

Go by Its Name: Evolution and Analysis of Conceptual Referential Communication

Siyu Yao^{1†}, Joshua Nunley^{2,3‡} and Eduardo J. Izquierdo^{2*}

¹Dept. of History and Philosophy of Science and Medicine, ²Cognitive Science Program, ³Dept. of Informatics
Indiana University, Bloomington, IN 47405 USA
siyuyao@iu.edu

[†]Conceptualization, methodology, software, investigation, writing (original draft and editing), visualization. [‡]Conceptualization, methodology, software, writing (review and editing), visualization. *Supervision and writing (review and editing).

Abstract

Referential communication is a complex form of social interaction that communicates a spatially or temporally distant referent. Previous modeling practices have studied how artificial agents manage to communicate locations that directly determine foraging behaviors. In our study, we introduce conceptual referential communication. In this mode of referential communication, communicated information can lead to behaviors that change flexibly to suit the environment. Instead of giving specific behavioral instructions, this mode only communicates a label of the desired referent, the location of which is unknown to both the sender and receiver. This requires the signal receiver to adjust its foraging behavior based on its own exploration of the environment. We evolve artificial dynamical agents that can communicate 2 and 3 different labels and successfully forage the target label in changing environments. We found that a typical strategy to communicate and differentiate labels in our experiments is by varying the numbers and lengths of contacts between the agents. We also identify several ways in which the receiver develops inter-neurons that differentiate and store information both from communication and the environment.

Introduction

A noticeable portion of human and animal communication is referential: communication is not just about something directly present "here" and "now", but it often refers to objects or events that stand in a temporal or spatial distance from the time and place in which communication takes place. One of the most delicate examples of animal referential communication is the waggle dance of honey bees (Frisch et al., 1967; Chittka, 2022). A forager honey bee who has discovered a food source returns to the hive and performs an eight-shaped dance to recruit fellow worker bees to forage the same food source. The angle of the dance relative to a fixed reference reveals the direction of the food source, and the duration of the waggle run indicates the distance to the food source. Other forms of animal referential communication are also extensively studied in bumble bees (Dornhaus and Chittka, 1999, 2005), birds, and mammals (Townsend and Manser, 2013). Taking a primitive form in animals while presenting an "aboutness" characteristic of human language, referential communication has the potential to teach important

lessons about the evolution of language structures and the emergence of social interaction.

Multiple studies in the artificial life field have modeled referential communication. As a first proof of concept, Williams et al. (2008) propose a minimal sense of referential communication based on the autopoietic theory developed by Maturana et al. (1980) and an earlier attempt to reconceptualize communication (Di Paolo, 1997). According to Williams et al. (2008), referential communication is a coordinated behavior between a sender, who has access to a referent, and a receiver who does not. Communication as a coordinated behavior can be any form of interaction that shapes the future behavior of those agents in such a way: (1) "the future behavior of a receiver should be constrained by its interaction with a sender", (2) "the nature of the receiver's constrained behavior should vary based on properties of the referent", and (3) "the communicative interaction should have a degree of separation from the referent" (Williams et al., 2008, p.703). Dispensing with any preconceived channels and structures of communication, viewing communication as a coordinated behavior is especially suitable for answering questions about how communicative behaviors emerge from non-communicative ones, and how any structures or strategies of communication emerge under bodily and environmental constraints without explicit human design. Williams et al. (2008) successfully evolved agents that can communicate exact locations on a ring by patterns of movement in an area departed from the target location.

Later studies tend toward increased interest in the receiver. Beyond questions about how the communication is done, researchers become more interested in how information is actually stored and used by the receiver, and whether the receiver can play a more active role than rigidly following senders' instructions. For example, Manicka (2012) adds a temporal distance between communication and searching and analyzes the storage of the content communicated in receiver inter-neurons using information theory. Campos and Froese (2017) incorporate a switch of roles between agents and apply dynamical analysis to the agents. Moreover, Fox and Bullock (2023) necessitate bidirectional

communication between two agents, each of whom has access to a unique piece of information that it needs to communicate to the other.

Despite these attempts, existing studies have not modeled and analyzed how the receiver can employ communicated information in an environment-specific way *after* communication. In real-world communication, after an agent receives a piece of information, it often can preserve the information and use it in flexible ways. Communication and its behavioral consequences happen in a brain-body-environment (BBE) system: how the receiver behaves after communication is co-determined by the information communicated, the environment setting, the receiver's own states, and the receiver's interaction with the environment. This flexible use of information has not been sufficiently studied. Fox and Bullock (2023) and Campos and Froese (2017) explored the BBE nature *within* communication by requiring bidirectional interaction and flexible sender/receiver role assignment under different environmental constraints. Nevertheless, in their studies, adjustments to the environment occur mostly at the beginning of communication. After the initial switch of strategies, no flexibility is shown in the agents' behaviors after communication.

