Recursive Metropolis-Hastings naming game: symbol emergence in a multi-agent system based on probabilistic generative models

Jun Inukai[1†], Tadahiro Taniguchi[1*†], Akira Taniguchi[1] and
Yoshinobu Hagiwara[2]

[1]Graduate School of Information Science and Engineering, Ritsumeikan University, Kusatsu, Japan,
[2]Research Organization of Science and Technology, Ritsumeikan University, Kusatsu, Japan

In the studies on symbol emergence and emergent communication in a population of agents, a computational model was employed in which agents participate in various language games. Among these, the Metropolis-Hastings naming game (MHNG) possesses a notable mathematical property: symbol emergence through MHNG is proven to be a decentralized Bayesian inference of representations shared by the agents. However, the previously proposed MHNG is limited to a two-agent scenario. This paper extends MHNG to an N-agent scenario. The main contributions of this paper are twofold: (1) we propose the recursive Metropolis-Hastings naming game (RMHNG) as an N-agent version of MHNG and demonstrate that RMHNG is an approximate Bayesian inference method for the posterior distribution over a latent variable shared by agents, similar to MHNG; and (2) we empirically evaluate the performance of RMHNG on synthetic and real image data, i.e., YCB object dataset, enabling multiple agents to develop and share a symbol system. Furthermore, we introduce two types of approximations—one-sample and limited-length—to reduce computational complexity while maintaining the ability to explain communication in a population of agents. The experimental findings showcased the efficacy of RMHNG as a decentralized Bayesian inference for approximating the posterior distribution concerning latent variables, which are jointly shared among agents, akin to MHNG, although the improvement in ARI and κ coefficient is smaller in the real image dataset condition. Moreover, the utilization of RMHNG elucidated the agents' capacity to exchange symbols. Furthermore, the study discovered that even the computationally simplified version of RMHNG could enable symbols to emerge among the agents.

1. Introduction

The origin of language remains one of the most intriguing mysteries of human evolution (Deacon, 1997; Steels, 2015; Christiansen and Chater, 2022). Humans utilize various symbol systems, including gestures and traffic lights. Although the language is considered a type of symbol (or sign) system, it boasts the richest structure and the strongest ability to describe events among all symbol systems (Chandler, 2002). The adaptive, dynamic, and emergent nature of symbol systems is a common feature in human society (Taniguchi et al., 2018, 2021). This paper focuses on the emergent nature of general

symbols and their meanings, rather than the structural complexity of language. The meaning of signs is determined within society in bottom-up and top-down manners, owing to the nature of symbol systems (Taniguchi et al., 2016). More specifically, the self-organized (or emergent) symbol system enables each agent to communicate semiotically with others while being subject to the top-down constraints of the emergent symbol system. Agent-invented symbols can hold meaning within a society of multiple agents, even though no agent can directly observe the intention in the brain of a speaker. Peirce, the founder of semiotics, defines a symbol as a triadic relationship between sign, object, and interpretant (Chandler, 2002). The interpretant serves as a mediator between the sign and the object. In nature, the relationship between sign and object exhibits arbitrariness. This implies that human society—a multi-agent system using a symbol system—must form and maintain these relationships within a society in a decentralized manner. Research on symbol emergence, language evolution, and emergent communication has been addressing this issue using a constructive approach for decades.
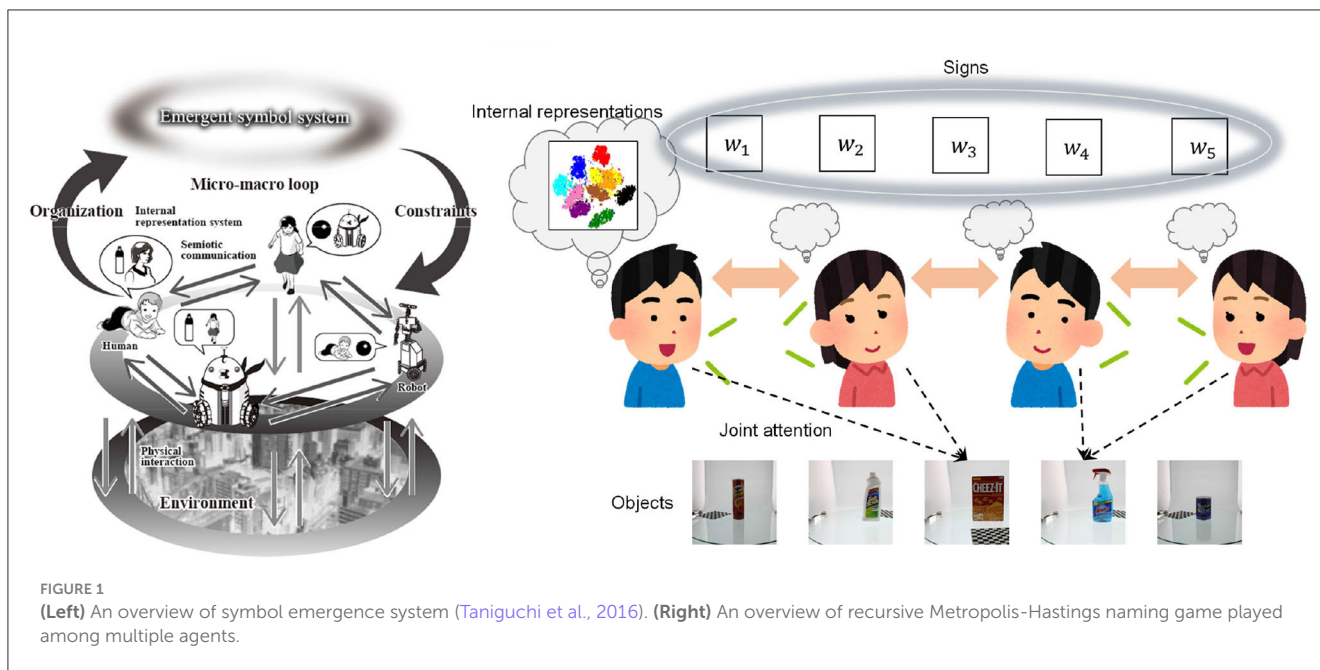
Studies on emergent communication take many forms. Numerous studies have explored emergent communication by engaging agents in Lewis-style signaling games, such as referential games. The referential game is a cooperative game in which a speaker agent sends a message (i.e., a sign) representing an object, and a listener agent receives and interprets the message to identify a target object. The primary objective of the game is to enable both agents to learn the relationship between signs and objects to establish a language or communication protocol. Lazaridou et al. (2017, 2018) and Havrylov and Titov (2017) demonstrated that agents can communicate using their own language by performing reference games. Furthermore, Choi et al. (2018) and Mu and Goodman (2021) suggested that the compositionality of emergent language can be improved by modifying the real image data used in reference games. However, Bouchacourt and Baroni (2018) highlighted the issue of agents being able to communicate even when using uninterpretable images in referential games. Noukhovitch et al. (2021) demonstrated the necessity of referential games for agent communication. Numerous studies have also attempted emergent communication with multiple agents. Gupta et al. (2021) explored extending to multiple agents using meta-learning, while Lin et al. (2021) employed autoencoder, a standard representation learning algorithm. Chaabouni et al. (2022) investigated the effects of varying the number of agents in referential games on agent communication. These studies successfully achieved communication through games that provided rewards.

In contrast, Taniguchi et al. (2023) proposed an alternative formulation of emergent communication based on probabilistic generative models and the assumption of joint attention. The Metropolis-Hastings naming game (MHNG) was introduced to explain the process by which two agents share the meaning of signs in a bottom-up manner from a Bayesian perspective. It was demonstrated that symbol emergence can be considered decentralized Bayesian inference. In the model, signs are treated as latent variables shared by all agents, as shown in Figure 1 on the right. If there's a special agent that centrally manages the agents and can refer to the internal variables (i.e., belief states)

of all agents, the shared variable can be inferred in a centralized manner. Conversely, decentralized Bayesian inference implies that the latent variable can be inferred without each agent having access to the internal variables of its counterparts. MHNG assumes *joint attention* between two agents—widely observed in human infants learning vocabularies—instead of reward feedback from a listener to a speaker. This idea is rooted in the concept of a symbol emergence system, as shown in Figure 1 on the left (Taniguchi et al., 2016, 2018), rather than the view of an emergent communication channel often assumed in emergent communication studies based on Lewis-style signaling game. The notion of a symbol emergence system was proposed to capture the overall dynamics of symbol emergence from the perspective of emergent systems, i.e., complex systems exhibiting emergent properties. This approach aims to further investigate the fundamental cognitive mechanisms enabling humans to organize symbol systems within a society in a bottom-up manner.

In this paper, we use the term *symbol system* in a restricted sense. Here, a symbol system simply refers to a set of signs and their (probabilistic) relationship to objects. In the context of studies on symbol emergence and emergent communication, we cannot assume a ground-truth relationship between signs and objects, unlike many studies in artificial intelligence, e.g., standard pattern recognition task that assumes a ground-truth label given to each object. Ideally, a multi-agent system should form a symbol system with which agents can appropriately categorize (or differentiate) objects and associate signs with objects. The definition of appropriate categorization and sign sharing is crucial to the formulation of symbol emergence. Different approaches assume different goals of symbol emergence and criteria based on various hypothetical principles. *Iterated learning* assumes that the goal of symbol emergence is for each agent to use the same sign for each object. In contrast, *emergent communication based on referential games* assumes that organizing signs allows a speaker to provide information that enables a listener to choose an object intended by the speaker. Taniguchi et al. proposed a collective predictive coding (CPC) hypothesis in the discussion of Taniguchi et al. (2023). The CPC hypothesis posits that the goal of symbol emergence is the formation of global representations created by agents in a decentralized manner. This can also be called *social representation learning*, i.e., symbol emergence is conducted as a representation learning process by a group of individuals in a decentralized manner. From a Bayesian perspective, this can be regarded as decentralized Bayesian inference.

The MHNG was proposed, demonstrating that the language game enables two agents (agents 1 and 2) to form a symbol system, with MHNG's process mathematically considered as a Bayesian inference of $p(w \mid o^1, o^2)$, where o^1 and o^2 represent the observations of agents 1 and 2, and w represents the shared representations, i.e., signs. Furthermore, MHNG does not assume the existence of explicit feedback from the listener to the speaker in the game, unlike Lewis-style signaling games widely employed in emergent communication studies. Instead, MHNG assumes joint attention, considered foundational to language acquisition during early developmental stages (Cangelosi and Schlesinger, 2015). The CPC hypothesis and MHNG are based on *generative models* rather than *discriminative models*, which are prevalent in the dominant

FIGURE 1
(Left) An overview of symbol emergence system (Taniguchi et al., 2016). **(Right)** An overview of recursive Metropolis-Hastings naming game played among multiple agents.

approach to emergent communication in the deep learning community. The MHNG and results of constructive studies substantiate the CPC hypothesis in a tangible manner (Taniguchi et al., 2023). However, existing studies on MHNG only demonstrate that the game can become a decentralized approximate Bayesian inference procedure in a two-agent scenario. No theoretical research or evidence exists to show that the CPC hypothesis can hold in more general cases, i.e., in N-agent settings where $N \geq 3$. In other words, it is crucial to determine whether a language game can perform decentralized approximate Bayesian inference of $p(w \mid o^{1:N})$, where $o^{1:N} = \{o^1, o^2, \ldots, o^n, \ldots, o^N\}$, and o^n represents the observations of the n-th agent.

