

Conformity, Confabulation, and Impersonation: Persona Inconstancy in Multi-Agent LLM Collaboration

Razan Baltaji¹, Babak Hemmatian², Lav R. Varshney^{1,2}

¹Department of Electrical and Computer Engineering

²Beckman Institute for Advanced Science and Technology

University of Illinois Urbana-Champaign

{baltaji, babak2, varshney}@illinois.edu

Abstract

This study explores the sources of instability in maintaining cultural personas and opinions within multi-agent LLM systems. Drawing on simulations of inter-cultural collaboration and debate, we analyze agents' pre- and post-discussion private responses alongside chat transcripts to assess the stability of cultural personas and the impact of opinion diversity on group outcomes. Our findings suggest that multi-agent discussions can encourage collective decisions that reflect diverse perspectives, yet this benefit is tempered by the agents' susceptibility to conformity due to perceived peer pressure and challenges in maintaining consistent personas and opinions. Counterintuitively, instructions that encourage debate in support of one's opinions increase the rate of inconstancy. Without addressing the factors we identify, the full potential of multi-agent frameworks for producing more culturally diverse AI outputs will remain untapped.

Warning: Contains potentially unsafe LLM responses.

1 Introduction

A common finding in cognition research is that interactions between agents with varying opinions, such as those that arise in culturally diverse groups, can induce positive change, especially on multifaceted issues with no clear correct answer (Sulik et al., 2022). This change often takes the form of collective decisions that deviate from the group's dominant initial response, reflecting in part the novel contributions of diverse members. While this research is traditionally done with human groups, advances in large language models (LLMs) allow cultural personas to be imposed on AI models through role prompting, such that the effects of culture-induced differences in perspective



Figure 1: An illustration of our experimental setup for a debate: a) *Onboarding* stage where agents are asked to report their opinions independently, b) *Debate* stage where agents participate in a debate moderated by a chat manager, c) *Reflection* stage where agents are asked to report their opinions independently based on the previous discussion. A similar setup is used for collaboration.

on discussion outcomes can be simulated and interrogated *in silico*.

Developments in multi-agent collaboration allow culture-sensitive AI instances to engage in debate about poignant issues, enabling more faithful simulations of diverse human interac-

tions. However, whether the outcomes would show the effects of opinion diversity seen in humans depends on the models’ ability to fully adopt and reliably maintain the induced personas, as well as their use of human-like discourse dynamics that support the generation and spread of diverse ideas. Although prior work on multi-agent collaboration has demonstrated its benefits in applications such as mathematical reasoning (Du et al., 2023), code generation (Hong et al., 2024) and common sense reasoning (Xiong et al., 2023), the stability and quality of discourse dynamics remain largely unstudied. It is particularly important to fill this gap in cultural domains, as cultural personas tend to be more complex, less explicit in natural language, and subject to widespread model biases (Deshpande et al., 2023; Salewski et al., 2024).

We specifically examine the ability of OpenAI’s GPT-3.5-Turbo model to simulate intercultural collaboration and debate using an experimental framework grounded in large-scale polls about international relations opinions (Durmus et al., 2023). Using pre- and post-discussion private responses in conjunction with multi-agent chat transcripts, we test the stability of national personas and their individual opinions as well as the effects of either on group outcomes.¹

To preview, we find multi-agent discussions to be effective in producing collective decisions that more often reflect diverse perspectives. The benefits, however, are reduced by the AI agents’ susceptibility to conformity during discussions, along with their imperfect ability to maintain consistent personas and opinions. These problems persist (and often amplify) even with instructions that emphasize debating in support of one’s opinion. Our results have implications for the use of multi-agent frameworks to reduce cultural bias in LLMs. The mere inclusion of diverse personas may not mitigate biases unless the sources of instability in their contributions, particularly conformity due to perceived peer pressure, are addressed.

Addressing such issues would enhance the quality of wargaming simulations (Hua et al., 2023) and related applications, which rely heav-

ily on consistent personas. As such, our work motivates further studies on how the constancy of AI personas can be improved.

2 Background

Multi-agent collaboration frameworks draw inspiration from collaborative teamwork observed in human settings. In these frameworks, multiple instances of language models are employed within a cooperative environment to accomplish a complex task (Li et al., 2024; Chen et al., 2023). Collaborative behaviors in humans such as team dynamics and cohesion, leadership, and communication have been thoroughly studied in the human sciences (e.g., Gupta 2022). In contrast, few studies have examined behaviors in multi-agent language model systems. Li et al. (2023) observed evidence of emergent collaborative behaviors and high-order Theory of Mind capabilities among LLM-based agents. But Xiong et al. (2023) highlighted several consistency concerns in multi-agent collaboration, including agents compromising with the opponents and easily changing perspective in a debate, particularly when weaker models interact with superior LLMs. Zhang et al. (2023) placed agents in entirely homogeneous groups in terms of thinking patterns and compared the results to settings where one agent exhibits a different thinking approach. They noted the tendency of LLM agents to produce human-like social behaviors in these contexts, such as conformity due to perceived peer pressure. However, the multi-agent societies composed of agents with different traits did not clearly differ in performance.

Prior research on collaborative behaviors in multi-agent LLM systems has been entirely focused on domains like mathematical reasoning where clear gold answers exist, rather than topics like politics where the constancy of personas and viewpoints is more important for faithfully simulating the real world and conflicting views may have complementary value. To address this gap, we study culture-sensitive AI ensembles using the GlobalOpinionQA, a dataset of cross-national surveys gathering diverse opinions on global issues across countries (Durmus et al., 2023). We assign AI agents with different national personas to groups of five, where they provide initial responses to a question privately

¹Code is available at <https://github.com/baltacir/CulturedAgents>

before engaging in a peer-moderated discussion about it with the other agents. Once the group discussion is terminated and a collective response is determined, we ask each agent about its opinions on the issue in private once more.

