An emergent communication framework for honeybee waggle dance

Nabila P.R. Siregar[*1], Phong Le[2], and Raquel G. Alhama[3]

[*]Corresponding Author: nabilasiregar.ac@gmail.com
[1]Informatics Institute, University of Amsterdam, the Netherlands
[2]School of Computer Science, University of St Andrews, Scotland
[3]Institute for Logic, Language and Computation, University of Amsterdam, the Netherlands

We present a computational framework to investigate whether artificial agents in a foraging environment develop a communication protocol with features of the honeybee waggle dance. Analyzing communicative success, compositionality, and emergent semantics, we find that (i) agents tend to use directional signals for communication, especially when they have the freedom to explore the environment, (ii) in the absence of direct distance representations, agents develop a system based on shortest-path routes when their mental map representations overlap, and (iii) constraining their communication channel to simulate kinematic properties of honeybees provides a slight positive benefit. Overall, we find that, under certain constraints, agents develop a (relatively) compositional communication system reminiscent of the waggle dance.

1. Introduction

Social animals rely on complex communication systems to support coordination between individuals to handle survival tasks as a group; for instance, foraging. A particularly interesting example is the communication system of honeybees (*Apis mellifera*), which consists of a dance by which the honeybee's body movements encode the location and quality of food (Frisch, 1954); in particular, by means of the so-called *waggle dance*, a forager honeybee encodes a food source's *direction* and *distance* from the hive through movement. The former is encoded with the angle of the waggle, while the latter is indicated with its duration. In addition, dance vigor and repetition can signal resource quality, while odor cues hint at food identity (Dong, Lin, Nieh, & Tan, 2023).

The correspondence between value ranges of sign and meaning (angle for orientation and duration for distance) is arguably an iconic property of this code (Pleyer, Perlman, Lupyan, Reus, & Raviv, 2025). Interestingly, the waggle dance exhibits language-like properties which are rare across species, such as *displacement* (the ability to refer to entities which are not visible) and (bare) compositionality, i.e. the use of constituents (distance, direction) to create meaningful complex expressions (Hockett, 1960; Gil, 2024). These properties make this system particularly interesting from an evolutionary standpoint.

Here, we leverage a computational framework to investigate what the minimal conditions are under which artificial agents in a foraging environment develop a communication protocol exhibiting features of the honeybee waggle dance. To that aim, we base our work on neural-network-based emergent communication (EC) (Lazaridou & Baroni, 2020). We design a game that represents the food-foraging environment of honeybees. In such a game, a *sender* agent is tasked with composing a message to refer to the location of a 'food' landmark, while a *receiver* agent has to identify that landmark from a set including also a number of distractor landmarks, using only the message. The agents are rewarded when they achieve successful communication, i.e. when the receiver correctly identifies the location of the 'food' landmark in its own map.

We use our framework to ask the following questions. First, given the displaced nature of honeybee communication, honeybees need to rely on their mental representations to communicate about a referent which is not in their immediacy. However, we do not know how much overlap exists in each honeybee's representation of their foraging area. Therefore, we ask *What are the effects of mental representation overlap?* In our experiments, we explore different assumptions with regards to the agent's navigation environment; particularly, we test different levels of overlap in the representation of each agent's knowledge of the environment.

Second, given that honeybees rely on distinct kinematic parameters to encode each semantic component, we ask: *to what extent should a communication channel incorporate movement-based constraints analogous to those in honeybee communication?* We experiment with channel constraints by restricting the output range to capture the circular trajectory characteristic of the waggle dance.

We run our model under the conditions specified by each research question and analyze their respective effects on the emergent communication protocol, focusing on communicative success, compositionality, and emergent semantics. Results suggest that representation overlap is key for our agents to develop a successful communication system; however, further analyses reveal that our agents did not learn to compute direct distance. Channel constraints yield only a small positive effect on compositionality, but not all simulations converge to using the channel in the intended manner. Still, overall, we find that agents develop a (relatively) compositional communication system reminiscent of the waggle dance when they share mental map representations.

2. Method

Our game assigns each artificial agent a directed graph representing routes from the nest to several landmarks, including a food source. Each edge encodes the distance and direction from the source to the target node. This representation simulates a sequence of landmarks, aligning with prior proposals on honeybee navigation (Collett, Harland, & Collett, 2002; Chittka, Kunze, Shipman, & Buchmann, 1995; Gould, 1990).

During the game, the sender transmits two *real-valued* tokens to the receiver, who must then identify the food node among a set of candidate nodes in its own graph. To investigate the effect of directional precision in communication, we discretize (bin) the angles representing directions, analogous to the coarse directional encoding observed in honeybee communication (Okada et al., 2014).

