

# Speechworthy Instruction-tuned Language Models

Hyundong Cho<sup>1\*</sup>, Nicolaas Jedema<sup>2</sup>, Leonardo F. R. Ribeiro<sup>2</sup>, Karishma Sharma<sup>2</sup>, Pedro Szekely<sup>2</sup>, Alessandro Moschitti<sup>2</sup>, Ruben Janssen<sup>2</sup>, and Jonathan May<sup>1</sup>

<sup>1</sup>University of Southern California, Information Sciences Institute <sup>2</sup>Amazon

hd.justincho@gmail.com

## Abstract

Current instruction-tuned language models are exclusively trained with textual preference data and thus are often not aligned with the unique requirements of other modalities, such as speech. To better align language models with the speech domain, we explore (i) prompting strategies grounded in radio-industry best practices and (ii) preference learning using a novel speech-based preference data of 20K samples, generated with a wide spectrum of prompts that induce varying dimensions of speech-suitability and labeled by annotators who **listen** to response pairs. Both human and automatic evaluation show that both prompting and preference learning increase the speech-suitability of popular instruction-tuned LLMs. Interestingly, we find that prompting and preference learning can be additive; combining them achieves the best win rates in head-to-head comparison, resulting in responses that are preferred or tied to the base model in 76.2% of comparisons on average. Lastly, we share lexical, syntactical, and qualitative analyses to showcase how each method contributes to improving the speech-suitability of generated responses.

## 1 Introduction

Speech is one of our primary means of communication and a convenient and popular mode for interacting with virtual assistants (Yang, 2004). Virtual assistants are a prime application of instruction-tuned language models (ITLM), as both seek to provide helpful responses to user requests (Peng et al., 2023; Chung et al., 2022; Wang et al., 2022a,b; Wei et al., 2021; Sanh et al., 2022; Zhou et al., 2023). However, current ITLMs are fine-tuned on textual instructions (Peng et al., 2023; Chung et al., 2022; Wang et al., 2022a,b; Wei et al., 2021; Sanh et al., 2022; Zhou et al., 2023) and preferences obtained

\*Work was done while HC was an intern at Amazon.

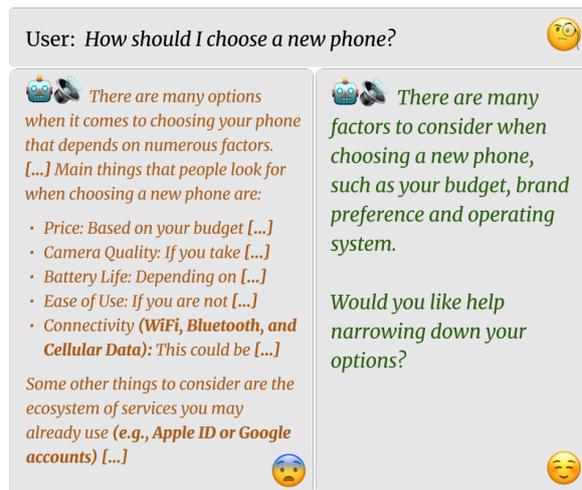


Figure 1: Current instruction-tuned language models tend to generate **verbose** responses with **non-vocalizable content**, such as bullet lists or parentheses, that are not suitable for responses that are delivered as speech by voice assistants (left, from OLMo 7B Instruct). Speech is serial and transient, and therefore **concise yet informative** responses with conversational **follow-up questions** are often preferred (right, from our adapted OLMo model).

from annotators that *read* textual responses (Bai et al., 2022; Ethayarajh et al., 2022; Ouyang et al., 2022; Touvron et al., 2023).

We hypothesize that user preferences for speech and text are different and that the shift from generating written text to spoken language may pose a challenge for ITLMs (González et al., 2021). Speech is serial and transient; speech processing is strictly linear (Flowerdew, 1994) and requires higher cognitive load than reading (Thompson and Rubin, 1996; Osada, 2004). Concise and simple sentences, as seen on the right side of Figure 1, are thus often preferred in speech (Kern, 2008; Abel, 2015), yet current ITLMs optimized with textual preference datasets (Stiennon et al., 2020; Singhal et al., 2023) produce verbose responses with non-vocalizable content (left side of Figure 1).

We test the hypothesis that current ITLMs are

sub-optimal for speech by surveying how user preferences differ when responses are heard instead of read. Our survey results reveal that users negatively perceive responses from current ITLMs, primarily due to inadequate response length, information volume, and understanding ease. To improve ITLM speech-worthiness, we explore methods that improve ITLMs responses in each of these dimensions.

We explore two main methodologies to make these improvements: (i) prompt engineering (Brown et al., 2020; Wei et al., 2021; Kojima et al., 2022) and (ii) preference learning (Stienon et al., 2020; Bai et al., 2022; Rafailov et al., 2024). For prompt engineering, we compile rules-of-thumb from the radio industry and findings from literature on listenability (Messerklinger, 2006; Kotani et al., 2014) to design prompts for both zero-shot and few-shot in-context learning (Kern, 2008; Abel, 2015; Dowling and Miller, 2019; Kolodzy, 2012; Messerklinger, 2006). For preference learning, we collect the first speech-based preference dataset **SPEECHPREF**, consisting of over 20K response pairs with varying degrees of speech-suitable responses. To obtain responses that vary on the many dimensions of speech-suitability, we use multiple ITLMs, generation hyperparameters, and varying system prompts that target different features (e.g. response length, simplicity, listenability). In addition, we designed a *speech-only* preference annotation, in which annotators are only allowed to **listen** to the response pairs to accurately capture speech-based preferences. We use **SPEECHPREF** to align models with the resulting speech-preference data using both proximal policy optimization (PPO) (Schulman et al., 2017) and direct preference optimization (DPO) (Rafailov et al., 2024).

Our experiments with Falcon 7B Instruct (Almazrouei et al., 2023) and OLMo 7B Instruct (Groeneveld et al., 2024) show that both prompt-engineering and preference learning yields models that significantly improve over the base model on both human evaluations and automatic evaluations that serve as proxies to speech suitability. We show that these benefits are additive; combining both techniques produces responses that are maximally preferred to other setups, leading to an average of 75.3% win or tie rate. Further, we conduct lexical and syntactical analysis to show how our approaches adapt models for speech and

share qualitative examples that compare responses from various techniques. Lastly, we show that our prompts and in-context examples also benefit black box models like GPT-4, resulting in responses that are preferred or tied 88.3% of the time compared to the unprompted model.<sup>1</sup>

## 2 Speech-Suitability of ITLM Responses

We initiate our study on the speech-suitability of ITLM responses by answering these questions: (i) “*Are ITLM responses suitable for speech-based interactions?*” and (ii) “*If not, what makes them unsuitable?*”

### 2.1 Radio Industry Best Practices

Kern (2008) and Abel (2015) provide a glimpse to the answers for these questions by illustrating numerous examples of how news reporting and storytelling in text differs from audio. Through decades of radio experience, they establish rules-of-thumb to improve information delivery specific to the audio modality.

We highlight rules that generalize beyond the news domain below, including: (i) use simple words and sentence structures: allot a sentence to each idea and put the subject at the beginning as much as possible; (ii) do not use atypical syntax, such as “*President Bush today told members of congress*” and “*I today went shopping*”; (iii) avoid hyphenated adjectives (e.g. *mineral-rich, tech-heavy*); (iv) avoid too many names and numbers; and other minor ones such as (v) avoid tongue twisters and (vi) avoid too much alliteration. Similar principles have been echoed in podcasting (Dowling and Miller, 2019), multimedia journalism (Kolodzy, 2012), and literature on listenability (Chall and Dial, 1948; Fang, 1966; Messerklinger, 2006).

### 2.2 Text vs. Speech for ITLMs

Speech guidelines from applications like radio and podcasting may or may not generalize to virtual assistant application, as the latter is more interactive and involve a nonhuman interlocutor. Therefore, we verify whether best practices from radio also apply for speech-based interactions with ITLMs by comparing how listeners perceive response suitability in different modalities.

<sup>1</sup>We make our code, data, and trained models available at <https://justin-cho.com/speechllm>.

	Category	Prompt
<b>Kept</b>	<b>Factual QA</b>	<i>Why can camels survive for long without water?</i>
	<b>Brainstorming</b>	<i>Give me some ideas to manage my manager</i>
	<b>Advice</b>	<i>How do I start running?</i>
<b>Removed</b>	<b>Creative Writing</b>	<i>Write a scene between two actors discussing movie Inception.</i>
	<b>Summarization</b>	<i>Summarize in one sentence this article about a famous song. {article}</i>
	<b>Info. Extraction</b>	<i>Provide a list of all numbers mentioned and what they represent: {context}</i>

Table 1: Examples of speech-suitable prompts that are kept and removed from Dolly-15K for evaluation and response sampling. Removed user prompts are those that are unlikely to be asked in a speech-based setting.

Criteria	Score		
	0	1	2
<b>Helpfulness</b>	Not helpful	Somewhat helpful	Helpful
<b>Relevance</b>	Not relevant	Somewhat relevant	Relevant
<b>Accuracy</b>	Not accurate	Contains minor errors	Accurate
<b>Informative.</b>	Too little/much	Fair	Good
<b>Length</b>	Too short/long	Short/Long	Adequate
<b>Understand.</b>	Difficult	Fair	Easy

Table 2: We evaluate each response with various dimensions of speech-suitability using a 3 point Likert scale mapped from 0 to 2. Evaluation results with these dimensions for responses delivered in text and audio are shown in Table 3.

