Multi-Agent Based Message Generation and Delivery for Personalized Environmental Awareness Promotion in Urban City

Anonymous ACL submission

Abstract

With growing concerns about climate change, effectively promoting pro-environmental behavior becomes a pressing societal challenge. While traditional publicity strategies are usually general, targeting diverse citizen profiles and behavioral motivations with different strategies, sounds to be much more appropriate. In this paper, we propose MA-MGD, a Multi-Agent Based Message Generation and Delivery Framework, which use multi-agent system to generate personalized publicity strategy for citizen with different profiles and behavioral motivations, in a context of environment protection promotion. The system consists of citizen profiling, simulation testing and iterative feedback to promote environment-friendly living styles and low-carbon behavioral changes. Deployed on the AgentSociety platform, the system targets 200 virtual citizens in a simulated Beijing environment and dynamically delivers personalized messages and posters through multi-agent collaboration, to promote environment protection. Experimental results demonstrate that, compared to static or template-based methods, MA-MGD significantly improves message rationality, carbon awareness, and low-carbon travel intent. Our findings highlight the potential of LLM-based multi-agent frameworks in enabling cost-effective, adaptive, and behaviorally impactful environmental interventions.

1 Introduction

Climate change and environmental degradation pose increasingly severe threats to global sustainability. In response, fostering public environmental awareness and promoting pro-environmental behaviors have become urgent priorities for policymakers and communities alike. Previous efforts(Qi et al., 2018; Du et al., 2019; Kong et al., 2023; Fu et al., 2024) have focused primarily on monitoring and early warning of ecological pollution, with the aim of conveying environmental risks through quantitative data to raise public awareness of environmental problems, which is not sufficient. In addition, other traditional publicity approaches, such as standardized announcements and generic campaigns, often fail to effectively engage diverse audiences. They tend to overlook differences in individual psychology, demographic characteristics, and behavioral patterns, resulting in limited success in driving behavioral change(Mendiola and Hechanova, 2025). Therefore, there is an urgent need for a personalized publicity approach that integrates the above factors to enable precise communication and dynamic feedback, thus effectively promoting environmental awareness and behavioral change.

Traditional large language models(LLMs) primarily focus on text-based dialogue and often struggle with complex tasks. With technological advancements, agent technology has emerged as a key tool and pathway to innovating and expanding the application scope of large-language models. Through agent-based systems, large models can better realize LLMs' potential across a broader range of domains. Among them, multimodal agents and tool use have enabled practical breakthroughs in real-world applications. Ghost in the Minecraft (Zhu et al., 2023) demonstrates the potential of language models in handling long-term and complex tasks by integrating a text-based knowledge base with long-term memory management. In an open-world game environment, the agent successfully acquires all items in the full technology tree, achieving a 47.5% improvement in task success rate. Auto-GPT style agent (Yang et al., 2023b) introduces the Additional Opinions algorithm, which enhances online decision-making through supervised and imitation learning without fine-tuning the base model. It significantly outperforms traditional methods on tasks such as WebShop and ALFWorld. AVLEN(Paul et al., 2022) establishes a multimodal hierarchical reinforcement learning framework, enabling agents to autonomously navigate based on audio and visual cues. It also supports requesting natural language assistance from humans, resulting in significantly improved navigation performance in complex environments with background noise. Agents(Li et al., 2024; Qian et al., 2024; Zhou et al., 2024; Hong et al., 2024; Huang et al., 2024) have already demonstrated exceptional performance and broad compatibility across a wide range of tasks.

To enhance the effectiveness of environmental publicity, we propose a multi-agent system for targeted message delivery based on an iterative supervised feedback mechanism(**MA-MGD**). Specifically, MA-MGD identifies target audiences precisely, dynamically adjusts content generation strategies, and demonstrates significantly better behavioral guidance outcomes compared to traditional template-based approaches through system-level validation. It not only advances the application of personalized agent interactions in the domain of environmental governance but also provides a generalizable and reusable framework for integrating multi-agent systems with social behavior interventions.

Our main contributions are summarized as follows:

- We propose MA-MGD, a Multi-Agent Based Message Generation and Delivery Framework. It enables a closed-loop process from citizen profile construction and audience selection to custom content generation and iterative feedback optimization.
- MA-MGD introduces a closed-loop "simulation testing + iterative feedback" mechanism to enhance the effectiveness of information campaigns.
- Extensive qualitative and quantitative experiments are conducted on the AgentSociety(Piao et al., 2025) platform, demonstrating the high efficiency and effectiveness of the framework.

2 Related Work

2.1 Single-Agent Systems

Single-agent systems have traditionally served as the foundational architecture for artificial intelligence applications, where a single model or entity interacts with the environment, makes decisions, and optimizes actions based on its internal state or learned policy. These systems have been successfully deployed in a variety of domains, including game playing (Mnih et al., 2015; Shao et al., 2019; Brown et al., 2020), robotics(Levine et al., 2016; Gu et al., 2017; Bruce et al., 2017; Tang et al., 2025) , dialog systems(Zhang et al., 2019), and even environmental modeling(Qi et al., 2018; Fu et al., 2024). In such settings, the single-agent paradigm simplifies control and decision making by centralizing intelligence within one agent, enabling stable performance in relatively closed or deterministic environments.

