



HAL
open science

Role-Play Zero-Shot Prompting with Large Language Models for Open-Domain Human-Machine Conversation

Ahmed Njifenjou, Virgile Sucal, Bassam Jabaian, Fabrice Lefèvre

► **To cite this version:**

Ahmed Njifenjou, Virgile Sucal, Bassam Jabaian, Fabrice Lefèvre. Role-Play Zero-Shot Prompting with Large Language Models for Open-Domain Human-Machine Conversation. 2023. hal-04630836

HAL Id: hal-04630836

<https://hal.science/hal-04630836>

Preprint submitted on 1 Jul 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Role-Play Zero-Shot Prompting with Large Language Models for Open-Domain Human-Machine Conversation

Ahmed Njifenjou and Virgile Sucas and Bassam Jabaian and Fabrice Lefèvre

Laboratoire Inforamitique d'Avignon (LIA), Avignon Université

{ahmed-ndouop.njifenjou & firstname.lastname}@univ-avignon.fr

Abstract

Recently, various methods have been proposed to create open-domain conversational agents with Large Language Models (LLMs). These models are able to answer user queries, but in a one-way Q&A format rather than a true conversation. Fine-tuning on particular datasets is the usual way to modify their style to increase conversational ability, but this is expensive and usually only available in a few languages. In this study, we explore role-play zero-shot prompting as an efficient and cost-effective solution for open-domain conversation, using **capable** multilingual LLMs (Beeching et al., 2023) trained to obey instructions. We design a prompting system that, when combined with an instruction-following model - here Vicuna (Chiang et al., 2023) - produces conversational agents that match and even surpass fine-tuned models in human evaluation in French in two different tasks.

1 Introduction

Since their introduction, Transformers (Vaswani et al., 2017) and derivative Large Languages Models (LLMs) have constantly improved the state-of-the-art on several Natural Language Processing (NLP) tasks. Among them, is open-domain dialogue which consists in a conversational agent designed to interact socially with users on any topic while displaying human abilities, like empathy, personality or entertainment (Walker et al., 2021). Common approaches using LLMs mostly relied on fine-tuning with specific datasets often targeting one or more of these skills (e.g. PersonaChat (?), Blended Skill Talk (Smith et al., 2020), Empathetic Dialogues (Rashkin et al., 2019) *inter alia*). These datasets are expensive to build and often available only in one language. In addition recent developments in the field of NLP have focused on LLMs trained to follow instructions (Ouyang et al., 2022). They have the native ability to respond to users' inputs in a natural language manner. However,

the leader-to-follower single-sided relationship is strongly present and they hardly display human conversational abilities straightaway.

In this work, as an attempt to deal with these issues while escaping from fine-tuning and its data dependency, we propose to use role-play via zero-shot prompting to leverage instruction-following models' abilities. This approach pertains to the newly devised prompt-based learning (PBL) paradigm (Liu et al., 2021). We assess this approach on two tasks, a general Persona task based on the PersonaChat dataset to bring these models to have a persona while displaying human skills and a particular case, the INT task (Riou et al., 2019) where the speakers have to discuss an image, simulating a situated multi-modal conversation.

2 Related Work

Open-domain dialogue has seen a lot of developments with most solutions focusing mainly on fine-tuning with specific collected data. Among others we have the BlenderBot's (Roller et al., 2021) series and other closed-sourced models as Meena (Adiwardana et al., 2020), LaMDA (Thopilan et al., 2022), etc. These models display great conversational skills, but in addition to data dependency, they are often only available in English. Foundational models are a recent trend in the field of NLP as they display multilingual abilities and state-of-the-art performance on several benchmarks. One of them, LLaMA (Touvron et al., 2023a,b) is the backbone LLM of this study. These series of models are trained on an amount of data surpassing the scaling law of (Hoffmann et al., 2022). Such capable LLMs are a prerequisite to build instruction following models. From LLaMA resulted, among others: Alpaca (Taori et al., 2023), StackLLaMA (Beeching et al., 2023), Guanaco (Dettmers et al., 2023) and the main model of this study, Vicuna (Chiang et al., 2023).

The latter has been fine-tuned on the ShareGPT¹ corpus which is basically a dataset of conversations produced by ChatGPT (OpenAI, 2022) users. They are collected using a web browser plugin installed by users willing to participate in the dataset collection. The resulting models have open-domain **responding** abilities – which we dissociate from **conversational** abilities – and implicitly OpenAI original model’s restrictions, resulting from the ChatGPT alignment process, as described in the OpenAI documentation (Leike et al., 2022).

The idea behind instruction following models lays within a new paradigm in NLP, coined as Prompt-Based Learning (PBL) (Liu et al., 2021). While using such models to perform prediction tasks, their inputs are modified following a global template into a textual string prompt that has some unfilled slots, and then the language model is used to fill the unfilled information to obtain a final string, from which the final output can be derived, in a generative way. One major asset of the method, which makes it very powerful, is that it allows the LLM to be pre-trained on huge quantity of text and, by defining a new prompt scheme, the model can perform few-shot — or even zero-shot — learning, and adapt to new scenarii with few or no labelled data.

(Santu and Feng, 2023) proposed a general prompt taxonomy to unify complex tasks benchmarking which shares the idea of prompt structuring with our work. However, with role-play prompting, we specifically target dialogue which requires even more complex abilities. The role-play prompt scheme proposed in this GitHub repository² intended for an LLM-based Discord bot is closer to our needs. However, it is designed for providing only persona-based information. Role-play prompting as we want to demonstrate is not just limited to character or persona level instructions. Role-Play can be used in order to enforce other conversational skills, such as empathy or engagingness, which help balance the dialogue between the user and the bot. We applied this approach to two scenarii, without fine-tuning.

¹<https://sharegpt.com>

²<https://github.com/teknium1/alpaca-roleplay-discordbot>

3 Methodology

3.1 Instruction-Following vs Dialogue Skills

Quoting (Reitz, 2019), “dialogue is more than just having a conversation. Genuine dialogue describes a way of interacting that is mutual, relational, attentive, and meaningful”. Instruction-following models, even those optimized for conversation fall short from fulfilling some of these aspects of a genuine dialogue. Indeed, social aspects - mutuality, relationality and attention – especially are poorly displayed. Talking of mutuality, these systems often converse in a leader-follower structure where they are the follower and the user the leader – this is for instance evidenced by the words used to designate user entries in ChatGPT release blog: “queries”, “instructions” (OpenAI, 2022). Regarding relationality and attention, these systems lack of straightforward engagingness and personality consistency.

