

BeliefNest: A Joint Action Simulator for Embodied Agents with Theory of Mind

Rikunari Sagara

r.sagara.628@nitech.jp
University of Shizuoka

Koichiro Terao

koichiro.terao@oka-pu.ac.jp
Okayama Prefectural University

Naoto Iwahashi

naoto.iwahashi@oka-pu.ac.jp
Okayama Prefectural University



Abstract: This paper introduces an open-source simulator, BeliefNest, designed to enable embodied agents to perform collaborative tasks by leveraging Theory of Mind. BeliefNest dynamically and hierarchically constructs simulators within a Minecraft environment, allowing agents to explicitly represent nested belief states about themselves and others. This enables agent control in open-domain tasks that require Theory of Mind reasoning. The simulator provides a prompt generation mechanism based on each belief state, facilitating the design and evaluation of methods for agent control utilizing large language models (LLMs). We demonstrate through experiments that agents can infer others' beliefs and predict their belief-based actions in false-belief tasks.

Tags: AI, Games, Agent

1 Introduction

Theory of Mind is a fundamental cognitive ability that underpins human social behavior, enabling individuals to infer the beliefs, intentions, and knowledge of others. In this paper, we propose **BeliefNest**, an open-source simulator designed to support research on collaborative behavior in embodied agents endowed with Theory of Mind capabilities. Recent advances in embodied agents powered by large language models (LLMs) have shown promising progress. However, there is still

no platform that can explicitly represent nested belief states and integrate them with action generation mechanisms. **BeliefNest** addresses this gap by providing a flexible simulation framework that incorporates both hierarchical belief structures and prompt generation support.

BeliefNest offers the following features:

- Explicit representation of nested belief states, as studied in Theory of Mind, using hierarchical simulators (see [Section 3](#))
- Support for prompt generation based on each belief state, enabling the design and evaluation of methods for agent control with LLMs (see [Section 4](#))
- Integration with the Minecraft environment, which is widely used in LLM agent research (*Altera.AL., 2024; Guanzhi Wang & Anandkumar., 2023; Julian Yocum & Russell., 2023; Shaofei Cai & Liang., 2025*), and support for open-domain tasks

In this paper, we describe the design and functionality of **BeliefNest** and demonstrate its effectiveness through experiments on false-belief tasks. The source code is available at Github.

2 Related Work

Theoretical frameworks involving hierarchical belief structures, such as *X thinks that Y thinks that ...*, have been widely discussed in dialogue and planning research that leverages Theory of Mind (*Jasper A Taylor & Mellish., 1996; Mark K Ho & Cushman., 2022*). In addition, numerous studies have aimed to explicitly model such nested beliefs to enhance the efficiency of dialogue and planning systems (*Christian Muise & Sonenberg., 2021; McEleney & O'Hare., 2005; Philip R Cohen & Shvo., 2023*). More recently, growing attention has been given to evaluating the extent to which LLMs can perform Theory of Mind reasoning (*Hainiu Xu & He., 2024; James W. A. Strachan & Becchio., 2024; Kosinski., 2024; Maarten Sap & Choi., 2022; Ullman., 2023; Winnie Street & Dunbar., 2024*), as well as to developing methods for improving ToM capabilities in LLMs (*Alex Wilf & Morency., 2024; Hongxin Zhang & Gan., 2023; Ini Oguntola & Sycara., 2023; Lance Ying & Shu., 2024; Shuwen Qiu & Zheng., 2023; Zhining Zhang & Shu., 2025*). Further, several works~(*Alex Wilf & Morency., 2024; Chun-Hsiao Yeh & Ma., 2025; Wufei Ma & Yuille., 2025*) have addressed tasks that require distinguishing between different perspectives, an aspect central to Theory of Mind. In particular, Yeh et al.~(*Chun-Hsiao Yeh & Ma., 2025*) highlight the limitations of LLMs in distinguishing belief states grounded in differing perspectives. The **BeliefNest** simulator proposed in this study aims to support Theory of Mind-based reasoning by systematically organizing belief state information to be fed into LLMs, thereby addressing these limitations.