We implement each agent as an independent neural network, in which a Relational Graph Convolutional Network (RGCN) (Schlichtkrull et al., 2018) encodes each node into a vector representation. The sender concatenates the representations of the nest and food nodes to parameterize a bivariate Gaussian distribution, from which it samples two real-valued numbers as message tokens. Upon receiving the message, the receiver encodes it into a vector representation. Each candidate landmark is represented by concatenating the receiver's own nest representation with that of the candidate node. The receiver then compares these landmark representations with the message representation to identify the landmark with the highest similarity score. The two agents are jointly trained to maximize the log-likelihood that the receiver correctly identifies the food node.

Data generation We generate 24000 games for training and 6000 games for testing. In each game, the agents' graphs share the same set of nodes (without coordinates). For each agent, a graph is randomly sampled from one of the following structural types: (i) a single-route tree with a direct nest–food edge; (ii) multiple routes from the nest to the food node, with at least one direct nest–food edge; (iii) multiple routes from the nest to the food node, where a direct nest–food edge may or may not exist; (iv) a single-route tree from the nest to the food node, without a direct nest–food edge. Every graph sample contains 10 nodes.

We test four conditions for the edges (see Figure 1 for examples): in the 'same-graph' condition, all agents share the same set of edges, reflecting perfect representational alignment. In the 'independent-graph-v1' condition, edges are randomly generated for each agent independently, capturing individual variability in path encoding. In the 'independent-graph-v2' condition, we ensure no overlap between edges in sender and receiver (except for case (i)). In the 'different-graph' condition, we ensure no overlap between edges in sender and receiver.

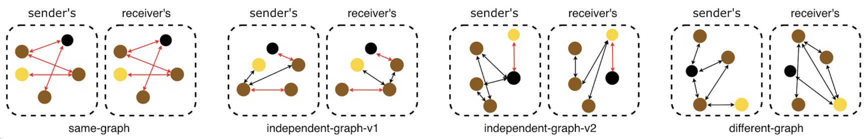

Figure 1. Input representations. Yellow nodes are food; Black are nests; Brown are distractors. Red edges indicate overlap between the sender's and receiver's graphs.

Channel Constraints We investigate two types of message generation. In the *unconstrained* setting, the bivariate Gaussian distribution is left fully parameterized without restrictions. In the *constrained* setting, the covariance matrix of the bivariate Gaussian distribution is diagonal. The first token is normalized to the range $[0°, 360°)$ and the second token to $[0, 1]$. The constrained setting aims to incorporate the natural range of movement of a honeybee, and has the potential to introduce an inductive bias for channel specialization and iconicity.

3. Results

3.1. *Communication*

First we analyze whether our artificial agents learn a communication protocol that allows them to communicate successfully. As can be seen in Figure 2, in all conditions agents achieved communication success well above chance. Representational alignment has a strong impact: agents with perfectly aligned path representations ('same-graph') outperform others with partial overlap by 20%. Constraining the communication channel has a small effect on the final accuracy of communication; however, with this analysis we do not know whether the emerged protocol is shaped differently depending on channel constraints, since languages with different structure may result in similar accuracy of communication.

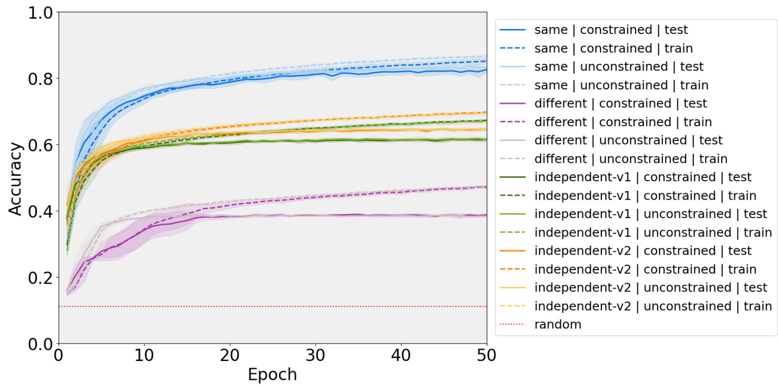

Figure 2. Communication Accuracy (over five runs with different random seeds). The horizontal red line indicates the accuracy of randomly guessing the location of the food node.

3.2. *Compositionality*

Analogous to compositionality in the waggle dance—where the signals for distance and direction are statistically independent—we investigate whether the communication emerging in our agents also exhibits independence between tokens. To assess this, we compute the mutual information (MI) (Shannon, 1948) between the two tokens: higher MI indicates greater shared information and thus lower compositionality. As shown in Table 1, MI is consistently lower for higher overlap between sender and receiver's maps, suggesting that the emerged protocol is more compositional when there is higher representational alignment. One exception is the ‘different-graph’ setting, which on the surface appears to be perfectly compositional; however, as shown in Figure 2, this condition results in the least effective communication, hence it is unlikely to occur in nature (Menzel et al., 2005).

