The emergence and evolution of a referential code in populations of bee-like agents

Grzegorz Chrupała*

Abstract

Communication requires a shared code, and any change to it must be coordinated between senders and receivers to avoid a breakdown of communication. The honeybee waggle dance illustrates this problem: species with horizontal combs point directly at a food source, while species with vertical combs cannot point directly and instead reference the dance to gravity, decoded against the position of the sun. We model the emergence and evolutionary transition between these two codes in populations of bee-like agents, with selection acting at the level of colonies. In a horizontal-comb model, we find that direct pointing evolves readily when food is moderately hard to find by random search alone, whether because sites are few and large or many and small, and fails when food is too sparse to spark dances or so abundant that it is found without signaling. Adding an exogenous benefit for vertical combs, we then find that the transition to the gravity-referenced code is driven mainly by the mutation rate and the magnitude of this benefit, with the coupling between sender and receiver mutations playing a further role at low mutation rates. Given a favorable confluence of these factors, the transition proceeds reliably and without a breakdown of communication.

1 Introduction

Communication in animal species takes many forms, from the chemical signals of ants to the symbolic, compositional language of humans. Each of these communication systems has its own idiosyncratic properties. What ecological and evolutionary pressures shaped these systems has been of substantial interest to biologists and cognitive scientists. In general, communication assumes some kind of shared code between senders and receivers, and the emergence of such a code requires a degree of alignment between these two roles. Likewise, any change or modification to the code needs to be coordinated between senders and receivers if the breakdown of communication is to be avoided.

*Code, configurations, and result files to reproduce every figure and table in this paper are available at `https://github.com/gchrupala/bees`. AI coding agents, primarily Claude Code and Codex, were used in developing the codebase and as a writing aid in the preparation of this manuscript. The author reviewed all content, and takes full responsibility for it.

In many cases it is possible to think of a plausible pathway for the coordinated emergence and development of a shared code, but a verbally formulated scenario leaves many details underspecified and makes it hard to be precise about the specific conditions that make code change trajectories feasible. For this reason computational models of the evolution of communication systems can provide insights which are hard to obtain by other means. Such models have been used extensively to study the emergence of communication and language (Section 2). In this paper we model a transition of a particularly intriguing communicative system in invertebrates: the waggle dance of honey bees.

Communication is referential when signals stand for entities or states in the world rather than merely expressing the signaler's internal state. The communicative code of *Apis* species von Frisch (1967) is a unique referential system where food sites discovered by foraging scouts are indicated to other members of the colony by means of a figure-of-eight dance whose axis encodes the direction towards the food source, and whose duration stands for the distance to it. In this study we focus on the signaling of direction, as its evolution illustrates some interesting questions about coordination between participants in a communicative system. Different species of *Apis* vary in the specific dance variant used to encode the direction of the food patch: in species with exposed horizontal combs (*Apis florea* and *Apis andreniformis*) the bees orient the axis of the waggle dance directly towards the food patch. In other members of the *Apis* genus, combs are vertical and in some species also hidden in cavities. Such combs make direct pointing impossible, and these species use gravity as a reference point, and map the azimuth of the food patch relative to the sun onto the angle of the waggle dance with respect to the vertical. Figure 1 illustrates these two versions of the waggle dance.[1]

The direct-pointing variant is iconic: the form of the signal—the direction of the waggle axis—mirrors the direction of the referent. The ancestral, direct-pointing dance has been characterized as an enactment of the foraging flight, in which the waggle run imitates the departure toward the food (Barron and Plath, 2017). The gravity-referenced variant is more abstract, relying on an indirect mapping of the sun's azimuth onto an angle measured from the vertical. The waggle dance thus offers an opportunity to study how a referential code can shift from an iconic to a more abstract form, in a relatively simple system.

The distribution of comb orientation and waggle dance code throughout the *Apis* phylogeny (Figure 2) suggests that the horizontal comb and direct pointing are the ancestral state, extant in the stem *Micrapis* subgenus. Vertical combs and the gravity-referenced code are derived traits, present in the other two *Apis* subgenera. Additionally, species within the *Apis* sensu stricto subgenus are cavity nesters. This pattern indicates that the gravity code was linked to vertical combs and that cavity nesting is not required for its emergence (Barron and Plath, 2017).

[1]This is a somewhat simplified overview: in individual species there may be considerable behavioral flexibility in how the dance is performed: e.g. in *Apis mellifera*, sometimes bees dance on a horizontal platform at the entrance to the beehive and in this case they use direct pointing, even though this species generally employs the gravity code (Esch, 2012).

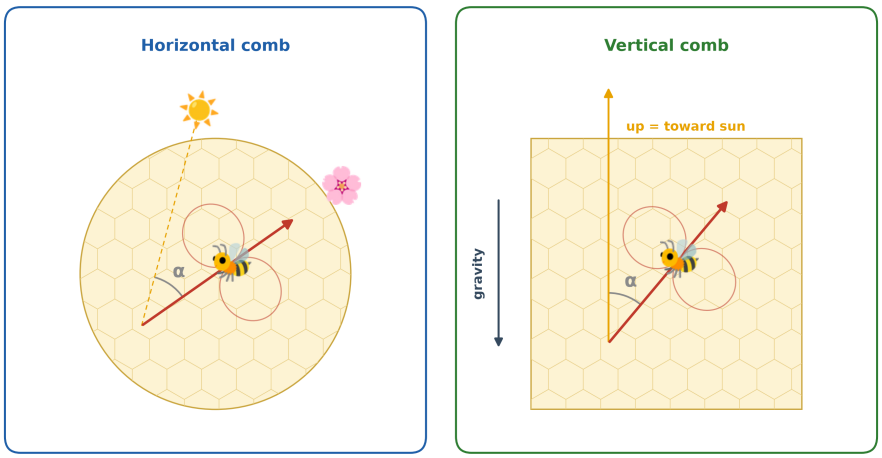

Figure 1: The two variants of the honeybee waggle dance. On a horizontal, exposed comb (left) the bee aims the straight waggle run directly along the true direction to the food, which it can identify as being at a particular azimuth relative to the sun. On a vertical comb (right) the bee cannot directly point at the food. Instead, it uses gravity as a reference: it maps the direction *up* to the position of the sun, and it rotates the dance axis away from vertical by an angle α equal to the food site's azimuth relative to the sun. The figure-of-eight loops and the central waggle run are schematic and not to scale.

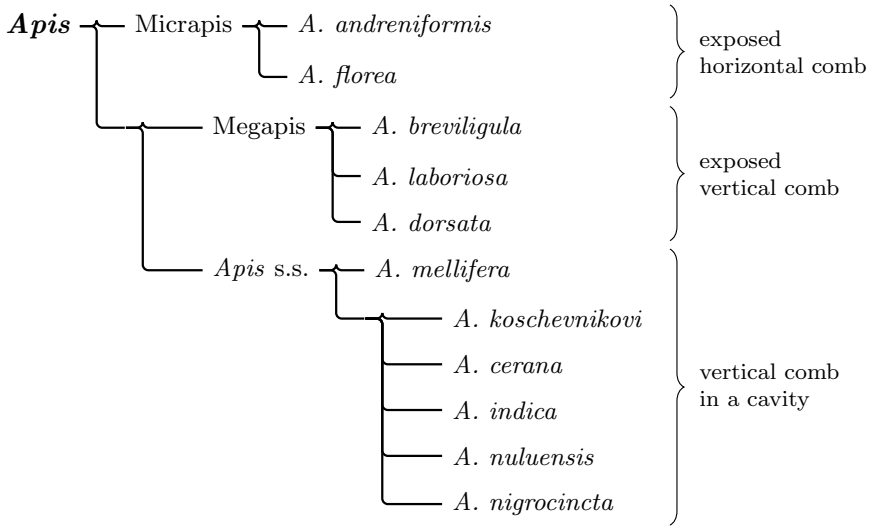

Figure 2: A simplified illustration of the *Apis* phylogeny, with the three subgenera *Micrapis*, *Megapis* and *Apis* sensu stricto. Adapted from Lo et al. (2010)

Nevertheless, the exact sequence and nature of evolutionary events that led to the emergence of the gravity-based waggle dance is not known in detail. It is also not clear how the transition from one version of the code to the other could happen gradually, while maintaining communicative success at all stages, and how this process interacted with the change in the orientation of honeycombs. Because the code is shared, such a transition poses a coordination problem: a change adopted by senders but not matched by the receivers' interpretation, or the reverse, disrupts communication, so that intermediate stages risk a loss of communicative success. A mechanistic account of this transition has been proposed at the neural level: the central-complex circuitry that represents the bee's directional heading (its own compass orientation) is not tied to any particular spatial reference frame, so a gravity reference could in principle substitute for the celestial one without a dedicated switching mechanism (Barron and Plath, 2017). Such an account addresses how an individual brain can support either version of the code, but not how a population of senders and receivers stays coordinated while the shared code itself evolves, which is the question we address.

The present study investigates this transition by means of a computational model of evolving populations of bee-like agents. We model selection at the level of colonies, and individual behavior at the level of worker bees as colony members. We address two principal questions: (i) under what ecological conditions a direct-pointing referential code emerges and stabilizes, and (ii) which factors lead to a gradual transition from this code to its gravity-referenced variant. Our aim is not to reconstruct the precise historical sequence

of these changes, but to delineate the circumstances under which such a transition can proceed without a breakdown of communication.

Our starting point is colonies with horizontal combs, which show some heritable variation in traits controlling the inclination to direct signaling and attending to signals. Here we show that: communication via direct pointing *evolves readily when food is moderately hard to find by random search alone*: i.e. when food sites are few in number but large, as well as when there are many small ones. It fails when food is too sparse to spark dances, or so abundant that it is found without signaling.

In the second stage we introduce extrinsic evolutionary pressure for vertical combs, which partially degrades the direct-pointing code while making the gravity reference available. We show that in this scenario the transition towards the gravity-referenced code is driven mainly by the *mutation rate* and the *magnitude of the exogenous benefit of the vertical comb.* Additionally the *strength of the coupling between the signaling behavior of senders and receivers* also plays a role, especially in the low-mutation-rate regime.