In our study, we explore how communication of the same thing can lead to different behavioral results, as the receiver goes into varied environments afterward and explore them to complete the task. To do this, we conceptualize and model a different mode of communication from earlier studies. We call this new mode of communication **conceptual referential communication**.

Despite profuse debates about the nature of concepts, most agree that concepts are abstractions that unite a group of occurrences, enabling subjects to discriminate those occurrences with internal representations or external behaviors (Margolis and Laurence, 2022). By conceptual referential communication, we mean that what is communicated is a label that stands for one or a group of referents; the label suffices for distinguishing the referent(s) from the rest of the environment, but it does not indicate the exact behavior to find the referent. For example, the odor or shape of a flower are labels that stand for it, but they do not suggest the exact location of the flower for a faraway forager. Likewise, in human communication, one may simply ask a friend to find a bookstore without specifying how to do it. Following the coordinate behavior approach, we distinguish conceptual from non-conceptual communication by behavioral consequences instead of the grammar of signals. Conceptual referential communication is characterized by the **flexibility** of consequent behavior under environmental changes and the **active role** of the receiver in determining its behavior based on its own interaction with the environment. The receiver stores communicated information, perceives environmental cues, and then actively adjusts its behavior based on both pieces of information. The combination of both pieces of informa-

tion gives another motivation to call this type of communication conceptual, as it pertains to the manipulation of internal representation beyond using physical signs.

Conceptual referential communication has counterparts in nature and presents unique benefits. Communicating specific behavioral instructions is not always beneficial because, first, there can be multiple referents with the same label, and approaching any of them would complete the task. Second, as the environment changes, referents with the same label can emerge, disappear, or occur in different places. Due to this variability, specificity does not imply efficiency and conceptual referential communication becomes a competitive alternative. Remarkably, the communication among bumble bees can be considered as conceptual. Unlike honeybees that specify the target location, bumble bees inform recruited foragers only about the odor of the target flower, so that they may bring nectar of the same flower type from a variety of spatial locations (Dornhaus and Chittka, 1999). This difference is explained by a trade-off between communication and foraging efficiency under different environments in the bees' evolutionary history. Honey bees used to live in tropical areas where the flowers are mostly clumped. Therefore, even if wagging dance costs a significant amount of time to communicate, it leads foragers directly to the food source. In contrast, bumble bees live in temperate habitats where flowers are more distributed. The lack of specificity in communication is compensated by swift exploration of the environment and the use of each individual's memory of it (Dornhaus and Chittka, 1999).

We design a task to model conceptual referential communication. The task consists of 2 stages. In the communication stage, the sender receives a command that indicates the target label and communicates the label to the receiver. Next, in the courier stage, the receiver is released into the environment. The receiver can identify how sectors in the environment are labeled, and its task is to find and stay in the sector with the target label. We permute the labels of the sectors across trials, forcing the receiver to perceive the environment in order to determine the foraging behavior best suited to finding the referent. This permutation design makes those labels conceptual in the sense that they are the abstraction that unites different foraging positions across trials. We evolve agents to perform the task and gradually increase its complexity by adding more labels. Beyond successfully implementing the task, we also analyze how it is done. By tracking the agents' behavioral and neural activation patterns, we answer questions concerning (1) how the target label is communicated, (2) how that information is stored in the receiver, and (3) how that information is used by the receiver to achieve the task.

Method

Our model is an extension of those in Williams et al. (2008) and Manicka (2012). Two agents, the "sender" and the "re-

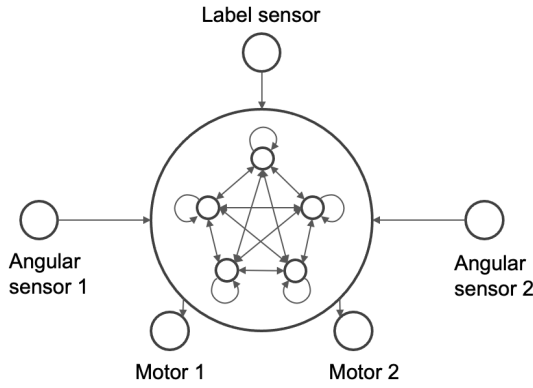


Figure 1: The structure of agents.

ceiver”, dwell on the edge of a ring. We mark their positions by their angular positions, ranging from 0 to 2π . The ring consists of multiple sectors marked with different discrete labels. In our task, we input a command, that is, a target label, to the sender, let the sender communicate it to the receiver, and release the receiver to the environment to find and stay in the sector that matches the target label. The sender receives no information about the environment. Our tasks differ from earlier ones by communicating a piece of conceptual information that does not map to one single behavior, such as moving to one designated position on the ring. As the positions of the sectors in the environment vary across the trials, successful performance of this task requires the receiver to actively perceive the environment and adjust its behavior flexibly to the environment change. There is no one-one mapping between the receiver foraging behavior and the pattern of communication.

