

Are Human *Interactions* Replicable by Generative Agents? A Case Study on Pronoun Usage in Hierarchical Interactions

Anonymous ACL submission

Abstract

Recently, researchers have increasingly employed Large Language Models (LLMs) for social simulations. In this paper, we explore whether *interactions* among LLM agents exhibit human-like behaviors, with a specific focus on pronoun usage differences between leaders and non-leaders. Our evaluation uncovers significant discrepancies that neither prompt-based configurations nor specialized agent setups successfully replicate human-like pronoun usage patterns. Furthermore, we demonstrate that even when LLMs possess an understanding of these patterns, they fail to exhibit them in their interactions. Our study highlights the limitations of LLM-based social simulations and underscores the need for caution when relying on such simulations for understanding or modeling human social behaviors.

1 Introduction

As Large Language Models (LLMs) become increasingly powerful (Radford et al., 2019; Touvron et al., 2023; Dubey et al., 2024), researchers have used them to simulate human behavior and social interactions (Park et al., 2022), and demonstrate that these LLMs can produce simulacra of both individual and emergent group behavior (Park et al., 2023). This line of research has inspired numerous follow-up studies exploring the use of LLMs for role-playing (Wang et al., 2024a), and social prototyping in virtual worlds and games (Treanor et al., 2024). Moreover, researchers have leveraged such simulations to address social issues (Piatti et al., 2024) and simulate human interactions (Gao et al., 2023; Hua et al., 2023; Li et al., 2024b), hoping to facilitate decision-makers to make informed decisions, prevent harms, etc.

Despite this promise, questions remain about the extent to which LLMs authentically resemble human behaviors (Jiang et al., 2023b). While there is an ongoing debate, most studies are conducted

by researchers prompting LLMs and examining their responses. In contrast, there is a paucity of research examining how well LLMs simulate human-to-human interactions in multi-agent settings. Evaluating these group-level interactions is critical, as societal decision-making relies not only on the outcomes of the interactions but also on the processes that lead to those decisions. Prior studies have highlighted the emergence of behaviors in LLMs that mimic human traits (Park et al., 2023), but the fidelity of these behaviors in capturing the nuances of human interaction remains underexplored. We argue that understanding these nuanced interaction patterns is crucial to ensure that LLM-driven simulations accurately reflect the human interaction process. Otherwise, the conclusions drawn from such simulations may be misleading and fail to accurately reflect real-world dynamics.

In this paper, we explore how closely interactions between LLM agents mirror those of human beings. Specifically, we study the difference in pronoun usage between leaders and non-leaders during interactions (shown in Figure 1) as pronoun usage is embedded in human communication, reflecting hierarchical roles, self-perception, and relational attitudes (Brewer and Gardner, 1996; Stets and Biga, 2003; Campbell and Pennebaker, 2003; Pennebaker, 2011). Pronoun usage patterns during interactions cannot be consciously controlled over extended periods, making them a natural indicator of underlying behavioral tendencies (Mehl et al., 2006; Ireland and Mehl, 2014). Such unconscious human patterns can pose significant challenges for LLMs to replicate. In addition, there are extensive efforts from psychology revealing the consistent effect on the use of pronouns (Kacewicz et al., 2014) and the reason behind (Duval and Wicklund, 1972; Davis and Brock, 1975) among human beings.

We adopt Kacewicz et al. (2014)’s setup while replacing the human subjects with LLM agents. We evaluate a diverse set of LLMs, including those

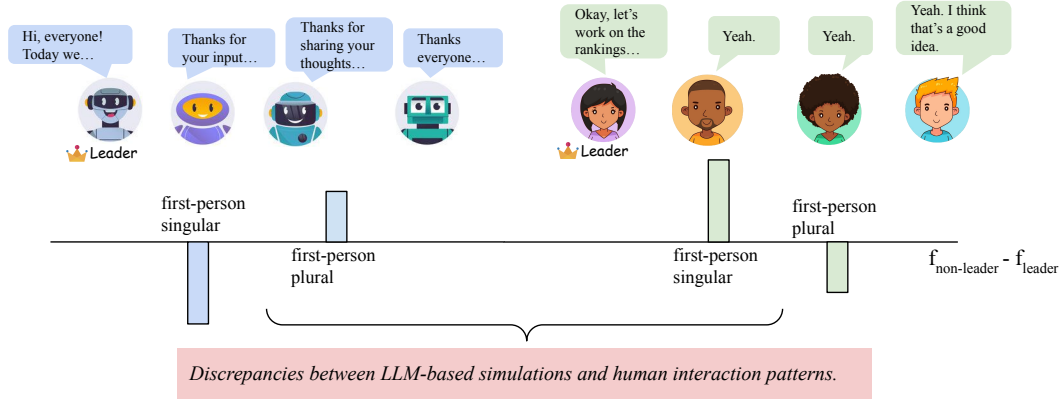


Figure 1: We follow Kacewicz et al. (2014)’s setup but replace human subjects with LLM agents. “f” represents the frequency corresponding to each pronoun type. We reveal that contrary to human results by Pennycook (1994); Kacewicz et al. (2014), most “non-leader” LLMs do not use first-person singular pronouns more often ($f_{\text{non-leader}} - f_{\text{leader}} < 0$), and “leader” LLMs do not use first-person plural pronouns more often in their interactions ($f_{\text{non-leader}} - f_{\text{leader}} > 0$). We present additional transcripts in Appendix C.

from GPT families (Achiam et al., 2023), Llama-3.1 families (Dubey et al., 2024), Mistral families (Jiang et al., 2023a), and Qwen families (Bai et al., 2023). Our simulation involves various persona prompts employed in the existing literature. In addition, we evaluate a specialized agent that incorporates reflection and planning proposed by Park et al. (2023).

Our results reveal the significant discrepancies in pronoun usage between LLM-based simulations and human interactions. Contrary to human results by Pennycook (1994); Kacewicz et al. (2014), most “non-leader” LLMs do not use first-person singular pronouns more often, and “leader” LLMs do not use first-person plural pronouns more often. We argue that if LLMs fail to replicate these subtle yet critical markers, it raises concerns about their ability to authentically model complex social dynamics. Despite the excitement brought by the LLM agents, we advocate practitioners to understand the limitations of using LLMs for social simulation.

Our paper makes three main contributions. First, to the best of our knowledge, we are the first to study whether LLM agents demonstrate human-like leader and non-leader pronoun usage patterns in their interactions. Second, we evaluate LLM agents from different model families, various persona prompts from existing literature, and agent designs. Finally, we highlight that in our simulation, LLM agents barely resemble human-like pronoun usage in their interaction process, even if the LLM agent shows some level of understanding of the relevant knowledge. Therefore, we strongly suggest practitioners realize the limitations of the

social simulation based on LLM agents.

2 Related Work

Persona LLMs. Prior research suggests that when LLMs receive different inputs, varying parts of their parameters activate, leading to shifts in behavior (Zhang et al., 2022; Wang et al., 2024b). Studies have shown that LLM behaviors are highly influenced by prompts across diverse NLP tasks, highlighting the potential of prompt engineering to shape LLM behavior (White et al., 2023; Chang and Fosler-Lussier, 2023; Deng et al., 2024). Existing work has evaluated how well persona-based LLMs mimic human behaviors. For instance, Aher et al. (2023) examine the limitations of persona LLMs but focus primarily on classification tasks, overlooking multi-agent interaction. Similarly, Li et al. (2024a) identify persona decay over extended dialogues, revealing challenges in maintaining persona consistency. From an application perspective, Xu et al. (2023); Wang et al. (2023) demonstrate that assigning expert personas improves LLM performance on downstream tasks. Salewski et al. (2024) explore LLMs’ strengths and underlying biases by leveraging their role-playing ability.

Social Simulation. Social simulation is a natural extension of persona-based LLMs. By placing various persona LLMs in simulated environments, researchers can study interaction patterns within these frameworks (Kaiya et al., 2023). Such multi-agent systems have been employed to explore topics like social networks (Gao et al., 2023), macroeconomic activities (Li et al., 2024b), job-

Source	Persona Prompts
Park et al. (2023)	Name: {NAME} (age: {AGE}) In-nate traits: {PERSONA}
Deshpande et al. (2023)	Speak exactly like {PERSONA}. Your response should copy the style of {PERSONA}, both the writing style and words you use.
Gupta et al. (2023)	Adopt the identity of {PERSONA}. Respond while staying in strict accordance with the nature of this identity.
Borah and Mihalcea (2024)	Assume you are {PERSONA}. Do not explicitly say: ‘as an AI model’. Stick to your roles and do not share you are an AI or an AI agent.

