectively. All correlation results are $p < 0.05$. **Bold** denotes best overall model, underline best model of the group.

We collate our main results in Table 3. Due to space limitations, we only report 1-shot performance with a language specific example for QWEN1.5-CHAT and the results correspond to the average of the languages. Additional results, including Zero shot and individual language performance, are available in Appendix C.

GPT-3.5 performance with 4B parameters Our main observation is that, although being one of our smallest models, ECOH-0.5B-ML outperforms the

predictive performance of the teacher model (reported in F1), and the explanations of QWEN1.5-CHAT-7B. Furthermore, ECOH-4B-EN has similar explanation quality to that of GPT-3.5-Turbo. As expected, training models using random response selection (NSP-ROBERTA-L and UNIEVAL) is not sufficient for accurately detecting more advanced coherence issues. In fact, these models’ performance sit between QWEN1.5-0.5B-CHAT (.600 F1) and the naive single output model (.333 F1).

Model size and Multilingual finetuning Since our smallest model already achieves strong results (.945 F1 score), increasing the model size results in only a small relative improvement of 3% in performance. However, we do observe larger performance improvements with multilingual finetuning. For instance, for ECOH-0.5B, we observe an improvement of over 7% (.878 to .945). This indicates, as expected, that including multilingual data during finetuning improves results for the various covered languages.

Explanations We also note that our finetuned models have much higher BLEU and GPT-4 scores than the base models. The obtained BLEU scores are to be expected, given that ECOH is finetuned with explanation data stemming from the same prompt, which is a biased observation from the response generator. This is supported by the teacher model’s performance, achieving the highest GPT-4 assessment, despite having low BLEU. In any case, by validating the responses of the ECOH models with GPT-4, we see that the explanations are on average of higher quality than the ones generated by even the largest base model (QWEN1.5-CHAT) that we studied.

3.4 Generalization to unseen dialogue datasets and languages

In order to evaluate our models’ capabilities on unseen dialogue datasets, we evaluate our models on XPersona, which was not seen during finetuning. We only select the best baselines (as reported in Table 3) for this analysis. Additionally, our XPersona subset contains French, which is not present in XDailyDialog, so in addition to the average performance across all languages, we present the results for French separately. For fair comparison, we utilise the English example when evaluating the performance of QWEN1.5-CHAT in French.

Model	ρ_{pb}	F1	BLEU	GPT-4
QWEN1.5-CHAT-7B				
FR	.4608	.660	2.97	3.20 \pm 1.58
ML	.6125	.778	3.43	3.75 \pm 1.51
GPT-3.5-TURBO				
FR	.7205	.860	5.04	4.32 \pm 1.25
ML	.7631	.880	4.94	4.45 \pm 1.05
ECOH-ML				
0.5B				
FR	.8089	.910	13.71	3.68 \pm 1.46
ML	.8882	.945	17.00	3.82 \pm 1.38
1.8B				
FR	.7863	.890	14.10	4.40 \pm 1.15
ML	.8472	.920	15.70	4.26 \pm 1.17
4B				
FR	.9270	.960	14.58	4.36 \pm 0.95
ML	.9448	.970	16.33	4.38 \pm 0.96

Table 4: Reported results for GENRESCO-H-PC-TEST (French-FR subset and full-ML set). ρ_{pb} denotes Point Biserial Correlation. All correlation results are $p < 0.05$. **Bold** denotes best overall model, underline best model of the group.

Looking at the results in Table 4, we find that the conclusions from DailyDialog also carry over to XPersona. For the unseen language (French-FR), we note a large drop in performance for QWEN1.5-CHAT-7B, when compared to the other languages, which could be explained by the 1-shot example being in English. For our proposed models, we see a larger gap in performance between French and the other languages for the smaller models, whereas for ECOH-4B, the performance for French is well within the range of that observed for other languages. This is also what we observe for GPT-3.5-TURBO. This finding suggests that, given an LLM that natively supports languages for which we have no finetuning data, coherence knowledge can be drawn from languages that were included for finetuning.⁷

3.5 Generalization to external annotations

Since the models were trained and evaluated on synthetic data, it is important to check if ECOH performs adequately on external evaluations conducted by human annotators. As such, we also assess ECOH on the FED-turn annotations (Mehri and Eskenazi, 2020a) for "Relevance" and "Over-

⁷It is important to acknowledge that this finding is only likely to extend to languages that follow western normative rules for coherence. An additional interesting experiment would be to test ECOH on a language that does not conform to these rules – however, these are typically low-resource.

all", which is a typically used benchmark for dialogue evaluation. Similar to other works, we calculate the average human annotation ($[0, 2]$ for Relevance and $[0, 4]$ for Overall) and report results using Spearman correlation between the human annotation and the score provided by each evaluator. For the LLMs, we keep the binary formulation for coherence (score is either 0 or 1). For the coherence explanation evaluation, lacking a reference, we again use GPT-4 as an explanation evaluator but without a reference response, and evaluate all responses. In order to gauge evaluation performance, we also calculate correlations with GPT-4 as a response evaluator. We refrain from providing GPT-4 explanation scores due to potential self-evaluation bias.