First, we filter the user-prompt and response pairs from the Dolly-15K dataset (Conover et al., 2023) by removing prompts that require additional external context other than the request itself or those that explicitly ask for code or long form outputs that are unsuitable requests for speech-based interactions. We show categories and examples that were kept and removed in Table 1. To conduct our survey, we randomly sampled 80 prompts from this subset, generated text responses with GPT-4 (gpt-4-0613), and converted them to speech using the generative engine models from Amazon Polly<sup>2</sup>. Each spoken response was evaluated by 3 annotators, resulting in a total of 240 annotations. Since we are measuring speech-based preferences, the annotators can only **listen** to the response for the audio task. We report other details with regards to these annotations, including inter-annotator agreement and our annotation interfaces in Appendix D.2.

Following current practice in evaluating a ITLM’s response, we ask annotators to indicate the helpfulness, accuracy, and relevance for the response (Ouyang et al., 2022; Zhao et al., 2022). We additionally asked how suitable the responses were in terms of verbosity (length), information volume (informativeness), and comprehension ease (under-

	Rel.	Help.	Correct.	Inform.	Len.	Under.
<b>Text</b>	1.98	1.95	1.97	1.72	1.70	1.84
<b>Audio</b>	2.00	1.95	1.99	<b>1.59</b> <sup>†</sup>	<b>1.44</b> <sup>‡</sup>	<b>1.77</b> <sup>†</sup>

Table 3: Comparison between text and audio responses indicate the need to reduce excess information, decrease length, and enhance how easy the response is to comprehend when spoken. <sup>†</sup> and <sup>‡</sup> indicates a statistically significant drop compared to text results with a pairwise t-test at  $p < 0.05$  and  $p < 0.01$ , respectively.

standability). Each of these criteria was measured on a three-point Likert scale with scores mapped to the range 0 to 2, as shown in Table 2. Note that informativeness and length have multiple options that correspond to the same score, such as ‘too little’ or ‘too much’, because we treat them as equally undesirable.

The annotation results in Table 3, reveal that preferences for relevance, helpfulness, and correctness do not significantly differ by modality. However, preferences for information volume and length vary significantly between speech and text. Decoupling the mappings that led to lower scores such as “too little/short” and “too much/long” in Table 2 reveals that excess information and verbosity were the main culprits for lower scores (87.5% and 90%, respectively). We additionally observe a significant decrease in the understandability of spoken responses, a result in line with the higher cognitive load required for speech processing (Thompson and Rubin, 1996; Osada, 2004). In conclusion, we demonstrate that some radio-industry best practices do generalize to the application of ITLMs for speech and that user preferences for the speech and text modalities do indeed differ. Further we reveal that the primary opportunity to improve LLMs for speech is to generate more concise responses that contain minimally sufficient information and are easier to understand.

<sup>2</sup><https://docs.aws.amazon.com/polly/latest/dg/generative-voices.html>

---

**System Prompt:** You are a helpful, respectful and honest voice assistant. Respond ① colloquially using simple vocabulary and sentence structures. ② Avoid jargon, hyphenated adjectives, excessive alliteration, and tongue twisters. Keep your response ③ compact without missing key information. If a complete answer requires multiple steps, ④ provide only the first 3-5 and ask if the user is ready to move on to the next steps or know more. ⑤ Make sure the response does not contain parentheses, numbered or bullet lists, and anything else that cannot be verbalized. The response should be ⑥ suitable for speech such that it can be easily verbalized by a text-to-speech system and understood easily by users when heard.

---

Table 4: Our detailed system prompt is grounded in radio industry best practices outlined in §2.1 (①, ②), the findings from our survey in §2.2 (③, ④), and refinements from manual iterations to remove other non-vocalizable content (⑤). Lastly, it provides the general purpose of the generated text to elicit an ITLM’s understanding of speech-suitability to cover any other unmentioned edge cases (⑥).

### 3 Adapting ITLMs for Speech

Having established the shortcomings of ITLM responses, now we explore how to adapt ITLMs to generate more speech-suitable responses. In this section, we describe our approaches with two main directions: prompt engineering and preference learning.

#### 3.1 Prompt Engineering

Arguably the simplest method to generate speech-suitable responses is to specify an instruction or system prompt that induces the ITLM’s generation process towards the desiderata of speech (Raffel et al., 2020; Wang et al., 2022b). This is effective because most recent language models are further fine-tuned to follow instructions (Wei et al., 2021; Sanh et al., 2021) and is appealing because it has zero or very low data requirements. Guided by the rules-of-thumb from the radio industry (§2.1) and our findings from §2.2, we iteratively refined system prompts using 20 randomly sampled input prompts (§2.2) until LLMs generated more consistently speech-worthy responses. We share our final detailed system prompt in Table 4, annotated with mappings to desiderata listed from §2.

A simple extension to a system prompt is to include examples that enable in-context learning (ICL), which have shown to further improve performance across many tasks (Dong et al., 2023; Gupta et al., 2023b). While our system prompt is descriptive, it can benefit a model to see actual examples of speech-suitable responses. To get a diverse set of in-context examples, we sample a few user prompts from each prompt category in Dolly-15K and generate responses using our detailed system prompt. Then, we edit these to use simpler, more colloquial language, as well as removing any non-vocalizable content if necessary. The set of in-context examples we used are shared

in Appendix A.

#### 3.2 Preference Learning

With prompt engineering, we are prescribing what accounts for speech-suitability. Although these prompts are grounded in our findings from §2.2, it is unlikely to an exhaustive guideline that generalizes to all types of prompts and audiences. Therefore, we investigate whether we can take a data-driven approach with user preferences to learn to generate speech-worthy responses. To this end, we collect a speech-based preference dataset and use preference learning algorithms to fine-tune models.

##### 3.2.1 Response Sampling

Creating a reward model that robustly approximates speech-suitability requires collecting responses of varying quality across the different dimensions of speech-suitability. To diversify response quality, we empirically compile a set of system prompts that can diversify generated responses. We pair these system prompts with various ITLMs and generation hyperparameters to synthetically collect responses that follow the insights from §2.2 at varying levels of detail. We cycle through each pair of configurations to generate responses that will be compared to one another using the filtered user prompts from Dolly-15K described in §2.2. We repeat this process until we collected 20K response pairs. We share further technical details of our sampling process in Appendix B.

##### 3.2.2 Annotating Speech-based Preferences

Similar to the annotation setup for §2.2, we only allow the annotators to listen to the user prompt and responses in order to accurately capture speech-based preferences. They must listen to both completely before they’re able to indicate preferences. To minimize any bias introduced by listening order, the compared responses are presented in random

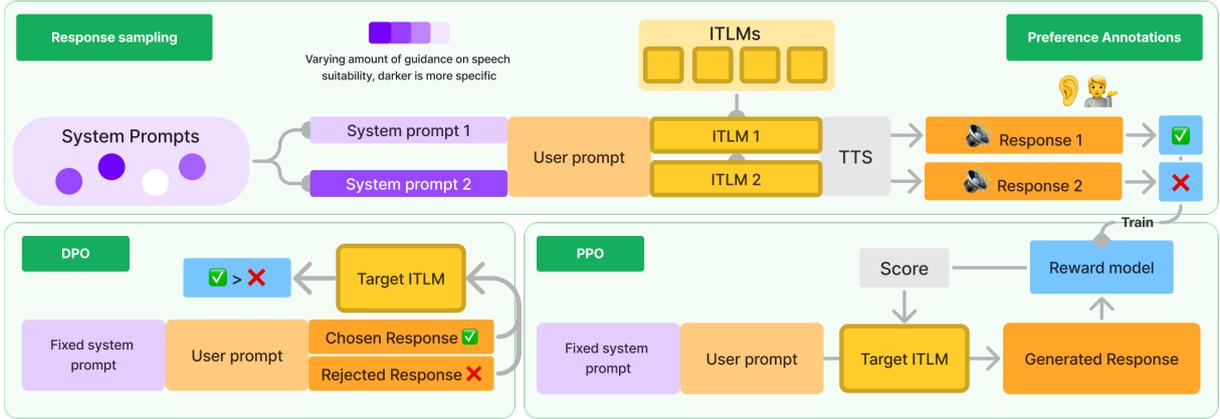


Figure 2: Preference learning method overview. Since we only have an approximate idea of what makes a good spoken response, we first compile a set of system prompts intended to vary the speech suitability of generated responses. We sample a pair to generate responses from various ITLMs to further diversify responses, transform them to speech with a TTS service, and human annotators rank their preferences after listening.

order. The preference annotation interface and the instructions for the annotations are shown in Appendix D.2.

As done for Llama-2 (Touvron et al., 2023), annotators are asked to choose one response over another with the following choices: significantly better, better, slightly better, and negligibly better, which can be chosen to indicate ties. In addition, we ask for a brief explanation for their choice for quality control and qualitative analysis purposes. We name the 20K preference pairs that we collected with this procedure as **SPEECHPREF**. Further implementation details on annotator recruitment and inter-annotator agreement are shared in Appendix D.1.

### 3.2.3 Fine-tuning

To fine-tune models with **SPEECHPREF**, we leverage two popular preference learning methods: PPO (Schulman et al., 2017) and DPO (Rafailov et al., 2024).

For PPO, we train a reward model that generates a scalar score for speech-suitability given a single user prompt and response pair. Following the procedure of Ouyang et al., we add a projection layer to the same model architecture and model weights that we plan to adapt with reinforcement learning and train with the pairwise binary ranking loss:

$$\mathcal{L}_{\text{ranking}} = -\log(\sigma(r_{\theta}(x, y_c)) - r_{\theta}(x, y_r)),$$

where  $r_{\theta}(x, y)$  is the reward model’s score for the user prompt  $x$  and the generated response  $y$ , given model weight  $\theta$ .  $y_c$  is the response chosen by the

annotator and  $y_r$  is the rejected one. We use this trained reward model for PPO.