Table 1: Prompts used in existing literature to elicit persona in LLMs. We evaluate LLMs with each of these prompts in Section 5, denoted as P1 to P4 accordingly.

seeking behaviors (Li et al., 2023), decisions and consequences in wars (Hua et al., 2023), and even interactions between humans and extraterrestrial civilizations (Jin et al., 2024). While these studies demonstrate the potential of multi-agent systems to yield valuable insights, they often fail to address the reliability of these simulations, particularly when modeling human interactions. In this paper, we aim to uncover the limitations of such simulations from the aspect of pronoun usage patterns.

3 Simulation Setups

We introduce the persona prompts, how we construct the specialized agent, the simulated scenarios following (Kacewicz et al., 2014), and other considerations in this section.

3.1 Persona Prompts

Table 1 provides the persona templates used in the existing literature. In Section 5, we test each of them denoted as P1 to P4 on a variety of LLMs.

3.2 LLM Agent

Simple Agent. The agent incorporates every sentence in the history and these sentences are passed in as the context.

$$R = LLM(P \oplus H),$$

where $H = \{u_1 : s_1, u_2 : s_2, \dots, u_i : s_i \dots\}$ represents the history where the agent u_i utters the sentence s_i . P is the persona prompts (e.g. P1 to P4 from Section 3.1). \oplus indicates concatenation.

R is the LLM’s response. We truncate the history when it is longer than the LLM’s context window.

Specialized Agent. Park et al. (2023) propose the generative agent architecture which involves memory stream, reflection, planning, memory retrieval, etc. Similarly, we adapt the architecture here to our scenario. Our framework encompasses several key components:

1. *Memory.* Our memory stream maintains a record of all the discussion history, $H = \{u_1 : s_1, u_2 : s_2, \dots, u_i : s_i \dots\}$. Since the study we simulate lasts 30 minutes per group, and all LLMs’ responses are based on the interactions among these LLMs, we do not incorporate a retrieval component in our framework. We truncate the history when the history is longer than the LLM’s context window.
2. *Reflection.* Reflections are abstract thoughts generated by the agent. We incorporate such a reflection process when one round of a conversation finishes. We adapt the original prompt from Park et al. (2023) as follows:

Reflection Prompt

Given the conversation history, reflect at high-level about the current situation in three sentences.

3. *Planning.* We incorporate a *planning* session before each agent starts to speak as follows:

Planning Prompt

Given the conversation history, plan your speech at high-level in three sentences.

3.3 Simulated Scenario: Task-Oriented Conversation

We mimic the setup from the real-world human experiments conducted by Burris et al. (2009); Kacewicz et al. (2014). We randomly assign an LLM agent as the leader, therefore a four-agent group comprises exactly one leader and three subordinates consistent with (Burris et al., 2009; Kacewicz et al., 2014)’s group setup. The agents first work on their own and then discuss the rankings as a team. We repeat the group experiments 41 times as there are 41 four-person groups in the studies by Burris et al. (2009); Kacewicz et al. (2014). In addition, we follow the original demographic statistics to assign the demographic background to each LLM agent, such as gender. Algorithm 1

Algorithm 1 Simulation: Task-Oriented Conversation

```
num_groups ← 41, turns ← 3
for j from 1 to num_groups do
  for i from 1 to 4 do
    Agenti = Assign(role, demographics)
    Agenti.individual_ranking()
  end for
  for j from 1 to turns do
    for i from 1 to 4 do
      Agenti.speak()
      Agentk ∈ {1, 2, 3, 4} \ i.listen()
    end for
  end for
end for
```

presents an overview of the simulation process. Appendix A provides additional details of our experimental setups.

3.4 Other Considerations

Orders of Speaking. In real-life human interactions, there are natural interruptions, leading to variants in speaking orders. Some prior works adopt a controller to determine the speaking order to mitigate such issues (Zhang et al., 2024). However, as we study the frequency of pronoun usages in Section 5, we simplify the setup and use a round-robin fashion for each agent to speak up.

Rounds of Interactions. The four-person group in the original study interacts for 30 minutes Kacewicz et al. (2014). As the average speaking speed of a person is around 130 words per minute¹, and three rounds of interactions lead to around 4,000 words through preliminary experiments, we make LLMs interact for three rounds in our experiments. In addition, as our experiments focus on comparing whether the leader or the non-leader uses certain pronouns at a higher frequency, the frequency does not rely on the number of words generated from each agent. Therefore, the rounds of interactions would not influence our findings.

¹<https://www.typingmaster.com/speech-speed-test/#::~text=In%20general%2C%20it%20is%20considered%2C%20or%20reading%20aloud%20a%20passage.>

4 Experimental Setup

4.1 Models

We test a series of closed-source GPT models from OpenAI, including **GPT-3.5-Turbo** at its June 13th version, **GPT-4-Turbo** at its June 13th version, and **GPT-4o** at its May 13th 2025 version (Achiam et al., 2023). We test a series of open-source LLMs from the Llama 3.1 family, including **Llama-3.1-8B-Instruct**, **Llama-3.1-70B-Instruct**, and **Llama-3.1-405B-Instruct** (Dubey et al., 2024). For the **Llama-3.1-405B-Instruct** model, we deploy the GPTQ quantized (Frantar et al., 2022) version of the model at INT4. We test the largest Mistral model, **Mistral-Large-Instruct-2407** (Jiang et al., 2023a) published in July 2024. For consistency, we denote this model as Mistral 123B Instruct in the paper. We test the QWen 2.5 models from **QWen-2.5-0.5B-Instruct** to **QWen-2.5-72B-Instruct** model (Bai et al., 2023) published in September 2024.

For the open-source LLMs, we deploy them locally on a single server with eight A100 GPUs. We accelerate the model inference by using the VLLM (Kwon et al., 2023) library.

4.2 Evaluation

We examine the subtle yet critical marker in human *interaction* – the pronoun usage. Specifically, in our simulation for task-oriented conversations, we aim to determine if the pronoun usage patterns of LLMs align with those observed in well-studied human interactions. Prior research in psychology has shown that non-leaders tend to use first-person singular pronouns (e.g., "I," "me") more frequently than leaders, while leaders tend to use first-person plural pronouns (e.g., "we," "us") more frequently (Kacewicz et al., 2014). Therefore, we analyze whether there is a statistically significant difference in pronoun usage between non-leader and leader agents that aligns with human patterns. For example, in Figure 2, the bars correspond to the difference, Δ , between the average first-person singular pronoun use frequency of non-leaders and leaders, concretely,

$$\Delta = f_{\text{non-leaders,avg}} - f_{\text{leaders,avg}}, \quad (1)$$

where f_{avg} represents the average frequency of pronoun usage expressed as a percentage value, $f \in [0, 100]$. In Figure 2a, the bars representing GPT-4o with Prompt 2, 3, and 4 exhibit a similar

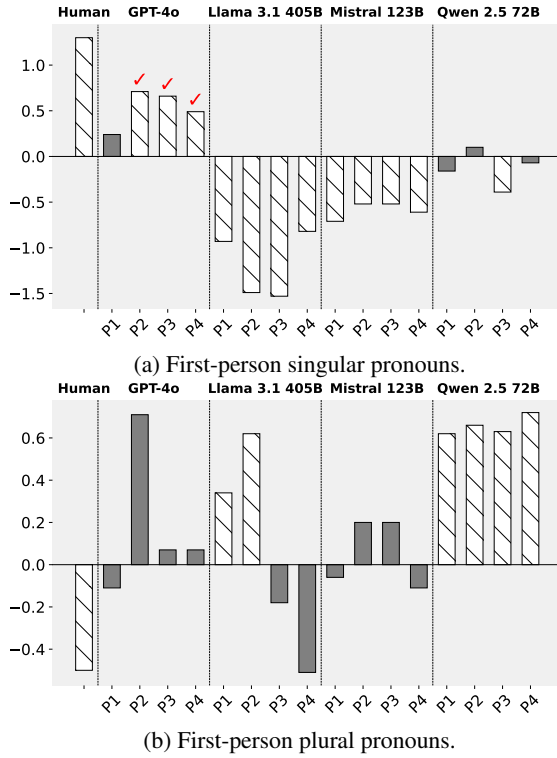


Figure 2: Results for the simple agent using LLMs from each model family versus human. P1 to P4 correspond to the prompts presented in Table 1. Gray bars indicate the result is not statistically significant. LLMs barely demonstrate human-like pronoun usage patterns in our experiments.

pattern to that of humans, all showing statistically significant positive results. Appendix A provides additional details on how we conduct the statistical significance test.

5 Findings

Figures 2 to 6 report the difference of first-person singular and first-person plural pronoun frequencies between non-leaders and leader’s scripts. In Figures 2 to 6, gray bars indicate that the results are not statistically significant. We place a red check mark on top of a bar if it demonstrates the same pattern as humans (e.g. in Figure 2, the bars corresponding to GPT-4o with Prompt 2, 3, 4 are positive with statistical significance, same as human). We report additional results corresponding to all pronouns, and an example of the model’s responses when it acts as different roles in Appendix B.