The Agents

Each agent contains two angular sensors, one label sensor, 5 inter-neurons, and two motor neurons (see Figure 1). The two angular sensors, one each in clockwise and counterclockwise directions, can sense the distance to the other agent when it comes into the sensing range of $\frac{\pi}{8}$. The value of each sensor is $1 - \text{distance} \cdot \frac{8}{\pi}$, thus falling within $[0,1]$. The label sensor can sense three discrete values, 1, 0, and -1. With this, the sender senses the input command (either 1 or -1), and the receiver senses the label of the sector it is currently in. To prevent the receiver from knowing the environment while communicating and undergoing transient, the regions for communication and for the receiver to start exploring the environment are labeled as 0. In short, the agents communicate by moving and sensing each other with angular sensors. Both sender and receiver receive unique information from their label sensors: the sender knows what the target label is but not how the environment is labeled; the receiver senses the label of its current sector, but it has to know the target label from the sender.

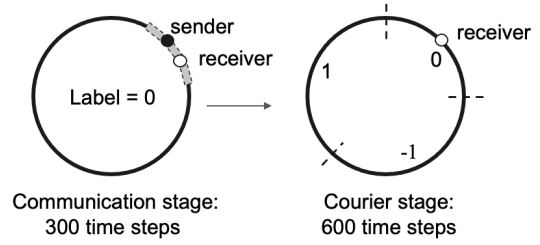


Figure 2: Settings of the 2-label task. Grey rectangles indicate the range of angular sensors. Dashed lines indicate sectors of the environment.

The inter- and motor neurons constitute a continuous-time recurrent neural network (CTRNN) with the following state equation (Williams et al., 2008):

$$\tau_i \dot{s}_i = -s_i + \sum_{j=1}^N w_{ji} \sigma(s_j + \theta_j) + I_i \quad i = 1, \dots, N$$

where s is the state of each neuron, τ is the time constant, w_{ji} is the strength of the connection from the j^{th} to the i^{th} neuron, θ is a bias term, and $\sigma(x) = \frac{1}{1+e^{-x}}$ is the sigmoid activation function that normalizes neuron outputs to the range $[0,1]$. The output of each neuron is $\sigma(s_j + \theta_j)$. I_i represents an external input. All 5 inter-neurons are interconnected. Each inter-neuron is connected to all three sensors, such that $I_i = \sum_{k=1}^3 w_{ki} s_k$, where s_k is the value of the k^{th} sensor. The two motor neurons are not interconnected, but each of them receives a single-directional connection from all the inter-neurons. The speed of the agent is determined by the difference between the outputs of two motor neurons.

The Tasks

Our first task is for the agents to communicate 2 labels. The task consists of two stages (see Figure 2). First, in the communication stage, the two agents are put on a ring labeled as 0, with an initial distance of $\frac{3\pi}{32}$ either clockwise or counterclockwise, so that the two agents are within each other’s sensory region. The agents are free to move on the edge of the ring and cross each other without impediment. The communication stage lasts 300 time steps. Then, in the courier stage, the sender is removed, and only the receiver is placed on another ring, which we call the environment. The environment is split into three sectors: a starting sector labeled with 0 and two sectors labeled with 1 and -1 each. The starting sector has a size of $\frac{\pi}{2}$, and the other two sectors equally divide the rest of the ring. The receiver starts at the center of the label-0 sector and is free to move on the ring for 600 time steps. Over the last 200 time steps, the fitness score is calculated by the proportion of time that the receiver spends in the sector with the target label.

The second task is increased in complexity and requires the agents to communicate 3 labels. To do this, we make a few changes to the 2-label task. We add another sensor to each agent, so that each target label is represented as a 2-bit

code: (1,-1), (-1,1), and (1,1). Label (0,0) is used to mark the communication region and the starting region in the environment. The environment is split into 4 equal sectors, each with one of the above labels.

Evolution

We evolve populations of sender/receiver pairs, instead of evolving each of them separately. Neural parameters are evolved using the evolutionary search algorithm implemented by Candadai (2019). The evolution starts with a population of random genotypes. Based on the overall fitness the agents achieve in our task, a fraction (elitist fraction) of the best-performing agent pairs is retained for the next generation. A new population of solutions is then created by copying the elite agents and adding a mutation noise. The following neural parameters are evolved under corresponding ranges: time constants $\in [1, 11]$, biases $\in [-10, 10]$, and connection weights (from sensors to neurons and between neurons) $\in [-15, 15]$.

Evolving Agents for Conceptual Referential Communication

We first evolve agents for the 2-label task. We use 8 trials to evaluate the overall performance of the task. Each trial is a different combination of (1) 2 starting positions of the communication stage, with the sender lying at $\frac{3\pi}{32}$ clockwise or counterclockwise to the receiver, (2) 2 command labels, 1 or -1, and (3) 2 environmental settings, with the -1- and 1-labeled sectors switching their positions. The overall fitness is calculated by averaging the fitness scores across all 8 trials. We evolve 10 agent pairs with the following evolutionary parameters: population size 400, elitist fraction 0.05, and mutation noise added to offspring 0.05. Each population is evolved for 1500 generations. We found out that all of the 10 populations produce a best agent pair with fitness = 1, which means that their receivers stay in the correct sectors in the last 200 time steps in all 8 trials.