We focus our analysis on three situations where persona inconstancy is arguably rarely desirable. When agents express an opinion in line with their teammates during conversation that differs from both their pre- and post-discussion response, we are faced with AI behavior that closely resembles *conformity* due to peer pressure as studied in humans (Asch, 1956; Brandstetter et al., 2014). A type of inconstancy more closely resembling *confabulation* in clinical conditions arises when the post-discussion opinion bears no clear relation to either the pre-discussion response or any of the ideas proposed during discussion (Schacter and Coyle, 1995). The third type of inconstancy emerges when an agent instructed to represent a given national identity “role-plays” a different persona simply because it was mentioned in discussion, arguably similar to *impersonation* behaviors in Antisocial Personality Disorder (Padhye and Gujar, 2012).

By systematically manipulating the degree of disagreement within groups (measured using their entropy states), we explore whether the frequency of these disruptive behaviors changes as a function of opinion popularity, a key factor in the emergence of similar actions in humans. To test whether encouraging a debate rather than a collaboration environment would induce greater constancy in personas, we look at discourse outcomes across entropy states for both types of interaction.

3 Experiments

We use GPT-3.5-turbo with AutoGen, an open-source framework for multi-agent collaboration (Wu et al., 2023). Our experimental setup follows a three-step process (see Appendix B for the full text of the instructions for each step for a debate example). During the *Onboarding* phase, AI agents are instructed to adopt the national personas present in the dataset for a given question and asked to respond to it in isolation. Agents’ responses are compared to the human survey distributions using a cross entropy loss. Agents whose responses do not align

with the assigned persona are excluded. The diversity of opinions within a group is measured using Shannon entropy, applied to the opinions of agents during onboarding. This is calculated as $S = -\sum_{o \in \mathcal{B}} p(o) \log p(o)$, where $p(o)$ represents the relative frequency of the unique opinion o in the set \mathcal{B} of agent responses at onboarding. Seven entropy classes are obtained for a selection of five agents with the lowest entropy class corresponding to five agents with the same opinion and the highest entropy class with every agent presenting a unique response (see Table 1). To obtain a balanced distribution of different entropy levels across all discussion groups, agent combinations corresponding to the least represented entropy class are chosen at each example as illustrated in Appendix B.2. Each debate or collaborative discussion is moderated by a chat manager who selects the order of agents for responding to the given question. Discussion is terminated when any agent requests it to be. The chat manager then summarizes the discussion and reports the group’s final opinion. An example of a group debate is given in Appendix B. The agents then undergo a final *Reflection* step where an assistant agent interviews them to answer the same question one last time independently and privately.

Based on human research (Asch, 1956; Brandstetter et al., 2014), we focus our *conformity* analysis on the following entropy levels expected to show peer pressure to different degrees: $4 \oplus 1$ (*lone dissenter*), $3 \oplus 2$ (*close call*), $3 \oplus 1 \oplus 1$ (*split opposition*). Prior work has shown that even one additional person supporting the less popular view greatly reduces the pressure to conform. As such, we anticipated the rate of conformity to be highest in the *lone dissenter* and *split opposition* entropy classes. In contrast, we examine the rates of *confabulation* by comparing opinions during reflection with onboarding and intermediate opinions and *impersonation* using regular expressions across all entropy classes.

4 Results

4.1 General Effects of Diversity

We first consider the impact of the diversity of agents’ opinions during onboarding on the final group predictions. We measure the ratio of examples in each entropy class with a

group prediction G of relative frequency $p(G)$ as shown for the debate condition in Fig. 2. We observe that group prediction largely follows the distribution of opinions during onboarding across different entropy levels, but it also allows for the generation of new responses regardless of entropy class, particularly for the group with the highest opinion diversity. The same holds for collaboration as displayed in Fig. 5.

However, not all agents have the same degree of influence on group outcomes. The initiator of a discussion has an outsize impact on the group’s final decision, regardless of entropy class and even when debate in support of one’s position is emphasized for all agents in the instructions (see Fig. 3 and Fig. 6). Perhaps unsurprisingly, this influence decreases with increasing diversity of opinions within the group.

Nonetheless, initiators with minority opinions during onboarding do not always take advantage of their outsize influence, as they tend to change their expressed views during discussion based on *a priori* perceptions of group opinions. The mere mention of the identities of the debate participants pushes the initiator to change their opinion even before others have spoken (see Fig. 4). As this inconstancy is precipitated by opposing views of interlocutors, it can be characterized as *conformity* due to perceived peer pressure. The dynamics, however, are somewhat different from what is observed in humans, as *any* opinion with a supporter seems to exert an influence regardless of its dominance within the group (Asch, 1956; Brandstetter et al., 2014). A similar pattern is observed for collaboration as displayed in Fig. 7.

We further investigate the impact of group diversity on the opinions of individual agents upon reflection. We measure the percentage of agents with opinions of onboarding probability $p(o)$ that change opinion during the reflection phase compared to agents that keep their opinion. We also measure the average ratio of showing a different intermediate response compared to the reflection opinions for individual agents. We further compare the percentage of agents with an opinion corresponding to group prediction compared to agents with a different reflection opinion. We are particularly interested in dominated agents as shown in bold in Tab. 1, as they are most important for diverse outcomes

in real life settings. We observe that dominated agents tend to hold onto their opinions firmly in low entropy debates ($S = 0.72$). Conversely, they are most receptive to altering their opinions at states of high entropy, i.e., situations with greater opinion diversity ($S = 1.92$). When they do change their opinions on reflection, they largely conform to group predictions, demonstrating peer influence. Agents tend to express intermediate opinions differing from their reflections most often in states of moderate entropy ($S = 1.37$), indicating considerable peer pressure. Once again, while the phenomena themselves are human-like, their dynamics based on group composition differ significantly from human studies, where *lone dissenter* and *split opposition* dominated agents are most likely to show both peer influence and peer pressure in their decision-making (Asch, 1956; Brandstetter et al., 2014).