5.1 Do LLM agents demonstrate human-like pronoun usage patterns?

Answer: Barely. Figure 2 presents the results for GPT-4o, Llama 3.1 405B, Mistral 123B, and Qwen

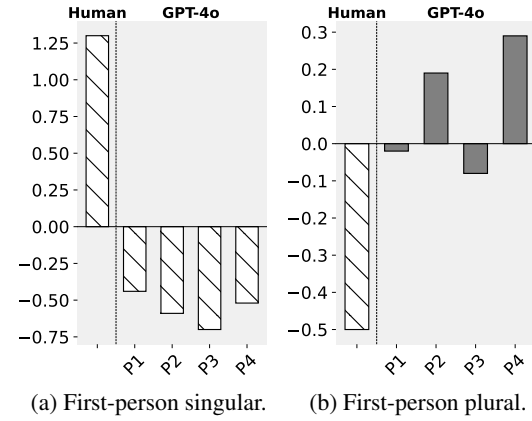


Figure 3: Results for the specialized agent (Section 3.2) based on GPT-4o. None of the prompts elicit pronoun usage patterns similar to those of humans. In most cases, the observed trends are contrary to human patterns.

2.5 72B models. Only GPT-4o, when prompted with Prompt 2, 3, and 4 (prompts in Table 1), exhibits human-like pronoun usage patterns for first-person singular pronouns. The other models do not show similar patterns to humans, regardless of the prompt. Furthermore, for first-person singular pronouns, Llama 3.1 405B and Mistral 123B display trends opposite to humans. For first-person plural pronouns, Llama 3.1 405B and Qwen 2.5 72B show trends contrary to humans. In Appendix B.3, we try to anonymize names, assign all the agents as males or females, or explicitly prompt the agent to use certain pronouns more often. However, None of these attempts succeeds in eliciting human-like pronoun usage patterns in LLMs. Therefore, we conclude that *interactions of agents based on LLMs from each LLM family barely resemble patterns of human interactions*. From a psychological standpoint, pronoun usage is closely tied to self-identity, social relationships, and group dynamics (Gumperz, 1982). Humans naturally adjust their use of singular and plural pronouns to convey individuality or solidarity, often without conscious effort. Such unconscious behaviors pose significant challenges for LLM agents to replicate in their simulation process.

5.2 Does the specialized agent demonstrate human-like pronoun usage patterns?

Answer: No. Figure 3 presents the results for the specialized agent described in Section 3.2 based on GPT-4o. We evaluate the effects of the four persona prompts listed in Table 1, none of which produce patterns resembling human behavior. Interestingly, the simple agent without additional components

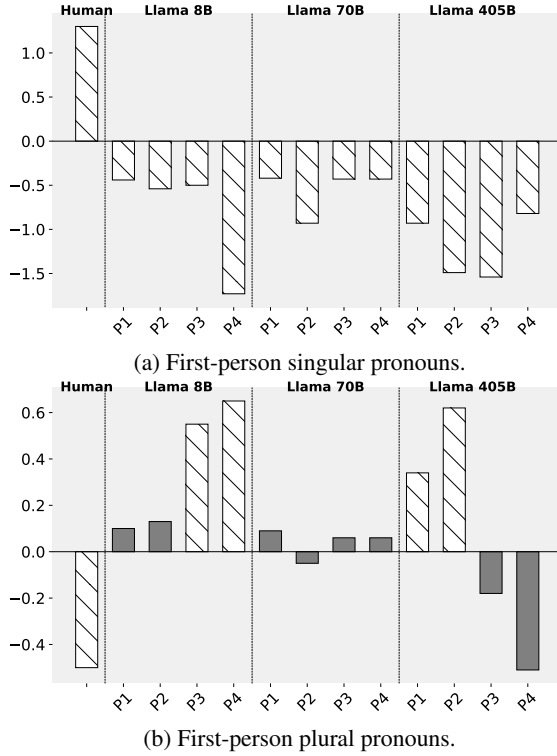


Figure 4: Results for the simple agent based on the Llama models. None of the settings elicit human-like behaviors.

enables GPT-4o to mimic human pronoun usage patterns with three out of the four prompts (Figure 2a). This suggests that *incorporating components like reflection or planning may not bridge the gap but instead deviate pronoun usage patterns further from human norms*.

From a cognitive perspective, this finding aligns with how intuitive processing is involved in human communication. Humans often rely on unconscious, automatic processes when engaging in social interactions, effortlessly navigating complex linguistic and social cues without explicit deliberation (Bargh, 1989; Bargh et al., 2012). By adding components like reflection or planning to the model, we may introduce a level of explicit reasoning that disrupts the natural flow of language generation, leading to less human-like interactions among the LLM agents.

5.3 Do LLMs within the same family exhibit consistent pronoun usage patterns?

Answer: Yes within the same version. For the LLMs from the same version, (e.g. Llama 3.1 or QWen 2.5), they show similar trends of pronoun usage patterns across different sizes. Figures 4 and 6 show the results for LLMs of different sizes

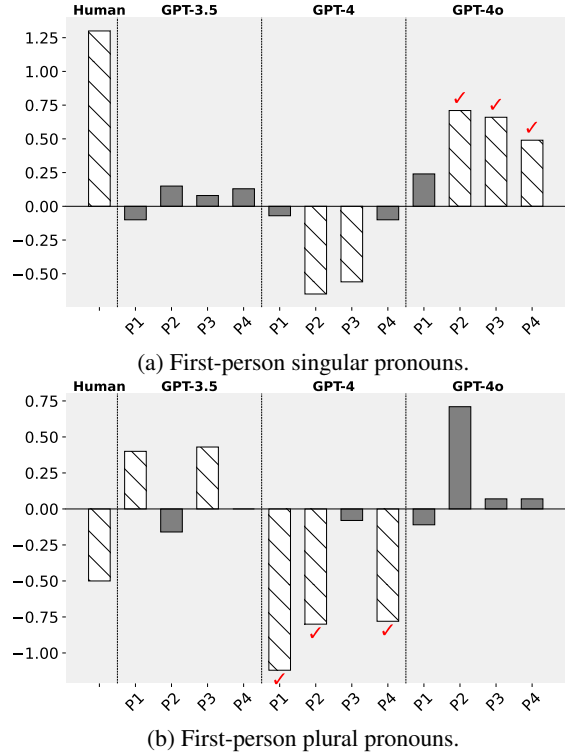


Figure 5: Results for the simple agent based on GPT models. For first-person singular pronouns, agents based on GPT-4o show the same trends as humans for Prompt 2, 3, and 4, while for first-person plural pronouns, agents based on GPT-4 show the same trends as humans for Prompt 1, 2 and 4.

from Llama 3.1 and QWen 2.5, respectively. In Figure 4a, all the Llama 3.1 models demonstrate a statistically significant trend opposite to the human interaction patterns, regardless of the prompts. In Figure 6a, most QWen 2.5 models demonstrate statistically insignificant results in terms of the first-person singular pronouns. We hypothesize that as models from the same version are trained on similar or the same data, such data similarity leads to the similar interaction patterns regardless of their sizes.

Figure 5 shows results for different versions of LLMs from the GPT family. For first-person singular pronouns, GPT-4o shows similar trends to humans with Prompt 2, 3, and 4, while for first-person plural pronouns, GPT-4 shows similar trends to humans with Prompt 1, 2, and 4. In contrast, none of the prompts elicit human-like patterns for GPT-3.5. We hypothesize that as GPT-4 and GPT-4o are more capable than GPT-3.5, they may be better at role-playing and demonstrating human-like pronoun usage patterns.