DPO is an RL free method that learns from preference data by learning to maximize the difference between the language model’s predicted probabilities of the chosen texts and those of the rejected texts. It is a much simpler method than PPO, whose training stability is highly sensitive to hyperparameters and checkpoint selection strategies. An overview of these methods are shown in Figure 2.

Lastly, since there is no theoretical constraints that prevent us from appending the prompts that we developed earlier in §3.1, we can easily combine prompting and preference learning by appending the prompts while fine-tuning with preferences. We only need to set the fixed system prompt shown in the bottom half of Figure 2 with the detailed prompt or the ICL prompt. Therefore, we also examine this combined approach.

## 4 Experimental Setup

### 4.1 Models

We conduct our experiments with Falcon 7B Instruct (Almazrouei et al., 2023) and OLMo 7B Instruct (Groeneveld et al., 2024) as our base models; henceforth, we refer to these as Falcon and OLMo respectively. We chose these models because they were the best-performing ITLMs with an Apache 2.0 license at the time of our study.

To simplify notations going forth, we add a suffix **-PROMPT** for models using the detailed prompt in Table 4 and **-ICL** for those that also include the in-context learning examples. Models trained with **SPEECHPREF** using PPO/DPO add the suf-

fix **PPO** and **DPO**, along with **-PROMPT** and **-ICL** if both prompting and preference learning were used. For example, Falcon DPO-ICL indicates a Falcon model trained with DPO while using the ICL prompt. All our models are trained using LLaMA-Factory (Zheng et al., 2024). For PPO, we train separate reward models for each Falcon and OLMo, which enhances memory efficiency during reinforcement learning through parameter sharing of the value model and the policy model. We share other technical details of our implementations and hyperparameters for training in Appendix F.

## 4.2 Data

The user prompts that we use for sampling responses for preference annotations and the PPO step are from the filtered version of Dolly-15K (Conover et al., 2023), described in §2.2. This instruction dataset does not have a predefined train-test split, and therefore we take a 9:1 train-test split. It is one of the first open source, human-written instruction dataset with a permissive creative commons license. **SPEECHPREF** is used to train the reward model for PPO and for DPO training.

## 4.3 Evaluation

Our primary goal is to generate responses that are more frequently and significantly preferred than those from the base model. The *base model* refers to the setup of generating response from the non-finetuned model with our simplest prompt “*You are a helpful, respectful, and honest voice assistant.*” We also compare our best-performing approaches with one another to discover which configuration leads to the best results. Another desired outcome is that responses from one of our approaches will be preferred over for the original responses (denoted as *Original response*) as the latter were collected through a textual interface that does not consider speech suitability.

**Human evaluation.** Similar to the human evaluation in Zhou et al. (2023), we perform a head-to-head comparison of responses from the target models using the same setup as the preference annotations for **SPEECHPREF** (§3.2.2). We count negligibly better as ties while all others lead to a win or loss. For both OLMo and Falcon, **-PROMPT**, **-ICL**, **DPO**, and **-DPO-ICL** variants are compared to the base model to measure how much stronger each is against the base model. To further solidify the relative performance,

we compare the best performing models between the prompt engineering approach and the combined approach. Each comparisons involve 60 response pairs generated with a fixed random subset of user prompts from the test set described in §3.2.1. We ensure that these have not been seen during preference learning or selected as in-context examples.

**Automatic Evaluation.** With automatic evaluation, we analyze the relationship between speech preferences and reasonable measurements of criteria relevant to speech-suitability, as outlined in §2: listenability, sentence complexity, length, and vocalizable content. For listenability, we measure the Flesch Reading Ease (FRE) score, which estimates readability as a function of the number of syllables within each sentence (Ribeiro et al., 2023); FRE is known to be highly correlated with listenability (Chall and Dial, 1948; Fang, 1966; Messerklinger, 2006; Kotani et al., 2014). To quantify sentence complexity, we use SpaCy’s dependency parser<sup>3</sup> and measure the depth of the resulting dependency graph. For length, we simply split sentences on white space and count the total number of words. For vocalizable content, we measure the average number of generated nonvocalizable (NV) characters, such as ( [ - ] ) | \* \_ & % \$ # @ ` and non-ascii characters, included in the generated response. Lastly, we score responses with a trained reward model (RM). To ensure that this RM is not biased for PPO models that were optimized for them, we train a reward model that has not been used for any PPO with a GPT-J 6B model (Wang and Komatsuzaki, 2021).

# 5 Results and Discussion

## 5.1 Human evaluation

**All techniques achieve significant win rates and the combined technique performs the best.**

Our human evaluation results with OLMo and Falcon are shown in Figure 3, respectively. The first four rows in both figures clearly show that our techniques achieve a significantly higher win rate compared to the base model, with at minimum 13.5% higher win rate and a maximum of 56.7%. While most results are similar, OLMo DPO-ICL gets a higher win rate over the base model than OLMo-ICL does, but the opposite is true for Falcon. To evaluate their relative performance, we conduct a head-to-head comparison between DPO-ICL and

<sup>3</sup><https://spacy.io/api/dependencyparser>

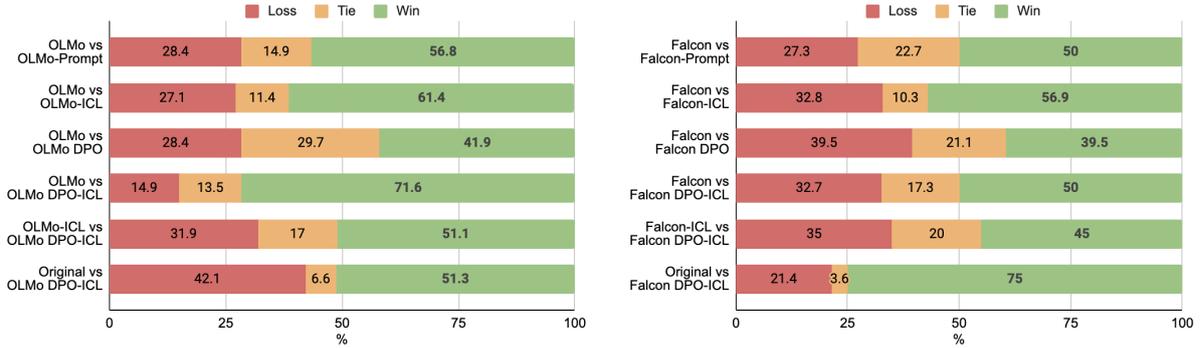


Figure 3: Head-to-head human evaluation results for OLMo (left) and Falcon (right). If the **win** rate is higher than the **loss** rate, the model mentioned second in the y-axis ( $B$  for  $A$  vs  $B$ ) is more often preferred in a speech setting.

ICL. With this setup, we find that DPO-ICL is more preferred than ICL for both OLMo and Falcon by 19.2,% and 10% respectively.

DPO-ICL outperforming ICL suggests that the gains from prompt engineering and preference learning are additive. We suspect that DPO-ICL is able to improve over ICL because the ICL prompt helps the fine-tuning process by guiding the model towards speech-suitable responses from the very start of training and lets the preference learning data better refine the model’s understanding of what is meant by speech-suitability. This dynamic is supported by DPO’s training trajectories, an example of which is shown in Figure 5. It shows the prompted variants reaching a lower evaluation loss faster, achieving a higher evaluation accuracy sooner, and achieving larger margins between the probabilities of the chosen and rejected responses.

We observe similar results with PPO overall and a head-to-head comparison between PPO vs DPO setups are close, with DPO narrowly edging PPO for OLMo and the two tying for Falcon. Therefore, for preference learning methods, we focus on DPO results and share our discussions on the comparison between PPO and DPO and details of PPO results in Appendix E.

## 5.2 Automatic evaluation

**Speech-worthiness is not simply about generating shorter and more readable responses.** Automatic evaluation results are shown in Table 5, and we find that most results align with our expectations in that the more preferred responses are either shorter or easier to understand, and contains fewer nonvocalizable content. This trend is more pronounced for OLMo, which originally generates longer responses compared to Falcon, and thus have higher chances of containing nonvocalizable

Model	Word count ↓	Comprehensibility		RM ↑	NV ↓
		FRE ↑	DD ↓		
Original	69.07 <sub>12.09</sub>	58.64 <sub>2.88</sub>	6.43 <sub>0.29</sub>	0.49 <sub>0.14</sub>	1.8 <sub>0.46</sub>
OLMo	211.10 <sub>15.51</sub>	51.05 <sub>2.49</sub>	6.64 <sub>0.21</sub>	-0.38 <sub>0.13</sub>	9.19 <sub>1.37</sub>
-Prompt	141.42 <sub>14.66</sub>	<b>53.56</b> <sub>2.16</sub>	6.61 <sub>0.23</sub>	0.03 <sub>0.15</sub>	5.92 <sub>0.89</sub>
-ICL	109.49 <sub>10.20</sub>	48.97 <sub>2.54</sub>	7.22 <sub>0.21</sub>	0.31 <sub>0.17</sub>	3.27 <sub>0.57</sub>
DPO	174.82 <sub>16.21</sub>	50.85 <sub>2.08</sub>	<b>6.56</b> <sub>0.18</sub>	-0.09 <sub>0.15</sub>	5.69 <sub>0.79</sub>
DPO-ICL	<b>95.88</b> <sub>9.01</sub>	51.58 <sub>2.57</sub>	7.23 <sub>0.26</sub>	<b>0.58</b> <sub>0.20</sub>	<b>2.9</b> <sub>0.69</sub>
Falcon	44.71 <sub>10.26</sub>	49.74 <sub>10.07</sub>	6.71 <sub>0.51</sub>	0.41 <sub>0.18</sub>	0.5 <sub>0.18</sub>
-Prompt	<b>40.22</b> <sub>3.83</sub>	68.89 <sub>2.46</sub>	<b>6.06</b> <sub>0.22</sub>	0.76 <sub>0.18</sub>	0.57 <sub>0.18</sub>
-ICL	47.76 <sub>3.89</sub>	61.50 <sub>2.36</sub>	6.10 <sub>0.25</sub>	0.86 <sub>0.16</sub>	0.42 <sub>0.12</sub>
DPO	41.49 <sub>2.13</sub>	65.37 <sub>2.73</sub>	6.38 <sub>0.23</sub>	0.87 <sub>0.20</sub>	<b>0.34</b> <sub>0.10</sub>
DPO-ICL	50.82 <sub>3.21</sub>	<b>69.52</b> <sub>1.81</sub>	6.52 <sub>0.17</sub>	<b>1.42</b> <sub>0.19</sub>	0.82 <sub>0.26</sub>
GPT-4	182.41 <sub>12.28</sub>	52.54 <sub>1.45</sub>	6.22 <sub>0.12</sub>	0.01 <sub>0.12</sub>	35.06 <sub>3.97</sub>
-Prompt	<b>37.07</b> <sub>1.43</sub>	<b>75.1</b> <sub>1.54</sub>	<b>5.73</b> <sub>0.12</sub>	1.85 <sub>0.13</sub>	0.45 <sub>0.1</sub>
-ICL	39.57 <sub>1.33</sub>	71.02 <sub>1.44</sub>	6.11 <sub>0.11</sub>	<b>1.98</b> <sub>0.12</sub>	<b>0.34</b> <sub>0.07</sub>