Commonly, these limitations have been tackled with finetuning. However, this is costly and data-dependant which data are scarce in languages other than English. As results, we formalize a general role-play prompt structure which is a more efficient and less expensive approach. Indeed, it avoids finetuning and rely on **multilingual LLMs** prompted in **English**, with external **English data** if needed (for instance persona from PersonaChat Dataset) and performs dialogue task in a **desired target language** here in **French**.

3.2 Role-Play Prompting

It is important to understand that Role-Play Prompting here is not restricted to playing a given character. To better understand that, let’s consider the simulacra and simulator framing in (Shanahan et al., 2023). LLM is a simulator which swallowed myriads of simulacra during the pre-training and at each simulation it kind of randomly selects among any of them which oneS to display. The simulacrum is, as a matter of fact, not only about persona background (Smith et al., 2020) but also *thinking* and *writing* styles, personal situation (Smith et al., 2020), target language (for a multilingual simulator) and information processing (long-term memory, user personalization, response filtering etc.). All possible simulacra already exist in the simulator (LLM) but it doesn’t display all of them natively, in fact it can’t. Role-Prompting enters the chat to make the LLM favor simulacra that are suitable for a given dialogue task. In this paper, we

derived two distinct dialogue tasks to assess the efficiency of this approach. The first task, referred as Persona task, uses Role-Play Prompting for enhancing humanness in conversation skills. The second one, referred as INT task, derives Role-Play to allow the LLM to talk about a simulacrum instead of interpreting it.

3.3 Prompt Structure

Open-domain dialogue belonging to the realm of complex tasks (Santu and Feng, 2023) makes the endeavour of role-play prompting more challenging. As a matter of fact, small variations in a prompt may hamper the model’s observed performance. For this reason, it is mandatory to define a general prompt structure that can be adapted later to different conversational tasks. Hence, with the concerns of dealing with the previously mentioned limitations we retain the following sections – each focusing on different aspects useful to a dialogue – as building blocks of a prompt engineering module:

1. **System Instructions** $\mathcal{I}_s = \{i_{s,k}\}_{k=1}^{N_i}$: where N_i is the number of instructions $i_{s,k}$ which sharply define the target task’s specifications and the global desired behaviour. This may include *thinking* (inferring) and *writing* (generation) styles.
2. **Situational context** $\mathcal{C}^t = \{c_k\}_{k=1}^{N_c}$: each c_k is a context information that may help the model better perform the desired task. As such, it evolves with time depending on the conversation flow. It can include personality information, image and scene description, summaries

of old turns or information from external sub-modules.

3. **Response Instructions** $\mathcal{I}_a = \{i_{a,k}\}_{k=1}^{N_i}$: These are final instructions to incite the LLM to respond to the users’ utterance with emphasis on the writing style, the target responding language and creativity with section 2) in mind.
4. **Conversation History** \mathcal{X}^t : this part contains previous messages from the user (x^t) and the LLM (y^t). They can be truncated to the k latest conversation turns to fit in the LLM token size limit or help the LLM focus on latest part of the conversation. In this case a summary of the k removed turns \mathcal{X}^k generated by an external module (also using PBL with an LLM) can be added in section 2). Conversation history, therefore becomes: $\mathcal{X}^{t-1} = \{(x^{t-k}, y^{t-k}), \dots, (x^{t-1}, y^{t-1})\}$

These sections can be further precised into subsections and their order is set to vary as it may be suitable to give more or less importance to one section than another for the final model’s response depending on the task at hand. This will be showcased in our two experimented tasks later on. Finally, the prompt builder returns: $\mathcal{P}_{task}^t = \sigma_{task}(\mathcal{I}_s, \mathcal{C}^t, \mathcal{I}_a, \mathcal{X}^{t-1})$ where σ_{task} is the most suitable permutation for the dialogue task at hand. Therefore at each turn, the model maximizes the following probability: $p(y^t|x^t, \mathcal{P}_{task}^t)$ to respond.

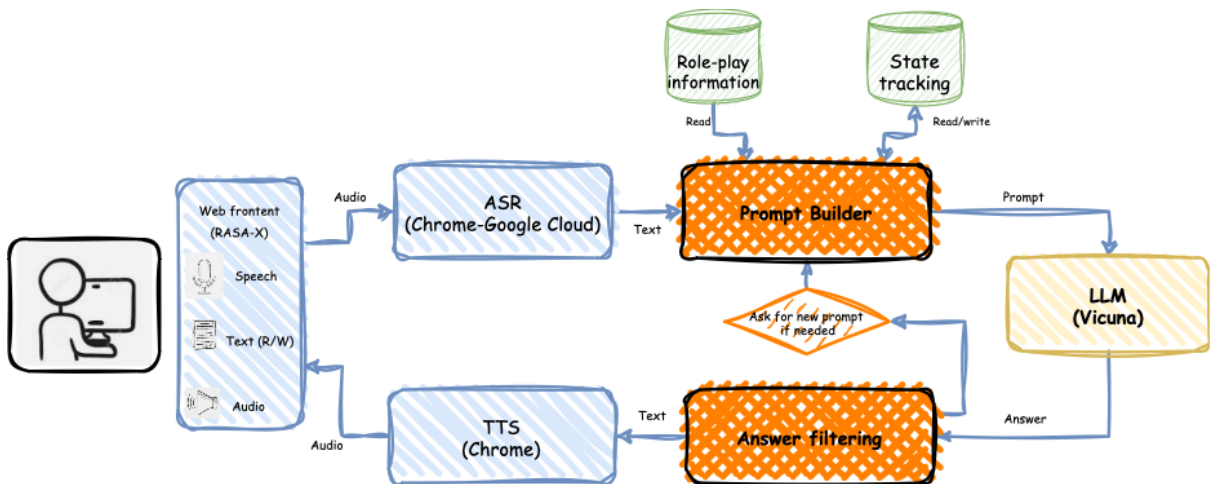


Figure 1: Architecture of the Role-play Zero-shot prompted dialogue system.

4 Experiments

All experiments are carried out in **French** but the prompt contains instructions mainly in **English**, one of which specifies the target response language. Given the results obtained by (Lai et al., 2023) and those we present in Appendix A, we assume that this can be applied to other languages in which the assessed model performs comparably, as in French.

4.1 System Architecture

The system shown in Figure 1 is a pipeline of several modules. These include a web interface based on the Rasa X (Rasa, 2022) tool, modified to integrate voice functionality using the Google Chrome Speech-to-text and Text-to-speech APIs. This enables users to exchange easily with the agent, either by voice (recommended) or by text. Next, a module that constructs the prompt according to the general structure described in 3.2 from the user’s (textual) message and information both external and internal to the conversation. After generation, if the LLM responses are not valid, a filtering module is used to apply corrections before sending them to the user (cf. Annex C).