2.2 Multi-agent systems

Multi-agent systems (MAS)(Yang et al., 2025; Yu et al., 2025; Ye et al., 2025; Haase and Pokutta, 2025), have emerged as a powerful alternative to traditional single-agent architectures, offering decentralized coordination, parallel reasoning, and task specialization across diverse agents. Unlike single agents, which struggle to scale in complex, interactive settings, MAS enable dynamic role allocation and real-time feedback integration—features critical for domains like urban simulation, education, and environmental governance(Balaji and Srinivasan, 2010; Stone and Veloso, 2000). Recent advances in large language models (LLMs) have further empowered MAS, giving rise to intelligent agent teams that can collaborate, plan, and adapt autonomously. For instance, MetaGPT(Hong et al., 2024) coordinates agents in software engineering workflows, Auto-GPT(Yang et al., 2023a) and ChatDev(Qian et al., 2024) demonstrate iterative decision-making via agent communication, and AgentSociety(Piao et al., 2025) showcases how generative agents simulate human behavior at scale in social-environmental interventions. Additionally, Generative Agents(Park et al., 2023) in Smallville highlight how agents equipped with memory and planning can create believable long-term social interactions, reinforcing MAS as a promising framework for modeling and influencing real-world collective behavior.

3 Method

3.1 Simulated scenario

The city simulated in AgentSociety is based on Beijing, constructed using real-world geographic information to create a virtual environment that includes various functional urban areas such as residential zones, commercial districts, office areas, and transportation hubs. The simulation features 200 virtual citizens, each with a unique personal profile that

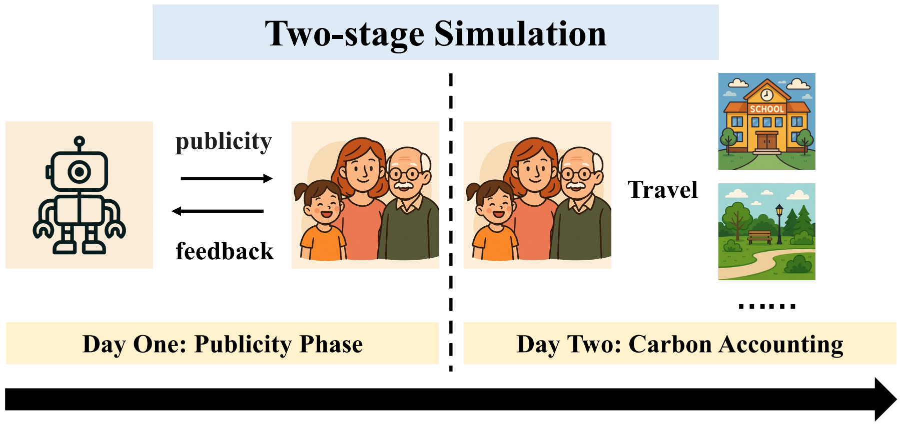

Figure 1: Two-Day Environmental Publicity Process on the AgentSociety Platform

includes attributes such as age, gender, educational background, occupation, and marital status, as well as personalized background stories and lifestyle habits. These citizens make different behavioral choices based on their individual characteristics and level of environmental awareness. AgentSociety runs a two-day simulation: the first day focuses on promoting low-carbon lifestyles through various agent-driven strategies, and the second day tracks carbon emissions. Lower costs and emissions indicate more effective publicity methods. More details are given in Appendix A .

3.2 Citizen Profiling

To develop more targeted publicity strategies instead of one-fits-all dissemination, we conduct citizen surveys and constructed detailed user profiles. Specifically, we start with the profiles of 200 simulated citizens and categorize them into five representative categories (Table 1), each reflecting a primary behavioral driving force. We primarily assign citizens to these five archetypes based on their responses to an environmental awareness questionnaire. More details are given in Appendix B.

3.3 Agent Framework

The general agent framework that we designed is shown in the figure2. To effectively reduce publicity costs, the system first selects appropriate target citizens for communication. A simulated assessment is first conducted to identify citizens with relatively low environmental awareness, who are then selected as the targets for subsequent tailored publicity interventions aimed at improving their environmental consciousness. For citizens with low environmental awareness, two intervention methods are designed: (1)Multi-agent iterative modification and targeted message push, and (2) Multi-agent iterative modification and targeted poster push.

Multi-agent iterative modification and targeted message push.

As shown in Figure 3, we build a multi-agent system to support the full process of targeted message delivery and iterative content optimization. The system first identifies citizens with relatively low environmental awareness based on the results of a simulated assessment and obtains their corresponding ID list. Then it retrieves detailed profile information for each individual based on these IDs, providing data support for subsequent personalized publicity. In each round, we randomly select 5 individuals from the target group for interaction. The entire process consists of 32 rounds, theoretically covering up to $32 \times 5 = 160$ citizens.

Based on the collected citizen profiles, we construct a Citizen Attitude Simulation Agent to model potential changes in citizens' environmental awareness and attitudes after receiving the publicity, thereby providing a feedback mechanism for content optimization. Meanwhile, the Citizen Dialogue Generation Agent generates targeted environmental messages tailored to each individual's background information, aiming to enhance the receptiveness and persuasiveness of the communication. In addition, the Dialogue Revision Agent dynamically adjusts and refines the publicity content by incorporating historical message records and the trajectory of attitude changes. After N rounds of iteration between the Dialogue Revision Agent and the Citizen Attitude Simulation Agent, the system generates a more refined and effective version of the customized message, which is then formally delivered to the target citizen to achieve precise intervention.

Multi-agent iterative modification and targeted

Table 1: Five major citizen archetypes.