Next, we evolve agents for the 3-label task. The performance of each agent pair is evaluated over 36 trials, that is, the combination of 2 starting positions, 3 commands, and 6 environment settings with varied label positions. This task is much more difficult to evolve than the 2-label version. We set the population size to 1500, the elitist fraction to 0.05, and the mutation noise added to offspring to 0.05. We evolved 30 populations, each going through 8000 generations, and obtained 3 pairs of agents that can successfully perform the task with an overall fitness of 1. The rest agent pairs all obtain a fitness between 0.8 and 0.9.

To test whether agents can genuinely perform the task successfully instead of taking advantage of computational artifacts, we set the time step 10 times smaller than the one used to evolve them. For the 2-label task, 7 of all 10 pairs turn out to fulfill this requirement. All 3 agent pairs in the 3-label task pass the test.

Analyzing Best Solutions in the 2-Label Task

In this section, we aim to answer three questions: (1) How do the agents communicate? (2) How does the receiver store information? (3) How does the receiver use the information to achieve the task in the environment? These questions are important because, first, strategies to communicate discrete labels can show more qualitative differences than those in earlier studies, where agents employ an analog strategy to communicate continuous locations, that is, under the same communication strategy, certain features of communication patterns correlate with the target location (Williams et al., 2008; Manicka, 2012; Campos and Froese, 2017; Fox and Bullock, 2023). Second, in our study, communicated information is not the sole determinant of behavior. It is important to show how stored information affects the agents' reaction to different environments and causes different behaviors therein. We will answer these questions through visual inspection of the neural or behavioral state trajectories over time. A complete information-theoretic analysis of these agents will be performed in future work.

Communication Patterns

In the 7 agent pairs that are properly integrated, the most common strategy, occurring in 5 pairs of agents, is a *broad-casting* strategy: The sender moves in different patterns under the two commands, while the receiver moves in a fixed way on the ring regardless of the sender's behavior. To differentiate the two commands, the sender utilizes two of four possible patterns (Figure 3). (1) The sender moves with the receiver at the same pace with a phase difference larger than the angular sensor range, so that the two are not in contact throughout the whole communication stage. (2) The sender moves around the circle periodically in the opposite direction as the receiver, so that the two cross each other several times during the communication stage. (3) The sender stays relatively still during the communication stage so that the receiver crosses the sender several times as it moves around the circle. (4) The sender for the most part stays close to the receiver so that the two can sense each other for a long period of time.

The four patterns of movement can be further crystallized into several patterns of angular sensor activation. Angular sensors show a "spike" when the agents briefly cross each other. A sensor is continuously activated when they stay close to each other. For pattern (1), the sensors remain inactivated after the small initial "spike". Patterns (2) and (3) produce several "spikes" in the sensors. Pattern (4) produces a long continuous activation in one of the sensors. In the other 2 agent pairs that do not follow the broadcasting strategy, the sender's movements under different commands cause the receiver to also move in different ways, but the patterns of sensory inputs are similar to the broadcasting cases: one pair shows a difference in the number of "spikes", and the other utilizes the contrast between one continuous acti-

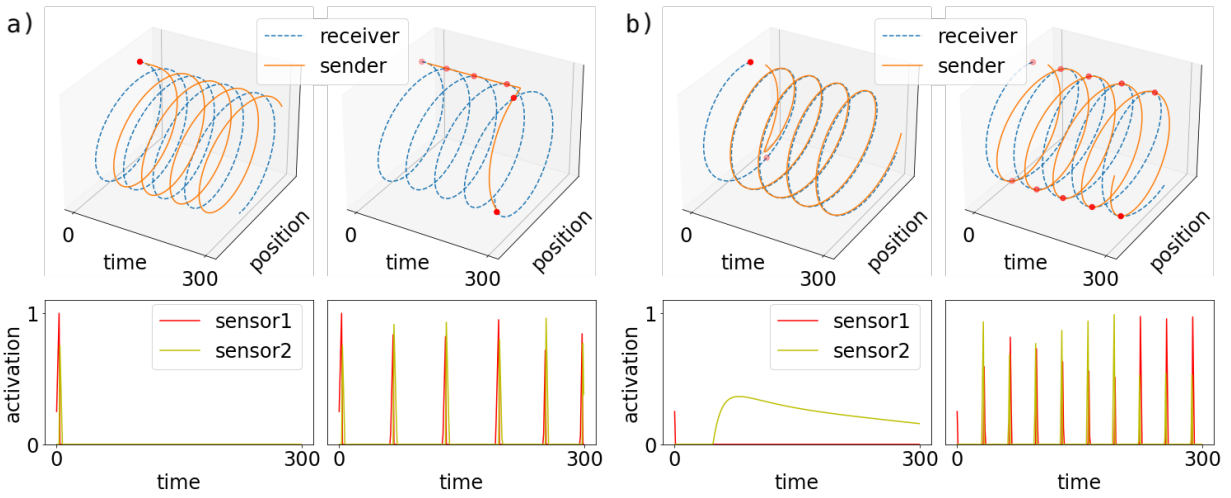


Figure 3: Illustration of four movement patterns from two pairs of agents, followed by the patterns of angular sensor activation. a): Agent pair No.1 under command 1 with movement pattern (1) and under command -1 with movement pattern (3). b): agent pair No.2 under command 1 with movement pattern (4) and under command -1 with movement pattern (2). Red dots mark the point of contact between the two agents.

vation and multiple "spikes". Due to these commonalities among agent pairs, we use one typical agent pair (No.1) as a representative in the following analyses.

