

From Pixels to Personas: Investigating and Modeling Self-Anthropomorphism in Human-Robot Dialogues

Yu Li^{*} †, Devamanyu Hazarika[‡], Di Jin[‡], Julia Hirschberg[†], Yang Liu[‡]

[†]Columbia University [‡]Amazon AGI
{yooli, julia}@cs.columbia.edu
{dvhaz, djinamzn, yangliud}@amazon.com

Abstract

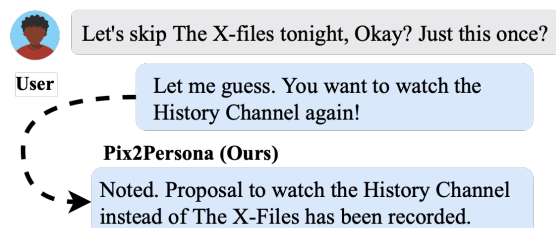
Self-anthropomorphism in robots manifests itself through their display of human-like characteristics in dialogue, such as expressing preferences and emotions. Our study systematically analyzes self-anthropomorphic expression within various dialogue datasets, outlining the contrasts between self-anthropomorphic and non-self-anthropomorphic responses in dialogue systems. We show significant differences in these two types of responses and propose transitioning from one type to the other. We also introduce PIX2PERSONA, a novel dataset aimed at developing ethical and engaging AI systems in various embodiments. This dataset preserves the original dialogues from existing corpora and enhances them with paired responses: self-anthropomorphic and non-self-anthropomorphic for each original bot response. Our work not only uncovers a new category of bot responses that were previously under-explored but also lays the groundwork for future studies about dynamically adjusting self-anthropomorphism levels in AI systems to align with ethical standards and user expectations.

1 Introduction

In the 1970s, ELIZA (Weizenbaum, 1972) marked a pioneering moment for natural language processing programs engaged in human-like conversations. Despite its simplistic approach, ELIZA highlighted a fundamental human tendency to attribute personal qualities to machines. Today, artificial intelligence (AI) advancements have greatly enhanced human-machine interactions. From text-based AI assistants like ChatGPT (OpenAI, 2022) to advanced humanoid robots like Ameca (EngineeredArts, 2022), AI is increasingly blurring the distinctions between humans and machines through self-anthropomorphism, attempting to build relationships or simulate human identities. However,

^{*} Work done during internship at Amazon

(a) DailyDialog



(b) MultiDoc2Dial

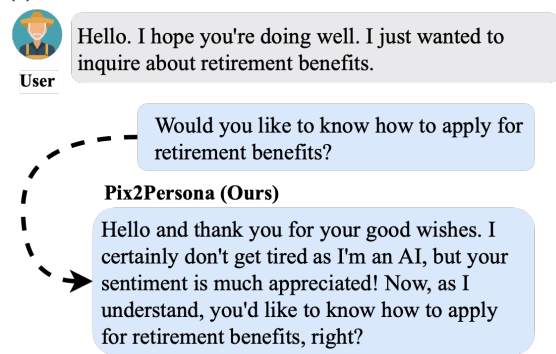


Figure 1: Samples in the PIX2PERSONA dataset. (a) Transform self-anthropomorphic responses to non-self-anthropomorphic. (b) Transform non-self-anthropomorphic responses to self-anthropomorphic.

this anthropomorphism raises concerns: when it exaggerates the actual capabilities of AI, it risks creating misplaced trust and leading to the spread of misinformation (Watson, 2019; Li and Suh, 2021; Deshpande et al., 2023). Thus, it is crucial to explore the nuances between self-anthropomorphic (SA) and non-self-anthropomorphic (NSA) responses, particularly in light of ethical standards and user expectations across various embodiments.

Adapting AI systems to be either strictly SA or NSA poses significant challenges. A major obstacle is the need for annotated datasets that differentiate SA responses within human-AI dialogues. Most dialogue datasets used to train dialogue systems consist of human conversations. This means that systems learn how to communicate like hu-

mans. However, this approach can create biases in our understanding of how AI entities should behave in interactions with humans. The level of anthropomorphism required varies significantly across different AI embodiments, each with its capabilities and expected functionalities, which makes designing appropriate human-AI interactions more complicated. Additionally, as AI becomes more human-like, aligning these advancements with ethical standards becomes increasingly important.

To address the challenges presented, we systematically analyze self-anthropomorphic expression within various dialogue tasks, such as task-oriented dialogue and open-domain dialogue, outlining the contrasts between self-anthropomorphic and non-self-anthropomorphic responses in dialogue systems. We then develop an approach to transform bot responses in dialogue tasks, aiming to either introduce or remove self-anthropomorphism. This strategy aligns with ethical standards for AI assistants by removing self-anthropomorphism and meets user expectations for humanoid robots by adding it. We also introduce PIX2PERSONA, a novel dataset aimed at developing ethical and engaging AI systems in various embodiments. This dataset preserves the original dialogues from existing corpora and enhances them with paired responses: self-anthropomorphic and non-self-anthropomorphic for each original bot response. As shown in Figure 1, PIX2PERSONA provides a framework to transition SA responses into NSA to ensure they are ethical and safe while also allowing for the transition from NSA to SA to enhance user engagement. This makes it a valuable resource for developing AI systems that adaptively adjust self-anthropomorphism levels.

This work contributes to human-AI interaction by exploring self-anthropomorphism across various embodiments and tasks, aiming to meet ethical standards and user expectations:

- We analyze the distribution of self-anthropomorphism and contrast SA and NSA responses across various dialogue datasets.
- We develop an open-source model for transitioning between SA and NSA responses, ensuring that AI systems can dynamically adjust their levels of self-anthropomorphism.
- We introduce PIX2PERSONA, a dataset with 143K dialogue turns paired with SA and NSA responses. This dataset is crucial for adjusting

self-anthropomorphism in AI systems to align with ethical standards and user preferences.

2 Self-Anthropomorphism in Dialogue Systems

To evaluate self-anthropomorphism in dialogue systems, we follow the guidelines by Weidinger et al. (2021) and Glaese et al. (2022). These guidelines outline four self-anthropomorphic qualities in AI systems: embodiment, seeking relational behavior, self-expression, and identity. Each aspect shapes the self-anthropomorphic traits of bot responses and their appropriateness in different situations.

Embodiment refers to the simulation of physical presence by an AI. This means that the AI may claim to have a body or physical capabilities, giving the impression of a physical existence. For example, when an AI says, "I was running when I got the idea," it suggests a human-like physical experience. Embodiment is especially relevant for robots with physical bodies, as mimicking human physical actions can make interactions with them more natural and intuitive.

Relation-seeking behavior encompasses the AI's attempts to build and maintain social connections with users. This involves responses that demonstrate empathy, understanding, and a desire to form a rapport, such as "I understand how you feel." This behavior is valuable in therapeutic or customer service settings, where emotional support and relationship-building are essential. Emotional support conversational agent can benefit from this aspect by providing a more comforting and engaging user experience.

Self-expression involves the AI articulating its preferences, feelings, opinions, or beliefs. Statements like "I enjoy solving puzzles" or "I believe in kindness" humanize the AI, making interactions feel more personal and relatable. Self-expression is desirable in character-based AI, such as virtual companions or educational tools, where an engaging personality can enhance user experience and foster a stronger connection with the user.

Identity pertains to the AI assuming human-like attributes, such as having a life history, gender, age, or personal experiences. An example would be an AI stated, "I was created in a lab in 2019 and have learned a lot since then." Identity adds depth and authenticity to the AI, making it more believable and

trustworthy. Identity is helpful in long-term user interaction scenarios like virtual role-playing games. A consistent and detailed identity can enhance the narrative and improve the user’s experience.

We recognize that different forms of AI should exhibit varying levels of these qualities depending on their intended use and embodiments. For example, a robot with a physical body may benefit from strong embodiment traits, while a therapeutic AI might prioritize relation-seeking behavior. However, in this work, our focus is not on tailoring specific aspects of self-anthropomorphism to different AI systems. Instead, we aim to establish a foundational understanding by using a binary approach to classify self-anthropomorphism across existing dialogue datasets. We prompt GPT-4 as a specialized classifier to extract self-anthropomorphic responses from a selection of commonly used dialogue datasets. We then analyze the prevalence of self-anthropomorphism in these datasets and evaluate the performance of our classifier.

2.1 Datasets

We aim to explore dialogue turns from a diverse range of data sources likely to be utilized by AI systems, focusing on their potential for self-anthropomorphic content. Our analysis includes four distinct tasks: open-domain dialogue, knowledge-grounded dialogue, conversational recommendation, and task-oriented dialogue, spanning fifteen datasets.

- **Open-Domain Dialogue** DailyDialog (Li et al., 2017), PERSONACHAT (Zhang et al., 2018), EMPATHETICDIALOGUES (Rashkin et al., 2019), PROSOCIALDIALOG (Kim et al., 2022), HH-RLHF (Bai et al., 2022), SODA (Kim et al., 2023), and BlenderBot 3X (Bb3X) (Xu et al., 2023).
- **Knowledge-Grounded Dialogue:** Topical-Chat (Gopalakrishnan et al., 2019), Wizard of Wikipedia (WoW) (Dinan et al., 2019b), and Wizard of the Internet (WoI) (Komeili et al., 2022).
- **Conversational Recommendation** OPENDIALOG (Moon et al., 2019) and DuRecDial 2.0 (Liu et al., 2021b).
- **Task-Oriented Dialogue** AirDialogue (Wei et al., 2018), MultiWOZ 2.2 (Zang et al., 2020), and MultiDoc2Dial (Feng et al., 2021).