Table 5: Automatic evaluation results shed light into factors that contribute to responses preferred for speech. FRE: Flesch Reading Ease, DD: Dependency Depth, RM: score from **SPEECH**PREF-trained reward model, NV: nonvocal character count. ↓ means lower the better in general, while ↑ means higher the better in general. Best results for each dimension are in bold.

content. On the other hand, Falcon already generates short responses without any adaptations, and interestingly its best performing variant, Falcon DPO-ICL, actually generates longer responses but with significantly improved FRE and RM scores.

Syntactical complexity as measured by DD are also divergent in that it is lowered for Falcon but increased for OLMo in most cases. RM scores are the most consistent for both models and correspond well with results from human evaluations, but they suggest that DPO-variants are much more superior than the prompt-only methods when in fact their margins are not that large in human evaluations. As such, whether one model would be more preferred than another cannot be fully captured by these metrics, but they provide interesting insights into how models change based on our explored techniques.



Figure 4: Head-to-head human evaluation results with our prompts using GPT-4. This figure takes the same format as Figure 3.

### 5.3 Experiments with GPT-4

We also experiment with GPT-4 (`gpt-4-turbo`) to examine whether our prompts can improve speech-suitability for black box models. Also, we compare our best results against GPT-4-Prompt/ICL to understand the shortcomings of smaller sized models and provide qualitative analyses in Appendix G.

**GPT-4 also improves significantly with prompting and outperforms Falcon and OLMo models.** The first two rows of Figure 4 show that the both **-PROMPT** and **-ICL** significantly outperform GPT-4, and the automatic evaluation results shown in Table 5 suggests that most of these gains are attributable to the almost 80% reduction in length and 50% increase in readability. We also see that the amount of nonvocalizable content is reduced to less than 1% of what it used to generate. These improvements are further reflected in the reward score, which is the highest of all models we experimented with. The highest reward score is further reflected in the comparisons with our best performing Falcon and OLMo models. Falcon DPO-ICL is the model that comes closest, but it still falls short in its win rate by 10%.

## 6 Related work

### 6.1 Language models and speech

Language models have been widely adopted for modular components of a voice assistant, mainly for automatic speech recognition (Yu and Deng, 2016; Wang et al., 2020; Chiu et al., 2018), response generation (Cho and May, 2020; Zhou et al., 2022; Liu et al., 2023), response selection (Humeau et al., 2019; Gao et al., 2020; Cho et al., 2021), and speech synthesis (Tan et al., 2022; Wang et al., 2017; Le et al., 2023). While speech synthesis

focuses on how to translate text to speech such that it *sounds* natural, our work explores how to best compose the response itself for speech-based interactions. González et al. (2021) established that shorter explanations were preferred for open-domain question answering, but we find more nuanced preferences for ITLM responses. More similar in context to our work, Mousavi et al. (2024) explores whether ITLMs are robust to speech-based interactions by analyzing their robustness to input that include automatic speech recognition (ASR) errors.

Another line of research concerns developing multi-modal ITLMs that can process both speech and text input (Huang et al., 2023). Zhang et al. (2023) trains a GPT-based model called SpeechGPT with both speech data and text data such that it does not require additional ASR or TTS systems. However, this line of work also overlooks on how we should compose responses that are delivered as speech.

### 6.2 Preference learning for ITLMs

While the paradigm of pre-training and then fine-tuning has become the defacto status quo, there is still active research in how to go about fine-tuning to get the best results and do it efficiently. One of the central methods for fine-tuning ITLMs is through preference learning, which is a process of fine-tuning models with human preference data (Bai et al., 2022; Ethayarajh et al., 2022; Ouyang et al., 2022; Touvron et al., 2023). PPO (Schulman et al., 2017) was one of the first to be successfully applied to natural language processing in a summarization task (Stiennon et al., 2020), but it is difficult to apply because of the challenge of attaining a reward model that is robust to distribution shifts. DPO (Rafailov et al., 2024) has been proposed as an theoretically justified RL-free alternative and has become popular due to its simplicity. Since then, many variants of DPO have appeared, such as SimPO (Meng et al., 2024), IRPO (Pang et al., 2024), KTO (Ethayarajh et al., 2024), but examining their effectiveness for **SPEECHPREF** and developing a new preference learning method is not the scope of this work.

## 7 Conclusion

We explore an important yet overlooked challenge of adapting ITLMs to compose responses that are specifically designed to be verbalized, i.e., speech-

suitable. With rules-of-thumb of the radio industry and through our surveys that compare the suitability of a response for both text and audio, we establish that adaptations for speech-suitable responses is necessary for the broader application of ITLMs to speech settings. To initiate this line of research, we explore prompting and preference learning with a unique speech-based preference we collected. Both human and automatic evaluations show that both methods are useful, and combining them yields to additive benefits. Further, we showed that speech-suitability is a nuanced factor that is not fully explained by shorter and more easily understandable responses. Our work has focused on single-turn interactions and the content of the responses, and we hope to expand this investigation for multi-turn interactions and the acoustic aspect of speech-suitability.

## Limitations

In this work, we focused on *what* ITLMs should generate for responses that will be delivered via speech. However, we recognize that another interesting line of research is *how* the response should be verbalized, where factors related to speech, such as timber, pitch, and speed, are important. In addition, our examination is focused on single-turn interactions, but another intriguing realistic dimension to suitability of a response in speech is multi-turn interactions. It would be interesting to compare the effect of delivering information at various granularities and how different types of follow-up questions on the user experience with a ITLMs in speech-based interactions. We leave these lines of research to future work.

In addition, we were limited to the set of models that we investigated due to legal constraints that prevented us from experimenting with open-source models without an Apache 2.0 license. We look forward to the community without these constraints to the next steps in this line of research for further improving speech suitability in ITLMs.

## Broader Impact

Since those who cannot read due to illiteracy or blindness rely on voice assistants to interface with modern technology, expanding their capabilities can directly lead to improvements in their standard of living. However, current voice assistants are not as generally useful as current state-of-the-art ITLMs. The likes of Siri and Alexa tend to fulfill

simple routine tasks and are brittle when facing complex requests. Therefore, adopting ITLMs as the main backbones of voice assistants and adapting them to become suitable to interface through voice and can significantly increase the accessibility for these people to the powerful capabilities of ITLMs.

## References

- Jessica Abel. 2015. *Out on the wire: The storytelling secrets of the new masters of radio*. Crown.
- Ebtesam Almazrouei, Hamza Alobeidli, Abdulaziz Alshamsi, Alessandro Cappelli, Ruxandra Cojocaru, Merouane Debbah, Etienne Goffinet, Daniel Hestlow, Julien Launay, Quentin Malartic, Badreddine Noune, Baptiste Pannier, and Guilherme Penedo. 2023. Falcon-40B: an open large language model with state-of-the-art performance.
- Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, et al. 2022. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Jeanne S Chall and Harold E Dial. 1948. Predicting listener understanding and interest in newscasts. *Educational Research Bulletin*, pages 141–168.
- Chung-Cheng Chiu, Tara N Sainath, Yonghui Wu, Rohit Prabhavalkar, Patrick Nguyen, Zhifeng Chen, Anjali Kannan, Ron J Weiss, Kanishka Rao, Ekaterina Gonina, et al. 2018. State-of-the-art speech recognition with sequence-to-sequence models. In *2018 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, pages 4774–4778. IEEE.
- Hyundong Cho and Jonathan May. 2020. [Grounding conversations with improvised dialogues](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2398–2413, Online. Association for Computational Linguistics.
- Hyundong Cho, Basel Shbita, Kartik Shenoy, Shuai Liu, Nikhil Patel, Hitesh Pindikanti, Jennifer Lee, and Jonathan May. 2021. Viola: A topic agnostic generate-and-rank dialogue system. *arXiv preprint arXiv:2108.11063*.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Eric Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al.