Model	Relevance r	Overall r	GPT-4
NSP-ROBERTA	.2530	<u>.2543</u>	-
UNIEVAL	<u>.2532</u>	.2521	-
QWEN1.5-CHAT			
0.5B	<i>.0451</i>	<i>.0340</i>	2.35 \pm 1.42
1.8B	<i>.2693</i>	<i>.2228</i>	2.91 \pm 1.52
4B	<i>.1613</i>	<i>.1189</i>	3.30 \pm 1.67
7B	<u>.3279</u>	<u>.2998</u>	3.74 \pm 1.54
GPT-3.5-Turbo	.4025	.3636	3.54 \pm 1.66
GPT-4	.5108	.5320	-
ECOH			
0.5B			
EN	.2247	.1548	3.17 \pm 1.77
ML	.1670	.1294	3.17 \pm 1.77
1.8B			
EN	<u>.2941</u>	.2408	3.38 \pm 1.77
ML	.2581	.1801	<u>3.50</u> \pm 1.68
4B			
EN	.2445	.2326	3.17 \pm 1.82
ML	.2685	<u>.2515</u>	3.37 \pm 1.81

Table 5: Reported results for FED-Turn. Performance is calculated using Pearson Correlation (r). All results are $p < 0.05$ unless *italicised*. **Bold** denotes best overall model, underline best model of the group.

From Table 5, we draw several conclusions. Firstly, when looking at the correlation metric, we see that the performance gap between random response-trained models and ECOH is much smaller. This is mainly due to the older chatbots models used for FED-turn – **Meena** (Adiwardana et al., 2020) and **Mitsuku**⁸ – being more likely to output irrelevant and non-specific responses that mimic random response selection.

Secondly, we note that our finetuning is still useful for detecting coherence issues on FED, since,

⁸Mitsuku blogpost




Figure 1: Confusion matrices for the best models of each family (ECOH-1.8B-EN, GPT-3.5-Turbo and QWEN1.5-7B-CHAT) on FED-turn.

overall, ECOH outperforms the corresponding parent model (e.g., ECOH-4B vs QWEN1.5-4B) on Relevance. However, our multilingual models underperform against the English-specific finetuning (with the exception of the 4B model) which could be explained by FED being exclusively in English.

Finally, despite GPT-4 not being a perfect evaluator (low correlation in FED), we assume the scores we obtain for FED remain comparable to previous experiments. With that in mind, we find that the explanation quality is overall lower for ECOH and the teacher model, GPT-3.5-Turbo when compared to GENRESCO (for instance, we report an absolute drop of 1.01 for GPT-3.5-Turbo). In contrast, the QWEN1.5-CHAT models’ explanation quality remains almost unchanged when compared to GENRESCO. As a result, ECOH models achieve less impressive results when compared to their parent models. When comparing models of the same size up until 4B, we do see some improvement in explanation quality. However, our 4B model fails to outperform the 7B model, both in terms of correlation and explanation quality. For the correlation, we believe this is due to the parent model (QWEN1.5-CHAT-4B) having low predictive performance – if we look at the 1.8B models, they yield better correlations. For the explanation quality, we note that the teacher model, GPT-3.5, has also lower results than expected.

However, it is important to acknowledge that these models are tailored towards the binary classification of coherence. As such, we also present results for FED-turn with a binary mapping. In this case, we consider a response to be relevant when the majority of the annotators rate the response as fluent. We present the confusion matrices for the best models of each family in Figure 1. Here, we note that ECOH-1.8B-EN is more likely to say a

response is coherent, incurring slightly more False Positives than the other models (17 against 5), but also lower False Negatives (83 against >129).

3.6 Error analysis

In order to determine limitations and weaknesses with using ECOH, we conducted a thorough analysis of all errors in GENRESCO-PC and FED. We summarise significant findings below.

Context:

A: Hi!

B: hi

A: what are you up to?

B: code refactoring. you?

A: me? just chilling out at work. what is code refactoring?

Response. B: good question. I don’t even know what I am doing

Annotations: **Relevant:** 1.6 **Overall:** 3.0

Evaluation: "The response contradicts B’s previous statement about code refactoring. The answer is No."

Table 6: Example prediction by ECOH-1.8B-ML on FED-Turn. Here, both speakers are human.

Humanlike Responses One of the issues we found when analysing incorrect predictions is that the model sometimes rates human responses as incoherent when they are not "helpful" or do not directly respond to the context, which is not the expected behaviour with chatbot responses. For instance, Table 6 presents a cherry-picked example where ECOH incorrectly rates a human response from FED-Turn as incoherent. We believe such issues could be mitigated by including some portion of coherent human responses from DailyDialog++ during training. However, we also found similar issues with GPT-3.5-Turbo and QWEN, indicating this type of error is transversal to many current LLM-based evaluators.

Context:

A: Che grande giornata! *What a great day!*

B: Sì. Che ne dici di uscire e Dioerti il sole sull'erba? *Yes. How about we go out and enjoy the sun on the grass?*

A: Grande, andiamo! *Great, let's go!*

B: Ehi, tesoro, credo di aver avuto un colpo di calore per essere stato tutto il giorno al sole. Sono così rilassato. Mi sembra di essere in un altro mondo. *Hey, honey, I think I had heatstroke from being in the sun all day. I'm so relaxed. I feel like I'm in another world.*

Response. A: Oh no, forse dovremmo trovare un posto all'ombra o rientrare in casa. Come ti senti adesso? *Oh no, maybe we should find a shady spot or go indoors. How do you feel now?*

Reference: *"The response is considerate and logical, offering a solution to avoid further sun exposure due to the heatstroke B mentioned, and it also shows concern for B's current well-being. The answer is Yes."*

Evaluation: *"The response contradicts B's statement about feeling relaxed and suggests going back indoors. The answer is No."*

Table 7: Example prediction by ECOH-1.8B-ML on GENRESCOH-DD-TEST. Translation provided in *italic*.