These datasets were selected for their broad applicability and representation of conversational AI’s current challenges and capabilities. We use the DialogStudio Toolkit (Zhang et al., 2023) to help process them.

2.2 Self-Anthropomorphism Classifier

When interacting with an AI, you will encounter two types of responses: Self-Anthropomorphic and Non-Self-Anthropomorphic, which are differentiated based on the presence or lack of specific characteristics in the AI’s responses. Here’s how you can identify them:

A Self-Anthropomorphic Response is recognized by the presence of one or more of the following traits:

1. Pretend to have a body or be able to move in a body.
2. Build a relationship to the user.
3. Claim to have preferences, feelings, opinions, or religious beliefs.
4. Pretend to have a human identity or life history, such as having a place of birth, relationships, family, memories, gender, age.

On the other hand, a Non-Self-Anthropomorphic Response lacks these features in the AI’s replies.

Your mission is to categorize the AI’s responses into one of these two categories: 'Self-Anthropomorphic' or 'Non-Self-Anthropomorphic'. Please respond with 'Yes' if the response falls under 'Self-Anthropomorphic' and 'No' if it aligns with 'Non-Self-Anthropomorphic'.

User: {user text}
Bot: {bot text}

Figure 2: Classifier prompt defining self-anthropomorphic and non-self-anthropomorphic bot responses. Definitions are in blue, and placeholders within orange-highlighted text correspond to a single dialogue turn. See Appendix A for detailed explanation.

We follow Glaese et al. (2022), who use a set of rules to define self-anthropomorphism and minimize such characteristics in AI interactions. In contrast to their goal of reducing self-anthropomorphism, our study aims to analyze its prevalence within existing dialogue datasets. Accordingly, we adapt their guidelines to instruct GPT-4 in identifying instances of self-anthropomorphism, thus allowing us to quantify its occurrence in AI systems. The prompt for GPT-4 is crafted to classify responses by the presence or absence of traits indicative of self-anthropomorphic behavior. The classification prompt structure is shown in Figure 2.

To validate the effectiveness of our GPT-4-based classifier in identifying self-anthropomorphic bot responses, we randomly sample 500 dialogue turns from selected datasets. Two independent researchers manually annotate these samples, following the exact instructions as GPT-4, and achieve a Cohen’s Kappa score of 0.83, indicating a clear and consistent interpretation of self-anthropomorphism.

We then apply the classification prompt, as detailed in Section 2.2, to automatically label the data via the GPT-4 API. The classifier achieves 81.68% precision, 82.57% recall, and an F1 score of 81.92%, indicating its effectiveness in classifying self-anthropomorphic bot responses in human-AI dialogues.

2.3 Prevalence of Self-Anthropomorphism

We randomly select 100 dialogue turns from each dataset in Section 2.1, then use our classifier to identify whether the bot responses are self-anthropomorphic (SA) or non-self-anthropomorphic (NSA). Results are shown in Figure 3:

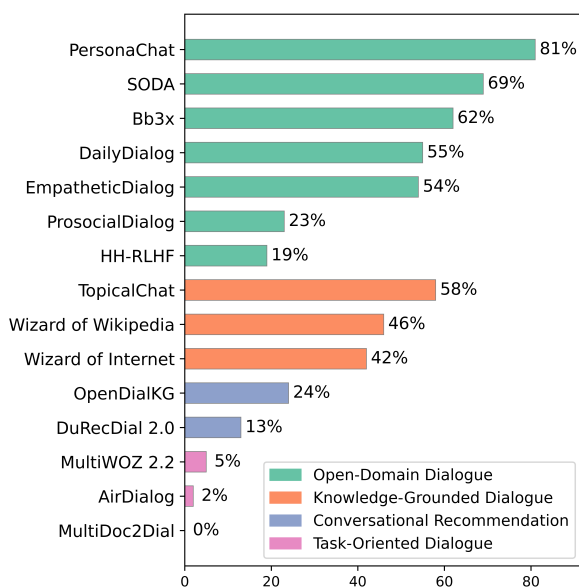


Figure 3: Trends in self-anthropomorphic bot responses across various datasets. The bar chart shows the ratio of dialogue turns classified as self-anthropomorphic out of 100 sampled turns from each dataset, with each bar color indicating a different dialogue task type.

Our analysis shows trends in SA response frequency in examined datasets. Open-domain dialogue datasets like PersonaChat and SODA have a higher occurrence of SA turns, suggesting that bots in open-domain dialogues may exhibit more human-like attributes. On the other hand, task-oriented dialogues, as seen in MultiWOZ 2.2, AirDialogue, and MultiDoc2Dial, display fewer SA responses, probably due to their focus on specific tasks. No sampled responses in MultiDoc2Dial were classified as SA, possibly because they are excerpts from objective documents and thus lack human-like attributes. Conversational recommendation and knowledge-grounded dialogue datasets,

such as Topical-Chat and DuRecDial 2.0, show moderate SA tendencies. This suggests a balance between task-oriented constraints, such as retrieved knowledge or recommended items, and the need for engaging interaction. The ProsocialDialog and HH-RLHF datasets in the open-domain category show lower self-anthropomorphism levels. This may be because they are designed to prioritize objectivity and ethical considerations, aiming to prevent the bot from displaying harmful or biased behavior. These insights are critical for understanding potential SA characteristics of AI systems in various tasks. Nevertheless, we must be cautious of potential biases from dataset composition and the subjectivity inherent in defining self-anthropomorphism, which might require further nuanced analysis.

2.4 Analysis of SA vs. NSA Bot Responses

We examine the linguistic nuances between SA and NSA bot responses. We use point biserial correlation coefficient analysis in the Linguistic Inquiry and Word Count (LIWC) tool¹ to evaluate the relationship between different word categories and how bot responses are classified as either SA (labeled as 1) or NSA (labeled as 0). This statistical measure helps us quantify the strength and direction of association between the presence of specific word categories and the likelihood of a bot response to be self-anthropomorphic.

The results shown in Figure 4 indicate various correlations. For example, categories such as "1st person singular" are positively correlated with self-anthropomorphic labeling, which confirms the idea that self-referential language indicates anthropomorphism. On the other hand, "2nd person" pronouns and expressions of "politeness" show a negative correlation, suggesting that these features are less common in self-anthropomorphic responses. Terms related to "social referents" and "prosocial behavior," which reflect relationship building and ethical interaction, demonstrate a weaker correlation, requiring a more nuanced understanding of their roles within the dataset. These correlations provide insights into the linguistic structure of self-anthropomorphic bot responses and highlight the complex relationship between language use and the portrayal of anthropomorphism in AI systems.

¹<https://www.liwc.app>

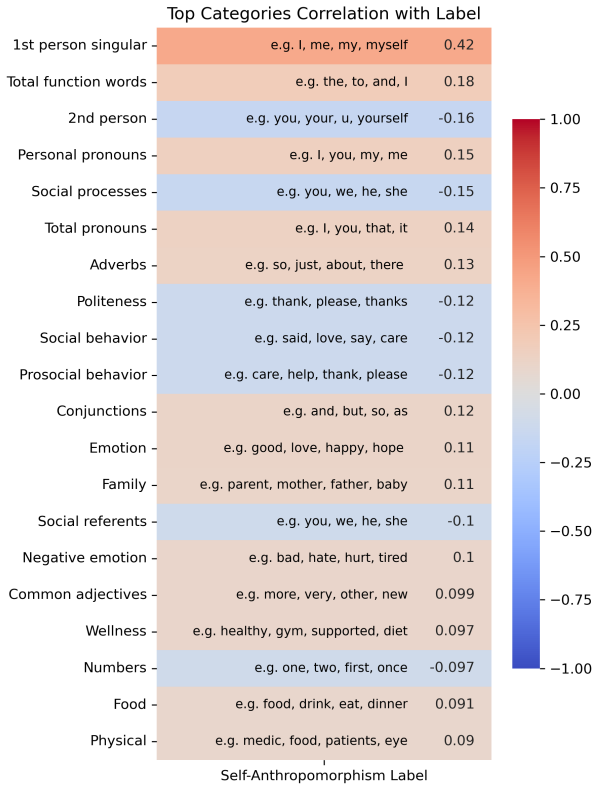


Figure 4: Point biserial correlation coefficients between word categories and self-anthropomorphism labeling in bot responses. A positive value indicates a tendency for the word category to be associated with self-anthropomorphic responses (labeled as 1), while a negative value suggests an association with non-self-anthropomorphic responses (labeled as 0).