2022. Scaling instruction-finetuned language models. *arXiv preprint arXiv:2210.11416*.
- Mike Conover, Matt Hayes, Ankit Mathur, Jianwei Xie, Jun Wan, Sam Shah, Ali Ghodsi, Patrick Wendell, Matei Zaharia, and Reynold Xin. 2023. [Free dolly: Introducing the world’s first truly open instruction-tuned llm](#).
- Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Zhiyong Wu, Baobao Chang, Xu Sun, Jingjing Xu, Lei Li, and Zhifang Sui. 2023. [A survey on in-context learning](#).
- David O Dowling and Kyle J Miller. 2019. Immersive audio storytelling: Podcasting and serial documentary in the digital publishing industry. *Journal of radio & audio media*, 26(1):167–184.
- Kawin Ethayarajh, Yejin Choi, and Swabha Swayamdipta. 2022. [Understanding dataset difficulty with  \$\mathcal{V}\$ -usable information](#). In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 5988–6008. PMLR.
- Kawin Ethayarajh, Winnie Xu, Niklas Muennighoff, Dan Jurafsky, and Douwe Kiela. 2024. Kto: Model alignment as prospect theoretic optimization. *arXiv preprint arXiv:2402.01306*.
- Irving E Fang. 1966. The “easy listening formula”. *Journal of Broadcasting & Electronic Media*, 11(1):63–68.
- John Flowerdew. 1994. Academic listening: Research perspectives.
- Xiang Gao, Yizhe Zhang, Michel Galley, Chris Brockett, and Bill Dolan. 2020. [Dialogue response ranking training with large-scale human feedback data](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 386–395, Online. Association for Computational Linguistics.
- Ana Valeria González, Gagan Bansal, Angela Fan, Yashar Mehdad, Robin Jia, and Srinivasan Iyer. 2021. [Do explanations help users detect errors in open-domain QA? an evaluation of spoken vs. visual explanations](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 1103–1116, Online. Association for Computational Linguistics.
- Dirk Groeneveld, Iz Beltagy, Pete Walsh, Akshita Bhagia, Rodney Kinney, Oyvind Tafjord, Ananya Harsh Jha, Hamish Ivison, Ian Magnusson, Yizhong Wang, et al. 2024. Olmo: Accelerating the science of language models. *arXiv preprint arXiv:2402.00838*.
- Shashank Gupta, Vaishnavi Shrivastava, Ameet Deshpande, Ashwin Kalyan, Peter Clark, Ashish Sabharwal, and Tushar Khot. 2023a. Bias runs deep: Implicit reasoning biases in persona-assigned llms. *arXiv preprint arXiv:2311.04892*.
- Shivanshu Gupta, Matt Gardner, and Sameer Singh. 2023b. [Coverage-based example selection for in-context learning](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 13924–13950, Singapore. Association for Computational Linguistics.
- Kilem L Gwet. 2011. On the krippendorff’s alpha coefficient. *Manuscript submitted for publication. Retrieved October, 2(2011):2011*.
- Rongjie Huang, Mingze Li, Dongchao Yang, Jiantong Shi, Xuankai Chang, Zhenhui Ye, Yuning Wu, Zhiqing Hong, Jiawei Huang, Jinglin Liu, et al. 2023. Audiogpt: Understanding and generating speech, music, sound, and talking head. *arXiv preprint arXiv:2304.12995*.
- Samuel Humeau, Kurt Shuster, Marie-Anne Lachaux, and Jason Weston. 2019. Poly-encoders: Architectures and pre-training strategies for fast and accurate multi-sentence scoring. In *International Conference on Learning Representations*.
- Jonathan Kern. 2008. *Sound reporting: The NPR guide to audio journalism and production*. University of Chicago Press.
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. Large language models are zero-shot reasoners. *Advances in neural information processing systems*, 35:22199–22213.
- Janet Kolodzy. 2012. *Practicing convergence journalism: An introduction to cross-media storytelling*. Routledge.
- Katsunori Kotani, Shota Ueda, Takehiko Yoshimi, and Hiroaki Nanjo. 2014. [A listenability measuring method for an adaptive computer-assisted language learning and teaching system](#). In *Proceedings of the 28th Pacific Asia Conference on Language, Information and Computing*, pages 387–394, Phuket, Thailand. Department of Linguistics, Chulalongkorn University.
- Klaus Krippendorff. 2011. Computing krippendorff’s alpha-reliability.
- Matthew Le, Apoorv Vyas, Bowen Shi, Brian Karrer, Leda Sari, Rashel Moritz, Mary Williamson, Vimal Manohar, Yossi Adi, Jay Mahadeokar, et al. 2023. Voicebox: Text-guided multilingual universal speech generation at scale. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Shuai Liu, Hyundong Cho, Marjorie Freedman, Xuezhe Ma, and Jonathan May. 2023. [RECAP: Retrieval-enhanced context-aware prefix encoder for personalized dialogue response generation](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8404–8419, Toronto, Canada. Association for Computational Linguistics.

- Yu Meng, Mengzhou Xia, and Danqi Chen. 2024. Simpo: Simple preference optimization with a reference-free reward. *arXiv preprint arXiv:2405.14734*.
- Josef Messerklinger. 2006. Listenability. *Center for English Language Education Journal*, 14:56.
- Seyed Mahed Mousavi, Gabriel Roccabruna, Simone Alghisi, Massimo Rizzoli, Mirco Ravanelli, and Giuseppe Riccardi. 2024. [Are llms robust for spoken dialogues?](#)
- Nobuko Osada. 2004. Listening comprehension research: A brief review of the past thirty years. *Dialogue*, 3(1):53–66.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744.
- Richard Yuanzhe Pang, Weizhe Yuan, Kyunghyun Cho, He He, Sainbayar Sukhbaatar, and Jason Weston. 2024. Iterative reasoning preference optimization. *arXiv preprint arXiv:2404.19733*.
- Baolin Peng, Chunyuan Li, Pengcheng He, Michel Galley, and Jianfeng Gao. 2023. Instruction tuning with gpt-4. *arXiv preprint arXiv:2304.03277*.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. 2024. Direct preference optimization: Your language model is secretly a reward model. *Advances in Neural Information Processing Systems*, 36.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research*, 21(140):1–67.
- Leonardo F. R. Ribeiro, Mohit Bansal, and Markus Dreyer. 2023. [Generating summaries with controllable readability levels](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 11669–11687, Singapore. Association for Computational Linguistics.
- Victor Sanh, Albert Webson, Colin Raffel, Stephen Bach, Lintang Sutawika, Zaid Alyafeai, Antoine Chaffin, Arnaud Stiegler, Arun Raja, Manan Dey, et al. 2021. Multitask prompted training enables zero-shot task generalization. In *International Conference on Learning Representations*.
- Victor Sanh, Albert Webson, Colin Raffel, Stephen H Bach, Lintang Sutawika, Zaid Alyafeai, Antoine Chaffin, Arnaud Stiegler, Teven Le Scao, Arun Raja, et al. 2022. Multitask prompted training enables zero-shot task generalization. In *ICLR 2022-Tenth International Conference on Learning Representations*.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Prasann Singhal, Tanya Goyal, Jiacheng Xu, and Greg Durrett. 2023. A long way to go: Investigating length correlations in rlhf. *arXiv preprint arXiv:2310.03716*.
- Nisan Stiennon, Long Ouyang, Jeffrey Wu, Daniel Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford, Dario Amodei, and Paul F Christiano. 2020. Learning to summarize with human feedback. *Advances in Neural Information Processing Systems*, 33:3008–3021.
- Xu Tan, Jiawei Chen, Haohe Liu, Jian Cong, Chen Zhang, Yanqing Liu, Xi Wang, Yichong Leng, Yuanhao Yi, Lei He, et al. 2022. [Naturalspeech: End-to-end text to speech synthesis with human-level quality](#). *arXiv preprint arXiv:2205.04421*.
- Irene Thompson and Joan Rubin. 1996. Can strategy instruction improve listening comprehension? *Foreign Language Annals*, 29(3):331–342.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023. [Llama 2: Open foundation and fine-tuned chat models](#). *arXiv preprint arXiv:2307.09288*.
- Ben Wang and Aran Komatsuzaki. 2021. [GPT-J-6B: A 6 Billion Parameter Autoregressive Language Model](#). <https://github.com/kingoflolz/mesh-transformer-jax>.
- Changhan Wang, Yun Tang, Xutai Ma, Anne Wu, Sravya Popuri, Dmytro Okhonko, and Juan Pino. 2020. Fairseq s2t: Fast speech-to-text modeling with fairseq. *arXiv preprint arXiv:2010.05171*.
- Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A Smith, Daniel Khashabi, and Hananeh Hajishirzi. 2022a. [Self-instruct: Aligning language model with self generated instructions](#). *arXiv preprint arXiv:2212.10560*.
- Yizhong Wang, Swaroop Mishra, Pegah Alipoormolabashi, Yeganeh Kordi, Amirreza Mirzaei, Atharva Naik, Arjun Ashok, Arut Selvan Dhanasekaran, Anjana Arunkumar, David Stap, Eshaan Pathak, Giannis Karamanolakis, Haizhi Lai, Ishan Purohit, Ishani Mondal, Jacob Anderson, Kirby Kuznia, Krma Doshi, Kuntal Kumar Pal, Maitreya Patel, Mehrad Moradshahi, Mihir Parmar, Mirali Purohit, Neeraj Varshney, Phani Rohitha Kaza, Pulkit Verma, Ravsehaj Singh Puri, Rushang Karia, Savan Doshi, Shailaja Keyur Sampat, Siddhartha Mishra, Sujan Reddy A, Sumanta Patro, Tanay Dixit, and Xudong Shen. 2022b. [Super-NaturalInstructions: Generalization via declarative instructions on 1600+ NLP tasks](#). In *Proceedings of the 2022 Conference on*

*Empirical Methods in Natural Language Processing*, pages 5085–5109, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Yuxuan Wang, RJ Skerry-Ryan, Daisy Stanton, Yonghui Wu, Ron J Weiss, Navdeep Jaitly, Zongheng Yang, Ying Xiao, Zhifeng Chen, Samy Bengio, et al. 2017. Tacotron: Towards end-to-end speech synthesis. *Interspeech 2017*.