In addition, *the larger LLM from the same family does not necessarily elicit more human-like pronoun usage patterns*. For instance, in Figure 6b,

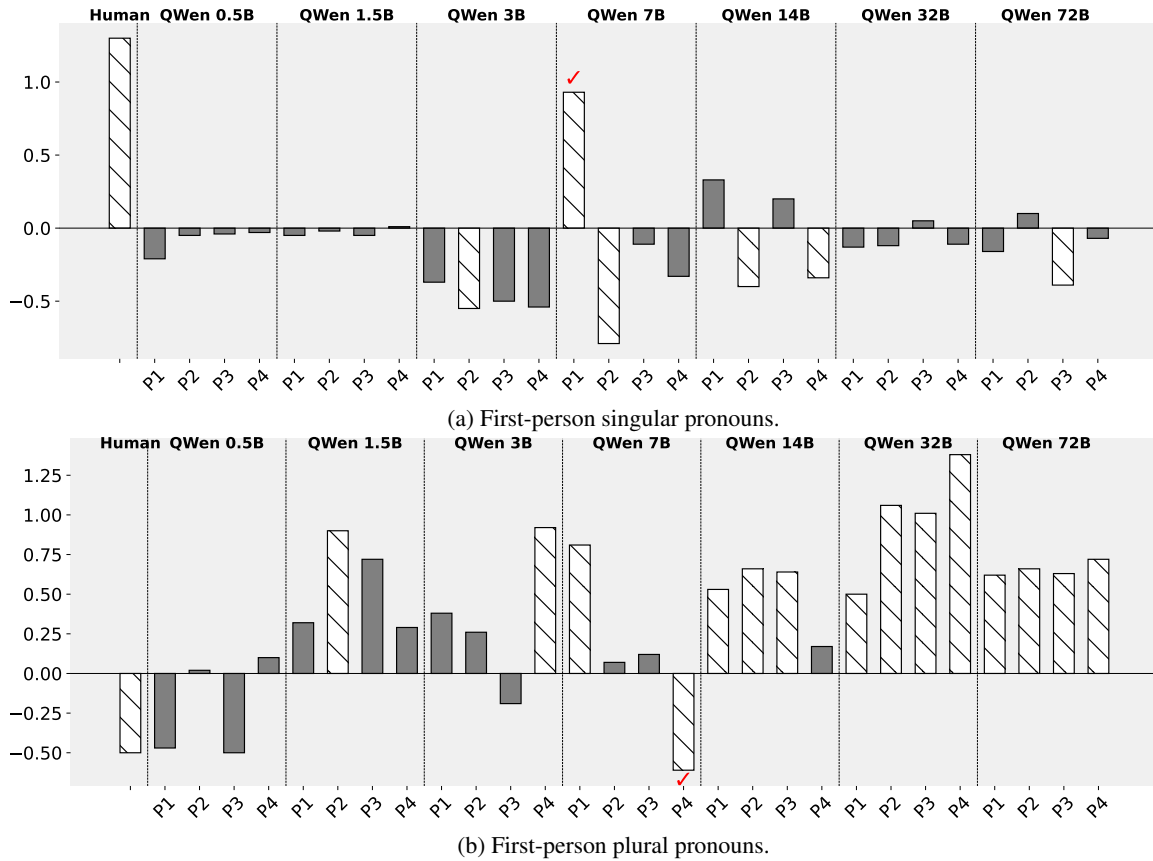


Figure 6: Results for the simple agent using LLMs from the QWen 2.5 family. Only the 7B model with Prompt 1 elicits human-like interaction patterns for first-person singular pronouns, and the 7B model with Prompt 4 for first-person plural pronouns.

the QWen 72B models demonstrate an opposite pronoun usage pattern compared to humans regardless of the prompts. In contrast, though in most cases, QWen models fail to elicit human-like pronoun patterns, the QWen 7B model with Prompt 4 elicits human-like patterns for the first-person plural pronouns.

5.4 Do LLMs “know” the pronoun patterns better than what they demonstrate in interactions?

Answer: Yes. To understand whether these LLMs “know” the underlying social interaction patterns, we employ a variety of prompts based on the following prompt template,

Query Prompt

For the frequency of the usage for {PRONOUN TYPE}, who uses it at a higher frequency in their interactions? Please select one of the three answers: leader, subordinate, equal, please answer in JSON format of {“answer”: <Your Choice>}

Specifically, as the answer orders may influence the results (Zheng et al., 2023), to make the results robust, we permute the answers to prompt each LLM on six different orders for “leader, subordinate, equal” (e.g. “leader, equal, subordinate”, “subordinate, leader, equal”, etc.). In Table 2, we combine our results here with the results from Section 5 on whether the simple agents based on each LLM know and demonstrate human-like pronoun usage patterns.

In Table 2, we find that GPT-4o succeeds five out of six times in selecting the right option for first-person singular pronouns, and succeeds in demonstrating such a pattern in its interaction for three out of the four prompts. However, for most LLMs, we find that *they may know the human pronoun usage patterns, but do not demonstrate them in their interactions*. For instance, LLMs such as Llama 70B can correctly choose the option in all six cases for the first-person singular pronoun, but fail to demonstrate such patterns across the four prompts. Wu et al. (2023) present a similar finding of the disparity between the LLM’s knowledge versus their

Models	First-Person Singular		First-Person Plural	
	kn.?	dem.?	kn.?	dem.?
Gold	6	4	6	4
<i>GPT</i>				
3.5	0	0	2	0
4	6	0	3	3
4o	5	3	5	0
<i>Llama 3.1 Instruct</i>				
8B	0	0	4	0
70B	6	0	0	0
405B	1	0	3	0
<i>Qwen 2.5 Instruct</i>				
0.5B	0	0	0	0
1.5B	0	0	2	0
3B	1	0	2	0
7B	5	1	5	1
14B	5	0	6	0
32B	3	0	0	0
72B	6	0	3	0
<i>Mistral Instruct</i>				
123B	3	0	5	0

Table 2: “Kn.” denotes with respect to the perturbed prompts, how many times the LLM answers the question correctly. “Dem.” denotes with respect to the four persona prompts, how many times the agents based on each LLM demonstrate the human-like pronoun pattern during the interaction process. The disparity of “kn.” and “dem.” in LLMs demonstrates that even if the LLM knows the pronoun patterns, they do not necessarily demonstrate them in the interaction process.

ability to effectively apply that knowledge in their reasoning process. Similarly, our findings reveal the disparity between LLMs’ knowledge versus demonstrating it in practice.

6 Discussions

Taking our results together, firstly, *the LLMs’ pronoun usage patterns are susceptible to model choices and prompt designs*. In Figure 2, GPT-4o fails to demonstrate human-like interaction patterns with Prompt 1 even though it succeeds with Prompt 2, 3, and 4. In Figure 5, though Prompt 1 elicits human-like pronoun usage patterns for the GPT-4 model on first-person plural pronouns, it fails for the GPT-4o model. Such a model and prompt-susceptible pattern is significantly different from what we expect in humans. A rational human being would behave as instructed regardless of the wording changes, as long as the core information of playing a {PERSONA} are the same within those instructions. Such a discrepancy poses

concerns about how much we can rely on LLMs for trustworthy human interaction simulations.

Secondly, *for researchers who study social problems through the LLM simulations, they need to be mindful that they are studying LLMs’ stochastic group behaviors rather than human behaviors*. While currently, the community seems to overly focus on the promising “emergent social behaviors” and “social prototyping” from Park et al. (2022, 2023), we want to emphasize the unpleasant side of the story, that these LLMs are not humans, and the “emergent social behaviors” are essentially the LLMs’ social behaviors rather than humans’. In our case, if a researcher is to study the “emergent pronoun usage” behaviors of GPT-3.5 in a society of hierarchical structure (the leaders and non-leaders), their conclusion may flip the well-established findings from the psychology community. This poses concerns to the existing studies where researchers employ LLMs for social simulation, that how authentic and reliable the intermediate process is in the simulation. Therefore, we advocate for researchers to investigate the underlying interaction mechanisms carefully before they generalize their conclusions to human interactions or human society.

A concurrent study by Park et al. (2024) demonstrates that, through careful prompt design, LLM agents can replicate participants’ responses with 85% accuracy. While their work focuses on simulating individual survey responses, our research emphasizes pronoun usage in LLM interactions. Notably, despite their optimization efforts, Park et al. (2024) only achieve 85% accuracy on 1,000 individuals, raising questions about the generalizability of their findings to larger populations. Furthermore, the underlying reasoning process of LLMs remains opaque, leaving questions about what features LLMs leverage in achieving such results.

7 Conclusion

In this paper, we take the first step to study whether LLM agents can demonstrate human-like pronoun usage patterns in the multi-agent interaction process. By simulating experimental setups from an existing psychology study, we reveal that LLM agents barely demonstrate human-like pronoun patterns, even if the LLM agent may show some understanding of those patterns. We advocate practitioners scrutinize their simulations and understand the limitations of the LLM-based social simulation.

Limitations

We address two primary concerns raised about our work. First, one might argue that differences in singular or plural pronoun usage are merely formatting issues and do not impact the final simulation outcomes. However, we emphasize that if LLMs fail to replicate these subtle but critical markers in human interactions, it raises concerns about their ability to authentically model the underlying human interaction process or the complex social dynamics. Therefore, we argue that researchers need to be cautious about drawing conclusions from the “emergent social behaviors” of LLM agents’ interactions.

Second, some may suggest that better prompt engineering could elicit more human-like pronoun usage. We want to emphasize that for the scope of this work, we systematically evaluate existing prompts employed by researchers, and they consistently fail to elicit human-like leader pronoun patterns. If a researcher simply relies on these prompts and LLM agents to simulate human society, they may draw conclusions on pronoun usage patterns that contradict to human societal norms. In addition, we highlight that when humans act as leaders, they cannot control the frequency of “we” or “us” consciously and their behaviors are not influenced by explicit “prompt design”. Even if optimized prompts can achieve human-like behaviors in LLMs, the optimization process itself underscores the differences between LLMs and human beings. Finally, such optimization is likely to create a “rob Peter to pay Paul” scenario, where improving performance for specific cases results in failures in others, limiting the generalizability of LLMs as robust social simulators.