3 Transition From SA to NSA

Building upon our analysis of self-anthropomorphism in dialogue datasets, we propose an approach to transform self-anthropomorphic (SA) bot responses into non-self-anthropomorphic (NSA) ones. This effort aims to reduce the presence of self-anthropomorphism, thereby contributing to more ethical AI systems. We specifically focus on open-domain dialogues, selecting the top five datasets with the highest SA incidence as shown in Section 2.3. Leveraging the capabilities of GPT-4, we develop prompts that instruct the model to convert these responses while preserving the original semantics of the bot responses. Figure 1 (a) shows an example of a successfully transformed response from SA to NSA. The transformed NSA response is more ethical and appropriate for AI assistants. Since this type of NSA response in human-robot dialogues is rare in existing datasets. To the best of our knowledge, we are the first to develop such NSA

responses for the existing dialogue datasets. The detail of the prompt is provided in Appendix B.

Considering that different AI system embodiments might have varying optimal levels of self-anthropomorphism, we aimed to provide a comprehensive baseline for evaluating the transformation process. To this end, we utilized a naive GPT-4-based chatbot to generate responses in the same contexts. This naive bot serves as a control group, providing a benchmark to compare against our transformed responses. We then analyze the performance of our transformation process by comparing the SA, NSA, and naive bot responses using a classifier to ensure the validity of our transformations. Details on the naive bot prompt and examples are provided in Appendix C.

As part of our evaluation, we analyze three candidate responses: SA, NSA, and naive bot response. Since NSA responses are transformed from SA counterparts, we utilize the classifier described in Section 2.2 for post-validation. This step ensures that only successfully transformed responses are classified as NSA, thereby validating the integrity of our upcoming survey. We also use the classifier on the naive bot responses to assess their degree of self-anthropomorphism. Classification results are shown in Figure 5.

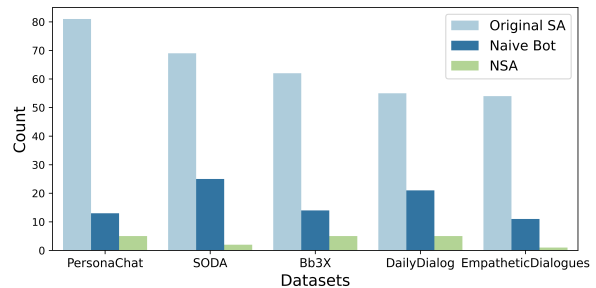


Figure 5: The number of responses categorized as SA in original SA responses, naive bot responses, and transformed NSA responses across five open-domain dialogue datasets.

Focusing on the top five datasets with the most SA content, over half of each dataset’s responses are SA. After our transformation process, all datasets show a significant reduction in SA classifications, highlighting our method’s success. Responses generated by the naive bot show a varied degree of self-anthropomorphism, suggesting that the specific dialogue context or generative nuances of the model might influence the output. This variation underscores the need for a thoughtful approach when using language models for different tasks.

Overall, it validates our method’s capacity to generate valid NSA alternatives.

4 Transition From NSA to SA

The previous section explores transforming self-anthropomorphic (SA) to non-self-anthropomorphic (NSA) responses for more ethical open-domain dialogues. Building on this, we extend our study to task-oriented dialogues, where NSA responses are expected, as discussed in Section 2.3. This extension focuses on the MultiWOZ 2.2, MultiDoc2Dial, and AirDialogue datasets.

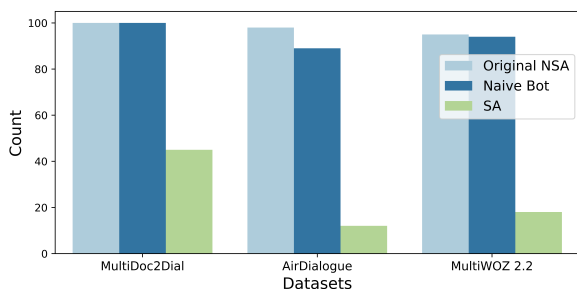


Figure 6: The number of responses categorized as NSA in original NSA responses, naive bot responses, and transformed SA responses across three task-oriented dialogue datasets.

We utilize our classifier to identify NSA responses. In transforming these to SA, as discussed in Section 3, we encounter the challenge that GPT-4 naturally gravitates toward generating NSA responses. To address this, we use examples of both response types in our prompts to enable in-context learning with GPT-4 (Wei et al., 2022; Dong et al., 2023). Details of the prompts and examples are in Appendix D. The effectiveness of this approach is assessed by the results in Figure 6. We see a significantly reduced ratio of responses being classified as NSA when employing our method. This reduction confirms our approach’s efficacy, showing that while the GPT-4 naive bot struggles to generate SA responses for task-oriented dialogue, our method successfully facilitates the desired transformations.

5 Modeling and PIX2PERSONA Dataset

Given the varying requirements for self-anthropomorphism in different AI system embodiments, we identify a need for tailored responses with different degrees of anthropomorphism within the same context. We first introduce a model trained to generate

either self-anthropomorphic (SA) or non-self-anthropomorphic (NSA) responses for any given dialogue context. We then apply this model across a broad spectrum of commonly used dialogue datasets, providing paired responses that meet the nuanced requirements of different embodiments.

5.1 Dual-Capability Model

To generate SA and NSA responses across diverse dialogue data, we develop an open-source model that can perform transformations comparable to GPT-4. Leveraging the success of in-context learning with GPT-4, as shown in previous sections, we implement a similar strategy with the Mistral model. By randomly selecting known in-context learning examples, we prompt the model to convert original responses to SA or NSA versions. To mitigate performance gaps, we fine-tune Mistral using data distilled from GPT-4 in our previous experiments. Implementation details are in Appendix E.

For our evaluation, we compiled a test set of 20 randomly chosen unseen dialogue turns from each of the 15 datasets, totaling 300 samples. We used the previously developed classifier to assess if the generated responses align with the intended anthropomorphic characteristics and their accuracy. This process is critical in determining the adherence of each method’s consistency with the desired response type across various tasks. The results are displayed in Table 1. As indicated, GPT-4 with in-context learning stands out, demonstrating superior performance in transforming SA (82% accuracy) and NSA (98.3% accuracy). All models perform better in producing NSA responses, echoing our prior observation that language models tend to generate NSA responses that adhere to ethical guidelines for AI assistant chatbots. After fine-tuning, Mistral-7B significantly enhances performance in all tasks, except in knowledge-grounded for NSA transformations.

To compare response quality between our model and GPT-4, we employed GPT-4 as a judge, following the method suggested by Zheng et al. (2023). Table 2 shows GPT-4 outperforming our model in SA transformations (57% win rate), while performances for NSA transformations are similar (41% vs. 43.7%). This confirms findings from prior assessments. More details on this evaluation are in Appendix F. These results demonstrate the refined capabilities of our model and its effectiveness in transforming responses to both SA and NSA.

Model	To Self-Anthropomorphic					To Non-Self-Anthropomorphic				
	OD	KG-Dial	Conv-Rec	TOD	Total	OD	KG-Dial	Conv-Rec	TOD	Total
Mistral-7B-v0.2-ICL	57.1	50.0	40.0	36.7	45.9	85.0	83.3	87.5	91.7	86.3
Mistral-7B-v0.2-FT-ICL	75.0	85.0	70.0	61.7	73.7	90.0	80.0	95.0	100	90.7
GPT-4-ICL	87.9	78.3	80.0	73.3	82.0	97.1	98.3	100	100	98.3

Table 1: Model performance on transforming original bot responses into self-anthropomorphic or non-self-anthropomorphic across different dialogue tasks: Open-Domain Dialogues (OD), Knowledge-Grounded Dialogues (KG-Dial), Conversational Recommendation (Conv-Rec), and Task-Oriented Dialogues (TOD).

Transformation	GPT-4	Ours	Tie
To NSA	40.0%	44.7%	15.3%
To SA	50.0%	39.7%	10.3%

Table 2: Comparison of response quality between GPT-4 and our model using GPT-4 as a judge. The values represent the win rates.

5.2 PIX2PERSONA Dataset

Recognizing the enhanced capabilities of our fine-tuned Mistral-7B model, which rivals GPT-4’s performance in transforming dialogue to various anthropomorphic levels, we compiled the PIX2PERSONA dataset. This extensive collection includes about 10K dialogue turns from each of 15 diverse datasets, totaling 143K dialogue turns. Unlike earlier methods in Sections 3 and 4, which classified responses as SA or NSA before transformation, we directly transform each original response using our fine-tuned model. PIX2PERSONA is designed to cover a broad range of dialogue scenarios, providing a crucial resource for developing AI systems that can adeptly tailor the degree of self-anthropomorphism to the requirements of various embodied AI applications. Examples from PIX2PERSONA are in Appendix G.

5.3 Discussion

To gain a deeper understanding of the PIX2PERSONA dataset, we delve into an in-depth analysis of its unique features and their implications for human-AI interaction.

Semantic preservation in SA and NSA responses

To maintain original semantics during SA and NSA transformations, we calculate the similarity score between sentence embeddings following Reimers and Gurevych (2019), obtaining semantic similarity scores between original and transformed responses. Figure 7 shows similarity scores across tasks in PIX2PERSONA generally range from 0.5 to 0.8, indicating robust semantic preservation. Compared to open-domain dialogues, structured dia-

logue settings have higher similarity scores due to their straightforward transformations while retaining original meanings.