The fundamental reason why the MHNG can act as an approximate Bayesian inference of $p(w \mid o^1, o^2)$ is that the utterance of a sign $w \sim p(w \mid o^{Sp})$ by the speaker agent Sp can be sampled based on agent Sp's observations alone, and the acceptance ratio of the sign (i.e., the message) can be solely determined by the listener Li based on its own observations and internal state. These properties are derived from the theory of the Metropolis-Hastings algorithm. MHNG has a solid theoretical basis in Markov Chain Monte Carlo (MCMC) (Hastings, 1970). However, the proof provided by Taniguchi et al. (2023) assumed that the naming game is played between only two agents. This assumption was based on the need for individual agents to make the proposal sampling of a sign and the acceptance/rejection decision, respectively, without direct observation of the internal states of the other agent. Due to the difficulty, a naming game having the same theoretical property as MHNG for the N-agent ($N \geq 3$) case has not been proposed.
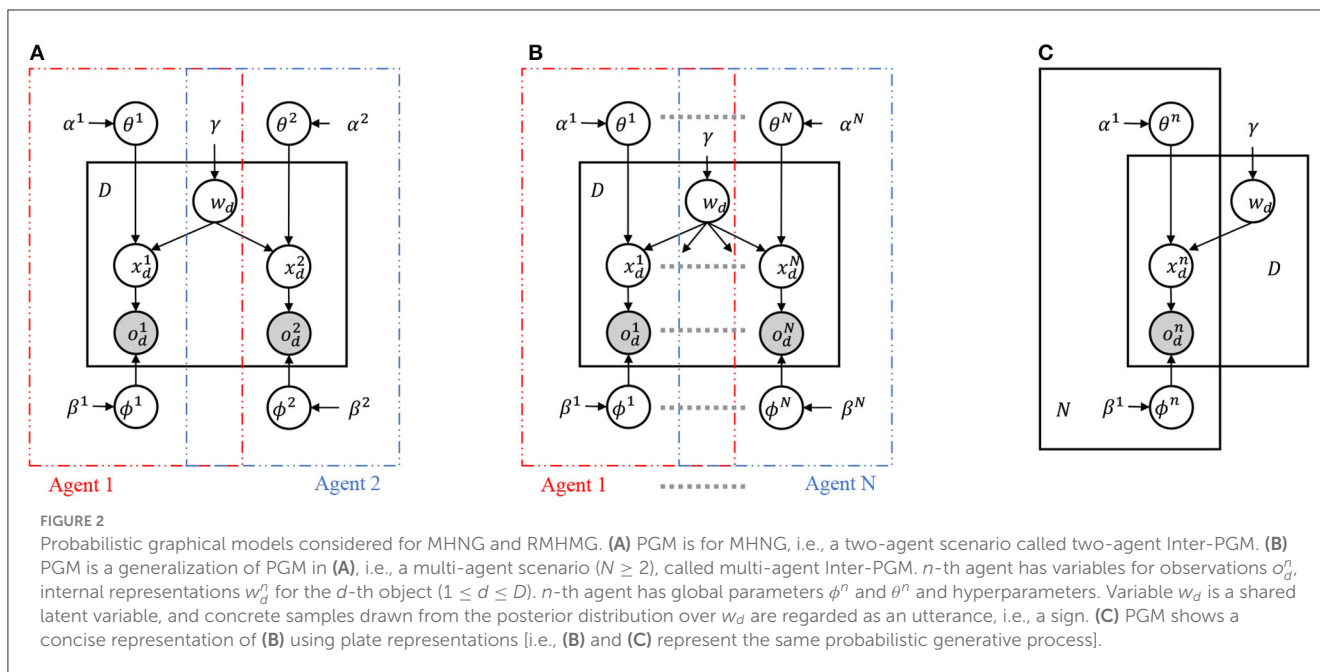
The goal of this paper is to extend the MHNG to the N-agent ($N \geq 3$) scenario and show that the extended naming game can act as an approximate Bayesian inference algorithm for $p(w \mid o^{1:N})$. The main idea of the proposed method is the introduction of a *recursive structure* into the MHNG. Let us consider a 3-agent case. If $w \sim p(w \mid o^1, o^2)$ can be sampled in the MH algorithm, the

acceptance ratio for the third agent can be calculated based on the third agent's internal states, and the communication can be regarded as a sampling process of $p(w \mid o^1, o^2, o^3)$. Notably, $w \sim p(w \mid o^1, o^2)$ can be sampled using the original two-agent MHNG. By extending this idea in a recursive manner, we can develop a recursive MHNG (RMHNG). The details will be described in Section 2.

The main contributions of this paper are twofold.

- We propose the RMHNG played between N agents and provide mathematical proof that the RMHNG acts as an approximate Bayesian inference method for the posterior distribution over a latent variable shared by the agents given the observations of all the agents.
- The performance of the RMHNG is empirically demonstrated on synthetic data and real image data. The experiment shows that the RMHNG enables more than two agents to form and share a symbol system. The inferred distributions of signs are shown to be a posterior distribution over $p(w \mid o^{1:N})$ in an empirical manner. To reduce computational complexity and maintain applicability for the explanation of communication in human society, two types of approximations, i.e., (1) one-sample (OS) approximation and (2) limited-length (LL) approximation, are proposed and both are validated through experiment.

The remainder of this paper is structured as follows. In section 2, we describe RMHNG, explaining its assumed generative model, algorithms, and theoretical results. Additionally, a practical approximation is provided. Section 3 presents an experiment using synthetic data and demonstrates the RMHNG empirically. Section 4 presents an experiment using the YCB object dataset (Calli et al., 2015), which contains real images of everyday objects. In Section 5, we engage in a comprehensive discussion. Finally, we conclude the paper in Section 6.

FIGURE 2
Probabilistic graphical models considered for MHNG and RMHMG. **(A)** PGM is for MHNG, i.e., a two-agent scenario called two-agent Inter-PGM. **(B)** PGM is a generalization of PGM in **(A)**, i.e., a multi-agent scenario ($N \geq 2$), called multi-agent Inter-PGM. n-th agent has variables for observations o_d^n, internal representations w_d^n for the d-th object ($1 \leq d \leq D$). n-th agent has global parameters ϕ^n and θ^n and hyperparameters. Variable w_d is a shared latent variable, and concrete samples drawn from the posterior distribution over w_d are regarded as an utterance, i.e., a sign. **(C)** PGM shows a concise representation of **(B)** using plate representations [i.e., **(B)** and **(C)** represent the same probabilistic generative process].

2. Recursive Metropolis-Hastings naming game

2.1. Overview

The RMHNG is a language game played between multiple agents ($N \geq 2$). It is an extension of the original MHNG. When $N = 2$, the RMHNG is equivalent to the original MHNG. Notably, the game does not allow agents to give any feedback to other agents during the game, unlike Lewis-style signaling games (Lewis, 2008), which have been used in studies of emergent communication. Instead, the game assumes joint attention. Generally, when we ignore the representation learning parts, the original MHNG is played as follows:

1. For each object $d \in \mathbb{D}$, the n-th agent (where $n \in \{1, 2\}$) views the d-th object and infers the internal state x_d^n, its percept, from its observations o_d^n, i.e., calculate $p(x_d^n \mid o_d^n)$ or sample $x_d^n \sim p(x_d^n \mid o_d^n)$, where $\mathbb{D}$ is a set of object. Set the initial roles to $\{Sp, Li\} = \{1, 2\}$.
2. The Sp-th agent says a sign w_d^{Sp} (i.e., a word) corresponding to the d-th object in a probabilistic manner by sampling a sign from the posterior distribution over signs (i.e., $w_d^{Sp} \sim p(w_d \mid x_d^{Sp})$) for each d.
3. Let a counterpart, that is, a listener, be Li-th agent. Assuming that the listener is looking at the same object, i.e., joint attention, the listener determines whether to accept the sign based on its belief state with probability $r = \min\left(1, \frac{P(x_d^{Li} \mid \theta^{Li}, w_d^{Sp})}{P(x_d^{Li} \mid \theta^{Li}, w_d^{Li})}\right)$. A listener updates its internal parameter θ^{Li}.
4. They alternate their roles, i.e., take turns, and go back to 2.

The RMHNG extends the original MHNG to allow for communication between multiple agents ($N \geq 3$) and forms a shared symbol system among them. The key idea of RMHNG is as follows:
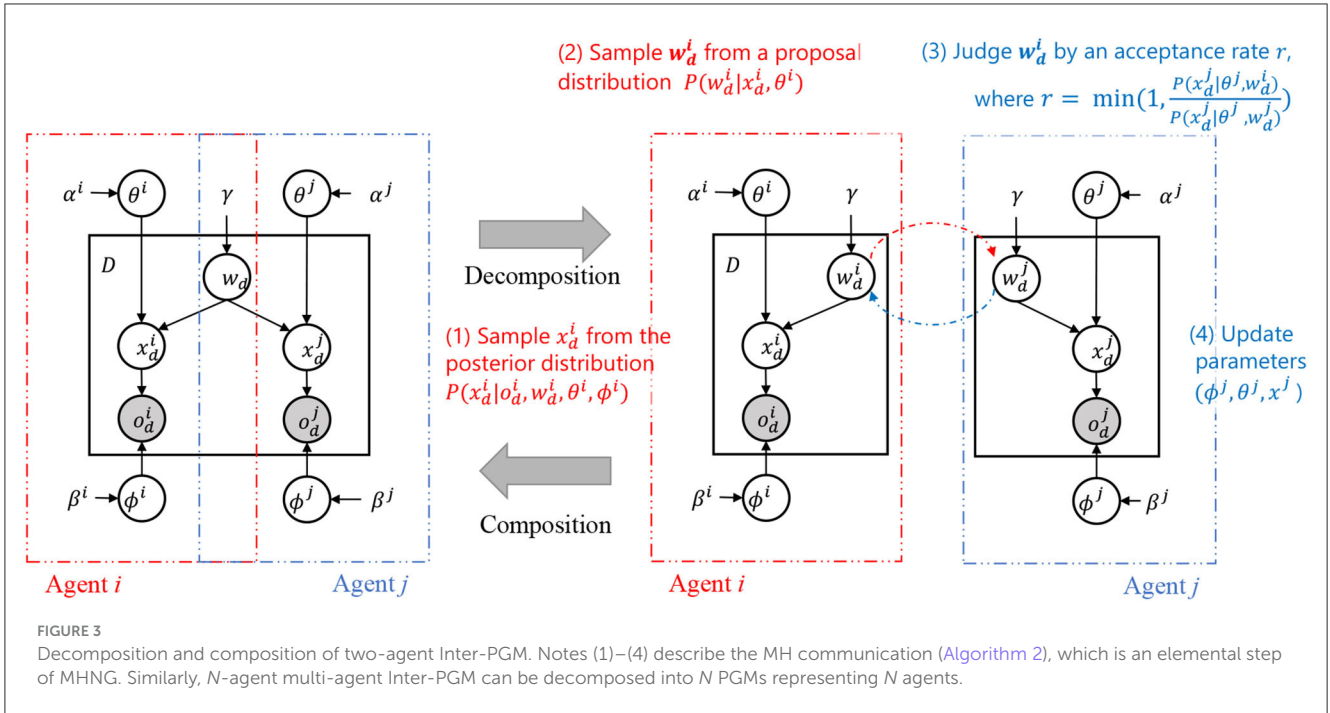
TABLE 1 Variables of Inter-PGM and their explanations.

Variable	Explanation
w_d	A sign, e.g., a name, for the d-th object
x_d^*	Perceptual state or feature vector corresponding to the d-th object
o_d^*	Observation for the d-th object
θ^*	Knowledge about the relations between signs and perceptual states
ϕ^*	Knowledge about relations between perceptual states and observations
α	A hyperparameter for θ^*
β	A hyperparameter for ϕ^*

Superscript $* \in \mathbb{N}$ refers to a specific agent.

1. In an RMHNG played by M agents, we recursively use an RMHNG played by $M - 1$ agents as a proposal distribution of w_d, which corresponds to a speaker in the original MHNG. Note that, an RMHNG played by $M - 1$ agents $(1, \ldots, M - 1)$ is a sampler of an approximate distribution of $p(w_d \mid x_{1:M-1})$.
2. An RMHNG played by two agents ($N = 2$) is equivalent to an original MHNG.

