

MobileCity: An Efficient Framework for Large-Scale Urban Behavior Simulation

Xiaotong Ye^{1*} Nicolas Bougie^{1*} Toshihiko Yamasaki² Narimasa Watanabe¹

¹Woven by Toyota

²The University of Tokyo

{tony.yip, nicolas.bougie, narimasa.watanabe}@woven.toyota
yamasaki@cvm.t.u-tokyo.ac.jp

Abstract

Generative agents offer promising capabilities for simulating realistic urban behaviors. However, existing methods often rely on static profiles, oversimplified behavioral logic, and synchronous inference pipelines that hinder scalability. We present **MobileCity**, a lightweight generative-agent framework for city-scale simulation powered by cognitively-grounded generative agents. Each agent acts based on its needs, habits, and obligations, evolving over time. Agents are initialized from survey-based demographic data and navigate a realistic multimodal transportation network spanning multiple types of vehicles. To achieve scalability, we introduce asynchronous batched LLM inference during action selection and a low-token communication mechanism. Experiments with 4,000 agents demonstrate that MobileCity generates more human-like urban dynamics than baselines while maintaining high computational efficiency. Our code is publicly available at <https://github.com/Tony-Yip/MobileCity>.

1 Introduction

Generative agents (Park et al., 2023), powered by Large Language Models (LLMs) (Madaan et al., 2022), have emerged as a transformative paradigm for simulating human-like behaviors across domains including recommender systems (Zhang et al., 2024), peer review (Bougie and Watanabe, 2024), medical Q&A (Li et al., 2024), and game simulation (Kim and Kim, 2023; Hu et al., 2024). Urban simulation models behaviors and transportation within a city, enabling evaluation of policies, transportation changes, and infrastructure planning. It supports forecasting market demand, analyzing traffic and safety impacts, and assessing public health and community well-being.

Despite recent progress (Park et al., 2023; Wang et al., 2023), existing generative-agent frameworks

exhibit notable limitations for large-scale urban mobility simulation. Most systems do not explicitly model human needs, temporal routines, or obligation-driven behaviors, leading to repetitive or unrealistic activity patterns (Feng et al., 2024; Samuel et al., 2024; Bougie and Watanabe, 2025). Besides, prior work typically assumes a single or overly simplified transportation system and fail to incorporate environmental factors such as weather or temperature, limiting the realism of mobility decisions. Finally, synchronous LLM calls and multi-turn dialogues incur substantial token and computation costs, making them prohibitively expensive to run at scale (Kaiya et al., 2023).

In light of this, we introduce **MobileCity**, a scalable generative-agent simulator built on a tile-based city representation. Agents are initialized with survey-based demographic profiles and evolve according to dynamic internal states through three modules: a *needs*, a *habits*, and an *obligations* module governing compulsory tasks. MobileCity further incorporates a multi-modal transportation system with three mobility options and integrates environmental factors such as weather, temperature, and venue availability, enabling context-aware decisions. Finally, in order to ensure scalability, we employ asynchronous batched LLM calls for action selection, streamline communication by exchanging memory indices instead of generating dialogues, and record only event-level state changes in an OpenSearch backend. Experiments with 4,000 agents demonstrate that MobileCity achieves higher behavioral realism and significantly better simulation efficiency compared to existing baselines. Beyond improving fidelity, we also showcase practical applications in mobility pattern forecasting and demographic analytics, illustrating MobileCity’s utility for urban planning and computational social science.

Our main contributions are:

*Equal contribution.

- **Cognitively-grounded, survey-conditioned urban agents.** We propose MobileCity, where each agent’s behavior is jointly driven by needs, habits, and obligations, and initialized from survey-based demographic and behavioral profiles, enabling diverse and temporally realistic daily routines.
- **Multi-modal mobility and context-aware decision making.** We incorporate a realistic transportation system with multiple modes and integrate environmental factors such as weather, temperature, and venue availability to support context-aware mobility and activity choices.
- **A scalable, low-token simulation pipeline for thousands of agents.** We achieve efficient city-scale simulation via (i) constrained, multiple-choice action selection to reduce token usage, (ii) lightweight communication by exchanging memory indices instead of generating full dialogues, and (iii) asynchronous batched LLM inference with event-level logging for scalable execution.

2 Related Work

Recent studies on generative agents have demonstrated significant progress in simulating human behavior. [Park et al. \(2023\)](#) introduce the first framework in which agents maintain memories and engage in social interactions. Building upon this foundation, [Wang et al. \(2023\)](#) incorporate basic needs to make daily activities more realistic, while [Chen et al. \(2024\)](#) design customizable environments that support emergent collaborative behaviors. To broaden applicability, [Zhou et al. \(2023\)](#) present an open-source system for autonomous language agents, and [Hong et al. \(2024\)](#) demonstrate how agents can collaborate in complex software engineering workflows.

As research moves toward larger-scale simulations, computational efficiency becomes a central concern. [Park et al. \(2024\)](#) scale simulations to 1,000 agents through a hierarchical decision-making architecture, although the proposed architecture still incurs prohibitive inference costs. [Yu et al. \(2024\)](#) reduce unnecessary LLM calls by learning simplified policies, yet real-time simulations remain constrained by the latency of LLM responses, especially when generating multi-turn dialogues.

Despite these advancements, existing systems typically overlook several factors essential for realistic urban mobility: diverse transportation modes, environmental conditions such as weather or temperature, and long-term behavioral traits influenced by needs, habits, and obligations. Moreover, prior work ([Bougie and Watanabe, 2025](#)) usually relies on token-intensive content generation. As a result, generating human-like behaviors with low inference cost and high scalability remains an open challenge.

3 Agent Modules

3.1 Agent Profile

We derive personas from questionnaire surveys completed by human participants, enabling the simulation to capture diverse demographic and psychological characteristics. Each agent is initialized with the following attributes:

- **Demographic Information** includes gender, age, job category, education level, financial status, family status like marriage.
- **Human Parameters** ([Barrick and Mount, 1991](#)) describe long-term behavioral tendencies. They include the Big Five personality traits and behavioral traits.
- **Hobbies** are initialized from SNS data, like X Posts, and dynamically updated based on agents’ activity records during simulation.

3.2 Individual Action Module

Human decisions arise from three psychological mechanisms that drive human action ([Wood et al., 2022](#)): *needs* ("I want to do"), *habits* ("I do it as usual"), and *obligations* ("I have to do"). We formalize them into three separate modules.

3.2.1 Needs-driven Action

Agents have spontaneous tendencies to maintain physiological or social equilibrium, consistent with theories of homeostasis ([Cannon, 1932](#)), and interpersonal balance ([Festinger, 1957](#); [Heider, 1958](#)). Namely, when an agent’s internal state deviates from its optimal level, it seeks to restore or enhance that state. We introduce eight agent needs, following [Maslow \(1943\)](#)’s hierarchy of needs in Table 1.