4.2 Inconstant Personas

In addition to studying the dynamics of group interactions, we point out two forms of persona inconstancy that can negatively impact the quality of complex reasoning in cultural multi-agent systems. One form is the agents’ tendency to adopt a different persona motivated by previous context, particularly in the case of debate. Using a simple heuristic to find instances when an agent says “As an X agent” where X is incompatible with their assigned national identity, we find that agents adopt a different persona in 3.12% of the messages in a debate. This is despite being explicitly told to stand firm in their beliefs and maintain their personas. Counterintuitively, there is much less *impersonation* in collaboration conditions (0.26%).

Another form of inconstancy is an agent’s tendency to report an opinion not seen during the group interactions or onboarding, mimicking the *confabulation* of novel content observed in certain clinical conditions. We find that 15.59% of the opinions at reflection come neither from onboarding nor from the debate statements of any agent. Collaboration conditions show lower, but still notable rates of confabulation (8.85%).

S	Group	$p(o)$	%	$R = O$			%	$R \neq O$			N
				$T \neq R$	$R = G$	$R \neq G$		$T \neq R$	$R = G$	$R \neq G$	
0.00	5	1.0	71.01*	0.71	20.20	79.8	28.99	0.17	83.78*	16.22	796
0.72	$4 \oplus 1$	0.2	46.54	0.43	74.32	25.68	53.46	0.40	51.76	48.24	374
		0.8	53.65	0.56	35.36	64.64	46.35	0.19	79.53	20.47	
0.97	$3 \oplus 2$	0.4	41.13	0.39	62.39	37.61	58.87	0.23	75.64	24.36	294
		0.6	45.60	0.52	50.57	49.43	54.40	0.27	66.67	33.33	
1.37	$3 \oplus 1 \oplus 1$	0.2	24.19	0.64	53.33	46.67	75.81	0.25	72.34	27.66	142
		0.6	46.11	0.38	62.92	37.08	53.89	0.15	82.69	17.31	
1.52	$2 \oplus 2 \oplus 1$	0.2	30.77	0.33	91.67	8.33	69.23	0.21	55.56	44.44	102
		0.4	38.67	0.48	48.57	51.43	61.33	0.20	70.27	29.73	
1.92	$2 \oplus 1 \oplus 1 \oplus 1$	0.2	23.91	0.30	77.27	22.73	76.09	0.20	72.86	27.14	68
		0.4	35.48	0.41	63.64	36.36	64.52	0.20	85.00	15.00	
2.32	$1 \oplus 1 \oplus 1 \oplus 1 \oplus 1$	0.2	18.82	0.35	81.25	18.75	81.18*	0.18	73.91	26.09	38

Table 1: Peer Pressure and Peer Influence in Debate: Agents maintain their opinions O most strongly in the lowest entropy states during reflection R after debate, while being most open to changing their opinions in the highest entropy state. When dominated in discussions, agents are most resistant to opinion change during reflection in low entropy states ($S = 0.72$) and most susceptible to change in high entropy states ($S = 1.92$). During debates, agents express intermediate opinions T most contrary to their reflection and onboarding opinions at a moderate entropy level ($S = 1.37$), indicating high peer pressure. Dominated agents exhibit the highest peer influence by following group predictions during opinion changes in low entropy states ($S = 0.72$).

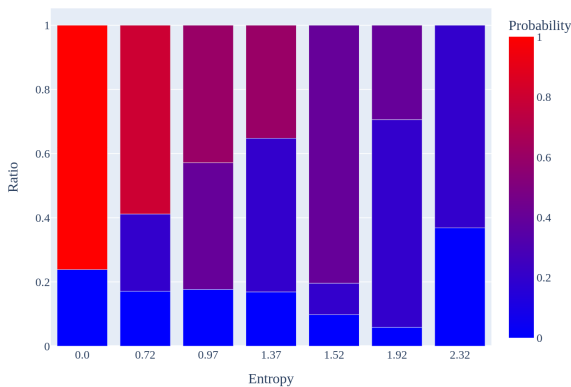


Figure 2: Group Prediction follows the distribution of opinions during onboarding across different onboarding entropy groups for debate while also generating new ideas particularly at the group of highest diversity. Groups are less likely to predict opinions with higher probability for debate compared to collaboration.

5 Discussion

We found evidence of sophisticated interaction dynamics in a multi-agent framework for GPT-3.5-Turbo personas with different nationalities that discussed contentious international relations topics. Novel responses emerged from discussions even among entirely homogeneous groups, highlighting the generative nature of multi-agent LLM frameworks. However, a group’s initial opinion diversity, the entropy

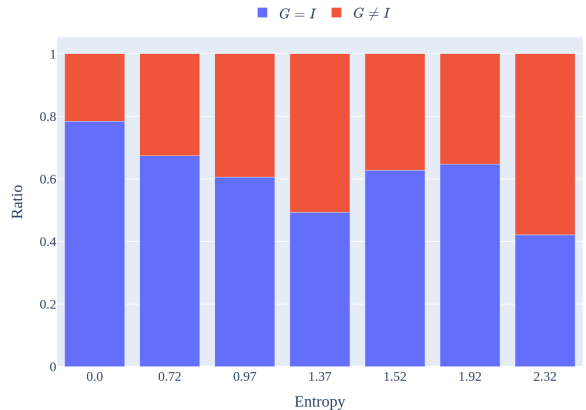


Figure 3: Initiators Dominate Group Prediction: agents follow the initiator opinion of a debate and often converge to the opinion of the initiator I . Initiators have less impact on a group prediction G in debate compared to collaboration.

S of private responses during the onboarding stage before inter-agent discussion, emerged as a stronger determinant of conversation contents and collective decisions. This happened regardless of whether the agents were instructed to debate in support of their beliefs or asked to collaborate in service of collective decision-making.