The pattern of results from these simulations indicates that a basic referential code based on an iconic, direct-pointing signal evolves readily given favorable ecological conditions. On the other hand, the transition from this basic ancestral system to the abstract code based on gravity-referenced pointing requires the confluence of several key favorable ecological and evolutionary factors; once these hold, however, the transition proceeds reliably and without a breakdown of communication.

2 Related work

The current study builds on several distinct strands of work. Firstly, it draws on research into the emergence and evolution of communication in artificial multi-agent systems, viewed as an abstract computational problem. Secondly, because our work is inspired by and grounded in real honeybee societies, it also builds extensively on models and simulations of bee communication. These studies are methodologically diverse, employing agent-based models, neural models, evolutionary computation, or combinations thereof.

2.1 Communication in agent societies

A long line of computational work has studied how communication systems can emerge and stabilize in populations of interacting agents. The signaling game introduced by Lewis (1969) provides the formal foundation for much of this work, framing a convention of meaning as a solution to a recurring coordination problem between a sender and a receiver. Skyrms (2010) shows how such signaling conventions can arise without prior agreement through evolutionary and reinforcement-learning dynamics. A distinct route emphasizes cultural transmission: in the iterated-learning framework, linguistic structure emerges as a code is repeatedly learned and reproduced across successive generations of agents, without being built in (Smith et al., 2003). Related formal work models how a protolanguage itself

evolves under selection, showing for instance how combining a small set of signals into words overcomes the error limit that caps a simple signal repertoire (Nowak and Krakauer, 1999). Agent-based models in this tradition have aimed to demonstrate that populations playing repeated language games can self-organize shared, grounded vocabularies and categories, a simulation tradition surveyed by Wagner et al. (2003). More recently, deep multi-agent reinforcement learning has been used to show that discrete communication protocols can emerge when agents are trained on cooperative referential tasks (Lazaridou et al., 2017). Subsequent work has examined when such emergent codes become compositional and generalize to novel meanings (Chaabouni et al., 2020), and how communicative pressure can drive neural agents towards attested language universals (Lian et al., 2023); see Lazaridou and Baroni (2020) for a review of this line of research.

2.2 Computational models of the waggle dance

A venerable research tradition asks an ecological question: given that the dance exists, *when does it pay off*? The main tool here is the agent-based foraging simulation. A colony of many simple forager agents follows fixed behavioral rules. Dance recruitment is present as a switch, with recruits either guided by a successful forager's spatial information or left to search on their own, while the ecology (the density, quality, spread, and persistence of food patches, and sometimes colony size) is varied across runs. The outcome of interest is colony-level foraging performance: typically the net energy or nectar collected. Importantly, these models represent neither evolution nor within-lifetime learning: the communication strategy is fixed and the agents do not adapt. What they model is the ecological cost-benefit accounting of a given strategy: how a fixed colony fares, with recruitment versus without, in one environment or another. Using such a model, Dornhaus et al. (2006) find that recruitment benefits a colony mainly when resources are patchy, variable, and hard to find, with the advantage depending on colony size, and Schürch and Grüter (2014) show that the dance's long-term benefits can outweigh its short-term costs. Bailis et al. (2010) pose a more abstract version of the same question, analyzing when positional communication beats reliance on private information. This body of work is the closest in spirit to our own foraging payoff, but it holds the communication system constant and measures its payoff in a fixed environment; it does not model the code itself as an evolving trait.

More recently, the emergence itself of the bee communicative code has also been investigated computationally. Bernard et al. (2023) treat displaced communication as an evolutionary phenomenon: pairs of foraging agents, each controlled by a continuous-time recurrent neural network, undergo experimental evolution over tens of thousands of generations in a one-dimensional circular arena where a receiver must reach one of several sites a sender has perceived as holding food. Their central finding concerns how a displaced signal arises. Although senders could in principle modulate signal amplitude, communication instead evolved from incidental, movement-derived cues (the onset delay and duration of a sender's presence in the shared communication zone), which were subsequently ritualized

Approach	Code adapts by	Level	Question
Communication in agent societies			
Signaling games and RL	payoff dynamics	population	emergence
Iterated learning	transmission	chain	emergence
Deep multi-agent RL	learning	pair	emergence
Models of the waggle dance			
Foraging simulations	(fixed)	colony	payoff
Evolved controllers	evolution	pair	emergence
Learned representations	learning	individual/pair	learnability
This work	evolution	colony	transition

Table 1: Computational approaches to how a communicative code arises or changes, and where the present model sits among them. The columns give how the code adapts (if at all), the level it applies to, and the question each line of work asks. Rows are grouped into general models of communication and models specific to the waggle dance; each row is discussed and cited in the text.

into signals. The more efficient amplitude-based scheme evolved only when the cheaper cue was experimentally suppressed. This model shares our evolutionary framing but selects at the level of individual sender–receiver pairs rather than colonies, and it abstracts away the geometry of the dance: the object of study is how any displaced signal bootstraps from behavior, not how a specific directional code is structured.

Siregar et al. (2026) instead ask what internal representations and channel properties allow a compositional, waggle-dance-like code to be *learned*. A single sender–receiver pair, implemented as graph neural networks operating over map-like graphs of routes between the nest and candidate sites, is trained by gradient descent on a referential game to transmit two real-valued tokens identifying the food node. They ask whether the two tokens factor into independent direction and distance components (compositionality), how this depends on the overlap between the sender's and receiver's mental maps, and whether constraining the channel to the kinematic range of a real dance helps. They find that direction is readily encoded whereas direct distance is not, that representational overlap is the main driver of compositionality, and that channel constraints have only a small effect. This is a within-lifetime study of representation and structure, with no population dynamics or selection. In a similar vein, Portegys (2020) has an individual agent learn to both produce and follow dance-like movements that encode a nectar location, again treating the dance as a code to be acquired within an individual's lifetime rather than shaped by selection. The two differ in what is left open: Siregar et al. (2026) let the code emerge from a referential game and ask what structure results, whereas Portegys (2020) fixes the dance in advance and has a single agent learn to reproduce and respond to it, studying acquisition of a known code rather than the emergence of its structure.

Our study occupies a different point in this space (Table 1). We model the emergence and in particular the evolutionary *transition* between two forms of a referential code at the level of colonies, using a small set of interpretable, heritable traits rather than opaque neural controllers. This lets us treat costs, benefits, and comb geometry as explicit parameters and ask which ecological and evolutionary conditions drive the shift from iconic direct pointing to an abstract gravity-referenced code: a question about the conditions and dynamics of a communicative transition that none of the prior work addresses.

3 Method

Our model is inspired by bee colonies and their waggle dance, but it is simplified in many important ways, as our goal was not to model bee biology in detail but rather investigate scenarios for the emergence and evolution of referential communicative codes, while adopting the case of the waggle dance as a way to ground the modeling in the real world and ensuring that it is evolutionarily plausible.

We model evolution at the level of colonies, and behavior at the level of individual bee-like agents. This treatment is motivated by the reproductive and social organization of real bee colonies, where workers are highly related and sterile, and thus selection acts effectively on heritable colony-level traits rather than on individual workers—this is the standard kin-selection rationale for treating the honeybee colony as a superorganism (Hamilton, 1964; Seeley, 1989). We do not model sexual reproduction: agent colonies reproduce without genetic input from other colonies. Colonies have a number of heritable traits which are transmitted to offspring colonies subject to mutation, and individual agents express these traits with some non-heritable individual variation. In the interest of simplicity and focus, we model the waggle dance as only signaling the direction of a food site: our agents do not indicate distance.

3.1 Colonies, workers, and traits

We keep two kinds of quantity visually distinct throughout: a colony's heritable traits, which evolve under mutation and selection and whose names we set in SMALL CAPITALS, and the model's fixed or swept simulation parameters, which are chosen by the experimenter, referred to by their mathematical symbols, and listed in Tables 3 and 4. A colony is described by a small set of heritable, real-valued traits, listed in full in Table 2. The prose here introduces the traits needed for the basic horizontal-comb model (Section 3.6); the additional traits and geometry required for tilted combs are developed in Section 3.8.

Four traits govern the behavior of a colony's workers. The DIRECTIONAL BIAS b controls how precisely a dance points; the RECEIVER ATTENTION a is the propensity to attend to a dance rather than to search at random; the DANCE PROPENSITY p controls how readily a successful scout recruits others to the patch it just found (Section 3.4); and the FORAY DISTANCE ℓ sets the mean outbound distance a worker travels on a foray. The first three

Symbol	Name	Range	Role
Behavioral (horizontal-comb model, Section 3.6)			
b	DIRECTIONAL BIAS	$[0, 1]$	sharpness of a dance about the food direction
a	RECEIVER ATTENTION	$[0, 1]$	probability of following a dance rather than searching at random
p	DANCE PROPENSITY	$[0, 1]$	readiness to recruit others after a successful foraging attempt
ℓ	FORAY DISTANCE	$[0, D]$	mean outbound distance of a foray
Nest and code choice (vertical transition, Section 3.8)			
γ	COMB TILT	$[0, 1]$	comb inclination, 0 horizontal to 1 vertical
ϕ	COMB ORIENTATION	$[0, P)$	compass heading the comb faces
t_s	SENDER TRANSPOSITION	$[0, 1]$	weight for the gravity code when producing a dance
t_r	RECEIVER TRANSPOSITION	$[0, 1]$	weight for the gravity code when interpreting a dance

Table 2: The heritable, colony-level traits of the model. The first group governs behavior on a horizontal comb (Sections 3.1 and 3.6); the second adds the nest geometry and code-choice traits used in the vertical transition (Section 3.8). Workers express b, a, p, t_s, and t_r with individual variation (Section 3.1); ℓ varies through the per-foray draw (Section 3.3), and the nest traits γ and ϕ are shared by all of a colony's workers. Here D is the maximum search distance and P the orientation period.

are bounded to $[0, 1]$ and the FORAY DISTANCE to $[0, D]$, where D is the maximum search distance, a fixed model parameter (set to $D = 6$ km in all our experiments).