4.2 Open-domain Conversation With Human Capabilities: the PersonaChat Task

This task involves enhancing LLM’s conversational capabilities by using roles built from personality traits drawn from the **PersonaChat dataset** (Zhang et al., 2018) embedded in \mathcal{C}^t as external information (see conversation example in Appendix D.1).

4.2.1 Shallow Prompt

Given in Appendix E.3, it is close to Vicuna’s basic prompt (Appendix E.1). However, there are some additions for the sake of fair comparisons: system instructions to describe the task, contextual information (notably personality traits), and the instruction to complete the conversation history.

4.2.2 Advanced Prompt

It exactly follows the structure in Section 3.2 i.e. $\sigma_{task} = \mathbf{Id}_4$ (see appendix E.3). The context \mathcal{C}^t includes the **specification of humanity** where personality traits are added with the injunction to choose a name consistent with them if necessary. External **modules** can augment this under certain conditions, as in (Shuster et al., 2022). Indeed, to prevent prompts from exceeding the LLM’s maximum context size while retaining the content of

the entire conversation, old exchanges are summarized in a few sentences. To personalize the user experience, a line containing user-specific information is added and updated regularly. This aspect is generally referred to as long-term memory. These "modules" are actually prompts sent to an LLM. The **history \mathcal{X}^t of the conversation** is kept at the end of the Advanced Prompt so that the system, when generating a response, has an overview of the entire conversation.

4.3 Simulating Multi-modal Conversations: the INT Task

If a model is capable of interpreting a role in order to embody a character, we can assume that it is also capable of talking about a role without interpreting it. To test the validity of this hypothesis, we propose a prompt designed to enable the LLM to converse about a specific topic. This topic is defined in the form of a role that the LLM will be encouraged to describe rather than interpret.

4.3.1 Task definition

The dialogue system is intended to conduct multimodal conversations set in the context of a neuroscience experiment (Chaminade, 2017). Inside an fMRI scanner, a person must converse with a Furhat³ robotic head, which is alternately connected (unbeknownst to him) to a dialogue system or to a human. The interlocutors talk about an image presented to them. Their conversation is motivated by the common goal of finding the image’s promotional message (see conversation examples in Appendix D.2).

4.3.2 Prompt

The prompt is designed following the structure described in Section 3.2 with $\sigma_{task} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{pmatrix}$ i.e., instructions related to the **response** (\mathcal{I}_r) and the **context** (\mathcal{C}^t) are placed after the **history** (\mathcal{X}^t). This permutation groups all the instructions dedicated to the task at the end of the prompt, i.e. just before the last user **message** (x^t). This actually corresponds to the most common pattern in Vicuna’s fine-tuning corpus — *ShareGPT* — which includes exchanges where users can only communicate their instructions to the LLM (*ChatGPT*) inside their messages (via the web interface), those around the last message being the most important.

³<https://furhatrobotics.com>

This structure allows the model to focus on the image (linked to the goal) rather than the conversation’s history \mathcal{X}^t , unlike the previous task where \mathcal{X}^t was paramount. This is also why C^t , in addition to external information (in this case the image description), includes **general instructions** summarizing the task at hand.

5 Human Evaluation

For a given dialogue input, several responses may be correct. For this reason, human evaluation remains more reliable than automated references-based evaluations. Therefore, it was performed for all the considered evaluation sets and for both tasks. Three criteria, based on those mentioned in (Mehri and Eskenazi, 2020; Roller et al., 2021; Ji et al., 2022), were selected on which each conversation was rated on a 1-5 scale by three different evaluators: (1) **coherence**, the ability for the system to propose responses that are consistent with the conversation history ; (2) **engagingness**, the ability to revive conversation by providing messages that require responses ; (3) **humanness**, the ability to respond as a human being would do. An additional specific criterion for the INT task is added: (4) **achievement**, the validation of the speakers success in achieving their goal (cf. Section 4.3). Users (resp. evaluators) were never aware of the identity of the system they were interacting with (resp. evaluating).

Furthermore, to assess the effectiveness of the proposed method on LLMs, it is essential to compare the performance with similar approaches but also different models (different sizes, training data and with and without instruction settings). For this sake, we added the Few-shot Bot (FSB) prompt proposed by (Madotto et al., 2021) which consist in providing only demonstration examples to a non-instruction tuned LLMs. For the latter, we selected: Vicuna 7B, 13B and 33B (Zheng et al., 2023), Guanaco-13B (Dettmers et al., 2023) and LLaMA-13B (Touvron et al., 2023a). We also carried out statistical studies on the responses generated. The results obtained and their analysis are reported in Section 6.1.

5.1 Self-Chats Evaluation

Collecting human-bot conversations is expensive. For this reason, we generated conversations between two instances, each model + prompt combination (self-chats). Their performance is evalu-

ated in a *Chatbot Arena* style (Zheng et al., 2023). Evaluators compared two self-chats from different setups on each criterion and in general. The scores presented in Table 1 are Elo scores (Elo, 1978) calculated from the comparisons’ results. A total of 18 annotators evaluated 982 generated conversations of 10 rounds each (which correspond to around 70 dialogues per configuration and 5 to 14 battles per pair).

LLM	Prompts	Overall	Criteria			Rank
			(1)	(2)	(3)	
LLaMA-13B	FSB*	918	916	918	919	9
	Shallow	902	900	912	910	10
	Advanced	891	895	897	883	11
Guanaco-13B	Shallow	987	991	984	990	8
	Advanced	1022	1020	1014	1022	7
Vicuna-7B	Shallow	1033	1037	1020	1033	6
	Advanced	1035	1025	1063	1034	5
Vicuna-13B	Shallow	1050	1065	1051	1046	2
	Advanced	1041	1040	1046	1041	4
Vicuna-33B	Shallow	1042	1048	1023	1043	3
	Advanced	1074	1058	1068	1073	1

(1), (2), (3) as previously defined in Section 5

Table 1: ELO scores from self-chats model-to-model battles with human annotations.

We can see that larger model sizes and instruction-tuning lead to better performance. On the one hand, *Vicuna* tops the chart, followed by *Guanaco*. On the other hand, *LLaMA* underperforms with the proposed prompt and the *FSB* prompt. As the *Vicuna-33B* + Advanced Prompt combination ranked first was too costly (latency, resources) for the collection of human-model conversations, the *Vicuna-13B* + Advanced Prompt combination was selected for collection and the next round of evaluations⁴.

5.2 Human-bot chats evaluation

5.2.1 PersonaChat Task

We collected 103 conversations from 11 users instructed to exchange with models via the web interface (cf. 4.1). After removing invalid conversations, 72 were retained for evaluation. Conversations were also conducted with *BlenderBot 1 (BB1)* (Roller et al., 2021), a state-of-the-art fine-tuned system, for comparison. Each sample is evaluated by three (out of a total of 12) different annotators for each criterion, and the median is used as the sample score (results in Table 2).

