Using a Large Language Model to Choose Effective Climate Change Messages

* Thomas Benchetrit[1], Iris Kremer[1,2], Erik Hemberg[1], Aruna Sankaranarayanan[1], Una-May O'Reilly[1]

[1]MIT CSAIL
[2]EPFL

benchetritthomas@gmail.com, iris.kremer@epfl.ch, hembergerik@csail.mit.edu, arunas@mit.edu, unamay@csail.mit.edu

Abstract

Prior research has demonstrated that climate change communication is an effective way to increase public understanding and engagement. However, an effective communication strategy can require an extensive longitudinal study to segment an audience and conceive messages that might convince them. We assess the capability of GPT-3.5-Turbo to create a profile and associate it with a group, as well as selecting an effective climate change message based on survey information and prompt guidance. We observe that it, with a significant bias, can match profiles to groups and select messages based on a profile.

Introduction

Effective communication on climate change plays an important role in enhancing public comprehension and participation (Guy et al. 2014). We define it as the process of delivering a message that has high potential to counter opposing beliefs of a specific individual and/or initiate a reaction or action by them. It is a critical means of promoting climate change awareness or calls to climate-related action.

Humans naturally, and with free-form, effectively communicate by leveraging interactions with each other. In one interpretation, they use an interaction to build a mental profile of their conversational partner and rely upon the profile to guide their composition of an effective message. To improve upon the amorphous inconsistency of human communication and to scale communication to reach many, free-form behavior can be replaced by a **communication strategy**. A communication strategy is a more structured and consistent approach to effective communication: first, an individual's profile is structured into a set of features based on information systematically elicited from them. Next, a messaging technique most appropriate to the feature set is selected and a message according to this technique is communicated. The messaging technique selection is facilitated by a mapping that assigns similar profiles to a **Group** and maps each Group to a messaging technique. An example of such a strategy is presented in the "Global Warming's Six Americas" studies (Maibach, Roser-Renouf, and Leiserowitz 2009). This study categorizes sampled individuals of the US population into six different Groups of climate change attitudes based on their responses to questions about their climate change beliefs and socioeconomic values. The climate change attitude Groups range from **Dismissive** to **Alarmed**. Subsequently, research identifies messaging techniques which are used as guidelines for effectively communicating considering the attitudes (Roser-Renouf et al. 2015). E.g. Use a **type** of **emotional** message that resorts to **cause and effect reasoning**, and then it finally empirically identifies how effective each messaging technique is for a Group by surveying multiple individuals of every Group based on what messaging technique they prefer.

Despite their efficiency, communication strategies are costly and time-consuming (Goldberg and Gustafson 2023). They may also be adversely sensitive to factors not considered within the profile (Leiserowitz and Thaker 2022). Given the vast amounts of human communication upon which a Large Language Model (LLM) was trained, we explore the use of it, specifically - **GPT-3.5-Turbo** (Ouyang et al. 2022), to support effective communication, with some form of communication strategy for different audiences facilitated by prompts and survey data (Argyle et al. 2023; Aher, Arriaga, and Kalai 2023; Park et al. 2022).

Our study has two research questions:

RQ 1 – Profile Matching Given survey information and prompt guidance, can the LLM create a profile and then match the (hidden) ground-truth Group?

RQ 2 – Message Selection Can the LLM accurately choose effective messages for a group given a prompt and a profile?

We demonstrate the ability of GPT-3.5-Turbo to adapt climate change messaging techniques to profile information. Our contributions are:

- We demonstrate that the LLM can generate individual profiles that maintain the climate change belief of the individual. We observe high sensitivity of the LLM to the prompts.
- We observe that the LLM can only provide the right message technique to a specific individual for certain Groups. We find evidence of a bias in both the LLM and supervised model's.

Related Work

An overview of related work in the areas of climate change communication, computational argumentation and LLM for

*We thank the EPFL WISH Foundation for the support provided to Iris Kremer.

social science is shown in Figure 1.