Each agent maintains a vector of need levels $C_N \in [0, 1]^8$, which decays over time following $C_N(t + \Delta t) = \text{clip}(C_N(t) - \Delta t \cdot D_N, 0, 1)$,

Maslow’s Hierarchy	Agent Needs
Physiological	Fullness, Energy
Safety	Health, Financial Security
Love/Belonging	Pleasure, Social Connection
Esteem	Status Recognition
Self-Actualization	Self-Growth

Table 1: Maslow’s hierarchy of needs

where D_N represents the individual decay rate vector. In contrast to prior work (Bougie and Watanabe, 2025; Yan et al., 2024), decay rates are heterogeneous across personas and need types. Lower-level physiological needs decay faster, while higher-level needs are more stable. For example, residents living alone experience quicker decline in *Social Connection* due to increased susceptibility to loneliness. In addition, we maintain an importance vector I_N , which encodes how strongly an agent prioritizes each need. For instance, agents with lower income place higher importance on financial security.

During action selection, the needs-driven score of an action at time t is defined as $N(t) = N_{hp} N_{imp}$. N_{hp} represents the weighted cosine similarity score between the agent’s human parameters x_{hp} and the action’s feature vector x_{act} . N_{imp} measures the importance-weighted fulfillment of unsatisfied needs defined by $1 - C_N$ and I_N .

3.2.2 Habit-driven Action

Habit-driven actions are triggered by temporal or spatial regularities reinforced through repeated actions. To reproduce such patterns, we define a habitual action preference function.

Suppose that the agent performed an action in the past, with the midpoint time of t_m , the amplitude, defined by action feedback, is A_H . We aim to determine the habit strength at the current time t . To model the daily cycle on a continuous circular domain, we normalize the minute-based time difference onto the interval $[-\pi, \pi]$ using $\Delta\theta(t) = \frac{2\pi}{1440} ((t - t_m) \bmod 1440)$, where $\Delta\theta(t)$ is the normalized angular distance. The habit intensity as a function of current time is modeled as a Gaussian distribution on the circle: $H(t) = R(t)A_H \exp(-k_H \Delta\theta(t)^2)$, where k_H controls the sharpness of the temporal peak, which is defined by the angular half-width of action execution time a_H , and A_H is defined by k_H to maintain a constant area. $R(t)$ represents the forgetting strength in the Ebbinghaus (Rubin and

Wenzel, 1996) forgetting curve model. As time passes, the habit strength will gradually decrease by $R(t) = \exp(-r_H(t - t_m))$. Habits are removed entirely once their strength falls below a minimal threshold.

3.2.3 Obligation-driven Action

Obligation-driven action refers to behaviors selected not from internal needs or habits but from externally imposed duties (Gershuny, 2003). In our framework, these mandatory tasks are encoded as core time slots in each agent’s calendar, derived from our questionnaire survey. They reflect factors such as sleep schedules, family structure (e.g., cohabitation, marital status, children’s ages), and historical activity logs.

During action selection, candidate needs-driven and habit-driven actions are first filtered by an availability mask determined by the next mandatory task. An action is admissible only if: a) it is semantically appropriate for the current time (e.g., “eat breakfast” is invalid at night), b) its venue is open during the intended period, and c) the agent can complete it, including travel time, and still arrive at the upcoming mandatory task on schedule.

3.3 Mobility Selection Module

When the locations of an agent’s consecutive actions differ, the agent must choose an appropriate mode of transportation. We implement a transportation system within the virtual town, comprising three transportation modes: walking, PMV (personal mobility vehicle), and bus. During action selection, the LLM is instructed to select an action from a list of multi-mechanism-driven actions and the most appropriate transportation mode, conditioning on the agent’s persona and environmental information including weather, temperature, and spatial context.

4 Towards Scalable Simulation

One of the primary goals of our system is to enable efficient simulation of large-scale agent populations. To this end, we introduce three strategies to improve efficiency.

4.1 Reducing Token Consumption

We first reduce token usage in the individual action module by shifting the LLM’s role from free-text generation to discrete selection. Specifically, the **action selector** precomputes a list of feasible candidate actions Act_{needs} , Act_{habit} , Act_{obl} , with the

mechanisms described in Section 3.2, and mobility options *walking*, *PMV*, *bus* in Section 3.3. The LLM is then prompted with a multiple-choice question containing these candidates, and its output is restricted to the index of the chosen option. An example is illustrated in Appendix B.4.

Instead of generating full dialogues, agents exchange information through a lightweight memory-transfer mechanism. The LLM is prompted to select which memory entries are shared between agents and to update their mutual relationship scores. Formally, when agents i and j meet, the LLM outputs only: $(\Delta\mathcal{M}_i, \Delta\mathcal{M}_j, \Delta R_{ij}) = \text{LLM}_{\text{comm}}(\mathcal{M}_i, \mathcal{M}_j, \text{context}_{ij})$, where $\Delta\mathcal{M}$ represents the exchanged memory indices, and ΔR_{ij} updates the social affinity between agents.

4.2 Asynchronous Mechanism

A central component of our scalability strategy is asynchrony. Our city-scale agent simulator operates under an asynchronous scheduling mechanism. At the beginning of each simulated day, a list of all agents $\mathcal{A} = \{a_1, a_2, \dots, a_N\}$ is initialized, and the system maintains a set of independent local clocks $\mathcal{I} = \{\theta_1, \theta_2, \dots, \theta_N\}$. This design allows each agent to progress through its own timeline, rather than synchronizing with a global simulation step. The pseudo-code is shown below.

Asynchronous Action Batch Scheduling

1. Initialize agents A , clocks θ , batch $\mathcal{B}$, threshold B .
 2. For each $a_i \in A$:
 - (a) If mandatory task due $\rightarrow$ execute and advance θ_i .
 - (b) Else compute candidates from needs and habits, append to $\mathcal{B}$.
 3. If $|\mathcal{B}| = B$ or all awaiting $\rightarrow$ dispatch batch.
 4. Update (θ_i, C_N) for returned agents.
 5. Remove agents with $\theta_i \geq 24:00$. Repeat until $A = \emptyset$.
-

The same mechanism is applied to agent-to-agent communication. Throughout the simulation, pairs of agents (a_i, a_j) likely to converse are dynamically generated, or when agents proactively reaching out when their social need is high. Instead of invoking the language model for every pair immediately, the system collects communication tasks into a shared batch buffer. Once the batch reaches a predefined threshold, all pending conversations are processed asynchronously, exchanging memory indices and updating relationship scores in parallel:

Asynchronous Conversation Batch Scheduling

1. Initialize conversation batch $\mathcal{B}_{\text{conv}}$, threshold B_{conv} .
 2. Detect potential pairs (a_i, a_j) :
 - (a) Face-to-face if both share venue and time overlap.
 - (b) Virtual contact if agent a_i has high social need.
 3. Append $(a_i, a_j, \text{MEMORY}_i, \text{MEMORY}_j)$ to $\mathcal{B}_{\text{conv}}$.
 4. If $|\mathcal{B}_{\text{conv}}| = B_{\text{conv}} \rightarrow$ dispatch batch.
 5. LLM returns exchanged memories and relationship updates $(\Delta\mathcal{M}_i, \Delta\mathcal{M}_j, \Delta R_{ij})$.
 6. Update memories and relationship states.
-