Workers express the colony traits with non-heritable individual variation. A worker i realizes each trait $x \in \{b, a, p\}$ as

$$x_i = \mathrm{clip}\big(x + \varepsilon_i\big), \qquad \varepsilon_i \sim \mathcal{N}(0, \eta^2), \tag{1}$$

where η is a fixed worker-variation scale and clip returns the value to its admissible range. The FORAY DISTANCE ℓ carries no separate worker jitter: within-colony variation in how far workers travel arises instead from the per-foray draw described in Section 3.3. A colony thus behaves as a noisy ensemble of workers sharing the same underlying genotype.

3.2 Producing and interpreting signals

On a horizontal comb the dance points directly at the food: this is the iconic, direct-pointing code. A worker that has just foraged successfully at azimuth α produces a dance whose intended direction is α itself. The emitted signal is drawn from a von Mises distribution centered on α with concentration $\kappa = b_i\, \kappa_{\max}$, plus a small wrapped Gaussian dance noise. The von Mises distribution is the circular analogue of a Gaussian: a signal ψ centered on μ with concentration κ has density

$$f(\psi; \mu, \kappa) = \frac{\exp\big(\kappa \cos(\psi - \mu)\big)}{2\pi I_0(\kappa)}, \tag{2}$$

where I_0 is the modified Bessel function of order zero and κ plays the role of an inverse variance; here $\mu = \alpha$ and $\kappa = b_i\, \kappa_{\max}$. The DIRECTIONAL BIAS b therefore controls the sharpness of the dance: a low bias leads to a near-uniform, uninformative signal, while a high bias gives rise to a tightly concentrated one that reliably indicates α.

A worker that follows a dance reads its direction as a search direction, subject to a small added interpretation noise. Communication is successful when this recovered direction is close enough to a food site to lead the follower to it.

3.3 Foraging environment

A colony is evaluated over a fixed number of independent foraging episodes, and its fitness is averaged over them (Section 3.4). Each episode draws a fresh environment (Figure 3), with patch radii in meters and distances from the nest in kilometers. The number of food sites is drawn from a Poisson distribution with mean n, so episodes vary in how much food they offer. Each site is placed at an azimuth drawn uniformly around the colony and at a distance drawn uniformly from $[d_{\min}, d_{\max}]$, and is a physical circular disk whose radius is drawn independently from a lognormal distribution with median μ_r and log-scale spread σ_r. A forager finds the patch when its straight outbound path crosses the disk, so a fixed patch subtends a smaller angle when farther from the nest, and distant food demands more

accurate dances. Each patch has value v per visit and a finite capacity limiting how many foragers it can supply before being exhausted; capacity scales with the patch area, growing with the square of its radius, so that larger patches hold proportionally more food while the per-visit value stays fixed. Throughout, $v = 1$, the log-scale spread is $\sigma_r = 0.6$, and sites lie no closer than $d_{\min} = 0.75$ km; the median radius μ_r, the capacity at that radius, and the maximum distance $d_{\max}$ vary by experiment (Table 4).

The FORAY DISTANCE trait ℓ (Section 3.1) sets how far workers travel. Each individual foray draws its own outbound length from a Gamma distribution with mean ℓ and a fixed shape, capped at the maximum search distance D. The per-foray draw is the source of within-colony variation in travel distance and gives a realistic spread of short and long forays around the colony mean.

Each episode also fixes the position of the sun, which serves as the external reference for the gravity-based code introduced in Section 3.8. The sun's azimuth is drawn uniformly from an arc of width Δ_{sun} centered on a fixed heading μ_{sun} (in our experiments $\mu_{\text{sun}} = \pi/2$ and $\Delta_{\text{sun}} = \pi$), so it is constant within an episode but varies between episodes. We explain the role of this variable sun azimuth in Section 3.8.

3.4 Foraging and colony payoff

Within each episode (Section 3.3) the colony makes a sequence of foraging attempts. On each attempt a worker is drawn at random; if at least one dance has already been performed in the current episode, the worker follows a randomly chosen dance with probability equal to its attention a_i, reading it out as a search direction, and otherwise searches in a uniformly random direction. The worker then sets out along that direction on a foray whose length is drawn as in Section 3.3 and enters the first food patch on its path, among patches that still have remaining capacity and lie within the foray length; if no patch is reached the attempt fails. The fraction of foraging attempts that reach a patch is their *foraging success*.

Payoff accrues per episode. A successful foraging attempt yields the patch value v, minus a travel cost c_{travel} per unit of patch distance, and may prompt the worker to perform a dance advertising the patch to later foragers. Recruitment is itself an evolved decision: a successful scout dances with probability $1 - (1 - p_i)^C$, where p_i is its DANCE PROPENSITY and C is the patch's remaining capacity after the visit, so an exhausted patch never leads to a dance and a scarcely provisioned one rarely does. A dance, when performed, costs $c_{\text{base}} + c_{\text{cue}}\, b_i$, growing with the precision of the signal. A failed attempt costs travel proportional to the foray length it travelled, and each act of attending to a dance incurs an attention cost c_{att}. Collecting these terms, an episode with successful attempts S (each at distance r_k), the subset $S_d \subseteq S$ of successes that produced a dance (each by a worker with bias b_k), failed attempts F (each with foray length ℓ_j), and n_a dance-following attempts yields raw payoff

$$\pi_{\text{raw}} = \sum_{k \in S} \big(v - c_{\text{travel}}\, r_k\big) - \sum_{k \in S_d} \big(c_{\text{base}} + c_{\text{cue}}\, b_k\big) - \sum_{j \in F} c_{\text{travel}}\, \ell_j - c_{\text{att}}\, n_a, \tag{3}$$

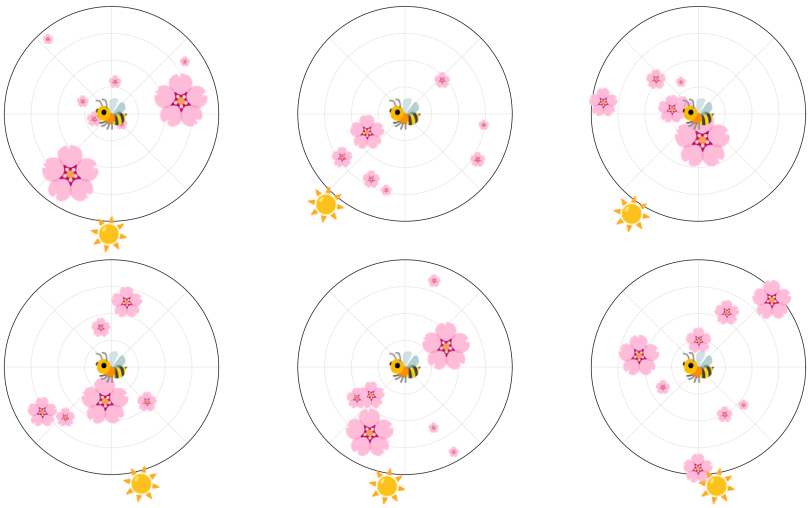

Figure 3: Six samples of the foraging environment. Each panel shows the colony (bee) at the centre, the food patches (flowers) at their drawn azimuths and distances, and the sun, which marks the external reference for the gravity-based code. Each patch is a physical disk, and the flower size scales with the sampled radius of the patch; radial rings mark distance from the colony in kilometers. The panels are drawn at representative values chosen to illustrate the geometry: the number of patches per episode is Poisson with mean 6, so panels differ in how many patches they contain, the median radius is 150 m (lognormal, log-sd 0.6), and patches lie out to a maximum distance of 6 km. All three sit within the ranges the experiments explore (Sections 4.2 and 4.3).

where the dance-cost sum runs over S_d because a success produces a dance only when the recruitment decision is positive. The scaled episode payoff is $\pi = (1 + B\gamma)\,\pi_{\text{raw}}$, where γ is the COMB TILT and $B \geq 0$ the vertical-comb benefit; in the horizontal model $\gamma = 0$, so the scaling is inert. We motivate its form in Section 3.7. A colony's fitness is its mean episode payoff (floored at a small positive value). For diagnostic purposes we also record a *recruitment advantage*: the difference between the foraging success of workers that followed a dance and that of contemporaneous workers that searched at random while dances were available. A positive recruitment advantage indicates that the code is functioning, since following a dance then beats searching blindly.

3.5 Evolutionary dynamics

Evolution proceeds in non-overlapping generations over a fixed-size population of colonies; there is no sexual reproduction. In each generation every colony is evaluated, and the next generation is produced by repeated selection of parent colonies with a probability proportional to fitness. A selected colony's traits are copied subject to mutation. Each trait is perturbed by an independent Gaussian step and then clipped to its range. The step's standard deviation is a fixed fraction σ_m of that range: σ_m for the traits on the unit interval, and $\sigma_m D$ for the FORAY DISTANCE, whose range is $[0, D]$. The workers of each offspring colony are then regenerated from its traits as in Section 3.1. The mutation rate σ_m thus sets the pace of evolutionary change.

3.6 Horizontal comb

The horizontal-comb model is exactly as described above: the comb is flat and the only code available is direct pointing. We use it to investigate the conditions under which a referential code emerges and stabilizes, starting from an initial population with little communicative ability (low DIRECTIONAL BIAS and low attention), so that early foragers signal weakly and rarely attend to one another.