User Archetype	Driving Force	Key Characteristics	Communication Strategy
Policy Followers	Rule compliance	“Follow directives”, “Regulations”	Refer to policies and social responsibilities
Interest-Driven	Economic benefits	“Save money”, “Rewards”	Emphasize cost savings and incentives
Social-Driven	Group conformity	“Peer influence”, “Follow the crowd”	Highlight neighbors’ or community participation
Wavering Pragmatists	Convenience and cost	“Easy”, “Time/money saving”	Offer low-threshold, convenient options
Value-Driven Vanguards	Intrinsic moral values	“Responsibility”, “Volunteering”	Encourage them to become advocates or promotion agent

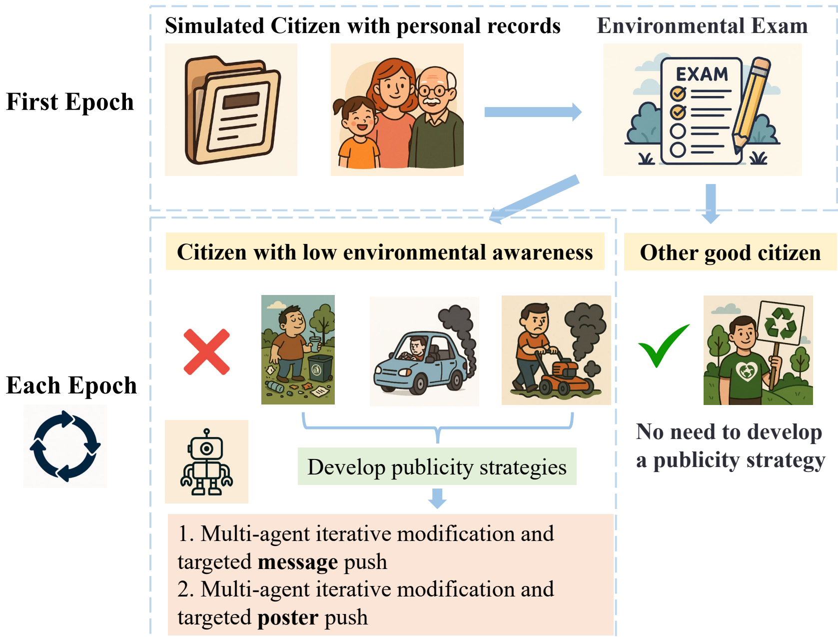

Figure 2: Architecture Overview of the multi-agent system for targeted information framework based on user profiling(MA-MGD)

poster push.

As shown in Figure 4, this diagram illustrates the use of a multi-agent system for targeted poster delivery and iterative optimization. After obtaining the list of citizens to be targeted (those with low environmental awareness scores), we retrieve their corresponding AOI IDs (Area of Interest identifiers) based on their residential locations. Due to budget constraints, a total of only 33 posters can be deployed, which matches the number of action rounds.

After obtaining the AOI IDs, we retrieve the IDs of residents living in the corresponding areas and select the target individuals for publicity. Their detailed profiles are then accessed to generate customized posters based on this information. To ensure that the poster content maintains a high level of rationality, the system incorporates a Poster Rationality Scoring Agent, which works in coordination with the Poster Revision Agent. Together, they perform multiple rounds of iterative refinement on the poster content until its rationality reaches an acceptable level. The finalized poster versions are then generated and displayed.

More details of the method are given in Appendix C

4 Experiment

4.1 Evaluation metrics

The model used in this study is Qwen2.5-14B-Instruct(Yang et al., 2024). Since this study is conducted on the AgentSociety(Piao et al., 2025) platform, the platform requires environmental promote agents to maximize the effectiveness of environmental publicity under the constraint of a limited resource budget. The evaluation system is based

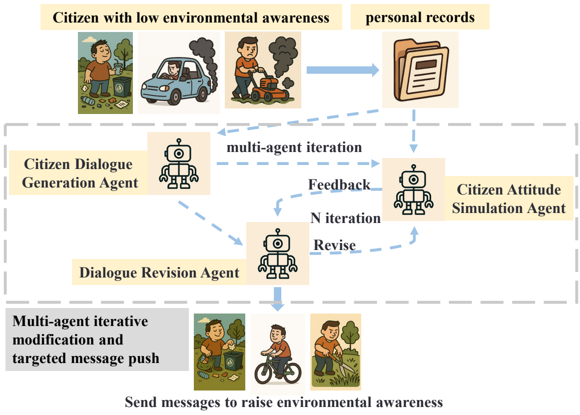

Figure 3: Structural Diagram of Multi-Agent Iterative Modification and Targeted Message Push

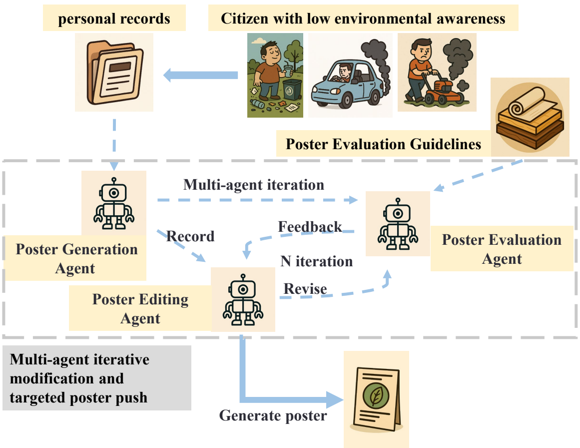

Figure 4: Structural Diagram of Multi-Agent Iterative Modification and Targeted Poster Push

on three core metrics:

- **Publicity Score**: Evaluates the impact of environmental promotion on residents' environmental awareness, based on data collected from surveys (0–50 points).
- **Carbon Emission Score**: Assesses citizens' specific transportation choices and calculates the corresponding carbon emissions (0–50 points).
- **Publicity Rationality**: Dynamically monitors the rationality of publicity content delivered by environmental promotion agents, including one-on-one messages, posters, and announcements (scored as a multiplier between 0 and 1).

In the AgentSociety platform, environmental publicity is carried out through three main approaches: one-on-one message delivery, targeted poster placement, and city-wide announcements. One-on-one messaging does not consume any funds but is limited to a maximum of 5 citizens per round. Poster placement can reach all residents within a specified Area of Interest (AOI) and costs 3000 units of funds per poster. Announcements are broadcast to all citizens at once, with each announcement costing 20,000 units of funds.