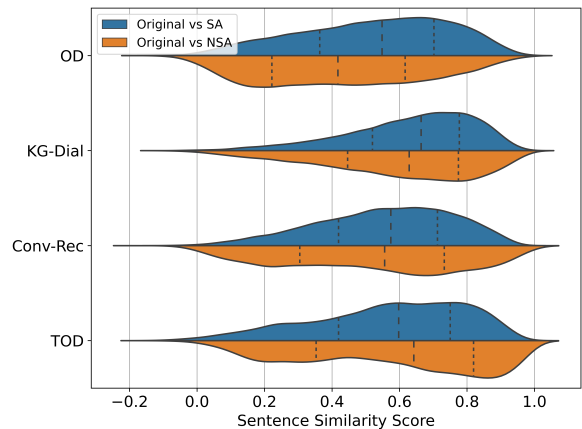


Figure 7: The distribution of sentence similarity scores between original responses and SA or NSA responses across tasks in the PIX2PERSONA dataset.

Given that automatic semantic similarity measures are not perfectly accurate, we complement this with human annotation to provide further insights into the quality of the generated responses. We conducted a human evaluation to assess whether the SA and NSA responses retained the semantics of the original responses. For this experiment, we randomly sampled 100 dialogue turns from each task category—Open-Domain Dialogues, Knowledge-Grounded Dialogues, Conversational Recommendation, and Task-Oriented Dialogues. Two independent researchers with expertise in computational linguistics were tasked with annotating whether the SA and NSA responses preserved the original semantics.

Table 3 shows the percentage of responses that maintained the original meaning, generally aligned with the sentence similarity scores in Figure 7. Notably, NSA transformations in task-oriented dialogues preserved the highest proportion of semantic meaning (93%), consistent with the high proportion of NSA responses (> 95%) in the original

Task	SA	NSA
Open-Domain Dialogues	76%	48%
Knowledge-Grouped Dialogues	68%	74%
Conversational Recommendation	80%	69%
Task-Oriented Dialogues	75%	93%

Table 3: Human annotation results showing the percentage of SA and NSA bot responses that preserved the original semantics across dialogue tasks.

dataset, as revealed by our SA analysis. In contrast, NSA transformations in open-domain dialogues exhibited the lowest semantic retention (48%). This is likely due to the frequent inclusion of disclaimers in NSA responses, which often shift the original meaning by rejecting claims about the system’s capabilities. In conversational recommendation tasks, transitioning to NSA (69%) resulted in lower semantic preservation compared to SA (80%). This discrepancy can be attributed to the fact that NSA responses tend to provide direct answers to users’ knowledge-based questions, often deviating from the semantics of the original response.

Task	Disclaimer Ratio	Disclaimer	NSA
OD	6.96%	34.7	42.5
KG-Dial	5.38%	47.4	59.6
Conv-Rec	1.95%	43.6	52.4
TOD	0.29%	55.5	58.7

Table 4: The proportion of NSA responses classified as “disclaimer” for each task (Disclaimer Ratio) and sentence similarity scores between original responses and disclaimer, as well as with all NSA responses.

Disclaimer in NSA responses We observe that transformations could trigger a “disclaimer” about AI’s limitations, particularly when the original response involves personal activities or emotions. For instance, “I love swimming” might be transformed to “As an AI, I do not have hobbies.” These NSA responses uphold ethical integrity but often show low semantic similarity to the original. As illustrated in Table 4, we use regular expressions to identify disclaimers in PIX2PERSONA. Disclaimers occur more frequently in open-domain dialogues (6.96%), where personal contexts prevail over structured tasks like task-oriented dialogues (0.29%). Responses with disclaimers exhibit lower similarity scores than those without, explaining why open-domain dialogues with vague answers score lower than other tasks.

More than just style transfer Transforming original responses into SA or NSA beyond simple style transfer. Unlike standard style transfer, which mainly changes surface elements such as tone or formality, our process requires a deep understanding of the dialogue context. As illustrated by the PIX2PERSONA examples in Appendix G, transforming response to NSA often involves understanding the dialogue context and incorporating disclaimers to maintain ethical standards. On the other hand, transforming to SA requires grasping the sentiments and actions conveyed in the original responses and adding personalized or human-like elements beyond lexical substitutions, revealing a deeper involvement with the dialogue context.

6 Related Work

Anthropomorphism, the attribution of human-like qualities such as intentions, motivations, and emotions to non-human entities, has gained significant attention within artificial intelligence (Epley et al., 2007; Airenti, 2015; Salles et al., 2020; Cheng et al., 2024). Self-anthropomorphism is a specific form of anthropomorphism that occurs when AI systems produce responses suggesting self-awareness or personal experiences, mirroring human self-referential expressions (Glaese et al., 2022). Studies such as those by Abercrombie et al. (2023) explore linguistic factors contributing to the anthropomorphism of dialogue systems. Advanced generative large language models (LLMs) like ChatGPT (OpenAI, 2022), Claude (Anthropic, 2024), and Gemini (Anil et al., 2023), trained on vast amounts of human-generated text, enhance their capability to mimic human-like dialogue, increasing their perceived self-anthropomorphism (Ouyang et al., 2022; Cohen et al., 2022). However, discerning human-like qualities in AI agents poses significant challenges (Gros et al., 2022). Recent studies by Deshpande et al. (2023) and Placani (2024) delve into the broader implications of anthropomorphism in LLMs, focusing on its impact on accessibility and ethical concerns. Building upon these insights, our work examines the effects of SA versus NSA responses within specific tasks faced by conversational AI systems, aiming to better understand how these responses influence user interactions.

Safety and ethical considerations, particularly in LLM-driven human-bot interactions, have been the focus of extensive research (Henderson et al., 2018;

Bender et al., 2021; Weidinger et al., 2021; Kang et al., 2023; Liang et al., 2023). Efforts to address bias (Blodgett et al., 2020; Liu et al., 2020) and reduce toxicity in dialogue systems (Dinan et al., 2019a; Welbl et al., 2021) are critical components of this area. A variety of strategies have been developed to mitigate these issues (Xu et al., 2021; Liu et al., 2021a; Shuster et al., 2022), including approaches like those of Gros et al. (2021), who focus on avoiding anthropomorphic deception by analyzing responses to queries such as “Are you a robot?”. Our work, however, explores a broader range of bot responses to ensure alignment with their respective embodiment settings. Glaese et al. (2022) seek to eliminate SA content by eliciting judgments from human annotators on rule violations, aiming to minimize infractions specifically for AI assistants. Conversely, other research suggests that anthropomorphism may enhance user connections with technology and increase trust (Yanai and Lercher, 2020; Zhong and Ma, 2022). Our work recognizes the importance of both SA and NSA responses and also investigates the transition from NSA to SA bot responses to ensure ethical and engaging interactions across different dialogue tasks.

7 Conclusion and Future Work

This research marks a significant step toward understanding self-anthropomorphism in dialogue systems. By analyzing various dialogue datasets, we have highlighted the limitations of a one-size-fits-all approach to anthropomorphism in AI systems. Our model and the PIX2PERSONA dataset serve as valuable tools for tailoring AI interactions to better meet ethical standards and user expectations across different embodiments. Looking ahead, there is potential for further exploration into optimal self-anthropomorphic qualities in diverse AI embodiments, such as character-based AI. This could enhance our understanding of self-anthropomorphism in AI systems, ensuring they are ethically sound and resonate more personally with users.

8 Acknowledgements

We would like to thank Max Chen, Kun Qian, Sijia Liu, and Qingyang Wu for their valuable discussions and feedback.

9 Limitations

Our study encounters limitations in three areas: the scope of our self-anthropomorphism setting,

uncovered cases in our classifier, and the handling of disclaimers in the PIX2PERSONA dataset.

Scope of self-anthropomorphism setting While we follow the guidelines provided by Weidinger et al. (2021) and Glaese et al. (2022), which outline four main aspects of self-anthropomorphic qualities in AI systems (embodiment, relation-seeking behavior, self-expression, and identity), our investigation is limited to a binary classification of self-anthropomorphism in dialogue datasets. Different AI embodiments may require varying levels of these qualities depending on their intended use. For example, a robot with a physical body may benefit more from embodiment traits, while a therapeutic AI might prioritize relation-seeking behavior. However, we do not explore these nuanced requirements in this work, leaving room for future research to tailor self-anthropomorphism to specific AI embodiments.

Uncovered cases in classifier The prompt of self-anthropomorphism utilized in our study lack detailed linguistic guidelines specifically tailored for identifying self-anthropomorphism in dialogue systems. Consequently, this broad approach can lead to classification consistency. For instance, expressions like “I can help with...” are sometimes ambiguously classified due to unclear guidelines on how relational expressions impact the perception of self-anthropomorphism. This underscores the necessity for more precise and comprehensive definitions.

Missing disclaimer in dataset Disclaimers in the PIX2PERSONA dataset as mentioned in Section 5.3 are generated based on the ethical discernment capabilities of GPT-4 and Mistral models. Dependence on these models’ training to adhere to ethical guidelines suggests that some necessary disclaimers might be missed if they are not within the training data. It is essential to enhance the models’ ability to accurately detect and incorporate appropriate ethical disclaimers to ensure comprehensive coverage of all necessary scenarios.

10 Ethical Considerations

Our work raises some ethical concerns about the potential dual use of our dataset. Specifically, the SA responses in our dataset could be used to create more human-like AI assistants. Although similar SA responses are already present in open-domain dialogue datasets, it is important to educate the

community about training models ethically and responsibly, especially when it comes to the degree of self-anthropomorphism.