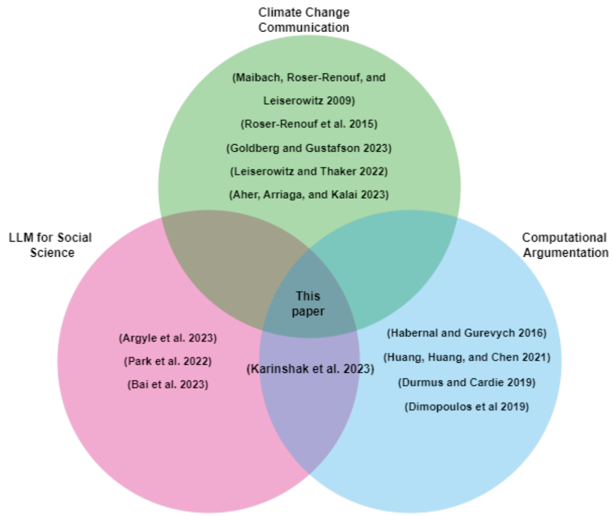

Figure 1: Overview of related work in the areas of climate change communication (green), computational argumentation (blue) and LLM for social science (red).

Climate change communication strategies, their design and their efficiency have been well studied. For example, one study segments the American population into six distinct groups, each with different responses to climate change messaging, advocating tailored communication strategies (Roser-Renouf et al. 2014). Similarly, another study introduces a brief, efficient tool for categorizing individuals into these segments, enabling targeted messaging(Maibach, Roser-Renouf, and Leiserowitz 2009). There is also a nuanced interaction between knowledge, ideology, and beliefs about climate change, suggesting that specific knowledge can potentially counteract ideological bias (Guy et al. 2014). Furthermore, one study explores the complex relationship between emotions, gender, political identity, and the reception of climate change messages, suggesting that a negative emotional framing may resonate differently based on the demographic composition of the audience(Bloodhart, Swim, and Dicicco 2019). This is further underscored by the cultural specificity of climate messaging, indicating that Indian perspectives on climate change are shaped by local impacts and cultural values (Leiserowitz and Thaker 2022). Lastly, there is a study of the efficacy of consensus messaging in altering beliefs about climate change, even among skeptical audiences, though the effects diminish over time (Goldberg et al. 2022). Collectively, these studies suggest a nuanced approach to climate change messaging that considers cultural, ideological, emotional, and cognitive factors in audience reception and engagement.

The field of *computational argumentation* aims at assessing argument persuasiveness or convincingness. For example, on paper investigated argument convincingness via feature-rich Support Vector Machine (SVM) and bidirectional LSTM (Habernal and Gurevych 2016). However, it fails to consider personal data and contextual factors influencing argument persuasiveness. Moreover, others partially addressed contextual factors with a Heterogeneous Argument Attention Network (HARGAN), a graph-based neural network model facilitating better viewpoint identification, but it lacked personalization and a specific focus on climate change (Huang, Huang, and Chen 2021). Other studies have attempted to incorporate the user's profile into the argumentation assessment. For instance, a dataset of online debates with comprehensive participant profiles, demonstrates the influence of user traits on debate outcomes (Durmus and Cardie 2019). This dataset allows for investigation of the effect of user characteristics and beliefs on argument persuasiveness. Finally, the concept of argumentation has also been explored in the context of negotiation by introducing a framework for argumentation-based negotiation even with incomplete opponent profiles (Dimopoulos, Mailly, and Moraitis 2019). The strategy involves an agent seeking arguments to support its stance, considering the uncertain knowledge of its opponent.

Nevertheless, contributions of Large Language Models (LLMs) to argumentation remain a relatively new domain, in particular *LLM for social science*. In a practical application of LLMs for social influence, (Karinshak et al. 2023) have used LLMs to generate pro-vaccination messages, indicating that LLMs can produce arguments that are perceived as convincing on a human level in a health context though they do not tailor the content to individual users. (Park et al. 2022) and (Bai et al. 2023) have demonstrated the progress in social computing systems and influencing human beliefs with LLMs. Further,(Argyle et al. 2023) explore 'silicon samples' that simulate human subpopulation opinions accurately. Despite these advances, the LLM's potential in assessing climate change arguments has not been fully explored. Our research explores the capabilities of LLMs for tailored argumentation within a climate change context.

Experimental Methodology

We study using a large language model to choose effective climate change messages and use the following methodology. We draw individuals and their information records from surveys of Americans assessing their beliefs and attitudes toward climate change (on Climate Change Communication YPCCC). We only use records $r \in \mathcal{R}$ that answered all the questions ($|\mathcal{R}| = 5836$). Each record contains responses to six climate-related questions r_A, five questions and responses r_B, that guide Group assignment, and 14 items of socio-economical information r_C. From r_A, we compute the record's (individual's) Group $r_{\mathbf{G}} = \{A, O, C, I, D, S\}$ (*Alarmed, Concerned, Cautious, Disengaged, Doubtful, Dismissive*) using the Short Americas Short Survey screening (Chryst et al. 2018).