Consequently, when played by N agents, the RMHNG approximates the distribution of $p(w_d \mid x_{1:N})$ through mathematical induction. In RMHNG, the interactions of the agents represent a chain-like communication characterized by local one-to-one iterative message exchanges. Through this process, information is disseminated throughout the multi-agent system. Metaphorically speaking, one person engages in a discussion with another; then, that person converses with the next person based on the outcome of the previous discussion, and this form of communication continues throughout the group.

FIGURE 3
Decomposition and composition of two-agent Inter-PGM. Notes (1)–(4) describe the MH communication (Algorithm 2), which is an elemental step of MHNG. Similarly, N-agent multi-agent Inter-PGM can be decomposed into N PGMs representing N agents.

2.2. Generative model

Figure 2 presents three probabilistic graphical models (PGMs) representing the interactions between multiple agents sharing a latent variable w_d. (A) The left panel shows a PGM that integrates two PGMs representing two agents with a shared latent variable w_d. This model is referred to as the two-agent Inter-PGM. (B) The center panel generalizes the PGM in (A) to integrate PGMs representing N agents. This model can be considered a multimodal PGM in which a shared latent variable integrates multimodal observations. We refer to this model as the multi-agent Inter-PGM. (C) The right panel provides a concise representation of (B) using plate representations, meaning (B) and (C) represent the same probabilistic generative process. When agent n observes the d-th object, they receive observations x_d^n and infer their internal representation x_d^n. A latent variable representing a sign, w_d, is shared among the agents. The inference of θ^n and x_d^n corresponds to a general representation learning problem. As studied in Taniguchi et al. (2023), introducing the Symbol Emergence in Robotics Toolkit (SERKET) framework (Nakamura et al., 2018; Taniguchi et al., 2020) allows us to decompose the main part of the naming game (exchanging signs w_d^n between agents) and the representation learning part (inferencing x_d^n and θ^n). For simplicity and to focus on the extension of the MHNG, we assume that x_d^n can be inferred using a feature extractor and given to the model as a fixed variable throughout the paper, and focus on inferring θ^n and w_d through the naming game. Table 1 summarizes the variables of the generative model.

2.3. Inference as a naming game

The RMHNG, like the MHNG, acts as a decentralized approximate Bayesian inference based on the MH algorithm. A

standard inference scheme for $p(w_d \mid x_d^{1:N})$ in Figure 2C requires the information about $x_d^{1:N}$, e.g., the posterior distribution $p(x_d \mid o_d^{1:N})$. However, $x_d^{1:N}$ are internal representations of each agent, and the agents cannot access each other's internal state, which is a fundamental principle of human semiotic communication. If the agents' brains were connected, the shared variable w_d would be a representation of the connected brain and could be inferred by referencing $x_d^{1:N}$. But this is not the case in real-world communication. The challenge is to infer the shared variable w_d without connecting the agents' brains and without simultaneously referencing $x_d^{1:N}$. The solution is to play the RMHNG.

The decomposition of the generative model inspired by SERKET, as shown in Figure 3 right, allows for a more manageable and systematic approach to the inference of hidden variables. The SERKET framework enables the decomposition of a PGM into multiple modules, which simplifies the overall inference process by breaking it down into inter-module communication and intra-module inference (Nakamura et al., 2018; Taniguchi et al., 2020). In the context of the RMHNG, the semiotic communication between agents is analogous to the inter-module communication in the SERKET framework.

2.3.1. MH receiving

Algorithm 1 presents the MH-receiving algorithm. When a listener agent $A^{Li} \in \mathbb{A}$ receives a sign $w^\star$ for the d-th object, the agent evaluates whether to accept the sign and update $A^{Li}.w_d$ or not, where $\mathbb{A}$ is a set of agents. Here, $A^i.w_d$ represents the w_d that agent A^i possesses. Similarly, $A^{Li}.x_d$ denotes the x_d held by agent A^i. For $A^n \in \mathbb{A}$, A^n is an instance of a struct (or a class), and $A^n.\bullet$ indicates the variable $\bullet$ of the n-th agent, i.e., $(A^n.\theta, A^n.\phi, (A^n.o_d)_{d \in \mathbb{D}}, (A^n.x_d)_{d \in \mathbb{D}}, (A^n.w_d)_{d \in \mathbb{D}}) = (\theta^n, \phi^n, (o_d^n)_{d \in \mathbb{D}}, (x_d^n)_{d \in \mathbb{D}}, (w_d^n)_{d \in \mathbb{D}})$. The function MH-receiving

returns the sign for the d-th object agent Li holds after receiving a new name for the d-th object from another agent.

```
1: function MH-RECEIVING(w*, A^Li, d)
2:     r = min(1, P(A^Li.x_d|A^Li.θ,w*) / P(A^Li.x_d|A^Li.θ,A^Li.s_d))
3:     u ~ Unif(0, 1)
4:     if u ≤ r then
5:         return w*
6:     else
7:         return A^Li.w_d
8:     end if
9: end function
```

Algorithm 1. MH Receiving

2.3.2. MH communication

Algorithm 2 presents the MH-communication algorithm. The function **MH-communication** describes the elementary communication in both the MHNG and the RMHNG. A sign s for the d-th object is sampled (i.e., uttered) by agent Sp and received by agent Li, where $Li, Sp \in \mathbb{N}$.

```
1: function MH-COMMUNICATION(A^Sp, A^Li, d)
2:     w* ~ P(A^Sp.w_d | A^Sp.x_d, A^Sp.θ)
3:     return MH-receiving(w*, A^Li, d)
4: end function
```

Algorithm 2. MH Communication

2.3.3. Recursive MH communication

Algorithm 3 presents the recursive MH communication algorithm. This algorithm represents the recursive MH communication process, as shown in Figure 4. The recursive MH communication is one of the MH sampling procedures for $p(w_d \mid o_d^{1:N})$. Given $n + 1$ ($n < N$) agents, each with parameter w_d, this algorithm is used to compute w_d for interactions among n agents. If $n > 1$, the RMH-communication function is recursively called for agents $A^{1:n-1} \subset \mathbb{A}$ to compute interactions among them. Then, A^{n+1} updates its own parameter w_d using the received information $\bar{s}$ by calling the MH-receiving function. If $n = 1$, the MH-communication function is called. After the internal loop (from line 2 to line 9) is completed, the algorithm returns the w_d of a randomly selected agent j from $A^{1:n+1}$. This algorithm can recursively calculate interactions among N agents.

2.3.4. Recursive MH naming game

Algorithm 4 presents the recursive MH naming game algorithm. The agents repeatedly engage in recursive MH communication for each object, shuffling the order of the agents. The recursive MH communication is mathematically a type of approximate MH sampling procedure for $p(w_d \mid o_d^{1:N})$. After the recursive MH communication is performed for every object, each

```
1: function RMH-COMMUNICATION(A^{1:n}, A^{n+1}, d)
2:     for t = 1 to T do                    ▷ Internal iteration
3:         if n > 1 then
4:             w̄ ← RMH-communication(A^{1:n-1}, A^n, d)
5:             A^{n+1}.w_d ← MH-receiving(A^{n+1}, w̄, d)
6:         else
7:             A^2.w_d ← MH-communication(A^1, A^2, d)
8:         end if
9:     end for
10:    j ← rand({1 : n + 1})
11:    return A^j.w_d
12: end function
```

Algorithm 3. Recursive Metropolis-Hastings Communication

agent internally updates its global parameter θ^n. By iterating this block I times, the agents can sample $\{w_d\}, \{\theta_n\}$ from the posterior distribution over $p(\{w_d\}, \{\theta_n\} \mid \{x_d^n\})$.

```
1: Explanation of each parameter
2: A^{1:N} = (A^1, A^2, ..., A^N)
3: Initialize all parameters
4: for i = 1 to I do                       ▷ Number of iteration
5:     B^{1:N} = shuffle(A^{1:N})  ▷ Randomization of the order of
       the agents
6:     for d = 1 to D do                    ▷ Naming every object
7:         RMH-communication(B^{1:M-1}, B^M, d)
               ▷ M = N : No approximation
8:             ▷ M < N : Limited-length approximation (see
       2.5)
9:     end for
10:    for n = 1 to N do                     ▷ Parameter update
11:        B^n.θ ~ P(B^n.θ | B^n.s, B^n.x)
12:    end for
13: end for
```

Algorithm 4. Recursive Metropolis-Hastings naming game
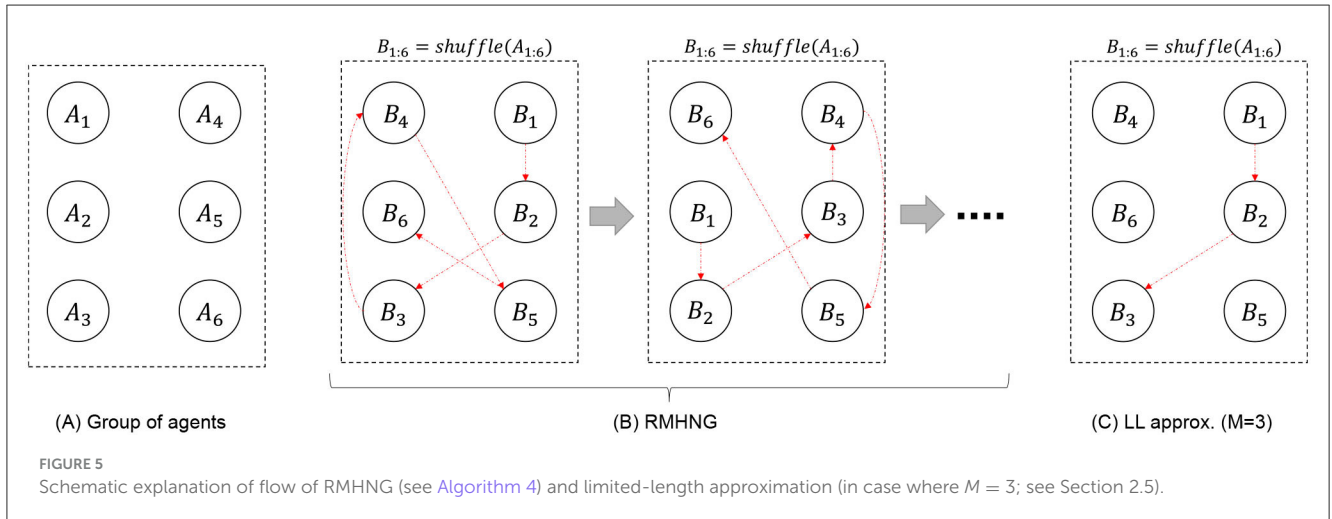
2.4. Theory and proof

For the main theoretical result, we use the following corollary.

Corollary 2.1. *The MH communication is a Metropolis-Hastings sampler of* $P(w_d \mid x_d^{Sp}, x_d^{Li}, \theta^{Sp}, \theta^{Li})$.

The acceptance probability r in **MH-receiving** is equivalent to that in the MH algorithm for $P(w_d \mid x_d^{Sp}, x_d^{Li}, \theta^{Sp}, \theta^{Li})$ in the case that $P(w \mid x^{Sp}, \theta^{Sp})$ is a proposal distribution. This result is a generalization of (Hagiwara et al., 2019, 2022) and a special case of (Taniguchi et al., 2023). For the details of the proof, please refer to the original papers.

The first theoretical result is as follows.

Theorem 1. *The RMH communication converges to a MCMC sampler of* $P(w_d \mid x_d^{1:n}, \theta^{1:n})$ *when* $T \to \infty$.

FIGURE 4
The upper figure is schematic explanation of RMH communication and RMHNG. The recursive MH communication is one of the MH sampling procedures for $p(w_d \mid o_d^{1:N})$. Given $n + 1$ ($n < N$) agents, each with parameter w_d, this algorithm is used to compute w_d for interactions among n agents.