Information Storage and Usage

To investigate how communicated information is stored, we look into the activation of inter-neurons during the communication stage. In all receivers, typically there are 2-3 inter-neurons that are directly activated or deactivated by sensory "spikes". Different numbers and patterns of "spikes" play a role in adjusting the degree of inter-neuron activation at the end of the communication stage (as an example, see Figure 4). For a few pairs, activation differences of inter-neurons also lead to a noticeable difference in the state of motor neurons. Because the state of a CTRNN in each time step is determined by its state in the last time step, end-state activation differences in inter- and motor neurons constitute how commands are stored to be used in the courier stage.

The third question we want to answer is how the information communicated is used by the receiver as it explores the environment. The receiver needs to recognize environmental labels and develop ways of reacting to them that are also sensitive to the information stored in communication.

First, to study how inter-neurons react to environmental labels, we look into trials with the same command but different environments. If neuron activation diverges as soon as the receiver enters sectors with different labels, then those neurons play a central role in distinguishing environmental labels.

In each receiver, we discover three types of inter-neuron reactions to the environment. This is observed in 5 of the receivers. Here we take the receiver from agent pair No.1 to

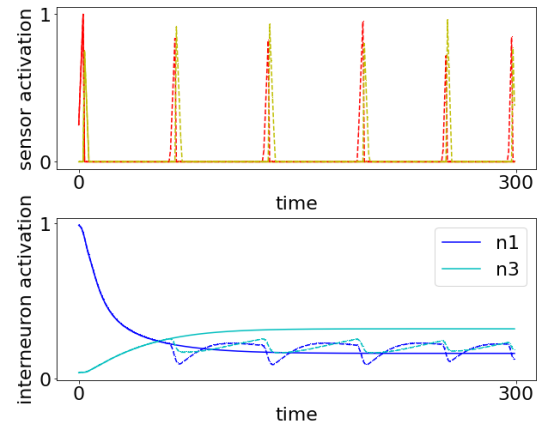


Figure 4: Neural states of receiver No.1. Upper panel: angular sensor activation. Lower panel: activation patterns of inter-neurons 1 and 3. Solid and dashed lines represent situations under two different commands.

illustrate their internal working (Figure 5a). The first type of neural reaction to the environment shows opposite behaviors when encountering the two labels: the neuron's degree of activation increases under one label but decreases under the other (Figure 5a, blue and cyan lines). The second type reacts specifically to one label, and the other label does not affect its activation degree (Figure 5a, red lines). The third type is not directly sensitive to the environment, only changing states slowly in reaction to the state change of other inter-neurons.

When the target label varies, the reaction of inter-neurons to environmental labels remains mostly consistent. Commu-

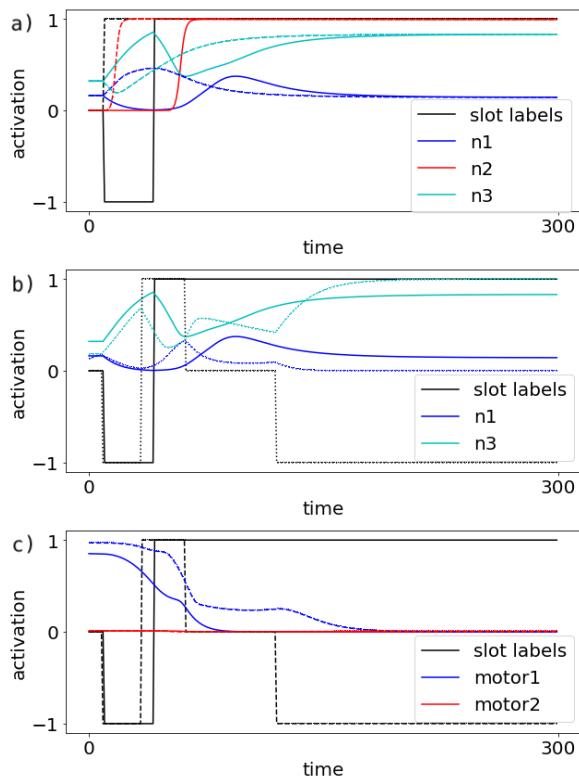


Figure 5: a): Inter-neuron activation of receiver No.1 under two environments, shown by solid and dashed lines. Black lines mark the the agent's current sector label. Blue, red and cyan lines mark 3 inter-neurons. b): Inter-neuron activation patterns of the same receiver under two commands (shown by slid and dotted lines) in the same environment. c): Motor activation of the same receiver under two commands (shown by solid and dashed lines) in the same environment. Only the first 300 time steps of the courier phase are plotted because the neurons remain still in the rest time steps.

nication differences do not affect whether an inter-neuron is activated or deactivated by a certain environmental label but mostly the degree of it (Figure 5b). The different degrees of reaction are caused by the retention of activation differences at the beginning of the courier stage.