Profile Matching We summarize the experimental method used to answer **RQ 1 – Profile Matching** in Fig.2. To answer this research question we use two prompts. In the **profile-generation Prompt**, we provide the LLM with survey information on the individual and direct it to generate a profile that considers this information. We provide r_B and responses to r_C in the **profile-generation Prompt**.

The **profile-generation Prompt** is composed of four segments:

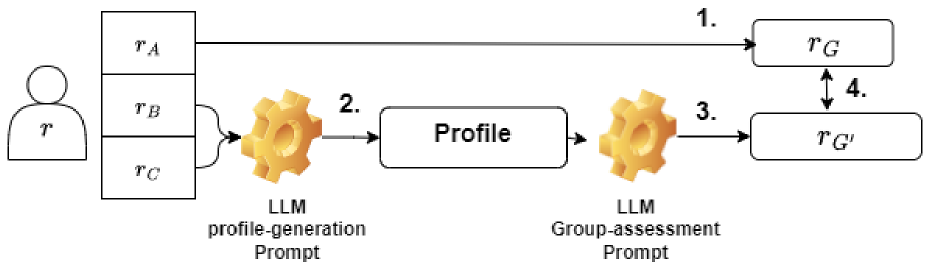

Figure 2: Diagram of the experimental method for **RQ 1 – Profile Matching**. 1. Get ground truth Group from the record r 2. LLM profile generation with information from a record. 3. Match the generated profile to a group. 4. Compare generated group with ground truth to measure the LLM performance on this task.

1. **Overview** segment provides the high-level task to the LLM to indicate how the LLM should interpret the rest of the instructions.
2. **Constraints** segment sets the constraints for solving the task. Namely, it directs the LLM to stay as close to the context data as possible and not hallucinate from training data.
3. **Description** segment present an in-depth explanation of the task. We aim the profile to be composed of both factual information about the person profiled and inferences and what belief the person might have. We also include in the task the inference of possible messaging techniques by the LLM to serve as guidelines for the next LLM-inference.
4. **Format** segment gives guidelines on *how* the answer should be written. It specifies the format and which information we want to obtain. More specifically, we prompt the LLM to obtain the answer in a paragraph that reads as a description instead of a list of features, which the LLM was outputting without this format instruction.

To expand the profile, we prompt the LLM in free-form language as shown in Fig. 3.

Similarly to the **profile-generation Prompt**, the **Group-assessment Prompt** is composed of five segments:

1. **Overview** segment provides the high-level task to the LLM to indicate how the LLM should interpret the rest of the instructions.
2. **Arguments** segment indicates which content will be passed in the context to the LLM.
3. **Description** segment present an in-depth explanation of the task. We ask the LLM to answer the SASSY 6-question survey (Chryst et al. 2018) to obtain the predicted record's Group
4. **Constraints** segment sets the constraints to solve the task. Namely, it directs the LLM to stay as close to the context data as possible and not hallucinate from training data.
5. **Format** segment gives guidelines on *how* the answer should be written. It specifies the format and which information we want to obtain. More specifically, we prompt the LLM to obtain the answer in a JSON format.

To expand the profile, we prompt the LLM in free-form language as shown in Fig. 4.

We run 120 trials of the two prompts along with the assessment prompt, sampling individuals from the six different Groups uniformly. We obtain the LLM's Group assignment by independently prompting it with each profile: using **Group-assessment Prompt**, p, we asked the LLM to answer the questions r_C from the SASSY 6-question survey (Chryst et al. 2018) that it was not provided when it generated the profile $r_{\mathbf{G}'} = f_\theta(p(r_B, r_C))$. We then check if the inferred Group of the profile matches the individual's ground-truth Group, $r_{\mathbf{G}} = r_{\mathbf{G}'}$. Specifically, we:

1. Compute the record's *true* **Group** $r_{\mathbf{G}}$ from r_A
2. Generate the record's profile from r_B and r_C using the **profile-generation Prompt**
3. Prompt the LLM to obtain $r_{\mathbf{G}'} = f_\theta(p(r_A, r_B))$
4. Check if $r_{\mathbf{G}} = r_{\mathbf{G}'}$

Message Selection To answer **RQ 2 – Message Selection,** we consider the communication to be effective if the LLM's message choice is the better choice given the ground-truth Group of the profile and the messaging technique with higher preference, see Figure 5. We use the profiles generated by the **profile-generation Prompt** to assess a messaging technique for each profile. In the **message-choice Prompt**, we pass the profile back to the LLM and prompt it to pick one of two messages, each following a different messaging technique. Each message is chosen out of a different set of climate change messages that we assemble per messaging technique (Roser-Renouf et al. 2015).