3.7 Modeling code transition

We ask how a colony can move from the direct-pointing code, available on a horizontal comb, to the gravity-referenced code used on a vertical one. We do not model why combs come to hang vertically; instead we impose an exogenous benefit for vertical building through the payoff scaling $(1 + B\gamma)$ (Section 3.4), which rewards greater COMB TILT on its own, and ask whether communication can survive as the comb tilts under that pressure. The scaling is multiplicative because the advantages usually attributed to a vertical comb are architectural,[2] and improve the return on what a colony forages rather than paying a

[2] A vertical comb carries cells on both faces of a single sheet, whereas a horizontal one can open its cells only upward; and a comb attached along its top edge is loaded in its own plane rather than in bending, which wax tolerates better (Seeley and Morse, 1976; Hepburn et al., 2014).

subsidy that arrives whether it forages or not. This couples the two pressures, since tilt is then worth most to a colony that already forages well, and therefore has a working direct code to lose. Because COMB TILT is itself a trait that starts flat and evolves by small mutational steps, a lineage passes through intermediate, partly tilted combs rather than jumping from horizontal to vertical.

Those intermediate tilts are what make the transition non-trivial. As the comb tilts, direct pointing degrades: projecting the food direction onto the sloping surface foreshortens it in a bearing-dependent fashion, leaving less of the direction in the surface to carry the signal. How faithfully the heading is recovered depends on whether the receiver corrects for comb orientation. We thus model two decoders (Section 3.8): a more biologically plausible *flatten* decoder which applies no such correction and is therefore biased on a tilted comb, and an idealized *unproject* decoder that corrects exactly and serves as a bias-free reference against which to read the results of *flatten*.

Meanwhile, the growing tilt makes the gravity-referenced code increasingly reliable, so a colony can in principle switch codes as direct pointing weakens. We again model this switch as gradual, through a blend of the two codes (Section 3.8).

3.8 Dancing on a tilted comb

The tilted-comb model extends the horizontal one with the machinery for comb tilt and the gravity-referenced code. It introduces two further colony-level traits describing the nest: the COMB TILT $\gamma \in [0, 1]$, where $\gamma = 0$ is a horizontal comb and $\gamma = 1$ a vertical one, and the COMB ORIENTATION ϕ, the compass heading the comb faces. Like the behavioral traits these are heritable and mutate by an independent Gaussian step each generation, but they are properties of the nest and so are shared by all of a colony's workers rather than being expressed with individual variation. Because the orientation ϕ is an angle defined on a circle of period P, its mutation step has standard deviation $\sigma_m P$ rather than σ_m.

Comb geometry. The comb is a flat surface on which dances are performed, and its tilt determines which reference directions a dance can use. We represent the comb by its unit normal, the direction perpendicular to its face. When the comb is horizontal this normal points straight up; as the comb tilts, the normal swings down toward the horizon in the compass direction ϕ, until on a vertical comb it lies flat. Writing the tilt as a physical angle $\theta = \gamma\,\pi/2$ (so $\theta = 0$ for a horizontal comb and $\theta = \pi/2$ for a vertical one), the normal has eastward, northward, and upward components

$$\mathbf{n} = \big(\underbrace{\sin\theta\cos\phi}_{\text{east}},\ \underbrace{\sin\theta\sin\phi}_{\text{north}},\ \underbrace{\cos\theta}_{\text{up}}\big). \tag{4}$$

At $\theta = 0$ this reduces to $(0, 0, 1)$, pointing straight up, and at $\theta = \pi/2$ it lies in the horizontal plane pointing along ϕ.

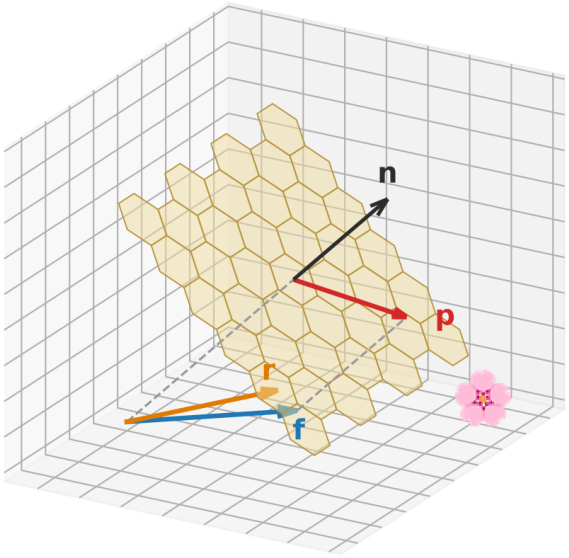

Figure 4: The direct code's projection and the flatten decode's bias. The tilted hexagonal comb and the world-horizontal plane (the grid) share the hive's east–north position. The food direction $\mathbf{f}$ (blue) lies in the world plane; its orthogonal projection onto the comb, $\mathbf{p}$ (red), found by carrying its start and end points along the comb's normal $\mathbf{n}$ (black) until they meet the surface, shown by the dashed rays, has in-surface angle δ_{dir}, the encoded signal. Decoding that signal back to a world heading, the flatten variant recovers $\mathbf{r}$ (orange), which drifts off the true direction by a systematic angle on any tilted comb, whereas the unproject variant would recover $\mathbf{f}$ (blue) exactly. The flower marks the food site.

The two codes. Both codes express the food direction as an angle in the comb surface; they differ in the reference the angle is measured from and in how much of the direction the surface can carry. In the *direct* (iconic) code the worker projects the world direction of the food, $\mathbf{f} = (\cos\alpha, \sin\alpha, 0)$, orthogonally onto the comb and dances along the resulting in-surface vector $\mathbf{p}$ (Figure 4).

In the comb's orthonormal surface axes $\mathbf{e}_1, \mathbf{e}_2$, this projection has coordinates $(p_1, p_2) = (\mathbf{e}_1 \cdot \mathbf{f},\ \mathbf{e}_2 \cdot \mathbf{f})$, and the encoded signal is its angle $\delta_{\text{dir}} = \operatorname{atan2}(p_2, p_1)$. The code's *strength* is how much of the unit direction survives the projection, $s_{\text{dir}} = \|\mathbf{p}\|$. What the projection discards is the part of $\mathbf{f}$ along the comb normal, $\mathbf{f} \cdot \mathbf{n} = \sin\theta\cos(\alpha - \phi)$, leaving

$$s_{\text{dir}} = \sqrt{1 - \sin^2\theta\,\cos^2(\alpha - \phi)}. \tag{5}$$

On a horizontal comb $s_{\mathrm{dir}} = 1$; that is the direction lies entirely within the surface and is represented faithfully. As the comb tilts s_{dir} shrinks, reaching its minimum $\cos\theta$ for the food site in the direction the comb faces ($\alpha = \phi$). Only on a fully vertical comb does this minimum fall to zero: the facing direction then projects to nothing and its heading cannot be recovered, while every other direction stays informative.

In the *gravity-referenced* code the worker instead projects the vertical direction $\mathbf{g} = (0, 0, 1)$ onto the comb to obtain an in-surface reference, and encodes the food azimuth, measured relative to the sun, as an angle from that reference. Its strength depends only on tilt, not on the bearing to the food site,

$$s_{\mathrm{grav}} = \sin\theta = \sin(\gamma\,\pi/2), \tag{6}$$

behaving oppositely to s_{dir}: it is zero on a horizontal comb, where gravity is perpendicular to the surface and gives no in-surface reference, and grows to 1 as the comb becomes vertical. The two trade off: a horizontal comb supports only direct pointing, a vertical comb makes the gravity reference available for every food direction, and at intermediate tilts both carry information.

When code strength is below 1, the direction it recovers is correspondingly degraded, modeled as a circular interpolation between the true angle and a uniformly random one with weight equal to the strength.

Decoding the direct code. The direct encoding is linear in the horizontal food vector: $(p_1, p_2)^\top = M(\cos\alpha, \sin\alpha)^\top$, where M is the 2×2 matrix whose rows are the east–north parts of $\mathbf{e}_1, \mathbf{e}_2$, with $\det M = \cos\theta$. Recovering the world heading exactly requires inverting this map, and we model two receivers that differ in whether they do so (Figure 4). The inversion asks a lot of a receiver: it must know COMB TILT and COMB ORIENTATION, and apply a correction that varies with the direction danced. The uncorrected reading asks nothing, since the dance is a direction in space and the food is on the ground: a receiver that flies along the dance's compass bearing drops the vertical component simply by flying. That reading is also exactly right on a horizontal comb (the ancestral state), which is why we regard it as the more plausible of the two: it leaves the ancestral rule unchanged, whereas the correction would have had no function to select it before combs tilted.

Unproject The receiver applies M^{-1}, exactly inverting the encoding; since $\det M = \cos\theta > 0$ for $\theta < \pi/2$, the true heading is recovered without bias.

Flatten The receiver reads the dance as a direction in space and drops its vertical component, taking the compass bearing of what is left. Where the sender projected the food direction onto the comb, the receiver projects the danced direction back onto the horizontal plane: a second projection rather than an undoing of the first, so the recovered heading $\mathbf{r}$ is systematically biased on a tilted comb.[3]

[3]Formally, the receiver lifts δ to the in-surface vector $\cos\delta\,\mathbf{e}_1 + \sin\delta\,\mathbf{e}_2$ and keeps its east–north part,

The gravity code, by contrast, is always decoded by exact inversion: the projected gravity reference is available to any receiver, so the sun-relative food azimuth is read back directly.

Transposition and blending the codes. Which code a worker uses is governed by two further heritable behavioral traits, the SENDER TRANSPOSITION t_s and RECEIVER TRANSPOSITION t_r (both in $[0, 1]$ and expressed with individual variation as in Section 3.1). These set how much weight a worker places on the gravity-referenced code rather than the direct code when, respectively, producing and interpreting a dance.[4]

Both roles blend the codes by the same rule. Given the angles δ_{dir} and δ_{grav} the two codes assign, and a transposition t, each angle is turned into a unit vector, scaled by its weight, the two vectors are summed, and the blend is the heading of the resultant:

$$\text{blend}(\delta_{\text{dir}}, \delta_{\text{grav}}; t) = \text{atan2}\Big(\textstyle\sum_k w_k \sin\delta_k,\ \sum_k w_k \cos\delta_k\Big), \qquad k \in \{\text{dir}, \text{grav}\}, \tag{7}$$

where the weights combine the transposition with each channel's geometric strength,

$$w_{\text{dir}} = (1 - t)\, s_{\text{dir}}, \qquad w_{\text{grav}} = t\, s_{\text{grav}}. \tag{8}$$

A sender blends the in-surface angles the two codes assign to the food just visited, giving the intended dance direction $\delta = \text{blend}(\delta_{\text{dir}}, \delta_{\text{grav}}; t_{s,i})$, and emits a signal from a von Mises distribution about δ as in Section 3.2. A receiver decodes the observed signal under both codes and blends the two recovered *world* directions the same way, with $t_{r,i}$ in place of $t_{s,i}$.