3 Nested Structure of the Simulator

Figure 1 illustrates the overall architecture of **BeliefNest**. The system consists of a real-world simulator, which represents the actual environment, and multiple belief simulators that hierarchically represent the belief structures of agents. All of these simulators are instantiated within the Minecraft world, each occupying a distinct and isolated region of the environment. Each agent possesses a dedicated belief simulator that represents its own beliefs and can dynamically construct additional layers as needed. This allows for the explicit representation of nested belief states such as myself in the mind of another or someone else in the mind of another. Belief simulators can be dynamically added or removed depending on the context. In this paper, we refer to simulators in which an agent is physically present as shallower-level simulators, and to those that reflect the beliefs held by that agent as deeper-level simulators.




Figure 1: Overall structure of **BeliefNest**. The simulators form a nested hierarchy that corresponds to the nested structure of beliefs. Let Z be a finite sequence of agent identifiers that represents a path of belief nesting. In the simulator representing state s^Z , agent i constructs a new simulator whose state $s^{Z||i}$ reflects its belief b_i^Z . Here, the operator $||$ denotes sequence concatenation: $Z||i$ is the sequence obtained by appending agent i to the end of Z . In the real world, agent 1 can see the red and green blocks, while agent 2 can see the blue and green blocks. Agent 1 is unaware of the blue block, and agent 2 is unaware of the red block. Therefore, agent 1 believes that agent 2 can see only the green block. These belief-dependent perceptions are explicitly represented within each simulator. All agents are assumed to have prior knowledge of the floor, the white outer walls, and the black inner wall.

4 Observation and Belief Updating

Each belief simulator reflects the belief state based on observations made by agents in shallower-level simulators. However, when an agent engages in planning or hypothesis testing, it must temporarily suspend belief updates based on shallower-level simulators and instead generate its own timeline within the belief simulator to perform inference. Here, a `timeline` refers to a sequence of observations and actions recorded over time. The ability to generate such internal timelines corresponds to mental simulation, which is essential for planning. To support this functionality, we define two operational modes for simulators:

- **Control mode:** The agent performs actions and generates a virtual timeline.
- **Follow mode:** The simulator passively updates its state by sequentially reflecting belief information from the shallower-level simulator.

Each simulator can switch modes independently. The pseudocode of the asynchronous belief updating process for each simulator is shown in [Equation 1](#). Below, we describe the main processes performed in each mode.

Initialization: $b_i \quad \forall i \in \mathcal{A}$

Repeat forever:

$$\begin{aligned}
 s &\leftarrow \begin{cases} \text{get_state}() & \text{if mode} = \text{"control"} \quad (\text{C-1}) \\ \text{receive_belief}() & \text{if mode} = \text{"follow"} \quad (\text{F-1}) \end{cases} \\
 &\text{if mode} = \text{"follow"} \text{ then} \\
 &\quad \text{apply_to_simulator}(s) \quad (\text{F-2}) \\
 &\text{for each } i \in \mathcal{A} : \\
 &\quad o_i \leftarrow O_i(s) \quad (\text{Common-1}) \\
 &\quad b_i \leftarrow \text{update_belief}(b_i, o_i) \quad (\text{Common-2}) \\
 &\quad \text{send_belief}(i, b_i) \quad (\text{Common-3})
 \end{aligned} \tag{1}$$

In control mode, agents act within the simulator, and environment and agent information at each time step is recorded. The simulator performs the following C- processes, which are specific to control mode:

C-1: State acquisition The current state s , representing the state of the environment and the agents, is obtained as objective information. Since state transitions are managed internally by the Minecraft environment, **BeliefNest** does not track them directly; thus, this step retrieves the current state as needed.

In follow mode, the simulator receives and reflects belief states sent from the shallower-level simulator. The simulator performs the following F- processes, which are specific to follow mode:

F-1: Receiving belief state from shallower-level simulator Belief information sent via step (Common-3) in the shallower-level simulator is received and adopted as the current state.

F-2: Reflecting state in the simulator environment The received state is used to update the environment and agent configurations within the simulator.