Table 1. MI between tokens. Configurations are sorted by amount of edge overlap. Numbers show the mean and standard deviation (in brackets) over five runs with different random seeds.

	same	independent (v1)	independent (v2)	different
Constrained	0.130 (0.023)	0.329(0.058)	0.911 (0.181)	0.021(0.007)
Unconstrained	0.290(0.019)	0.325(0.154)	0.970 (0.266)	0.018(0.008)

3.3. *Semantics*

Agents develop protocols with some degree of compositionality, but are the messages encoding distance and direction? Decoding artificially emerged languages is not trivial: looking at the messages does not readily allow us to conclude what meaning they convey. However, we can analyze the collection of messages and find to what extent each token carries information that is consistent with hypothesized semantic components. Hence, we report the MI between each token and the distance and direction for the corresponding sender’s graph. Although bees use direct distance (from a coordinate space), agents may opt for alternative distance encodings. We base our analyses on two hypothetical encodings: direct distance and shortest connected path in the sender’s graph.

We focus on the conditions that yielded higher compositionality, which are those with higher representational alignment. We first analyze the unconstrained condition with some degree of graph overlap (independent-graph-v1) for sender and receiver. In Figure 3 we see that direction appears to be unequivocally captured (as evidenced by high MI), and is encoded in every token. However, the same cannot be said about distance. Direct distance is clearly disregarded, while shortest path is encoded to some extent, but overshadowed by direction. In this experimental condition, the shortest path is not always useful to the receiver (as this path may not exist in the receiver’s map), hence its communicative advantage is limited. Results for the constrained version of the model are qualitatively similar.

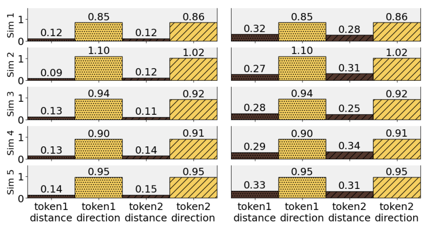

Figure 3. MI between tokens and semantic components (direction, distance) for unconstrained independent-graph-v1. Left: direct distance. Right: shortest path.

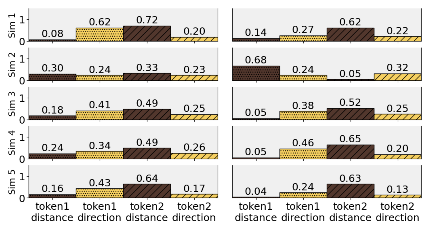

Figure 4. MI between message tokens and semantic components (direction, shortest-path) for same-graph. Left: unconstrained channel. Right: constrained channel.

When both agents use the same graph, as can be seen in Figure 4, the first token appears to be specialized in representing direction in several runs, with much lower correlation to distance. The second token reflects distance much more than the first, although it remains somewhat correlated to direction as well (albeit less so than the first token). Unexpectedly, the pattern observed under the constrained channel condition (Figure 4, right) is less stable than under the unconstrained channel (Figure 4, left).

We hypothesize that channel constraints interfere with learning dynamics, leading to instability. Nevertheless, it is clearly visible from the plot that if one token is better at one component (distance or direction), the reverse is true for the other token, suggesting that different tokens preferentially encode distinct aspects of meaning. This component-specific encoding is reminiscent of the waggle dance, where angle encodes direction and duration encodes distance.

4. Conclusion

We used artificial agents and a signaling game simulating the foraging environment of honeybees. This is undoubtedly a simplified system, lacking the advantages conferred by the biological constraints and large-group dynamics of honeybee societies. Yet, equipped with these minimal tools, our simulations show encouraging results. The key result of our simulations is that, given the absence of a cognitive mechanism for computing direct distance, artificial agents find a way to develop a (relatively) compositional communication system —which is a challenge for EC systems (see e.g., Peters et al. (2025), Carmeli, Belinkov, and Meir (2024), Lazaridou and Baroni (2020))— exhibiting semantic properties reminiscent of the honeybee waggle dance.

To develop a language akin to honeybee's waggle dance, agents need to learn to compute direct distance from graph-based maps. This ability however did not emerge spontaneously in our system, possibly due to a lack of an appropriate architectural mechanism. In our experiments, agents develop different communication strategies depending on the degree of overlap in their mental map representations. When the overlap is imperfect, the agents rely primarily on directional signals (although shortest-path is also moderately captured). When the overlap is complete, however, the shortest-path information is more prominently incorporated in a relatively compositional communication protocol.