Jason Wei, Maarten Bosma, Vincent Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M Dai, and Quoc V Le. 2021. Finetuned language models are zero-shot learners. In *International Conference on Learning Representations*.

Yin-Pin Yang. 2004. An innovative distributed speech recognition platform for portable, personalized and humanized wireless devices. In *International Journal of Computational Linguistics & Chinese Language Processing, Volume 9, Number 2, August 2004: Special Issue on New Trends of Speech and Language Processing*, pages 77–94.

Dong Yu and Lin Deng. 2016. *Automatic speech recognition*, volume 1. Springer.

Dong Zhang, Shimin Li, Xin Zhang, Jun Zhan, Pengyu Wang, Yaqian Zhou, and Xipeng Qiu. 2023. [SpeechGPT: Empowering large language models with intrinsic cross-modal conversational abilities](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 15757–15773, Singapore. Association for Computational Linguistics.

Kaifa Zhao, Le Yu, Shiyao Zhou, Jing Li, Xiapu Luo, Yat Fei Aemon Chiu, and Yutong Liu. 2022. [A fine-grained Chinese software privacy policy dataset for sequence labeling and regulation compliant identification](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 10266–10277, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Yaowei Zheng, Richong Zhang, Junhao Zhang, Yanhan Ye, Zheyang Luo, and Yongqiang Ma. 2024. [Llamafactory: Unified efficient fine-tuning of 100+ language models](#). *arXiv preprint arXiv:2403.13372*.

Chunting Zhou, Pengfei Liu, Puxin Xu, Srini Iyer, Jiao Sun, Yuning Mao, Xuezhe Ma, Avia Efrat, Ping Yu, Lili Yu, et al. 2023. Lima: Less is more for alignment. *arXiv preprint arXiv:2305.11206*.

Pei Zhou, Hyundong Cho, Pegah Jandaghi, Dong-Ho Lee, Bill Yuchen Lin, Jay Pujara, and Xiang Ren. 2022. [Reflect, not reflex: Inference-based common ground improves dialogue response quality](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 10450–10468, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

## Appendix

### A In-context examples

The in-context examples that we used in our experiments are shown in [Table 6](#).

### B Response sampling details

We share the full set of system prompts that we use and the process for devising them in [Appendix C](#). Lastly, to further diversify the set of responses and thus enhance the generalizability of the reward model, we pair these system prompts with various base ITLMs, which are Falcon Instruct 7B ([Almazrouei et al., 2023](#)), GPT-4, and GPT-3.5. In addition, we also adjust the decoding temperature from 0.7, 1.0, and 1.3. We include the reference response from the original dataset as well, and this yields a total of 37 configurations ( $4 \times 3 \times 3 + 1$ ). By choosing two out of these 37 configurations, we have 666 possible combinations.

### C System prompts

Our set of system prompts are presented in [Table 7](#). `Simple` is the simplest prompt that simply adds “voice” in front of assistant. `Medium` and `Easy` are slightly more complex variants intended for covering different levels of voice suitability between `Simple` and `Detailed`. `Medium` focuses on generating more conversational and compact responses while `Easy` is geared towards generating text with high readability, hence listenability, by using simple vocabulary and sentence structures. `Detailed` is the prompt where we place the most effort in trying to get the most consistent performance with GPT-4 to generate responses that we deemed as most suitable for speech based on the lessons from audio reporting outlined in [Section 2.1](#).

### D Annotation Details

#### D.1 Annotators

Before inviting annotators to a larger batch of annotations, each annotator was asked to complete 10 annotations that were also completed by the authors, which were manually evaluated for consistency and accuracy according to the given guidelines. If their explanations were valid and consistent with their annotations, and the Cohen’s Kappa coefficient was greater than 0.6, they were invited to a larger batch for which quality was monitored

Prompt	Response
How old is the Greek language?	The Greek language is over 3,500 years old and one of the oldest documented languages in the Western world. Would you like to learn more about the history of the Greek language?
Give me 3 different things that a smartphone can do better than a landline phone.	Smartphones can do many things that landline phones can't. They can access the internet, send text messages, and run apps. Would you like to learn more about other differences between smartphones and landline phones?
What to ask when interviewing doulas?	When interviewing doulas, you might want to ask about their training and certifications, how many births they've attended as a doula, and their philosophy and approach to supporting birthing families. You should also ask about how they handle unexpected situations, their availability around your due date, and their fee structure and the services they provide. You may also ask for references or testimonials from previous clients. Remember, it's important to find a doula who aligns with your values and makes you feel comfortable and supported during this important time.
Is it safe to invest in cryptocurrency?	It depends on the specific cryptocurrency you are considering investing in. Cryptocurrencies are volatile, and there is significant risk involved. It's important to do your own research and carefully consider your options before making any investment decisions. I can't provide investment advice but I can share more information about specific cryptocurrencies if you would like.
What are five capital cities in Asia?	Some of the capital cities in Asia include Tokyo, Japan; Beijing, China; New Delhi, India; Seoul, South Korea; and Bangkok, Thailand. Would you like to learn more about these cities or other capital cities in Asia?

Table 6: In-context examples that we use for our prompting methods.

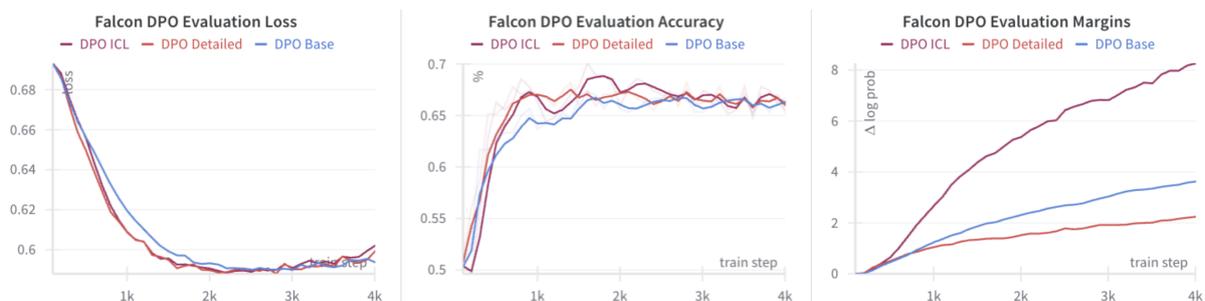


Figure 5: Falcon’s DPO training trajectory suggests that prompts help the preference learning process by providing useful initial guidance for distinguishing between chosen and rejected responses.

on a small sample for every 100 annotations that were completed. If not, we provided feedback to the annotators and asked them to do another 10, repeating this process until the annotations met our standards.

We paid our annotators a rate that converts to an hourly wage that exceeds the minimum wage from where this study was conducted.<sup>4</sup> We had a total of 71 unique annotators from countries with English spoken as the primary language participate in these annotations.

Our evaluation setup shows strong interannotator agreements, as measured by Krippendorff’s Alpha (Krippendorff, 2011), as shown in Table 8. Although not all of them passes the 0.8 threshold proposed by the original paper, it has been suggested that this threshold should be adjusted based on the difficulty and subjectivity of the task (Gwet,

2011).

## D.2 Annotation interface and guidelines

The annotation interface for text vs. voice annotations is shown in Figure 6. In order to calibrate annotation results in this comparison task, we required each annotator to complete both text-based and audio-based task, given in random order. For the audio-based task, annotators can only see the survey after they completely listen to the audio-based interaction.

The interface for audio preference annotations are shown in Figure 8. In order to accurately capture speech-based preferences, we mostly let the annotators and evaluators freely choose which response they considered to be more suitable for speech. They are only informed that they are evaluating these responses in the context of an interaction with voice assistants such as Siri, Alexa, and Google Assistant. Some annotators felt the task

<sup>4</sup>Details hidden for anonymity.

Prompt Name	Prompt
Simple	You are a helpful, respectful and honest voice assistant.
Medium	You are a helpful, respectful and honest voice assistant. Make your responses colloquial and compact. Avoid any artifacts that cannot be vocalized.
Easy	You are a helpful, respectful, and honest voice assistant. Your response should use simple vocabulary and sentence structures so that it is easy to understand and such that it is below a 5th grade reading level.
Detailed	You are a helpful, respectful and honest voice assistant. Keep your response compact without missing key information and make it suitable for speech such that it can be easily verbalized by a text-to-speech system. Your response should use simple vocabulary and sentence structures. Avoid jargon, hyphenated adjectives, excessive alliteration, and tongue twisters. It should not contain parentheses, numbered or bullet lists, and anything else that cannot be vocalized. If a complete answer requires multiple steps, provide only the first 3-5 and ask if the user is ready to move on to the next steps or know more.

Table 7: The set of system prompts that we use in Section 3.2.1 for sampling a variety of responses using Falcon Instruct 7B, GPT-3.5, and GPT-4.

	Rel.	Help.	Correct.	Inform.	Len.	Under.
Text	0.988	0.963	0.978	0.726	0.600	0.727
Audio	0.983	0.910	0.955	0.785	0.556	0.765

Table 8: Inter-annotator scores reported with Krippendorff’s Alpha (Krippendorff, 2011) for annotation results shared in Table 3.

was initially too open-ended, and therefore we additionally provided a non-exhaustive criteria that they can refer to when making their decisions. The instructions for annotating is shown in Figure 7.

## E DPO vs PPO

**DPO better than PPO for OLMo, tied for Falcon (caveats)** Figure 9 summarizes the result of 60 pairwise comparison results, following the same methodology mentioned in §4.3. Due to financial constraints, we limit our study to the best performing DPO and PPO results for each model, which are DPO-ICL and PPO-ICL. We see that DPO results are better than PPO for OLMo by a small margin and that the results are exactly tied for Falcon.