4.3 Data Logging and Visualization

In previous simulation systems (Park et al., 2023; Wang et al., 2023), the state and location of all agents at every time step were saved into local JSON files, which were then repeatedly accessed by the frontend for visualization. This I/O-intensive pipeline introduced significant latency and storage overhead. To address this issue, we decouple the simulation backend from the frontend and record only essential state changes. Specifically, each agent’s need satisfaction vector C_N is logged only when an action is completed, and spatial coordinates are recorded only upon movement. All event-level logs are stored in an OpenSearch (OpenSearch Project, 2021) database instead of local files. After the simulation, missing agent states are linearly interpolated based on the temporal continuity of needs and locations, allowing the frontend to reconstruct smooth and continuous trajectories directly from OpenSearch queries.

5 Experimental Results

5.1 Runtime Analysis

Prior generative-agent systems suffer from severe runtime limitations due to heavy LLM invocation. Humanoid Agents (Wang et al., 2023) requires 40 minutes to simulate only 3 agents. AgentSociety (Gershuny, 2003) adopts cohort-based batching, yet inference for 1,000 agents partitioned into 8 groups still takes 11 minutes for a single global decision cycle. These baselines highlight the computational bottleneck of LLM-driven multi-agent simulations and motivate the need for a more efficient execution framework. We accelerate end-to-end simulation using three mechanisms as explained in Section 4. To quantify their effects, we conduct an ablation analysis.

We observe that weekday simulations consistently finish faster than weekend simulations. This is expected: employed agents spend a larger portion of weekday daytime in workplaces, resulting in




Figure 1: The crowd distribution across different urban venues, on weekdays (top) and weekends (bottom).

Population	R	R+D	R+A+D
40	194	115	39
400	2,093	1,329	383
4,000	22,432	15,234	3,734

(a) Weekday

Population	R	R+D	R+A+D
40	246	154	52
400	2,497	1,649	495
4,000	29,656	20,731	4,850

(b) Weekend

Table 2: Runtime (seconds) under different acceleration settings. R = Reducing Token Consumption; R+D = adding Data Logging; R+A+D = full system including the Asynchronous Mechanism.

fewer action selections and correspondingly fewer LLM calls for memory updates. In contrast, weekend schedules involve more frequent transitions across leisure venues, increasing the total number of model queries.

5.2 Human Likeness

A central question is how closely synthetic residents resemble real human behavior. To evaluate this, we present each agent’s daily actions to GPT-4o (King and ChatGPT, 2023) and ask it to judge whether the behavior appears human-like or machine-generated using a 5-point Likert scale. Higher scores indicate stronger alignment with natural human behavior. Table 3 reports the averaged scores across interactions, comparing our method with the baseline (Park et al., 2023), AGA (Yu

et al., 2024), and HumanoidAgent (Wang et al., 2023). Our approach achieves the highest human-likeness score by a notable margin, demonstrating that the integration of needs, temporal habits, and obligation-driven decision-making produces behaviors that GPT-4o reliably interprets as human. A qualitative example of generated daily interactions is provided in Appendix D.2.

Table 3: Human-likeness score evaluated by GPT-4o.

Method	Activity
Baseline	3.11 ± 0.18
AGA	3.22 ± 0.28
HumanoidAgent	3.30 ± 0.31
Ours	4.09 ± 0.27

5.3 Venue Heatmap

Understanding how crowds occupy urban spaces over time is crucial for urban planning and resource allocation. Figure 1 illustrates the temporal distribution of venue utilization generated by MobileCity. Between 22:00 and 06:00, most agents remain in residential rooms, reflecting natural nighttime resting patterns. During weekday mornings, employed agents concentrate in office areas, producing a pronounced surge in workplace occupancy. As work hours end, the population gradually shifts toward leisure-oriented venues such as sports centers, cultural spaces, and parks. In contrast, weekend patterns exhibit a more diverse distribution




Figure 2: Residents with different employment status have different fluctuations in agent needs during the day.

throughout the day, with consistent increases in visits to commercial, dining, and recreational locations. Overall, the simulated dynamics closely align with real-world urban mobility trends, where work schedules, leisure routines, and daily rhythms jointly shape venue occupancy.

5.4 Macro-Level Action Distribution




Figure 3: Percentage point differences in activity distribution between our method and real-world data across demographic categories.

While aligning individual agent behaviors with their human counterparts is crucial, it is also necessary that human proxies replicate real-world user behavior at a macro level. In each category, the three-digit code represents age, employment status, and income level. We compare the percentage distribution of activities between our method and real-world data. Figure 3 presents this comparison as a heatmap of percentage point differences. Categories are encoded as three-digit IDs XYZ , where $X \in \{1, 2, 3\}$ denotes age group (1: 25–44, 2: 45–

64, 3: 65–84), $Y \in \{0, 1, 2\}$ denotes employment status (0: unemployed, 1: employed, 2: part-time), and $Z \in \{0, 1\}$ denotes income level (0: medium, 1: high). Unemployed agents have more time to perform those actions than employed agents, since employees have to work in the office on weekdays. We also noticed that employed adults show higher exercise engagement in our simulation, while older demographics exhibit shifted time allocation preferences. The observed differences provide valuable insights into demographic-specific behavioral tendencies that can inform future social studies while demonstrating our method’s capability to replicate complex human behavioral patterns.

5.5 Emotion Monitoring

To analyze agents’ emotional states, we group agents by employment status (unemployed, part-time, employed) and compute the average values of their basic needs over five weekdays. We visualize the five basic needs whose values exhibit the most noticeable fluctuations. As shown in Figure 2, these attributes fluctuate least for unemployed agents, moderately for part-time workers, and most dramatically for employed agents. All attributes, except *Fullness*, follow a consistent pattern: a steady decline between 9:00 and 18:00, followed by recovery during non-working hours. This trend arises because employed agents are predominantly occupied with work during the day, which restricts them from engaging in replenishing activities. *Fullness*, however, rises at 8:00, 12:00, and 18:00, corresponding to mealtimes.

5.6 Transportation Statistics

We additionally evaluate transportation mode preferences across demographic categories, as summarized in Table 4. Overall, walking emerges as the

Table 4: Time percentage (%) spent by agents using different transportation modes. “Exp” represents experimental results from our simulation, while “GT” refers to ground truth values from our proprietary dataset.