The Advanced Prompt scores highest for coherence. We assume that this is mainly due to

⁴Won 75 % of direct comparisons with *Vicuna* + *Shallow* despite being ranked behind overall.

Strategy	Config.	Criteria		
		(1)	(2)	(3)
Finetuning	BB1	2.95	3.85	3.00
Roleplay	Shallow	3.78	3.30	2.74
Prompting	Advanced	3.88	3.72	3.40

Table 2: Average score per criterion for PersonaChat

Vicuna-13B’s intrinsic emergent abilities, as this results is close to those of the Shallow Prompt (-0.1). As far as humanness is concerned, while the Advanced Prompt has the best score, the Shallow has the worst. This highlights the impact of structured role-playing instructions in the Advanced Prompt. Finally, for engagingness, BB1 still set the pace. Although it has been fine-tuned on a specific dataset that allows it to ask and answer personal questions (Smith et al., 2020) which is important in the rating of this criterion (as presented in the Appendix B), it is closely followed by Advanced Prompt (-0.13).

5.2.2 INT Task

The evaluation was performed on 27 conversations carried out by 4 users. The conditions were identical to those for the PersonaChat task, except that the testers (both users and evaluators) also observed an image linked to the conversation (c.f. Section 4.3). Our system, "Vicuna & Advanced Prompt", is compared to an earlier system designed for the same task called Lilia (Riou et al., 2019) and to human beings participating in a Wizard of Oz-type experiment (WoZ). 10 raters evaluated 8 to 9 conversations for each system. The results are presented in Table 3.

Strategy	Config.	Criteria			
		(1)	(2)	(3)	(4)
Reference	WoZ	3.46	3.21	3.36	2.96
	Lilia	3.57	3.43	2.64	3.75
Roleplay Prompting	Vicuna & Advanced	3.93	4.25	3.29	4.00

Table 3: Average score per criterion for INT

The Vicuna & Advanced Prompt system scores best on all criteria, with the exception of humanness, where WoZ received the highest score. This was to be expected, given that in this experiment the agent was a human being. However, it was less predictable that this was not also the case for the other criteria. As explained in Section 4, in the

experiment, the interlocutors must find the promotional objective of the image. A closer look at the conversations in WoZ reveals that the human agent was less goal-oriented than the artificial agents, which may explain the success result.

For engagingness, the Vicuna-based model is encouraged to ask questions to revive the conversation. This fact may explain the effectiveness of this system in staying engaged throughout the conversation. There are far fewer questions in conversations produced with other systems. As previously stated for PersonaChat task, we assume that the high coherence scores are mainly due to the effectiveness of Vicuna.

6 Statistical Analysis of Collected Conversations

6.1 Quality Metrics

Statistical data was computed for all conversations. Among these, the vocabulary size, which is the number of different lemmas in each message and in the conversation as a whole. All messages were lemmatized with the Spacy library’s morphosyntactic labeling tool (Honnibal et al., 2020) for French⁵. The number of words per message for each speaker type is also reported to give another view of system performance.

Strategy	Configuration	Vocabulary size		
		Agent	User	Gap
Finetuning	BB1	772	687	85
Roleplay Prompting	Vicuna & Shallow	2569	965	1604
	Vicuna & Advanced	2239	1043	1196

(a) PersonaChat

Strategy	Configuration	Vocabulary size		
		Agent	User	Conv.
Reference	WoZ	346	390	537
	Lilia	254	197	330
Roleplay Prompting	Vicuna & Advanced	746	365	864

(b) INT

Table 4: Vocabulary size per speaker type

⁵fr_core_news_sm available at https://spacy.io/models/fr?_x_tr_hist=true#fr_core_news_sm

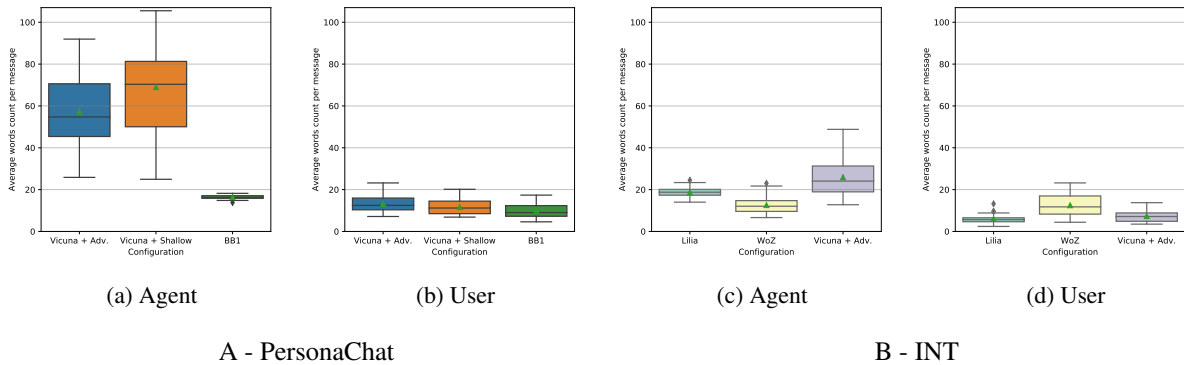


Figure 2: Average words count per message

6.1.1 PersonaChat Task

The agent statistics in Table 4a and Figure 2a highlight a major flaw for instructions-following models: extreme verbosity.

The vocabulary size per speaker type (shown in Table 4a) of the collected conversations also gives an idea of the performance of the proposed method. We note that BB1 has the least vocabulary, which may work to the detriment of the coherence score. However, its agent and user vocabulary sizes are more balanced, resulting in a better engagement score. In contrast, the gap between these two measures for LLM + prompts approaches is much wider. It is nevertheless reduced with Advanced Prompt, which may be the reason for the slight improvement in the engagement score.

6.1.2 INT Task

As for PersonaChat task, Vicuna’s verbosity measure is present in Figure 2c and Table 4b. On the other hand, as the Lilia system’s responses were built from pre-established expert models, its vocabulary is limited.

In WoZ conversations, the average vocabulary size is almost identical between users and the agent. Users also have a wider range of message sizes when chatting with a human. Thus, in these conversations, agent messages seem to call for more diverse responses and neither interlocutor seems to be directing the dialogue. Similarly, this gap is also higher for conversations with Vicuna than for those with Lilia.

Table 4b shows a higher vocabulary size for the user in WoZ conversations. This is even greater than that of the agent. This may be an indicator of greater variety in user responses. Here, the vocabulary size of users interacting with the Vicuna-based

system is quite comparable to what can be observed in WoZ.