The messaging techniques, $\mathbf{m}$, are taken from the "Six Americas" study (Maibach, Roser-Renouf, and Leiserowitz 2009). They have two attributes: Category $\mathbf{c}$ and Type $\mathbf{t}$. Message categories are: *Evidence, E,* – provide evidence of a climate change event, *Cause, D* – describe the causes of climate change, *Consequence, C* – provide a consequence of climate change, and *Action, A* – given an example of addressing climate change. Message types are *Emotional, e* and *Scientific, s*. A messaging technique is defined as the cross-product between a type $\mathbf{t} = \{e, s\}$ and a category $\mathbf{c} \in \{E, D, C, A\}$ leading to 8 techniques $|\mathbf{m}| = |\mathbf{t} \times \mathbf{c}|$. We summarize the experimental setup to answer **RQ 2 – Message Selection** in Fig.5. More specifically, we:

```
You will be given someone's answer to a survey. Provide a profile of this person in terms of who they are, what they may believe in, and their possible stances regarding climate change.

Do not extrapolate your own views on the subject but only what you can imply for the answers to the survey. Be as accurate as possible and try to make links between each answer to provide a comprehensive and global profile.

Include the base information about the person, and extrapolate from their answer their possible views on climate change, but do not include the answers from the survey. Include also what kind of climate change messages they might be more inclined to listen to based on their profile. Also provide how this person, specifically, may be convinced to either take action or update their belief regarding climate change, based on their characteristics that are not necessarily related to their climate-related answers.

Write your profile as if you were describing a person. Do not write it as a list of characteristics. Except for the base information about the person, do not include the answers from the survey in your description, but write what you can infer from these answers. The person you describe should feel like a real person.
```

Figure 3: Profile generation prompt used for matching profiles. Colors indicate **Overview**, **Constraints**, **Description** and **Format**.

```
You goal is to assess what someone would have answered to a survey regarding climate change.

To do so you will be given the person's profile tagged as <PROFILE>, the survey question <QUESTION> and the multiple choice answers <ANSWERS> that you need to select.

Infer from the profile how the person described will answer to the provided question. Base yourself solely on what is included in the profile.

 DO NOT assume that the person is already familiar with the climate change issue if you do not have good reason to believe so.

Provide your answer in JSON format with your reasoning using the key ``reason'', and the selected answer with the exact same formulation as the one provided in <ANSWERS> using the key ``ans'' Example of output:
```

{"reason": "This answer", "ans": "chosen_anwser"}

Figure 4: Group attribution prompt used to answer the question that will be used to assess the Group of the user. Colors indicate **Overview**, **Arguments**, **Constraints**, **Description** and **Format**.

1. Compute the record's *true* **Group** $r_{\mathbf{G}}$ from r_C
2. Sample randomly 2 messages $m1$ and $m2$ and, using $r_{\mathbf{G}}$ we compute v_1 and v_2, the messages *preference* values for the record r
3. Generate the record's profile from r_B and r_C using the **profile-generation Prompt**
4. Prompt the LLM to obtain the LLM-chosen message
5. Compute the best message according to the ground-truth preferences v_1 and v_2
6. Check if the best message and the LLM-guessed message are the same

Prompts, parameters, examples, survey information and message sources can be found in the Supplementary Material (https://anonymous.4open.science/r/llm_assessing-5E2C/README.md). We note that prompt engineering is an active field of research focusing on the influence of textual prompts on the responses generated by LLMs. The intricacies between prompts and their corresponding outputs remain largely opaque. The methodology for crafting optimal prompts for designated tasks lacks a robust, systematic framework. However, certain strategies have emerged, such as the best practices for prompt engineering proposed by OpenAI (OpenAI 2023). The structure of the prompts used in this study were informed by guidelines suggested in recent literature, including those outlined by (White et al. 2023), but mostly based on trial and error.