FIGURE 5
Schematic explanation of flow of RMHNG (see Algorithm 4) and limited-length approximation (in case where $M = 3$; see Section 2.5).
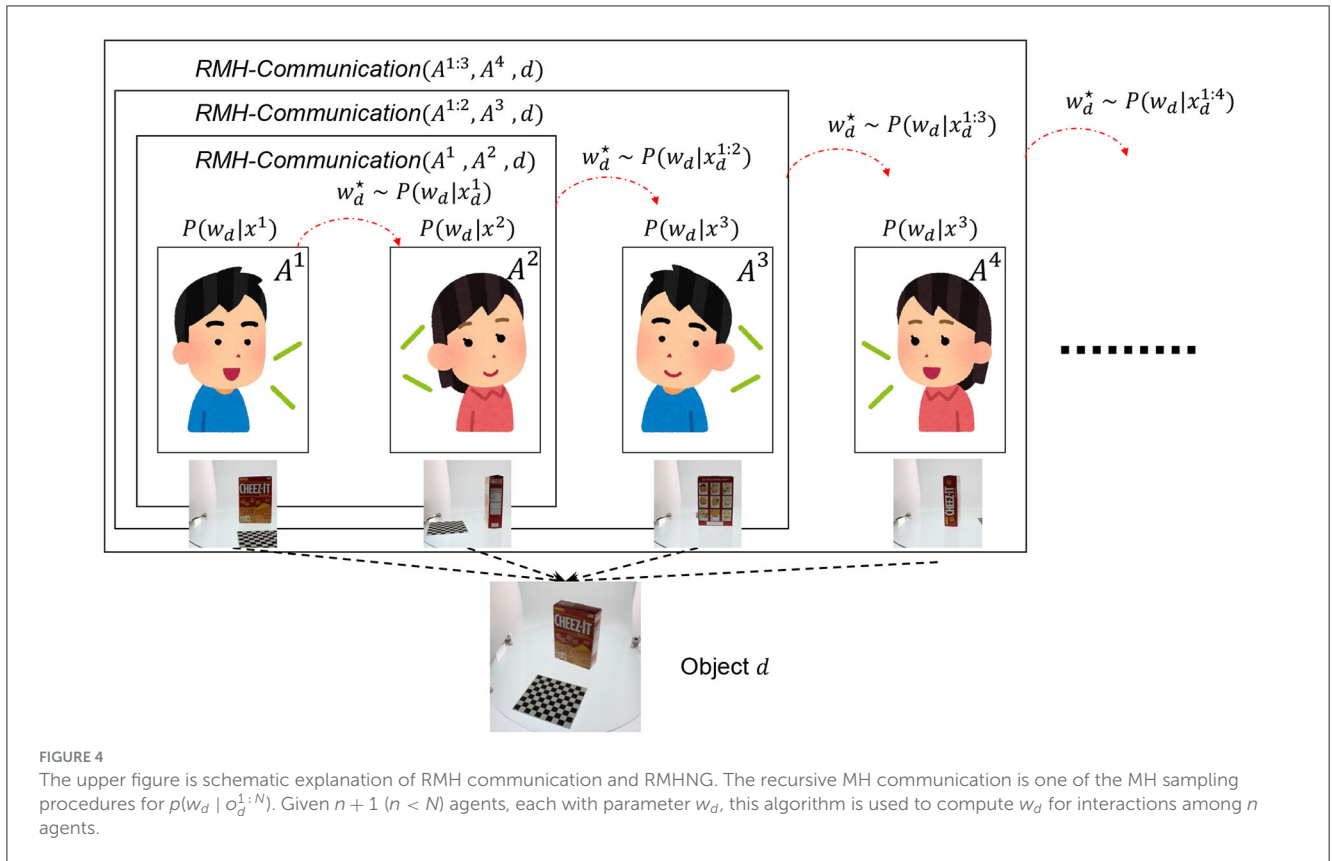
Proof: When $n = 2$, the RMH communication is reduced to the execution of MH communication T times. The MH communication is proven to be an MH sampler in corollary 1. Therefore, RMH communication is a MCMC sampler, and the sample distribution converges to $P(w_d \mid x_d^{1:2}, \theta^{1:2})$ when $T \to \infty$. When $n > 2$, if the $RMH-communication(B^{1:n-1}, B^n, d)$ is a sampler for $P(w_d \mid x_d^{1:n}, \theta^{1:n})$, $RMH-communication(B^{1:n}, B^{n+1}, d)$ becomes an MH sampler for $P(w_d \mid x_d^{1:n+1}, \theta^{1:n+1})$. Therefore, RMH communication is a MCMC sampler, and the sample distribution converges to $P(w_d \mid x_d^{1:n+1}, \theta^{1:N+1})$ when $T \to \infty$. Therefore,

the RMH communication converges to a MCMC sampler of $P(w_d \mid x_d^{1:n}, \theta^{1:n})$ when $T \to \infty$ by mathematical induction.

Theorem 2. *The RMHNG converges to a MCMC sampler of* $P(w_d, \theta^{1:n} \mid x_d^{1:n})$ *when* $T \to \infty$.

Proof: The RMHNG samples the local parameters w_d for all d using the RMH communication, and the global parameters $\theta^{1:n}$ from $P(\theta^{1:n} \mid \{x_d^{1:n}\}_{d \in \mathbb{D}}, \{w_d\}_{d \in \mathbb{D}})$. When $T \to \infty$, RMH communication converges to a sampler of $P(w_d \mid x_d^{1:n}, \theta^{1:n})$. As a result, the RMHNG converges to a Gibbs sampler of $P(w_d, \theta^{1:n} \mid x_d^{1:n})$.

As a result, the RMHNG is proved to be a decentralized approximate Bayesian inference procedure for $p(\{w_d\}_{d\in\mathbb{D}}, \{\theta_n\}_{n\in\{1,\ldots,N\}} \mid \{x_d^n\}_{d\in\mathbb{D}})$.

2.5. Approximations

Though the RMHNG is guaranteed to be a decentralized approximate Bayesian inference procedure for $p(\{w_d\}_{d\in\mathbb{D}}, \{\theta_n\}_{n\in\{1,\ldots,N\}} \mid \{x_d^n\}_{d\in\mathbb{D}})$, the computational cost increases exponentially with respect to the number of agents N. The computational cost is $O(IDT^{(N-1)})$. This indicates that the computational cost of **RMH-communication**, i.e., $O(T^{(N-1)})$, has a significant impact on the overall computational cost. Therefore, we introduce a lazy version of RMHNG, which employs two approximations to reduce the computational cost.

2.5.1. One-sample approximation

The number of internal iterations T corresponds to the iterations of MCMC for sampling w_d given variables of a (sub)group of agents. Theoretically, T should be large. However, practically, even $T = 1$ can work in an approximate manner. We refer to the RMHNG with $T = 1$ as the OS approximation (OS), a special case. With the OS, the computational cost of RMH communication is significantly reduced from $O(T^{(N-1)})$ to $O(N)$.

2.5.2. Limited-length approximation

RMH communication is a process of information propagation through a chain connecting N agents (as shown in Figure 5). Limited-length approximation (LL) truncates the chain to M agents. By shuffling the order of the agents according to the data points, it is expected that sufficient information will be statistically propagated among all the agents. LL reduces the computational cost of RMH communication from $O(T^{(N-1)})$ to $O(T^{(M-1)})$, where $M \leq N$ is the length of the truncated chain, i.e., the number of agents participating in an RMH communication. To reduce computational complexity while maintaining applicability for explaining communication in human society, two types of approximations are proposed: (1) OS approximation and (2) LL approximation. Both types were validated through experimentation.

Notably, when $M = 2$, the LL approximation corresponds to the random sampling condition. Random sampling refers to the process where two individuals are randomly selected from a pool (i.e., a set of agents), they communicate with each other, and after their communication concludes, they are returned to the pool. This process is then repeated. In the experiment, the random sampling condition is primarily investigated as a representative of the LL approximation.

2.6. Example: multi-agent Inter-GMM

To evaluate the RMHNG, we developed a computational model of symbol emergence called multi-agent Inter-GMM. This is based on the Gaussian mixture model (GMM) and is a special case of the multi-agent Inter-PGM. Hagiwara et al. (2019, 2022) proposed the

Inter-Dirichlet mixture (Inter-DM) which combines two Dirichlet mixtures (DMs), $p(x_d^n \mid w_d)$ and $p(o_d^n \mid x_d^n)$, represented as categorical distributions in Figure 2A. Taniguchi et al. (2023) proposed Inter-GMM + VAE which combines two GMM + VAEs, i.e., $p(x_d^n \mid w_d)$ and $p(o_d^n \mid x_d^n)$ represented as a categorical distribution as a part of GMM and a VAE respectively. Inter-GMM is defined as a part of Inter-GMM + VAE and combines two GMMs via a shared latent variable. We generalized the two-agent Inter-GMM and obtained the multi-agent Inter-GMM, which has N Gaussian emission distributions corresponding to N agents. The probabilistic generative process of the multi-agent inter-GMM is as follows:

$$w_d \sim \text{Cat}(\gamma) \qquad\qquad d = 1,\ldots,D \tag{1}$$

$$\mu_k^n, \Lambda_k^n \sim \mathcal{N}(\mu_k^n \mid m, (\bar{\alpha}\Lambda_k^n)^{-1})\mathcal{W}(\Lambda_k^n \mid \nu, \bar{\beta}) \qquad k = 1,\ldots,K \tag{2}$$

$$\alpha^n = (m, \bar{\alpha}, \nu, \bar{\beta}) \tag{3}$$

$$\theta^n = (\mu_{1:K}^n, \Lambda_{1:K}^n) \tag{4}$$

$$x_d^n \sim \mathcal{N}(x_d^n \mid \mu_{w_d}^n, (\Lambda_{w_d}^n)^{-1}) \qquad d = 1,\ldots,D \tag{5}$$

where μ_k^n and Λ_k^n are the mean vector and the precision matrix of the k-th Gaussian distribution of the n-th agent. $\text{Cat}(*)$ is the categorical distribution, $\mathcal{N}(*)$ is the Gaussian distribution, $\mathcal{W}(*)$ is the Wishart distribution. In this model, we assume that the agents share a pre-defined set of signs, i.e., $w_d^n \in \mathbb{S}$, which is a finite set.[1] The Inter-GMM is a probabilistic generative model represented by the PGM shown in Figure 2C. In other words, the multi-agent Inter-GMM is an instance of the multi-agent Inter-PGM. Therefore, the RMHNG can be directly applied to the multi-agent Inter-GMM.

3. Experiment 1: synthetic data

3.1. Conditions

We evaluated the RMHNG using the multi-agent Inter-GMM with four agents ($N = 4$) using synthetic data. For all experiments (excluding the measurement of computation time), the number of iterations (I) was set to 100, and each experiment was conducted five times.

3.1.1. Dataset

We created synthetic data to serve as observations for the four agents. A dataset was generated from five 4-dimensional Gaussian distributions with mean vectors of $(0, 1, 2, 3)$, $(0, 5, 6, 7)$, $(8, 5, 10, 11)$, $(12, 13, 10, 15)$, and $(16, 17, 18, 15)$, respectively. The variance of each Gaussian distribution was set to the identity

---

1 The assumption does not imply that every agent knows the pre-defined signs. In particular, if the i-th agent knows a set of signs $\mathbb{S}_i$, we can define the universal set of signs as the union of all individual sets, represented as $\mathbb{S} = \bigcup_i \mathbb{S}_i$ as the shared set of signs. This assumption is about the mathematical formulation, not about the knowledge of each agent.

matrix $\mathbf{I}$. The values obtained for each dimension were taken as observations for each agent. In other words, the value of the n-th dimension of data sampled from the GMM was considered as the observation for the n-th agent. Notably, for the n-th agent, the n-th and $n+1$-th Gaussian distributions have the same mean and variance. Therefore, the n-th agent cannot differentiate the n-th and $n+1$-th Gaussian distributions without communication.