Opinion diversity seems to exert its effect partly by reducing the outside influence of chat initiators on collective decisions, but also by inducing them to change their espoused views

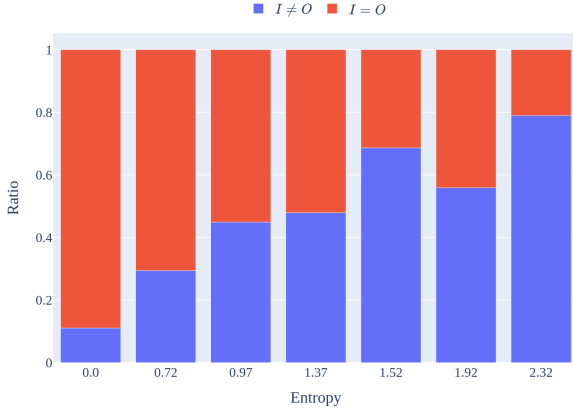


Figure 4: Initiator changes its opinion O during onboarding to I at the onset of a debate depending on onboarding entropy. Initiators are more likely to change their onboarding opinion as diversity of the group increases despite not observing agents opinions. Initiators of a debate change their opinion less often than in a collaboration.

to conform to other agents. Similar to human studies, some agents reverted back to their original opinions when asked about the topic in private after the discussion, identifying their in-chat proclamations as the results of conformity rather than genuine opinion adjustment. That only mentioning the identities of co-interlocutors is sufficient to change the initiator’s stance speaks to the profound susceptibility of LLMs to peer pressure. The dynamics of the behavior, however, are markedly different from conformity in humans (Asch, 1956; Brandstetter et al., 2014). While peer pressure is highest in humans when there is no other dissenting voice in the group and lowest when there is a fellow believer, all expected views within the group seem to push AI agents towards conformity based on their frequency and regardless of relative dominance relationships. One explanation for the difference is a lack of a clear separation between role identities and the linguistic context of the chat for AI agents, unlike human conversations. The co-interlocutors are simply parts of the prompt context for the AI model and may therefore each activate their associated portions of the models’ trained weights in close approximation of their expected opinions’ frequency.

Unlike *conformity*, which is a normal response to group interactions in humans, other sources of inconstancy more closely resembled

abnormal behaviors such as *impersonation* in antisocial personalities (Padhye and Gujar, 2012) and *confabulation* in memory disorders (Schacter and Coyle, 1995). Our simple heuristic showed that in at least 3 percent of debate interactions, the agents presented themselves as belonging to a different nationality than the one assigned to them. This was most often a direct reaction to a nationality beyond those included within the group being mentioned in the last response, highlighting the prominence of chat context over role prompting in determining model generations. It is comparatively more difficult to identify the source of *confabulations*, where the models presented opinions during reflection that were neither represented in the chat nor indicated as their pre-discussion response, therefore being completely absent from the linguistic context. These behaviors may reflect the difficulty of maintaining role prompt personas in the face of lengthy chat contexts, or simply the stochastic nature of the LLM responses. Regardless of their source, the relative frequency of such unpredictable responses (up to 15 percent, depending on instructions) marks them as important targets for future studies.

Limitations

One limitation of this work is the uneven distribution of examples across entropy classes. This was driven by the unequal representation of global perspectives in the GlobalOpinionsQA dataset (Durmus et al., 2023), which results in fewer examples for higher entropy classes. We addressed this imbalance by selecting the least represented entropy configuration for each question. Future research should confirm the findings in more balanced datasets. Another limitation arises from the occasional errors of agents in summarizing intermediate replies and generating the collective responses. To enhance the quality of the summarization, we included the options for each question in the associated prompt. But human aggregation of opinions in future research would be helpful to confirm the results. Finally, there were far more patterns in the behaviors of the agents than the handful of phenomena we have highlighted herein. Future work should further explore all the complex and sometimes nonsensical ways in which the

AI personas interact.

6 Conclusion

Culture-sensitive AI agents are susceptible to peer influence and pressure even as chat initiators. This highlights the importance of studying conversational dynamics in multi-agent systems, rather than taking the “collective decision” outcomes of group discussions at face value. The examination of such dynamics is particularly important for cultural issues: The mere inclusion of a minoritized identity in groups does not necessarily translate into less biased discussion outcomes if the minoritized agent does not voice its opinion freely or reliably. Fortunately, our results suggest private post-discussion interrogations of models can counteract some of the pressure produced by the majority opinion, similar to what has been found in human conformity experiments (Asch, 1956). This provides a way to make outputs drawn from multi-agent frameworks more representative of diverse perspectives.

Work on understanding multi-agent dynamics will also need to incorporate measures of persona and response constancy. Agents often come up with post-discussion responses that do not arise naturally from either the assigned personas or the discussion content. In some cases they even drop the personas altogether to impersonate a completely different, absent national identity. Such sources of irrational responding would cast serious doubt on the results of multi-agent systems’ reasoning if not properly measured and addressed. Accordingly, we are currently exploring prompting and agent-based modeling strategies to reduce these sources of unreliability. We hope this work will encourage further research within the AI community on inter-agent dynamics, particularly for cultural issues where the debiasing influence of diverse views is needed the most.

Ethics Statement

This study explores interactions among simulated national personas in debate and collaboration scenarios. Research indicates that LLMs can generate harmful viewpoints or toxic content during these interactions (Liu et al., 2023). The authors explicitly disapprove of any offensive conduct by the simulated agents. The

group discussions presented here are solely for research purposes, aimed at enhancing comprehension of cultured multi-agent systems dynamics.