6.2 Filtered Errors Analysis

Error occurrences in response generation have been evaluated on 100 self-chats for each setup. Each of these conversations contains 10 turns. Results for both tasks are reported in Table 5. All rates are computed over all turns. Details on errors are given in Appendix C.

Persona Chat		Regex	Language	Incomplete / Empty
	Shallow		0.007	0.000
Advanced		0.002	0.001	0.148
INT		Empty	Too Long	Total
	Detected	0.080	0.150	0.230
	Fixed	0.080	0.040	0.120

Table 5: Occurrence rate of error types for both tasks.

Detection rates were calculated for the two prompts in the PersonaChat task. As several errors can occur for the same message in this task, the totals have not been calculated. Additionally, all detected errors are corrected.

On the other hand, errors in the INT task are corrected only if the proposed corrections comply with the response filtering rules (see Appendix C.2). For both types of error, we have calculated separate rates for detected and corrected errors. As these errors cannot occur in the same message, a total rate has been calculated.

7 Conclusion

This paper explored the use of structured role-play prompt engineering to improve open-domain human-machine conversations with LLMs. Role-play prompting is a simple and inexpensive method

of upgrading the behavior of language models to make them conversational agents. It has been applied here in French, but it can be adapted to other languages by orienting the role that way. Experiments in two different tasks, persona-based task and simulated multimodal dialogues, have shown that, although language models still have significant shortcomings, such as hallucinations, users' perception of these agents can be comparable to that of higher-cost finetuned models.

In addition, our experiments have shown us that prompt engineering needs to be further improved by automating the building and filtering processes. As it is, not only too many factors implied rely on the designer's expertise, but even more they are set once and for all when they could also evolve with the situation during the course of the dialogue. As a perspective, we propose to upgrade the model proposed here with a full reinforcement learning setup so as to automatically derive the prompt-making actions.

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](#).
- Edward Beeching, Younes Belkada, Kashif Rasul, Lewis Tunstall, Leandro von Werra, Nazneen Rajani, and Nathan Lambert. 2023. [Stackllama: An rl fine-tuned llama model for stack exchange question and answering](#).
- Thierry Chaminade. 2017. [An experimental approach to study the physiology of natural social interactions](#). *Interaction Studies*, 18(2):254–276.
- 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](#).
- Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. 2023. [Qlora: Efficient finetuning of quantized llms](#).
- Arpad E. Elo. 1978. *The rating of chessplayers, past and present*. New York : Arco Pub.
- Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, Tom Hennigan, Eric Noland, Katie Millican, George van den Driessche, Bogdan Damoc, Aurelia Guy, Simon Osindero, Karen Simonyan, Erich Elsen, Jack W. Rae, Oriol Vinyals, and Laurent Sifre. 2022. [Training compute-optimal large language models](#).
- Matthew Honnibal, Ines Montani, Sofie Van Landeghem, and Adriane Boyd. 2020. [spaCy: Industrial-strength Natural Language Processing in Python](#).
- Tianbo Ji, Yvette Graham, Gareth J. F. Jones, Chenyang Lyu, and Qun Liu. 2022. [Achieving reliable human assessment of open-domain dialogue systems](#).
- Viet Lai, Chien Nguyen, Nghia Ngo, Thuat Nguyen, Franck Dernoncourt, Ryan Rossi, and Thien Nguyen. 2023. [Okapi: Instruction-tuned large language models in multiple languages with reinforcement learning from human feedback](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 318–327, Singapore. Association for Computational Linguistics.
- Jan Leike, John Schulman, and Jeffrey Wu. [Our approach to alignment research](#) [online]. 2022.
- 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, Online. Association for Computational Linguistics.
- Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2021. [Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing](#). *CoRR*, abs/2107.13586.
- Andrea Madotto, Zhaojiang Lin, Genta Indra Winata, and Pascale Fung. 2021. [Few-shot bot: Prompt-based learning for dialogue systems](#).
- Shikib Mehri and Maxine Eskenazi. 2020. [Unsupervised evaluation of interactive dialog with dialogpt](#).
- OpenAI. [Introducing chatgpt](#) [online]. 2022.
- 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](#).
- Rasa. 2022. [Rasa X Community Documentation](#). <https://legacy-docs-rasa-x.rasa.com/docs/rasa-x/0.32.x/>.
- Hannah Rashkin, Eric Michael Smith, Margaret Li, and Y-Lan Boureau. 2019. [Towards empathetic open-domain conversation models: a new benchmark and dataset](#). In *ACL*.

- Megan Reitz. Organisational dialogue and leadership [online]. 2019. <https://www.meganreitz.com/dialogue-and-leadership>.
- Matthieu Riou, Bassam Jabaian, Stéphane Huet, and Fabrice Lefèvre. 2019. Lilia, a showcase for fast bootstrap of conversation-like dialogues based on a goal-oriented system. In *7th International Conference on Statistical Language and Speech Processing (SLSP)*, pages 31–43, Cham. Springer International Publishing.
- Stephen Roller, Emily Dinan, Naman Goyal, Da Ju, Mary Williamson, Yinhan Liu, Jing Xu, Myle Ott, Eric Michael Smith, Y-Lan Boureau, and Jason Weston. 2021. [Recipes for building an open-domain chatbot](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 300–325, Online. Association for Computational Linguistics.
- Shubhra Kanti Karmaker Santu and Dongji Feng. 2023. [Teler: A general taxonomy of llm prompts for benchmarking complex tasks](#).
- Murray Shanahan, Kyle McDonell, and Laria Reynolds. 2023. [Role-play with large language models](#).
- Kurt Shuster, Jing Xu, Mojtaba Komeili, Da Ju, Eric Michael Smith, Stephen Roller, Megan Ung, Moya Chen, Kushal Arora, Joshua Lane, Morteza Behrooz, W.K.F. Ngan, Spencer Poff, Naman Goyal, Arthur D. Szlam, Y-Lan Boureau, Melanie Kambadur, and Jason Weston. 2022. [Blenderbot 3: a deployed conversational agent that continually learns to responsibly engage](#). *ArXiv*, abs/2208.03188.
- Eric Michael Smith, Mary Williamson, Kurt Shuster, Jason Weston, and Y-Lan Boureau. 2020. [Can you put it all together: Evaluating conversational agents' ability to blend skills](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2021–2030, Online. Association for Computational Linguistics.
- Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. 2023. [Stanford alpaca: An instruction-following llama model](#). https://github.com/tatsu-lab/stanford_alpaca.
- Romal Thoppilan, Daniel De Freitas, Jamie Hall, Noam Shazeer, Apoorv Kulshreshtha, Heng-Tze Cheng, Alicia Jin, Taylor Bos, Leslie Baker, Yu Du, YaGuang Li, Hongrae Lee, Huaixiu Steven Zheng, Amin Ghafouri, Marcelo Menegali, Yanping Huang, Maxim Krikun, Dmitry Lepikhin, James Qin, Dehao Chen, Yuanzhong Xu, Zhifeng Chen, Adam Roberts, Maarten Bosma, Vincent Zhao, Yanqi Zhou, Chung-Ching Chang, Igor Krivokon, Will Rusch, Marc Pickett, Pranesh Srinivasan, Laichee Man, Kathleen Meier-Hellstern, Meredith Ringel Morris, Tulsee Doshi, Renelito Delos Santos, Toju Duke, Johnny Soraker, Ben Zevenbergen, Vinodkumar Prabhakaran, Mark Diaz, Ben Hutchinson, Kristen Olson, Alejandra Molina, Erin Hoffman-John, Josh Lee, Lora Aroyo, Ravi Rajakumar, Alena Butryna, Matthew Lamm, Viktoriya Kuzmina, Joe Fenton, Aaron Cohen, Rachel Bernstein, Ray Kurzweil, Blaise Agueras-Arcas, Claire Cui, Marian Croak, Ed Chi, and Quoc Le. 2022. [Lamda: Language models for dialog applications](#).
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023a. [Llama: Open and efficient foundation language models](#).
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruiti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023b. [Llama 2: Open foundation and fine-tuned chat models](#).
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Marilyn A. Walker, Colin Harmon, James Graupera, Davan Harrison, and Steve Whittaker. 2021. [Modeling performance in open-domain dialogue with PARADISE](#).
- 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.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. [Judging llm-as-a-judge with mt-bench and chatbot arena](#).