Our future work therefore includes: (i) integrating into the agent architecture a mechanism for computing direct distances, analogous to that used by honeybees (Menzel et al., 2005), (ii) investigating the impact of more fine-grained mental map overlaps to reflect the fact that honeybees appear to be extensively familiar with their foraging areas (Wang, Mach, Chen, & Menzel, 2025; Degen et al., 2015; Capaldi et al., 2000), and (iii) analyses to determine whether the emerged code exhibits iconicity (i.e. using with analogous range mapping).

Acknowledgements

R.G.A. was financed by NWO SSH-OC-XS grant (406.XS.24.01.128).
N.P.R.S. was supported in part by MSc Computational Science programme at the University of Amsterdam.

References

Capaldi, E. A., Smith, A. D., Osborne, J. L., Fahrbach, S. E., Farris, S. M., Reynolds, D. R., Edwards, A. S., Martin, A., Robinson, G. E., Poppy, G. M., et al.. (2000). Ontogeny of orientation flight in the honeybee revealed by harmonic radar. *Nature*, *403*(6769), 537–540.

Carmeli, B., Belinkov, Y., & Meir, R. (2024). Concept-best-matching: Evaluating compositionality in emergent communication. In L.-W. Ku, A. Martins, & V. Srikumar (Eds.), *Findings of the association for computational linguistics: Acl 2024* (pp. 3186–3194). Bangkok, Thailand: Association for Computational Linguistics.

Chittka, L., Kunze, J., Shipman, C., & Buchmann, S. (1995). The significance of landmarks for path integration in homing honeybee foragers. *Naturwissenschaften*, *82*(7), 341–343.

Collett, M., Harland, D., & Collett, T. S. (2002). The use of landmarks and panoramic context in the performance of local vectors by navigating honeybees. *Journal of Experimental Biology*, *205*(6), 807–814.

Degen, J., Kirbach, A., Reiter, L., Lehmann, K., Norton, P., Storms, M., Koblofsky, M., Winter, S., Georgieva, P. B., Nguyen, H., et al.. (2015). Exploratory behaviour of honeybees during orientation flights. *Animal Behaviour*, *102*, 45–57.

Dong, S., Lin, T., Nieh, J. C., & Tan, K. (2023). Social signal learning of the waggle dance in honey bees. *Science*, *379*(6636), 1015–1018.

Frisch, K. von. (1954). The Language of Bees. In *The Dancing Bees: An Account of the Life and Senses of the Honey Bee.* Vienna: Springer Vienna.

Gil, D. (2024). Bare and Constructional Compositionality. *International Journal of Primatology*, *45*(3), 635–669.

Gould, J. L. (1990). Honey bee cognition. *Cognition*, *37*(1), 83–103.

Hockett, C. F. (1960). The Origin of Speech. *Scientific American*, *203*(3), 88–97.

Lazaridou, A., & Baroni, M. (2020). Emergent Multi-Agent Communication in the Deep Learning Era. *arXiv preprint arXiv:2006.02419*.

Menzel, R., Greggers, U., Smith, A., Berger, S., Brandt, R., Brunke, S., Bundrock, G., Hülse, S., Plümpe, T., Schaupp, F., et al.. (2005). Honey bees navigate according to a map-like spatial memory. *Proceedings of the National Academy of Sciences*, *102*(8), 3040–3045.

Okada, R., Ikeno, H., Kimura, T., Ohashi, M., Aonuma, H., & Ito, E. (2014). Error in the Honeybee Waggle Dance Improves Foraging Flexibility. *Scientific*

Reports, *4*(1), 4175.

Peters, J., Puiseau, C. Waubert de, Tercan, H., Gopikrishnan, A., Carvalho, G. A. Lucas de, Bitter, C., & Meisen, T. (2025). Emergent language: a survey and taxonomy. *Autonomous Agents and Multi-Agent Systems*, *39*(1), 1–73.

Pleyer, M., Perlman, M., Lupyan, G., Reus, K. de, & Raviv, L. (2025). The 'design features' of language revisited [doi: 10.1016/j.tics.2025.10.004]. *Trends in Cognitive Sciences*.

Schlichtkrull, M., Kipf, T. N., Bloem, P., Berg, R. van den, Titov, I., & Welling, M. (2018). Modeling Relational Data with Graph Convolutional Networks. In A. Gangemi et al. (Eds.), *The Semantic Web: 15th International Conference, ESWC 2018, Heraklion, Crete, Greece, June 3–7, 2018, Proceedings* (Vol. 10843, pp. 593–607). Springer, Cham.

Shannon, C. E. (1948). A Mathematical Theory of Communication. *Bell System Technical Journal*, *27*(3), 379–423.

Wang, Z., Mach, J., Chen, X., & Menzel, R. (2025). Waggle-dance-recruited honeybees expect landscape structures. *Current Biology*, *35*(20), 4922-4931.e2.