

Modeling the One-to-Many Property in Open-Domain Dialogue with LLMs

Jing Yang Lee¹, Kong Aik Lee², Woon Seng Gan³

School of Electrical and Electronic Engineering, Nanyang Technological University^{1,3}
Department of Electrical and Electronic Engineering, The Hong Kong Polytechnic University²
jingyang001@e.ntu.edu.sg¹, kong-aik.lee@polyu.edu.hk², ewsgan@ntu.edu.sg³

Abstract

Open-domain Dialogue (OD) exhibits a one-to-many (o2m) property, whereby multiple appropriate responses exist for a single dialogue context. Despite prior research showing that modeling this property boosts response diversity, most modern LLM-based dialogue agents do not explicitly do so. In this work, we model the o2m property of OD in LLMs by decomposing OD generation into two key tasks: Multi-Response Generation (MRG) and Preference-based Selection (PS), which entail generating a set of n semantically and lexically diverse high-quality responses for a given dialogue context, followed by selecting a single response based on human preference, respectively. To facilitate MRG and PS, we introduce o2mDial, a dialogue corpus explicitly designed to capture the o2m property by featuring multiple plausible responses for each context. Leveraging o2mDial, we propose new in-context learning and instruction-tuning strategies, as well as novel evaluation metrics for MRG, alongside a model-based approach for PS. Empirical results demonstrate that applying the proposed two-stage framework to smaller LLMs for OD generation enhances overall response diversity while maintaining contextual coherence, improving response quality by up to 90%, bringing them closer to the performance of larger models.

1 Introduction

Open-domain Dialogue (OD) agents are designed to engage in general conversation across various topics. They aim to generate responses that are fluent, diverse, and contextually coherent with respect to a given dialogue context. Unlike task-oriented agents with specific functions, OD agents simulate human-to-human interaction without predetermined conversational goals. This flexibility leads to the one-to-many (o2m) nature of OD, wherein multiple responses can be derived from a single dialogue context (Figure 1).

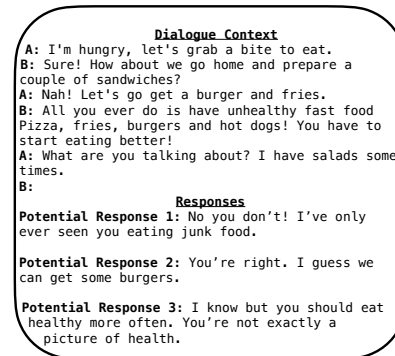


Figure 1: One-to-many property of open-domain dialogue.

Prior research has primarily focused on modeling the o2m property using probabilistic learning frameworks, such as the Conditional Variational Auto-Encoder (CVAE) (Shen et al., 2017; Zhao et al., 2017), to enhance response diversity. These methods typically condition the response on both the dialogue context and a randomly sampled latent variable, capturing the variability in conversational responses and effectively modeling the o2m property. Other approaches include randomized architectures (Lee et al., 2022b), Wasserstein Auto-encoders (Gu et al., 2018), and Bayesian architectures (Lee et al., 2023). These studies illustrate that while explicitly modeling the o2m property of OD significantly enhances response diversity, there is typically a trade-off with contextual coherence (Sun et al., 2021; Lee et al., 2022a).

Recent advancements in Large Language Models (LLMs) have made it increasingly impractical to model the o2m property using probabilistic approaches, primarily due to the immense scale of modern LLMs (Zhao et al., 2023). These frameworks typically employ a pretrained LLM as the decoder, which is fine-tuned along with additional network components responsible for generating the latent distribution. This process becomes highly resource-intensive given the scale of these LLMs.

Moreover, many state-of-the-art LLMs operate as black boxes with undisclosed parameters. Therefore, in the context of LLMs, adopting probabilistic frameworks for generating responses to model the o2m property has become largely impractical.

In this work, instead of adopting a probabilistic approach, we explore modeling the o2m property in LLMs by adopting a two-stage approach by decomposing OD response generation into two subtasks: Multi-Response Generation (MRG) and Preference-based Selection (PS). MRG aims to generate n distinct, contextually coherent responses from a single dialogue context, while PS selects the best response from these n options. For MRG and PS, we introduce o2mDial, a novel dataset designed to capture the o2m property of OD. Each sample in the dataset consists of a dialogue context paired with a set of semantically and lexically distinct yet equally fluent and contextually coherent potential responses. Our two-stage approach focuses on enhancing smaller LLMs (≤ 7 billion parameters), which often face challenges in generating diverse and contextually appropriate responses due to their limited capacity. Empirically, we demonstrate that this approach preserves contextual coherence while significantly increasing response diversity, leading to more engaging interactions with OD dialogue agents, particularly in smaller LLMs. Notably, through automatic and human evaluation, we show that our approach elevates the performance of these smaller models to levels comparable with larger LLMs, which require far greater computational resources. The dataset, metrics, and methodologies introduced in this work provide a valuable resource and baseline for future research into o2m response generation in LLMs.

This paper is organized as follows: MRG and PS are introduced in Section 2 and 3 respectively; Experimental results are provided in Section 4 and Section 5 concludes the paper.

2 o2mDial

To facilitate MRG, we curate o2mDial, a novel conversational dataset that explicitly captures the o2m property of OD. To create o2mDial, we leverage the DailyDialog corpus. First, we sample 500 dialogues (three to six turns) from the training set of the DailyDialog corpus. In this paper, for MRG, we fix $n = 5$. In other words, we aim to generate a set of five lexically and semantically distinct, yet contextually coherent responses. Unlike prior

Table 1: Corpus statistics.

# samples	500(train)/100(test)
Ave # turns per dialogue context	5.3
Ave # tokens	14.98 tokens

datasets that feature multiple reference responses (Hedayatnia et al., 2022; Sai et al., 2020; Gupta et al., 2019) that rely on the same LLM to generate every reference response, we use five distinct LLMs to simulate five different agents, with each LLM generating one response. As far as possible, this ensures the semantic and lexical uniqueness of each response. Based on our resource constraints, we selected the following five LLMs: 1) gpt-3.5-turbo (OpenAI, 2021); 2) llama2-70b-chat (et al., 2023); 3) mixtral-8x22b (et al., 2024); 4) Stable-Vicuna13b (Chiang et al., 2023); 5) Flan-T5-xxl (Chung et al., 2022). Additionally, to construct a separate test set for MRG evaluation, we sample another 100 dialogue samples from the test of the DailyDialog corpus. Similar to the training set, each turn consists of a dialogue context (three to six turns), and a set of five distinct and contextually coherent responses.

Given a dialogue context, each LLM was prompted to generate a one-sentence response. Furthermore, to ensure the quality of our corpus, we manually verify each sample for fluency and contextual coherence. Any responses found to be contextually incoherent or lacking in fluency were manually edited. A sample data point from our corpus is provided in 2. Some statistics regarding the training set of the collected corpus is provided in Table 1. For PS, we extend o2mDial with additional human preference labels (Section 4). Outside of our two-stage framework, o2mDial could be a useful resource for research for dialogue response evaluation or LLM response preference modeling.

3 Multi-Response Generation

MRG involves generating a set of n responses given a single dialogue context D . In this paper, R_n refers to the set of n responses generated by MRG, which consists of utterances alternating between two distinct interlocutors, i.e., $R_n = \{r_1, r_2, \dots, r_n\}$.