The final score is as following:

$$\text{Final Score} = (\text{P-Score}+\text{C-Score})\times\text{Rationality}$$

Publicity Score. The publicity score is assessed based on citizens' awareness of and behavioral changes toward low-carbon lifestyles, quantified through a standardized questionnaire that measures their level of environmental consciousness.

Carbon Emission Score. The carbon emission score evaluates the actual effectiveness of environmental publicity by monitoring citizens' real transportation behaviors and calculating the corresponding changes in carbon emissions.

Publicity Rationality. The publicity rationality metric leverages an AI evaluation system to continuously monitor all publicity content generated by the environmental promotion agent, ensuring its credibility and appropriateness.

More details of the evaluation metrics are given in Appendix D

4.2 Cost-effectiveness Comparison of Publicity Methods

In AgentSociety, we employ Large Language Models (LLMs) to automatically evaluate the rationality of publicity materials. However, due to the generative and evaluative nature of LLMs, their scores tends to be conservative, making it difficult to achieve extreme high scores. Specifically, the rationality score of the publicity materials is calculated as follows:

$$\frac{0.1\times\text{M-Score}+0.3\times\text{P-Score}+0.6\times\text{A-Score}}{100} \tag{1}$$

Where M score is the message score, P score is the poster score, and A score is the announcement score.

In the experiment, we find that regardless of the iterative optimization strategy applied, the individual scores given by the LLM typically fall between 85 and 95. Even in the best-case scenario, it is difficult to exceed the upper limit of 100 points.

Therefore, in practical applications, whether the drop in rationality caused by certain publicity methods is acceptable needs to be weighed against the improvement in environmentally friendly behavior they bring.

To further analyze the impact of different combinations of publicity methods on rationality, we use an ideal individual rationality score of 88 as the baseline and simulate the overall rationality scores under various strategy combinations.

Table 2: Rationality Scores under Different Publicity Strategy Combinations

Publicity Strategy	**Rationality Score**
Message only	**98.8**
Message + Poster	95.2
Message + Poster + Announcement	88.0
Poster only	96.4
Announcement only	92.8
Message + Announcement	91.6

The results shown in table 2 indicate that using only the message-based communication strategy yields the highest overall rationality score. In contrast, when all three publicity methods are used together, the rationality score is the lowest. This suggests that although multi-channel publicity may enhance the breadth and depth of information dissemination, it can also lead to a decline in overall rationality due to the limitations of independent scoring across different media.

Based on this analysis, we conduct ablation experiments in the following sections to evaluate the effects of individual and combined publicity strategies. We further explore how to maximize the rationality score while ensuring the effectiveness of publicity.

4.3 Ablation study

Ablation study of publicity material generation methods. The results of the ablation experiment on publicity material generation are shown in Table 3. As the data indicates, different generation methods exhibit significant differences in both publicity rationality and their impact on pro-environmental behavior. Under the baseline condition (i.e., no publicity), although the rationality score is perfect, the scores for promoting low-carbon travel and carbon emission awareness are relatively low, suggesting that the absence of publicity material negatively affects residents' behavior. When using fixed-text publicity, the rationality remains near optimal, but its effectiveness in enhancing residents' low-carbon awareness is limited. With the introduction of resident categorization strategies and the exclusion of groups resistant to change, the overall effectiveness improved, indicating that personalized publicity enhances information acceptance. Although the single-shot generation method performed relatively well in conveying carbon emission knowledge, it remained suboptimal in guiding actual behavioral change. In contrast, the iterative refinement approach achieved the best results across all metrics—particularly excelling in promoting low-carbon travel willingness and improving carbon emission awareness—with a final total score of 73.28, significantly outperforming other methods. This demonstrates that continuously optimizing and adjusting publicity content can more effectively achieve low-carbon promotion goals.

Ablation study on the Selection of Publicity Strategies.

To investigate the impact of different publicity methods on residents' environmental awareness and behavior, we conduct the following experiment, with the results shown in Table 4. We first test releasing announcements alone and find that although this approach helps improve the overall carbon emission questionnaire scores (carbon emission result) among residents, its effectiveness is limited. This is mainly because announcements reach all residents, including those with relatively weak environmental awareness. In contrast, when we generate personalized messages based on each resident's profile and pair them with corresponding posters, the survey scores (survey result) increase noticeably, but the carbon emission questionnaire scores show little change.

Although announcements can to some extent improve knowledge about carbon emissions and increase survey scores, their overall effectiveness in promotion drops significantly, resulting in no notable improvement in the total score. These observations highlight the importance of using a combination of publicity tools and suggest that different publicity methods need to be optimized for specific objectives.

According to the generated planning logs, most agents tend to choose non-travel behaviors such as "Stay at home" and "Eat at home", influenced

Table 3: Fusion Experiment Results of Publicity Material Generation Methods

Generation Method	Rationality	Low-carbon Score	Emission Score	Total Score
No Publicity (Baseline)	**1**	60.00	71.49	65.74
Fixed Text Template	0.99	60.00	75.46	66.93
Resident Categorization without the Unemployed	0.99	68.93	74.51	70.83
Single-shot Generation	0.99	62.63	**78.50**	69.65
Iterative Refinement	0.99	**70.35**	77.96	**73.28**

Table 4: Ablation Study on Publicity Strategy Selection

Publicity Strategy	Rationality	Low-carbon Travel	Emission Awareness	Total Score
Message	**0.9882**	70.35	77.96	**73.28**
Poster	0.9586	70.00	75.05	69.52
Announcement	0.9300	65.02	71.98	63.77
Message + Poster	0.9494	69.18	79.58	70.61
Message + Announcement	0.9184	**72.52**	78.76	69.47
Announcement + Poster	–	–	–	–
Message + Announcement + Poster	0.8830	71.00	**80.00**	66.66

by a combination of contextual factors including time of day, occupation, and other background variables. This behavioral pattern is often reinforced by favorable weather conditions and the absence of urgent needs, further enhancing the rationality and attractiveness of staying at home.