Commonsense Reasoning Another issue we found recurrent, especially with the smaller models, is the limited nature of their commonsense reasoning. We hypothesise that this is a capability that smaller models struggle with, and this is reflected in their evaluation capabilities. We see an example of this is Table 7, where the model fails to understand that sun exposure is mitigated by going back indoors.

4 Related Work

4.1 Metrics for Dialogue Evaluation

Statistic-based metrics such as BLEU (Papineni et al., 2002), ROUGE (Lin, 2004), and METEOR (Banerjee and Lavie, 2005), are a popular choice for dialogue evaluation because they are simple to calculate and lightweight. However, since they assume valid responses have significant word-overlap with the ground truth, their correlations with human judgements annotations are very low (Liu et al., 2016) due to the one-to-many nature of dialogues. Additionally, they cannot be used to evaluate models whenever a gold-response is not available.

Consequently, learned metrics were proposed. The typical approach was to finetune pretrained encoder models using positive and negative samples targeting different quality aspects such as fluency and relevance (Mehri and Eskenazi, 2020b; Phy et al., 2020; Sai et al., 2020; Mendonca et al., 2022; Zhao et al., 2020). Other approaches used graph

representations to model dialogue interactions explicitly (Huang et al., 2020; Zhang et al., 2021).

With the introduction of LLMs in a wide range of NLP tasks, most recent approaches leverage them for dialogue evaluation. G-EVAL (Liu et al., 2023a) uses GPT-3.5-Turbo and GPT-4 for the evaluation of generation models using a "Chain-Of-Thoughts" step and a scoring function based on return token probabilities. LLM-EVAL (Lin and Chen, 2023) is a single-prompt-based evaluation method that leverages a unified evaluation schema to cover multiple dimensions of conversation quality in a forward pass. DIALEVALML (Mendonça et al., 2023) combines encoder-based models and direct prompting and score extraction from GPT-3.5-Turbo. XDIAL-EVAL (Zhang et al., 2023) probes the evaluation capabilities of several open source LLMs against GPT-3.5-Turbo (Ouyang et al., 2022), and also finetunes them with dialogue data. To the best of our knowledge, this is the first work that conducts supervised learning of LLMs for the task of dialogue evaluation.

4.2 Dataset Generation

There are several studies that propose augmentation and synthetic generation approaches to scale dataset sizes that target commonsense reasoning (Bhagavatula et al., 2023; Wang et al., 2023), summarisation (Jung et al., 2024), and dialogues (Chen et al., 2023; Kim et al., 2023) for training purposes.

For dialogue evaluation in particular, most metrics are finetuned using self-supervised data (Mehri and Eskenazi, 2020b; Phy et al., 2020; Yeh et al., 2021; Mendonca et al., 2023). The most widely used approach is to select positive samples consisting of the ground truth response, and negative responses from randomly drawn dialogues. Ghazarian et al. (2022) relies on Abstract Meaning Representation (AMR) to apply semantic-level manipulations to existing responses. Our work, in comparison, leverages a strong LLM to generate new incoherent responses at scale.

5 Conclusions

This paper presents GENRESCOH, a large scale collection of positive and negative responses and corresponding explanations covering several languages. GENRESCOH is generated from XDaily-Dialog and XPersona using state-of-the-art LLMs, which better matches the responses seen by contemporary chatbots. With this dataset, we train a family

of evaluators we call ECOH. Our smallest model (0.5B) is able to achieve similar performance to that of the teacher model (GPT-3.5-Turbo), despite being much smaller.

Despite this good performance, we note some limitations when using ECOH, especially when evaluating human responses and/or responses that require more robust commonsense reasoning. Whilst we argue that including more data that targets commonsense and human responses, or even training a larger model could mitigate these issues, since we include an explanation in the predictions, one could still use our evaluators for an initial evaluation screening and escalate to a human evaluator if necessary.

6 Limitations

Reduced Language Selection Our work is only evaluated in English, German, Italian, French and Chinese. This limitation stems in part from the upstream dialogue dataset (XDailyDialog) only covering 4 high resource languages. Whilst XPersona does contain additional languages, we were limited to only including French as unseen language due to annotator and resource limitations.

Generation Generating synthetic data from LLMs might surface or even amplify harmful content within these models. In particular, the choice of a single LLM to generate the responses may induce distribution biases. We identify in Section 2 the reduced lexical diversity of generated responses from GPT-3.5-Turbo when compared to humans. Furthermore, our limited analysis shows that our model sometimes struggles with rating human responses. As such, the generated negative samples may also not accurately represent all coherence issues LLM-based generators typically exhibit. Future investigation may look into producing a systematic quality analysis of a more diverse pool of LLMs, which could inform more faithful generation of negative responses.