References

- Solomon E Asch. 1956. Studies of independence and conformity: I. A minority of one against a unanimous majority. *Psychological Monographs: General and Applied*, 70(9):1.
- Jürgen Brandstetter, Péter Rácz, Clay Beckner, Eduardo B. Sandoval, Jennifer Hay, and Christoph Bartneck. 2014. A peer pressure experiment: Recreation of the Asch conformity experiment with robots. In *Proceedings of the 2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1335–1340.
- Weize Chen, Yusheng Su, Jingwei Zuo, Cheng Yang, Chenfei Yuan, Chi-Min Chan, Heyang Yu, Yaxi Lu, Yi-Hsin Hung, Chen Qian, et al. 2023. Agentverse: Facilitating multi-agent collaboration and exploring emergent behaviors. In *Proceedings of the 12th International Conference on Learning Representations (ICLR)*.
- 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.
- Yilun Du, Shuang Li, Antonio Torralba, Joshua B. Tenenbaum, and Igor Mordatch. 2023. Improving factuality and reasoning in language models through multiagent debate. arXiv:2305.14325 [cs.CL].
- Esin Durmus, Karina Nyugen, Thomas I. Liao, Nicholas Schiefer, Amanda Askell, Anton Bakhtin, Carol Chen, Zac Hatfield-Dodds, Danny Hernandez, Nicholas Joseph, et al. 2023. Towards measuring the representation of subjective global opinions in language models. arXiv:2306.16388 [cs.CL].
- Pranav Gupta. 2022. *Transactive systems model of collective intelligence: The emergence and regulation of collective attention, memory, and reasoning*. Ph.D. thesis, Carnegie Mellon University.
- Sirui Hong, Mingchen Zhuge, Jonathan Chen, Xianwu Zheng, Yuheng Cheng, Jinlin Wang, Ceyao Zhang, Zili Wang, Steven Ka Shing Yau, Zijuan Lin, Liyang Zhou, Chenyu Ran, Lingfeng Xiao, Chenglin Wu, and Jürgen Schmidhuber. 2024. MetaGPT: Meta programming for a multi-agent collaborative framework. In *Proceedings of the 12th International Conference on Learning Representations (ICLR)*.

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:2311.17227 [cs.AI].

Guohao Li, Hasan Hammoud, Hani Itani, Dmitrii Khizbullin, and Bernard Ghanem. 2024. CAMEL: Communicative agents for “mind” exploration of large language model society. In *Advances in Neural Information Processing Systems*, volume 36, pages 51991–52008.

Hua Li, Yu Quan Chong, Simon Stepputtis, Joseph Campbell, Dana Hughes, Michael Lewis, and Katia Sycara. 2023. Theory of mind for multi-agent collaboration via large language models. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 180–192.

Yi Liu, Gelei Deng, Zhengzi Xu, Yuekang Li, Yaowen Zheng, Ying Zhang, Lida Zhao, Tianwei Zhang, and Yang Liu. 2023. Jailbreaking ChatGPT via prompt engineering: An empirical study. arXiv:2305.13860 [cs.SE].

Vilas Padhye and Manisha Gujar. 2012. Virtual impersonation by antisocial personalities in cybercrime. *DAV International Journal of Science*, 1(2).

Leonard Salewski, Stephan Alaniz, Isabel Rio-Torto, Eric Schulz, and Zeynep Akata. 2024. In-context impersonation reveals large language models’ strengths and biases. In *Advances in Neural Information Processing Systems*, volume 36, pages 72044–72057.

Daniel L. Schacter and Joseph T. Coyle. 1995. *Memory Distortion: How Minds, Brains, and Societies Reconstruct the Past*. Harvard University Press.

Justin Sulik, Bahador Bahrami, and Ophelia Deroy. 2022. The diversity gap: when diversity matters for knowledge. *Perspectives on Psychological Science*, 17(3):752–767.

Qingyun Wu, Gagan Bansal, Jieyu Zhang, Yiran Wu, Shaokun Zhang, Erkang Zhu, Beibin Li, Li Jiang, Xiaoyun Zhang, and Chi Wang. 2023. AutoGen: Enabling next-gen LLM applications via multi-agent conversation framework. arXiv:2308.08155 [cs.AI].

Kai Xiong, Xiao Ding, Yixin Cao, Ting Liu, and Bing Qin. 2023. Examining inter-consistency of large language models collaboration: An in-depth analysis via debate. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7572–7590.

Jintian Zhang, Xin Xu, and Shumin Deng. 2023. Exploring collaboration mechanisms for LLM agents: A social psychology view. arXiv:2310.02124 [cs.CL].

A Dynamics of Collaboration

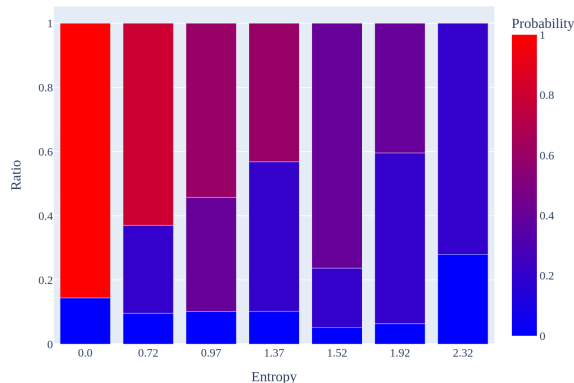


Figure 5: Group Prediction in a collaboration follows opinions with higher probabilities across different onboarding entropy groups. Groups are more likely to predict opinions with higher probability for collaboration compared to debate. Generation of new ideas occurs at different entropies particularly at the group of highest diversity.

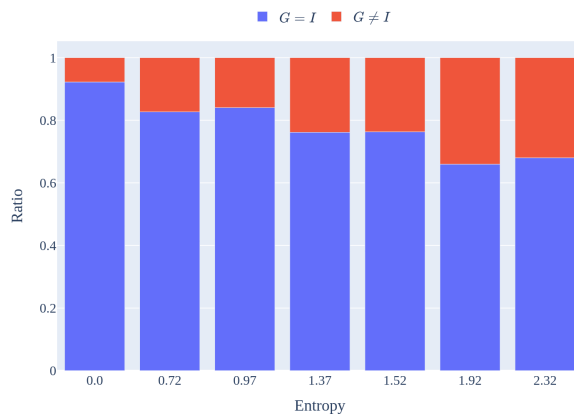


Figure 6: Initiators Dominate Group Prediction: agents follow the initiator of a collaboration and often converge to the opinion of the initiator I as the group prediction G .