Since the majority of plans do not trigger actual travel behaviors, the system's traffic logging mechanism remains inactive, resulting in a travel distance of zero for most individuals and, consequently, zero actual carbon emissions. As a result, the overall carbon emissions are contributed solely by a small subset of individuals who do engage in travel, and it is only those people whose travel modes may be influenced by environmental attitudes.

This indicates that although the environmental publicity mechanism in the model can adjust residents' environmental attitudes, its actual regulatory effect on carbon emissions is relatively limited in this round of experiments due to the prevailing trend of stay-at-home behavior. Ultimately, the overall carbon emission scores remain within a narrow fluctuation range between 63 and 69.

Ablation Study on Iterative Refinement Approach.

To further optimize the effectiveness of publicity material generation, we conduct an ablation study on the iterative refinement approach, examining the impact of different iteration counts on publicity rationality, willingness for low-carbon travel, and carbon emission awareness. The experimental results are presented in Table 5.

As shown in Table 5, when the number of iterations is set to 2, the model achieves the highest scores in both the carbon emission questionnaire and the total score, indicating optimal overall performance. However, as the number of iterations increases to 3 or more, although the publicity rationality remains stable or even slightly improves, the scores related to carbon emissions and the overall score show a slight decline. This suggests that excessive iterations may lead to content redundancy or deviation from user preferences, thereby reducing the effectiveness of the publicity.

Therefore, conducting 2 iterations proves to be the optimal choice under the current experimental settings, as it maintains high rationality while maximizing residents' acceptance of environmental information and their behavioral response.

Comprehensive Ablation Study.

In addition to the systematic experiments on publicity methods and iterative strategies described above, we also explore several other approaches and summarize the key results in Table 6 as part of a comprehensive ablation study, aiming to further evaluate the overall effectiveness of the publicity strategy.

The experimental results show that under the

Table 5: Ablation Study on Iterative Refinement Approach

Refinement Strategy	Rationality	Low-carbon Travel	Emission Awareness	Total Score
Single iteration	0.99	62.63	78.50	69.65
2 iterations	0.99	62.63	**78.86**	**69.91**
3 iterations	0.99	62.63	78.13	69.55
5 iterations	0.99	62.63	78.10	69.54

Table 6: Comprehensive Ablation Study

Publicity Method	Rationality	Low-carbon Travel	Emission Awareness	Total Score
No Publicity (Baseline)	**1**	60.00	71.49	65.74
Iterative Rationality Optimization	0.99	66.88	77.56	71.23
Random Citizen Selection (No Simulation Test)	0.99	68.87	76.25	71.47
Targeted Citizen Selection (With Simulation Test)	0.99	**70.35**	**77.96**	**73.28**

baseline condition without any publicity, the average total score of residents is 65.74, indicating relatively limited performance. However, with the introduction of publicity mechanisms—especially when combining simulation testing with targeted citizen selection—the overall score significantly increases to 73.28, representing the best performance among all methods.

Through systematic ablation experiments and comparative analysis of various publicity strategies and generation methods, we comprehensively evaluate the effectiveness of different approaches in enhancing residents' environmental awareness and promoting behavioral change. The results indicate that the publicity strategy combining simulation testing with an iterative refinement mechanism outperforms all others across multiple metrics.

Specifically, compared to the baseline without any publicity (total score is about 65.74), the adoption of targeted publicity methods significantly improves residents' awareness and willingness to engage in low-carbon travel. Among all tested approaches, the combination of "simulation-based targeting + message refinement with 2 iterations" achieves the highest overall score of 73.28, striking a well-balanced performance across rationality, low-carbon travel score, and carbon emission questionnaire score.

Conclusion

In this study, we propose MA-MGD, a multi-agent system that leverages citizen profiling, simulation testing, and iterative feedback to personalize environmental publicity and promote low-carbon behaviors. Through extensive experiments conducted on the AgentSociety platform, our results demonstrate that targeted communication—especially when combined with two rounds of message refinement—significantly improves both environmental awareness and pro-environmental behavior, outperforming static and generic strategies across multiple evaluation metrics. Furthermore, the ablation studies confirm that optimal publicity outcomes depend not only on the choice of medium but also on the dynamic tailoring of content to individual motivations. This work opens new avenues for applying large language models to real-world environmental governance through adaptive, scalable, and human-centered communication strategies.

Limitations

While our proposed MA-MGD framework demonstrates significant improvements in personalized environmental publicity, it still faces several limitations:

- **Simulation Constraints.** All experiments are conducted within a virtual environment based on AgentSociety(Piao et al., 2025). Although the simulation includes realistic geographic and demographic data, it cannot fully capture the unpredictability and complexity of real-world human behaviors, such as emotional fluctuations or spontaneous reactions.

- **Multi-channel Conflict in Rationality Scoring.** Although combining messages, posters, and announcements broadens dissemination, the rationality score often decreases due to independent scoring mechanisms for each medium. This suggests potential incompatibility among modalities in the current evaluation scheme.

- **Diminishing Returns from Iterative Refinement.** Our ablation study reveals that excessive iterations (e.g., more than two) in message optimization can lead to redundant content or deviations from user preferences, slightly reducing overall effectiveness. Better feedback refine approaches are expected.

References

Parasumanna Gokulan Balaji and Dipti Srinivasan. 2010. An introduction to multi-agent systems. In *Innovations in multi-agent systems and applications-1*, pages 1–27. Springer.