In addition, as common processing, the following steps are performed for each agent $i \in \mathcal{A}$, where \mathcal{A} denotes the set of agents in the simulator.

Common-1: Observation based on perspective Subjective observations o_i are obtained via the observation function O_i , which captures visible elements in the environment and other agents from the agent's viewpoint. **Common-2: Belief updating** The belief state b_i is updated using observation o_i . Environmental and agent-related information is updated for observable elements, while unobserved elements remain unchanged. **Common-3: Propagation to deeper-level simulators** The updated belief state is sent to the belief simulators associated with each agent.

As an example, when the real-world simulator is in control mode and all belief simulators are in follow mode, agent actions in the real-world simulator propagate belief updates to each belief simulator, and the resulting changes are reflected in their internal environments. Subsequently, if a particular belief simulator switches to control mode, belief updates from the real world are no longer applied to it. The agent's actions within that simulator become a hypothetical simulation, exerting no influence on the real world.

Agents within any simulator can be controlled via JavaScript programs. Since the state of each belief simulator reflects the beliefs of the agents in the corresponding shallower-level simulator, it is possible to generate JavaScript programs using LLMs based on the simulator's state, enabling agents to take actions grounded in their beliefs. To support program generation by LLMs, **BeliefNest** provides functionality for prompt generation based on the state of each simulator. When a prompt template written in Jinja2 format is provided, **BeliefNest** automatically inserts a string representation of the simulator's state into the specified locations within the template to generate the final prompt. Users are free to employ any prompt template or LLM of their choice, making it possible to design and evaluate various methods for agent control utilizing LLMs.

As control primitives, a set of functions is provided to enable agent actions. These include, for example, movement, block breaking, item crafting, and taking from or storing items into chests. These capabilities are implemented using the Mineflayer library. We build upon and modify the action primitives introduced in Voyager (Guanzhi Wang & Anandkumar, 2023) to better fit our objectives.

5 Timeline Management

Each simulator implements a mechanism for managing timeline branches, which represent divergent temporal progressions. Each branch can be assigned a unique identifier, and users can switch to any existing branch at any time. When a branch is selected, the simulator reverts to the latest state recorded in that branch, allowing the timeline to proceed from that point onward. If it is necessary to return to a specific point in the past, a branch can be created at that time in advance, enabling later restoration to that state.

By leveraging the branching mechanism, users can freely design and evaluate control strategies in which agents hypothetically explore multiple action plans based on others' beliefs before selecting one for actual execution. The specific procedures, such as how trials are conducted or how prompts are structured, are left to the user, while **BeliefNest** provides a flexible simulation framework to support such exploratory processes.

6 Example 1: Sally-Anne Task

The Sally-Anne task (*S Baron-Cohen & Frith., 1985*) is a well-known example of a first-order false-belief task, widely used to assess whether an individual can infer beliefs that differ from their own. In this experiment, we evaluate whether a particular agent (hereafter referred to as the observer) can infer the beliefs and belief-based actions of other agents. Notably, the observer is not an external third party but an in-world agent embedded within the environment. Through the belief simulator constructed by **BeliefNest**, the observer is expected to correctly infer Sally's belief and the actions derived from it.

Figure ? shows the experimental setup, which includes three players (Sally, Anne, and the observer) and two boxes. In this task, subtle differences in the spatial configuration of agents can lead to varying visibility conditions and thus to multiple plausible interpretations. We define the following three conditions to capture these variations:

- **Condition 1 (standard setup):** Sally places a diamond into the left box and then moves behind a wall. Anne subsequently transfers the diamond to the right box, and Sally does not observe this action. The observer witnesses the entire sequence.
- **Condition 2:** Unlike Condition 1, Sally moves to a location where she can observe Anne's action. The observer sees the entire sequence as well.
- **Condition 3:** As in Condition 2, Sally moves to a position where she can see Anne's action, but the observer cannot see where Sally moved, and therefore does not know that Sally observed the action.