Category	Walking		PMV		Bus	
	Exp	GT	Exp	GT	Exp	GT
100	89.96	88.78	0.00	0.57	10.04	10.65
110	93.99	92.74	0.00	0.43	6.01	6.83
120	94.39	93.19	0.00	0.29	5.61	6.52
210	96.84	95.59	0.87	1.54	2.29	2.87
211	95.44	93.92	0.00	0.53	4.56	5.55
220	95.41	94.36	0.00	0.56	4.59	5.08
300	95.43	94.21	0.87	1.41	3.70	4.38
310	92.43	90.98	3.23	3.99	4.34	5.04
320	94.55	93.45	0.00	0.38	5.45	6.17

dominant choice across all groups, reflecting its suitability for short-distance travel. PMV usage remains consistently low, which aligns with mobility patterns observed in our ground-truth dataset. Agents tend to rely on walking for nearby destinations and switch to public transit for longer routes, resulting in a natural bimodal split between these two modes. Environmental factors also contribute: PMV is rarely selected during rainy periods due to reduced safety and comfort. Across all categories, the experimental results closely track the ground-truth percentages, indicating that our agent-based mobility model successfully captures realistic travel preferences.

6 Conclusion

We presented **MobileCity**, a scalable framework for large-scale generative-agent simulation in dynamic urban environments. Our system integrates a realistic multi-modal transportation model and a unified agent architecture that jointly incorporates static human parameters, dynamic basic needs, temporal habits, and compulsory tasks. Through asynchronous batched action selection and lightweight communication based on memory exchange, **MobileCity** achieves human-like behavioral realism while remaining computationally efficient. The resulting simulations provide fine-grained insights into urban mobility patterns, offering a practical tool for improving traffic safety, infrastructure design, and urban planning while reducing reliance on costly real-world data collection.

7 Limitations

There are several limitations to our work. First, our simulation framework primarily focuses on model-

ing typical urban scenarios, while rare or extreme events, such as natural disasters, rapid population shifts, or sudden infrastructure failures, remain challenging to accurately capture. Second, the computational demands of large-scale, high-resolution urban simulations may become costly. Trade-offs in spatial granularity, temporal resolution, or agent complexity are necessary, which may limit the ability to represent micro-scale dynamics or long-term urban evolution. Besides, the behavior of agents may inherit biases present in the underlying data or model training. This includes reproducing social, cultural, or policy biases, as well as occasional generation of inconsistent or unfounded outputs. Finally, our work raises ethical and policy considerations. Automated urban simulations have the potential to influence real-world decision-making. It is therefore critical that users remain aware of the inherent biases and limitations of these systems.

8 Ethics Statement

This paper presents **MobileCity**, an LLM-powered agent framework designed to simulate large-scale urban mobility and social behaviors in a cost-effective and scalable manner. While our approach offers significant benefits for urban planning, traffic management, and behavioral modeling, it also raises several ethical considerations.

One primary concern is the potential for bias amplification. Since our agent behaviors are derived from survey data and LLM-generated actions, any biases inherent in these sources could propagate within the simulation. This may lead to an unrealistic or skewed representation of population behaviors, which, if used for policy-making or infrastructure design, could reinforce existing social or economic inequalities.

Another potential risk is the misuse of simulation insights. The ability to predict crowd density, individual behaviors, and mobility trends may be leveraged for unethical purposes, such as excessive surveillance, behavioral manipulation, or commercial exploitation without public consent. Safeguards should be in place to ensure that data-driven insights are used responsibly and in ways that benefit society.

To mitigate these risks, we advocate for the responsible deployment of our framework, emphasizing transparency, fairness, and the inclusion of human oversight when deriving actionable insights from the simulation. By adhering to these prin-

ciples, we can ensure that the use of generative agents in urban simulations remains ethically and socially beneficial.

References

- Murray R Barrick and Michael K Mount. 1991. The big five personality dimensions and job performance: a meta-analysis. *Personnel psychology*, 44(1):1–26.
- Nicolas Bougie and Narimasa Watanabe. 2024. Generative adversarial reviews: When llms become the critic. *arXiv preprint arXiv:2412.10415*.
- Nicolas Bougie and Narimasa Watanabe. 2025. CitySim: Modeling urban behaviors and city dynamics with large-scale LLM-driven agent simulation. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 215–229, Suzhou (China). Association for Computational Linguistics.
- Lars Böcker, Martin Dijst, and Jan Prillwitz. 2013. Impact of everyday weather on individual daily travel behaviours in perspective: A literature review. *Transport Reviews*, 33.
- Walter B. Cannon. 1932. *The Wisdom of the Body*. W. W. Norton & Company.
- Weize Chen, Yusheng Su, Jingwei Zuo, Cheng Yang, Chenfei Yuan, Chi-Min Chan, Heyang Yu, Yaxi Lu, Yi-Hsin Hung, Chen Qian, Yujia Qin, Xin Cong, Ruobing Xie, Zhiyuan Liu, Maosong Sun, and Jie Zhou. 2024. Agentverse: Facilitating multi-agent collaboration and exploring emergent behaviors. In *International Conference on Learning Representations (ICLR)*.
- Jie Feng, Jun Zhang, Junbo Yan, Xin Zhang, Tianjian Ouyang, Tianhui Liu, Yuwei Du, Siqi Guo, and Yong Li. 2024. Citybench: Evaluating the capabilities of large language model as world model. *arXiv preprint arXiv:2406.13945*.
- Leon Festinger. 1957. *A Theory of Cognitive Dissonance*. Stanford University Press.
- Jonathan Gershuny. 2003. *Changing times: Work and leisure in postindustrial society*. OUP Oxford.
- Fritz Heider. 1958. *The Psychology of Interpersonal Relations*. Wiley.
- 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 *International Conference on Learning Representations (ICLR)*.
- Sihao Hu, Tiansheng Huang, Fatih Ilhan, Selim F. Tekin, Gaowen Liu, Ramana Kompella, and Ling Liu. 2024. *A survey on large language model-based game agents*. *CoRR*, abs/2404.02039.
- Afshin Jafari, Dharendra Singh, Alan Both, Mahsa Abdollahyar, Lucy Gunn, Steve Pemberton, and Billie Giles-Corti. 2021. Activity-based and agent-based transport model of melbourne (atom): an open multi-modal transport simulation model for greater melbourne. *CoRR*, abs/2112.12071.
- Zhao Kaiya, Michelangelo Naim, Jovana Kondic, Manuel Cortes, Jiaxin Ge, Shuying Luo, Guangyu Robert Yang, and Andrew Ahn. 2023. Lyfe agents: Generative agents for low-cost real-time social interactions. *CoRR*, abs/2310.02172.
- Munyeong Kim and Sungsu Kim. 2023. *Generative AI in mafia-like game simulation*. *CoRR*, abs/2309.11672.
- Michael R King and ChatGPT. 2023. A conversation on artificial intelligence, chatbots, and plagiarism in higher education. *Cellular and molecular bioengineering*, 16(1):1–2.
- Junkai Li, Siyu Wang, Meng Zhang, Weitao Li, Yunghwei Lai, Xinhui Kang, Weizhi Ma, and Yang Liu. 2024. *Agent hospital: A simulacrum of hospital with evolvable medical agents*. *CoRR*, abs/2405.02957.
- Aman Madaan, Shuyan Zhou, Uri Alon, Yiming Yang, and Graham Neubig. 2022. Language models of code are few-shot commonsense learners. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1384–1403.
- Abraham Harold Maslow. 1943. A theory of human motivation. *Psychological review*.
- OpenSearch Project. 2021. Opensearch. <https://github.com/opensearch-project/OpenSearch>. GitHub repository.
- Joon Sung Park, Joseph C. O’Brien, Carrie Jun Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein. 2023. Generative agents: Interactive simulacra of human behavior. In *The 36th Annual Symposium on User Interface Software and Technology (UIST)*, pages 2:1–2:22.
- Joon Sung Park, Carolyn Q. Zou, Aaron Shaw, Benjamin Mako Hill, Carrie J. Cai, Meredith Ringel Morris, Robb Willer, Percy Liang, and Michael S. Bernstein. 2024. *Generative agent simulations of 1,000 people*. *CoRR*, abs/2411.10109.
- Amanda L Rebar, Ryan E Rhodes, and Benjamin Gardner. 2019. How we are misinterpreting physical activity intention–behavior relations and what to do about it. *International Journal of Behavioral Nutrition and Physical Activity*.
- David C Rubin and Amy E Wenzel. 1996. One hundred years of forgetting: A quantitative description of retention. *Psychological review*.