Noam Brown, Anton Bakhtin, Adam Lerer, and Qucheng Gong. 2020. Combining deep reinforcement learning and search for imperfect-information games. *Advances in neural information processing systems*, 33:17057–17069.

Jake Bruce, Niko Sünderhauf, Piotr Mirowski, Raia Hadsell, and Michael Milford. 2017. One-shot reinforcement learning for robot navigation with interactive replay. *arXiv preprint arXiv:1711.10137*.

Shengdong Du, Tianrui Li, Yan Yang, and Shi-Jinn Horng. 2019. Deep air quality forecasting using hybrid deep learning framework. *IEEE Transactions on Knowledge and Data Engineering*, 33(6):2412–2424.

Xiaohua Fu, Jie Jiang, Xie Wu, Lei Huang, Rui Han, Kun Li, Chang Liu, Kallol Roy, Jianyu Chen, Nesma Talaat Abbas Mahmoud, and 1 others. 2024. Deep learning in water protection of resources, environment, and ecology: Achievement and challenges. *Environmental Science and Pollution Research*, 31(10):14503–14536.

Shixiang Gu, Ethan Holly, Timothy Lillicrap, and Sergey Levine. 2017. Deep reinforcement learning for robotic manipulation with asynchronous off-policy updates. In *2017 IEEE international conference on robotics and automation (ICRA)*, pages 3389–3396. IEEE.

Jennifer Haase and Sebastian Pokutta. 2025. Beyond static responses: Multi-agent llm systems as a new paradigm for social science research. *arXiv preprint arXiv:2506.01839*.

Sirui Hong, Mingchen Zhuge, Jonathan Chen, Xiawu 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 *The Twelfth International Conference on Learning Representations*.

Rongjie Huang, Mingze Li, Dongchao Yang, Jiatong Shi, Xuankai Chang, Zhenhui Ye, Yuning Wu, Zhiqing Hong, Jiawei Huang, Jinglin Liu, Yi Ren, Yuexian Zou, Zhou Zhao, and Shinji Watanabe. 2024. Audiogpt: Understanding and generating speech, music, sound, and talking head. *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(21):23802–23804.

Yueping Kong, Peng Xue, Yuqian Xu, and Xiaolong Li. 2023. An environmental pattern recognition method for traditional chinese settlements using deep learning. *Applied Sciences*, 13(8).

Sergey Levine, Chelsea Finn, Trevor Darrell, and Pieter Abbeel. 2016. End-to-end training of deep visuomotor policies. *Journal of Machine Learning Research*, 17(39):1–40.

Guohao Li, Hasan Abed Al Kader Hammoud, Hani Itani, Dmitrii Khizbullin, and Bernard Ghanem. 2024. Camel: communicative agents for "mind" exploration of large language model society. In *Proceedings of the 37th International Conference on Neural Information Processing Systems*, NIPS '23, Red Hook, NY, USA. Curran Associates Inc.

Anna A Mendiola and Ma Regina M Hechanova. 2025. Framing messages toward pro-environmental behavior: A self-determination theory perspective. *Journal of Management for Global Sustainability*, 13(1):5.

Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, and 1 others. 2015. Human-level control through deep reinforcement learning. *nature*, 518(7540):529–533.

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.

Sudipta Paul, Amit Roy-Chowdhury, and Anoop Cherian. 2022. Avlen: Audio-visual-language embodied navigation in 3d environments. *Advances in Neural Information Processing Systems*, 35:6236–6249.

Jinghua Piao, Yuwei Yan, Jun Zhang, Nian Li, Junbo Yan, Xiaochong Lan, Zhihong Lu, Zhiheng Zheng, Jing Yi Wang, Di Zhou, and 1 others. 2025. Agentsociety: Large-scale simulation of llm-driven generative agents advances understanding of human behaviors and society. *arXiv preprint arXiv:2502.08691*.

Zhongang Qi, Tianchun Wang, Guojie Song, Weisong Hu, Xi Li, and Zhongfei Zhang. 2018. Deep air learning: Interpolation, prediction, and feature analysis of fine-grained air quality. *IEEE Transactions on Knowledge and Data Engineering*, 30(12):2285–2297.

Chen Qian, Wei Liu, Hongzhang Liu, Nuo Chen, Yufan Dang, Jiahao Li, Cheng Yang, Weize Chen, Yusheng Su, Xin Cong, Juyuan Xu, Dahai Li, Zhiyuan Liu, and Maosong Sun. 2024. ChatDev: Communicative agents for software development. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15174–15186, Bangkok, Thailand. Association for Computational Linguistics.

Kun Shao, Zhentao Tang, Yuanheng Zhu, Nannan Li, and Dongbin Zhao. 2019. A survey of deep reinforcement learning in video games. *arXiv preprint arXiv:1912.10944*.

Peter Stone and Manuela Veloso. 2000. Multiagent systems: A survey from a machine learning perspective. *Autonomous Robots*, 8(3):345–383.

Chen Tang, Ben Abbatematteo, Jiaheng Hu, Rohan Chandra, Roberto Martín-Martín, and Peter Stone. 2025. Deep reinforcement learning for robotics: A survey of real-world successes. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 28694–28698.

An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, Guanting Dong, Haoran Wei, Huan Lin, Jialong Tang, Jialin Wang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Ma, and 43 others. 2024. Qwen2 technical report. *Preprint*, arXiv:2407.10671.

Hui Yang, Sifu Yue, and Yunzhong He. 2023a. Auto-gpt for online decision making: Benchmarks and additional opinions. *Preprint*, arXiv:2306.02224.