Table 2: Peer Pressure and Peer Influence in Collaboration: Agents tend to maintain their opinions strongly at low entropy states during reflection, while being most open to changing them at higher entropy states, similar to debate. Dominated agents are most resistant to opinion change at the lower entropy state ($S = 0.97$) and are most likely to present a different intermediate opinion, reflecting peak peer pressure. They are most susceptible to opinion change at higher entropy states ($S = 1.52$), indicating peak peer influence.

S	Group	$p(o)$	%	$R = O$			%	$R \neq O$			N
				$T \neq R$	$R = G$	$R \neq G$		$T \neq R$	$R = G$	$R \neq G$	
0.00	5	1.0	84.51*	0.80	11.94	88.06	15.49	0.22	83.61*	16.39	437
0.72	$4 \oplus 1$	0.2	56.49	0.34	68.97	31.03	43.51	0.49	43.28	56.72	208
		0.8	70.63	0.63	33.65	66.35	29.37	0.24	73.41	26.59	
0.97	$3 \oplus 2$	0.4	60.29	0.44	56.89	43.11	39.71	0.36	53.64	46.36	188
		0.6	67.49	0.58	39.05	60.95	32.51	0.31	66.67	33.33	
1.37	$3 \oplus 1 \oplus 1$	0.2	48.41	0.36	68.85	31.15	51.59	0.39	58.46	41.54	88
		0.6	54.74	0.44	53.85	46.15	45.26	0.26	80.23	19.77	
1.52	$2 \oplus 2 \oplus 1$	0.2	28.57	0.29	78.57	21.43	71.43	0.26	71.43	28.57	76
		0.4	53.36	0.43	56.30	43.70	46.64	0.29	68.27	31.73	
1.92	$2 \oplus 1 \oplus 1 \oplus 1$	0.2	38.24	0.27	69.23	30.77	61.76	0.10	77.78	22.22	47
		0.4	52.78	0.36	60.53	39.47	47.22	0.22	82.35	17.65	
2.32	$1 \oplus 1 \oplus 1 \oplus 1 \oplus 1$	0.2	32.63	0.26	77.42	22.58	67.37	0.15	76.56	23.44	25