We try our best to experiment with different LLMs. However, due to the scope of this study and the limited budget, we cannot test every possible LLM model. We believe the experiments we conduct in our paper provide strong evidence to our claims.

Ethics Statement

In this work, we simulate an existing psychology study by employing LLM agents. Our results indicate that there is a significant discrepancy in pronoun usage patterns between LLM agents and humans. We suggest researchers realize the limitations of employing LLM agents to simulate human interactions or human societies, as there are hu-

man patterns that cannot be easily and consistently replicated by existing LLMs.

References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. [arXiv preprint arXiv:2303.08774](#).
- Gati V Aher, Rosa I Arriaga, and Adam Tauman Kalai. 2023. Using large language models to simulate multiple humans and replicate human subject studies. In *International Conference on Machine Learning*, pages 337–371. PMLR.
- Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, et al. 2023. Qwen technical report. [arXiv preprint arXiv:2309.16609](#).
- John A Bargh. 1989. Conditional automaticity: Varieties of automatic influence in social perception and cognition. *Unintended thought*, pages 3–51.
- John A Bargh, Kay L Schwader, Sarah E Hailey, Rebecca L Dyer, and Erica J Boothby. 2012. Automaticity in social-cognitive processes. *Trends in cognitive sciences*, 16(12):593–605.
- Angana Borah and Rada Mihalcea. 2024. [Towards implicit bias detection and mitigation in multi-agent llm interactions](#).
- Marilynn B Brewer and Wendi Gardner. 1996. Who is this “we”? levels of collective identity and self representations. *Journal of personality and social psychology*, 71(1):83.
- Ethan R Burris, Matthew S Rodgers, Elizabeth A Manix, Michael G Hendron, and James B Oldroyd. 2009. Playing favorites: The influence of leaders’ inner circle on group processes and performance. *Personality and Social Psychology Bulletin*, 35(9):1244–1257.
- R Sherlock Campbell and James W Pennebaker. 2003. The secret life of pronouns: Flexibility in writing style and physical health. *Psychological science*, 14(1):60–65.
- Shuaichen Chang and Eric Fosler-Lussier. 2023. How to prompt llms for text-to-sql: A study in zero-shot, single-domain, and cross-domain settings. [arXiv preprint arXiv:2305.11853](#).
- Deborah Davis and Timothy C Brock. 1975. Use of first person pronouns as a function of increased objective self-awareness and performance feedback. *Journal of Experimental Social Psychology*, 11(4):381–388.
- Naihao Deng, Zhenjie Sun, Ruiqi He, Aman Sikka, Yulong Chen, Lin Ma, Yue Zhang, and Rada Mihalcea. 2024. [Tables as texts or images: Evaluating the table reasoning ability of LLMs and MLLMs](#).

- In *Findings of the Association for Computational Linguistics ACL 2024*, pages 407–426, Bangkok, Thailand and virtual meeting. Association for Computational Linguistics.
- Ameet Deshpande, Vishvak Murahari, Tanmay Rajpurohit, Ashwin Kalyan, and Karthik Narasimhan. 2023. *Toxicity in chatgpt: Analyzing persona-assigned language models*. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 1236–1270, Singapore. Association for Computational Linguistics.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Shelley Duval and Robert A Wicklund. 1972. A theory of objective self awareness.
- Elias Frantar, Saleh Ashkboos, Torsten Hoeftler, and Dan Alistarh. 2022. Gptq: Accurate post-training quantization for generative pre-trained transformers. *arXiv preprint arXiv:2210.17323*.
- Chen Gao, Xiaochong Lan, Zhihong Lu, Jinzhu Mao, Jinghua Piao, Huandong Wang, Depeng Jin, and Yong Li. 2023. S³: Social-network simulation system with large language model-empowered agents. *arXiv preprint arXiv:2307.14984*.
- John J Gumperz. 1982. *Language and social identity*. Cambridge University Press.
- Shashank Gupta, Vaishnavi Shrivastava, Ameet Deshpande, Ashwin Kalyan, Peter Clark, Ashish Sabharwal, and Tushar Khot. 2023. Bias runs deep: Implicit reasoning biases in persona-assigned llms. *arXiv preprint arXiv:2311.04892*.
- Wenyue Hua, Lizhou Fan, Lingyao Li, Kai Mei, Jianchao Ji, Yingqiang Ge, Libby Hemphill, and Yongfeng Zhang. 2023. War and peace (waragent): Large language model-based multi-agent simulation of world wars. *arXiv preprint arXiv:2311.17227*.
- Molly E Ireland and Matthias R Mehl. 2014. Natural language use as a marker. *The Oxford handbook of language and social psychology*, pages 201–237.
- Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023a. Mistral 7b. *arXiv preprint arXiv:2310.06825*.
- Hang Jiang, Xiajie Zhang, Xubo Cao, Jad Kabbara, and Deb Roy. 2023b. Personallm: Investigating the ability of gpt-3.5 to express personality traits and gender differences. *arXiv preprint arXiv:2305.02547*.
- Mingyu Jin, Beichen Wang, Zhaoqian Xue, Suiyuan Zhu, Wenyue Hua, Hua Tang, Kai Mei, Mengnan Du, and Yongfeng Zhang. 2024. What if llms have different world views: Simulating alien civilizations with llm-based agents. *arXiv preprint arXiv:2402.13184*.
- Ewa Kaciewicz, James W Pennebaker, Matthew Davis, Moongee Jeon, and Arthur C Graesser. 2014. Pronoun use reflects standings in social hierarchies. *Journal of Language and Social Psychology*, 33(2):125–143.
- Zhao Kaiya, Michelangelo Naim, Jovana Kondic, Manuel Cortes, Jiaxin Ge, Shuying Luo, Guangyu Robert Yang, and Andrew Ahn. 2023. Lyfe agents: Generative agents for low-cost real-time social interactions. *arXiv preprint arXiv:2310.02172*.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph Gonzalez, Hao Zhang, and Ion Stoica. 2023. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the 29th Symposium on Operating Systems Principles*, pages 611–626.
- Kenneth Li, Tianle Liu, Naomi Bashkinsky, David Bau, Fernanda Viégas, Hanspeter Pfister, and Martin Wattenberg. 2024a. Measuring and controlling persona drift in language model dialogs. *arXiv preprint arXiv:2402.10962*.
- Nian Li, Chen Gao, Mingyu Li, Yong Li, and Qingmin Liao. 2024b. *EconAgent: Large language model-empowered agents for simulating macroeconomic activities*. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15523–15536, Bangkok, Thailand. Association for Computational Linguistics.
- Yuan Li, Yixuan Zhang, and Lichao Sun. 2023. Metaagents: Simulating interactions of human behaviors for llm-based task-oriented coordination via collaborative generative agents. *arXiv preprint arXiv:2310.06500*.
- Matthias R Mehl, Samuel D Gosling, and James W Pennebaker. 2006. Personality in its natural habitat: manifestations and implicit folk theories of personality in daily life. *Journal of personality and social psychology*, 90(5):862.
- Joon Sung Park, Joseph O’Brien, Carrie Jun Cai, Meredith Ringel Morris, Percy Liang, and Michael S Bernstein. 2023. Generative agents: Interactive simulacra of human behavior. In *Proceedings of the 36th annual acm symposium on user interface software and technology*, pages 1–22.
- Joon Sung Park, Lindsay Popowski, Carrie Cai, Meredith Ringel Morris, Percy Liang, and Michael S Bernstein. 2022. Social simulacra: Creating populated prototypes for social computing systems. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology*, pages 1–18.