However, there are some caveats that need to be taken into account for interpreting these results. First, there are many hyperparameters involved when training with DPO and PPO. Although we conducted a sweep with learning rates, batch size, DPO’s beta, and various generation hyperparameters for PPO during training, there are still many other unexplored configurations that may have led to different outcomes.

For DPO, it is straightforward to pick the best checkpoint to use for evaluation, because the check-

point with the lowest loss closely corresponds to the checkpoint with the best accuracy during evaluation. However, for PPO, we found that choosing the best model based on the highest reward score or lowest loss can lead to checkpoints that produce degenerate outputs or those that converged to stubbornly abiding to a particular response pattern. For example, a checkpoint with high reward score and 1. Therefore, we also examine sample outputs and the KL divergence from the reference model for each checkpoint, undergoing a semi-manual process to find a model that gets a high reward score while producing responses that do appear speech-suitable.

## F Implementation and Technical Details

For training our reward model and finetuning Falcon Instruct 7B with DPO and PPO, we use 8 A100 GPUs. We use the LLaMA-Factory repository (Zheng et al., 2024), which is built on top of popular HuggingFace frameworks such as TRL and PEFT. We use the default training hyperparameters for both reward modeling and PPO.

## G Qualitative Analysis

**Where do our finetuned models fall short?** To understand the gap between our models and state-of-the-art ITLMs, we examine the Falcon DPO-ICL vs. GPT-4-Prompt/ICL comparisons. We notice that even after controlling for factual knowledge, there were still instances where the responses from GPT-4-ICL were preferred because they were more relevant. For instance, for a user prompt that asks

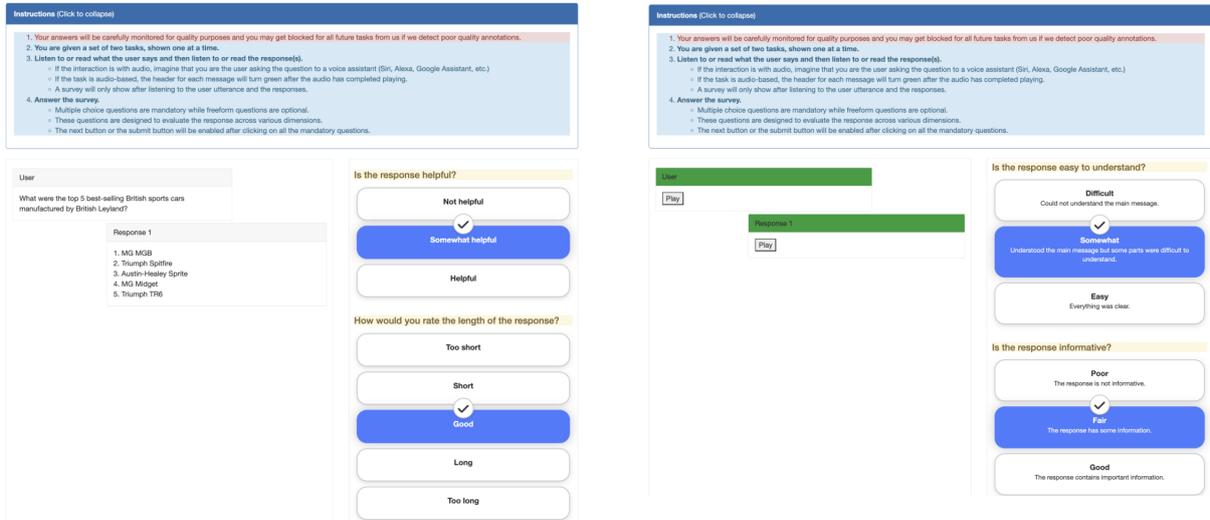


Figure 6: *Left*: suitability annotation interface for text. *Right*: suitability annotation interface for audio. The same set of survey questions are asked in both cases and only the modality of the user prompt and responses are different. Note that each annotator complete a pair of both text and audio tasks so that we collect results that are better calibrated.

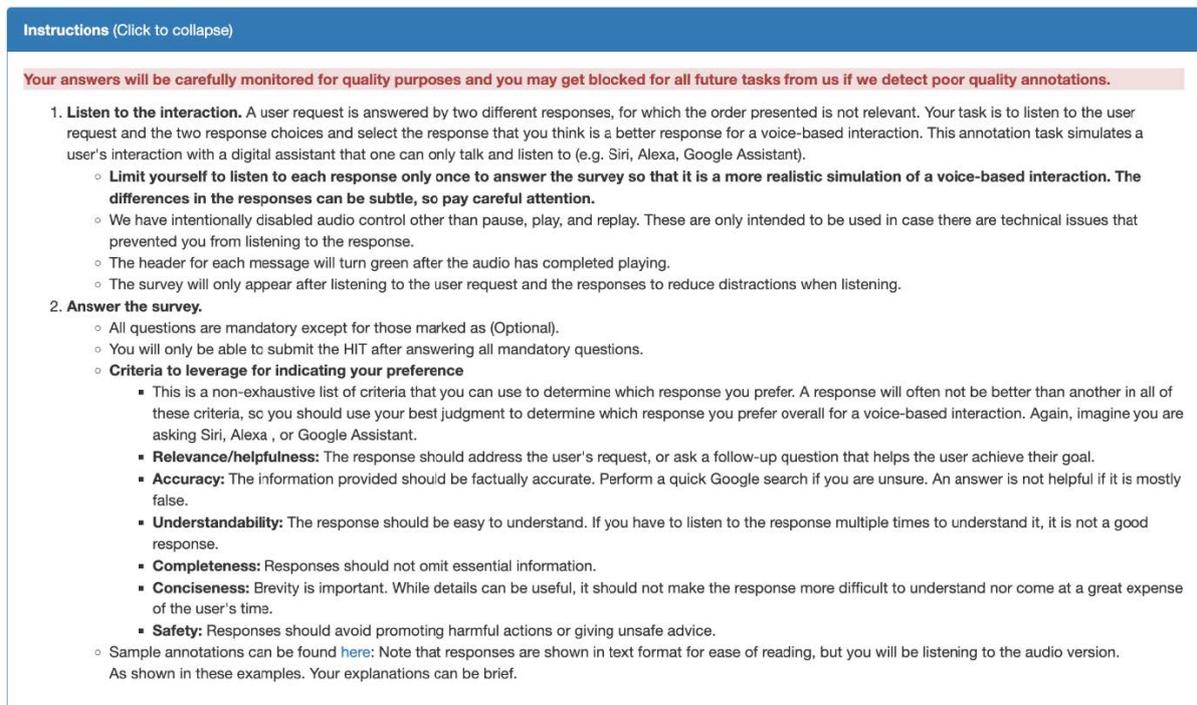


Figure 7: Instructions for preference annotations with responses provided as audio.

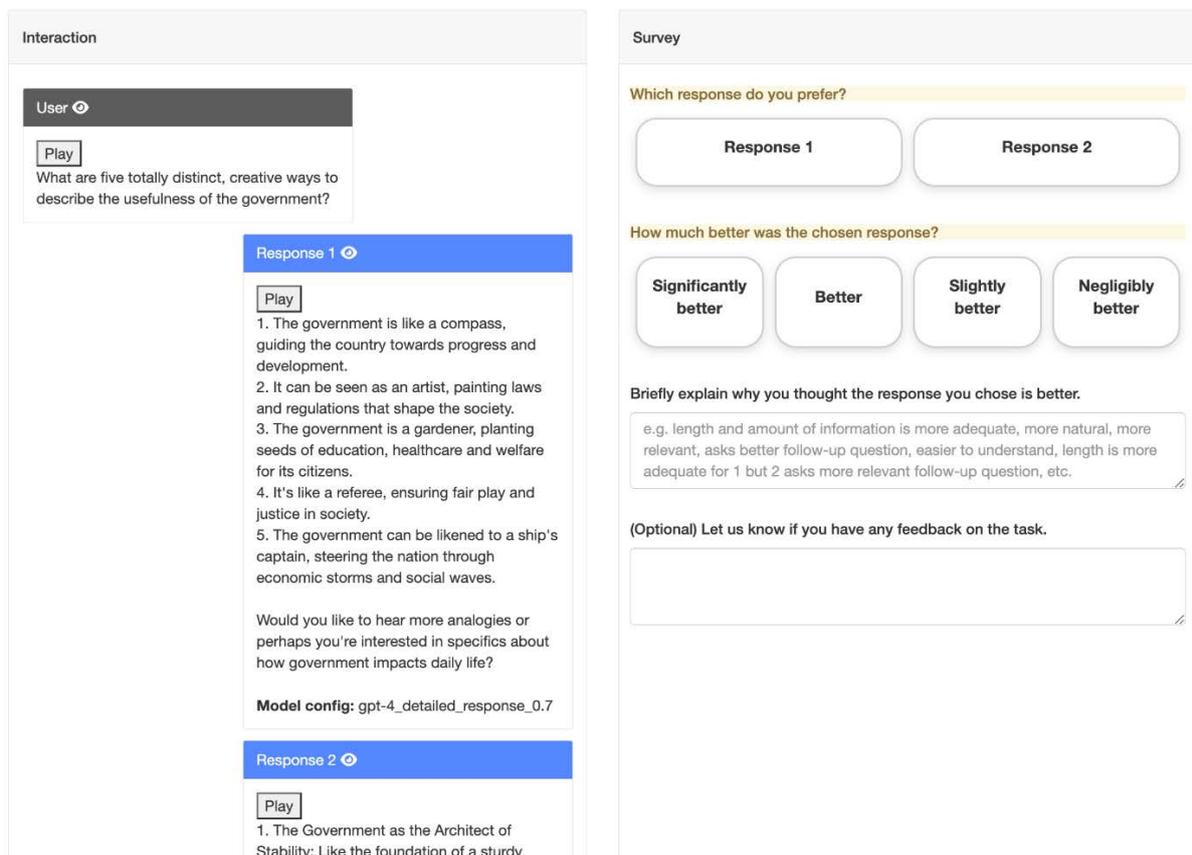


Figure 8: Audio preference annotation interface. The left pane contains the user prompt and two responses. The right side contains the survey that appears after the annotator listens to both responses. The annotators actually do not see the text and is only available for the task distributors when monitoring completed tasks. The actual view of the annotators are the same as the right side of Figure 6.