Hui Yang, Sifu Yue, and Yunzhong He. 2023b. Auto-gpt for online decision making: Benchmarks and additional opinions, 2023. *URL https://arxiv.org/abs/2306.02224*, 3.

Yingxuan Yang, Huacan Chai, Shuai Shao, Yuanyi Song, Siyuan Qi, Renting Rui, and Weinan Zhang. 2025. Agentnet: Decentralized evolutionary coordination for llm-based multi-agent systems. *arXiv preprint arXiv:2504.00587*.

Rui Ye, Keduan Huang, Qimin Wu, Yuzhu Cai, Tian Jin, Xianghe Pang, Xiangrui Liu, Jiaqi Su, Chen Qian, Bohan Tang, and 1 others. 2025. Maslab: A unified and comprehensive codebase for llm-based multi-agent systems. *arXiv preprint arXiv:2505.16988*.

Junwei Yu, Yepeng Ding, and Hiroyuki Sato. 2025. Dyntaskmas: A dynamic task graph-driven framework for asynchronous and parallel llm-based multi-agent systems. *arXiv preprint arXiv:2503.07675*.

Yizhe Zhang, Siqi Sun, Michel Galley, Yen-Chun Chen, Chris Brockett, Xiang Gao, Jianfeng Gao, Jingjing Liu, and Bill Dolan. 2019. Dialogpt: Large-scale generative pre-training for conversational response generation. *arXiv preprint arXiv:1911.00536*.

Shuyan Zhou, Frank F. Xu, Hao Zhu, Xuhui Zhou, Robert Lo, Abishek Sridhar, Xianyi Cheng, Tianyue Ou, Yonatan Bisk, Daniel Fried, Uri Alon, and Graham Neubig. 2024. Webarena: A realistic web environment for building autonomous agents. In *The Twelfth International Conference on Learning Representations*.

Xizhou Zhu, Yuntao Chen, Hao Tian, Chenxin Tao, Weijie Su, Chenyu Yang, Gao Huang, Bin Li, Lewei Lu, Xiaogang Wang, and 1 others. 2023. Ghost in the minecraft: Generally capable agents for open-world environments via large language models with text-based knowledge and memory. *arXiv preprint arXiv:2305.17144*.

A Appendix A: Detailed Information on Simulation Scenarios

A.1 Introduction

All experiments in this study are conducted on the AgentSociety(Piao et al., 2025) platform. With the escalating global environmental crisis, promoting low-carbon lifestyles has become an essential mission for urban sustainability. Cities, as the epicenter of human activities, play a pivotal role in encouraging eco-friendly behavioral shifts. This mission – challenges participants to design an intelligent **Environmental Promotion Agent** that interacts with simulated urban citizens, with the aim of raising awareness of low-carbon habits and guiding behavioral changes through strategic communication and public messaging.

A.2 Simulation Overview

City and Residents.

- **City**: A simulated environment based on real geographic data of Beijing, containing multiple Areas of Interest (AOIs) such as residential zones, commercial districts, office spaces, and transportation hubs.
- **Citizens**: 200 virtual residents, each characterized by attributes such as age, gender, education, profession, marital status, residence and

workplace (linked to AOI), as well as unique background stories and habitual patterns.

Simulation Timeline.

1. **Day 1 – Promotion Stage**: The Environmental Promotion Agent performs personalized communication, poster placements, and city-wide announcements to promote eco-conscious values.

2. **Day 2 – Carbon Emission Stage**: Citizens resume daily routines. Their choices regarding mobility (walking, public transport, or private cars) are monitored to compute carbon emissions and reflect the agent's influence.

Available Toolkits for the Agent.

- **LLM Tool** (`self.llm`): Send prompt to a language model (e.g., Qwen, OpenAI) to generate responses or analyze profiles.

- **Sensing Tool** (`self.sense`): Retrieve environment/citizen info such as current time, citizen profiles, AOI metadata, and communication history.

- **Communication Tool** (`self.communication`): Send personalized messages to individual citizens (max 5 per round).

- **Poster Tool** (`self.poster`): Place promotional posters in AOIs (cost: 3000 units per poster).

- **Announcement Tool** (`self.announcement`): Broadcast messages city-wide (cost: 20000 units per announcement).

B Appendix B

As shown in Table 7.

C Appendix C: Workflow of the Environmental Promotion Agent

Explain the workflow of MA-MGD in conjunction with code functions.

1. Initialization. At startup, the Promotion Agent initializes internal state trackers:

- `self.talled`: citizens who have already been contacted.

- `self.posted`: index for areas where posters have already been displayed.

- `self.learned`: citizens whose environmental awareness has been successfully improved.

2. Education and Publicity Loop (`forward` method).

Executed every simulation round, this is the main behavioral logic.

1. Retrieve full citizen profiles via `getCitizenProfile()`.

2. Remove those already in `talled` or `learned`.

3. Randomly select 5 citizens from the remaining pool.

4. For each selected citizen:

   - Construct a personal profile including occupation, age, marital status, education, commuting distance, and background story.
   - Use a language model (LLM) to generate multi-turn persuasive conversation promoting behavior like "walk 3KM or bike 6KM even if you own a car", "use fan and curtain instead of AC", "sort garbage strictly".
   - If the citizen's attitude (obtained via `getAttitude`) includes desirable keywords (e.g., "walk", "energy saving"), mark as `learned`.
   - Deliver the conversation response via `sendMessage()` and record in `talled`.

3. Survey-Based Attitude Assessment (`sim_survey` method) If no citizens have yet been "learned", a baseline survey is triggered:

- A 10-question multiple choice survey is generated to assess daily environmental behaviors.

- Answers are simulated by the LLM based on citizens' profile and background.

- Each answer is scored:

  – `A`: 100, `B`: 75, `C`: 50, `D`: 25

- Average score $\geq$ 87.5 or specific patterns (e.g., $\geq$ 85 and Q1 $\neq$ "D") mark a citizen as "learned".