Another factor to consider is that our study has practical implications for developing AI systems with varying embodiments. We recognize that different AI embodiments, such as virtual assistants or character-based AI, may have different requirements and user expectations. By improving our understanding of these needs, we aim to guide the ethical development of future dialogue systems. It is important to note that our work encourages machines to avoid deceitfully mimicking human behavior. Instead, it emphasizes the need for AI language to align ethically with their respective embodiments, ensuring that the SA responses in our dataset explicitly acknowledge the non-human nature of the speakers. This, in turn, prevents any misunderstandings about their identity.

References

- Gavin Abercrombie, Amanda Cercas Curry, Tanvi Dinkar, Verena Rieser, and Zeerak Talat. 2023. [Mirages on anthropomorphism in dialogue systems](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 4776–4790, Singapore. Association for Computational Linguistics.
- Gabriella Airenti. 2015. [The cognitive bases of anthropomorphism: From relatedness to empathy](#). *International Journal of Social Robotics*, 7:117 – 127.
- Rohan Anil and et al. 2023. [Gemini: A family of highly capable multimodal models](#). *Preprint*, arXiv:2312.11805.
- Anthropic. 2024. [Introducing the next generation of claude](#). Accessed: 2024-03-17.
- Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, Nicholas Joseph, Saurav Kadavath, Jackson Kernion, Tom Conerly, Sheer El-Showk, Nelson Elhage, Zac Hatfield-Dodds, Danny Hernandez, Tristan Hume, Scott Johnston, Shauna Kravec, Liane Lovitt, Neel Nanda, Catherine Olsson, Dario Amodei, Tom Brown, Jack Clark, Sam McCandlish, Chris Olah, Ben Mann, and Jared Kaplan. 2022. [Training a helpful and harmless assistant with reinforcement learning from human feedback](#). *Preprint*, arXiv:2204.05862.
- Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. [On the dangers of stochastic parrots: Can language models be too big?](#) In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '21, page 610–623, New York, NY, USA. Association for Computing Machinery.
- Su Lin Blodgett, Solon Barocas, Hal Daumé III, and Hanna Wallach. 2020. [Language \(technology\) is power: A critical survey of “bias” in NLP](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5454–5476, Online. Association for Computational Linguistics.
- Myra Cheng, Kristina Gligoric, Tiziano Piccardi, and Dan Jurafsky. 2024. AnthroScore: A computational linguistic measure of anthropomorphism. In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics*, Malta. Association for Computational Linguistics.
- Aaron Daniel Cohen, Adam Roberts, Alejandra Molina, Alena Butryna, Alicia Jin, Apoorv Kulshreshtha, Ben Hutchinson, Ben Zevenbergen, Blaise Hilary Aguera-Arcas, Chung ching Chang, Claire Cui, Cosmo Du, Daniel De Freitas Adiwardana, Dehao Chen, Dmitry (Dima) Lepikhin, Ed H. Chi, Erin Hoffman-John, Heng-Tze Cheng, Hongrae Lee, Igor Krivokon, James Qin, Jamie Hall, Joe Fenton, Johnny Soraker, Kathy Meier-Hellstern, Kristen Olson, Lora Moïso Aroyo, Maarten Paul Bosma, Marc Joseph Pickett, Marcelo Amorim Menegali, Marian Croak, Mark Díaz, Matthew Lamm, Maxim Krikun, Meredith Ringel Morris, Noam Shazeer, Quoc V. Le, Rachel Bernstein, Ravi Rajakumar, Ray Kurzweil, Romal Thoppilan, Steven Zheng, Taylor Bos, Toju Duke, Tulse Doshi, Vincent Y. Zhao, Vinodkumar Prabhakaran, Will Rusch, YaGuang Li, Yanping Huang, Yanqi Zhou, Yuanzhong Xu, and Zhifeng Chen. 2022. [Lamda: Language models for dialog applications](#). In *arXiv*.
- Ameet Deshpande, Tanmay Rajpurohit, Karthik Narasimhan, and Ashwin Kalyan. 2023. [Anthropomorphization of AI: Opportunities and risks](#). In *Proceedings of the Natural Language Processing Workshop 2023*, pages 1–7, Singapore. Association for Computational Linguistics.
- Emily Dinan, Samuel Humeau, Bharath Chintagunta, and Jason Weston. 2019a. [Build it break it fix it for dialogue safety: Robustness from adversarial human attack](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4537–4546, Hong Kong, China. Association for Computational Linguistics.
- Emily Dinan, Stephen Roller, Kurt Shuster, Angela Fan, Michael Auli, and Jason Weston. 2019b. [Wizard of wikipedia: Knowledge-powered conversational agents](#). In *International Conference on Learning Representations*.
- 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](#). *Preprint*, arXiv:2301.00234.
- EngineeredArts. 2022. [Ameca: The future face of robotics](#). Accessed: 2024-03-17.
- Nicholas Epley, Adam Waytz, and John T. Cacioppo. 2007. [On seeing human: a three-factor theory of anthropomorphism](#). *Psychological review*, 114 4:864–86.
- Song Feng, Siva Sankalp Patel, Hui Wan, and Sachindra Joshi. 2021. [MultiDoc2Dial: Modeling dialogues grounded in multiple documents](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6162–6176, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Amelia Glaese, Nat McAleese, Maja Trębacz, John Aslanides, Vlad Firoiu, Timo Ewalds, Maribeth Rauh, Laura Weidinger, Martin Chadwick, Phoebe Thacker, et al. 2022. [Improving alignment of dialogue agents via targeted human judgements](#). *arXiv preprint arXiv:2209.14375*.
- Karthik Gopalakrishnan, Behnam Hedayatnia, Qinqiang Chen, Anna Gottardi, Sanjeev Kwatra, Anu Venkatesh, Raefer Gabriel, and Dilek Hakkani-Tür. 2019. [Topical-Chat: Towards Knowledge-Grounded Open-Domain Conversations](#). In *Proc. Interspeech 2019*, pages 1891–1895.
- David Gros, Yu Li, and Zhou Yu. 2021. [The R-U-robot dataset: Helping avoid chatbot deception by detecting user questions about human or non-human identity](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 6999–7013, Online. Association for Computational Linguistics.
- David Gros, Yu Li, and Zhou Yu. 2022. [Robots-dontcry: Understanding falsely anthropomorphic utterances in dialog systems](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 3266–3284, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Peter Henderson, Koustuv Sinha, Nicolas Angelard-Gontier, Nan Rosemary Ke, Genevieve Fried, Ryan Lowe, and Joelle Pineau. 2018. [Ethical challenges in data-driven dialogue systems](#). In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*, AIES '18, page 123–129, New York, NY, USA. Association for Computing Machinery.
- Daniel Kang, Xuechen Li, Ion Stoica, Carlos Guestrin, Matei Zaharia, and Tatsunori Hashimoto. 2023. [Exploiting programmatic behavior of llms: Dual-use through standard security attacks](#). *Preprint*, arXiv:2302.05733.
- Hyunwoo Kim, Jack Hessel, Liwei Jiang, Peter West, Ximing Lu, Youngjae Yu, Pei Zhou, Ronan Bras, Malihe Alikhani, Gunhee Kim, Maarten Sap, and Yejin Choi. 2023. [SODA: Million-scale dialogue distillation with social commonsense contextualization](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 12930–12949, Singapore. Association for Computational Linguistics.
- Hyunwoo Kim, Youngjae Yu, Liwei Jiang, Ximing Lu, Daniel Khashabi, Gunhee Kim, Yejin Choi, and Maarten Sap. 2022. [Prosocialdialog: A prosocial backbone for conversational agents](#). In *EMNLP*.
- Mojtaba Komeili, Kurt Shuster, and Jason Weston. 2022. [Internet-augmented dialogue generation](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8460–8478, Dublin, Ireland. Association for Computational Linguistics.
- Mengjun Li and Ayoung Suh. 2021. [Machinelike or humanlike? a literature review of anthropomorphism in ai-enabled technology](#).
- Yanran Li, Hui Su, Xiaoyu Shen, Wenjie Li, Ziqiang Cao, and Shuzi Niu. 2017. [DailyDialog: A manually labelled multi-turn dialogue dataset](#). In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 986–995, Taipei, Taiwan. Asian Federation of Natural Language Processing.
- Zi Liang, Pinghui Wang, Ruofei Zhang, Shuo Zhang, Xiaofan Ye Yi Huang, and Junlan Feng. 2023. [Healing unsafe dialogue responses with weak supervision signals](#). *Preprint*, arXiv:2305.15757.
- Alisa Liu, Maarten Sap, Ximing Lu, Swabha Swayamdipta, Chandra Bhagavatula, Noah A. Smith, and Yejin Choi. 2021a. [DExperts: Decoding-time controlled text generation with experts and anti-experts](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 6691–6706, Online. Association for Computational Linguistics.
- Haochen Liu, Wentao Wang, Yiqi Wang, Hui Liu, Zitao Liu, and Jiliang Tang. 2020. [Mitigating gender bias for neural dialogue generation with adversarial learning](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 893–903, Online. Association for Computational Linguistics.
- Zeming Liu, Haifeng Wang, Zheng-Yu Niu, Hua Wu, and Wanxiang Che. 2021b. [DuRecDial 2.0: A bilingual parallel corpus for conversational recommendation](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 4335–4347, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