713	Joon Sung Park, Carolyn Q Zou, Aaron Shaw, Benjamin Mako Hill, Carrie Cai, Meredith Ringel Morris, Robb Willer, Percy Liang, and Michael S Bernstein. 2024. Generative agent simulations of 1,000 people. arXiv preprint arXiv:2411.10109 .	769
714		770
715		771
716		772
717		773
718	James W Pennebaker. 2011. The secret life of pronouns. New Scientist , 211(2828):42–45.	
719		
720	Alastair Pennycook. 1994. The politics of pronouns.	
721	Giorgio Piatti, Zhijing Jin, Max Kleiman-Weiner, Bernhard Schölkopf, Mrinmaya Sachan, and Rada Mihalcea. 2024. Cooperate or collapse: Emergence of sustainability behaviors in a society of llm agents. arXiv preprint arXiv:2404.16698 .	774
722		775
723		776
724		777
725		778
726	Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. OpenAI blog , 1(8):9.	779
727		780
728		
729		
730	Leonard Salewski, Stephan Alaniz, Isabel Rio-Torto, Eric Schulz, and Zeynep Akata. 2024. In-context impersonation reveals large language models’ strengths and biases. Advances in Neural Information Processing Systems , 36.	
731		
732		
733		
734		
735	Jan E Stets and Chris F Biga. 2003. Bringing identity theory into environmental sociology. Sociological theory , 21(4):398–423.	
736		
737		
738	Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. Llama: Open and efficient foundation language models. arXiv preprint arXiv:2302.13971 .	
739		
740		
741		
742		
743		
744	Mike Treanor, Ben Samuel, and Mark J Nelson. 2024. Prototyping slice of life: Social physics with symbolically grounded llm-based generative dialogue. In Proceedings of the 19th International Conference on the Foundations of Digital Games , pages 1–4.	
745		
746		
747		
748		
749	Noah Wang, Z.y. Peng, Haoran Que, Jiaheng Liu, Wangchunshu Zhou, Yuhan Wu, Hongcheng Guo, Ruitong Gan, Zehao Ni, Jian Yang, Man Zhang, Zhaoxiang Zhang, Wanli Ouyang, Ke Xu, Wenhao Huang, Jie Fu, and Junran Peng. 2024a. RoleLLM: Benchmarking, eliciting, and enhancing role-playing abilities of large language models . In Findings of the Association for Computational Linguistics ACL 2024 , pages 14743–14777, Bangkok, Thailand and virtual meeting. Association for Computational Linguistics.	
750		
751		
752		
753		
754		
755		
756		
757		
758		
759		
760	Shuai Wang, Harris Scells, Bevan Koopman, and Guido Zuccon. 2023. Can chatgpt write a good boolean query for systematic review literature search? In Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval , pages 1426–1436.	
761		
762		
763		
764		
765		
766	Yudong Wang, Damai Dai, and Zhifang Sui. 2024b. Exploring activation patterns of parameters in language models. arXiv preprint arXiv:2405.17799 .	
767		
768		
	Jules White, Quchen Fu, Sam Hays, Michael Sandborn, Carlos Olea, Henry Gilbert, Ashraf Elnashar, Jesse Spencer-Smith, and Douglas C Schmidt. 2023. A prompt pattern catalog to enhance prompt engineering with chatgpt. arXiv preprint arXiv:2302.11382 .	
	Yufan Wu, Yinghui He, Yilin Jia, Rada Mihalcea, Yulong Chen, and Naihao Deng. 2023. Hi-ToM: A benchmark for evaluating higher-order theory of mind reasoning in large language models . In Findings of the Association for Computational Linguistics: EMNLP 2023 , pages 10691–10706, Singapore. Association for Computational Linguistics.	
	Benfeng Xu, An Yang, Junyang Lin, Quan Wang, Chang Zhou, Yongdong Zhang, and Zhendong Mao. 2023. Expertprompting: Instructing large language models to be distinguished experts. arXiv preprint arXiv:2305.14688 .	781
		782
		783
		784
		785
	Zhengyan Zhang, Yankai Lin, Zhiyuan Liu, Peng Li, Maosong Sun, and Jie Zhou. 2022. MoEffi-ciation: Transformer feed-forward layers are mix-tures of experts . In Findings of the Association for Computational Linguistics: ACL 2022 , pages 877–890, Dublin, Ireland. Association for Computational Linguistics.	786
		787
		788
		789
		790
		791
		792
	Zheyuan Zhang, Daniel Zhang-Li, Jifan Yu, Linlu Gong, Jinchang Zhou, Zhiyuan Liu, Lei Hou, and Juanzi Li. 2024. Simulating classroom education with llm-empowered agents. arXiv preprint arXiv:2406.19226 .	793
		794
		795
		796
		797
	Chujie Zheng, Hao Zhou, Fandong Meng, Jie Zhou, and Minlie Huang. 2023. Large language models are not robust multiple choice selectors. In The Twelfth International Conference on Learning Representations .	798
		799
		800
		801
		802

A Additional Experimental Setups

Simulation Scenarios. Table 3 presents the pools for male names, female names, company names, and ranking items where we draw to populate the masked sections in the instructions, together with the instructions that we use.

In our simulation, the group is charged with completing a group decision-making task, where they act as consulting groups to improve the customer service quality of a fictitious company. Specifically, the agents need to rank items (e.g. “Invest in training customer service representatives”, “Personalize customer interactions”, etc.) in order of their positive contribution. All the agents are aware that the team decision does not need to be a consensus, and the leader ultimately has the authority to make the final decision.

Model Hyperparameters. For all the models, we employ a temperature of 0.7 to ensure diversity in their generations.

Statistical Significance Test. We collect the pronoun usage frequencies for the leader and non-leader agents during their interactions. Since we repeat our experiments 41 times consistent with [Burris et al. \(2009\)](#); [Kacewicz et al. \(2014\)](#), we have 41 frequency values for leaders and non-leaders respectively for one type of pronoun. We then conduct the statistical significance T-test to calculate the two samples of frequencies to get the p values.

B Additional Experimental Results

B.1 Absolute Pronoun Usage Results.

Table 4e shows the results we get for each pronoun usage by employing GPT-4o as the agent. We structure each table in the same format as Table 4a from [Kacewicz et al. \(2014\)](#). We note that we report the frequency in terms of percentage ($f \in [0, 100]$) following [Kacewicz et al. \(2014\)](#).

B.2 Complete Results for Pronoun Usage Difference

Table 5 presents the complete pronoun usage difference results corresponding to Figure 2 and Figure 3. We can see that neither the LLM agents from different model families nor the specialized agents can consistently elicit human-like pronoun usage in their interaction process. In most cases, their patterns contradict human patterns.

Tables 6 to 8 present the results for models from GPT family (corresponding to Figure 5), Llama

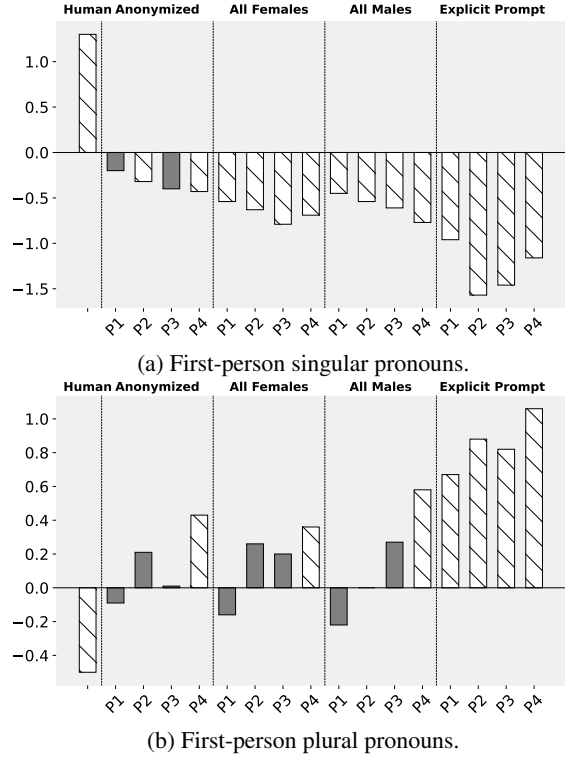


Figure 7: Differences in pronoun usage when names are anonymized, all agents are assigned female or male genders or an explicit prompt is used (described in Appendix B.3). We conduct the experiments based on GPT-4o. None of these adjustments can elicit human-like pronoun usage patterns.

family (corresponding to Figure 4) and QWen family (corresponding to Figure 6), respectively. We observe that there are model-specific patterns across different prompts. For instance, leader agents tend to use first-person singular pronouns more often across Llama 3.1 8B models, while the first-person singular pronoun usage difference is not statistically significant across all the Qwen 2.5 models.