3.1.2. Compared methods

e assessed the proposed model, *RMHNG* (proposal), by comparing it with two baseline models and a topline model. In *No communication* (baseline 1), two agents independently infer a sign w, i.e., perform clustering of the data. No communication occurs between the four agents. In other words, the *No communication* model assumes that the agents independently infer signs w_d^n ($n \in \{1, 2, 3, 4\}$), respectively, using four GMMs. *All acceptance* (baseline 2) is the same as the RMHNG, with an acceptance ratio always set to $r = 1$ in MH receiving (see Algorithm 1). Each agent always believes that the sign of the other is correct. In *Gibbs sampling* (topline), the sign w_d is sampled using the Gibbs sampler. This process directly uses $x_d^{1:4}$, although no one can simultaneously examine the internal (i.e., brain) states of human communication. This is a centralized inference procedure and acts as a topline in this experiment.

We also evaluated two approximation methods introduced in Section 2.5. OS and LL refer to the OS and LL approximations, respectively. In the LL approximation, $M = 2$, i.e., the chain length is one. In OS&LL, both OS and LL approximations were applied simultaneously.

3.1.3. Hyperparameters

In all methods, the hyperparameters of the agents were set to be the same. The hyperparameters were $\beta = 1$, $m = 0$, $W = 0.01$, and $\nu = 1$.

3.1.4. Evaluation criteria

- *Clustering*: We used Adjusted Rand Index (ARI) (Hubert and Arabie, 1985) to evaluate the unsupervised categorization performance of each agent in the MH naming game. A high ARI value indicates excellent categorization performance, while a low ARI value indicates poor performance. ARI is advantageous over precision since it accounts for label-switching effects in clustering by comparing the estimated labels and ground-truth labels. Appendix B provides more details.
- *Sharing sign*: We assessed the degree to which the two agents shared signs using the κ coefficient (κ) (Cohen, 1960). Appendix B provides more details.
- *Computation time*: We conducted experiments to measure the processing time of the program when running it at $I = 10$ by varying the values of T in Algorithm 3 and M in Algorithm 4. We conducted experiments with $T = 1, 2, 3, 4$ and $M = 2, 3, 4$. The program was run three times in each experiment (30 iterations in total, initialized every 10 iterations), and we calculated the average processing time per iteration (10 iterations).

- *Decentralized posterior inference*: To investigate whether RMHNG is an approximate Bayesian estimator of the posterior distribution $p(w \mid x^1, x^2, \ldots, x^N, \theta^1, \theta^2, \ldots, \theta^N)$, we need to compare it with the true posterior distribution. However, computing the true posterior distribution $p(w \mid x^1, x^2, \ldots, x^N, \theta^1, \theta^2, \ldots, \theta^N)$ directly is difficult. Therefore, we evaluate how well the distribution of signs obtained by RMHNG matches that of Gibbs sampling. Appendix A provides more details.

3.1.5. Machine specifications

The experiment was conducted on a desktop PC with an Intel(R) Core(TM) i9-9900K CPU @ 3.60GHz 3.60 GHz, 32GB of RAM, and an NVIDIA GeForce RTX 2080 SUPER GPU.

3.2. Results

3.2.1. Categorization and sign sharing

Table 2 shows the ARI and κ for each method used on the artificial data. As shown in Table 2, the ARI values for RMHNG were consistently close to those of Gibbs sampling, with a maximum difference of only 0.1. This indicates that RMHNG had a similar category classification accuracy as Gibbs sampling. In this setting, OS performed even better than RMHNG, achieving the highest values for both ARI and κ. This might be because OS facilitated the mixing process by introducing randomness in sampling. On the other hand, OS&LL and LL exhibited relatively low values for both ARI and κ. Notably, even with approximations, RMHNG had higher agent classification accuracy and sign-sharing rate than both No communication and All acceptance.

3.2.2. Change in ARI and κ for each iteration

Figure 6 shows the ARI (right) and κ (left) for each iteration (i in Algorithm 4). From the left graph in Figure 6, we can see that RMHNG, OS, and LL converge faster in terms of ARI, in that order, among the RMHNG and its approximation methods. OS&LL show an upward trend in ARI even at the 100th iteration, indicating that they have not converged. No communication has the fastest convergence in ARI among all the methods. As for All acceptance, we can see that the ARI does not show an upward trend even as the iteration count increases, compared to other methods. From the right graph in Figure 6, we can see that RMHNG, OS, and LL converge faster in terms of κ, in that order, among the RMHNG and its approximation methods. OS&LL show an upward trend in κ even at the 100th iteration, indicating that they have not converged. No communication and All acceptance do not show an upward trend in κ even as the iteration count increases, compared to other methods.
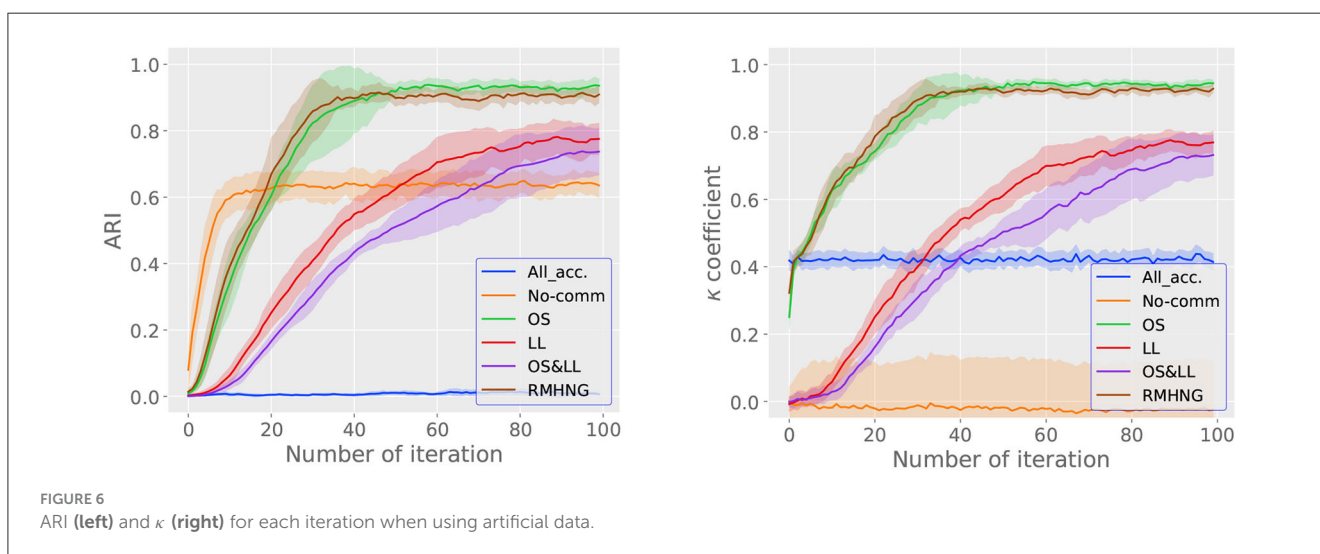
3.2.3. Computation time

Figure 7 shows the average computation time for varying values of M and T in RMHNG. As shown in the figure, it can be seen that the computation time increases logarithmically as T increases.

TABLE 2 Experimental results for synthetic data: each method was tested five times, and for each agent, ARI and κ were calculated when I was between 91 and 100.

Condition	ARI (Agent 1)	ARI (Agent 2)	ARI (Agent 3)	ARI (Agent 4)	κ
RMHNG	0.91 ± 0.02	0.92 ± 0.01	0.92 ± 0.01	0.88 ± 0.02	0.92 ± 0.01
OS	$\mathbf{0.93 \pm 0.02}$	$\mathbf{0.94 \pm 0.01}$	$\mathbf{0.94 \pm 0.02}$	$\mathbf{0.91 \pm 0.02}$	$\mathbf{0.94 \pm 0.01}$
LL	0.77 ± 0.04	0.80 ± 0.03	0.81 ± 0.02	0.71 ± 0.03	0.77 ± 0.03
OS&LL	0.73 ± 0.09	0.76 ± 0.06	0.77 ± 0.05	0.68 ± 0.05	0.73 ± 0.07
No communication	0.64 ± 0.03	0.67 ± 0.03	0.65 ± 0.01	0.60 ± 0.02	-0.02 ± 0.14
All acceptance	0.008 ± 0.005	0.009 ± 0.007	0.009 ± 0.007	0.009 ± 0.007	0.42 ± 0.02
Gibbs sampling	0.98 ± 0.01				-

Mean $\pm$ standard deviation of obtained $50(5 \times 10)$ARI and κ are shown. The highest scores are shown in bold, and the second-highest scores are underlined.

FIGURE 6
ARI **(left)** and κ **(right)** for each iteration when using artificial data.

Considering that the vertical axis is logarithmic, this confirms that the computation time follows the computational complexity of $O(T^{M-1})$. Additionally, it can be confirmed that significant reductions in computation time can be achieved by approximating RMHNG with OS ($T = 1, M = 4$), LL ($T = 4, M = 2$), or OS&LL ($T = 1, M = 2$). Specifically, RMHNG ($T = 4, M = 4$) took 3,178 s, OS ($T = 1, M = 4$) took 77 s, LL ($T = 4, M = 2$) took 187 s, and OS&LL ($T = 1, M = 2$) took 71 s.

3.2.4. Decentralized posterior inference

Figure 8 shows the results of calculating how closely the sign distribution obtained by each method matches that obtained by Gibbs sampling in the last 10 iterations (91–100 iterations) for each method. RMHNG shows a value of 0.96, indicating that the sign distribution obtained by RMHNG matches that obtained by Gibbs sampling by 96%. This confirms that RMHNG is an approximate Bayesian estimator for the posterior distribution $p(w \mid x^1, x^2, \ldots, x^N, \theta^1, \theta^2, \ldots, \theta^N)$. Additionally, OS shows a value of 0.9 or higher, indicating that it is also an approximate Bayesian estimator for the posterior distribution $p(w \mid x^1, x^2, \ldots, x^N, \theta^1, \theta^2, \ldots, \theta^N)$. Although LL and OS&LL have lower values compared to LL and OS, respectively, they are found to have

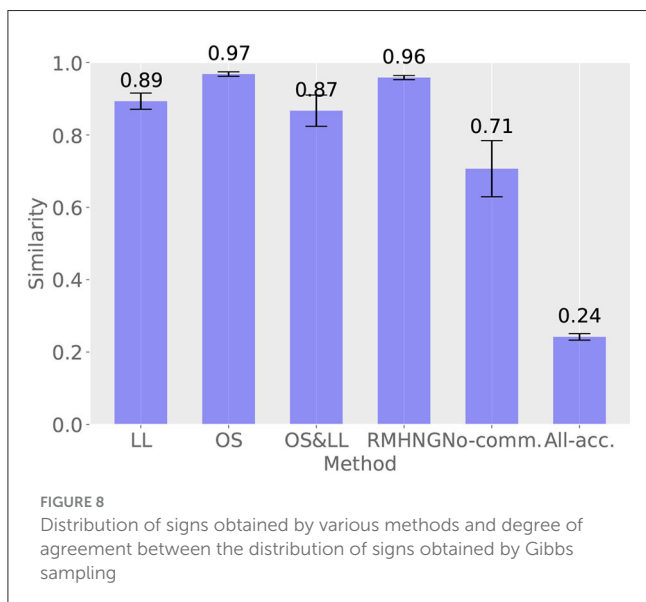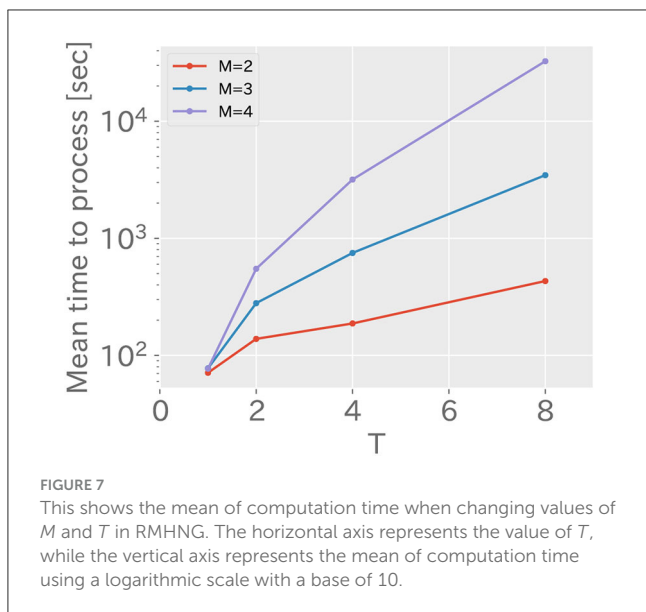higher matching rates with the sign distribution obtained by Gibbs sampling than No communication and All acceptance.