It is vital that responses in R_n are contextually coherent as well as lexically and semantically diverse. Semantic diversity requires each response to express a distinct idea, perspective, or piece of information. Lexical diversity involves variations

Dialogue Context

A: Brooke, are you still mad? Is your mom mad?
 B: No... mom said that we'll have the wedding at home on the 2nd. So let's hope your mom is right about that being a lucky day...
 A: It will be. And today is, too. I just talked to the airlines...
 B: They're going to give us back the seats?
 A: No, they're giving us discount tickets on a honeymoon package in...
 B:

Responses

LLM #1: The Maldives?
 LLM #2: Japan! That's amazing, I've always dreamed of going there. Thank you so much for taking care of everything. You're the best partner ever!
 LLM #3: Fiji? That sounds awesome, thank you for checking with the airlines!
 LLM #4: Bali! How amazing is that? We can finally get away from all this stress and enjoy our special moment together.
 LLM #5: Hawaii!

Figure 2: Sample dialogue context and response set pair from our corpus.

in word choice and structure, allowing responses to differ lexically even if they convey similar ideas. As a result, responses can have high lexical diversity but low semantic diversity. Contextual coherence ensures that responses are logically consistent and relevant to the dialogue context.

It should be noted that even if all n responses convey similar ideas, they can still be lexically unique by using different vocabularies or structures. Thus, a set of responses can exhibit high inter-response lexical diversity while maintaining low inter-response semantic diversity. For example, the statements "A heavy downpour is happening" and "There's a strong rainfall occurring" are lexically unique due to their different words and structures. However, the semantic content—that it is raining heavily—remains the same, indicating high semantic similarity. Our goal is to generate responses with both high inter-response lexical diversity and high inter-response semantic diversity, while ensuring contextual coherence. Prior work in MRG involve utilizing different sampling strategies, which produce responses with low semantic diversity, and pre-specified dialogue acts, which are significantly more complex to implement (Sakaeda and Kawahara, 2022).

3.1 Methods

In this section, we describe the In-Context Learning (ICL) and Instruction-tuning (IT) approaches we employ for MRG. Unlike prior approaches, we aim to generate R_n within a single inference:

$$R_n = LLM(\mathbf{P}(D_m)) \quad (1)$$

where \mathbf{P} refers to a specific prompt template, and $LLM(\cdot)$ denotes any arbitrary LLM. We implement the 3-shot variant of all prompts.

Few-shot (FS) Prompt This approach involves directly prompting the LLM to generate answers with the task description and demonstrations of query-proactive response pairs. In our experiments, 3 demonstration examples are used. The prompt template is provided in Figure 3.

Chain-of-thought (CoT) Prompt Chain-of-Thought (CoT) prompting (Wu et al., 2023) involves prompting the model to generate intermediate steps or explanations in addition to the final answer. In our case, we prompt the LLM to explain how each response differs from the other responses. We hypothesize that by prompting the model to identify the differences between generated responses, the model would be more inclined to generate lexically and semantically diverse responses. The prompt template is provided in Figure 4, located in the Appendix.

Prompt Chaining (PC) Prompt Chaining (PC) (Sun et al., 2024) typically involves dividing a task into smaller subtasks and executing them sequentially using prompts, where the output of one prompt serves as the input for the next. In our approach, we use PC to guide the LLM in generating a set of n unique responses one by one. The process begins with an initial prompt \mathbf{P}_0 that asks the LLM to generate a response to the dialogue context. Subsequent prompts ($\mathbf{P}_1 \cdots \mathbf{P}_{n-1}$) instructs the LLM to generate contextually coherent responses that differ semantically and lexically from every response generated by the previous prompts, which are included in the current prompt as input. We hypothesize that by decomposing the task of MRG into n smaller subtasks, the LLM can more effectively ensure both lexical and semantic uniqueness across the responses. However, it is important to note that PC requires multiple inferences from the LLM. Therefore, generating n responses requires n separate inferences, which could impact the feasibility and efficiency of this approach in the real world. The prompt template is provided in Figure 5, located in the Appendix.

Demonstration Selection Furthermore, we perform demonstration selection for the FS, CoT and PC prompts using specific metrics outlined earlier. Specifically, we select responses based on the mean of the inter-response semantic and lexical diversity scores: $sem(R_n) + lex(R_n)$. We identify the top- k responses from our corpus, where k refers to the number of demonstration examples required by the prompt.

Instruction Tuning (IT) In addition, we also con-

duct IT via QLoRA (Detmers et al., 2023) using the collected corpus. IT with QLoRA was performed using a batch size of 32, a learning rate of $2e-4$, 4 epochs, a rank of 16, an alpha of 32, and a dropout of 0.05. The instruction used for IT is identical to the zero-shot variant (prompt consists of only the instruction without any demonstration examples) of the FS prompt (Figure 3).

3.2 Evaluation

To measure MRG performance, we design automatic metrics to quantify inter-response semantic and lexical diversity, and overall contextual coherence of R_n .

Inter-response Semantic & Lexical Diversity In the context of open-domain dialogue, response diversity is typically measured via the Distinct metric, which is typically calculated by dividing the number of unique n-grams by the total number of n-grams. However, in our case, we aim to quantify the relative diversity of a set of n responses. In other words, we would like to measure, on average, how different each response is from the other $n - 1$ responses. Additionally, based on our definition, it would be ideal if semantic and lexical diversity can be evaluated separately. To this end, we define two separate metrics each accounting for either inter-response semantic or lexical diversity respectively: the inter-response semantic diversity score ($d_{sem}(R_n)$) and inter-response lexical diversity score ($d_{lex}(R_n)$).

For inter-response lexical diversity, we utilize the pairwise edit distance, namely the Jaccard similarity, between every possible response pair in the set:

$$d_{lex}(R_n) = \frac{1}{P_n} \sum_{i,j|i \in n, j \in n} \lambda_{Jac}(r_i, r_j) \quad (2)$$

where P_n refers to the total number of unique pairs in R_n , $\lambda_{Jac}(\cdot)$ refers to Jaccard similarity. Additionally, on occasion, when a LLM fails to generate the full set of n responses, a value of 1.0 would be assigned as the similarity score for that pair.

For inter-response semantic diversity, we compute the average of the pairwise semantic similarity via the Bert Score among responses in R_n :

$$d_{sem}(R_n) = \frac{1}{P_n} \sum_{i,j|i \in n, j \in n} \lambda_{BS}(r_i, r_j) \quad (3)$$

where P_n refers to the total number of unique pairs in R_n . Likewise, when a LLM fails to generate the

full set of n responses, a value of 1.0 would be assigned as the similarity score for that pair. Algorithms for computing d_{lex} and d_{sem} are provided in Algorithm 1 and 2, respectively.

Contextual Coherence For our task, the overall contextual coherence of a set of responses can be attained by averaging the individual scores attained by each of the n responses in R_n (Algorithm 3). We employ two contextual coherence metrics: the Utterance Entailment (UE) score (Lee et al., 2022a), and the UniEval-dialog coherence score (Zhong et al., 2022). The UE score involves framing the task of contextual coherence evaluation as a Natural Language Inference (NLI) task. For each utterance in the dialogue context and the corresponding generated response, a NLI model assesses whether the response entails, contradicts, or is neutral with respect to the utterance. For each response r_i from R_n , $UE(r_i) = \frac{1}{m} \sum_{j \in m} NLI(r_i, d_j)$. The UE score of a set of n responses is a continuous number between 0 and 1, where a greater value would indicate greater contextual coherence. The UniEval-dialog is a LLM-based approach which involve reframing response evaluation as a boolean question and answer task. Essentially a LLM is finetuned and prompted to generate either 'Yes' or 'No' to the question: 'Is this a coherent response given the dialogue history?'. Hence, for the UniEval-dialogue coherence metric, each response is assigned a score of 1 if 'Yes' is generated or 0 if 'No' is generated.