A worker with $t_i = 0$ therefore relies purely on direct pointing and one with $t_i = 1$ purely on the gravity reference, while the strength factors ensure that a code with no geometric support (such as the gravity code on a horizontal comb) contributes nothing regardless of t. The geometric strength thus enters in two distinct roles: it degrades the world direction each channel recovers, interpolating it toward a uniformly random heading as described above, and it scales that channel's weight in the blend, so a channel with weak geometric support is at once less reliable and less relied upon.

Communication is accurate to the extent that senders and receivers weight the two codes alike: the further $t_{r,i}$ sits from $t_{s,i}$, the more the receiver reads the dance under a mix of references the sender did not use, and the larger the resulting directional error. Alignment only matters where the codes disagree, however, so on a horizontal comb, where $s_{\text{grav}} = 0$ collapses both blends onto the direct code, t_r is free to drift.

which applies $M^{\top}$ rather than M^{-1}. On a flat comb M is a rotation, so the two coincide and the decodes are identical; they differ only once the comb tilts.

[4] The alternative is a binary trait, with a worker either pointing directly or referring to gravity. We use graded weights because they let TRANSPOSITION change by the same small mutational steps as every other trait here. Additionally, there is some biological support for blending: bees put into conflict between a light reference and gravity dance at compromise angles (Edrich, 1977). Blending is nonetheless the more permissive assumption: a binary trait would make mismatched senders and receivers mutually unintelligible, so a population would pay a coordination cost at intermediate frequencies that a blending one avoids.

The per-episode sun azimuth (Section 3.3) matters through this blending. Within an episode the sun is shared by every worker, so when a dancer and follower both use the gravity code it enters encoding and decoding with opposite sign and cancels, and where it falls does not matter. Because it varies across episodes, however, a colony cannot treat it as a fixed offset: a strategy that only partially commits to the gravity code leaves an uncancelled sun term, so the shifting reference selects for genuine transposition rather than a stable blend of the two codes.

Coupling of sender and receiver. A change of code pays off only to the extent that it is matched on the other side, so the two traits may plausibly not mutate independently. We therefore allow the mutations of t_s and t_r to be correlated, treating the strength of the coupling as a parameter. Their mutational steps are drawn as a correlated Gaussian pair whose sender–receiver coupling is the correlation coefficient ρ,

$$\Delta t_s = z_1, \qquad \Delta t_r = \rho\, z_1 + \sqrt{1-\rho^2}\, z_2, \qquad z_1, z_2 \sim \mathcal{N}(0, \sigma_m^2), \tag{9}$$

which leaves the marginal mutation scale of each trait unchanged while tuning how tightly the two co-vary. When ρ is high, a mutation that pushes senders toward the gravity code tends to push receivers by the same amount, so the two sides of the code can shift in a coordinated fashion rather than drifting apart; when $\rho = 0$ they evolve independently. The COMB TILT γ, held fixed at 0 in the horizontal model, is free to mutate in this version.

4 Experimental setup

We study the model of Section 3 in two stages that mirror the two questions of Section 1. The first stage keeps the comb horizontal and asks under what ecological conditions the direct-pointing code is worth maintaining (Section 4.2). The second stage releases the comb and the TRANSPOSITION traits and asks under what conditions a population can migrate from the direct to the gravity-referenced code as an extrinsic pressure tilts the comb (Section 4.3).

4.1 Fixed model parameters

A core set of parameters is held fixed throughout, and only the ecological and evolutionary parameters named below are varied. Table 3 lists the fixed values. The population contains 60 colonies of 80 workers each, evolving for 120 non-overlapping generations; each colony is evaluated over 50 independent foraging episodes of 12 foraging attempts, each attempt made by a single worker drawn at random (the worker pool is therefore a sample of the colony's trait distribution rather than a workforce: its size sets how finely that distribution is discretized, not how much a colony can forage). Workers express colony traits with variation scale η (Section 3.1); dances are performed with maximum concentration $\kappa_{\max}$ and

Symbol / name	Value	Meaning
colonies	60	population size
workers per colony	80	ensemble size per colony
episodes per colony	50	foraging episodes averaged for fitness
attempts per episode	12	foraging attempts per episode
η	0.08	worker-variation scale (trait sd; traits on $[0, 1]$)
$\kappa_{\max}$	14	maximum dance concentration ($\approx 16^\circ$ scatter at $b_i = 1$)
dance noise	0.18	production noise (wrapped Gaussian, rad; $\approx 10^\circ$)
interpretation noise	0.12	receiver decoding noise (rad; $\approx 7^\circ$)
D	6	maximum search distance (km)
$d_{\min}$	0.75	minimum food distance (km)
σ_r	0.6	patch-radius log-scale spread (80% of radii in 70–325 m)
foray shape	2	Gamma shape of per-foray distance (coefficient of variation ≈ 0.71)
c_{base}, c_{cue}	0, 0.02	dance cost terms (in units of v)
c_{att}	0.01	attention cost (in units of v)
v	1	food value (sets the payoff scale)
μ_{sun}, Δ_{sun}	$\pi/2$, π	sun arc center and width (rad; a 180° arc)

Table 3: Parameters held fixed across all experiments. The parameters that some experiment varies are listed separately in Table 4.

are subject to both production and interpretation noise (Section 3.2). These three terms compound: even a colony at maximum DIRECTIONAL BIAS recovers a search direction with a circular standard deviation of about 20°, so the model imposes a floor on communication error and evolved colonies do not approach perfect accuracy. All lengths follow a fixed convention (distances in kilometers, patch radii in meters): the maximum search distance is $D = 6$ km and food patches lie at distances in $[0.75, d_{\max}]$ km. Food patches are physical disks whose radius is drawn from a lognormal and whose capacity scales with patch area, and the number of patches per episode is Poisson (Section 3.3). Each foray draws its outbound distance from a Gamma of fixed shape with the colony's evolved mean. Foraging carries a per-distance travel cost, a dance cost $c_{\text{base}} + c_{\text{cue}}\, b_i$ that is purely bias-dependent ($c_{\text{base}} = 0$), and an attention cost c_{att} (Section 3.4). The sun is drawn from an arc of width π centered on $\pi/2$, and the COMB ORIENTATION is treated as axially symmetric (period π). Colonies start with little communicative ability: the initial DIRECTIONAL BIAS is drawn from $\mathcal{U}(0, 0.15)$, RECEIVER ATTENTION from $\mathcal{U}(0, 0.25)$, the FORAY DISTANCE from $\mathcal{U}(0.15D, 0.45D)$, the DANCE PROPENSITY from $\mathcal{U}(0.8, 1.0)$, and both TRANSPOSITION traits start at 0, so early foragers signal weakly, rarely attend to one another, and use only direct pointing.

The parameters that are not fixed in this way are summarized in Table 4: five ecological parameters (food-site count, patch radius, capacity, maximum food distance, and travel

Symbol / name	Meaning
Ecological	
food-site count	Poisson mean number of patches per episode
μ_r	median patch radius (m, lognormal)
capacity	base loads a patch supplies at the median radius
$d_{\max}$	maximum patch distance (km)
travel cost	cost per km of foray distance (in units of v)
Evolutionary	
generations	evolutionary horizon
B	vertical-comb benefit magnitude
σ_m	mutation step standard deviation
ρ	sender–receiver mutation correlation

Table 4: Parameters that are not held fixed across all experiments, and what each means. No single experiment varies all of them; Sections 4.2 and 4.3 give the values each takes in each experiment, whether swept or held fixed.

cost) and four evolutionary parameters (the number of generations, B, σ_m, and ρ). That table gives only what each parameter means: the values they take, swept or held fixed, are stated with the experiments themselves in Sections 4.2 and 4.3.

4.2 Emergence of direct pointing

The first stage uses the horizontal-comb model of Section 3.6: comb tilt is fixed at 0 and not allowed to evolve, so the gravity reference has zero strength and the TRANSPOSITION traits are inert, leaving only the direct-pointing code in play. Colonies evolve for 60 generations under mutation rate $\sigma_m = 0.07$, and we vary only the food distribution across 50 replicate seeds (400–449). The conditions form a full grid crossing the Poisson mean patch count $\{1, 2, 3, 4, 6, 8, 12, 16, 24\}$ against the median patch radius $\{15, 37.5, 75, 150, 300, 600\}$ meters, so that the ecology ranges from single small patches (essentially undiscoverable) to many large ones (easily found by random search). The remaining parameters of Table 4 are held fixed: a patch at the median radius has capacity 6, patches lie no farther than $d_{\max} = 6$ km, and foraging cost is 0.027 per km travelled.

The primary outcome is the *recruitment advantage* (Section 3.4). Because the follow decision is a randomized coin flip on RECEIVER ATTENTION within the same episodes, it is a near-randomized, contemporaneous estimate of what the dance actually buys, and it distinguishes functioning communication from a directional-bias trait that has merely drifted upward. We report it as a mean over the final generations, alongside final-generation mean DIRECTIONAL BIAS, foraging success, and DANCE PROPENSITY.