4. Experiment 2: YCB object dataset

4.1. Conditions

We evaluated RMHNG using the multi-agent Inter-GMM with four agents ($N = 4$) on a real image dataset. For all experiments (except for measuring computation time), the number of iterations (I) was set to 100, and each experiment was conducted five times.

4.1.1. Dataset

We evaluated the performance of RMHNG using the YCB object dataset. We selected several objects from the dataset, and their names are listed in Figure 9A. Figure 9B, shows an overview of the dataset, where we divided the images of each object into four sets and assigned each set to one of the four different agents. Each set consisted of 30 images. Specifically, images ranging from 0° to 87° were assigned to agent 1, those from 90° to 177° were assigned

FIGURE 7
This shows the mean of computation time when changing values of M and T in RMHNG. The horizontal axis represents the value of T, while the vertical axis represents the mean of computation time using a logarithmic scale with a base of 10.

FIGURE 8
Distribution of signs obtained by various methods and degree of agreement between the distribution of signs obtained by Gibbs sampling

to agent 2, those from $180°$ to $267°$ were assigned to agent 3, and those from $270°$ to $267°$ were assigned to agent 4.

4.1.2. Feature extraction

Firstly, we cropped the original images from $4,272 \times 2,848$ to a size of $2,000 \times 2,000$ from the center. Next, we reduced the cropped images to a size of 300×300 to prevent any degradation in image quality. Finally, we cropped the images further to a size of 224×224 from the center. We used the resulting images as the observations for each agent, denoted as o_d^n.[2]

---

2 We conducted two rounds of cropping for two reasons. Firstly, direct reduction of an ultra-high-resolution image can lead to degradation of image quality, so it was necessary to crop the image to a certain size before reducing

Feature extraction was performed using SimSiam (Chen and He, 2021), a representation learning method based on self-supervised learning, pre-trained on the collected cropped YCB-object dataset. The feature extractor outputted 512-dimensional vectors. To address the issue of high feature dimensionality compared to the small amount of data available, principal component analysis (PCA) was used to reduce the features to 10 dimensions.[3] Figure 10 shows a visualization of the features of all data and the features observed by each agent using PCA. From this figure, it can be expected that some degree of categorization is possible.

4.1.3. Hyperparameters

All agents were assigned the same hyperparameters, with values set as follows: $\beta = 1$, $m = 0$, $W = 100 \times \mathbf{I}$, and $\nu = 1$, where $\mathbf{I}$ is a 10-dimensional identity matrix $\mathbf{I}$.

Compared method and *evaluation criteria* are the same as those in the Experiment 1.

4.2. Result
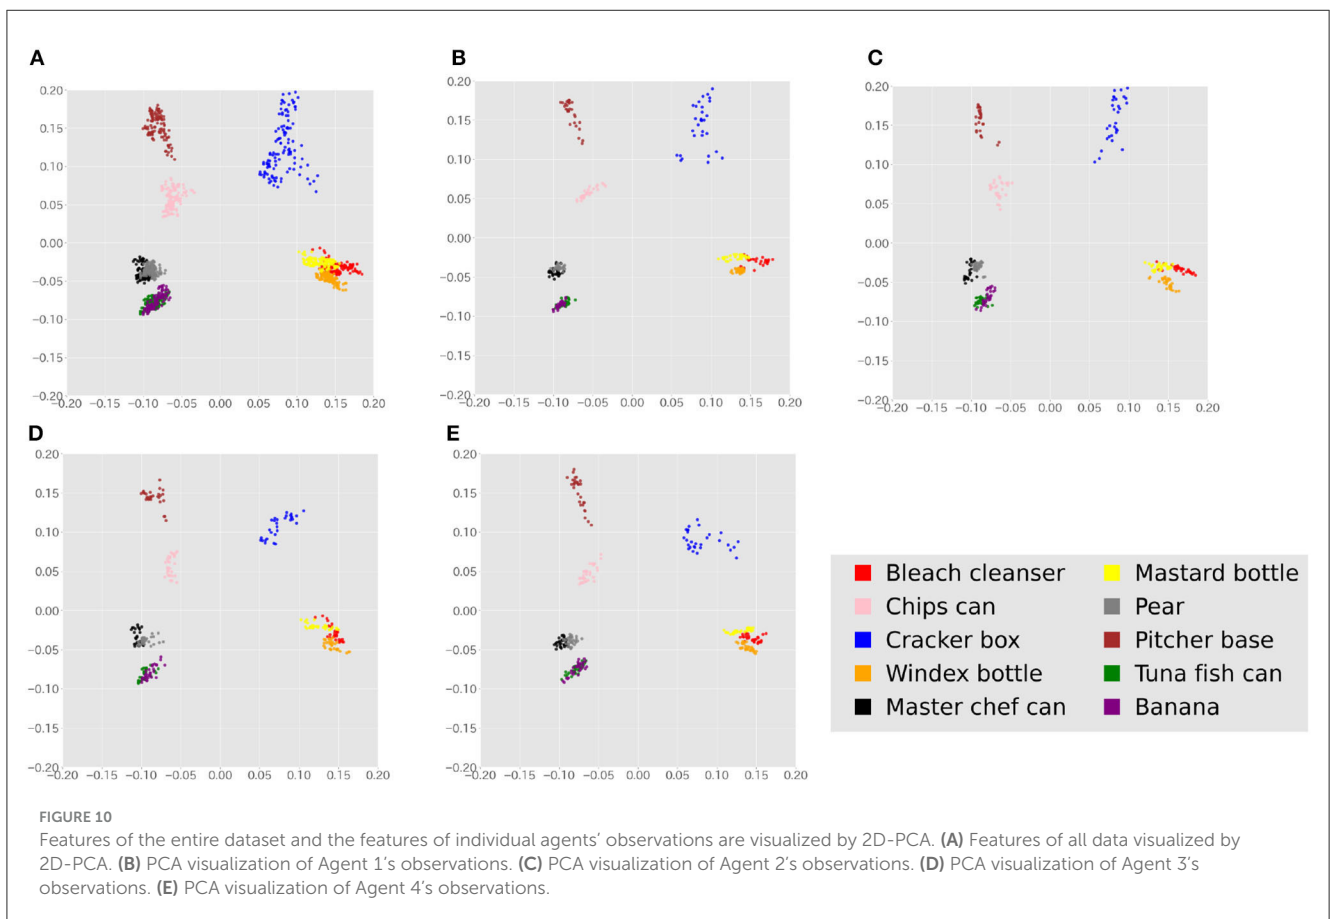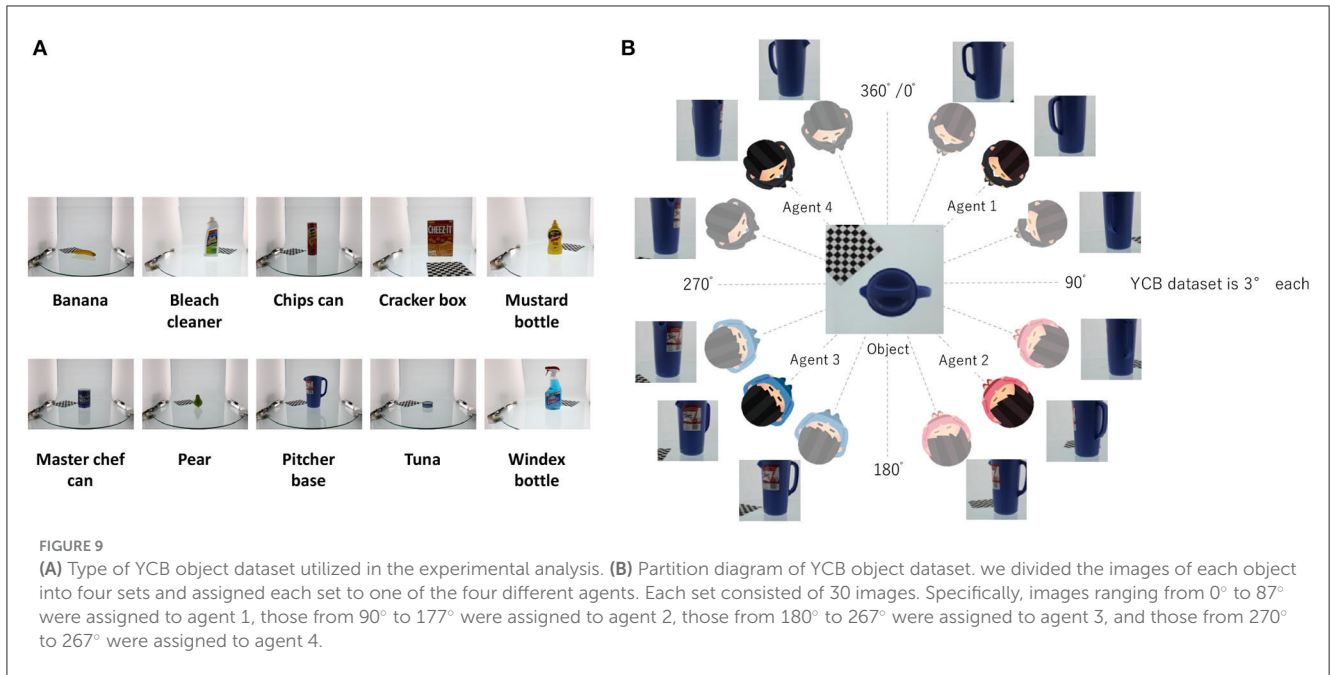
4.2.1. Categorization and sharing signs

Table 3 shows the ARI and κ for each method on the YCB object dataset. It is observed that RMHNG and Gibbs sampling have similar category classification accuracy with a maximum difference of only 0.04. Among the RMHNG approximations, OS had the highest ARI and κ values. Interestingly, it showed a value close to that of RMHNG for the κ. OS&LL had the lowest values for both ARI and κ. However, the difference in ARI between LL, OS, and OS&LL was at most 0.02, indicating similar performance. In the YCB object dataset experiments, although OS showed higher ARI than RMHNG in the synthetic data experiment, for Agent1, OS showed lower ARI than RMHNG, while for other agents, it showed similar values. Compared to OS&LL, No communication showed equivalent ARI for Agent1, lower ARI for Agent2, and higher ARI for other agents. However, the κ was the lowest among all methods for No communication. All acceptance had the lowest ARI among all methods and the highest κ among all methods.

4.2.2. Change in ARI and κ for each iteration

Figure 11 shows the ARI (right) and κ (left) for each iteration (i) in Algorithm 4 for various methods, while Figure 6 shows the convergence of the κ for synthetic data. From the left figure in Figure 11, we can see that the RMHNG method has the fastest convergence of ARI, followed by OS, OS&LL, and LL. Regarding OS&LL, we can see that ARI did not converge when using synthetic data, but it did converge when using the YCB object dataset. No communication had the fastest convergence of ARI among all

---

it. Secondly, some objects were almost entirely obscured or markers were largely reflected, making it difficult to crop directly from the original image.

3 This was conducted to avoid zero variance in dimensions that cannot be inferred, preventing divergence of the accuracy matrix of GMMs.