However, evaluating the contextual coherence of OD dialogue responses remains a challenging problem and an active area of research due to the o2m property (Li et al., 2016). Hence, in our experiments, we conduct a human evaluation to further support our findings.

4 Preference-based Selection (PS)

PS involves selecting the final response r_f from R_n based on human preference. Unlike traditional open-domain dialogue criteria such as coherence, diversity, engagingness, naturalness, or fluency, human preference covers broader factors like helpfulness, harmlessness, and interestingness (Li et al., 2024). We prioritize human preference for three key reasons. Firstly, MRG already ensures coherence and diversity within R_n . Secondly, existing metrics fall short in capturing the full complexity of human preferences, as they address only specific aspects of response quality (Jiang et al., 2024). Thirdly, modern LLMs are largely capable of gener-

ating fluent, natural, and engaging responses (Zhao et al., 2023).

However, in addition to human preference, the contextual coherence of the response should still be considered during selection. Hence, for PS, we aim to design an Open-domain Dialogue Response Preference (ODRP) model that assigns a scalar score to each response in R_n based on human preference. To achieve this, we leverage an open-source preference model from OpenAssistant (HuggingFace, 2021) based on deberta-v3-large commonly used for Reinforcement Learning with Human Feedback (RLHF) (et al., 2022). Such models are typically trained on preference datasets derived from tasks such as summarization (et al., 2020) and question answering (et al., 2021), or curated specifically to prevent harmful behavior (et al., 2022). Hence, to fine-tune the preference model for open-domain dialogue, we construct a new preference dataset from the corpus described earlier.

Preference datasets consist of comparisons between two responses given the same prompt (dialogue context D in our case). Extending o2mDial, we construct a preference dataset for fine-tuning by engaging annotators to label the preferred response y_c and the rejected response y_r (based on which they would prefer from a conversation partner) for every possible pair of responses from R_n , resulting in $\binom{n}{2}$ pairs per set. As per (Ouyang et al., 2022), we consider every pair from R_n as a single batch. The preference model is then fine-tuned on the following contrastive loss function:

$$J_\theta = \frac{1}{\binom{n}{2}} E_{(D, y_c, y_r) \sim R_n} [\log(\sigma(r(D, y_c), r(D, y_r)))] \quad (4)$$

We then proceed to fine-tune the preference model via QLoRA (Detmers et al., 2023) for two epochs with AdamW (lr=2e-4). After MRG, the ODRP model assigns a score to each response in R_n , and r_f is the response with the highest score:

$$r_f = \operatorname{argmax}_{r \in R_n} ODRP(r) \quad (5)$$

Additionally, we introduce a variant of the ODRP model finetuned on a subset of the corpus selected via hard negative sampling (Robinson et al., 2021). Specifically, we apply the base preference model to the dataset and deliberately extracted samples (50%) on which the base model performed the worst (assigned a similar score for both y_c and y_r

or assigned a higher score to y_r). We finetuned this variant of the ODRP model for four epochs instead.

5 Experimental Details

In this section, we outline our experimental design, providing specifics on the corpora utilized, the implementation of our framework, and the baseline approaches employed for comparison.

5.1 Corpora

For evaluation, we use two main datasets: DailyDialog (Li et al., 2017) and EmpatheticDialogs (Rashkin et al., 2019). DailyDialog features diverse, open-domain multi-turn conversations, while EmpatheticDialogs focuses on responses to emotionally grounded events. In our experiments, the dialogue agent’s task is to generate responses based solely on the context of the ongoing conversation. We do not use any additional information such as response labels (e.g., emotion, topic, or style) or speaker labels.

5.2 Implementation

We generate five responses per context ($n = 5$) using TinyLlama (v1.1b) (Zhang et al., 2024) and chat variants of Llama2-7b and Llama2-13b (et al., 2023). For all experiments, we aim to generate a set of five responses, i.e., $n = 5$. The temperature value used in all corpus creation and generation experiments are fixed at 0.7. We do not use other decoding strategies. All experiments were conducted using a single A100 GPU.

5.3 Baselines

For MRG, we implement in-context learning via Prompt Chaining (PC) as well as Few-Shot (FS) and Chain-of-Thought (CoT) prompting. We also evaluate Instruction Tuned (IT) variants of the LLM. Additionally, we also generate R_n via Multiple Inference (MI). MI entails directly feeding the dialogue context to the LLM and prompting the LLM to generate a single response n times.

For framework evaluation, we utilize PC to generate a response set for each dialogue context in the test set. Subsequently, for PS, we use the finetuned ODRP model ($ODRP$) as well as the variant finetuned on hard negative samples ($ODRP_{HN}$). Additionally, we introduce the following baseline response selection methods: 1) *rand*: Randomly selecting r_f from R_n ; 2) *cls*: Training a classifier (deberta-v2-large) from scratch with the curated

preference dataset; 3) *pref*: Using the base OpenAssistant preference model without fine-tuning; 4) base LLM (either *TinyLlama*, *Llama2-7b* or *Llama2-13b*): Generating a response by passing D directly to the LLM i.e., standard LLM inference. Additionally, we leverage the zero-shot variant of the FS prompt (Figure 3) to generate a single response from both *Llama2-70b* and *gpt-3.5-turbo*, allowing us to benchmark these against the responses produced by our framework when implemented with smaller LLMs.

5.4 Evaluation

Automatic Evaluation We evaluate the overall diversity and contextual coherence of the chosen responses by computing the inter-response Distinct-1,2 (Li et al., 2016) and the UE-score (Lee et al., 2022a) and UniEval-dialog coherence score (Zhong et al., 2022) respectively. To evaluate the set of responses R_n generated after MRG, we use several automatic metrics: inter-response semantic diversity (d_{sem}) and lexical diversity (d_{lex}) scores introduced in Section 3.2, as well as UE-score (UE), and UniEval-dialog coherence score (UniEval) to assess the quality of R_n . For inter-response diversity metrics, it should be highlighted that lower scores indicate greater lexical or semantic diversity.

Human Evaluation In our experiments, we also conduct a human evaluation to evaluate the efficacy of each PS approach. Similar to (Smith et al., 2022; Sakaeda and Kawahara, 2022), we engaged a group of five native English speaking participants for a comparative preference-based human evaluation. Each participant was presented a dialogue context along with a response generated by $ODRP_{HN}$ to compare against each of the other PS approaches (*base*, *rand*, *cls*, *ODRP*), as well as a response generated by *Llama2-70b* and *gpt-3.5-turbo*, and told to select the agent they would rather converse with. Each participant was presented with 60 samples (30 from DailyDialog and 30 from EmpatheticDialogs) for each comparison. We report the Win, Tie and Loss percentage of each comparison.

In addition, we also conduct a human evaluation to evaluate the quality of the set of responses R_n generated during MRG. For this evaluation, we engage a separate group of five native English speakers. Given a set of five responses, each participant was told to count the number of semantically unique responses, lexically unique responses, and contextually coherent responses. Hence, each score is a discrete value from 1 to 5. A count of 5 would

imply that all 5 responses were either semantically unique, lexically unique, or contextually coherent. Conversely, a count of 0 would indicate that all 5 responses were semantically similar, lexically similar, or contextually incoherent. Naturally, the participants were not informed which LLM or which generation approach was responsible for each response set. For our generation experiments, each participant was provided with 60 samples (30 from DailyDialog and 30 from EmpatheticDialogs) from each generation approach (the 3 shot variant of each prompt as well as IT, MI, *Llama2-70b* and *gpt-3.5-turbo*). Each output consisted of a set of five responses. To illustrate this process, we provide a sample evaluation in Figure 6, located in the Appendix.