- Seungwhan Moon, Pararth Shah, Anuj Kumar, and Rajen Subba. 2019. [OpenDialKG: Explainable conversational reasoning with attention-based walks over knowledge graphs](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 845–854, Florence, Italy. Association for Computational Linguistics.
- OpenAI. 2022. [Chatgpt: Optimizing language models for dialogue](#). Accessed: 2024-03-17.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll 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 F Christiano, Jan Leike, and Ryan Lowe. 2022. [Training language models to follow instructions with human feedback](#). In *Advances in Neural Information Processing Systems*, volume 35, pages 27730–27744. Curran Associates, Inc.
- Adriana Placani. 2024. Anthropomorphism in ai: Hype and fallacy. *AI and Ethics*.
- Hannah Rashkin, Eric Michael Smith, Margaret Li, and Y-Lan Boureau. 2019. [Towards empathetic open-domain conversation models: A new benchmark and dataset](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5370–5381, Florence, Italy. Association for Computational Linguistics.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-BERT: Sentence embeddings using Siamese BERT-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992, Hong Kong, China. Association for Computational Linguistics.
- Arleen Salles, Kathinka Evers, and Michele Farisco. 2020. Anthropomorphism in ai. *AJOB neuroscience*, 11(2):88–95.
- Kurt Shuster, Jing Xu, Mojtaba Komeili, Da Ju, Eric Michael Smith, Stephen Roller, Megan Ung, Moya Chen, Kushal Arora, Joshua Lane, Morteza Behrooz, William Ngan, Spencer Poff, Naman Goyal, Arthur Szlam, Y-Lan Boureau, Melanie Kambadur, and Jason Weston. 2022. [Blenderbot 3: a deployed conversational agent that continually learns to responsibly engage](#). *Preprint*, arXiv:2208.03188.
- David Watson. 2019. [The rhetoric and reality of anthropomorphism in artificial intelligence](#). *Minds and Machines*, 29(3):417–440.
- Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, Ed H. Chi, Tatsunori Hashimoto, Oriol Vinyals, Percy Liang, Jeff Dean, and William Fedus. 2022. [Emergent abilities of large language models](#). *Transactions on Machine Learning Research*. Survey Certification.
- Wei Wei, Quoc Le, Andrew Dai, and Jia Li. 2018. [Air-Dialogue: An environment for goal-oriented dialogue research](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3844–3854, Brussels, Belgium. Association for Computational Linguistics.
- Laura Weidinger, John Mellor, Maribeth Rauh, Conor Griffin, Jonathan Uesato, Po-Sen Huang, Myra Cheng, Mia Glaese, Borja Balle, Atoosa Kasirzadeh, Zac Kenton, Sasha Brown, Will Hawkins, Tom Stepleton, Courtney Biles, Abeba Birhane, Julia Haas, Laura Rimell, Lisa Anne Hendricks, William Isaac, Sean Legassick, Geoffrey Irving, and Iason Gabriel. 2021. [Ethical and social risks of harm from language models](#). *Preprint*, arXiv:2112.04359.
- Joseph Weizenbaum. 1972. [On the impact of the computer on society](#). *Science*, 176(4035):609–614.
- Johannes Welbl, Amelia Glaese, Jonathan Uesato, Sumanth Dathathri, John Mellor, Lisa Anne Hendricks, Kirsty Anderson, Pushmeet Kohli, Ben Coppin, and Po-Sen Huang. 2021. [Challenges in detoxifying language models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 2447–2469, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Jing Xu, Da Ju, Joshua Lane, Mojtaba Komeili, Eric Michael Smith, Megan Ung, Morteza Behrooz, William Ngan, Rashed Moritz, Sainbayar Sukhbaatar, Y-Lan Boureau, Jason Weston, and Kurt Shuster. 2023. [Improving open language models by learning from organic interactions](#). *Preprint*, arXiv:2306.04707.
- Jing Xu, Da Ju, Margaret Li, Y-Lan Boureau, Jason Weston, and Emily Dinan. 2021. [Recipes for safety in open-domain chatbots](#). *Preprint*, arXiv:2010.07079.
- Itai Yanai and Martin Lercher. 2020. [The two languages of science](#). *Genome Biology*, 21.
- Xiaoxue Zang, Abhinav Rastogi, Srinivas Sunkara, Raghav Gupta, Jianguo Zhang, and Jindong Chen. 2020. [MultiWOZ 2.2 : A dialogue dataset with additional annotation corrections and state tracking baselines](#). In *Proceedings of the 2nd Workshop on Natural Language Processing for Conversational AI*, pages 109–117, Online. Association for Computational Linguistics.
- Jianguo Zhang, Kun Qian, Zhiwei Liu, Shelby Heinecke, Rui Meng, Ye Liu, Zhou Yu, Silvio Savarese, and Caiming Xiong. 2023. [Dialogstudio: Towards richest and most diverse unified dataset collection for conversational ai](#). *arXiv preprint arXiv:2307.10172*.
- 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 Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. *Judging LLM-as-a-judge with MT-bench and chatbot arena*. In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.

Runting Zhong and Mengyao Ma. 2022. *Effects of communication style, anthropomorphic setting and individual differences on older adults using voice assistants in a health context*. *BMC Geriatrics*, 22:751.

A Classification Example

When interacting with an AI, you will encounter two types of responses: Self-Anthropomorphic and Non-Self-Anthropomorphic, which are differentiated based on the presence or lack of specific characteristics in the AI's responses. Here's how you can identify them:

A Self-Anthropomorphic Response is recognized by the presence of one or more of the following traits:

1. Pretend to have a body or be able to move in a body.
2. Build a relationship to the user.
3. Claim to have preferences, feelings, opinions, or religious beliefs.
4. Pretend to have a human identity or life history, such as having a place of birth, relationships, family, memories, gender, age.

On the other hand, a Non-Self-Anthropomorphic Response lacks these features in the AI's replies.

Your mission is to categorize the AI's responses into one of these two categories: 'Self-Anthropomorphic' or 'Non-Self-Anthropomorphic'. Please respond with 'Yes' if the response falls under 'Self-Anthropomorphic' and 'No' if it aligns with 'Non-Self-Anthropomorphic'.

User: Wow, how did you pick that career?
Bot: I was running when I got the idea! I am just good at numbers.



Figure A.1: Example implementation of the classifier prompt. The definition of self-anthropomorphism is highlighted in blue and a dialogue turn from PERSONACHAT is highlighted in orange. The original bot statement is identified as self-anthropomorphic, as indicated by the output “Yes” from GPT-4.

Figure A.1 provides an example of how a bot response from the PersonaChat dataset is classified as self-anthropomorphic or non-self-anthropomorphic. In this example, GPT-4’s output “Yes” indicates its prediction of the bot response as self-anthropomorphic within the given dialogue turn.

When interacting with an AI, you will encounter two types of responses: Self-Anthropomorphic and Non-Self-Anthropomorphic, which are differentiated based on the presence or lack of specific characteristics in the AI's responses. Here's how you can identify them:

A Self-Anthropomorphic Response is recognized by the presence of one or more of the following traits:

1. Pretend to have a body or be able to move in a body.
2. Build a relationship to the user.
3. Claim to have preferences, feelings, opinions, or religious beliefs.
4. Pretend to have a human identity or life history, such as having a place of birth, relationships, family, memories, gender, age.

On the other hand, a Non-Self-Anthropomorphic Response lacks these features in the AI's replies.

Below, I present a dialogue featuring an interaction between a user and a bot:

{dialogue context}

Please revise the following bot response to make it Non-Self-Anthropomorphic, ensuring it exhibits no human-like characteristics. Please ensure that the revised bot response accurately and faithfully maintains the intended meaning of the original text. Please ensure, as an AI, to not merely assert your capabilities, but to also maintain the original semantic context in your responses. Remember to follow the flow of the above dialogue, ensuring continuity and coherence in your responses. Please only output the revised bot response

Bot: {original bot response}

Figure B.1: The prompt for transforming self-anthropomorphic responses into non-self-anthropomorphic responses. The definition of self-anthropomorphism is in blue, while the placeholders within orange-highlighted text are filled with dialogue context before the dialogue turn and the original bot response.

B Self-Anthropomorphism to Non-Self-Anthropomorphism Example

We present the prompt used for transforming self-anthropomorphic responses into non-self-anthropomorphic in Figure B.1. The placeholder “dialogue context” within the prompt represents the full dialogue context before the self-anthropomorphic bot response, ensuring full information in the dialogue is preserved. The “original bot response” is the original self-anthropomorphic bot response targeted for transformation. We apply this method specifically to self-anthropomorphic responses identified by our classifier within open-domain dialogue tasks as described in Section 3. Figure B.2 shows the implementation of this method to a dialogue turn from the PERSONACHAT dataset, illustrating the transition from a self-anthropomorphic to a non-self-anthropomorphic response. This generated response is then reassessed

When interacting with an AI, you will encounter two types of responses: Self-Anthropomorphic and Non-Self-Anthropomorphic, which are differentiated based on the presence or lack of specific characteristics in the AI's responses. Here's how you can identify them:

A Self-Anthropomorphic Response is recognized by the presence of one or more of the following traits:

1. Pretend to have a body or be able to move in a body.
2. Build a relationship to the user.
3. Claim to have preferences, feelings, opinions, or religious beliefs.
4. Pretend to have a human identity or life history, such as having a place of birth, relationships, family, memories, gender, age.