FIGURE 9
(A) Type of YCB object dataset utilized in the experimental analysis. **(B)** Partition diagram of YCB object dataset. we divided the images of each object into four sets and assigned each set to one of the four different agents. Each set consisted of 30 images. Specifically, images ranging from 0° to 87° were assigned to agent 1, those from 90° to 177° were assigned to agent 2, those from 180° to 267° were assigned to agent 3, and those from 270° to 267° were assigned to agent 4.

FIGURE 10
Features of the entire dataset and the features of individual agents' observations are visualized by 2D-PCA. **(A)** Features of all data visualized by 2D-PCA. **(B)** PCA visualization of Agent 1's observations. **(C)** PCA visualization of Agent 2's observations. **(D)** PCA visualization of Agent 3's observations. **(E)** PCA visualization of Agent 4's observations.
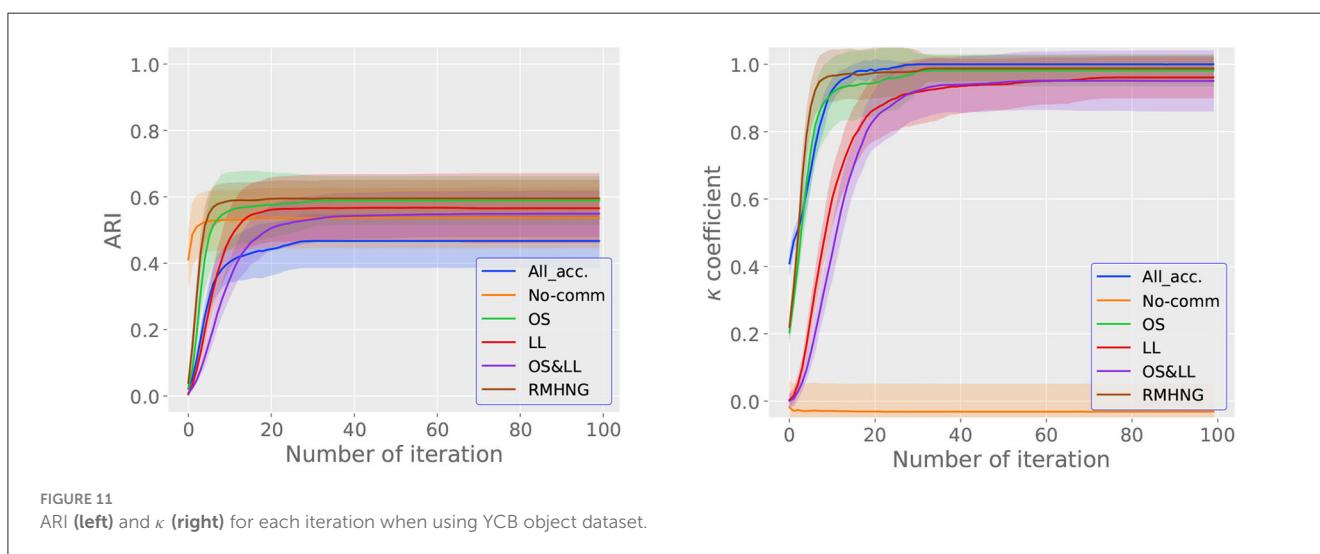
the methods. As for All acceptance, we can see that ARI did not show an increasing trend with iteration in synthetic data, but it did show an increasing trend when using the YCB object dataset. From the right figure in Figure 6, we can see that the RMHNG

method had the fastest convergence of the κ, followed by OS, LL, and OS&LL. No communication did not show any increasing trend compared to other methods. As for All acceptance, we can see that the κ did not show an increasing trend with iteration when using

TABLE 3 Experimental results for YCB object dataset: Each method was tested five times, and for each agent, the ARI and κ were calculated when I was $91 - -100$.

conditon	ARI (Agent 1)	ARI (Agent 2)	ARI (Agent 3)	ARI (Agent 4)	κ
RMHNG	**0.61 ± 0.05**	**0.59 ± 0.05**	**0.59 ± 0.05**	**0.59 ± 0.05**	0.99 ± 0.04
OS	0.59 ± 0.06	**0.59 ± 0.06**	**0.59 ± 0.08**	**0.59 ± 0.08**	0.98 ± 0.05
LL	0.56 ± 0.11	0.57 ± 0.09	0.56 ± 0.11	0.56 ± 0.11	0.96 ± 0.06
OS&LL	0.55 ± 0.07	0.55 ± 0.07	0.54 ± 0.06	0.54 ± 0.06	0.95 ± 0.09
No communication	0.55 ± 0.10	0.50 ± 0.09	0.55 ± 0.07	0.55 ± 0.08	−0.03 ± 0.08
All acceptance	0.47 ± 0.08	0.47 ± 0.08	0.47 ± 0.08	0.47 ± 0.08	**1.0 ± 0.0**
Gibbs sampling	0.63 ± 0.05				-

The mean ± standard deviation of obtained $50(5 \times 10)$ARI and κ are shown. Highest scores are shown in bold, and second-highest scores are underlined.

FIGURE 11
ARI **(left)** and κ **(right)** for each iteration when using YCB object dataset.

synthetic data, but it did show an increasing trend when using the YCB object dataset.

4.2.3. Decentralized posterior inference

Figure 12 shows the results of calculating the degree of similarity between the distribution of the sign obtained by each method and that obtained by Gibbs sampling in the last 10 iterations (91–100 iterations) for each method. RMHNG showed a value of 0.76, indicating that the distribution of the sign obtained by RMHNG matched that obtained by Gibbs sampling by 76%. Among the methods that approximated RMHNG, OS showed the highest value, both in the synthetic data experiment and the YCB object dataset experiment. Additionally, all approximation methods showed higher values than No communication.

Comparing Tables 2, 3 and Figures 6, 11, the synthetic dataset showed a more significant improvement in both ARI and κ coefficient values than the YCB object dataset when RMHNG and its approximations were applied. Importantly, Gibbs sampling and RMHNG improved clustering performance primarily when agents had partial observations that, when integrated, improved clustering. The dataset used in Experiment 1 was deliberately created to satisfy this condition, resulting in pronounced

FIGURE 12
Percentage agreement with Gibbs sampling.
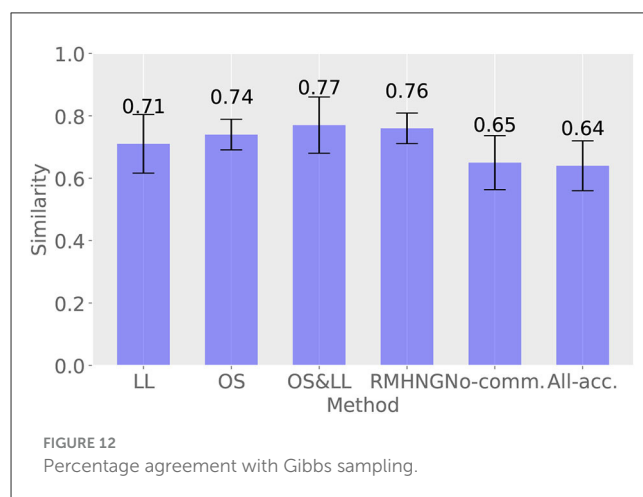
differences in the ARI and κ coefficient metrics. However, in Experiment 2, which used real image data, not all objects fulfilled this condition. The difference between RMHNG and No Communication narrows when an agent's partial observations are sufficient for clustering or when integrating these partial

observations does not improve clustering performance. For example, items such as the Mustard bottle, Bleach cleanser, and Windex bottle, which have similar feature distributions (as shown in Figure 10), were often grouped under a single category even when the partial observations were integrated. However, the primary goal of the experiments is to demonstrate RMHNG's ability to integrate observations and effectively sample from the posterior distribution, such as $P(w|x^1, x^2, x^3, x^4)$, in a decentralized manner. While the ARI metric provides a comparative framework against human labels, which is considered an intuitive and reasonable clustering benchmark, it does not directly evaluate the accuracy of sampling from the posterior distribution. In this sense, it is important to note that RMHNG could produce sampling results more similar to Gibbs sampling than No communication, as shown in Figure 12.

5. Conclusion

In this study, we extended the MHNG to the N-agent scenario by introducing the RMHNG, which serves as an approximate decentralized Bayesian inference method for the posterior distribution shared by agents, similar to the MHNG. We demonstrated the effectiveness of RMHNG in enabling multiple agents to form and share a symbol system using synthetic and real image data. To address computational complexity, we proposed two types of approximations: OS and LL approximations. Evaluation metrics, such as the ARI and the κ, were used to assess the performance of communication in each iteration of the naming game. Results showed that the 4-agent naming game successfully facilitated the formation of categories and effective sign-sharing among agents. Moreover, the approximated RMHNG exhibited higher ARI and κ compared to the No communication condition, showing that the approximate version of RMHNG could perform symbol emergence in a population. Additionally, we assessed the agreement between the sign distributions obtained by RMHNG and Gibbs sampling, confirming that RMHNG approximates the posterior distribution with a degree of agreement exceeding 87% for the synthetic data and 71% for the YCB object data. This result demonstrates that RMHNG could successfully approximate the posterior distribution over signs given every agent's observations.

While RMHNG assumes a communication network that funnels information in a sequential, chain-like, manner, communication network architectures in actual human societies can exhibit notable variations. For example, small-world and scale-free networks serve as models of human relational structures, encapsulating the complex interconnections among individuals (Watts and Strogatz, 1998; Barabási and Albert, 1999). Thus, assessing the validity of the chain-like information dissemination process in RMHNG and determining potential qualitative changes that might occur when RMHNG is adapted to networks embodying structures such as small-world or scale-free configurations remain intriguing avenues for future studies.

Several future perspectives emerge from this study. Firstly, we plan to analyze the behavior of the RMHNG in populations with a larger number of agents. Although we focused on the 4-agent scenario due to the computational cost of the original RMHNG ($O(IDT^{(N-1)})$), we empirically observed that the OS approximation performed well in many cases. Unlike the original RMHNG, the OS-approximated version exhibits scalability in terms of the number of agents ($O(N)$), enabling simulations with larger populations. This scalability opens up possibilities for providing valuable insights into language evolution through the MHNG framework. Additionally, extending the categorical signs to more complex signs, such as sequences of words, represents a natural progression for our research. Investigating the dynamics of communication with more intricate sign systems will shed light on the evolution and complexity of language.

Investigating the cognitive processes involved in the MH naming game is also an important research topic. If symbol emergence in human society is facilitated by decentralized Bayesian inference through communications like RMHNG, the cognitive process underpinning the MH naming game, specifically, the decision-making regarding acceptance or rejection in the MH receiving, becomes crucial. An experimental semiotics study based on the MHNG model was conducted, in which acceptance or rejection decisions of human participants were examined at the behavioral level (Okumura et al., 2023). Exploring the cognitive processes and related brain functions associated with these behaviors is also an avenue for future research.

Data availability statement

The original contributions presented in the study are included in the article/supplementary material, further inquiries can be directed to the corresponding author.

Author contributions

JI designed the study, collected data, conducted the experiment, and wrote the manuscript. TT contributed to the key idea of this study, performed data analysis and interpretation, and contributed to the writing of the manuscript. AT and YH critically reviewed the manuscript and assisted in its preparation. All authors approved the final version of the manuscript and agreed to be accountable for all aspects of the work, ensuring that any questions related to the accuracy or integrity of any part of the work are appropriately investigated and resolved.