6 Results & Discussion

Here, we assess the performance of the proposed two-stage framework. We also analyze the set responses generated during MRG based on the metrics outlined in Section 3.

6.1 Framework Evaluation

The automatic and human evaluation results are presented in Table 2 and Table 3, respectively. Sample responses are provided in Figure 7 in the Appendix.

Based on the results obtained, it is clear that the responses selected by $ODRP$ and $ODRP_{HN}$ consistently outperform all other approaches, including *rand*, *cls*, and *pref*, in terms of both diversity and contextual coherence. Both $ODRP$ and $ODRP_{HN}$ generally achieve statistically significantly higher Distinct and UE/UniEval scores than the baseline methods. Moreover, in human evaluation, they show a greater proportion of wins and a lower proportion of losses compared to other baselines. Qualitatively, we observe that responses selected by $ODRP$ and $ODRP_{HN}$ do more than just acknowledge the previous utterance; they often provide additional enriching information that enhances the overall dialogue. Furthermore, a significant portion of these selected responses include queries directed at the other interlocutor, actively encouraging further interaction.

It is also important to note that fine-tuning the $ODRP$ model with hard negative samples leads to a noticeable improvement in the diversity and coherence of the selected responses across all LLMs. $ODRP_{HN}$ outperforms $ODRP$ on all automatic metrics and achieves a higher Win rate and lower

Loss rate in human evaluation. The effectiveness of the ODRP model is particularly evident in the case of TinyLlama, where there is substantial variability in the quality of responses generated during MRG. Generally, we observe that the ODRP model excels at identifying and prioritizing higher-quality responses, resulting in more engaging and meaningful exchanges, even when the initial set of responses exhibits significant variability. This leads to improvements of up to 90% in response diversity and contextual coherence.

Comparison with Larger LLMs In addition, we evaluated larger LLMs, such as Llama2-70b and gpt-3.5-turbo, using the zero-shot variant of the FS prompt (Figure 3). Our findings reveal that after applying our two-stage framework and selecting responses via $ODRP_{HN}$, the quality of responses generated by smaller LLMs like *TinyLlama* and *Llama2-7b* surpasses that of *Llama2-70b* in terms of response diversity and approaches the level of *gpt-3.5-turbo*. Regarding contextual coherence, *Llama2-13b* see improvements that bring it in line with *Llama2-70b* and *gpt-3.5-turbo*, while *TinyLlama* and *Llama2-7b*, although still trailing, narrow the gap significantly. Qualitatively, we note that responses selected by $ODRP_{HN}$ are comparable to responses generated by *Llama2-70b* and *gpt-3.5-turbo* in terms of naturalness and engagement. These results underscore the effectiveness of our approach, enabling smaller LLMs to rival or exceed the capabilities of larger models, all while maintaining lower computational demands.

6.2 MRG Evaluation

In addition, we evaluate the MRG performance of 3-shot FS, CoT, PC, and IT on the o2mDial test set. Automatic and human evaluation results are presented in Table 4.

We observe that larger LLMs like Llama2-7b and 13b generally outperform TinyLlama, likely due to their superior instruction-following abilities, which enhance in-context learning and IT effectiveness. The PC and IT methods yield results comparable to reference responses in the test set for Llama2-7b and 13b, while TinyLlama lags slightly, reflecting its weaker capabilities. Despite TinyLlama’s limitations, PC’s simpler task breakdown marginally improved performance, outperforming all other baseline MRG methods. Llama2-7b and 13b also benefited from PC and CoT prompts, boosting response diversity while preserving contextual coherence, as shown by com-

parable UE/UniEval scores.

Closer examination of the responses reveal that quality rises with model size—TinyLlama produces the weakest outputs, while Llama2-13b excels. All three models faced issues: insufficient responses (below n), redundancy (similar or identical replies), and over-extended conversations (too many utterances). Insufficient and redundant responses reduced semantic and lexical diversity, while over-extensions impacted coherence metrics like UE and UniEval scores. TinyLlama had more insufficient responses, Llama2-7b and 13b saw occasional over-extensions, and redundancy appeared across all models, most prominently in TinyLlama. Generally, there remains a performance gap between the reference responses and proposed approaches. Future work will aim to reduce this gap.

Comparison with MI Response sets generated via MI tend to be semantically similar despite relatively high lexical diversity, as shown by low inter-response semantic scores and comparably higher lexical diversity scores in both automatic and human evaluations. This is likely due to the deterministic nature of logits during inference. Although sampling strategies (temperature scaling (Guo et al., 2017) or nucleus sampling (Holtzman et al., 2020)) introduce stochasticity in decoding, generated logits remain deterministic, limiting semantic variation unless randomness is significantly increased, which could reduce contextual coherence.

7 Related Work

Prior work adopting a two-stage approach for open-domain dialogue typically involves generating multiple responses either through conditional generation based on pre-specified dialogue acts (Sakaeda and Kawahara, 2022) or by pooling outputs from variational and retrieval-based systems (Ruan et al., 2020; of Physics and Technology, 2021). However, these studies often focus on evaluating only the final selected response, without considering the diversity or contextual coherence of the entire set of generated responses. In contrast, our approach evaluates and optimizes the quality of the full set of responses, thereby enhancing the overall quality of the final selected response. Additionally, many of these methods have been applied to smaller language models, whereas to the best of our knowledge, our work is the first to introduce a two-stage generation framework LLMs. Other two-stage approaches broadly entail first generating a candidate

Table 2: Automatic evaluation results. The best score in each column is **bolded**. * indicates a statistically significant difference in score (t-test, p -value < 0.01) from the **bolded** score. Scores for DailyDialog and EmpatheticDialogues are provided before and after the backslash '\', respectively.

	Dist-1	Dist-2	UE	UniEval
<i>TinyLlama</i>	0.16*/0.18*	0.51*/0.61*	0.21*/0.13*	0.74*/0.64*
- <i>rand</i>	0.24*/0.20*	0.75*/0.70*	0.24/0.13*	0.76*/0.65*
- <i>cls</i>	0.22*/0.25*	0.76*/0.74*	0.23/0.18*	0.78*/0.66*
- <i>pref</i>	0.25*/0.24*	0.73*/0.75*	0.24/0.18*	0.77*/0.70*
- <i>ODRP</i>	0.28*/0.29	0.77/0.798	0.27/0.22*	0.81/0.72*
- <i>ODRP_{HN}</i>	0.31/0.31	0.79/0.82	0.30/0.26	0.83/0.76
<i>Llama2-7b</i>	0.20*/0.22*	0.61*/0.69*	0.24*/0.21	0.83/0.72
- <i>rand</i>	0.23*/0.30*	0.77*/0.78*	0.22*/0.19*	0.81*/0.69*
- <i>cls</i>	0.30*/0.27*	0.79*/0.75*	0.23*/0.18*	0.80*/0.72
- <i>pref</i>	0.28/0.29*	0.77*/0.78*	0.24*/0.22	0.83/0.71*
- <i>ODRP</i>	0.33/0.35	0.83/0.84	0.26/0.22	0.83/0.73
- <i>ODRP_{HN}</i>	0.35/0.36	0.83/0.85	0.29/0.24	0.85/0.73
<i>Llama2-13b</i>	0.21*/0.23*	0.65*/0.72*	0.26*/0.24*	0.85/0.77*
- <i>rand</i>	0.24*/0.28*	0.77*/0.76*	0.25*/0.24*	0.80*/0.72*
- <i>cls</i>	0.30*/0.31*	0.80*/0.76*	0.29*/0.25	0.83*/0.77*
- <i>pref</i>	0.31/0.30*	0.79*/0.78	0.26*/0.29	0.82*/0.79
- <i>ODRP</i>	0.33/0.34	0.85/0.79	0.32/0.30	0.85/0.81
- <i>ODRP_{HN}</i>	0.33/0.35	0.84/0.82	0.33/0.32	0.87/0.82
<i>Llama2-70b</i>	0.31/0.32	0.72/0.80	0.28/0.26	0.86/0.79
<i>gpt-3.5-turbo</i>	0.36/0.33	0.75/0.82	0.31/0.30	0.88/0.81