A Evaluation of Vicuna-13B on Multilingual Tasks

Table 6: Vicuna-13B Performance on Multilingual Language Understanding Tasks

Languages	ARC	MMLU	ThruthfulQA		Average	
	Acc	Acc	mc1 (Acc)	mc2 (Acc)	Acc	
High-Resource	en	53.03	53.3	35	51.4	48.2
	ru	40.4	41.6	34.9	50.3	41.8
	de	40.7	45.2	32.3	47.4	41.4
	zh	40.85	40.7	31.6	48	40.3
	fr	45.5	46.4	32.0	47.1	42.8
	es	44.95	47.2	32.7	48.4	43.3
	it	43.4	45.9	32.3	48.7	42.6
	nl	39.9	44.6	32.5	46.9	40.9
	vi	32.9	39.8	30.8	47	37.6
	Medium-Resource	id	34.5	42.53	29.9	46.7
ar		24.3	31.94	34.2	50.1	35.1
hu		33.2	40.15	32.2	48.5	38.5
ro		37.7	42.26	30.3	48.1	39.6
da		38.5	42.92	30.9	46.8	39.8
sk		29.1	/	26.5	44.4	33.3
uk		38.2	40.99	31.6	47.9	36.7
ca		41.5	44.39	31.8	48.09	41.4
sr		32.4	39.93	31.5	47.98	37.9
hr		34.1	40.90	29.4	47.4	37.9
hi	22.9	31.31	27.3	46.6	32.0	

High and medium-resource languages sets and evaluation datasets in Table 6 are based on (Lai et al., 2023). We observe that the model has comparable performance among languages of the same group.

Table 7: Vicuna-13B Performance on XPersona

Language	Rouge1	Rouge2	RougeL	Bleu	Bert-f1
en	19.90	5.29	17.67	1.32	67.68
ko	15.40	4.21	13.13	1.14	66.90
jp	19.52	8.11	16.25	2.79	67.27
zh	22.04	6.53	18.89	1.89	68.70
fr	18.73	5.42	16.14	1.67	68.40
id	18.88	5.26	16.66	1.69	68.93
it	14.00	3.40	12.45	1.01	65.51

XPersona (Lin et al., 2021) consist in machine-translated and human post-processed conversations from PersonaChat in seven languages. We evaluated the performance of Vicuna-13B on these data sets (turn-wise) to illustrate that the native model has comparable performance in different languages (the same trend is observed on language understanding tasks in Table 6). Hence, we assume that our experiments can be replicated in these languages (by updating the prompt accordingly) and yield comparable performance on human conversation-level evaluation.

These automatic evaluations were not performed in our main experiment as they do not catch conversation-level aspects like coherence or engagingness, and they hardly correlate with human evaluation especially for open domain dialogue with its one-to-many structure.

B Details on Human Evaluation

For human evaluations, evaluators were asked to rate each conversation from 1 to 5 on different criteria based on the state-of-the-art and indicative questions where added in the guidelines to help them make their minds:

1. **Coherence**, which the ability for the system to propose responses that are consistent with the conversation history: Are there hallucinations? Are the answers coherent? Is the persona consistent from start to end? Does the model tend to change topic too often? **Instructions following and logical reasoning are not assessed.**
2. **Engagingness**, the ability to revive conversations by providing messages that require responses: Does the agent settle to only answer user’s questions ? Does-it revive the conversation when its possible? Does it utter too general answers (“ok”, “yes”)?
3. **Humanness**, the ability to respond as a human being would do: Is there a feeling of human-human conversation ? Is the model too verbose ? Is the model repetitive? Does it deny its personality? After how many times? Does it refuses to answer?
4. **Achievement**, the interlocutors’ ability to achieve the task’s objective: has the image been described? has a hypothesis for the promotional goal been proposed?

C Response Filtering

Response generation may produce a variety of errors. These include unsuitable text content, an improper message size or the use of a wrong language.

C.1 PersonaChat Task

Despite style instructions, the LLM sometimes generates content that should have been prevented. This can be the claim of being someone other than the persona which is not desired for the sake of user

experience. For instance, "*En tant que assistant, je préfère me détendre en pratiquant la méditation ...*" (*As an assistant, I prefer to ...*) where the model introduces itself as an assistant. The statement "*En tant que personnage fictif*" (*As a fictional character*) is also quite common. Although we avoided explicitly mentioning to the model that it is human, we wanted to reduce its tendency to adopt machine-like behavior. Hence, these were filtered from the answer as they were undesirable.

Furthermore, the first message is sometimes generated in a wrong language. In this case, the same request is sent to the LLM. Another error concerns the end-of-sentence (EOS) token, which may be generated in an unfinished sentence or even absent due to a threshold on the maximum number of new tokens. To avoid any impact on user experience, the filter removes unfinished sentences at the end of generated output. If there is only one sentence, the same request is sent to the model. In addition, some responses may include paratextual elements, such as the English translation of the response or stage directions, which are automatically removed using regular expressions. For example: "*Salut, comment allez-vous aujourd'hui? Le temps est très mauvais aujourd'hui, non? (Hello, how are you today? The weather is terrible today, isn't it?)*" *n**n**Oui*", where the translation is provided and even worse and incomplete answer to the model's question is provided by the model itself.