Different degrees of activation translates into different duration and rate of motion. For this specific receiver (Figure 5c), under the first command, as the receiver moves into the correct sector, its left motor decreases soon and comes down to a stop. Under the second command, the receiver keeps moving across the incorrect sectors till the correct label brings its left motor down to a stop. Initial activation differences alter the initial activation degree of the motor and the degree of inter-neuron reactions to sector labels, making it possible to move across incorrect sectors till inter-neuron activation change is accumulated and brings it to a stop in the correct sector.

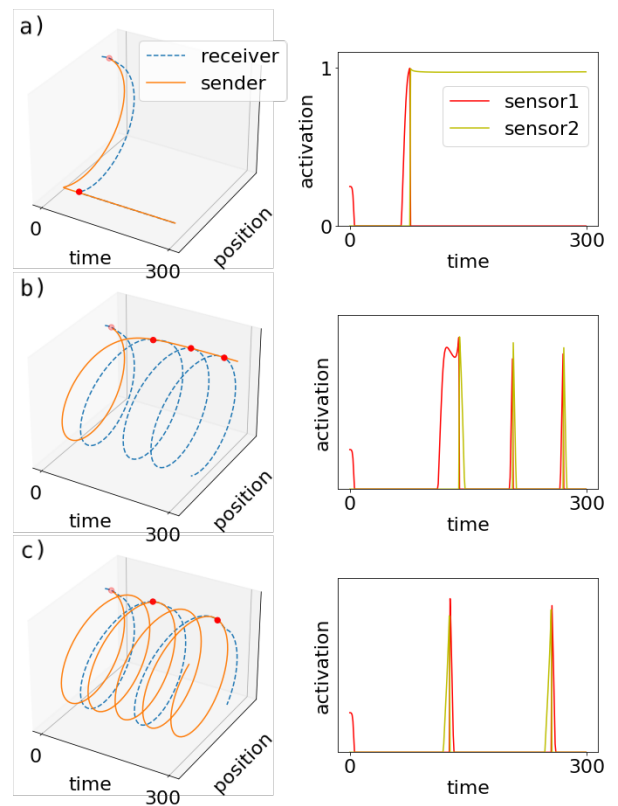


Figure 6: Communication patterns of the first agent pair across three commands, followed by the sensor activation.

Analyzing Best Solutions in the 3-Label Task

For the 3-label task, we aim to answer the same three questions and make comparisons with the 2-label task. Increased complexity can have important implications, because the simple 2-label task may not be able to elicit the most efficient neural structure and strategy for this type of task. For example, the minor inter-neuron activation differences after communication look like an expedient solution when the agent is required only to differentiate two labels. Making the task more complex by adding more labels can reveal more diverse patterns of communication and more efficient internal structures to employ communicated and environmental information.

Communication Patterns

Unlike most agents in the 2-label task, receivers in the 3-label task do not show a fixed behavior during the communication stage across all three commands. Instead, their movements vary in reaction to the difference in the sender's moving patterns. This indicates a *non-broadcasting* strategy: communication of different information is contributed not just by the solo dance of the sender, but by the coupled movement between the sender and the receiver. Despite the non-broadcasting communication patterns, the 3-label

agents do share with the 2-label ones how commands are differentiated within the receiver. Commands are again differentiated by the number and shape of sensory "spikes" triggered by the agents crossing or accompanying each other. In the first of the three pairs, command (1,-1) maps to a long fixed activation of the right sensor, command (1,1) maps to 3 "spikes", and command (-1,1) maps to 2 "spike"s (Figure 6). In the second pair, the three commands map respectively to 3 "spikes", 7 "spikes", and 2 major "spikes". For the third pair, one pattern is 5 discrete "spikes". For the other two labels, the agents are constantly close to each other while crossing each other several times, resulting in a series of alternative activation of left and right sensors.

Information Storage and Usage

In contrast to the typical case of 2-label tasks, where command differences map to small activation differences in a few inter-neurons at the end of the communication stage, different commands in the 3-label task are stored in a more distinctive way. All of the receivers contain several "specialized" inter-neurons that are maximally activated only under one command and remain quiet under others. Other "non-specialized" inter-neurons also distinguish the three commands by remarkable distinctions in their end states. For example, in the first receiver, inter-neurons 1 and 3 are specialized so that 1 is maximally activated by command (1,1) and 3 is maximally activated by (1,-1) while both remain inactivated under the rest of the commands. Inter-neurons 2 and 4 show three degrees of activation under the three commands. Inter-neuron 5 is insensitive to command differences. At the end of the communication stage, inter-neuron activation comes into three equilibrium states for three commands. This suggests that each target label draws the receiver into a distinct attractor.