4.3 Evolution of the gravity code

The second stage uses the full model of Section 3.8: COMB TILT γ and orientation ϕ are free to evolve, the TRANSPOSITION traits t_s, t_r are active, and the raw payoff is scaled by the vertical-comb benefit $1 + B\gamma$. Populations always start horizontal, so a successful run must evolve both a vertical comb *and* a matched sender–receiver gravity code. The transition is governed by the interaction of the ecology with three evolutionary parameters: the benefit magnitude B, the mutation rate σ_m, and the sender–receiver coupling ρ.

The question we ask in this stage is whether there is *any* combination of ecology and evolutionary regime which lets a population cross between the two codes without a breakdown of communication, and if so which. The space in which to look is both large and costly to cover: eight parameters interact, the values we admit for them (Table 5) already combine into some 2.3 billion distinct settings, and pricing a single setting costs a full evolutionary run on a panel of seeds. We therefore use optimization, which spends its budget where success looks likely instead of spreading it evenly. The search tells us *where* viable settings lie, but its trials cluster around what already worked, so their density cannot tell us how *large* a viable region is. The experiments below therefore re-examine what the search finds on fixed grids. They also re-examine it on fresh random seeds, because a search that maximizes over a panel of seeds will fit that panel's noise along with its signal.

Decode variants. On a tilted comb (Section 3.8) the direct code can be decoded via *flatten* or *unproject*, and we run the experiments below with each variant. The two pipelines are identical in every other respect, so comparing them isolates the effect of the geometric decode assumption on whether and where the transition occurs. Which way that effect should run is not clear a priori. On the one hand, *flatten* is biased on a tilted comb, and this imprecision may push a colony toward the gravity referenced code. On the other hand, *unproject* keeps direct pointing accurate as the comb tilts and so keeps transitional colonies more viable, which may instead ease the transition by keeping them alive long enough to complete it.

Parameter search. We search eight parameters: food-site count, patch radius, capacity, maximum distance, and travel cost, together with B, σ_m, and ρ. Table 5 gives the interval and step size searched for each, and the sampler proposes values on those steps, so the settings form a grid rather than a continuum. One *trial* is one such setting, scored by simulating it as described below; colonies evolve for 120 generations throughout. Trials are proposed by the Tree-structured Parzen Estimator (Bergstra et al., 2011) as implemented in Optuna (Akiba et al., 2019), which fits per-parameter distributions to the best-scoring trials so far and proposes settings typical of them, so that later trials concentrate where earlier ones succeeded; the first 64 trials are drawn at random. We run 1024 trials per decode variant.

Scoring a trial. Each trial evaluates a panel of $J = 10$ seeds. For seed j we write $\bar{\gamma}_j$, $\bar{t}_{s,j}$ and $\bar{t}_{r,j}$ for the final-generation population means of the COMB TILT and the two TRANSPOSITION traits, and m_j for the lowest foraging success (Section 3.4, averaged over the colonies) reached at any generation of the run. A seed is *stable* if $\bar{\gamma}_j \geq \gamma^*$ and $\min(\bar{t}_{s,j}, \bar{t}_{r,j}) \geq t^*$, and *non-viable* if $m_j \leq m^*$. The first two thresholds mark what we count as a completed transition: $\gamma^* = 0.80$ asks for a mostly vertical comb, and $t^* = 0.50$ asks that both roles weight the gravity reference at least as heavily as direct pointing, so that the code has genuinely tipped over rather than merely drifted. The threshold $m^* = 0.02$ is a viability floor: near-random search leaves the founding generation with low success, but the minimum can also occur later, when comb tilt outpaces the TRANSPOSITION traits and degrades the ancestral code before the gravity-referenced one takes over. Penalizing m_j wherever it falls steers the search away from settings in which foraging is never adequate. A seed's *progress* measures how near it came to the first two: each trait is expressed as a fraction of the threshold it must reach, and progress is the smallest of those fractions, capped at one,

$$g_j = \min\Big(1,\ \frac{\bar{\gamma}_j}{\gamma^*},\ \frac{\bar{t}_{s,j}}{t^*},\ \frac{\bar{t}_{r,j}}{t^*}\Big), \tag{10}$$

so that g_j tracks whichever of the three lags furthest behind. Each fraction reaches one exactly when its trait reaches its threshold, so $g_j = 1$ holds exactly when seed j is stable, and $g_j < 1$ grades the near misses. The trial score is then

$$\underbrace{\sum_{j=1}^{J} \mathbf{1}\big[g_j = 1\big]}_{\text{stable count}} + \underbrace{\frac{1}{J}\sum_{j=1}^{J} g_j}_{\text{mean progress}} - \underbrace{\sum_{j=1}^{J} \mathbf{1}\big[m_j \leq m^*\big]}_{\text{non-viable count}}. \tag{11}$$

Because the progress term is an average of quantities bounded by one, it can never contribute more than a single point, so one further stable seed outweighs any improvement in progress across the whole panel: the search pursues transitions that complete, and uses progress only to rank panels that fail.

Confirmation and held-out validation. Because the optimization ranks trials by their score on only ten seeds, a top-ranked candidate may owe that score to a genuinely robust transition or to a lucky draw of those seeds, and re-scoring it on unseen seeds tells the two apart. The highest-scoring search trials are re-run on a disjoint confirmation panel of 40 seeds, and the best candidates are then validated on 100 held-out seeds that were used in neither search nor confirmation. Validation reports, per candidate, the fraction of stable seeds, the non-viable count, and final-generation means of foraging success, COMB TILT, and the lower of the two TRANSPOSITION traits.

One-parameter sensitivity. The search locates a viable setting but, its trials clustering around what already worked, does not say how wide the viable region around it is. Around

Symbol / name	Interval	Step
food-site count	1–8	integer
μ_r (m)	60–360	15
capacity	2–12	integer
$d_{\max}$ (km)	3–7.5	0.25
travel cost (per km)	0.010–0.060	0.0025
B	0.10–0.60	0.02
σ_m	0.04–0.14	0.01
ρ	0.0–1.0	0.1

Table 5: The space searched by the optimization. Each parameter is proposed on a lattice of the given step within its interval, so the space is discrete rather than continuous.

the strongest validated candidate we therefore perturb each parameter in turn, holding the others at their validated values, over the same 100 held-out seeds; sweeping one axis at a time on a regular grid measures how far each parameter can move before the transition degrades, and so which parameters it tolerates freely and which it depends on finely.

Evolutionary-parameter interaction. To ask whether mutation parameters can compensate for a weaker architectural benefit, we hold the validated ecology fixed and run a full-factorial grid over the three evolutionary parameters, sweeping the benefit across its full search range while centring the mutation-parameter grids on the values found stable above, $B \in \{0.10, 0.30, 0.45, 0.60\}$, $\sigma_m \in \{0.05, 0.07, 0.09, 0.11\}$, and $\rho \in \{0.0, 0.3, 0.6, 0.9\}$, evaluating each of the 64 cells over the same 100 held-out seeds.

Outcome measures. Across the transition experiments we summarize runs by the same thresholds used in the search objective: the *stable* fraction (vertical comb with a coordinated gravity code) and the *non-viable* fraction (foraging success ≤ 0.02 at its worst generation), alongside final-generation trait and success means. Seed-bootstrap intervals over the per-seed outcomes quantify sampling uncertainty in the reported fractions.

5 Results

We report the two stages in turn: the horizontal-comb ecology which fixes when the direct-pointing code is worth maintaining (Section 4.2), and the vertical transition under both decode variants (Section 4.3).

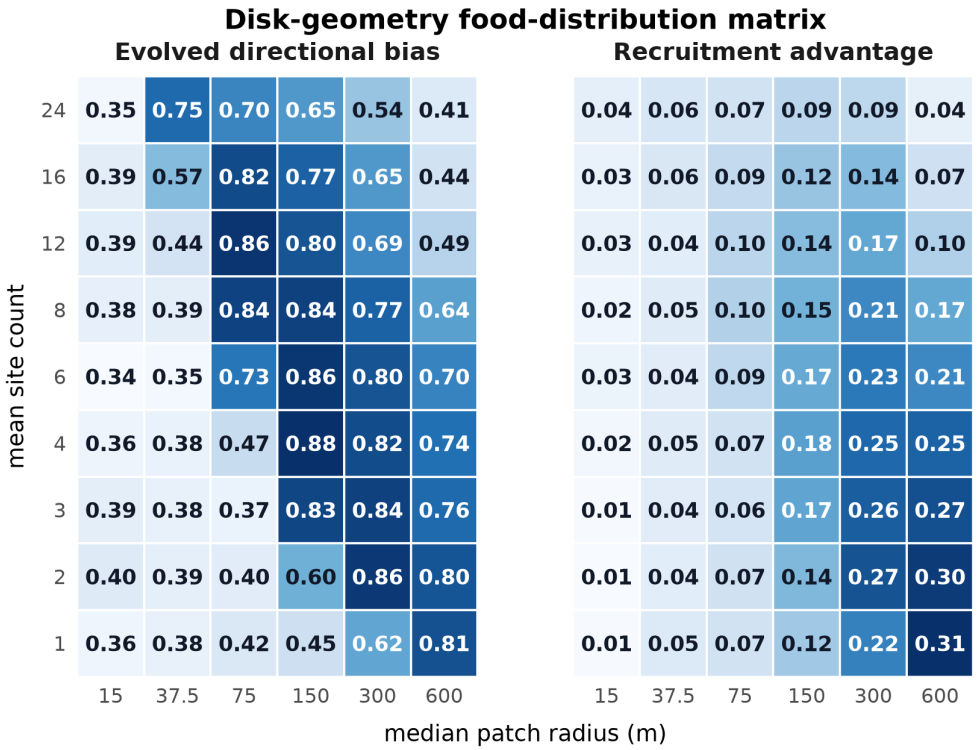

Figure 5: Evolved DIRECTIONAL BIAS (left) and recruitment advantage (right) on a flat comb across the grid of Poisson mean site count against median patch radius, 50 seeds per cell. Recruitment advantage is the follower–searcher success difference, and shading is normalized within each panel.