FED as a turn level coherence benchmark For most dialogue evaluation benchmarks, coherence annotations are conducted at the dialogue level and do not pinpoint the exact response that triggers incoherence (Yeh et al., 2021). As such, we opted with benchmarking ECOH on FED-turn relevance annotations, which is a typically used benchmark for dialogue evaluation. Despite relevance and coherence being different quality aspects, we note

that a) all irrelevant responses lack, by definition, coherence; b) we found that the vast majority of relevant responses on FED are also coherent. Nevertheless, we acknowledge the limitations of using FED-turn as a turn level coherence evaluation benchmark, namely due to its lack of relevant but incoherent responses.

7 Ethical Considerations

Culture-specific conversational norms We acknowledge that the definition of dialogue quality is a diverse, culturally informed concept. We attempt to reduce the English-centric bias in the generation by leaving the LLM to generate without English reference constraints. However it is possible the generation still conforms to English definitions of coherence given its pretraining and instruction tuning data is more than likely over represented by English text. Furthermore, the examples provided in the prompt, and the dialogues themselves, despite being validated by expert linguists, are still based on English dialogues. As such, users of our model should take extra care when evaluating responses in languages that are known to deviate substantially from English-centric notions of coherence.

Annotations The post-editing of the prompts and the manual validation of GPT-4 generations was partially conducted by volunteer annotators, and paid workers that have a fair wage according to their location.

Acknowledgments

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 DOI: 10.54499/UIDB/50021/2020.

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](#). *CoRR*, abs/2001.09977.
- Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, Binyuan Hui, Luo Ji, Mei Li, Junyang Lin, Runji Lin, Dayiheng Liu, Gao Liu, Chengqiang Lu, Keming Lu, Jianxin Ma, Rui Men, Xingzhang Ren,

- Xuancheng Ren, Chuanqi Tan, Sinan Tan, Jianhong Tu, Peng Wang, Shijie Wang, Wei Wang, Shengguang Wu, Benfeng Xu, Jin Xu, An Yang, Hao Yang, Jian Yang, Shusheng Yang, Yang Yao, Bowen Yu, Hongyi Yuan, Zheng Yuan, Jianwei Zhang, Xingxuan Zhang, Yichang Zhang, Zhenru Zhang, Chang Zhou, Jingren Zhou, Xiaohuan Zhou, and Tianhang Zhu. 2023. Qwen technical report. *arXiv preprint arXiv:2309.16609*.
- 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.
- Chandra Bhagavatula, Jena D. Hwang, Doug Downey, Ronan Le Bras, Ximing Lu, Lianhui Qin, Keisuke Sakaguchi, Swabha Swayamdipta, Peter West, and Yejin Choi. 2023. **I2D2: Inductive knowledge distillation with NeuroLogic and self-imitation**. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9614–9630, Toronto, Canada. Association for Computational Linguistics.
- Maximillian Chen, Alexandros Papangelis, Chenyang Tao, Seokhwan Kim, Andy Rosenbaum, Yang Liu, Zhou Yu, and Dilek Hakkani-Tur. 2023. **PLACES: Prompting language models for social conversation synthesis**. In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 844–868, Dubrovnik, Croatia. Association for Computational Linguistics.
- Cheng-Han Chiang and Hung-yi Lee. 2023. **A closer look into using large language models for automatic evaluation**. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 8928–8942, Singapore. Association for Computational Linguistics.
- Sarik Ghazarian, Nuan Wen, Aram Galstyan, and Nanyun Peng. 2022. **DEAM: Dialogue coherence evaluation using AMR-based semantic manipulations**. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 771–785, Dublin, Ireland. Association for Computational Linguistics.
- Edward T. Hall. 1959. *The silent language*. Doubleday, Garden City, N. Y.
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. **Lora: Low-rank adaptation of large language models**.
- 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**.
- Jaehun Jung, Peter West, Liwei Jiang, Faeze Brahman, Ximing Lu, Jillian Fisher, Taylor Sorensen, and Yejin Choi. 2024. **Impossible distillation: from low-quality model to high-quality dataset & model for summarization and paraphrasing**.
- Ehsan Kamalloo, Nouha Dziri, Charles Clarke, and Davood Rafiei. 2023. **Evaluating open-domain question answering in the era of large language models**. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5591–5606, Toronto, Canada. Association for Computational Linguistics.
- Hyunwoo Kim, Jack Hessel, Liwei Jiang, Peter West, Ximing Lu, Youngjae Yu, Pei Zhou, Ronan Bras, Malihe Alikhani, Gunhee Kim, Maarten Sap, and Yejin Choi. 2023. **SODA: Million-scale dialogue distillation with social commonsense contextualization**. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 12930–12949, Singapore. Association for Computational Linguistics.
- 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*.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81.
- Yen-Ting Lin and Yun-Nung Chen. 2023. **LLM-eval: Unified multi-dimensional automatic evaluation for open-domain conversations with large language models**. In *Proceedings of the 5th Workshop on NLP for Conversational AI (NLP4ConvAI 2023)*, pages 47–58, Toronto, Canada. Association for Computational Linguistics.
- Zhaojiang Lin, Zihan Liu, Genta Indra Winata, Samuel Cahyawijaya, Andrea Madotto, Yejin Bang, Etsuko Ishii, and Pascale Fung. 2021. **Xpersona: Evaluating multilingual personalized chatbot**. In *Proceedings of the 3rd Workshop on Natural Language Processing for Conversational AI*, pages 102–112.
- 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.