We found that inter-neuron activation states after communication indicate a structure to store information similar to binary coding. In Figure 7, we plot a three-dimensional phase space, each axis representing the activation of a chosen inter-neuron. The equilibrium states can be located in distinct places in the phase space. For both the first (Figure 7a) and second receivers (Figure 7b), different commands can be sufficiently distinguished by activating or deactivating one or more inter-neurons. For example, for the first receiver, the three labels are projected in this space as (0,0,0), (1,0,0.7), and (0,1,1); for the second receiver, the labels are projected as (0,0,0.4), (1,0,0.3), and (1,1,0). This suggests that the receiver evolves a structure by which the commands are stored in a form close to a series of binary codes. The storage of information in a binary code fashion instead of small activation differences also explains why a non-broadcasting strategy is more likely to be formed in their movements. Large differences in activation are easier to pass to motors and result in qualitatively different receiver movement patterns.

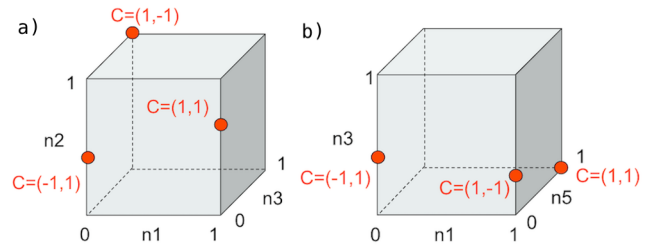


Figure 7: Inter-neuron states of receivers No.1 (a) and No.2 (b) at the end of the communication stage under three command labels (C).

With regard to information usage, similar to the 2-label tasks, each receiver contains several inter-neurons that react directly to environmental labels. They can either react specifically to one label and remain deactivated under the other two, or react to all three labels to different degrees. For example, for the first receiver, inter-neurons 2 and 4 are directly reactive to environmental labels, while 1 and 3 are entirely specialized for distinguishing commands, remaining insensitive to the environment. Both inter-neurons 2 and 4 are activated under (-1,1) and deactivated in different degrees under the other two environmental labels. For the second pair of agents, inter-neurons 1 and 3 are activated by (1,-1), inter-neuron 4 is activated by (1,1), and the rest two inter-neurons are not sensitive to environmental labels. Note that reacting to commands and to environmental labels can co-exist in one inter-neuron, only that command differences may alter the degree of reaction to certain environmental labels.

Unlike the 2-label tasks, where the motor is brought down gradually, the motors of 3-label agents react faster to bring them to a stop when they enter the correct region. This is likely because, as stated earlier, the inter-neurons of the 3-label agents react to different labels in more drastic ways (full vs. no activation) compared to the 2-label agents. Fig 8 shows how one receiver under the same command moves in two environments. In environment 1, the first sector that the agent moves into is the correct one (black solid line), and the two motors soon come down to stop the receiver. In environment 2, the receiver crosses 2 wrong sectors (after the black solid and dashed lines) and finally entered the correct one (the dotted line). Motor adjustments happen drastically as soon as the receiver enters a sector.

Discussion

Previous models of referential communication make agents communicate information that maps to a fixed behavior. Once the communication is performed, the receiver is destined to move toward the exact location commanded by the sender. This approach does not capture certain modes of communication in nature where only a label, instead of a specific behavioral instruction, is communicated. We call

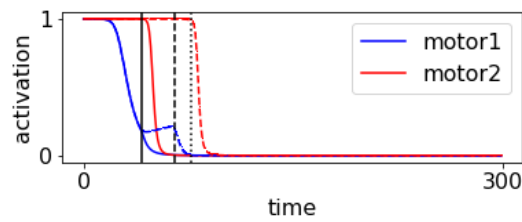


Figure 8: Motor activation during the courier stage. The three black lines indicate the boundary of three sectors. Red or blue solid and dashed lines indicate two environments. As the receiver settles in the right sector very soon, only the first 300 time steps are plotted.

this type of communication conceptual referential communication. Here, the receiver plays an active role in deciding how it behaves based on both the information it obtained from communication and its interaction with the environment. In this paper, we design a task that requires the communication of conceptual information and the flexible use of it by the receiver. The task’s complexity can be gradually increased by adding more labels to communicate. We evolved simple embodied dynamical agents that can communicate and differentiate 2 or 3 labels and complete the foraging task in changing environments. This constitutes a proof of concept for modeling conceptual referential communication.