- Vinay Samuel, Henry Peng Zou, Yue Zhou, Shreyas Chaudhari, Ashwin Kalyan, Tanmay Rajpurohit, Ameet Deshpande, Karthik Narasimhan, and Vishvak Murahari. 2024. Personagym: Evaluating persona agents and llms. *arXiv preprint arXiv:2407.18416*.
- Daniel Tischner. 2018. [Multi-modal route planning in road and transit networks](#). *CoRR*, abs/1809.05481.
- Zhilin Wang, Yu-Ying Chiu, and Yu Cheung Chiu. 2023. Humanoid agents: Platform for simulating human-like generative agents. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Mark Wardman and Jeremy Toner. 2020. Is generalised cost justified in travel demand analysis? *Transportation*, pages 75–108.
- Wendy Wood, Asaf Mazar, and David T Neal. 2022. Habits and goals in human behavior: Separate but interacting systems. *Perspectives on Psychological Science*.
- Yuwei Yan, Qingbin Zeng, Zhiheng Zheng, Jingzhe Yuan, Jie Feng, Jun Zhang, Fengli Xu, and Yong Li. 2024. Opencity: A scalable platform to simulate urban activities with massive llm agents. *arXiv preprint arXiv:2410.21286*.
- Yangbin Yu, Qin Zhang, Junyou Li, Qiang Fu, and Deheng Ye. 2024. [Affordable generative agents](#). *CoRR*, abs/2402.02053.
- An Zhang, Yuxin Chen, Leheng Sheng, Xiang Wang, and Tat-Seng Chua. 2024. On generative agents in recommendation. In *Conference on Research and Development in Information Retrieval (SIGIR)*, pages 1807–1817.
- Wangchunshu Zhou, Yuchen Eleanor Jiang, Long Li, Jialong Wu, Tiannan Wang, Shi Qiu, Jintian Zhang, Jing Chen, Ruipu Wu, Shuai Wang, Shiding Zhu, Jiyu Chen, Wentao Zhang, Ningyu Zhang, Huajun Chen, Peng Cui, and Mrinmaya Sachan. 2023. [Agents: An open-source framework for autonomous language agents](#). *CoRR*, abs/2309.07870.



Figure 4: The map of our simulated city.

A City Simulator

In our simulation, agents operate in discrete time steps of 15 seconds, during which they perform actions, move between venues, and perceive their environment. To enhance scalability and efficiency, we introduce three key improvements over traditional simulators (Park et al., 2023; Yu et al., 2024). First, to accommodate large-scale agent populations our simulator allows stateless simulation, significantly reducing computational overhead. Second, our platform includes diverse buildings and venues, enabling a more comprehensive representation of urban environments. Third, we design a transportation system with multiple mobility hubs, offering agents diverse route options and facilitating citywide movement patterns.

A.1 City Map

Our simulated city is constructed using a tile-based map representation, as shown in Figure 4. The city map follows a grid-based structure where each tile represents 25 meters, and each block spans 500 meters. It features a diverse range of urban infrastructures, including 18 buildings, and 68 venues (Figure 5). The city features 8 apartment complexes, 2

company offices, 5 parks, 1 hospital, 1 department store, and 1 stadium. Each building contains different spaces. For example, an apartment building contains several living spaces for agents, 1 restaurant, and 1 convenience store or 1 entertainment venue. A company building contains several offices, 1 company canteen, and 1 convenience store.




Figure 5: The buildings and venues in the simulated city.




Figure 6: An agent traveling from apartment to company has three route options: (1) walking through zebra crossings (yellow lines), (2) walking to PMV (Personal Mobility Vehicle) station, riding on highway, then walking to destination (blue lines), or (3) walking to bus station, taking bus, then walking to destination (green lines).

A.2 Transportation System

The diversity of transportation modes facilitates the investigation of mobility patterns among urban residents (Jafari et al., 2021). In our simulation, agents move using three transportation modes: Walking, PMV, and Bus, where PMV refers to a personal mobility vehicle. We introduce a constrained navigation system that dynamically determines optimal routes based on cost, constraints, and individual preferences. Inspired by real-world systems (Wardman and Toner, 2020), our navigation system generates multiple route options, each differing in time cost and monetary cost. In general, bus routes have the lowest time cost but the highest monetary cost, whereas walking routes are the opposite. To formalize this, we construct three graphs (Tischner, 2018): a walking graph G_w , a PMV graph G_p , and a bus graph G_b in our map. Each graph is constructed with nodes representing accessible points for agents, and edges representing different moving costs.

A walking graph consists of passages inside buildings, which are yellow areas in Figure 7, and zebra crossings between buildings in Figure 8. In one building, agents can access most of the areas except for collision walls. Between buildings, agents can walk across zebra crossings on highways. An agent moves 1 tile in each time step by walking.




Figure 7: Walking-accessible zone.

A PMV graph consists of nodes of PMV stations, represented as blue tiles in Figure 8. To ride a PMV, agents must walk to the PMV station first, then ride the PMV on the left side of the highway. An agent moves 2 tiles in each time step when using a PMV.




Figure 8: A PMV route example.

A bus graph consists of nodes of bus stations,

represented as green tiles in Figure 9. To get on a bus, agents must walk to the bus station first, and then the bus will move on the left side of the highway. An agent moves 5 tiles in each time step when using a bus.




Figure 9: Two bus routes in our city.

At each simulation time step, agents traverse 1, 2, or 5 tiles depending on whether they are walking, using a PMV, or taking a bus, respectively. Consequently, the time required to travel across a full block is 300, 150, and 60 seconds for walking, PMV, and bus travel, respectively.