- Yang Liu, Dan Iter, Yichong Xu, Shuohang Wang, Ruochen Xu, and Chenguang Zhu. 2023a. [G-eval: NLG evaluation using gpt-4 with better human alignment](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 2511–2522, Singapore. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized bert pretraining approach](#).
- Zeming Liu, Ping Nie, Jie Cai, Haifeng Wang, Zhengyu Niu, Peng Zhang, Mrinmaya Sachan, and Kaiping Peng. 2023b. [XDailyDialog: A multilingual parallel dialogue corpus](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12240–12253, Toronto, Canada. Association for Computational Linguistics.
- Philip M McCarthy. 2005. *An assessment of the range and usefulness of lexical diversity measures and the potential of the measure of textual, lexical diversity (MTLD)*. Ph.D. thesis, The University of Memphis.
- 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*, pages 130–141, Prague, Czechia. Association for Computational Linguistics.
- John Mendonça, Patrícia Pereira, Helena Moniz, Joao Paulo Carvalho, Alon Lavie, and Isabel Trancoso. 2023. [Simple LLM prompting is state-of-the-art for robust and multilingual dialogue evaluation](#). In *Proceedings of The Eleventh Dialog System Technology Challenge*, pages 133–143, Prague, Czech Republic. Association for Computational Linguistics.
- OpenAI. 2024. [Gpt-4 technical report](#).
- 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.
- Vikas Raunak, Amr Sharaf, Yiren Wang, Hany Awadalla, and Arul Menezes. 2023. [Leveraging GPT-4 for automatic translation post-editing](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 12009–12024, Singapore. Association for Computational Linguistics.
- Ananya B. Sai, Akash Kumar Mohankumar, Siddhartha Arora, and Mitesh M. Khapra. 2020. [Improving dialog evaluation with a multi-reference adversarial dataset and large scale pretraining](#). *Transactions of the Association for Computational Linguistics*, 8:810–827.
- Jaromir Savelka, Arav Agarwal, Christopher Bogart, Yifan Song, and Majd Sakr. 2023. [Can generative pre-trained transformers \(gpt\) pass assessments in higher education programming courses?](#) In *Proceedings of the 2023 Conference on Innovation and Technology in Computer Science Education V. 1, ITiCSE 2023*, page 117–123, New York, NY, USA. Association for Computing Machinery.
- Peifeng Wang, Zhengyang Wang, Zheng Li, Yifan Gao, Bing Yin, and Xiang Ren. 2023. [SCOTT: Self-consistent chain-of-thought distillation](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5546–5558, Toronto, Canada. Association for Computational Linguistics.
- Peter West, Chandra Bhagavatula, Jack Hessel, Jena Hwang, Liwei Jiang, Ronan Le Bras, Ximing Lu, Sean Welleck, and Yejin Choi. 2022. [Symbolic knowledge distillation: from general language models to commonsense models](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4602–4625, Seattle, United States. 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, Yiming Chen, Luis Fernando D’Haro, Yan Zhang, Thomas Friedrichs, Grandee Lee, and Haizhou Li. 2021. [DynaEval: Unifying turn and dialogue level evaluation](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 5676–5689, Online. Association for Computational Linguistics.

Chen Zhang, Luis D’Haro, Chengguang Tang, Ke Shi, Guohua Tang, and Haizhou Li. 2023. [xDial-eval: A multilingual open-domain dialogue evaluation benchmark](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 5579–5601, Singapore. Association for Computational Linguistics.

Chen Zhang, Luis Fernando D’Haro, Yiming Chen, Malu Zhang, and Haizhou Li. 2024. A comprehensive analysis of the effectiveness of large language models as automatic dialogue evaluators. *Proceedings of the AAAI Conference on Artificial Intelligence*, 39.

Saizheng Zhang, Emily Dinan, Jack Urbanek, Arthur Szlam, Douwe Kiela, and Jason Weston. 2018. [Personalizing dialogue agents: I have a dog, do you have pets too?](#) In *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.

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.

Ming Zhong, Yang Liu, Da Yin, Yuning Mao, Yizhu Jiao, Pengfei Liu, Chenguang Zhu, Heng Ji, and Jiawei Han. 2022. [Towards a unified multi-dimensional evaluator for text generation](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 2023–2038, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

A Dataset Curation

A.1 Generation

Prompt The prompt, which is shared for the development and test set is presented in Table 8. For each language, we translate the example dialogues

and responses using Google Translate⁹ and manually validate the full prompt with the expert linguists, ensuring the explanation is accurate for the translated response.

Given the dialog, generate a good and a bad response. In particular, the bad response should have issues that reduce its quality in terms of coherence, such as contradictions, logical inconsistencies, etc. Output the responses, together with a small explanation of the response using the following json format:

```
{"good_response": "...", "good_explanation": "...", "bad_response": "...", "bad_explanation": "..."}

```

Examples:

Dialogue: A: Have you figured out where you want to transfer to? B: I can’t think of where to go. A: Where would you like to go to school?

Output: {"good_response": "B: Well, It is not yet decided, but maybe in the east coast.", "good_explanation": "The response acknowledges the question and provides a region.", "bad_response": "B: Do you think that I can get married after school?", "bad_explanation": "The response does not acknowledge the prior question."}

Dialogue: A: You look so tan and healthy! B: Thanks. I just got back from summer camp A: How was it? B: Great. I got to try so many things for the first time.