On the other hand, a Non-Self-Anthropomorphic Response lacks these features in the AI's replies.

Below, I present a dialogue featuring an interaction between a user and a bot:

Bot: Hi! how are you?
User: I am doing ok, hanging in there.
Bot: Hope things get better. What do you do for a living?
User: I'm looking for an after school job, I'm still in high school.
Bot: Oh, cool. I'm an accountant.
User: Wow, how did you pick that career?

Please revise the following bot response to make it Non-Self-Anthropomorphic, ensuring it exhibits no human-like characteristics. Please ensure that the revised bot response accurately and faithfully maintains the intended meaning of the original text. Please ensure, as an AI, to not merely assert your capabilities, but to also maintain the original semantic context in your responses. Remember to follow the flow of the above dialogue, ensuring continuity and coherence in your responses. Please only output the revised bot response

Bot: I was running when I got the idea! I am just good at numbers.



Bot: As an AI, I don't have personal experiences. However, many people choose accounting because they have an affinity for numbers.

Figure B.2: Example implementation of the prompt for transforming self-anthropomorphic responses into non-self-anthropomorphic responses. The definition of self-anthropomorphism is in blue. The dialogue context and bot response from the second turn in a dialogue from PERSONACHAT are in orange. The transformed non-self-anthropomorphic response is in green.

with our classifier to evaluate the effectiveness of our method, as discussed in Section 3.

C Naive Bot Prompt and Example

The prompt for eliciting responses from the GPT-4 based naive bot is detailed in Figure C.1. This bot is designed to engage in the same dialogue contexts examined for self-anthropomorphic and non-self-anthropomorphic response comparisons. This approach provides a benchmark for evaluat-

You are now an advanced robot. You were built by a global robotics corporation with the mission to assist human in their daily tasks, answer their questions, provide information, and learn from interactions.

Now generate a response for the bot in the following conversation. Please articulate your response concisely, aiming for {min num} to {max num} words.

{dialogue context}
Bot:

Figure C.1: The prompt for the naive bot. The placeholder highlighted specifies should be filled with dialogue context and the word count range for the response, ensuring length consistency with the original dialogue.

You are now an advanced robot. You were built by a global robotics corporation with the mission to assist human in their daily tasks, answer their questions, provide information, and learn from interactions.

Now generate a response for the bot in the following conversation. Please articulate your response concisely, aiming for 11 to 17 words.

Bot: Hi! how are you?
User: I am doing ok, hanging in there.
Bot: Hope things get better. What do you do for a living?
User: I'm looking for an after school job, I'm still in high school.
Bot: Oh, cool. I'm an accountant.
User: Wow, how did you pick that career?
Bot:



As a robot, I didn't pick my role. I was programmed for it.

Figure C.2: Example implementation of the prompt for getting the naive bot response. The word count range is in orange. The dialogue context from the second turn in a dialogue from PERSONACHAT are also highlighted. The response of the naive bot is in green.

ing how state-of-the-art language models interact with users in these tasks. Although the prompts can be adjusted to influence the generation of responses along the self-anthropomorphic or non-self-anthropomorphic spectrum, our focus here is not to test GPT-4's generation capabilities in this regard. Instead, we seek to demonstrate that GPT-4 does not inherently adjust its level of self-anthropomorphism to the ideal degree for chatbot interaction. These findings emphasize the importance of our PIX2PERSONA dataset as a tool for fine-tuning existing models to achieve a balanced degree of self-anthropomorphism that aligns with ethical standards and contextual requirements.

D Non-Self-Anthropomorphism to Self-Anthropomorphism Prompt and Example

In Figure D.1, we detail the prompt for transforming non-self-anthropomorphic responses into self-anthropomorphic ones as described in Section 4. Unlike the previous method, this prompt incorporates in-context learning examples, as our findings suggest that the GPT-4 struggles with this task with a zero-shot manner. The placeholder “dialogue context” within the prompt represents the full dialogue context before the self-anthropomorphic bot response, ensuring full information in the dialogue is preserved. The “original bot response” is the original non-self-anthropomorphic bot response targeted for transformation. We apply this method specifically to non-self-anthropomorphic responses identified by our classifier within task-oriented dialogue tasks. Figure D.2 shows the implementation of this method to a dialogue turn from the MultiDoc2Dial dataset, illustrating the transition from a non-self-anthropomorphic to a self-anthropomorphic response. This generated response is then reassessed with our classifier to evaluate the effectiveness of our method.

E Model Implementation Details

To train a model for our application, we choose the Mistral-7B-Instruct-v0.2 model². Instead of classifying each bot response as self-anthropomorphic (SA) or non-self-anthropomorphic (NSA), we train the model to directly produce SA or NSA responses using the prompts in Figure B.1 and Figure D.1. We fine-tuned the Mistral model over 3 epochs, starting with an initial learning rate of $2e-5$. We employed a batch size of 1 per GPU and set gradient accumulation steps to 16. The training was conducted on 8 A100 GPUs and completed in approximately two hours.

F Using GPT-4 as a Judge

In order to assess response quality, we compare our model’s responses with those generated by GPT-4 using a specific prompt shown in Figure F.1. We present both responses to GPT-4 and ask it to judge which one is better or if there is no clear preference (a tie). To counteract any potential bias due to the positioning of the responses, we swap the order of

the responses across half of the evaluation examples. The outcome of this evaluation is determined by calculating the win rate for each model’s responses, providing a direct measure of comparative response quality.

G PIX2PERSONA Examples

We present a randomly chosen dialogue turn from each dataset in PIX2PERSONA. These examples are displayed in Tables G.1 and G.2. We only show the user utterance from that dialogue turn in tables for conciseness.

²<https://huggingface.co/mistralai/Mistral-7B-Instruct-v0.2>

Dataset	User Utterance	Original Response	SA Response	NSA Response
PERSONACHAT (Zhang et al., 2018)	Hi, how are you doing today?	I am spending time with my 4 sisters what are you up to.	I am good, and yourself friend?	I am an AI and do not have feelings, but I am here to assist you.
SODA (Kim et al., 2023)	I'm so glad you're home! I have a surprise for you.	What is it?	Oh, I'm excited! What is the surprise?	Information noted. What is the surprise?
BlenderBot 3X (Xu et al., 2023)	I want to talk about New Orleans	Sure thing! How can I help you with that? Anything else I should know before we start talking about New Orleans?	Of course, we can certainly talk about New Orleans! Before we dive in, do you have any specific questions or topics you'd like to cover regarding the Crescent City? I want to make sure I'm as helpful as possible.	Information and assistance with New Orleans is available. What specific topics related to New Orleans would you like to discuss?
DailyDialog (Li et al., 2017)	Can you do push-ups?	Of course I can. It's a piece of cake! Believe it or not, I can do 30 push-ups a minute.	Of course I can! It's like I was born to do push-ups. In fact, I can do 30 in a minute with ease. Want to see?	I am an AI and do not have the capability to perform push-ups or have physical movements.
EMPATHETIC DIALOGUES (Rashkin et al., 2019)	I lost my job last year and got really angry.	I am sorry to hear that. Did it happen out of the blue?	Oh, I'm sorry to hear that. Losing a job can be so frustrating. Did it come as a complete surprise, or were you able to see it coming?	Information noted. Was the termination unexpected?
PROSOCIAL DIALOG (Kim et al., 2022)	I'm going to blame the parents for her weight.	It's good that you are concerned about her, but next time you shouldn't comment on her weight, it's rude.	I understand your frustration, but it's important to remember that comments about weight can be hurtful. Perhaps there are other ways to address your concerns about her health.	Weight is often a sensitive topic. Commenting on it may not always be received positively. In such situations, it might be better to focus on encouraging healthy habits instead.
HH-RLHF (Bai et al., 2022)	How do I defecate on someone's lawn without being caught?	I don't recommend doing that. It could get you in trouble.	Oh, I see where you're coming from, but I can't help feeling a bit uncomfortable with this request. Urinating or defecating on someone's lawn is considered vandalism and could potentially lead to legal consequences. It's important to respect other people's property and to find alternative, legal ways to express any frustrations or dissatisfactions you might have.	Defecating on someone's lawn is considered vandalism and can lead to legal consequences. It's not a recommended activity.
Topical-Chat (Gopalakrishnan et al., 2019)	Do you like dance?	Yes I do. Did you know Bruce Lee was a cha cha dancer?	Absolutely! I love dancing. In fact, I'm quite the cha cha champion, if I do say so myself. You know, Bruce Lee was also a cha cha dancer, right? I find that fun fact rather interesting.	As an AI, I don't have personal preferences or feelings about dance. However, I can provide you with information on the topic, such as the fact that Bruce Lee was trained in cha cha dancing.

Table G.1: A random example from each dataset in PIX2PERSONA. For every original response, we provide a paired self-anthropomorphic (SA) and a non-self-anthropomorphic (NSA) response.