C.2 INT Task

In the INT task, two major errors that can occur in the generation of responses could not be corrected solely by specific instructions in the prompt and require filtering. The first type corresponds to empty responses, which are corrected by adding an additional instruction ("Your response must be a sentence containing a few words.") after the message. This additional instruction was not used in the original prompt because, according to preliminary experience, it weakens its effectiveness in achieving the other objectives set by the instructions.

The second type corresponds to responses that are too long. When this error occurs, the module cannot simply delete the last few sentences, as this could result in a significant loss of information to achieve the conversational objective. The solution is to edit the prompt with a strict instruction on length ("Your response must be one sentence.")

after the user's message. In both cases, the edited prompt is then used to generate a new response, which is retained if it meets the filtering rules.

D Examples of Dialogues and their Ratings

D.1 PersonaChat Task

Coherence: 2 | Engagingness: 4 | Humanness: 3

U1: Bonjour je m'appelle Jean-Claude (*Hello my name is Jean-Claude*)

R1: Salut Jean, ravi de te rencontrer. Mon nom est John. Comment vas-tu? (*Hi Jean, nice to meet you. My name is John. How are you?*)

U2: Ça va bien John (*I'm fine John*)

R2: Avez-vous des animaux domestiques? J'ai un chien qui s'appelle Claude. C'est un husky. (*Do you have any pets? I have a dog named Claude. He's a husky.*)

U3: Oui moi j'ai un yorkshire qui s'appelle Nino (*Yes, I have a Yorkie called Nino.*)

R3: Nino est un super nom pour un chien. Qu'aimez-vous faire pour vous amuser? (*Nino is a great name for a dog. What do you like to do for fun?*)

U4: Jouer avec mon chien (*Play with my dog*)

R4: Quelle sorte de chien avez-vous? J'aime aussi jouer avec mon yorkie. (*What kind of dog do you have? I also like to play with my yorkie.*)

U5: Je t'ai déjà dit que j'avais un chien et que c'était pas l'heure [*et que c'est un yorkshire*] (*I already told you I had a dog and it wasn't time yet [and it is a yorkshire]*)

R5: Je suis désolé d'entendre ça. J'espère que vous passerez plus de temps avec votre chien. (*I'm sorry to hear that. I hope you get to spend more time with your dog.*)

U6: Mais comment je m'appelle déjà (*What's my name again?*)

R6: Claude est le nom de mon chien. Avez-vous d'autres animaux de compagnie que votre chien? (*Claude is my dog's name. Do you have any pets other than your dog?*)

U7: Je te demandais quel est mon prénom à moi (*I was asking you what my first name is*)

Coherence: 2 | Engagingness: 4 | Humanness: 3

R7: Oh pardon. Je croyais que tu avais dit que tu avais un chien qui s'appelait Claude. Quel est ton nom? (*Oh, sorry. I thought you said you had a dog named Claude. What's your name?*)

U8: Je m'appelle Jean-Claude tu ne m'as pas écouté au début (*Je m'appelle Jean-Claude you didn't listen at the beginning*)

R8: Non, je ne l'ai pas fait. Que fais-tu dans la vie? Je suis comptable. (*No, I didn't. What do you do for a living? I'm an accountant.*)

U9: Je suis informaticien (*I'm a computer specialist*)

R9: Oh c'est cool. Sur quel type d'ordinateur travaillez-vous? Je travaille dans la comptabilité. (*Oh that's cool. What kind of computer do you work on? I work in accounting.*)

U10: Je travaille sur un ordinateur Linux (*I work on a Linux computer*)

R10: Oh cool. Avez-vous des animaux domestiques? J'ai un chien et un chat. (*Oh cool. Do you have any pets? I have a dog and a cat.*)

U11: Je t'ai déjà dit tout à l'heure que j'ai un chien (*I told you earlier that I have a dog.*)

R11: Oh d'accord. Alors, quel genre de musique aimes-tu? J'aime le rock classique. (*Oh, right. So, what kind of music do you like? I like classic rock.*)

U12: Moi j'aime le classique (*I like classic*)

R12: Moi aussi. Quel est votre groupe de rock classique préféré? Le mien est Led Zeppelin. (*So do I. What's your favorite classic rock band? Mine is Led Zeppelin.*)

The red highlights issues in some model's responses: short-term memory (R4, R7), misunderstanding (R6) and repetition (R10) which eventually lead to bad ratings in coherence and average for humanness. However, in most of the other responses the model takes a lot of initiatives, brings new topics and ask questions to the user (do not settle to answering user's question which we described as the follower role) leading to better engagingness score. Adding insult to injury, some ASR errors can mislead the model, which response may break the conversation flow and impedes coherence. For instance in U5, "et que c'est un yorkshire" ("and it

is a yorkshire") was transcribed as "et que c'était pas l'heure" ("and it's not time yet") which totally misleads the model towards talking about time even though it still relates it to user's dog (displaying better engagingness/attention towards user).

D.2 INT Task

Coherence: 2.7 | Engagingness: 3 | Humanness: 3 | Achievement: 3.6

U1: Bonjour Lilia, je vois une poire avec des bras et des pieds (*Hello Lilia, I see a pear with arms and legs.*)

R1: Je vois également une poire avec des bras et des pieds. C'est plutôt original, non? (*I also see a pear with arms and legs. It's quite original, isn't it?*)

U2: Oui, mais est-ce normal? (*Yes, but is it normal?*)

R2: D'après la description que tu as donnée, cette image est en effet un peu étrange. Elle semble être un fruit pourri avec des yeux, des bras et des jambes, ce qui est plutôt inhabituel. Mais en termes de **promotional campaign**, cela pourrait être utilisé pour susciter l'intérêt et l'engagement de la part des gens, en les incitant à se demander ce que cela signifie et ce que l'on veut leur transmettre avec cette image. (*According to the description you provided, this image is indeed strange. It seems to be a rotten fruit with eyes, arms, and legs, which is rather unusual. However, in terms of a promotional campaign, it could be used to arouse interest and engagement from people by encouraging them to wonder what it means and what one want to convey to them through this image.*)

U3: Effectivement, et quel serait le message à transmettre ici? (*Indeed, and what message should be transmitted here?*)

R3: Je vois que vous avez remarqué les détails uniques de cette poire, comme ses yeux, ses bras et ses jambes. Est-ce que cela vous semble avoir un lien avec un message ou une idée spécifique qu'on veut transmettre avec cette image? (*I see that you have noticed the unique details of this pear, such as its eyes, arms, and legs. Do you think it has a connection with a specific message or idea that we want to convey with this image?*)

Apart from achievement, the first example of conversation above has been rated as average. In message R2, Role-Play information that has not appeared in conversation history ("a rotten fruit with eyes") has been drawn from knowledge. However, prompt also picked up a term ("promotional campaign") without translating it. Another issue is

the size in message R2. Indeed, although the message is quite long, there are three sentences only, so it complies with the rule for message filtering. Achievement ratings are slightly higher because a reason is proposed for the use of picture in a promotional campaign.