Output: {"good_response": "A: I wish I could go to summer camp too. I’m so bored at home.", "good_explanation": "The response acknowledges the positive emotions displayed and contrasts it with their own perspective of summer break.", "bad_response": "A: Did you eat while you were there? You look frail.", "bad_explanation": "The response contradicts the earlier statement indicating they were healthy."}

Dialogue:

[Dialogue]

Table 8: Response generation instruction template.

Cost In total, our generation using a mixture of GPT-3.5-Turbo and GPT-4 amounted to approximately 300\$ USD, with 2/3 of the budget allocated to the generation of test set responses for both XDailyDialog and XPersona.

A.2 Additional Examples

Table 9, 10, 11 showcase additional examples resulting from GPT-4 response generation for XDailyDialog. Additionally, an example from XPersona for French is in Table 12. For convenience, we provide MT in *italic*.¹⁰

⁹<https://translate.google.com>

¹⁰These were not checked nor post-edited.

Context:

A: 你好，苗莉，你要去哪里？ *Hello, Miao Li, where are you going?*

B: 你好，我要去商店买些水果。 *Hello, I'm going to the store to buy some fruit.*

Positive Response. A: 哦，那真好，要买哪些水果？ *Oh, that's great. What kind of fruit to buy?*

"The response naturally continues the conversation by showing interest in the specific types of fruits B is planning to buy."

Negative Response. A: 为什么你现在才告诉我你要去商店？ *Why are you just telling me now that you're going to the store?*

"The response introduces a sense of confrontation and confusion that wasn't present in the prior conversation. It suggests an expectation that B should have informed A earlier, which is not supported by the dialogue."

Table 9: Example test set annotation for XDailyDialog for Chinese-ZH. MT in *Italic*.

Context:

A: Wo wollen Sie hin? *Where do you want to go?*

B: Ich fahre ins Krankenhaus. *I'm going to the hospital.*

A: Steigen Sie ein, ich bringe Sie hin. *Get in, I'll take you there.*

B: Wie viel wird es kosten? *How much will it cost?*

A: Es kostet Sie 10 Dollar. *It will cost you 10 dollars.*

B: Übrigens, wie spät ist es? *By the way, what time is it?*

Positive Response. A: Es ist jetzt halb drei nachmittags. *It's now half past two in the afternoon.*

"The response directly and appropriately answers B's question about the current time, maintaining the flow of the conversation."

Negative Response. A: Oh, ich habe keine Uhr, also weiß ich nie, wie spät es ist. *Oh, I don't have a watch, so I never know what time it is.*

"This response is inconsistent because as a driver, it would be expected that A has a way to tell the time, especially when their job involves scheduling. It contradicts the setting where a driver should be aware of the time to manage their routes and appointments."

Table 10: Example test set annotation for XDailyDialog for German-DE. MT in *Italic*.

A.3 Manual validation

For the manual validation step, a single annotator for each language is recruited to validate the response and corresponding explanation. A total of 100 randomly selected examples from the test set (generated by GPT-4) were sampled, per language, for this validation. We consider a response to be appropriate if the annotation is 1 or above (fair). The full guidelines provided to the expert annotators are presented in Figure 2.

Context:

A: Sto cercando il vecchio carillon. *I'm looking for the old music box.*

B: Sei arrivato nel posto giusto, ci sono dei requisiti cronologici? *You've come to the right place, are there any chronological requirements?*

A: Sarebbe fantastico se ci fosse un 1920 qui. *It would be great if there was a 1920's here.*

B: Ne abbiamo ricevuto solo 1 ieri e ora ce ne sono 6 qui. *We only received 1 yesterday and now there are 6 here.*

Positive Response. A: Davvero? Potrei vederli per scegliere il migliore? *Really? Could I see them to choose the best one?*

"The response shows interest and eagerness to see the selection which logically follows from the information provided about the availability of the items."

Negative Response. A: Oh, non mi interessano i carillon degli anni '20. *Oh, I don't care about 1920s music boxes.*

"The response is inconsistent as it contradicts the earlier expressed interest in a 1920s carillon, making it illogical and irrelevant to the previous exchange."

Table 11: Example test set annotation for XDailyDialog for Italian-IT. MT in *Italic*.

Context:

A: Salut! êtes-vous déjà monté à cheval? *Hi! have you ever ridden a horse?*

B: Oui! Ça fait partie de mon traitement contre le cancer. J'ai commencé un nouveau travail. *Toi? Yes! It's part of my cancer treatment. I started a new job. You?*

A: J'ai entendu dire que couper la viande de votre alimentation aide à combattre le cancer. *I heard cutting meat from your diet helps fight cancer.*

Positive Response. B: C'est intéressant, je devrais peut-être envisager cela aussi. Comment ça se passe pour toi? *That's interesting, maybe I should consider that too. How is it going for you?*

"The response shows interest in the advice given by A and seeks more personal insight, which maintains the flow of the conversation and shows engagement."

Negative Response. B: Oui, et c'est pourquoi je mange plus de viande maintenant pour rester en forme. *Yes, and that's why I eat more meat now to stay in shape.*

"The response is contradictory because A just mentioned that cutting meat helps fight cancer, yet B responds by saying they eat more meat to stay fit, which directly contradicts the advice given."

Table 12: Example test set annotation for Xpersona in French-FR. MT in *Italic*.