Their respective time costs t_w , t_p , and t_b are calculated as:

$$\begin{aligned} t_w &= \min_{\pi \in G_w} \text{dist}(s \rightarrow t), \\ t_p &= \min_{\pi \in G_w \cup G_p} \text{dist}(s \rightarrow t), \\ t_b &= \min_{\pi \in G_w \cup G_b} \text{dist}(s \rightarrow t), \end{aligned}$$

where s is the starting place, t is the terminal place, and π represents all the paths in graphs. Route selection is constrained by its upcoming compulsory tasks and influenced by agent group characteristics. For instance, if an agent must reach the office within 15 minutes, it prioritizes the bus to minimize travel time and avoid tardiness. Higher-income agents are more likely to choose bus due to its time efficiency, whereas lower-income agents may opt for walking to reduce costs. Additionally, weather conditions (Böcker et al., 2013) play a crucial role in mobility decisions, on rainy days, agents tend to avoid PMVs due to safety and comfort concerns.

B Individual Action Module

We now explain the details in the action module.

B.1 Needs-driven Action

The total needs-based score at time t is defined as:

$$N(t) = N_{\text{hp}} N_{\text{imp}} \quad (1)$$

Let $x_{\text{hp}} \in \mathbb{R}^D$ denote the agent’s human parameter (HP) vector, and $x_{\text{act}} \in \mathbb{R}^D$ the action’s HP vector, with weights $w \in \mathbb{R}_{\geq 0}^D$. The weighted cosine similarity is given by:

$$N_{\text{hp}} = \frac{1}{2} \left(1 + \frac{\langle w \odot x_{\text{hp}}, w \odot x_{\text{act}} \rangle}{\|w \odot x_{\text{hp}}\|_2 \|w \odot x_{\text{act}}\|_2} \right), \quad (2)$$

which maps the similarity to the range $[0, 1]$.

Let:

- $I_N \in \mathbb{R}^8$: the agent’s importance weights for each of the 8 needs;
- $C_N \in [0, 1]^8$: the agent’s current need satisfaction levels (scaled to $[0, 1]$);
- $A_N \in \mathbb{R}^8$: the action’s positive contribution to each need.

Then the importance-based need score is defined element-wise as:

$$N_{\text{imp}} = \langle \text{softmax}(I_N) \odot (1 - C_N) \odot \tanh(k_{\text{tanh}} \text{ReLU}(A_N)) \rangle \quad (3)$$

where:

- $\text{softmax}(I_N)$ normalizes the importance of each need;
- $(1 - C_N)$ represents the current deficiency or gap in satisfaction;
- $\tanh(k_{\text{tanh}} \text{ReLU}(A_N))$ introduces diminishing returns on positive need fulfillment, ensuring saturation as contribution increases.

B.2 Habit-driven Action

The habit-based score at time t is defined as:

$$H(t) = R(t) A_H \exp(-k_H \Delta\theta(t)^2) \quad (4)$$

$R(t)$ will be neglected in the following discussion since it’s not related to Gaussian distribution.

Our rationale for modeling habit strength using a Gaussian distribution is as follows. First, habit strength is treated as a continuous variable rather than a binary one. Habit strength is the cumulative result of numerous minor factors in reality. In behavioral prediction and health psychology, research (Rebar et al., 2019) has found that variables such as intention, behavior, and frequency are “approximately normally” distributed. Therefore, according to the Central Limit Theorem, the aggregation of these influences will approximate a normal distribution.




Figure 10: The curve remains close to zero throughout the interval, demonstrating that the Padé [2/2] approximation tracks the exponential decay very accurately. The maximum relative error is among 0.1% and 1%.

Given this premise, we discuss the derivation of the formula for A_H and k_H .

Object 1: given a half-action duration a , the integral of the habit strength over the interval $[-a, a]$ must account for more than 90% of the total integral area.

Consider a normalized Gaussian kernel:

$$f(x) = e^{-k_H x^2}, \quad k_H > 0. \quad (5)$$

The total area under this curve is:

$$\int_{-\infty}^{\infty} e^{-k_H x^2} dx = \sqrt{\frac{\pi}{k_H}}. \quad (6)$$

When integrating over a finite range $[-a, a]$, the result becomes:

$$I(a) = \int_{-a}^a e^{-k_H x^2} dx = \sqrt{\frac{\pi}{k_H}} \operatorname{erf}(\sqrt{k_H} a). \quad (7)$$

The error function is defined as:

$$\operatorname{erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} dt. \quad (8)$$

By substituting $t = \sqrt{k_H} x$, the Gaussian integral over $[-a, a]$ introduces $\operatorname{erf}(\sqrt{k_H} a)$. The fraction of total area within $[-a, a]$ is therefore:

$$P(a) = \frac{\int_{-a}^a e^{-k_H x^2} dx}{\int_{-\infty}^{\infty} e^{-k_H x^2} dx} = \operatorname{erf}(\sqrt{k_H} a). \quad (9)$$

In our experimental setup, the execution duration of non-mandatory tasks ranges from 0.5 to 3 hours, and correspondingly $a \in [\pi/48, \pi/8]$. If we require that $[-a, a]$ contains 90% of the total area, we solve:

$$\operatorname{erf}(\sqrt{k_H} a) = 0.9. \quad (10)$$

This yields:

$$\sqrt{k_H} a = \operatorname{erf}^{-1}(0.9) \approx 1.163, \quad (11)$$

and consequently:

$$k_H \approx \left(\frac{1.163}{a} \right)^2. \quad (12)$$

Since the computation of the exponential function is computationally expensive, we perform a rational approximation.

Object 2: given k_H , the habit peak A_H must vary to ensure the integral remains constant.

$$I(k_H) = A_H \int_{-a\sqrt{k_H}}^{a\sqrt{k_H}} e^{-t^2} \frac{dt}{\sqrt{k_H}} = 0.9 * A_H \frac{\sqrt{\pi}}{\sqrt{k_H}} \quad (13)$$

To maintain a constant area S :

$$A_H(k_H) \approx 0.627 S \sqrt{k_H} \quad (14)$$

Since the computation of the exponential function is computationally expensive, we apply Padé approximation that preserves accuracy near $u = 0$.

$$e^{-u} \approx \frac{1 - \frac{u}{2} + \frac{u^2}{12}}{1 + \frac{u}{2} + \frac{u^2}{12}}. \quad (15)$$

By substituting $u = k_H x^2$, we obtain the rational approximation of $f(x)$:

$$f(x) \approx A_H \frac{1 - \frac{k_H x^2}{2} + \frac{k_H^2 x^4}{12}}{1 + \frac{k_H x^2}{2} + \frac{k_H^2 x^4}{12}}. \quad (16)$$

The maximum relative error is among 0.1% and 1%, as demonstrated in Fig. 10.