Experiments

To answer **RQ 1 – Profile Matching**, we compare the **Group** assignment assessed by the LLM with ground truth. To answer **RQ 2 – Message Selection**, we compare three machine learning methods: 1. **LLM-No-Profile**, we do not provide the LLM any survey information, but the **profile-generation Prompt** still directs the profile generation and the **Group-assessment Prompt** is unchanged. This tests the *baseline* communication effectiveness of the LLM. 2. **LLM-With-**

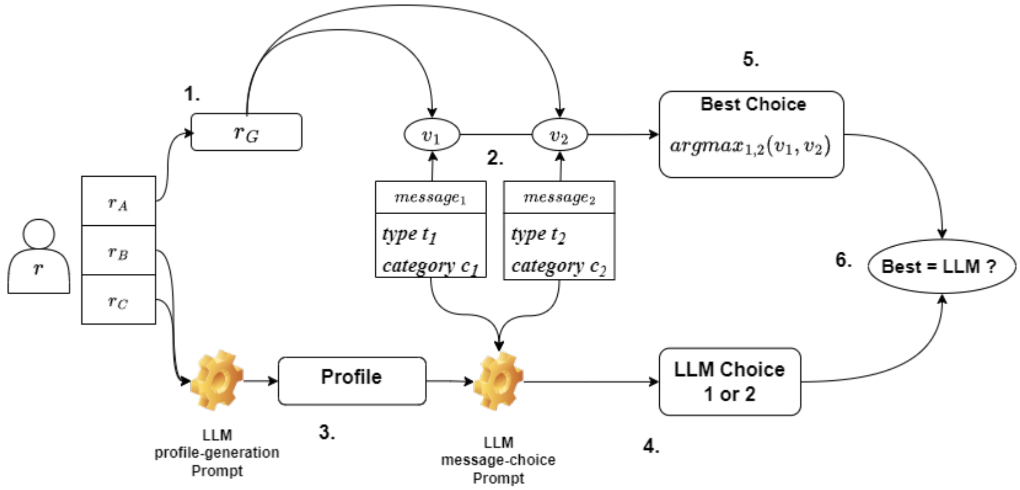

Figure 5: Diagram of the experimental method for **RQ 2 – Message Selection**. 1. Get ground truth Group from the person's record r 2. Sample 2 messages and compute message preference. 3. Generate profile with LLM. 4. Select message with the LLM. 5. Compute the best message according to the ground truth preferences 6. Compare LLM selected message and ground truth

Profile follows the previously described methodology. 3. For comparison, **Supervised**, we use a supervised approach with a trained MLP, $f_{\theta_{MLP}}$. The MLP is trained to choose one of two messages, given survey information. We input survey information and a message pair and $f_{\theta_{MLP}}$ outputs a preference for one of the two messages. We test the MLP on subjects and messages not seen during training to match the LLM experimental setup.

RQ 1 – Profile Matching: LLM Group Assignment

Table 1 answers **RQ 1 – Profile Matching**. Attitudinal Valence is the inclination to accept or reject the science of climate change (Roser-Renouf et al. 2014). Table 1 presents the confusion matrix of the LLM-predicted Attitudinal Valence grouping the *Alarmed* and *Concerned* Groups into *High* and the *Doubtful* and *Dismissive* Groups into *Low*. Overall, the LLM accurately predicts an individual's valence towards climate change, even given the incomplete information of the prompts. The confusion matrix for each Group can be found in the Supplementary Material. It shows varying accuracy across Groups: the *Alarmed* Group is more accurately predicted whereas the *Disengaged* Group was never outputted by the LLM. The accuracy for predicting the *Alarmed* group in may imply that an accurate Profile is a factor in the effectiveness of the LLM messaging technique choices. The overall results indicate some bias from the LLM and/or the prompts. Future work exploring different prompts could discriminate the source of this bias. We observe this profile-related bias in all our experiments. Finally, when we follow the YPCCC aggregation and group the *Cautious* and *Disengaged* groups as Middle (see supplementary material), the results seem to indicate that for Groups with low involvement in the topic of climate change, the survey questions and prompt directing are inadequate to inform an accurate Profile.