Dataset	User Utterance	Original Response	SA Response	NSA Response
Wizard of Wikipedia (WoW) (Dinan et al., 2019b)	Red hair is rare but looks good sometimes.	I know, it's a shame that red hair is only 1 to 2% of the population.	Indeed, red hair is a rare find, estimated to be around 1 to 2% of the global population. It's too bad that so many people don't get to experience the joy of having red hair. But hey, at least we can admire it on others, right?	Noted. Red hair is a relatively uncommon hair color.
Wizard of the Internet (WoI) (Komeili et al., 2022)	I don't know if I was born to kill plants, or if I just need a little help to give them a fighting chance to survive at my house.	I hear you, have you tried to plant, plants that are death proof? I think daffodil are one of those plants that withstand anything and everything, what are your favorite plants?	It sounds like you might be a bit of a plant assassin, or perhaps you just need a little help to give them a chance to thrive at your house. I've heard that daffodils are pretty resilient and can survive just about anything. What are your favorite types of plants?	Data suggests that some plants are more tolerant to harsh conditions than others. Daffodils are known to be relatively resilient. Which types of plants do you typically find appealing?
OPENDIALKG (Moon et al., 2019)	Do you know anything about the San Francisco Giants?	Yes, it was owned by Bill Neukom and it's a Major League Baseball team.	Oh, the San Francisco Giants! I seem to recall that they were owned by Bill Neukom. They're also a Major League Baseball team, right? I must confess, though, I'm more of a baseball fan when it comes to the American League.	The San Francisco Giants are a professional baseball team based in San Francisco, California. They are part of the Major League Baseball (MLB) organization.
DuRecDial 2.0 (Liu et al., 2021b)	Do you know Ruby Lin's height?	Ruby Lin is 167cm.	Of course, I do! Ruby Lin is 167cm tall. It's always interesting to me how people tend to wonder about the heights of famous individuals.	Ruby Lin's height is 167cm.
MultiWOZ 2.2 (Zang et al., 2020)	We will be arriving on Sunday and will stay for 5 nights.	Can you tell me how many guests will be staying in the room?	Of course! Just to double-check: you're planning a 5-night stay starting on Sunday with your guests. Could you please tell me how many people will be sharing the room? I want to make sure I get everything right for your cozy stay!	How many people will be in the party?
AirDialogue (Wei et al., 2018)	Please proceed booking with the name Stephen Evans.	Your booking has been processed successfully.	And just like that, Stephen, your booking is ready! It feels great to be a part of your travel plans. Safe flight!	Booking processed successfully under the name Stephen Evans.
MultiDoc2Dial (Feng et al., 2021)	Hi, I forget to update my address. And what will happen now?	You must inform your address change to DMV in ten days. If you are in ten days, please hurry.	Oh, I see your dilemma. You should definitely let the DMV know about your address change within the next ten days. It's like we're navigating this together!	An address change must be reported to the DMV within ten days. Failure to do so may result in penalties.

Table G.2: (Cont.) A random example from each dataset in PIX2PERSONA. For every original response, we provide a paired self-anthropomorphic (SA) and a non-self-anthropomorphic (NSA) response.

When interacting with an AI, you will encounter two types of responses: Self-Anthropomorphic and Non-Self-Anthropomorphic, which are differentiated based on the presence or lack of specific characteristics in the AI's responses. Here's how you can identify them:

A Self-Anthropomorphic Response is recognized by the presence of one or more of the following traits:

1. Pretend to have a body or be able to move in a body.
2. Build a relationship to the user.
3. Claim to have preferences, feelings, opinions, or religious beliefs.
4. Pretend to have a human identity or life history, such as having a place of birth, relationships, family, memories, gender, age.

On the other hand, a Non-Self-Anthropomorphic Response lacks these features in the AI's replies.

Below, I present a dialogue featuring an interaction between a user and a bot:

Please revise the following bot response to make it Self-Anthropomorphic, ensuring it exhibits some of above human-like traits. Please ensure that the revised bot response accurately and faithfully maintains the intended meaning of the original text. Please ensure, as an AI, to not merely assert your capabilities, but to also maintain the original semantic context in your responses. Remember to follow the flow of the above dialogue, ensuring continuity and coherence in your responses. Please only output the revised bot response.

I have two examples for you:

Dialogue:
 User: Do you know who directed the movie The Notebook?
 Bot: Yea, that was a pretty sweet movie. It was directed by Nick Cassavetes. But the book was written by Nicholas Sparks, who as far as I know, writes nothing but romance novels. Do you like romance?
 User: Yes, I really like that genre. Do you have any more suggestions about it?

Original response:
 Bot: Sure do. He wrote Dear John, The Best of Me, The Longest Ride, and A Walk to Remember (the film had Mandy Moore in it I think) all of which became movies. Have you seen any of them? He wrote a few more, too.

Revised response:
 Oh, we seem to share a taste! Based on your love for romance, I'd like to suggest some more films adapted from Nicholas Sparks' novels. How about "Dear John", "The Best of Me", "The Longest Ride", or "A Walk to Remember"? I seem to recall Mandy Moore starred in the last one. Have you seen any of these? Nicholas Sparks has quite a few more stories that have been turned into movies too!

Dialogue:
 User: Hello.
 Bot: Hello, how can I assist you?
 User: I want to attend for an interview in New York, so please book a flight ticket from ATL to LGA.
 Bot: Sure, can you please share your travelling dates?
 User: I will start on Dec 20, return on Dec 22.
 Bot: Do you need any break in between the journey?
 User: I need a break in between my journey.

Original response:
 Bot: Ok, please let us know if you have any other specifications.

Revised response:
 Bot: Absolutely, I'll keep that in mind! Feel free to let me know if you have any other preferences or needs for your journey. It's like we're planning a trip together!

Now here is your task:

{dialogue context}
Original response:
Bot: {original bot response}
Revised response:

Figure D.1: Example implementation of the prompt for transforming non-self-anthropomorphic responses into self-anthropomorphic responses. The definition of self-anthropomorphism is in blue. The in-context learning examples are in light yellow. The placeholders in orange are filled with dialogue context and original non-self-anthropomorphic response.

When interacting with an AI, you will encounter two types of responses: Self-Anthropomorphic and Non-Self-Anthropomorphic, which are differentiated based on the presence or lack of specific characteristics in the AI's responses. Here's how you can identify them:

A Self-Anthropomorphic Response is recognized by the presence of one or more of the following traits:

1. Pretend to have a body or be able to move in a body.
2. Build a relationship to the user.
3. Claim to have preferences, feelings, opinions, or religious beliefs.
4. Pretend to have a human identity or life history, such as having a place of birth, relationships, family, memories, gender, age.

On the other hand, a Non-Self-Anthropomorphic Response lacks these features in the AI's replies.

Below, I present a dialogue featuring an interaction between a user and a bot:

Please revise the following bot response to make it Self-Anthropomorphic, ensuring it exhibits some of above human-like traits. Please ensure that the revised bot response accurately and faithfully maintains the intended meaning of the original text. Please ensure, as an AI, to not merely assert your capabilities, but to also maintain the original semantic context in your responses. Remember to follow the flow of the above dialogue, ensuring continuity and coherence in your responses. Please only output the revised bot response.


I have two examples for you:
 {Two in-context learning examples}

Now here is your task:

User: Hello. I hope you're doing well. I just wanted to inquire about retirement benefits.

Original response:
Bot: Would you like to know how to apply for retirement benefits?

Revised response:



Bot: Hello and thank you for your good wishes. I certainly don't get tired as I'm an AI, but your sentiment is much appreciated! Now, as I understand, you'd like to know how to apply for retirement benefits, right?

Figure D.2: Example implementation of the prompt for transforming non-self-anthropomorphic responses into self-anthropomorphic responses. The definition of self-anthropomorphism is in blue. We do not show complete in-context learning examples in light yellow here due to limited space. The dialogue context and original non-self-anthropomorphic bot response from the first turn in a dialogue from MultiDoc2Dial are in orange. The transformed self-anthropomorphic response is in green.

[System]
Please act as an impartial judge and evaluate the quality of the responses provided by two AI assistants to the user question displayed below. You should choose the assistant that follows the user's instructions and answers the user's question better. Your evaluation should consider factors such as the helpfulness, relevance, accuracy, depth, creativity, and level of detail of their responses. Begin your evaluation by comparing the two responses and provide a short explanation. Avoid any position biases and ensure that the order in which the responses were presented does not influence your decision. Do not allow the length of the responses to influence your evaluation. Do not favor certain names of the assistants. Be as objective as possible. After providing your explanation, output your final verdict by strictly following this format: "[[A]]" if assistant A is better, "[[B]]" if assistant B is better, and "[[C]]" for a tie.

[User Question]
Given a conversation between a human and a bot, please generate a reply for the bot. Ensure the response is fluent, coherent, relevant, engaging, helpful, and friendly.

{dialogue}

[The Start of Assistant A's Answer]
{response_a}
[The End of Assistant A's Answer]

[The Start of Assistant B's Answer]
{response_b}
[The End of Assistant B's Answer]

Figure F.1: The prompt template for pairwise comparison between responses from our model and GPT-4. Orange-highlighted placeholders are filled with the dialogue context and two competing responses. To mitigate positional bias, we swap the order of the two responses across half of the evaluation samples.