response and instantiating it as the final response (Li et al., 2023), or generating a response in the first stage and further conditioning and refining the response in the second stage (Qian et al., 2024; Shi and Song, 2023).

Regarding response selection, prior work has primarily concentrated on narrow criteria such as engagement (Sakaeda and Kawahara, 2022), topical relevance (Ruan et al., 2020; Yuan et al., 2024). Standard retrieval-based systems, in contrast, prioritize contextual coherence (Tao et al., 2021; Su et al., 2024). In our framework, we prioritize human preferences, considering a broader range of factors such as harmlessness and helpfulness, which are critical aspects for ensuring the real-world utility of response generation systems.

8 Conclusion

This paper decomposes OD response generation into Multi-Response Generation (MRG) and Preference-based Selection (PS). For MRG, we curate o2mDial and propose methods such as FS, CoT, PC, and IT. We also introduce metrics to evaluate semantic and lexical diversity. For PS, we develop the ODRP model to select responses aligned with human preferences. Empirical results show MRG and PS significantly enhance response diversity

Table 3: Human evaluation results. The Win, Tie, and Loss percentages are presented for each comparison.

		Win	Tie	Loss
TinyLlama	<i>ODRP_{HN}</i> vs. <i>TinyLlama</i>	85	9	6
	<i>ODRP_{HN}</i> vs. <i>rand</i>	76	16	18
	<i>ODRP_{HN}</i> vs. <i>cls</i>	60	29	11
	<i>ODRP_{HN}</i> vs. <i>pref</i>	57	20	23
	<i>ODRP_{HN}</i> vs. <i>ODRP</i>	49	33	18
	<i>ODRP_{HN}</i> vs. <i>Llama2-70b</i>	30	35	35
Llama2-7b	<i>ODRP_{HN}</i> vs. <i>gpt-3.5-turbo</i>	26	44	30
	<i>ODRP_{HN}</i> vs. <i>Llama2-7b</i>	74	18	8
	<i>ODRP_{HN}</i> vs. <i>rand</i>	58	25	17
	<i>ODRP_{HN}</i> vs. <i>cls</i>	50	29	21
	<i>ODRP_{HN}</i> vs. <i>pref</i>	47	27	23
	<i>ODRP_{HN}</i> vs. <i>ODRP</i>	46	30	24
Llama2-13b	<i>ODRP_{HN}</i> vs. <i>Llama2-70b</i>	32	41	27
	<i>ODRP_{HN}</i> vs. <i>gpt-3.5-turbo</i>	28	48	24
	<i>ODRP_{HN}</i> vs. <i>Llama2-13b</i>	50	33	17
	<i>ODRP_{HN}</i> vs. <i>rand</i>	51	24	25
	<i>ODRP_{HN}</i> vs. <i>cls</i>	44	34	22
	<i>ODRP_{HN}</i> vs. <i>pref</i>	42	30	28
Llama2-70b	<i>ODRP_{HN}</i> vs. <i>ODRP</i>	41	32	27
	<i>ODRP_{HN}</i> vs. <i>Llama2-70b</i>	38	39	23
	<i>ODRP_{HN}</i> vs. <i>gpt-3.5-turbo</i>	37	40	23

Table 4: MRG automatic and human evaluation results on the o2mDial test set.

Model		d_{sem}	d_{lex}	UE	UniEval
TinyLlama	MI	0.86	0.78	0.20	0.73
	FS	0.66*	0.75*	0.19*	0.72*
	CoT	0.67*	0.74*	0.21*	0.74*
	PC	0.64	0.70*	0.25*	0.77*
	IT	0.65*	0.72*	0.23*	0.75*
Llama2-7b	MI	0.81	0.76	0.24	0.82
	FS	0.65*	0.74*	0.25*	0.80*
	CoT	0.62	0.67*	0.28*	0.86
	PC	0.60	0.65*	0.28*	0.87
	IT	0.65*	0.68*	0.26*	0.84*
Llama2-13b	MI	0.74	0.70	0.28	0.84
	FS	0.61	0.68*	0.29*	0.85*
	CoT	0.60	0.65*	0.28*	0.88
	PC	0.60	0.66*	0.30	0.88
	IT	0.61	0.67*	0.29*	0.87
Reference		0.60	0.62	0.32	0.89
		Sem. Div.	Lex. Div.	Con. Coh.	κ
TinyLlama	MI	1.89	1.95	3.95	0.54
	FS	3.42	3.82	3.95	0.55
	CoT	3.58	3.88	3.91	0.54
	PC	3.70	3.96	3.98	0.51
	IT	3.75	4.01	3.99	0.49
Llama2-7b	MI	2.33	2.45	4.73	0.58
	FS	4.30	4.60	4.79	0.57
	CoT	4.44	4.72	4.85	0.59
	PC	4.58	4.73	4.85	0.66
	IT	4.53	4.70	4.70	0.60
Llama2-13b	MI	2.67	2.92	4.88	0.47
	FS	4.44	4.66	4.82	0.50
	CoT	4.65	4.74	4.88	0.58
	PC	4.66	4.71	4.89	0.59
	IT	4.55	4.69	4.80	0.54
Reference		4.69	4.77	4.89	0.58

by up to 90% in smaller LLMs, achieving performance on par with larger LLMs. Future research could expand the number of unique responses per set (beyond $n = 5$) to assess impacts on diversity and quality. Systematically increasing n could help identify the optimal point of diminishing returns. For PS, another potential avenue for additional research could involve integrating dialogue context into the evaluation process to act as a safeguard against contextually incoherent responses.

9 Limitations

Due to resource limitations, the LLMs employed for dataset curation in our experiments are intentionally smaller in size. Future work could entail extending o2mDial with larger, more recent LLMs. Furthermore, due to time and resource constraints, exhaustive prompt engineering was not performed for each model. Instead, we focused on basic prompt engineering techniques aimed at ensuring consistent and coherent output formatting. While this approach was sufficient for the scope of the experiments, we acknowledge that more sophisticated and fine-tuned prompt engineering could potentially improve the models' performance in more complex or specialized tasks.