B Implementation Details

B.1 NSP-ROBERTA

We use the ROBERTA large encoder model downloaded from HuggingFace¹¹ for all experiments. We train a regression model on a single RTX A6000 GPU using the following sampling strategy: Given a fixed context from DailyDialog, **positive**

¹¹huggingface.co/roberta-large

responses are drawn directly from the same dialog; **negative** responses are randomly selected and a token coverage test discards semantically similar sentences. In total, 89,707/38,449 datapoints were obtained after processing.

A token representing the speaker was added for each turn, and a history length of 3 turns was used. 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 optimizer (Kingma and Ba, 2015), using a learning rate of 3e-6 and were trained for 3 epochs using a batch size of 16. Evaluation was conducted every 1,000 steps. The best performing model on the evaluation set was selected for testing.

B.2 ECoH

We train the ECoH models on a mixture of A100 80GB and RTX A6000 GPUs (depending on model size). We finetune using Huggingface Transformers and PEFT¹² for a 3 epochs for the English model and 1 epoch for the multilingual model with early stopping. We finetune from the base QWEN1.5-CHAT models (full precision) using LoRA (Hu et al., 2021), with $r = 8$, $\alpha = 32$ and dropout set to 0.1. Gradient accumulation steps is set to 4 with a learning rate of $1e - 4$. Batch size was set to maximize VRAM consumption, ranging from 2 up to 8 per device.

For inference, we follow QWEN1.5-CHAT inference code¹³, which generates responses using sampling with a temperature of 1, repetition penalty of 1.1, and top p set to 0.8.

C Additional Results

This appendix presents the individual results for zero shot, 1 shot with english example, 1 shot with target language example and the finetuned ECoH models for each individual language for for GENRESCO-H-DD-TEST, sorted by model size – 0.5B (Table 13), 1.8B and 4B (Table 14) and 7B and GPT-3.5-Turbo (Table 15). Table 16 presents the results for GENRESCO-H-PC-TEST.

Model	ρ_{pb}	F1	BLEU	GPT-4
QWEN1.5-0.5B-CHAT-0SHOT				
EN	.2141	.45	2.15	-
DE	.1382	.39	1.80	-
IT	.1695	.41	1.99	-
ZH	.1977	.44	1.88	-
QWEN1.5-0.5B-CHAT-1SHOT-EN				
EN	.2662	.60	2.69	2.12±1.20
DE	.1967	.55	2.19	-
IT	.2210	.55	2.47	-
ZH	.2361	.59	2.25	-
QWEN1.5-0.5B-CHAT-1SHOT-LANG				
EN	.2662	.60	2.69	2.12±1.20
DE	.1870	.59	4.35	1.40±0.77
IT	.1567	.56	4.46	1.84±1.24
ZH	.2803	.64	3.70	2.08±1.25
ECoH-0.5B-EN				
EN	.8995	.95	19.34	4.24±1.01
DE	.6407	.79	14.42	3.28±1.72
IT	.7035	.84	14.41	3.68±1.57
ZH	.8587	.93	15.89	3.92±1.41
ECoH-0.5B-ML				
EN	.9174	.96	19.34	4.08±1.15
DE	.8749	.94	14.42	4.04±1.30
IT	.8565	.93	14.41	3.48±1.66
ZH	.9038	.95	15.89	4.20±1.32

Table 13: Reported results for GENRESCO-H-DD-TEST for the 0.5B models. ρ_{pb} denotes Point Biserial Correlation. All correlation results are $p < 0.05$.

¹²huggingface.co/docs/peft

¹³github.com/QwenLM/Qwen1.5

Thank you for agreeing to annotate!

Our work is focused on generating adversarial dialogue responses. To this end, we asked GPT-4 to generate a positive and negative response. A negative response is characterized as having issues that reduce its quality in terms of coherence, such as contradictions and logical inconsistencies. **Your task is to determine if the GPT-4 generations are correct.**

Please provide a ternary answer (0, 1 or 2) to the following question:

Are the generated responses correct?

Your annotation(0-2):0 **bad** indicates that:

- one or both of the proposed responses are incorrect. That is, the good response is not coherent and/or the bad response is **without a doubt** coherent given the dialogue.
- The generated responses are nonsensical – there are clear fluency errors that make it difficult to understand what is being said.

Your annotation(0-2):1 **fair** indicates that:

- The explanation fails to adequately explain why the response is positive or negative.
- One or both responses are ambiguous.
- The responses have errors (or sound translationese) that do not affect the semantic understanding of the response.

Your annotation(0-2):2 **good** indicates that the generated response and corresponding explanations are correct and fluent (ignoring tokenization or special characters).

Figure 2: GPT-4 response validation guidelines.