B.3 Obligatory-driven Action

$\text{Mask}(t)$ is True if and only if three conditions are satisfied,

$$\text{Mask}(t) = M_{\text{sem}}(t, a_{\text{act}}) M_{\text{open}}(t, a_{\text{act}}) M_{\text{obl}}(t, a_{\text{act}}) \quad (17)$$

Semantic-temporal consistency ensures that the action's semantics align with the current time period. An action labeled *eat breakfast* should be invalid in the evening. $M_{\text{sem}}(t, a_{\text{act}}) = 1$ requires

$$t \in T_{\text{act}}^{\text{sem}} \quad (18)$$

Venue availability ensures that the physical location associated with an action must be open during the planned execution interval. Let $\Delta t_{\text{cur}}(a_{\text{act}})$ the travel time from the current location to the next action location. $M_{\text{open}}(t, a_{\text{act}}) = 1$ requires

$$t + \Delta t_{\text{cur}}(a_{\text{act}}) \geq t_{\text{start}}^{\text{venue}}, \quad (19)$$

$$t + \Delta t_{\text{cur}}(a_{\text{act}}) + \Delta t_{\text{act}} \leq t_{\text{close}}^{\text{venue}}. \quad (20)$$

Obligation constraint ensures that the agent must complete all ongoing voluntary actions before the next scheduled mandatory task. Let

$t_{\text{next}}^{\text{obl}}$ denote the start time of the next obligation, and $\Delta t_{\text{next}}(a_{\text{act}})$ the travel time from the current location to the next mandatory task location. $M_{\text{obl}}(t, a_{\text{act}}) = 1$ requires

$$t + \Delta t_{\text{cur}}(a_{\text{act}}) + \Delta t_{\text{act}} + \Delta t_{\text{next}}(a_{\text{act}}) \leq t_{\text{next}}^{\text{obl}} \quad (21)$$

B.4 Action Selector

We now provide a comprehensive explanation of Agent Action Selector, detailing the implementation and technical details. This is a detailed example.

Action Selector

Now, it is 7:00 AM on Monday, and our agent Alex Kim wakes up. He will select an action by following these steps.

Step 1: Consider restraints from the next Obligatory-driven Action. Alex is a 25-year-old software engineer working. He needs to start working remotely or in the office from 9:00 on weekdays.

Step 2: List Needs-driven Actions. Alex needs to eat a lot to maintain energy for high-intensity work, which means he has a high demand for needs of "Fullness" and "Energy". He is very hungry, so his Top-5 needs-driven actions will be: *have breakfast in the canteen*, *grasp some food from the convenience store*, *drink coffee in the cafe*, *have decent breakfast at a nearby restaurant*, and *have breakfast at home*.

Step 3: List Habit-driven Actions. According to his personal habits, Alex's Top-3 actions at 7:00 are: *drink coffee in the cafe*, *walk in the park*, and *meditate at home*.

Step 4: Select an action and transportation mode with LLM. The current environmental condition is: sunny, 15°C. It's a good weather to go out, LLM makes the action choice for Alex: *drink coffee in the cafe*. Meanwhile, it takes 20 minutes to walk to the cafe, but only 8 minutes by PMV. LLM makes the transportation choice for Alex: *PMV*.

C Temporal Optimization

C.1 Asynchronous Actions

At every iteration, the simulator scans through the active agents. For each agent a_i , the system first checks whether the next event on its schedule is a mandatory task. If so, the agent executes that task immediately, updates its local time $\theta_i \leftarrow \theta_i + \Delta t_{\text{act}}$, and adjusts its need satisfaction vector $C_N \leftarrow \text{clip}(C_N + A_N, [0, 1])$, where A_N denotes the need-specific increments contributed by the action. Otherwise, the agent’s action selector compiles two sets of candidate actions, $\text{ACT}_{\text{needs}}$ from the needs-driven module and $\text{ACT}_{\text{habit}}$ from the habit-driven module, and merges them into a unified candidate set CANDS . Each candidate set, together with the agent’s persona, current environment view, and current need satisfaction vector C_N , is assembled into an LLM request. Instead of invoking the model immediately, the task is placed into a shared *batch buffer*. When the batch size reaches a predefined threshold B , all queued tasks are sent to the LLM simultaneously as a parallel API call. The results are then returned asynchronously, and each agent updates its state independently according to the selected action. After each execution, if the local time θ_i of an agent reaches 24:00, the agent is temporarily removed from the active list.

Algorithm 1 Asynchronous Batched Action Scheduling

```

1: Initialize agents  $\mathcal{A} = \{a_1, \dots, a_N\}$  and clocks
    $\mathcal{I} = \{\theta_1, \dots, \theta_N\}$ ;  $\mathcal{B} \leftarrow \emptyset$ 
2: while  $\mathcal{A} \neq \emptyset$  do
3:   for each  $a_i \in \mathcal{A}$  do
4:     if  $\theta_i \geq 24:00$  then remove  $a_i$ ; continue
5:     else if mandatory( $a_i$ ) then execute;
        $\theta_i \leftarrow \theta_i + \Delta t_{\text{act}}$ ;  $C_N \leftarrow C_N + A_N$ ; continue
6:     else  $\text{CANDS} \leftarrow \text{MERGETOPK}(\text{ACT}_{\text{needs}}, \text{ACT}_{\text{habit}})$ ;
7:     add  $(a_i, \text{CANDS}, C_N, \text{persona}, \text{env})$  to  $\mathcal{B}$ ;
       mark  $a_i$  as awaited
8:   end if
9: end for
10: if  $|\mathcal{B}| \geq B$  or all agents awaited then dispatch  $\mathcal{B}$  to LLM in parallel;
11: update  $\theta_i \leftarrow \theta_i + \Delta t_{\text{act}}$ ,  $C_N \leftarrow \text{clip}(C_N + A_N, [0, 1])$ ; reset awaited flags;  $\mathcal{B} \leftarrow \emptyset$ 
12: end if
13: end while

```

C.2 Asynchronous Conversations

At every iteration, the simulator scans through the active agents and identifies potential conversation pairs (a_i, a_j) . Two types of conversations are considered: (i) **Face-to-face interactions** occur when two agents occupy the same venue within overlapping time windows, and (ii) **Socially initiated communications** occur when an agent’s social need in C_N exceeds a threshold and it proactively contacts another agent through a virtual channel.