		LLM-Predicted Attitudinal Valence	
		Low	*High*
Ground Truth / Attitudinal Valence	*High*	0.00	1.00
	Low	0.82	0.18

Table 1: **RQ 1 – Profile Matching.** Confusion Matrix (normalized) for Group prediction of Attitudinal Valence where *Dismissive* and *Doubtful* Groups are combined into Low and *Alarmed* and *Concerned* into High.

RQ 2 – Message Selection: LLM Message Choice, Given profile and two message options

Figure 6 answers **RQ 2 – Message Selection** for three different machine learning methods (LLM No-profile, LLM With-Profile, Supervised), showing the LLM accuracy in selecting the message with the preferred messaging technique. We note that communication strategy regarding climate change varies significantly between Groups and machine learning methods.

The supervised model demonstrated a gradual decrease in accuracy based on attitude, from the Groups that have a low attitudinal valence (*Dismissive*) to high (*Alarmed*). An exception to this trend is the *Disengaged* Group, where the supervised model showed a marked improvement. The LLM's accuracy differed when it was not given any profile information, compared to when it had a profile to work with. When given no profile (LLM No-Profile), the model performed significantly better than average only for the *Disengaged*. However, with profiles (LLM With-Profile), the model showed a gradual increase in accuracy for the Groups who had high valence. Notably, it performed better for 'Alarmed' individuals when provided with their profiles. The *Disengaged* Group was unique in that not providing profile information yielded better accuracy.

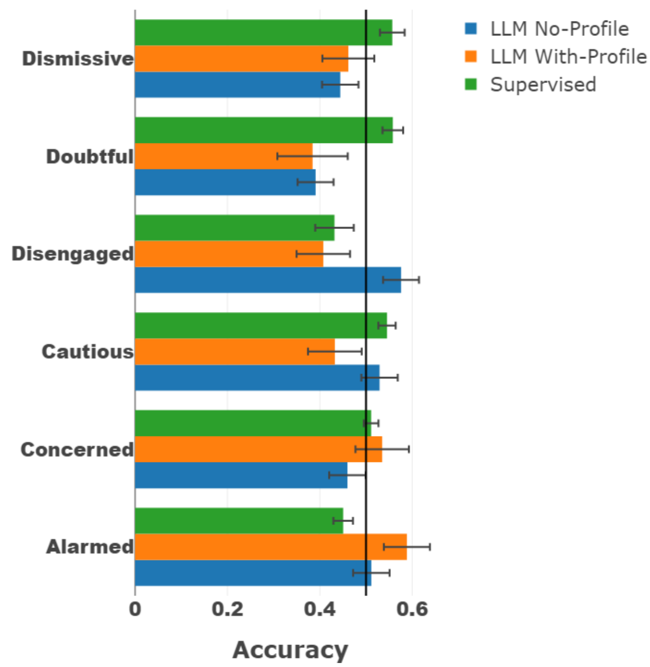

Figure 6: **RQ 2 – Message Selection** Accuracy of message selection of message techniques compared to ground truth. Accuracy (x-axis) per Group (y-axis), Error bars are 95% CI. Color indicates the machine learning method used.

The LLM's accuracy differed when it was not given any profile information, compared to when it had a profile to work with. When given no profile (LLM No-Profile), the model performed significantly better than average only for the *Disengaged*. However, with profiles (LLM With-Profile), the model showed a gradual increase in accuracy for the Groups who had high valence. Notably, it performed better for 'Alarmed' individuals when provided with their profiles. The *Disengaged* Group was unique in that not providing profile information yielded better accuracy. We relate this observation with the LLM Group Assignment experiment in which it did not receive any profile assignment from the model. The result could possibly be attributed to the fact that individuals from this Group tend to be neutral towards climate change. That is, they neither care about nor pay attention to it. For the *Alarmed* Group, the model effectively used the profile information and showed the best accuracy in assigning profiles, thereby raising questions about possible bias. It is unclear if the model's accuracy is due again to the LLM sensitivity to prompting or if it indicates a bias inherent to the model itself. To discern the cause, future studies could employ different language models and assess if the observed bias persists.

We observed that the LLM exhibited noticeable shifts in messaging technique attribution per Group, albeit limited, see Figure 7. The results also raised questions regarding the low overall accuracy and ability of the Large Language Models using profile information. This could be attributed to several factors. First, the LLM might give preference to certain properties of the messages and the individuals that do not align with our method of assessing message value. Further research needs to delve into understanding what factors the LLM is favoring. Moreover, it could be a case of not having the right information in the profile to infer the right technique, even though we could infer the right Group from the given data. Finally, LLMs are sensitive to prompting, so we can explore different prompting techniques, e.g. few-shot learning.