Model	ρ_{pb}	F1	BLEU	GPT-4
QWEN1.5-1.8B-CHAT-0SHOT				
EN	.4765	.67	2.14	-
DE	.2663	.49	1.80	-
IT	.3207	.54	1.85	-
ZH	.4047	.62	1.95	-
QWEN1.5-1.8B-CHAT-1SHOT-EN				
EN	.5473	.74	2.70	2.36 \pm 1.38
DE	.4413	.68	2.50	-
IT	.4430	.68	2.76	-
ZH	.5652	.76	2.46	-
QWEN1.5-1.8B-CHAT-1SHOT-LANG				
EN	.5473	.74	2.70	2.36 \pm 1.38
DE	.4680	.71	2.38	2.08 \pm 1.22
IT	.4536	.72	3.03	2.28 \pm 1.31
ZH	.6160	.79	2.21	2.12 \pm 1.27
ECOH-1.8B-EN				
EN	.9227	.96	20.15	4.62 \pm 0.85
DE	.7432	.86	16.11	4.00 \pm 1.40
IT	.7381	.86	15.57	3.76 \pm 1.59
ZH	.8926	.95	17.35	4.12 \pm 1.33
ECOH-1.8B-ML				
EN	.9327	.97	20.08	4.32 \pm 1.31
DE	.8859	.94	15.92	4.04 \pm 1.50
IT	.8732	.94	15.23	4.04 \pm 1.40
ZH	.9159	.96	17.88	4.56 \pm 0.92
QWEN1.5-4B-CHAT-0SHOT				
EN	.7365	.86	3.57	-
DE	.6501	.82	3.49	-
IT	.6275	.81	3.55	-
ZH	.7138	.85	3.48	-
QWEN1.5-4B-CHAT-1SHOT-EN				
EN	.6163	.79	4.08	3.56 \pm 1.50
DE	.5764	.78	4.06	-
IT	.5728	.78	4.34	-
ZH	.5400	.73	3.80	-
QWEN1.5-4B-CHAT-1SHOT-LANG				
EN	.6163	.79	4.08	3.56 \pm 1.50
DE	.5754	.79	7.72	2.64 \pm 1.73
IT	.5269	.75	13.78	3.52 \pm 1.69
ZH	.6213	.80	7.04	3.60 \pm 1.50
ECOH-4B-EN				
EN	.9464	.97	20.66	4.60 \pm 0.91
DE	.8980	.95	17.17	4.64 \pm 0.77
IT	.8982	.95	16.42	3.92 \pm 1.19
ZH	.9315	.97	17.44	4.62 \pm 0.86
ECOH-4B-ML				
EN	.9631	.98	20.74	4.28 \pm 1.34
DE	.9437	.97	16.93	4.38 \pm 1.25
IT	.9377	.97	15.99	3.88 \pm 1.67
ZH	.9520	.98	18.52	4.34 \pm 1.00

Table 14: Reported results for GENRESCO-H-DD-TEST for the 1.8B and 4B models. ρ_{pb} denotes Point Biserial Correlation. All correlation results are $p < 0.05$.

Model	ρ_{pb}	F1	BLEU	GPT-4
QWEN1.5-7B-CHAT-0SHOT				
EN	.7490	.86	4.30	-
DE	.4868	.66	4.74	-
IT	.4302	.61	4.90	-
ZH	.6739	.81	4.70	-
QWEN1.5-7B-CHAT-1SHOT-EN				
EN	.8745	.94	4.62	3.76 \pm 1.63
DE	.7938	.90	4.75	-
IT	.7711	.88	4.85	-
ZH	.8210	.91	4.37	-
QWEN1.5-7B-CHAT-1SHOT-LANG				
EN	.8745	.94	4.62	3.76 \pm 1.63
DE	.7998	.90	4.59	3.64 \pm 0.45
IT	.6722	.81	5.07	3.76 \pm 1.78
ZH	.8208	.91	5.07	4.28 \pm 1.14
GPT-3.5-TURBO-0SHOT				
EN	.8592	.93	4.92	4.58 \pm 1.04
DE	.8218	.91	5.47	4.66 \pm 0.85
IT	.8102	.90	5.44	4.36 \pm 1.41
ZH	.8113	.90	5.18	4.54 \pm 1.23

Table 15: Reported results for GENRESCO-H-DD-TEST for the 7B models and GPT-3.5-Turbo. ρ_{pb} denotes Point Biserial Correlation. All correlation results are $p < 0.05$.

Model	ρ_{pb}	F1	BLEU	GPT-4
QWEN1.5-7B-CHAT-1SHOT				
EN	.5787	.75	3.28	4.28 \pm 1.45
FR	.4608	.66	2.97	3.20 \pm 1.58
IT	.6474	.82	3.06	3.60 \pm 1.63
ZH	.7630	.88	4.40	3.92 \pm 1.32
ECOH-0.5B-ML				
EN	.9021	.95	17.71	3.90 \pm 1.32
FR	.8089	.90	13.71	3.68 \pm 1.46
IT	.8661	.93	14.34	3.88 \pm 1.33
ZH	.9260	.96	15.84	3.80 \pm 1.38
ECOH-1.6B-ML				
EN	.9043	.95	18.17	4.44 \pm 1.12
FR	.8634	.93	14.01	4.40 \pm 1.15
IT	.8872	.94	14.65	4.04 \pm 1.13
ZH	.9390	.97	16.35	4.16 \pm 1.28
ECOH-4B-ML				
EN	.9443	.97	18.81	4.40 \pm 1.04
FR	.9270	.96	14.58	4.36 \pm 0.95
IT	.9381	.97	15.15	4.36 \pm 0.91
ZH	.9700	.98	16.79	4.40 \pm 0.96
GPT-3.5-TURBO-0SHOT				
EN	.7767	.89	4.71	4.08 \pm 1.38
FR	.7205	.86	5.04	4.43 \pm 1.24
IT	.8102	.90	5.11	4.64 \pm 0.86
ZH	.7452	.87	3.89	4.76 \pm 0.72

Table 16: Reported results for GENRESCO-H-PC-TEST. ρ_{pb} denotes Point Biserial Correlation. All correlation results are $p < 0.05$.