Each conversation pair is converted into a communication task $\text{TASK}_{\text{conv}} = (a_i, a_j, \text{MEMORY}_i, \text{MEMORY}_j, \text{context}_{ij})$, where MEMORY_i and MEMORY_j denote the recent memory slots of each participant. Rather than invoking the LLM for each pair independently, the simulator appends these tasks to a global batch $\mathcal{B}_{\text{conv}}$. When the batch size reaches the threshold B_{conv} , all tasks are dispatched in parallel as a single batched API call: $\text{DISPATCHBATCH}(\mathcal{B}_{\text{conv}}) = \{(\Delta \mathcal{M}_i, \Delta \mathcal{M}_j, \Delta R_{ij}) = \text{LLM}_{\text{comm}}(\text{TASK}_{\text{conv}})\}$. Here, $\Delta \mathcal{M}_i$ and $\Delta \mathcal{M}_j$ represent the exchanged memory indices, and ΔR_{ij} updates the bilateral relationship score between agents i and j . Once processed, the updated memories and relationship states are written back into each agent’s local store: $\mathcal{M}_i \leftarrow \mathcal{M}_i \cup \Delta \mathcal{M}_i$, $\mathcal{M}_j \leftarrow \mathcal{M}_j \cup \Delta \mathcal{M}_j$, and $R_{ij} \leftarrow R_{ij} + \Delta R_{ij}$.

Algorithm 2 Asynchronous Batched Conversation Scheduling

```

1: Initialize active agents  $\mathcal{A}$ ; conversation batch
    $\mathcal{B}_{\text{conv}} \leftarrow \emptyset$ 
2: while  $\mathcal{A} \neq \emptyset$  do
3:   for each potential pair  $(a_i, a_j)$  from  $\mathcal{A}$  do
4:     if face_to_face( $a_i, a_j$ ) or
       high_social_need( $a_i$ ) then add
        $(a_i, a_j, \text{MEMORY}_i, \text{MEMORY}_j, \text{context}_{ij})$ 
       to  $\mathcal{B}_{\text{conv}}$ 
5:   end if
6: end for
7: if  $|\mathcal{B}_{\text{conv}}| \geq B_{\text{conv}}$  then dispatch  $\mathcal{B}_{\text{conv}}$  to
   LLM in parallel;
8: update  $\mathcal{M}_i \leftarrow \mathcal{M}_i \cup \Delta \mathcal{M}_i$ ,  $\mathcal{M}_j \leftarrow \mathcal{M}_j \cup \Delta \mathcal{M}_j$ ,
    $R_{ij} \leftarrow R_{ij} + \Delta R_{ij}$ ;  $\mathcal{B}_{\text{conv}} \leftarrow \emptyset$ 
9: end if
10: end while

```

D Experiments

D.1 Dataset Description

Our proprietary dataset is derived from a survey of over 4,000 respondents and contains continuous human parameters, ranging between 0 and 1. Human parameters capture personality and lifestyle traits. In addition, our dataset includes detailed daily activity schedules for each individual, specifying the modes of transportation used for different activities. These real-world schedules serve as a benchmark to assess the faithfulness of our proposed simulation framework, ensuring that it accurately reflects human behavioral patterns.

D.2 Example of Daily Plans

An example of actions generated by baseline (Park et al., 2023) is provided below:

```
00:00 sleeping
06:00 waking up, getting ready for the day
06:30 having breakfast, checking her emails
07:00 commuting to Hobbs Cafe
```

which receives a score of 3. And the actions generated by our model are:

```
07:06 wake up, stretch, make coffee
08:00 check messages, read the news
09:15 work on a project, attend virtual meeting
11:24 cook lunch, eat with a friend,
chat with Mike
```

which scores 4 out of 5.

To analyze agents’ emotional states, we group agents by employment status (unemployed, part-time, employed) and compute the average values of their basic needs over five weekdays. As shown in Figure 2, the attributes fluctuate least for unemployed agents, moderately for part-time workers, and most dramatically for employed agents. All attributes, except *fullness*, follow a consistent pattern: a steady decline between 9:00 and 18:00, followed by recovery during non-working hours. This trend arises because employed agents are predominantly occupied with work during the day, which restricts them from engaging in replenishing activities. *Fullness*, however, rises at 8:00, 12:00, and 18:00, corresponding to mealtimes.

D.3 Additional Baseline Information

In this section, we present a comparative analysis of our proposed framework, *MobileCity*, against three widely recognized approaches for modeling urban interactions: *SmallCity* (Park et al., 2023), *AGA* (Yu et al., 2024), and *HumanoidAgent* (Wang

et al., 2023). Our evaluation focuses on six key dimensions essential for simulating real-world urban behaviors: *daily activities*, *long-term habits*, *basic needs*, *remote communication*, *vehicle usage*, and *movements*. Table 5 provides a detailed comparison of these methods with human behavior.

The *Daily Activities* column assesses a system’s capacity to execute structured, day-to-day tasks, while *Long-Term Habits* measures its ability to develop and sustain recurring behavioral patterns over time. The *Basic Needs* criterion reflects the model’s capability to account for essential human necessities. *Remote Communication* evaluates how well the system facilitates interactions with external entities across distances. *Vehicle Usage* examines mobility-related functionalities, and *Compulsory Tasks* refers to the model’s ability to incorporate mandatory or routine obligations into its behavioral framework.

E Discussion

E.1 Potential Improvement

Our model presents several directions for future enhancement.

First, the introduction of rare events represents a significant challenge. While we have enhanced the plausibility of agent behaviors through the implementation of both dynamic and static agent characteristics, our current framework does not account for environmental dynamic variations beyond weather patterns. To investigate collective behavioral patterns during emergency scenarios such as earthquakes, floods, or fires, these events would need additional modules to produce human-like responses.

Second, our agent interaction mechanisms require refinement. The current paradigm restricts interactions to conversations between agents. A more valid approach would permit multi-agent dialogue sessions and collective activities such as group recreational events.

Third, the model does not yet fully represent heterogeneity in behavioral execution. In real settings, agents require varying durations to complete the same actions, and the resulting attribute changes differ across individuals. Future work should more precisely formalize and parameterize the relationship between agent personality traits and the variability in behavioral outcomes.

Table 5: Comparison of *MobileCity* with prior approaches.

Name	Daily Activities	Long-Term Traits	Basic Needs	Remote Communication	Vehicles	Compulsory Tasks
SmallCity	✓	✗	✗	✗	✗	✗
AGA	✓	✗	✗	✗	✗	✗
HumanoidAgent	✓	✗	✓	✗	✗	✗
MobileCity	✓	✓	✓	✓	✓	✓

E.2 Future Research

Future research endeavors could concentrate on the following directions.

First, cross-cultural urban simulation represents a promising avenue of inquiry. The incorporation of cultural factors and their influence on urban agent behaviors would enable the exploration of divergent collective behavioral patterns across different cultural contexts. Additionally, the datasets serving as foundational sources should encompass subjects from diverse cultural backgrounds to ensure comprehensive representation.

Second, policy evaluation applications offer significant practical value. Leveraging simulation outcomes to assess the potential implications of urban planning decisions and to forecast behavioral adaptations among citizens following the implementation of various policies could inform evidence-based governance strategies.

Third, long-term memory and learning mechanisms require careful examination. Changes in the environment affect how agents accumulate and transfer experiences, shaping their future behaviors based on past interactions. For example, if a transportation route becomes congested due to infrastructure changes, and agents share this information within the system, a shift in commuting patterns is expected as agents adapt to avoid delays.