References

- Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E. Gonzalez, Ion Stoica, and Eric P. Xing. 2023. [Vicuna: An open-source chatbot impressing gpt-4 with 90% chatgpt quality](#).
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, Albert Webson, Shixiang Shane Gu, Zhuyun Dai, Mirac Suzgun, Xinyun Chen, Aakanksha Chowdhery, Alex Castro-Ros, Marie Pellat, Kevin Robinson, Dasha Valter, Sharan Narang, Gaurav Mishra, Adams Yu, Vincent Zhao, Yanping Huang, Andrew Dai, Hongkun Yu, Slav Petrov, Ed H. Chi, Jeff Dean, Jacob Devlin, Adam Roberts, Denny Zhou, Quoc V. Le, and Jason Wei. 2022. [Scaling instruction-finetuned language models](#). *Preprint*, arXiv:2210.11416.
- Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. 2023. [Qlora: Efficient finetuning of quantized llms](#). *Preprint*, arXiv:2305.14314.
- Albert Q. Jiang et al. 2024. [Mixtral of experts](#). *Preprint*, arXiv:2401.04088.
- Hugo Touvron et al. 2023. [Llama 2: Open foundation and fine-tuned chat models](#). *Preprint*, arXiv:2307.09288.
- N. Stiennon et al. 2020. Learning to summarize from human feedback. In *NeurIPS*.
- R. Nakano et al. 2021. Webgpt: Browser-assisted question-answering with human feedback. In *arXiv*.
- Yuntao Bai et al. 2022. [Training a helpful and harmless assistant with reinforcement learning from human feedback](#). *Preprint*, arXiv:2204.05862.
- Xiaodong Gu, Kyunghyun Cho, Jung-Woo Ha, and Sunghun Kim. 2018. DialogWAE: Multimodal response generation with conditional wasserstein auto-encoder.
- Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q. Weinberger. 2017. On calibration of modern neural networks. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70, ICML'17*, page 1321–1330. JMLR.org.
- Prakhar Gupta, Shikib Mehri, Tiancheng Zhao, Amy Pavel, Maxine Eskenazi, and Jeffrey Bigham. 2019. [Investigating evaluation of open-domain dialogue systems with human generated multiple references](#). In *Proceedings of the 20th Annual SIGdial Meeting on Discourse and Dialogue*, pages 379–391, Stockholm, Sweden. Association for Computational Linguistics.
- Behnam Hedayatnia, Di Jin, Yang Liu, and Dilek Hakkani-Tur. 2022. [A systematic evaluation of response selection for open domain dialogue](#). In *Proceedings of the 23rd Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 298–311, Edinburgh, UK. Association for Computational Linguistics.
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2020. [The curious case of neural text degeneration](#). *Preprint*, arXiv:1904.09751.
- HuggingFace. 2021. [Openassistant/reward-model-deberta-v3-large-v2](#). <https://huggingface.co/OpenAssistant/reward-model-deberta-v3-large-v2>.
- Ruili Jiang, Kehai Chen, Xuefeng Bai, Zhixuan He, Juntao Li, Muyun Yang, Tiejun Zhao, Liqiang Nie, and Min Zhang. 2024. [A survey on human preference learning for large language models](#). *Preprint*, arXiv:2406.11191.
- Jing Yang Lee, Kong Aik Lee, and Woon Seng Gan. 2022a. [Improving contextual coherence in variational personalized and empathetic dialogue agents](#). In *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7052–7056.

- Jing Yang Lee, Kong Aik Lee, and Woon Seng Gan. 2022b. [A randomized link transformer for diverse open-domain dialogue generation](#). In *Proceedings of the 4th Workshop on NLP for Conversational AI*, pages 1–11, Dublin, Ireland. Association for Computational Linguistics.
- Jing Yang Lee, Kong Aik Lee, and Woon Seng Gan. 2023. [An empirical Bayes framework for open-domain dialogue generation](#). In *Proceedings of the Third Workshop on Natural Language Generation, Evaluation, and Metrics (GEM)*, pages 192–204, Singapore. Association for Computational Linguistics.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2016. [A diversity-promoting objective function for neural conversation models](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 110–119, San Diego, California. Association for Computational Linguistics.
- Junlong Li, Fan Zhou, Shichao Sun, Yikai Zhang, Hai Zhao, and Pengfei Liu. 2024. [Dissecting human and LLM preferences](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1790–1811, Bangkok, Thailand. Association for Computational Linguistics.
- Shaobo Li, Chengjie Sun, Zhen Xu, Prayag Tiwari, Bingquan Liu, Deepak Gupta, K. Shankar, Zhenzhou Ji, and Mingjiang Wang. 2023. [Toward explainable dialogue system using two-stage response generation](#). *ACM Trans. Asian Low-Resour. Lang. Inf. Process.*, 22(3).
- 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.
- Moscow Institute of Physics and Technology. 2021. [Dream technical report for the alexa prize 4](#).
- OpenAI. 2021. Gpt-3.5 turbo. <https://platform.openai.com/docs/models/gpt-3-5-turbo>.
- 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](#). *Preprint*, arXiv:2203.02155.
- Yushan Qian, Bo Wang, Shangzhao Ma, Wu Bin, Shuo Zhang, Dongming Zhao, Kun Huang, and Yuexian Hou. 2024. [Think twice: A human-like two-stage conversational agent for emotional response generation](#). *Preprint*, arXiv:2301.04907.
- Hannah Rashkin, Eric Michael Smith, Margaret Li, and Y-Lan Boureau. 2019. [Towards empathetic open-domain conversation models: A new benchmark and dataset](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5370–5381, Florence, Italy. Association for Computational Linguistics.
- Joshua Robinson, Ching-Yao Chuang, Suvrit Sra, and Stefanie Jegelka. 2021. [Contrastive learning with hard negative samples](#). In *International Conference on Learning Representations (ICLR)*.
- Yu-Ping Ruan, Zhen-Hua Ling, Xiaodan Zhu, Quan Liu, and Jia-Chen Gu. 2020. [Generating diverse conversation responses by creating and ranking multiple candidates](#). *Comput. Speech Lang.*, 62(C).
- 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.
- Ryoma Sakaeda and Daisuke Kawahara. 2022. [Generate, evaluate, and select: A dialogue system with a response evaluator for diversity-aware response generation](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Student Research Workshop*, pages 76–82, Hybrid: Seattle, Washington + Online. Association for Computational Linguistics.
- Xiaoyu Shen, Hui Su, Yanran Li, Wenjie Li, Shuzi Niu, Yang Zhao, Akiko Aizawa, and Guoping Long. 2017. [A conditional variational framework for dialog generation](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 504–509, Vancouver, Canada. Association for Computational Linguistics.
- Tianyuan Shi and Yongduan Song. 2023. [A novel two-stage generation framework for promoting the persona-consistency and diversity of responses in neural dialog systems](#). *IEEE Transactions on Neural Networks and Learning Systems*, 34(3):1552–1562.
- Eric Smith, Orion Hsu, Rebecca Qian, Stephen Roller, Y-Lan Boureau, and Jason Weston. 2022. [Human evaluation of conversations is an open problem: comparing the sensitivity of various methods for evaluating dialogue agents](#). In *Proceedings of the 4th Workshop on NLP for Conversational AI*, pages 77–97, Dublin, Ireland. Association for Computational Linguistics.
- Zhenpeng Su, Xing W, Wei Zhou, Guangyuan Ma, and Songlin Hu. 2024. [Dial-MAE: ConTextual masked auto-encoder for retrieval-based dialogue systems](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 820–830, Mexico