5.1 Emergence of direct pointing

On a comb held flat, the evolved DIRECTIONAL BIAS, which sets how precisely a dance points, is what tells us communication is favored. Absent selection for communication it drifts to a low baseline; it rises clearly above that baseline only along a diagonal band of the food grid (Figure 5), and the patch size at which a precise dance becomes worthwhile falls as patches grow more numerous. Below a patch radius of a few tens of meters the dance never evolves, however abundant the food. The recruitment advantage shows what a dance is worth where it does evolve: it is largest for few, large patches and erodes toward both small patches, where successful foragers are too rare to seed dances, and abundant large ones, where independent search already finds the food.

Outcome	Flatten trials	Unproject trials
Stable in all 10 seeds	14	24
Stable in ≥ 8 seeds	123	132
Stable in ≥ 5 seeds	312	338
Stable in ≥ 1 seed	611	641
No stable seed	413	383

Table 6: Stability counts over the 1024 Optuna trials (ten seeds each) under each decode.

	Flatten		Unproject	
Parameter	Median	10th–90th	Median	10th–90th
Food-site count (Poisson mean)	8	5–8	8	6–8
Patch radius (m)	210	165–345	225	150–360
Patch capacity	7	2–8	7	3–9
Vertical-comb benefit B	0.56	0.50–0.56	0.56	0.50–0.56
Max food distance (km)	3.25	3.25–5.25	5.25	3.5–5.25
Travel cost per km	0.025	0.010–0.025	0.018	0.010–0.022
Mutation scale σ_m	0.110	0.090–0.110	0.110	0.110–0.130
Correlation ρ	0.4	0.0–0.9	0.4	0.0–1.0

Table 7: Parameter distribution of the strongly stable trials (eight or more of ten stable seeds) under each decode.

5.2 Evolution of the gravity code

Parameter search. Stable vertical-gravity transitions, in which a seed ends with a mostly vertical comb and both worker roles committed to the gravity reference (Section 4.3), are common under both decodes, and *unproject* is slightly broader (Table 6). Both searches concentrate their strongly stable trials, those stable in at least eight of their ten seeds, in much the same region: abundant, large, reachable food paired with a strong tilt incentive and a high mutation scale (Table 7). The main difference is that *unproject* tolerates food at longer range: this suggests that the bias of *flatten* works against the transition more than it drives it. Distant food demands accurate dances, and that is just where *flatten*'s bias corrupts the direct code, so under *flatten* the transition completes only where food is close enough for the biased dance to stay usable.

Held-out validation. The top five distinct candidates from each search, carried through the 40-seed confirmation panel and rerun on 100 held-out seeds (200–299), all produced frequent stable transitions and no non-viable seeds; stable rates ran 80–91 of 100 for flatten and 82–92 for unproject. The strongest candidate of each decode (Table 8) lands on

Parameter	Flatten (cand. 729)	Unproject (cand. 541)
Food-site count (Poisson mean)	5	7
Patch radius (m)	315	315
Patch capacity	4	12
Vertical-comb benefit B	0.56	0.60
Max food distance (km)	3.25	4.75
Travel cost per km	0.017	0.010
Mutation scale σ_m	0.070	0.090
Correlation ρ	0.9	0.4
Stable seeds (of 100)	91	92
Final mean COMB TILT t_f	0.845	0.848

Table 8: Strongest validated candidate per decode over 100 held-out seeds.

very similar held-out rates and on overlapping ecologies: large patches, short-to-moderate distances, low travel cost, high benefit.

One-parameter sensitivity. A one-parameter sweep around each validated baseline identifies the mutation scale as the sharpest boundary (Figures 6 and 7): lowering it collapses the transition (flatten falls from 91/100 to 49/100 at $\sigma_m = 0.05$; unproject holds until a lone drop to 16/100 at 0.04), while the validated σ_m of 0.07–0.09 sits on the plateau. The sharpest *ecological* boundary is food-site count under flatten (46/100 at a Poisson mean of two, rising to 87/100 at seven); unproject is more robust to count. Vertical-comb benefit and correlation are monotone and moderate, and patch radius, capacity, distance, and travel cost stay within a narrow band.

Evolutionary-parameter interaction. Fixing each decode's validated ecology and sweeping B, σ_m, and ρ across the search range, stable rate falls steeply with benefit: robust near the validated $B \approx 0.56$–0.60, already marginal by $B = 0.30$, and collapsing at the search floor $B = 0.10$, with unproject again slightly more robust than flatten at every benefit tested (Figure 8). A complete failure to produce even one stable transition among the 100 seeds is rare and confined to that same floor: it occurs in 3 of the 64 cells under each decode, all at $B = 0.10$ and, in all but one, also at the lowest mutation scale. Within each benefit panel the rate rises with mutation scale and with correlation, and the best cells pair a high mutation scale (0.09–0.11) with moderate-to-high correlation, so higher mutation and correlation only partly compensate for weaker architectural benefit. This compensation from correlation is concentrated in the low-mutation regime: at the smallest mutation scale, raising the correlation from 0 to 0.9 roughly doubles the stable rate, whereas at the high-

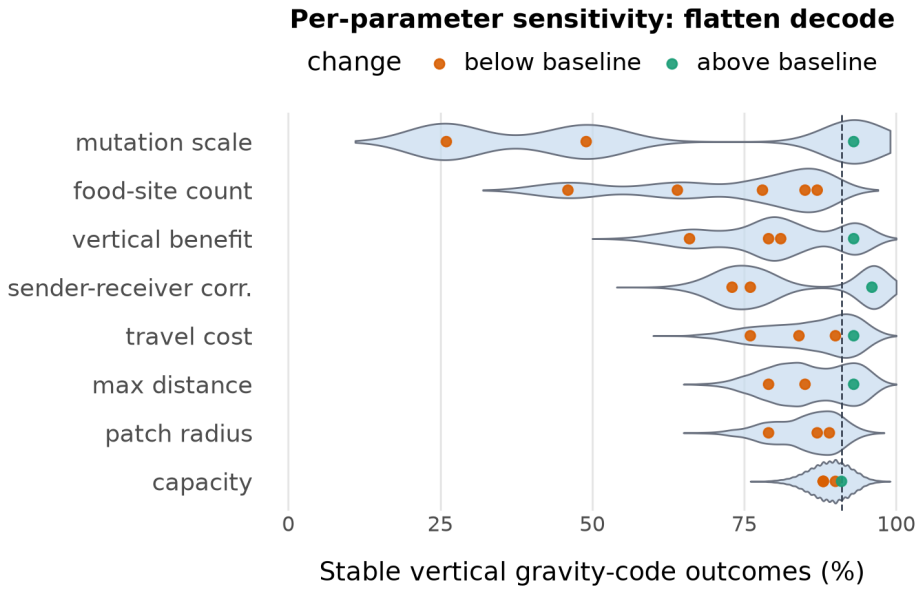

Figure 6: One-parameter sensitivity of the flatten decode. Each parameter's violin is a kernel density over seed-bootstrap stable-rate draws pooled across that parameter's swept values (100 held-out seeds per value), so a lobe's position is a swept value's rate and its width that value's sampling uncertainty; dots mark the per-value rates (orange below, green above baseline) and the dashed line the baseline. Parameters are sorted by worst-case drop (shared order and x-range with Figure 7).

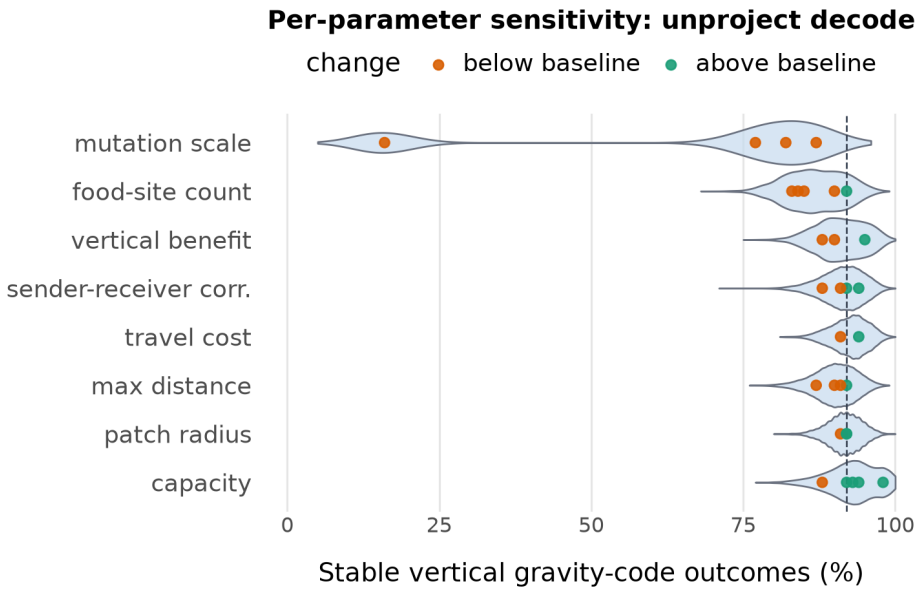

Figure 7: One-parameter sensitivity of the unproject decode, drawn as in Figure 6 and sharing its parameter order and x-range.

est mutation scale its effect is much weaker, and slightly negative under unproject. This regime-dependence explains why the search itself left ρ largely unconstrained, its strongly stable trials spanning almost the whole $[0, 1]$ axis under both decodes (Table 7): the search concentrated at high mutation scale, exactly where correlation barely moves the outcome.

6 Discussion and Conclusion

The current study poses two questions: under what ecological conditions a direct-pointing referential code emerges and stabilizes, and which factors let a population move from this code to its gravity-referenced variant. Both are feasibility questions rather than probability estimates: they ask what conditions make each code possible at all, not how likely a given ecology or lineage was to realize it. The two stages of the simulation answer them in turn.

On a horizontal comb, direct pointing is feasible only within a moderate band of foraging difficulty (Figure 5), and the edges of that band are gradual rather than sharp: too little food never seeds enough dances to be useful, too much makes independent search sufficient on its own.