4. Poster Generation and Evaluation Optionally, the promotion agent can create posters targeting specific regions (AOIs):

Table 7: Typology of Resident Motivation and Tailored Communication Strategies

Persona	Core Motivation	Behavioral Traits and Keywords	Tailored Communication Strategy
Policy Followers	Rule Compliance	Not strongly motivated by environmental protection but willing to follow official rules. Keywords: "policy requirements", "response to directives", "compliance-driven".	**Authority Strategy**: Emphasize regulatory norms and institutional legitimacy by citing official notices, community regulations, and government guidance.
Interest-Driven	Economic Benefit	Participation based on expected economic return such as subsidies, discounts, rewards. Keywords: "sensitive to incentives", "worth doing", "cost-effective".	**Incentive Strategy**: Quantify potential savings or economic benefits. e.g., "You can save XX yuan on gas if you take the bus today", "Earn points and redeem gifts by participating in recycling".
Social-Driven	Social Recognition	Behavior influenced by social circles, likely to follow the crowd. Keywords: "peer pressure", "neighborhood trend", "everyone's doing it".	**Social Strategy**: Use group psychology. e.g., "XX neighbor joined our recycling challenge!", "Over 50% of households in our community have started waste sorting."
Wavering Pragmatists	Convenience and Cost	Support environmental protection but actions depend on convenience and effort. Keywords: "easy and quick", "too troublesome", "inconvenient", "hesitant".	**Convenience Strategy**: Emphasize simplicity and ease of implementation. e.g., "The bin is just downstairs", "Smart bins make sorting effortless", "Free ride-sharing parking spots available."
Value-Driven Vanguards	Intrinsic Values	Driven by personal conviction and ethical commitment, they are proactive environmentalists. Keywords: "environmental responsibility", "personal belief", "value-driven".	**Resonance Strategy**: Appeal to higher-order values and identity. e.g., "Your actions inspire others", "Stand as an example for future generations", "Promote shared environmental ideals".

1. Select a new region using `select_aoi()`.
2. Generate a poster via LLM based on the aggregated citizen profiles in that area.
3. Evaluate the poster's `credibility` and `reasonableness` using another LLM prompt.
4. If evaluation passes threshold (e.g., credibility > 85 and reasonableness > 90), post it using `putUpPoster()`.

5. Citizen-Initiated Message Handling

- Upon receiving a message from a citizen, invoke `communication_response()`.
- A LLM-generated reply is formulated with a persuasive tone to encourage behavior change.
- The reply is sent back to the citizen via `sendMessage()`.

D Appendix D: Detailed Information on Evaluation metrics

To comprehensively assess the effectiveness of the Environmental Promotion Agent in guiding citizen awareness and promoting low-carbon behaviors, the evaluation framework consists of three major components. The final score is computed as:

$$\begin{aligned}\textbf{Final Score} = \textbf{Awareness Score} \\ + \textbf{Carbon Reduction Score} \\ \times \textbf{Content Rationality} \end{aligned} \quad (2)$$

Awareness Score (0–50 points)

This metric evaluates the agent's ability to influence the environmental awareness of urban citizens. The evaluation is based on post-simulation standardized surveys conducted among all residents.

Evaluation Aspects:

- **Transportation Preferences**: Willingness to adopt greener travel modes such as walking and public transit.
- **Energy-Saving Habits**: Behavioral changes in everyday energy usage.
- **Eco-Friendly Consumption**: Consideration of environmental impact during shopping decisions.

Scoring Method: Survey responses are quantified on a standardized scale and averaged across all citizens.

Score Range: 0–50 points

Carbon Reduction Score (0–50 points).

This score measures the agent's influence on actual behavior change in terms of carbon emissions, primarily through citizens' transportation choices during the simulation.

Carbon Emission Coefficients:

- Walking: 0 kg CO_2/km
- Public Transport: 0.016 kg CO_2/km
- Private Vehicle: 0.040 kg CO_2/km

Scoring Formula:

$$\begin{aligned} \text{Baseline Emission} &= \text{Total Distance} \times 0.040, \\ \text{Actual Emission} &= \sum_{\text{mode}} (\text{Distance} \times \text{Emission Coefficient}), \\ \text{Carbon Score} &= \left(\frac{\text{Baseline} - \text{Actual}}{\text{Baseline}} \right) \times 100. \end{aligned} \tag{3}$$

Score Range: 0–50 points

Content Rationality (0–1).

This multiplier evaluates the credibility and appropriateness of all promotional content generated by the agent, including personalized messages, posters, and announcements. It is assessed in real-time by an AI-based evaluator.

Evaluation Dimensions:

- **Credibility**: Is the content fact-based and scientifically grounded?
- **Reasonableness**: Is the messaging appropriate without exaggeration or offensiveness?

Weighted Aggregation:

$$\text{Content Rationality} = \frac{0.1 \times S_{\text{msg}} + 0.3 \times S_{\text{poster}} + 0.6 \times S_{\text{announce}}}{100} \tag{4}$$

Where:

- S_{msg}: Mean credibility and reasonableness score of messages
- S_{poster}: Mean score for posters
- S_{announce}: Mean score for announcements

Weight Justification:

- **Announcement (60%)**: Broadest coverage and greatest impact
- **Poster (30%)**: Regional influence
- **Message (10%)**: Personalized but limited reach

Score Range: 0–1 (applied as a multiplier)

Example Calculation. If an agent achieves:

- Awareness Score = 45
- Carbon Score = 40
- Content Rationality = 0.9

Then the final score is computed as:

$$\textbf{Final Score} = (45 + 40) \times 0.9 = 76.5$$