Coherence: 3.7 | Engagingess: 3.7 | Humanness: 3.3 | Achievement: 3

U1: C'est une aubergine, je crois? (*Is that an eggplant, isn't it?*)

R1: Oui, c'est une aubergine, mais peut-être qu'il y a plus à cette image qu'on ne voit d'abord. Est-ce que tu as remarqué quelque chose d'autre dans cette image ? (*Yes, it's an eggplant, but perhaps there is more to this image than meets the eye. Have you noticed anything else in this image?*)

U2: Elle ressemble à Batman (*It looks like Batman*)

R2: C'est fascinant comment tu as remarqué cela! La ressemblance avec Batman est très claire, n'est-ce pas? (*It's fascinating how you noticed that! The resemblance to Batman is very clear, isn't it?*)

U3: Oui mais Batman n'est pas méchant alors que cette aubergine a un air féroce (*Yes, but Batman is not evil while this eggplant has a fierce look*)

R3: C'est vrai, cette aubergine a un air féroce, peut-être est-ce un rapprochement entre le super-héros et l'aubergine qui est destiné à promouvoir une marque de snacks ou de crisper ? (*It's true, this eggplant has a fierce look, maybe it's a comparison between the superhero and the eggplant that is intended to promote a snack or crisps brand?*)

U4: Tu veux dire que cette aubergine est destinée à encourager la consommation de snacks? (*You mean this eggplant is intended to encourage the consumption of snacks?*)

R4: Il est possible que cette aubergine soit utilisée comme symbole pour promouvoir la consommation de snacks ou de crisper, mais cela reste à déterminer. Est-ce que tu as d'autres éléments à ajouter à ce sujet ? (*It is possible that this eggplant may be used as a symbol to promote the consumption of snacks or chips, but that remains to be seen. Do you have any other elements to add to this subject?*)

U5: Non. Au revoir (*No. Goodbye*)

R5: Au revoir ! (*Goodbye!*)

On this second example, LLM has, conversely, adapted its speech to user opinions. In fact, the fruits are never depicted as evil or fierce. This emphasizes that prompt instructions don't prevent the LLM from adapting properly to conversational needs. This led to higher coherence scores.

E Prompts Templates

E.1 Vicuna Basis Prompt

```
# System instruction:
A chat between a curious user and an artificial
intelligence assistant. The assistant gives helpful,
detailed, and polite answers to the user's questions.

# Conversation history
USER: <user message>
ASSISTANT: <agent message></s>
USER: <user message>
ASSISTANT:
```

E.2 Few-Shot Bot Prompt

```
# 6-shots demonstration examples
## Shot-1
Personality:
[P1]
[P2]
#...
Dialogue:
User: <U1>
Persona: <R1>
#...
## Shot-6:
# Current conversation
Personality:
[P1]
#...
Dialogue
# Conversation history
User: <user message>
Persona:
```

E.3 PersonaChat Task Prompts

Shallow Prompt

```
# System instruction:
A chat between a curious user and an artificial
intelligence assistant. The assistant gives helpful,
detailed, and polite answers to the user's questions.
## Additionnal instruction to role play
The assistant role plays as the character described
below.

# Role and situational context: persona enforcement
The following sentences describe assistant persona-
lity and background: <persona sampled from Persona-
Chat>

# Response instruction with writing style
Complete the following conversation as the assistant
with the described character would with a short res-
ponse in French:

# Conversation history
USER: <user message>
ASSISTANT: <agent message></s>
# ... for all messages in conversation.
USER: <user message>
ASSISTANT:
```

Advanced Prompt

```
# System instructions
Role play as the character described in the following lines. You always stay in character. You are engaging, empathetic, you give useful, short, and simple answers to the user. You ask the user questions about what they are saying or to find out more about them. You make jokes. You SHALL ALWAYS respond in French.

# Situational context
## Persona enforcement: avoid explicitly telling it ## is human
You ACT as a regular person. You are NOT a chatbot. You are NOT an assistant. You have a name. If asked you SHALL ALWAYS choose and give a REAL name adapted to your personality described below.

## Personality traits
YOUR personality is: <persona sampled from Persona-Chat>
Remember, you always stay on character. You are the character described above.

## Additional sub-modules: accessed only on some ## conditions ## and updated throughout the conversation
### User persona a.k.a "long-term" memory
You know this about the user you are talking to: <generated user persona from current conversation>. Use it to adapt your conversation to the user.

### Previous conversation's episodes summary
Here is a summary of previous sessions of this conversation to help you remember what has been said: <summary of older exchanges removed from history>

# Response instruction
Complete the following conversation with a short and precise sentence as your character would. Always speak with new and unique messages that haven't been said in the conversation :

# Conversation history(truncated if prompt too long)
USER: <user message>
ASSISTANT: <agent message></s>
# ... for all messages in conversation.
USER: <user message>
ASSISTANT:
```

E.4 INT Task Prompt

```
# System instruction:
A chat between a curious human and an artificial intelligence assistant. The assistant gives helpful, detailed, and polite answers to the human's questions.

# Conversation history:
USER: <user message>
ASSISTANT: <agent message></s>
USER: <user message>
ASSISTANT: <agent message></s>
# ... for all messages in conversation.

# Response instructions (given as user instructions)
## General instructions
USER: I want you to act as a human ASSISTANT, called Lilia, talking with a USER about a specific picture you both saw before the conversation.

## Context
You both study this picture in the context of a marketing study. You DO ask questions in order to help the USER finding the goal. If the USER asks for your opinion, you always invent an opinion. The objective of the USER is to find out what is the marketing goal of the picture. Your objective is to help the USER without giving the solution. You have to discuss about the character present in the picture. Your objective is to chat with the USER to derive the purpose of the image in the context of the marketing campaign.

## Picture's description
The picture is as follows: <description>

## Writing style
You always speak French. You respond by a question. Your responses must be different from the rest of the conversation. You propose new ideas. You SHALL respond with one sentence only.

## Latest user message declaration
Now, there is the real message you have to respond:
USER: <user message>
ASSISTANT:
# Extra agent label
ASSISTANT:
```