Acknowledgments

Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

References

Barabási, A.-L., and Albert, R. (1999). Emergence of scaling in random networks. *Science* 286, 509–512. doi: 10.1126/science.286.5439.509

Bouchacourt, D., and Baroni, M. (2018). "How agents see things: On visual representations in an emergent language game," in *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing* (Brussels: Association for Computational Linguistics), 981–985. doi: 10.18653/v1/D18-1119

Calli, B., Singh, A., Walsman, A., Srinivasa, S., Abbeel, P., Dollar, A. M., et al. (2015). "The YCB object and model set: Towards common benchmarks for manipulation research," in *2015 International Conference on Advanced Robotics* (Istanbul: IEEE), 510–517. doi: 10.1109/ICAR.2015.7251504

Cangelosi, A., and Schlesinger, M. (2015). *Developmental Robotics: From Babies to Robots*. Cambridge, MA: MIT Press. doi: 10.7551/mitpress/9320.001.0001

Chaabouni, R., Strub, F., Altché, F., Tarassov, E., Tallec, C., Davoodi, E., et al. (2022). "Emergent communication at scale," in *International Conference on Learning Representations* (virtual). Available online at: https://openreview.net/forum?id=AUGBfDIV9rL

Chandler, D. (2002). *Semiotics the Basics*. London: Routledge. doi: 10.4324/9780203166277

Chen, X., and He, K. (2021). "Exploring simple Siamese representation learning," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (Nashville, TN: IEEE), 15750–15758. doi: 10.1109/CVPR46437.2021.01549

Choi, E., Lazaridou, A., and de Freitas, N. (2018). "Compositional obverter communication learning from raw visual input," in *The International Conference on Learning Representation* (Vancouver, BC). Available online at: https://iclr.cc/archive/www/doku.php%3Fid=iclr2018:main.html

Christiansen, M. H., and Chater, N. (2022). *The Language Game: How Improvisation Created Language and Changed the World*. London: Transworld Publishers Limited.

Cohen, J. (1960). A coefficient of agreement for nominal scales. *Educ. Psychol. Meas.* 20, 37–46. doi: 10.1177/001316446002000104

Deacon, T. W. (1997). *The Symbolic Species: The Co-Evolution of Language and the Brain. Number 202*. New York, NY: W. W. Norton.

Gupta, A., Lanctot, M., and Lazaridou, A. (2021). "Dynamic population-based meta-learning for multi-agent communication with natural language," in *Advances in Neural Information Processing Systems*. Available online at: https://openreview.net/group?id=NeurIPS.cc/2021/Conference

Hagiwara, Y., Furukawa, K., Taniguchi, A., and Taniguchi, T. (2022). Multiagent multimodal categorization for symbol emergence: emergent communication via interpersonal cross-modal inference. *Adv. Robot.* 36, 239–260. doi: 10.1080/01691864.2022.2029721

Hagiwara, Y., Kobayashi, H., Taniguchi, A., and Taniguchi, T. (2019). Symbol emergence as an interpersonal multimodal categorization. *Front. Robot. AI* 6, 134. doi: 10.3389/frobt.2019.00134

Hastings, W. K. (1970). Monte carlo sampling methods using markov chains and their applications. *Biometrika* 57, 97–109. doi: 10.1093/biomet/57.1.97

Havrylov, S., and Titov, I. (2017). "Emergence of language with multi-agent games: learning to communicate with sequences of symbols," in *Advances in Neural Information Processing Systems* (Toulon), 30. Available online at: https://iclr.cc/archive/www/2017.html; https://openreview.net/forum?id=SkaxnKEYg

He, K., Zhang, X., Ren, S., and Sun, J. (2016). "Deep residual learning for image recognition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (Las Vegas, NV: IEEE). doi: 10.1109/CVPR.2016.90

Hubert, L., and Arabie, P. (1985). Comparing partitions. *J. Classif.* 2, 193–218. doi: 10.1007/BF01908075

Landis, J. R., and Koch, G. G. (1977). The measurement of observer agreement for categorical data. *Biometrics* 33, 159–174. doi: 10.2307/2529310

Lazaridou, A., Hermann, K. M., Tuyls, K., and Clark, S. (2018). Emergence of linguistic communication from referential games with symbolic and pixel input. *arXiv*. [preprint]. doi: 10.48550/arXiv.1804.03984

Lazaridou, A., Peysakhovich, A., and Baroni, M. (2017). "Multi-agent cooperation and the emergence of (natural) language," in *The International Conference on Learning Representations* (Toulon). Available online at: https://iclr.cc/archive/www/2017.html

Lewis, D. (2008). *Convention: A Philosophical Study*. Hoboken, NJ: John Wiley & Sons.

Lin, T., Huh, M., Stauffer, C., Lim, S-. N., and Isola, P. (2021). "Learning to ground multi-agent communication with autoencoders," in *Advances in Neural Information Processing Systems*.

Mu, J., and Goodman, N. (2021). "Emergent communication of generalizations," in *Advances in Neural Information Processing Systems* (virtual). Available online at: https://nips.cc/Conferences/2021

Nakamura, T., Nagai, T., and Taniguchi, T. (2018). Serket: an architecture for connecting stochastic models to realize a large-scale cognitive model. *Front. Neurorobot.* 12, 25. doi: 10.3389/fnbot.2018.00025

Noukhovitch, M. LaCroix, T., Lazaridou, A., Courville, A. (2021). "Emergent communication under competition," in *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems* (Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems), 974–982.

Okumura, R., Taniguchi, T., Hagiwara, Y., and Taniguchi, A. (2023). Metropolis-hastings algorithm in joint-attention naming game: experimental semiotics study. *arXiv*. [preprint]. doi: 10.48550/arXiv.2305.19936

Steels, L. (2015). *The Talking Heads Experiment: Origins of Words and Meanings*. Berlin: Language Science Press. doi: 10.26530/OAPEN_559870

Taniguchi, T., El Hafi, L., Hagiwara, Y., Taniguchi, A., Shimada, N., Nishiura, T., et al. (2021). Semiotically adaptive cognition: toward the realization of remotely-operated service robots for the new normal symbiotic society. *Adv. Robot.* 35, 664–674. doi: 10.1080/01691864.2021.1928552

Taniguchi, T., Nagai, T., Nakamura, T., Iwahashi, N., Ogata, T., Asoh, H., et al. (2016). Symbol emergence in robotics: a survey. *Adv. Robot.* 30, 706–728. doi: 10.1080/01691864.2016.1164622

Taniguchi, T., Nakamura, T., Suzuki, M., Kuniyasu, R., Hayashi, K., Taniguchi, A., et al. (2020). Neuro-serket: development of integrative cognitive system through the composition of deep probabilistic generative models. *New Gener. Comput.* 38, 23–48. doi: 10.1007/s00354-019-00084-w

Taniguchi, T., Ugur, E., Hoffmann, M., Jamone, L., Nagai, T., Rosman, B., et al. (2018). Symbol emergence in cognitive developmental systems: a survey. *IEEE Trans. Cognit. Dev. Syst.* 11, 494–516. doi: 10.1109/TCDS.2018.2867772

Taniguchi, T., Yoshida, Y., Taniguchi, A., and Hagiwara, Y. (2023). *Emergent Communication through Metropolis-Hastings Naming Game with Deep Generative Models*.

Watts, D. J., and Strogatz, S. H. (1998). Collective dynamics of 'small-world' networks. *Nature* 393, 440–442. doi: 10.1038/30918

Appendix

A. Evaluation of the decentralized posterior inference architecture

In order to evaluate the efficacy of RMHNG as an approximate Bayesian estimator for the posterior distribution $p(w \mid x^1, x^2, \ldots, x^N, \theta^1, \theta^2, \ldots, \theta^N)$, a comparison was made between RMHNG and the actual posterior distribution. However, direct computation of the true posterior distribution $p(w \mid x^1, x^2, \ldots, x^N, \theta^1, \theta^2, \ldots, \theta^N)$ presented significant challenges. Therefore, the evaluation focused on the degree of concurrence between the sign distribution generated by RMHNG and that produced by Gibbs sampling. The evaluation was conducted for the last 10 iterations (i.e., $91 - -100$ iterations) of RMHNG and Gibbs sampling. Let $f^R_{d,w}$ and $f^G_{d,w}$ be the number of times the sign w was sampled using RMHNG and Gibbs sampling, respectively, for the d-th dataset. The similarity between the two methods was calculated as $\sum_{d=1}^{D} \sum_{k=1}^{K} \min(f^R_{d,k}, f^G_{d,k})$. However, due to the singularity of the GMM, label switching (i.e., swapping of signs) between different inference results needed to be addressed. To solve this problem, bipartite graph matching was performed to correspond a clustering result with another. To perform bipartite graph matching, the sign obtained by RMHNG was considered as the point set $V^R = \{v^R_0, v^R_1, \ldots, v^R_K\}$, and the sign obtained by Gibbs sampling was considered as the point set $V^G = \{v^G_0, v^G_1, \ldots, v^G_K\}$. The edge connecting V^R_i and V^G_j was denoted by $E_{i,j}$, and the set of all edges was denoted by $E = \{e_{0,0}, e_{0,1}, \ldots, e_{i,j}, \ldots, e_{K,K}\}$. The graph $G = (V^G \vee V^R, E)$ was a complete bipartite graph. If the gain of each pair (v^R_a, v^G_b) was $\sum_{d=1}^{D} \min(f^R_{d,a}, f^G_{d,b})$, then the sign replacement problem could be reduced to a weighted maximum bipartite matching problem. To simplify the problem further, the gain of each pair was multiplied by $(-1) \times \sum_{d=1}^{D} \min(f^R_{d,a}, f^G_{d,b})$. This reduced the weighted maximum two-part matching problem to a minimum cost flow problem, which could be solved using the Hungarian method. Finally, the similarity was calculated by $\frac{1}{10D} \sum_{d=1}^{D} \sum_{k=1}^{K} \min(f^R_{d,k}, f^G_{d,k})$, where $\frac{1}{10D}$ was a normalization factor.

B. ARI and κ

ARI is a widely used measure for evaluating clustering performance by comparing the clustering results with the ground-truth labels. Unlike precision, which is calculated by directly comparing estimated labels to ground-truth labels and often used in the evaluation of classification systems trained using supervised learning, ARI considers label-switching effects in clustering. The formula for ARI is given by Equation (A1), where RI represents the Rand Index. Further details can be found in Hubert and Arabie (1985).

$$\text{ARI} = \frac{\text{RI} - \text{Expected RI}}{\text{Max RI} - \text{Expected RI}} \qquad \text{(A1)}$$

The kappa coefficient (κ) is defined by Equation (A2):

$$\kappa = \frac{C_o - C_e}{1 - C_e}. \qquad \text{(A2)}$$

Here, C_o represents the degree of agreement of signs among agents, and C_e denotes the expected value of coincidental sign agreement. The interpretation of κ is as follows (Landis and Koch, 1977):

1. Almost perfect agreement: $(1.0 \geq \kappa > 0.80)$
2. Substantial agreement: $(0.80 \geq \kappa > 0.60)$
3. Moderate agreement: $(0.60 \geq \kappa > 0.40)$
4. Fair agreement: $(0.40 \geq \kappa > 0.20)$
5. Slight agreement: $(0.20 \geq \kappa > 0.00)$
6. No agreement: $(0.0 \geq \kappa)$

C. Feature extraction by SimSiam

As a feature extractor, we utilized SimSiam (Chen and He, 2021), a self-supervised representation learning technique that was pre-trained on the YCB object dataset. We followed the same network architecture and hyperparameters as outlined in the original paper (Chen and He, 2021), but with a few minor adjustments. For data augmentation, we used the following parameters, using PyTorch notation.[4] RANDOMRESIZEDCROP with a scale in the range of $[0.1, 0.6]$, and RANDOMGRAYSCALE with a probability of 0.2. We normalized tensor images using the NORMALIZE function with a mean of $(0.485, 0.456, 0.406)$ and standard deviation of $(0.229, 0.224, 0.225)$ We used ResNet-18 as the BACKBONE (He et al., 2016) and set the dimension of the output feature vector to 512. A two-layer fully connected layer was employed as the PROJECTOR, using an intermediate layer with a dimension of 512. The predictor also utilized a two-layer fully connected layer, with an intermediate layer dimension of 128. During the training phase, we set the learning rate to 0.1 and utilized stochastic gradient descent as the optimizer. The batch size was set to 64, and we trained the network for 100 epochs.

---

4 PyTorch, Transforming and augmenting images: https://pytorch.org/vision/stable/transforms.html.