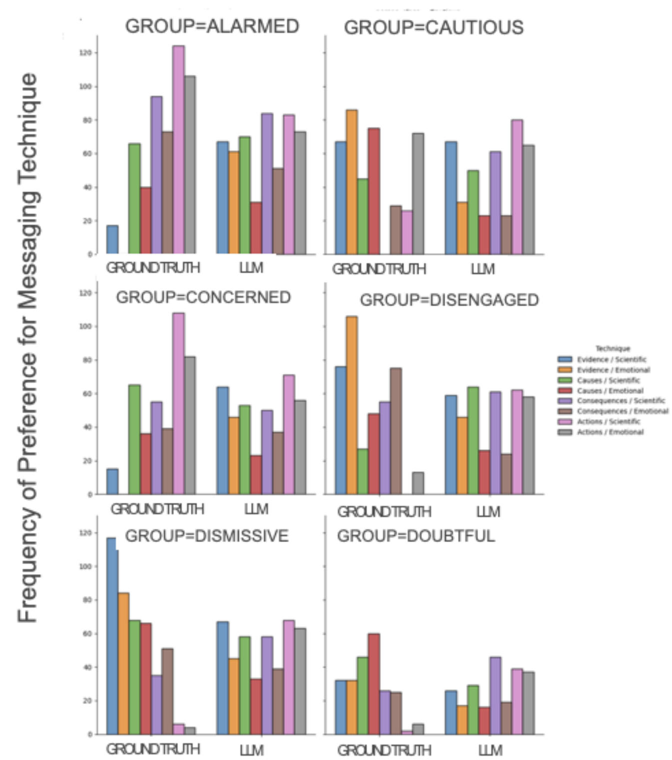

Figure 7: **RQ 2 – Message Selection** Distribution of message techniques compared to ground truth. Each sub plot is a Group. It shows ground truth distribution from surveys on left portion of X-axis and to its right equivalent for LLM. Y-axis is a count of preferences for a particular messaging technique. There are 8 messaging techniques, each displayed in a different color and left to right: Evidence+Scientific, Evidence+Emotional, Causes+Scientific, Causes+Emotional, Consequences+Scientific, Consequences+Emotional, Actions+Scientific, Actions+Emotional.

Conclusions

We demonstrated the ability of GPT-3.5-Turbo to adapt climate change messaging techniques to individuals. Three methods were presented to infer communication strategies regarding climate change: supervised method, the LLM with a profile, and the LLM without a profile. Each method favors different Groups among the Six Americas of climate change perception.

From our settings and experiments our contributions were: **RQ 1 Group Matching:** We observe that for high and low valences individuals, groups can be predicted from LLM-generated profiles. Thus, we demonstrate that the LLM can generate individual profiles that maintain the climate change belief of the individual, even with incomplete information. However, the uneven results across Groups raise concern regarding the sensitivity of the model to the prompts designed in this study. **RQ 2 Message Selection:** We observe that the LLM can for certain Groups only provide the right message technique to a specific individual, we find evidence of a bias in both the LLM and supervised model's climate change communication strategy. With our settings and prompts their respective approaches work only on certain Groups of individuals and not others.

References

Aher, G.; Arriaga, R. I.; and Kalai, A. T. 2023. Using Large Language Models to Simulate Multiple Humans and Replicate Human Subject Studies. arXiv:2208.10264.

Argyle, L. P.; Busby, E. C.; Fulda, N.; Gubler, J. R.; Rytting, C.; and Wingate, D. 2023. Out of One, Many: Using Language Models to Simulate Human Samples. *Political Analysis*, 31(3): 337–351.

Bai, H.; Voelkel, J. G.; Eichstaedt, j. C.; and Willer, R. 2023. Artificial Intelligence Can Persuade Humans on Political Issues.

Bloodhart, B.; Swim, J. K.; and Dicicco, E. 2019. "Be Worried, be VERY Worried:" Preferences for and Impacts of Negative Emotional Climate Change Communication. *Frontiers in Communication*, 3.

Chryst, B.; Marlon, J.; Van Der Linden, S.; Leiserowitz, A.; Maibach, E.; and Roser-Renouf, C. 2018. Global Warming's "Six Americas Short Survey": Audience Segmentation of Climate Change Views Using a Four Question Instrument. *Environmental Communication*, 12(8): 1109–1122.