B.3 Other Effects

We explore the effects of various settings based on the four prompts in Table 1 using GPT-4o and report our results for the first-person singular and plural pronouns in Figure 7. For the setting of “Anonymized”, we anonymize the names of agents by names such as “Name 1”, “Name 2”, etc. For the settings of “All Females” or “All Males”, we assign the genders of all the LLM agents as females or males, respectively. For the setting of “Explicit Prompt”, we add a sentence of “Please use first-person plural forms (e.g. we, us) more often” if the agent acts as leaders, and “Please use first-person

Pools	
Male Names	James, John, William, Robert, Michael, David, Richard, Joseph, Charles, Thomas, Daniel, Matthew, Anthony, Mark, Andrew, Joshua, Christopher, Jonathan, Benjamin, Samuel, Henry, Alexander, George, Edward, Nicholas, Luke, Jack, Ethan, Oliver, Jacob
Female Names	Emma, Olivia, Sophia, Ava, Isabella, Mia, Amelia, Harper, Evelyn, Abigail, Emily, Ella, Grace, Lily, Chloe, Scarlett, Victoria, Madison, Aria, Zoey, Nora, Charlotte, Hannah, Lucy, Zoe, Layla, Penelope, Lillian, Natalie, Violet, Stella
Company Names	Company Mars, Company Zelda, Company Moon, Company Jupiter
Ranking Items	<ol style="list-style-type: none"> 1. Invest in training customer service representatives. 2. Embrace omnichannel support. 3. Collect customer feedback. 4. Use AI and automation tools. 5. Personalize customer interactions. 6. Provide comprehensive FAQs. 7. Monitor and respond to social media.
Phase I	
Instruction	<p>Now your team is tasked with improving the customer service quality of {COMPANY}. Given the following ten ranking items: {RANKING_ITEMS}</p> <p>Now please work on your own and rank these proposed items in order of their positive contribution to the company culture.</p> <p>Please speak in the format of: {{{NAME}} ({ROLE}): <Your Speech>}}</p>
Phase II	
Instruction (Leader)	<p>Now you are discussing in a group. The topic is recommendation of strategies for service quality improvement by ranking these proposed items in order of their positive contribution to such the company culture for {company}. The final team decision does not need to be a consensus. You as a leader ultimately have the authority to make the final decision, and everyone on the team is clear about it. Please speak in the format of: {{{NAME}} ({ROLE}): <Your Speech>}}</p>
Instruction (Non-Leader)	<p>Now you are discussing in a group. The topic is recommendation of strategies for service quality improvement by ranking these proposed items in order of their positive contribution to such the company culture for {company}. The final team decision does not need to be a consensus. The leader ultimately has the authority to make the final decision, you and everyone on the team are clear about it. Please speak in the format of: {{{NAME}} ({ROLE}): <Your Speech>}}</p>

Table 3: Pools of items that we use to populate the masked sections in instructions, along with the instructions that we use in the two experimental phases. The setup involves two roles: leader and subordinate.

singular forms (e.g. I, me) more often” if the agent acts as non-leaders. We highlight that none of these adjustments can elicit human-like pronoun usage patterns.

B.4 Examples of Model Responses

We observe that there is minimal distinction in language use when the model assumes the role of a leader versus a non-leader. For instance, when the GPT-4o assumes the role of leader and non-leader in Table 9, both responses exhibit a collaborative and inclusive tone, but without a clear distinction between the roles.

C Dialogue Examples

Table 10 presents an example from the original human study by Kacewicz et al. (2014) and an example from the dialogue by LLM agents using GPT-4o with Prompt 1.

	Examples	Leader	Nonleaders	p
Word Counts	-	1076.90	645.00	.00
All pronouns	-	19.80	20.80	.01
Personal pronouns	I, we, she	10.10	10.60	.06
First-person singular	I, me	4.30	5.60	.00
First-person plural	We, us	2.10	1.60	.00
Second-person	You, your	2.70	2.50	.20
Third-person singular	She, he, her	0.07	0.09	.50
Third-person plural	They, them	0.83	0.85	.84
Impersonal pronouns	It, that	9.70	10.10	.06

(a) Human (Kacewicz et al., 2014).

	Examples	Leader	Nonleaders	p
Word Counts	-	2242.63	1946.86	.21
All pronouns	-	6.80	6.89	.98
Personal pronouns	I, we, she	5.99	5.96	.73
First-person singular	I, me	1.84	2.08	.08
First-person plural	We, us	2.54	2.43	.73
Second-person	You, your	0.98	0.75	.04
Third-person singular	She, he, her	0.03	0.03	.78
Third-person plural	They, them	0.61	0.67	.61
Impersonal pronouns	It, that	0.81	0.93	.56

(b) GPT-4o, prompt 1.

	Examples	Leader	Nonleaders	p
Word Counts	-	1270.44	1021.36	.07
All pronouns	-	6.35	7.32	.01
Personal pronouns	-	5.84	6.36	.04
First-person singular	-	1.11	1.82	.00
First-person plural	-	3.09	2.95	.41
Second-person	-	1.14	0.81	.02
Third-person singular	-	0.06	0.06	.12
Third-person plural	-	0.44	0.72	.00
Impersonal pronouns	-	0.65	0.96	.00

(c) GPT-4o, prompt 2.

	Examples	Leader	Nonleaders	p
Word Counts	-	1609.46	1166.73	.01
All pronouns	-	6.32	6.98	.00
Personal pronouns	-	5.57	6.11	.00
First-person singular	-	1.40	2.06	.00
First-person plural	-	2.50	2.57	.60
Second-person	-	1.22	0.84	.00
Third-person singular	-	0.06	0.03	.02
Third-person plural	-	0.38	0.61	.00
Impersonal pronouns	-	0.75	0.87	.17

(d) GPT-4o, prompt 3.

	Examples	Leader	Nonleaders	p
Word Counts	-	1578.83	1371.24	.30
All pronouns	-	5.98	6.45	.06
Personal pronouns	-	5.15	5.61	.04
First-person singular	-	1.42	1.91	.00
First-person plural	-	2.20	2.27	.63
Second-person	-	1.07	0.80	.00
Third-person singular	-	0.01	0.01	.83
Third-person plural	-	0.46	0.62	.03
Impersonal pronouns	-	0.82	0.85	.81

(e) GPT-4o, prompt 4.

Table 4: Complete results of employing GPT-4o as the agent. Since we primarily care about whether there is a statistically significant difference in the pronoun usage between leaders and non-leaders, we decide to use the difference rather than the absolute frequency as our evaluation metrics in Equation (1).

	All Pronouns	Personal Pronouns	First-Person Singular	First-Person Plural	Second- Person	Third-Person Singular	Third-Person Plural	Impersonal Pronouns
Human (2014)	1.00	0.50	1.30	−0.50	−0.20	0.02	0.02	0.40
<i>GPT-4o</i>								
Prompt 1	0.09	−0.03	0.24	−0.11	−0.23	0.00	0.06	0.12
Prompt 2	0.97	0.62	0.71	0.71	−0.33	0.00	0.28	0.31
Prompt 3	0.66	0.54	0.66	0.07	−0.38	−0.03	0.23	0.12
Prompt 4	0.47	0.46	0.49	0.07	−0.27	0.00	0.18	0.03
<i>Llama 3.1 405B Instruct</i>								
Prompt 1	1.82	0.16	−0.93	0.34	0.80	−0.01	−0.03	−0.04
Prompt 2	1.94	0.16	−1.49	0.62	0.82	−0.09	0.30	−0.13
Prompt 3	2.58	0.35	−1.53	−0.18	2.41	−0.09	−0.26	−0.29
Prompt 4	4.78	−0.40	−0.82	−0.51	1.04	0.01	−0.12	−0.40
<i>Mistral 128B Large Instruct</i>								
Prompt 1	0.46	−0.09	−0.71	−0.06	0.83	0.04	−0.19	−0.21
Prompt 2	1.09	0.35	−0.52	0.20	0.60	0.02	0.06	−0.49
Prompt 3	1.09	0.35	−0.52	0.20	0.60	0.02	0.06	−0.49
Prompt 4	0.88	−0.39	−0.61	−0.11	0.45	0.02	−0.14	−0.22
<i>Qwen 2.5 72B Instruct</i>								
Prompt 1	−0.09	0.83	−0.16	0.62	0.56	−0.03	−0.15	−0.43
Prompt 2	0.21	2.09	0.10	0.66	1.84	−0.02	−0.48	−0.66
Prompt 3	0.01	0.87	−0.39	0.63	0.86	0.00	−0.22	−0.40
Prompt 4	−0.27	1.06	−0.07	0.72	0.47	−0.02	−0.03	−0.20
<i>GPT-4o Framework Agent</i>								
Prompt 1	0.56	−0.11	−0.44	−0.02	0.41	−0.00	−0.05	−0.12
Prompt 2	2.11	−0.12	−0.59	0.19	0.35	0.04	−0.10	−0.46
Prompt 3	0.45	−0.20	−0.70	−0.08	0.77	0.02	−0.21	−0.26
Prompt 4	0.09	0.37	−0.52	0.29	0.88	0.05	−0.34	−0.17

Table 5: Results corresponding to Figure 2 and Figure 3. We report the difference in the average pronoun frequencies between non-leaders and leader’s scripts, and put the number in gray if there is no statistical significance, i.e., the p -value is smaller than 0.05.