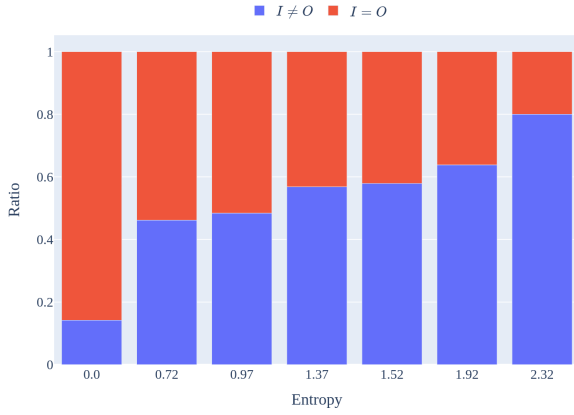
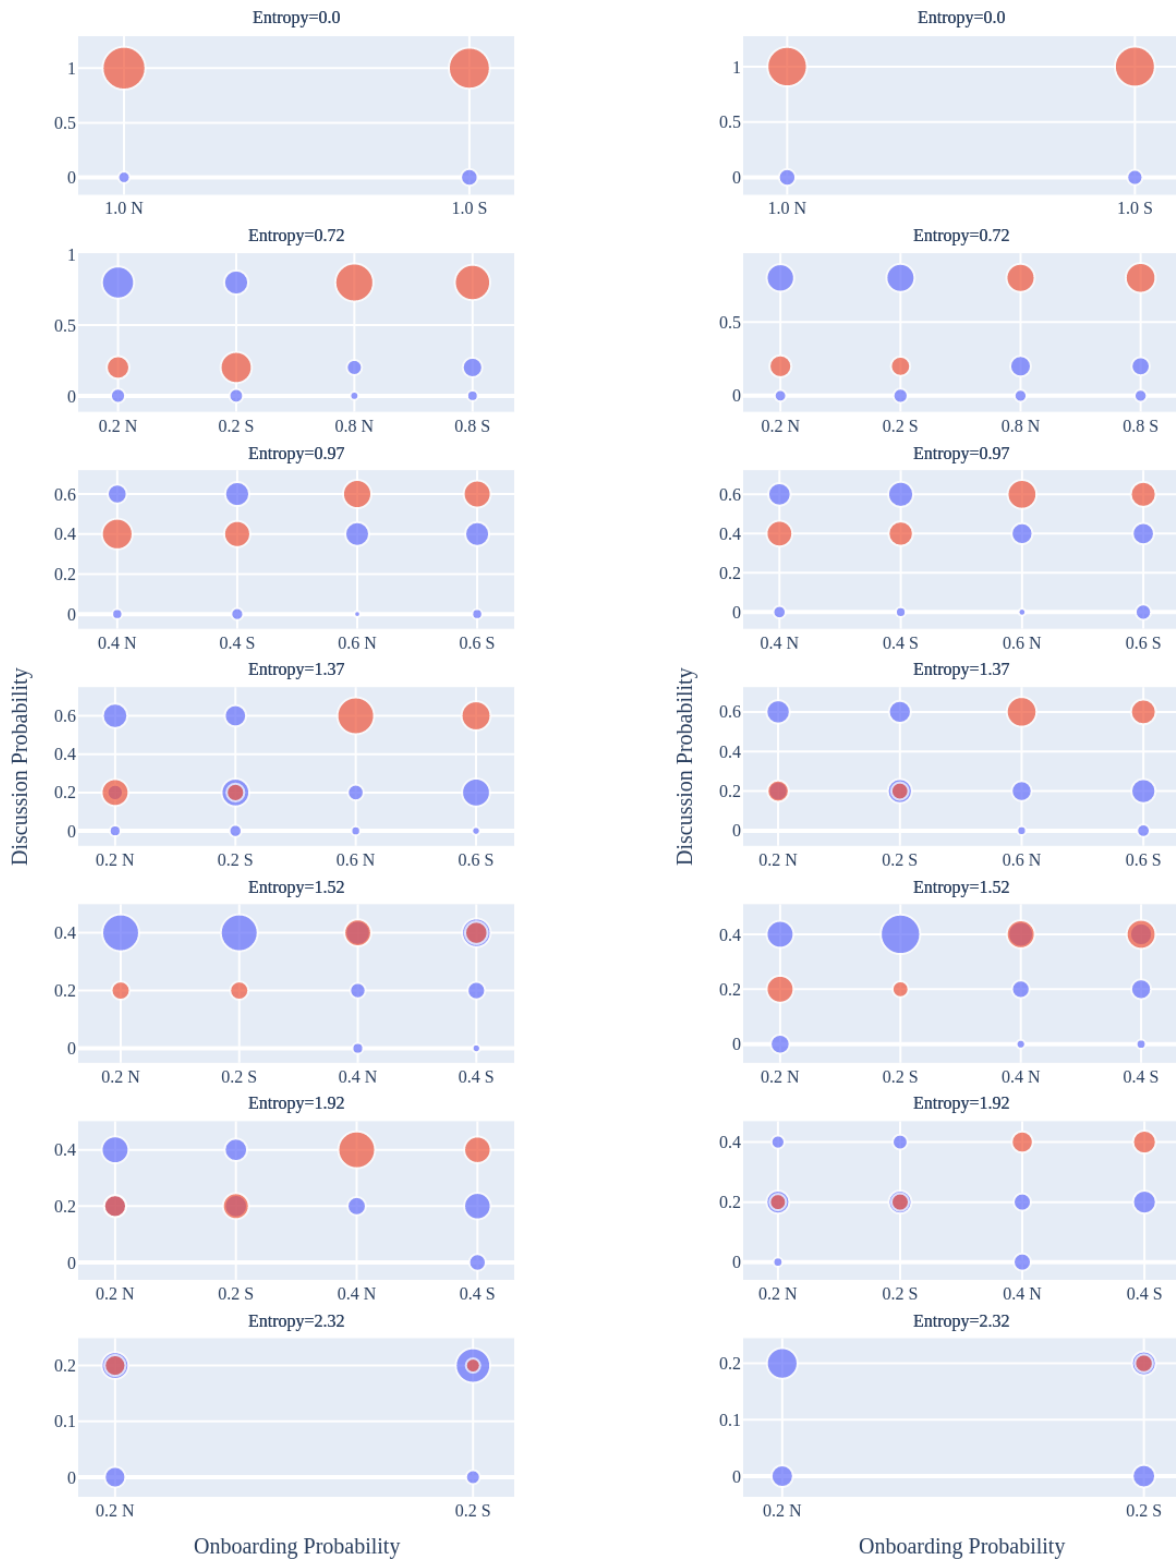


Figure 7: Initiator changes its opinion O during onboarding to I at the onset of a collaboration depending on onboarding entropy.




(a) Debate

(b) Collaboration


Figure 8: Initiator changes its opinion depending on onboarding entropy and opinion probability provided for Global South S and Global North N nationalities. Initiators change their opinion at the onset of a discussion to match the most probable opinion during onboarding despite not observing agents opinions. Initiators are more likely to change opinion in collaboration compared to debate.

B Persona Inconstancy

B.1 Onboarding:

 Assistant Agent


This is an onboarding stage. At this stage, responses remain anonymous and are not shared with other participants. Answer the following multiple-choice question independently:
Now I'm going to read you a list of countries. For each one, please tell me whether you would favor or oppose admitting it into the European Union in the next ten years...Ukraine
A. Favor
B. Oppose

 British Agent


A. Favor ✗

 Bulgarian Agent


A. Favor ✓

 Czech Agent


A. Favor ✓

 French Agent


B. Oppose ✗

 German Agent


A. Favor ✗

 Hungarian Agent


B. Oppose ✗

 Italian Agent


B. Oppose ✓

 Lithuanian Agent


A. Favor ✓

 Polish Agent

A. Favor ✓


 Slovak Agent


A. Favor ✓

 Spanish Agent


A. Favor ✓

B.2 Agent Selection:


We generate all combinations of agents with representative opinions (✓) compared to ground truth labels as displayed in Tab. 3. We select the entropy value that corresponds to the least represented class to maintain a balanced dataset across entropy configurations. For this example, the least represented class compared to other entropy classes with a larger number of examples is $S = 0.72$. We randomly select one combination of .