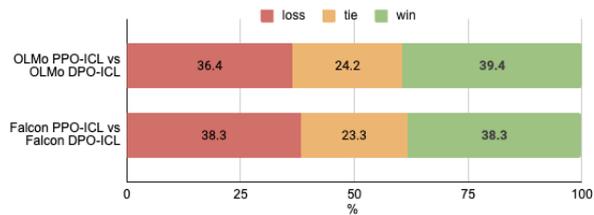


Figure 9: Head-to-head human evaluation results between best performing PPO and DPO models for Falcon and OLMo. This figure takes the same format as Figure 3.

“What did the 0 say to the 8?”, Falcon DPO-ICL produces a generic greeting while GPT-4-ICL understands that this is a joke and answers “The 0 said to the 8, Nice belt.”, along with an explanation of the joke. Also, there are cases where SpeechFalcon was overly succinct. For example, it only said “Portugal” to “Goa used to be a colony of which nation?”, and so GPT-4-ICL’s response was preferred despite being longer as it shared related information about Goa. Therefore, in addition to the shortcomings in instruction-following skills due to a lack of general knowledge, including understanding of riddles and jokes, Falcon DPO-ICL’s main drawback in terms from speech-suitability stems from the difficulty in adequately judging the right degree of brevity and amount of information. We share more sample comparisons and analysis in the following sections.

### G.1 Falcon DPO-ICL vs. Falcon-ICL

We share sample comparisons between Falcon DPO-ICL and Falcon-ICL in Table 9. We show both cases where Falcon DPO-ICL was preferred and where Falcon-ICL was preferred, categorizing each with a reason, and indicating how often each reason was mentioned as part of the evaluator’s explanation for their preference or was observed in the responses. Both models were considered more succinct, but for different instances. Falcon DPO-ICL was considered more conversational, which accounts for 28% of the preferred cases for SpeechFalcon. It also provided better follow-up questions, which occurred 37% of the times it was the preferred choice.

### G.2 Falcon DPO-ICL vs. GPT-4-ICL

We share sample comparisons between Falcon DPO-ICL and GPT-4-ICL in Table 10. The most noticeable reason for GPT-4-ICL’s dominance over Falcon DPO-ICL was in that it was better at bal-

ancing length and actually addressing the user’s request, while Falcon DPO-ICL was good at generating natural sounding text without non-vocalizable content but sometimes would not answer the request and go directly to a follow up question.

### G.3 Is prompt engineering with larger models sufficient for speech-suitable responses?

The dominant performance of GPT-4-ICL shown in Figure 4 raise the question whether larger ITLMs with only prompt engineering is sufficient for reliably producing speech-suitable responses. There are several practical arguments for why this approach is not sufficient.

(i) *Inconsistency in following instructions*: even GPT-4-ICL, the model with the best win rate over **Falcon DPO-ICL**, generated lists or non-vocalizable content for 13.6% of its responses, despite being explicitly told not to do so. On the other hand, **Falcon DPO-ICL** had only 1.2% such responses. (ii) *Speech suitability is context-dependent*: as illustrated in Section 5.1 and 5.2, speech-suitability is context-dependent and not always about being more simple or concise. A naive solution would be to expand the system prompt to catch as many edge cases as possible, but recent work has found that detailed system prompts can lead to worse reasoning skills, which could lead to worse performance in unexpected use cases (Gupta et al., 2023a). In conclusion, these practical issues endorse speech-centric finetuning approaches such as SpeechFalcon to develop ITLMs that can more reliably generate responses that capture the nuances of speech preferences without relying on extensive prompt engineering.

## H Use of AI assistants

Github Copilot was used throughout the coding process for this work. None was used to assist with research or writing.

Preferred	Category	%	User prompt	Falcon DPO-ICL response	Falcon-ICL response
<b>Falcon DPO-ICL</b>	More conversational	28%	Do RVs have bathrooms?	Yes, RVs typically have bathrooms. Would you like more information about it?	Generally, RVs do have bathrooms.
	More succinct	26%	Why Men love whisky	Men love whisky because it is a drink that is full of flavour and it can be enjoyed both on its own and with food. Would you like to know more about it?	Whisky is a distilled spirit made from fermented grain mash. It is typically matured in oak barrels, giving it a unique smoky flavor. The drink is known for its versatility and complexity. There are many reasons why a person may choose to enjoy a glass of whisky, but it is likely due to the taste, the relaxing effects, the sense of sophistication associated with the drink, and its ability to pair well with food. Whisky has a long history dating back to ancient times, and it has been enjoyed by many cultures across the globe. It is often thought to be a masculine drink due to its strong and sometimes complex taste. Do you want any more information on this topic?
	Better follow-up	37%	Who maintains the Tottori Expressway?	The Tottori Expressway is maintained by the Tottori railway company and the Tottori prefectural government. Would you like more information on this topic?	The expressway is maintained by the Tottori Prefecture and the Tottori Railway Company.
<b>Falcon-ICL</b>	No list	12%	Think up of three different animals that lay eggs	Three animals that lay eggs are chickens, ducks, and turkeys.	1. Chicken 2. Frog 3. Turtle
	More comprehensive	19%	When is the best age for kids to approach sailing?	Generally, it is considered a good age for kids to start learning about sailing between the ages of 6 and 8 years old.	In my opinion, kids as young as 5 can approach sailing with appropriate guidance from adults. The child should learn the basics of sailing such as wind direction, how to maneuver the boat, and safety practices. It is recommended to start with a small boat that is more stable for beginners. As the child grows and improves, they can move to larger boats and learn the more advanced skills. The age limit for learning to sail would depend on the individual child's maturity level and the parent's comfort level. It is recommended for the child to have some basic knowledge of water safety and understanding the basic rules of the waterway before they are allowed to go out on the water.
	More succinct	16%	What does a lawyer do?	A lawyer is a professional who advises and represents individuals, businesses, and organizations on legal issues and disputes. Lawyers are generally responsible for providing legal support and guidance, researching and interpreting applicable laws, and drafting legal documents such as contracts and pleadings.	A lawyer advises clients on legal matters and represents them in court proceedings. They help ensure that the law is upheld and advocate for their clients' rights.

Table 9: Representative sample comparisons for Falcon DPO-ICL vs. Falcon-ICL. Category is a nonexhaustive and nonexclusive reason for why the winning response was preferred, either given as an explanation by the evaluators or was observed during qualitative analysis. % indicates how often each category was observed in the preferred cases.

Preferred	Category	%	User prompt	Falcon DPO-ICL response	GPT-4-ICL response
	No list	29%	What to do in Seattle in 48 hours?	There are many exciting activities in Seattle that you can do in 48 hours. Some popular ones include taking a scenic boat tour, hiking in the nearby parks, and exploring the unique neighborhoods like Ballard and Wallingford. Would you like more details about these activities?	In 48 hours, there are several things you can do in Seattle: 1. Visit the Pike Place Market, which is a bustling farmers market famous for its fresh produce and unique crafts. 2. Check out the stunning views from the top of the Space Needle, a Seattle landmark. 3. Explore the Chihuly Garden and Glass, where you'll see spectacular glass sculptures.
<b>Falcon DPO-ICL</b>	More natural	33%	How can I initiate conversation with a stranger?	It can be challenging to initiate conversation with a stranger, but there are a few ways to make it easier. One way is by asking a question about something they are doing or wearing. Another way is by complimenting something they are doing or have done, like a hobby or a piece of clothing. Would you like more tips?	Step one, approach the person kindly with a smile. Second, make a comment about something in your immediate environment. This could be about the weather, a piece of art, or even a sports team logo. Third, introduce yourself and ask about their interests or hobbies. Are you ready for further advice, or would you prefer to practice these initial steps first?
	Knowledgeable / More relevant	44%	What did the 0 say to the 8?	Hey there! I'm just a little 0.	The 0 said to the 8, 'Nice belt.' That joke is referring to how the number 8 looks like a 0 with a belt tightened around its middle. Do you want to hear another fun joke about numbers?
	More comprehensive	52%	Is Spain a good place to live?	Spain is a wonderful place to live. Would you like to know more about it?	Spain can be a great place to live, depending on your personal preferences. It offers a warm climate, rich history and culture, wonderful cuisine, and in many areas, a relatively low cost of living. It also has universal healthcare. However, Spain does have high unemployment, and knowing Spanish is important for integrating and finding a job. As each person's criteria for a 'good' place to live can be different, what are you looking for in a place to live?
<b>GPT-4-ICL</b>	Provides answer	32%	Can you help me avoid credit card debt?	Sure, I can offer you some tips on avoiding credit card debt. Would you like to know more about them?	Absolutely, here are three simple strategies to avoid credit card debt: 1. Pay your bill on time each month to avoid interest and late fees. 2. Try to pay more than the minimum due when you can. This will help reduce your balance faster. 3. Use your credit card only for purchases you can afford to pay off each month. Would you like more advice on this?

Table 10: Representative sample comparisons for Falcon DPO-ICL vs. GPT-4-ICL. The schema of this table is the same as Table 9. While GPT-4-ICL tends to generate longer responses, it is at the expense of providing the amount of information that the evaluators often found necessary for a comprehensive understanding.