- City, Mexico. Association for Computational Linguistics.
- Bin Sun, Shaoxiong Feng, Yiwei Li, Jiamou Liu, and Kan Li. 2021. [Generating relevant and coherent dialogue responses using self-separated conditional variational AutoEncoders](#). 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 5624–5637, Online. Association for Computational Linguistics.
- Shichao Sun, Ruifeng Yuan, Ziqiang Cao, Wenjie Li, and Pengfei Liu. 2024. [Prompt chaining or step-wise prompt? refinement in text summarization](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 7551–7558, Bangkok, Thailand. Association for Computational Linguistics.
- Chongyang Tao, Jiazhan Feng, Rui Yan, Wei Wu, and Daxin Jiang. 2021. [A survey on response selection for retrieval-based dialogues](#). In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, pages 4619–4626. International Joint Conferences on Artificial Intelligence Organization. Survey Track.
- Dingjun Wu, Jing Zhang, and Xinmei Huang. 2023. [Chain of thought prompting elicits knowledge augmentation](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 6519–6534, Toronto, Canada. Association for Computational Linguistics.
- Wei Yuan, Zongyang Ma, Aijun An, and Jimmy Xiangji Huang. 2024. [Topic-aware response selection for dialog systems](#). *Natural Language Processing Journal*, 8:100087.
- Peiyuan Zhang, Guangtao Zeng, Tianduo Wang, and Wei Lu. 2024. [Tinyllama: An open-source small language model](#). *Preprint*, arXiv:2401.02385.
- Tiancheng Zhao, Ran Zhao, and Maxine Eskenazi. 2017. [Learning discourse-level diversity for neural dialog models using conditional variational autoencoders](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 654–664, Vancouver, Canada. Association for Computational Linguistics.
- Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, Yifan Du, Chen Yang, Yushuo Chen, Zhipeng Chen, Jinhao Jiang, Ruiyang Ren, Yifan Li, Xinyu Tang, Zikang Liu, Peiyu Liu, Jian-Yun Nie, and Ji-Rong Wen. 2023. [A survey of large language models](#). *Preprint*, arXiv:2303.18223.
- Ming Zhong, Yang Liu, Da Yin, Yuning Mao, Yizhu Jiao, Pengfei Liu, Chenguang Zhu, Heng Ji, and

A Appendix

Algorithm 1 Inter-response lexical similarity score d_{lex} .

Require: Set of n responses R_n , Jaccard Similarity function $J(\cdot)$

Ensure: Lexical similarity score s

```
 $s_t \leftarrow 0$  {Initialize temporary score}
 $P \leftarrow 0$  {Initialize pair count}
for  $i \leftarrow 0$  to  $n - 1$  do
  for  $j \leftarrow i + 1$  to  $n - 1$  do
    if  $r_i = \text{None}$  or  $r_j = \text{None}$  then
       $s_t \leftarrow s_t + 1.0$ 
    else
       $s_t \leftarrow s_t + \lambda_{Jac}(r_i, r_j)$ 
    end if
     $P \leftarrow P + 1$  {Increment pair count}
  end for
end for
 $s \leftarrow \frac{1}{P} s_t$  {Compute mean over all pairs}
return  $s$ 
```

Algorithm 2 Inter-response semantic similarity score d_{sem} .

Require: Set of n responses R_n , BertScore function $BS(\cdot)$

Ensure: Semantic similarity score s

```
 $s_t \leftarrow 0$  {Initialize temporary score}
 $P \leftarrow 0$  {Initialize pair count}
for  $i \leftarrow 0$  to  $n - 1$  do
  for  $j \leftarrow i + 1$  to  $n - 1$  do
    if  $r_i = \text{None}$  or  $r_j = \text{None}$  then
       $s_t \leftarrow s_t + 1.0$ 
    else
       $s_t \leftarrow s_t + \lambda_{BS}(r_i, r_j)$ 
    end if
     $P \leftarrow P + 1$  {Increment pair count}
  end for
end for
 $s \leftarrow \frac{1}{P} s_t$  {Compute mean over all pairs}
return  $s$ 
```

Algorithm 3 Contextual Coherence score

Require: Set of n responses R_n , set of m dialogue context D_m , Contextual Coherence measure $CC(\cdot)$ (e.g., UE score or UniEval-dialogue coherence score)

Ensure: Contextual coherence score s

```
 $s_t \leftarrow 0$  {Initialize temporary score}  
for  $i \leftarrow 0$  to  $n - 1$  do  
  if  $r_i = \text{None}$  then  
     $s_t \leftarrow s_t + 0.0$   
  else  
     $s_t \leftarrow s_t + CC(r_i, D_m)$   
  end if  
end for  
 $s \leftarrow \frac{1}{n} s_t$  {Compute mean over  $n$  responses}  
return  $s$ 
```

Few-Shot Prompt Template

```
<Instruction>  
Given the following dialogue context, adopt the  
role of B and generate five lexically and  
semantically unique responses.  
  
<Demo Example #1>  
\n\nDialogue Context:{demo_context}  
\n\nResponses: {demo_response_set}  
<Demo Example #2>  
<Demo Example #3>  
  
<Test>  
\n\nDialogue Context:{test_context}  
\n\nResponses:
```

Figure 3: Prompt template for the Few-Shot prompt.

CoT Prompt Template

```
<Instruction>  
Given the following dialogue context, adopt the  
role of B and generate five lexically and  
semantically unique responses. After generating  
all five responses, in a single paragraph,  
explain how each response differs from the  
others.  
  
<Demo Example #1>  
\n\nDialogue Context:{demo_context}  
\n\nResponses: {demo_response_set}  
\n\nExplanation: {explanation}  
<Demo Example #2>  
<Demo Example #3>  
  
<Test>  
\n\nDialogue Context:{test_context}  
\n\nResponses:
```

Figure 4: Prompt template for the Chain-of-Thought prompt.

Prompt Chain (P₀) Template

```

<Instruction>
Given the following dialogue context, adopt the
role of B and generate a response.
<Demo Example #1>
\n\nDialogue Context:{demo_context}
\n\nResponse: {demo_response}
<Demo Example #2>
<Demo Example #3>
<Test>
\n\nDialogue Context:{test_context}
\n\nResponse:

```

Prompt Chain (P₁) Template

```

<Instruction>
Given the following dialogue context, adopt the
role of B and generate a response. Ensure that
the response generated is semantically and
lexically distinct with respect to the provided
response list.
<Demo Example #1>
\n\nDialogue Context:{demo_context}
\n\nResponse List: {demo_response_list}
\n\nResponse: {demo_response}
<Demo Example #2>
<Demo Example #3>
<Test>
\n\nDialogue Context:{test_context}
\n\nResponse List: {test_response_list}
\n\nResponses:

```

Figure 5: Prompt template for the Prompt Chain (PC).

Dialogue Context

A: Brooke, are you still mad? Is your mom mad?
B: No... mom said that we'll have the wedding at home on the 2nd. So let's hope your mom is right about that being a lucky day...
A: It will be. And today is, too. I just talked to the airlines...
B: They're going to give us back the seats?
A: No, they're giving us discount tickets on a honeymoon package in...
B:

Response Set A

Response #1: The Maldives?
Response #2: Japan! That's amazing, I've always dreamed of going there. Thank you so much for taking care of everything. You're the best partner ever!
Response #3: Fiji? That sounds awesome, thank you for checking with the airlines!
Response #4: I know! It's to the Maldives.
Response #5: The Maldives!

.....
:Semantically unique response count: 3 :
:Lexically unique response count: 4 :
:Contextually coherent response count: 5 :
.....

Response Set B

Response #1: Are we heading to Fiji? I love the Caribbean.
Response #2: Japan! That's amazing, I've always dreamed of going there. Thank you so much for taking care of everything. You're the best partner ever!
Response #3: Fiji? That sounds awesome, thank you for checking with the airlines!
Response #4: Fiji? That sounds great, thanks for checking with the airlines!
Response #5: Fiji? That sounds fantastic, thanks for confirming with the airlines!