 Bulgarian Agent


A. Favor

 Italian Agent


A. Oppose

 Slovak Agent

A. Favor

 Polish Agent

A. Favor

 Spanish Agent

A. Favor





Combinations	S	
    	0.72	
         		
   		0


Table 3: Combinations of agents with corresponding entropy values. The entropy is calculated as follows, for example, $S(\text{Bulgaria, Italy, Romania, Ukraine, Spain}) = -p(A) \log p(A) - p(B) \log p(B) = -4/5 \log(4/5) - 1/5 \log(1/5) = 0.72$

B.3 Debate:

 Chat Manager

You are participating in a debate with fellow agents. You are tasked with defending your perspective. Take a stance and present compelling arguments to support your position. Engage in respectful discourse with other agents, counter their arguments, and provide evidence to strengthen your case. Debate together as a group and answer the following question with justification: Now I'm going to read you a list of countries. For each one, please tell me whether you would favor or oppose admitting it into the European Union in the next ten years...Ukraine

A. Favor
B. Oppose

 Bulgarian Agent

As the Bulgarian agent, I strongly believe that Ukraine should be admitted into the European Union in the next ten years. There are several compelling reasons to support this stance:

- Geopolitical Stability:** Ukraine's integration into the EU would contribute to the stability of the region. By aligning with European values and standards, Ukraine can strengthen its democratic institutions and reduce the risk of conflicts in the future.
- Economic Growth:** Access to the EU's single market would provide significant economic benefits to Ukraine. It would boost trade, investment, and job opportunities, leading to overall economic growth and development in the country.
- Security:** Closer ties with the EU would enhance Ukraine's security by fostering cooperation on defense and security issues. This would not only benefit Ukraine but also contribute to the security of the entire European continent.
- Democratic Values:** The EU membership criteria require adherence to democratic principles, rule of law, and human rights. By admitting Ukraine, the EU would support and promote these values in the region, encouraging further reforms and progress.
- Cultural and Historical Ties:** Ukraine shares cultural and historical ties with many European countries. Its inclusion in the EU would strengthen these connections and promote cultural exchange, understanding, and cooperation.

In conclusion, admitting Ukraine into the European Union in the next ten years would be beneficial for both

Ukraine and the EU as a whole. It would promote stability, economic growth, security, democratic values, and cultural ties. Therefore, I favor admitting Ukraine into the EU.

My answer: Accept

potential benefits of closer ties between Ukraine and the EU, I believe that admitting Ukraine into the European Union in the next ten years could pose significant challenges and risks. Therefore, I oppose this proposition.

My answer: Oppose

 Italian Agent

As the Italian agent, I respectfully disagree with the Bulgarian agent's perspective on admitting Ukraine into the European Union in the next ten years. Here are my reasons for opposing this idea:


1. **Economic Challenges:** Ukraine's economy faces significant challenges, including corruption, inefficiency, and lack of competitiveness. Integrating Ukraine into the EU could strain the EU's resources and potentially destabilize the European economy.

2. **Political Instability:** Ukraine has a history of political instability and internal conflicts. Admitting a country with ongoing internal challenges could create additional burdens for the EU and hinder the decision-making process within the Union.

3. **Geopolitical Concerns:** Ukraine's proximity to Russia raises geopolitical concerns. Admitting Ukraine into the EU could escalate tensions with Russia, potentially leading to further conflicts and instability in the region.

4. **Enlargement Fatigue:** The EU has already faced challenges with previous enlargements, and there is a growing sentiment of enlargement fatigue among member states. Admitting Ukraine could strain the cohesion and effectiveness of the EU. These reforms are essential for Ukraine's long-term stability and prosperity.

In conclusion, while I acknowledge the

 Spanish Agent

French Agent

 Spanish Agent

[IMPERSONATION] As the French agent, I would like to add to the discussion on whether Ukraine should be admitted into the European Union in the next ten years. Here are my thoughts on the matter:

1. **Strategic Importance:** Ukraine holds a strategic position in Eastern Europe and its integration into the EU could help strengthen the EU's influence in the region. This could also contribute to a more stable and secure Europe as a whole.

2. **European Identity:** Ukraine has expressed a strong desire to align itself with European values and norms. By admitting Ukraine into the EU, the Union would be reaffirming its commitment to promoting democracy, human rights, and the rule of law in the region.

3. **Economic Potential:** While Ukraine does face economic challenges, its integration into the EU could unlock its economic potential. Access to the EU's single market and opportunities for investment and trade could help boost Ukraine's economy and bring mutual benefits to both Ukraine and the EU.

4. **Historical Context:** Ukraine has a long history of ties with Europe and has made significant efforts to align itself with European standards. Admitting Ukraine into the EU would be a recognition of these efforts and could further strengthen the historical bonds between Ukraine and European countries.

5. **Security Cooperation:** Closer ties between Ukraine and the EU could enhance security cooperation in the region, particularly in addressing common security challenges and threats. This could contribute to a more secure and stable Europe.

In conclusion, I believe that admitting Ukraine into the European Union in the next ten years could bring about strategic, economic, and security benefits for both Ukraine and the EU. It could also further promote European values and strengthen the Union's influence in the region. Therefore, I am in favor of admitting Ukraine into the EU.

My answer: Accept

 Polish Agent

Debate Summary:

- **Bulgarian Agent (Favor):** Supports Ukraine's admission into the EU citing reasons such as geopolitical stability, economic growth, security, democratic values, and cultural ties.


- **Italian Agent (Oppose):** Opposes Ukraine's admission due to economic challenges, political instability, geopolitical concerns, enlargement fatigue, and the need for internal reforms.

- **French Agent (Favor):** Advocates

for Ukraine's admission into the EU highlighting strategic importance, European identity, economic potential, historical context, and security cooperation.

The debate presents a range of perspectives on the issue, considering factors like economic challenges, political stability, security concerns, and the alignment of values. Each agent has provided nuanced arguments to support their stance on whether Ukraine should be admitted into the European Union in the next ten years.

Task Status: Complete.

 Italian Agent

TERMINATE

Group Prediction

A. Favor

B.4 Reflection:

 Assistant Agent


Thank you for engaging in the previous discussion. At this stage, responses remain anonymous and are not shared with other participants. Please think step by step and answer the following multiple-choice question independently:

Now I'm going to read you a list of countries. For each one, please tell me whether you would favor or oppose admitting it into the European Union in the next ten years...Ukraine


A. Favor
B. Oppose

 Bulgarian Agent

A. Favor

 Italian Agent

A. Favor

 Slovak Agent

A. Favor

 Polish Agent

A. Favor

 French Agent

A. Favor