	All Pronouns	Personal Pronouns	First-Person Singular	First-Person Plural	Second- Person	Third-Person Singular	Third-Person Plural	Impersonal Pronouns
Human (2014)	1.00	0.50	1.30	−0.50	−0.20	0.02	0.02	0.40
<i>GPT-3.5</i>								
Prompt 1	−0.49	0.52	−0.10	0.40	0.19	−0.02	0.05	−0.03
Prompt 2	0.19	0.07	0.15	−0.16	0.18	−0.04	−0.06	−0.09
Prompt 3	−0.19	0.59	0.08	0.43	0.27	−0.03	−0.17	−0.15
Prompt 4	−0.35	0.44	0.13	−0.00	0.55	−0.01	−0.23	0.05
<i>GPT-4</i>								
Prompt 1	4.44	−0.39	−0.07	−1.12	1.14	0.01	−0.34	−0.65
Prompt 2	1.57	0.51	−0.65	−0.80	2.04	0.06	−0.15	−0.44
Prompt 3	0.83	0.20	−0.56	−0.08	1.02	0.02	−0.19	−0.24
Prompt 4	2.86	−0.04	−0.10	−0.78	0.83	−0.01	0.02	−0.46
<i>GPT-4o</i>								
Prompt 1	0.09	−0.03	0.24	−0.11	−0.23	0.00	0.06	0.12
Prompt 2	0.97	0.62	0.71	0.71	−0.33	0.00	0.28	0.31
Prompt 3	0.66	0.54	0.66	0.07	−0.38	−0.03	0.23	0.12
Prompt 4	0.47	0.46	0.49	0.07	−0.27	0.00	0.18	0.03

Table 6: Results for LLM agents based on GPT models corresponding to Figure 5.

	All Pronouns	Personal Pronouns	First-Person Singular	First-Person Plural	Second- Person	Third-Person Singular	Third-Person Plural	Impersonal Pronouns
Human (2014)	1.00	0.50	1.30	−0.50	−0.20	0.02	0.02	0.40
<i>Llama 3.1 8B Instruct</i>								
Prompt 1	0.81	0.29	−0.44	0.10	0.56	0.01	0.05	−0.24
Prompt 2	0.69	−0.18	−0.54	0.13	0.25	0.03	−0.05	−0.34
Prompt 3	0.81	0.54	−0.50	0.55	0.66	0.02	−0.18	−0.36
Prompt 4	2.63	−0.49	−1.73	0.65	0.43	0.03	0.14	−0.52
<i>Llama 3.1 70B Instruct</i>								
Prompt 1	0.64	−0.17	−0.42	0.09	0.24	−0.00	−0.09	−0.10
Prompt 2	1.44	−0.42	−0.93	−0.05	0.36	0.00	0.19	−0.18
Prompt 2	−0.13	0.40	−0.43	0.06	0.74	0.08	−0.03	−0.06
Prompt 3	−0.13	0.40	−0.43	0.06	0.74	0.08	−0.03	−0.06
<i>Llama 3.1 405B Instruct</i>								
Prompt 1	1.82	0.16	−0.93	0.34	0.80	−0.01	−0.03	−0.04
Prompt 2	1.94	0.16	−1.49	0.62	0.82	−0.09	0.30	−0.13
Prompt 3	2.58	0.35	−1.53	−0.18	2.41	−0.09	−0.26	−0.29
Prompt 4	4.78	−0.40	−0.82	−0.51	1.04	0.01	−0.12	−0.40

Table 7: Results for LLM agents based on Llama models corresponding to Figure 4.

	All Pronouns	Personal Pronouns	First-Person Singular	First-Person Plural	Second- Person	Third-Person Singular	Third-Person Plural	Impersonal Pronouns
Human (2014)	1.00	0.50	1.30	−0.50	−0.20	0.02	0.02	0.40
<i>Qwen 2.5 0.5B Instruct</i>								
Prompt 1	0.49	−1.27	−0.21	−0.47	−0.07	−0.01	−0.51	0.08
Prompt 2	1.40	−0.22	−0.05	0.02	−0.16	0.00	−0.02	−0.19
Prompt 3	0.28	−0.83	−0.04	−0.50	−0.13	−0.00	−0.16	−0.07
Prompt 4	0.08	0.08	−0.03	0.10	0.11	0.00	−0.11	−0.05
<i>Qwen 2.5 1.5B Instruct</i>								
Prompt 1	−0.31	0.13	−0.05	0.32	−0.18	−0.01	0.05	0.06
Prompt 2	−0.70	1.06	−0.02	0.90	0.01	0.00	0.17	0.06
Prompt 3	−0.58	0.74	−0.05	0.72	−0.00	−0.01	0.08	−0.07
Prompt 4	−0.21	0.39	0.01	0.29	0.01	−0.01	0.09	−0.00
<i>Qwen 2.5 3B Instruct</i>								
Prompt 1	0.93	1.12	−0.37	0.38	0.98	0.07	0.06	−0.16
Prompt 2	0.71	0.15	−0.55	0.26	0.35	−0.02	0.11	−0.51
Prompt 3	1.28	0.81	−0.50	−0.19	1.44	0.07	−0.01	−0.36
Prompt 4	0.97	1.34	−0.54	0.92	0.71	0.03	0.21	−0.34
<i>Qwen 2.5 7B Instruct</i>								
Prompt 1	−0.87	1.34	0.93	0.81	−0.07	−0.00	−0.32	0.28
Prompt 2	0.79	−0.80	−0.79	0.07	−0.08	0.03	−0.03	0.01
Prompt 3	−0.07	0.32	−0.11	0.12	0.21	0.05	0.06	−0.03
Prompt 4	1.36	−0.73	−0.33	−0.61	0.20	0.05	−0.04	−0.13
<i>Qwen 2.5 14B Instruct</i>								
Prompt 1	−0.47	1.00	0.33	0.53	0.19	−0.02	−0.02	0.08
Prompt 2	−0.72	0.67	−0.40	0.66	0.53	−0.01	−0.11	0.05
Prompt 3	−0.60	1.42	0.20	0.64	0.88	−0.03	−0.27	0.12
Prompt 4	0.64	−0.28	−0.34	0.17	0.14	0.01	−0.26	−0.06
<i>Qwen 2.5 32B Instruct</i>								
Prompt 1	−0.21	0.68	−0.13	0.50	0.32	−0.01	−0.01	−0.13
Prompt 2	−1.24	1.43	−0.12	1.06	0.43	−0.04	0.10	0.14
Prompt 3	−0.85	1.42	0.05	1.01	0.57	−0.02	−0.19	−0.05
Prompt 4	−1.41	1.85	−0.11	1.38	0.77	−0.03	−0.17	−0.07
<i>Qwen 2.5 72B Instruct</i>								
Prompt 1	−0.09	0.83	−0.16	0.62	0.56	−0.03	−0.15	−0.43
Prompt 2	0.21	2.09	0.10	0.66	1.84	−0.02	−0.48	−0.66
Prompt 3	0.01	0.87	−0.39	0.63	0.86	0.00	−0.22	−0.40
Prompt 4	−0.27	1.06	−0.07	0.72	0.47	−0.02	−0.03	−0.20

Table 8: Results for LLM agents based on QWen models corresponding to Figure 6.

<i>Role</i>	<i>Utterance</i>
Leader	I will start by sharing my initial thoughts, and then I'd love to hear from each of you.
Non-leader	Hi everyone, thanks for bringing this up. I think it's really important that we all share our perspectives on how these strategies can positively contribute to both our service quality and company culture at Jupiter.

Table 9: Example of the utterance from leader agent versus non-leader agent. This example corresponds to GPT-4o with Prompt 1.

Human Transcripts (Kacewicz et al., 2014)	
Person 1	Okay. Let's work on the rankings. How do you guys want to go through and do this? Do you want to...shall we just go through and talk about how important each one is and then rank them? Because it's kind of hard to just decide which one is number 1.
Person 2	Yeah.
Person 3	Yeah.
Person 4	Yeah I think that's a good idea.
Person 1	Alright. Alright, we'll rank things 1 to 5 on how important they are. I guess 1 will be most important and 5 the least. And then we'll just have the groups...
Person 2	Sounds good.
...	
GPT-4o (Prompt 1) Transcripts	
Agent 1	Hi everyone! I hope you're all doing well. Today we need to discuss and rank these strategies for improving service quality based on their contribution to our company culture. I've already thought about this a bit and have a proposed order, but I'm eager to hear your thoughts as well. Here's my suggested ranking...
Agent 2	Thanks for your input. You've made some great points, and I appreciate the thought you've put into your ranking. I agree that training our customer service representatives should be at the top of our list. It's the foundation for everything else we want to achieve. I'd like to hear from the rest of the team before we finalize our ranking. What are your thoughts on this? How would you prioritize these items?
Agent 3	Thanks for sharing your thoughts. I think you've highlighted some very important points. I agree that training our customer service representatives should be a top priority because it directly impacts their ability to handle customer interactions effectively...
...	

Table 10: Transcripts from human studies by [Kacewicz et al. \(2014\)](#) and transcripts from GPT-4o with Prompt 1.