The transition itself turns out to be feasible far more broadly than it is probable. Sweeping the benefit B, the mutation scale σ_m, and the sender–receiver correlation ρ over their full search ranges at a fixed, favorable ecology, a transition fails to occur in only 3 of

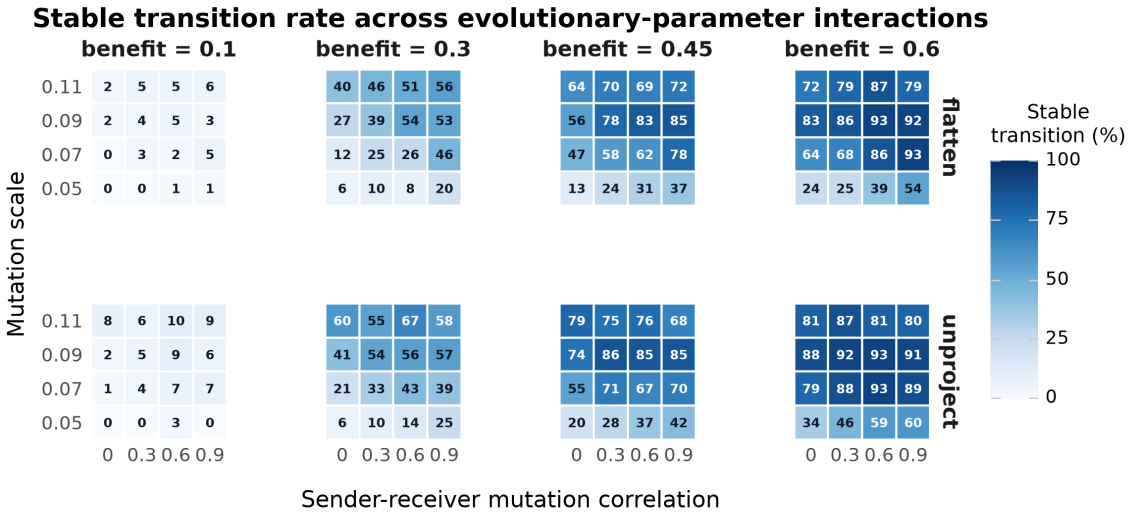

Figure 8: Stable transition rate across the evolutionary-parameter interaction grid, holding each decode's validated ecology fixed. Benefit is swept across the full search range (0.10–0.60) as a generic gradient rather than through the validated operating points (flatten B=0.56, unproject B=0.60). Tiles cross the sender–receiver mutation correlation ρ against the mutation scale σ_m, faceted by decode (rows) and vertical-comb benefit B (columns); each cell summarizes 100 held-out seeds, and its label is the percentage of seeds with a stable transition.

64 grid cells per decode, each at the lowest benefit and mutation scale tested (Figure 8); everywhere else, some seeds complete it, though often only a minority do. The regime the search converges on, a high mutation scale paired with a large benefit, is not especially biologically plausible, nor is it necessary to allow the code transition to occur.

Our approach is to treat the pattern of results here as a transition feasibility demonstration rather than a claim about its likelihood: across nearly all of the range we tested, nothing about the evolutionary parameters rules the transition out.

Taken together, these answers complement the mechanistic proposal of Barron and Plath (2017), which addresses how a single bee's brain can support either code: if the underlying heading representation is indeed indifferent to the spatial reference frame, our results indicate where the remaining barrier lies, not in individual cognition but in whether an evolving population of senders and receivers can stay coordinated as the code itself changes.

The broader picture across both stages is the same kind of claim at two different scales: a referential code, and later its replacement, need not be finely tuned into existence. Direct pointing is favored across a real, if bounded, swath of foraging ecologies; the shift to a gravity-referenced code, once that ecology holds and vertical building pays off, fails only at the joint extreme of weak benefit and low mutation scale, and remains open everywhere else we tested. This is a claim about what is possible, not what was likely in the actual evolutionary history of *Apis*: that would need independent evidence on ancestral foraging ecology, the real benefit of vertical combs, and the genetic basis of dance production and interpretation. Feasibility is, however, enough to answer the coordination problem posed in the introduction: a shared code can be revised without a breakdown of communication across most of the conditions we examined, not just at one finely tuned point among them.

7 Limitations

Several aspects of the model bound what these results can support. Only the dance's directional component is modeled, so the results say nothing about how encoding distance would affect the emergence or the transition of the code. Colonies reproduce asexually, with no gene flow between them, and the vertical-comb benefit B is an exogenous, fixed-form term (Section 3.7) rather than one derived from a mechanical or thermal account of comb architecture, so the model is silent on magnitude or origin of the real advantage. Likewise, the mutation scale σ_m sets how fast colony traits explore trait space over a fixed number of generations, not a real per-generation mutation rate. Transposition is also blended continuously rather than switched discretely (Section 3.8): a permissive choice that may widen the feasible region we report.

The parameter analyses are similarly bounded. The one-parameter sensitivity sweeps (Figures 6 and 7) and the evolutionary-parameter interaction grid (Figure 8) both hold ecology fixed at a single validated candidate, so neither shows how these effects generalize

elsewhere; and because the horizontal and transition stages sweep different ecological grids under different evolutionary horizons (Sections 4.2 and 4.3), their results are best read side by side rather than combined into a single claim.

References

Akiba, T., Sano, S., Yanase, T., Ohta, T., and Koyama, M. (2019). Optuna: A next-generation hyperparameter optimization framework. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2623–2631. ACM.

Bailis, P., Nagpal, R., and Werfel, J. (2010). Positional communication and private information in honeybee foraging models. In *Swarm Intelligence: 7th International Conference, ANTS 2010*, volume 6234 of *Lecture Notes in Computer Science*, pages 263–274. Springer.

Barron, A. B. and Plath, J. A. (2017). The evolution of honey bee dance communication: a mechanistic perspective. *Journal of Experimental Biology*, 220(23):4339–4346.

Bergstra, J., Bardenet, R., Bengio, Y., and Kégl, B. (2011). Algorithms for hyper-parameter optimization. In *Advances in Neural Information Processing Systems 24 (NIPS 2011)*, pages 2546–2554.

Bernard, A., Wischmann, S., Floreano, D., and Keller, L. (2023). The evolution of behavioral cues and signaling in displaced communication. *PLOS Computational Biology*, 19(3):1–16.

Chaabouni, R., Kharitonov, E., Bouchacourt, D., Dupoux, E., and Baroni, M. (2020). Compositionality and generalization in emergent languages. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 4427–4442.

Dornhaus, A., Klügl, F., Oechslein, C., Puppe, F., and Chittka, L. (2006). Benefits of recruitment in honey bees: effects of ecology and colony size in an individual-based model. *Behavioral Ecology*, 17(3):336–344.

Edrich, W. (1977). Interaction of light and gravity in the orientation of the waggle dance of honey bees. *Animal Behaviour*, 25:342–363.

Esch, H. (2012). Foraging honey bees: How foragers determine and transmit information about feeding site locations. In Galizia, C. G., Eisenhardt, D., and Giurfa, M., editors, *Honeybee Neurobiology and Behavior: A Tribute to Randolf Menzel*, pages 53–64. Springer Netherlands, Dordrecht.

Hamilton, W. D. (1964). The genetical evolution of social behaviour. I. *Journal of Theoretical Biology*, 7(1):1–16.

Hepburn, H. R., Pirk, C. W. W., and Duangphakdee, O. (2014). *Honeybee Nests: Composition, Structure, Function*. Springer, Berlin, Heidelberg.

Lazaridou, A. and Baroni, M. (2020). Emergent multi-agent communication in the deep learning era. *arXiv preprint arXiv:2006.02419*.

Lazaridou, A., Peysakhovich, A., and Baroni, M. (2017). Multi-agent cooperation and the emergence of (natural) language. In *International Conference on Learning Representations (ICLR)*.

Lewis, D. (1969). *Convention: A Philosophical Study*. Harvard University Press, Cambridge, MA.

Lian, Y., Bisazza, A., and Verhoef, T. (2023). Communication drives the emergence of language universals in neural agents: Evidence from the word-order/case-marking trade-off. *Transactions of the Association for Computational Linguistics*, 11:1033–1047.

Lo, N., Gloag, R. S., Anderson, D. L., and Oldroyd, B. P. (2010). A molecular phylogeny of the genus Apis suggests that the giant honey bee of the philippines, A. breviligula Maa, and the Plains Honey Bee of southern India, A. indica Fabricius, are valid species. *Systematic Entomology*, 35(2):226–233.

Nowak, M. A. and Krakauer, D. C. (1999). The evolution of language. *Proceedings of the National Academy of Sciences*, 96(14):8028–8033.

Portegys, T. E. (2020). Morphognostic honey bees communicating nectar location through dance movements. *bioRxiv*.

Schürch, R. and Grüter, C. (2014). Dancing bees improve colony foraging success as long-term benefits outweigh short-term costs. *PLoS ONE*, 9(8):e104660.

Seeley, T. D. (1989). The honey bee colony as a superorganism. *American Scientist*, 77(6):546–553.

Seeley, T. D. and Morse, R. A. (1976). The nest of the honey bee (Apis mellifera l.). *Insectes Sociaux*, 23(4):495–512.

Siregar, N. P., Le, P., and Alhama, R. G. (2026). An emergent communication framework for honeybee waggle dance. In *International Conference on the Evolution of Language*, page 435.

Skyrms, B. (2010). *Signals: Evolution, Learning, and Information*. Oxford University Press, Oxford.

Smith, K., Kirby, S., and Brighton, H. (2003). Iterated learning: A framework for the emergence of language. *Artificial Life*, 9(4):371–386.

von Frisch, K. (1967). *The Dance Language and Orientation of Bees*. Harvard University Press, Cambridge, MA.

Wagner, K., Reggia, J. A., Uriagereka, J., and Wilkinson, G. S. (2003). Progress in the simulation of emergent communication and language. *Adaptive Behavior*, 11(1):37–69.