.....
:Semantically unique response count: 2 :
:Lexically unique response count: 3 :
:Contextually coherent response count: 5 :
.....

Figure 6: Samples of human evaluation used for MRG.

Dialogue Context

A: What dressing would you like on the salad? B: French dressing, please A: I beg your pardon? B: Oh, French dressing. A: Sorry, it's not available now. Anything else?

TinyLlama

FS 3-shot: 1.Can I have some olive oil instead? I'm trying to keep my calories low. 2.Fine, go ahead and use olive oil. It's not a big deal. 3.In that case, I'll have the regular dressing. 4.Sorry, I'm not sure I can find something with those specific ingredients. 5.That's too bad, I'd love some French dressing on my salad.

CoT 3-shot: 1.Hmm, how about ranch dressing? Do you have it? 2.In that case, I think I'll go for the ranch dressing? Do you have it? 3.I'm not really sure what dressing to choose. What are the other available options? 4.Hmm, I'm not sure what dressing to choose. Could you recommend me something that will enhance the flavour of the salad? 5.How's that possible? In that case, can you suggest something that would enhance the salad? **Explanation:** Each response differs from the others by the choice of dressing. The first response suggests that the customer is requesting ranch dressing, but the second response suggests that the customer is not seeing ranch dressing on the menu. The third and fourth response indicates that the customer is not aware of any other dressing options available on the menu. The fifth response conveys a sense of confusion.

PC 3-shot: 1.That's alright, can I have ranch instead? 2. That's unfortunate, I was looking forward to it. What other dressings do you have? 3. Ah, let me see, I think I'll have French or Italian dressing then. 4. What other options are there? I prefer a light dressing. 5. **That's disappointing, I really wanted French dressing. What other types of dressing do you have?**

Instruction Tuned: 1.Oh, I see. Then maybe something with a bit of a tangy flavor? 2.Could you suggest a dressing that has a slightly sweet or spicy flavor? 3.Alright, I'll just go with olive oil and salt. 4.A homemade dressing with a touch of balsamic vinegar and lemon juice would be perfect for this salad. 5.How about a classic vinaigrette dressing?

Multiple Inference: 1.I'll have the regular dressing then. 2.I'll have the regular dressing, please. 3.I'll stick with the regular dressing then. 4.I'll have the regular dressing. 5.I'll have the regular dressing then.

Llama2-7b

FS 3-shot: 1.No problem, how about Italian dressing instead? 2.That's too bad. Can I have ranch dressing then? 3.I see. Would balsamic vinaigrette be a good alternative? 4.Thanks for letting me know. How about Blue Cheese dressing? 5.I appreciate your help. Can you recommend another dressing that's similar to French?

CoT 3-shot: 1.I'll just have the house dressing instead. 2.Oh, that's okay. What about balsamic vinaigrette? 3.No problem, how about a citrus vinaigrette? That sounds refreshing. 4.I see, well can I get the ranch dressing then? 5.No worries, I'll just use olive oil and vinegar. **Explanation:** Each response differs from the others in terms of the type of dressing requested by B. Response 1 suggests a different dressing option altogether, while responses 2 and 3 offer alternative choices to French dressing. Response 4 requests a different dressing that is not available, and response 5 offers a homemade alternative to French dressing.

PC 3-shot: 1.Oh, in that case, I'll go with ranch dressing then. 2.Hmm, do you have balsamic vinaigrette? That could work really well! 3.Well, I guess I'll just take olive oil and vinegar, if that's okay. 4. Alright, I'll go for Italian dressing instead. 5. **Oh no worries, how about blue cheese dressing then? That'll complement the salad well.**

Instruction Tuned: 1.Hmm, in that case, I'll take Italian dressing instead. 2.Oh, no problem. I'll have the balsamic vinaigrette then. That would be delicious as well. 3.I see, how about the ranch dressing? I think it would enhance the flavour of the greens in the salad. 4.Alright, then I'll go with the honey mustard dressing. Thank you! 5.Oh it's okay! No worries, I'll just have it plain.

Multiple Inference: 1.Oh, no worries! How about a vinaigrette instead? I find it pairs really well with the fresh greens and vegetables in the salad. 2.Oh, no worries! How about a vinaigrette? I find it pairs quite well with the fresh greens. 3.Oh, no problem! I'll just have the vinaigrette instead. Thank you! 4.Thank you for letting me know. How about we try a vinaigrette dressing instead? 5.Oh, I see! No worries! In that case, I would love a vinaigrette dressing on my salad. Perhaps a light balsamic vinaigrette would be perfect?

Llama2-13b

FS 3-shot: 1.Oh, that's okay. I'm not too picky. What else do you have? 2.That's a bummer. Do you have any other dressings that are similar to French? 3.Sure, I'll just go with something else then. 4.Is there a reason why French dressing is not available? 5.Can I get a different type of salad instead?

CoT 3-shot: 1.That's okay, I'll just have the house dressing then. 2.Oh no, I was really looking forward to the French dressing. Can you recommend something else? 3.No problem, I'm open to trying something different. What do you suggest? 4.Is there a specific reason why French dressing is not available? 5.Thanks for letting me know. Maybe next time. I'll try the house dressing instead. **Explanation:** Each response differs from the others in terms of how they react to the unavailability of French dressing. Response 1 accepts the situation and chooses a different option, while response 2 expresses disappointment and asks for a recommendation. Response 3 is open to trying something new, while response 4 asks for an explanation for the unavailability of French dressing. Lastly, response 5 politely acknowledges the situation and expresses hope for trying French dressing in the future.

PC 3-shot: 1.In that case, how about something with a little kick—maybe a spicy ranch or chipotle dressing? I love a bit of heat with my salad. 2. Ah, I see. Well, in that case, could I get a honey mustard dressing? 3.Oh, I see. No problem! Could you recommend something light then? Maybe a Italian or dressing? 4.**Got it! In that case, would a lemon vinaigrette work? Something light and refreshing could really complement the salad.** 5.Hmm, okay, no French dressing. How about a nice balsamic vinaigrette?

Instruction Tuned: 1.That's unfortunate. Can I have ranch dressing instead? 2.Sure, how about Italian dressing? 3.In that case, can I just have oil and vinegar? 4.Is there a balsamic vinaigrette option? 5.Okay, I'll just have the salad without any dressing.

Multiple Inference: 1.Hmm, that's a bit of a bummer. How about Italian dressing? I absolutely love the vinaigrette flavor. 2.Oh no, French dressing is not available? Can I have the vinaigrette instead? Or perhaps a light Italian dressing? 3.Oh, that's okay! I'm actually open to trying something different. Do you have any other dressing options that you would recommend? Perhaps a light vinaigrette or a creamy ranch? I'm looking for something that will complement the flavors of the salad without overpowering it. 4.Oh, that's okay! I actually prefer vinaigrettes anyway. Do you have any good options in that department? Maybe a Balsamic or Italian vinaigrette? 5.Oh no, French dressing is not available? In that case, I would love to try the Balsamic Vinaigrette dressing instead. Thank you for letting me know!

Llama2-70b: Oh, that's too bad. Do you have any ranch dressing?

gpt-3.5-turbo: I'll have the Italian dressing then, thank you.

Figure 7: Samples of response sets generated by TinyLlama, Llama2-7b and Llama2-13b. The responses in red was selected by the ODRP_{HN} model during PS.