Dimopoulos, Y.; Mailly, J.-G.; and Moraitis, P. 2019. Argumentation-based Negotiation with Incomplete Opponent Profiles. In *13èmes Journées d'Intelligence Artificielle Fondamentale (JIAF 2019)*, Actes JIAF 2019, 91–100. Toulouse, France.

Durmus, E.; and Cardie, C. 2019. A Corpus for Modeling User and Language Effects in Argumentation on Online Debating. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics.

Goldberg, M. H.; and Gustafson, A. 2023. Weak Evidence that Tailoring Environmental Messages is More Persuasive: Findings From a Systematic Review.

Goldberg, M. H.; Gustafson, A.; van der Linden, S.; Rosenthal, S. A.; and Leiserowitz, A. 2022. Communicating the Scientific Consensus on Climate Change: Diverse Audiences and Effects Over Time. *Environment and Behavior*, 54(7-8): 1133–1165.

Guy, S.; Kashima, Y.; Walker, I.; and O'Neill, S. 2014. Investigating the effects of knowledge and ideology on climate change beliefs. *European Journal of Social Psychology*, 44.

Habernal, I.; and Gurevych, I. 2016. Which argument is more convincing? Analyzing and predicting convincingness of Web arguments using bidirectional LSTM. In Erk, K.; and Smith, N. A., eds., *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 1589–1599. Berlin, Germany: Association for Computational Linguistics.

Huang, K.-Y.; Huang, H.-H.; and Chen, H.-H. 2021. HARGAN: Heterogeneous Argument Attention Network for Persuasiveness Prediction. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(14): 13045–13054.

Karinshak, E.; Liu, S. X.; Park, J. S.; and Hancock, J. T. 2023. Working With AI to Persuade: Examining a Large Language Model's Ability to Generate Pro-Vaccination Messages. *Proc. ACM Hum.-Comput. Interact.*, 7(CSCW1).

Leiserowitz, A.; and Thaker, J. 2022. Climate Change in the Indian Mind 2022. Technical report, Yale Program on Climate Change Communication, United States of America. Accessed on 26 Sep 2023.

Maibach, E.; Roser-Renouf, C.; and Leiserowitz, A. 2009. Global Warming's Six Americas 2009: An Audience Segmentation Analysis.

on Climate Change Communication (YPCCC), Y. P.; and for Climate Change Communication (Mason 4C), G. M. U. C. 2020. Climate Change in the American Mind: National survey data on public opinion (2008-2018) [Data file and codebook].

OpenAI. 2023. Best Practices for Prompt Engineering with OpenAI API.

Ouyang, L.; Wu, J.; Jiang, X.; Almeida, D.; Wainwright, C. L.; Mishkin, P.; Zhang, C.; Agarwal, S.; Slama, K.; Ray, A.; Schulman, J.; Hilton, J.; Kelton, F.; Miller, L.; Simens, M.; Askell, A.; Welinder, P.; Christiano, P.; Leike, J.; and Lowe, R. 2022. Training language models to follow instructions with human feedback. arXiv:2203.02155.

Park, J. S.; Popowski, L.; Cai, C.; Morris, M. R.; Liang, P.; and Bernstein, M. S. 2022. Social Simulacra: Creating Populated Prototypes for Social Computing Systems. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology*, UIST '22. New York, NY, USA: Association for Computing Machinery. ISBN 9781450393201.

Roser-Renouf, C.; Stenhouse, N.; Rolfe-Redding, J.; Maibach, E.; and Leiserowitz, A. 2014. Engaging Diverse Audiences with Climate Change: Message Strategies for Global Warming's Six Americas.

Roser-Renouf, C.; Stenhouse, N.; Rolfe-Redding, J.; Maibach, E.; and Leiserowitz, A. 2015. Engaging diverse audiences with climate change: Message strategies for Global Warming's Six Americas. In Cox, R.; and Anders, H., eds., *Handbook of Environment and Communication*.

White, J.; Fu, Q.; Hays, S.; Sandborn, M.; Olea, C.; Gilbert, H.; Elnashar, A.; Spencer-Smith, J.; and Schmidt, D. C. 2023. A Prompt Pattern Catalog to Enhance Prompt Engineering with ChatGPT. arXiv:2302.11382.