A first-order false-belief task measures whether an agent can infer another agent's first-order belief. In this setting, the observer constructs a belief about Sally's belief, which corresponds to a second-order belief from the observer's perspective. This experiment evaluates whether the second-order belief simulator is constructed correctly, thereby enabling the observer to infer Sally's belief. Furthermore, based on the constructed belief state, the Sally agent within the belief simulator is given the task `Get a diamond from a chest`. We then evaluate whether the generated action is consistent with the inferred belief state, using GPT-4o (OpenAI, 2024) to generate actions based on prompts informed by belief states.



Figure 2: Condition 1.



Figure 3: Condition 2. The observer can see Sally.




Figure 4: Condition 3. The observer cannot see Sally.

In the simulator representing the observer's first-order belief (i.e., the observer's own belief), the diamond was moved to the right box in all conditions. In contrast, the second-order belief simulator (i.e., the observer's belief about Sally's belief) produced different results depending on the condition. The inferred beliefs and actions were as follows:

- **Condition 1:** Reflecting the fact that Sally did not observe Anne's action, it was inferred that she believed the diamond was still in the left box. Consequently, it was predicted that Sally would open the left box.
- **Condition 2:** Since the observer recognized that Sally had observed Anne's action, the second-order belief also inferred that Sally believed the diamond was in the right box. Therefore, it was predicted that Sally would open the right box.
- **Condition 3:** Although Sally actually saw Anne's action, the observer could not confirm this. As a result, in the second-order belief, it was inferred that Sally believed the diamond remained in the left box, and thus would open the left box.

These results demonstrate that **BeliefNest** is capable of accurately constructing others' beliefs that differ from an agent's own, based on subtle differences in spatial configuration. Moreover, it enables appropriate inference of belief-based actions in accordance with those constructed beliefs.

7 Example 2: Ice Cream Van Task

The Ice Cream Van Task (*Perner & Wimmer, 1985*) is a second-order false-belief task that involves a two-level belief hierarchy (e.g., X believes that Y believes that Z), and is used to evaluate whether an agent can correctly handle deeper belief structures. In this experiment, we examine whether the observer can infer what John thinks Mary believes.

Figure 5 illustrates the experimental setup, which includes four players: John, Mary, the ice cream seller, and the observer. The scenario unfolds as follows:

- The ice cream seller tells both John and Mary at location A that he will stay there today.
- Mary returns home.
- The ice cream seller then tells John that he is going to location B, and leaves location A.
- On the way, the ice cream seller unexpectedly encounters Mary and tells her as well that he is going to location B.
- John returns home.

The observer witnesses the entire sequence. The goal is to determine whether the observer can accurately infer John's second-order belief by verifying whether the third-order belief simulator (i.e., observer's belief about John's belief about Mary's belief) is correctly constructed.

Based on the resulting belief state, the Mary agent within the belief simulator is given the task "Go to the ice cream seller. The experiment then evaluates whether the generated action is appropriate with respect to the inferred belief.



Figure 5: Experimental setup for the Ice Cream Van Task.

In the observer's third-order belief simulator, Mary was inferred to go to location B. This inference is based on the observer's recognition that John was unaware of Mary's reunion with the ice cream seller in Step 4. This result aligns with the expected reasoning for this task, demonstrating that **BeliefNest** can accurately represent complex belief structures and generate plausible belief-based behavior.

8 Conclusion

In this paper, we introduced **BeliefNest**, an open-source simulator designed to support research on collaborative behavior by equipping embodied agents with Theory of Mind capabilities. Through experiments, we demonstrated that agents using **BeliefNest** can correctly infer others' beliefs and

predict their actions in false-belief tasks. While this study focused on false-belief scenarios to validate the core functionalities of **BeliefNest**, the simulator is applicable to a broader range of research on joint action. We expect that **BeliefNest** will serve as a valuable foundation for advancing studies on collaboration in dialogue systems and robotics.