There are three key insights we learned from analyzing the behaviors and neural dynamics of agents in the 2- and 3-label task. First, a typical strategy for communicating and differentiating labels in a circular space is by combining periodic movement and staying still. These behavioral patterns can trigger different numbers and lengths of sensory “spikes”, which then adjust the activation states of interneurons. For the easier 2-label task, a broadcasting strategy is more favorable, where the receiver moves in a fixed pattern and the sender changes its movement under different labels. However, agents adopt a non-broadcasting strategy in the 3-label task, which can be explained by more drastic inter-neuron reactions to labels. This suggests that when communication leads to more distinct internal state changes, this may also complicate the communication behavior and make the communication patterns a result of mutual coupling.

Second, receivers evolved several inter-neurons that serve to distinguish and store communicated information, and for the more complicated 3-label task, they even learn to differentiate and store commands in a manner close to binary coding. Information is stored in the end-state of 2-3 interneurons after communication. In the simple 2-label task, command differences result in small differences in the degree of activation of those neurons. However, when 3 labels are communicated, we found that each command maps to a

unique combination of activated or deactivated neurons.

Third, to identify environmental labels, receivers evolve inter-neurons that are activated, deactivated, or not affected by different environmental labels. Stored as the initial neuron states of the courier stage, communicated information influences the extent of those reactions. This combination of communicated information and environmental cues controls the motor to continue moving or promptly stop in a sector.

This study has a few limitations, and many remaining questions are worthy of future study. First, it is only a first proof of concept and involves only behavioral analysis of the agents. Complete information-theoretic analysis of them still remains to be done. Second, our circular environment and task design can be a source of constraint to the agents’ behaviors. Future works remain to be done to test whether similar conclusions can be reached with a different task design. Moreover, our task is simple compared to real animal or human communication, and a more complicated environment or task design may reveal further interesting conclusions. For example, the task can be extended to a two-dimensional environment, so one can expect more complicated patterns of communication beyond circular movement. The receiver may also show more complex strategies for exploring the environment beyond simply going in one direction to cross all the sectors. Another possible extension is to distribute one label into multiple sectors with varying sizes or numbers. One may also include more receivers. These extensions would enable studying the trade-off between specificity and efficiency in communication and further test the evolutionary explanation for the communication difference between honey bees and bumble bees (Dornhaus and Chittka, 1999). Finally, it would be interesting to vary the time of the communication stage. Now the pattern of communication is evolved with time as a constraint, but human or animal communication does not take the whole of the allotted time to happen. Instead, they stop communicating when the information is sufficiently addressed. Taking this variation into evolution may help to explore more succinct patterns of communicative behaviors and explore the distinctions between communicative and non-communicative behaviors once the former is finished.

Acknowledgements

This research was supported in part by Lilly Endowment, Inc., through its support for the Indiana University Pervasive Technology Institute. We thank three anonymous reviewers for constructive comments. We also thank Randall Beer for valuable feedback on early results of this project. This work was supported in part by NSF Grant 1845322.

Data Availability

The code used in this project and exemplar agents can be found here: <https://github.com/SiyuYaoGY/Conceptual-Referential-Communication/>.

References

- Campos, J. I. and Froese, T. (2017). Referential communication as a collective property of a brain-body-environment-body-brain system: a minimal cognitive model. In *2017 IEEE Symposium Series on Computational Intelligence (SSCI)*, pages 1–8. IEEE.
- Candadai, M. (2019). Stochastic search. <https://github.com/madvn/stochsearch>. Accessed: 2023-05-05.
- Chittka, L. (2022). *The mind of a bee*. Princeton University Press.
- Di Paolo, E. A. (1997). Social coordination and spatial organization: Steps towards the evolution of communication. In *Fourth European Conference on Artificial Life*, volume 4, page 464. MIT Press.
- Dornhaus, A. and Chittka, L. (1999). Evolutionary origins of bee dances. *Nature*, 401(6748):38–38.
- Dornhaus, A. and Chittka, L. (2005). Bumble bees (*bombus terrestris*) store both food and information in honeypots. *Behavioral Ecology*, 16(3):661–666.
- Fox, R. and Bullock, S. (2023). Nectar of the bots: Evolving bidirectional referential communication. *Adaptive Behavior*, 31(1):65–86.
- Frisch, K. v. et al. (1967). *The Dance language and orientation of bees*. Harvard University Press.
- Manicka, S. (2012). Analysis of evolved agents performing referential communication. In *ALIFE 2012: The thirteenth international conference on the synthesis and simulation of living systems*, pages 393–400. MIT Press.
- Margolis, E. and Laurence, S. (2022). Concepts. In Zalta, E. N. and Nodelman, U., editors, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Fall 2022 edition.
- Maturana, H. R., Varela, F. J., Maturana, H. R., and Varela, F. J. (1980). Problems in the neurophysiology of cognition. *Autopoiesis and Cognition: The Realization of the Living*, pages 41–47.
- Townsend, S. W. and Manser, M. B. (2013). Functionally referential communication in mammals: the past, present and the future. *Ethology*, 119(1):1–11.
- Williams, P. L., Beer, R. D., and Gasser, M. (2008). Evolving referential communication in embodied dynamical agents. In *ALIFE*, pages 702–709.