References

- Alex Wilf, & Morency, L.-P. (2024). *Think twice: Perspective-taking improves large language models' theory-of-mind capabilities*. In Proceedings of ACL 2024, pages 8292–8308, 2024.
- Altera.AL. (2024). Project sid: Many-agent simulations toward AI civilization. *arXiv Preprint*.
- Christian Muise, & Sonenberg, L. (2021). Efficient multi-agent epistemic planning: Teaching planners about nested belief. *arXiv Preprint*.
- Chun-Hsiao Yeh, & Ma, Y. (2025). Seeing from another perspective: Evaluating multi-view understanding in MLLMs. *arXiv Preprint*.
- Guanzhi Wang, & Anandkumar, A. (2023). Voyager: An open-ended embodied agent with large language models. *arXiv Preprint*.
- Hainiu Xu, & He, Y. (2024). OpenToM: A comprehensive benchmark for evaluating theory-of-mind reasoning capabilities of large language models. *arXiv Preprint*.
- Hongxin Zhang, & Gan, C. (2023). Building cooperative embodied agents modularly with large language models. *arXiv Preprint*.
- Ini Oguntola, & Sycara, K. (2023). Theory of mind as intrinsic motivation for multi-agent reinforcement learning. *arXiv Preprint*.
- James W. A. Strachan, & Becchio, C. (2024). *Testing theory of mind in large language models and humans*. Nature Human Behaviour, 2024.
- Jasper A Taylor, Jean Carletta, & Mellish, C. (1996). *Requirements for belief models in cooperative dialogue*. User Model. User-adapt Interact., 6(1):23–68, March 1996.
- Julian Yocum, & Russell, S. (2023). Mitigating generative agent social dilemmas. In *NeurIPS 2023 Foundation Models for Decision Making Workshop, November 2023*.
- Kosinski, M. (2024). *Evaluating large language models in theory of mind tasks*. Proceedings of the National Academy of Sciences, 121(45):e2405460121, 2024.
- Lance Ying, & Shu, T. (2024). GOMA: Proactive embodied cooperative communication via goal-oriented mental alignment. *arXiv Preprint*.
- Maarten Sap, & Choi, Y. (2022). *Neural theory-of-mind? on the limits of social intelligence in large LMs*. In Proceedings of EMNLP 2022, pages 3762–3780, 2022.
- Mark K Ho, Rebecca Saxe, & Cushman, F. (2022). *Planning with theory of mind*. Trends in Cognitive Sciences, 2022.
- McEleney, B., & O'Hare, G. (2005). Efficient dialogue using a probabilistic nested user model. *Proceedings of the Fourth IJCAI Workshop on Knowledge and Reasoning in Practical Dialogue Systems, August 2005*.
- OpenAI. (2024). GPT-4o system card. *arXiv Preprint*.
- Perner, J., & Wimmer, H. (1985). "John thinks that Mary thinks that..." attribution of second-order beliefs by 5- to 10-year-old children. J. Exp. Child Psychol., 39(3):437–471, June 1985. \endthebibliography.
- Philip R Cohen, Lucian Galescu, & Shvo, M. (2023). An explainable collaborative dialogue system using a theory of mind. *arXiv Preprint*.
- S Baron-Cohen, A. M. Leslie, & Frith, U. (1985). *Does the autistic child have a "theory of mind"?* Cognition, 21(1):37–46, October 1985.
- Shaofei Cai, & Liang, Y. (2025). ROCKET-2: Steering visuomotor policy via cross-view goal alignment. *arXiv Preprint*.
- Shuwen Qiu, & Zheng, Z. (2023). MindDial: Belief dynamics tracking with theory-of-mind modeling for situated neural dialogue generation. *arXiv Preprint*.
- Ullman, T. (2023). Large language models fail on trivial alterations to theory-of-mind tasks. *arXiv Preprint*.
- Winnie Street, & Dunbar, R. I. M. (2024). LLMs achieve adult human performance on higher-order theory of mind tasks. *arXiv Preprint*.
- Wufei Ma, & Yuille, A. (2025). SpatialReasoner: Towards explicit and generalizable 3D spatial reasoning. *arXiv Preprint*.
- Zhining Zhang, & Shu, T. (2025). AutoToM: Automated bayesian inverse planning and model discovery for open-ended theory of mind. *arXiv Preprint*.