



HAL
open science

Emergent kin selection of altruistic feeding via non-episodic neuroevolution

Max Taylor-Davies, Gautier Hamon, Timothé Boulet, Clément Moulin-Frier

► **To cite this version:**

Max Taylor-Davies, Gautier Hamon, Timothé Boulet, Clément Moulin-Frier. Emergent kin selection of altruistic feeding via non-episodic neuroevolution. International Conference on the Applications of Evolutionary Computation (Part of EvoStar), 2025, Trieste, Italy. hal-04865103

HAL Id: hal-04865103

<https://hal.science/hal-04865103v1>

Submitted on 5 Jan 2025

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Emergent kin selection of altruistic feeding via non-episodic neuroevolution

Max Taylor-Davies¹, Gautier Hamon², Timothé Boulet², and Clément Moulin-Frier²

¹ School of Informatics, University of Edinburgh, Edinburgh, Scotland
m.taylor-davies@sms.ed.ac.uk

² Flowers team, Inria Center of the University of Bordeaux, France.

Abstract. Kin selection theory has proven to be a popular and widely accepted account of how altruistic behaviour can evolve under natural selection. Hamilton’s rule, first published in 1964, has since been experimentally validated across a range of different species and social behaviours. In contrast to this large body of work in natural populations, however, there has been relatively little study of kin selection *in silico*. In the current work, we offer what is to our knowledge the first demonstration of kin selection emerging naturally within a population of agents undergoing continuous neuroevolution. Specifically, we find that zero-sum transfer of resources from parents to their infant offspring evolves through kin selection in environments where it is hard for offspring to survive alone. In an additional experiment, we show that kin selection in our simulations relies on a combination of kin recognition and population viscosity. We believe that our work may contribute to the understanding of kin selection in minimal evolutionary systems, without explicit notions of genes and fitness maximisation.

Keywords: Kin selection · Neuroevolution · Multi-agent systems · Artificial ecosystems

1 Introduction

At first glance, it seems difficult to square the phenomenon of purely altruistic behaviour (acts which confer a benefit to the recipient at a cost to the actor) with the basic principle of natural selection: how can a gene be selected for when it decreases, rather than increases, the fitness of its host? One plausible account can be made through the theory of inclusive fitness. Key to this theory is the recognition that individual organisms within a social environment are not isolated from their conspecifics in terms of fitness. Whether a given gene is selected for is thus determined by its effect(s) on the fitness of any bearers of copies of that gene. Under this view, we can think of an altruistic act as an exchange of fitness from one agent to another. If the exchange is positive-sum and both sides are bearers of the gene in question, then from the gene’s perspective the behaviour confers a fitness benefit—even while it decreases the fitness of the acting individual.

Hamilton’s rule, published in 1964, offers a simple and elegant formalisation of this idea [6,7]. Imagine a particular gene produces some social behaviour which confers fitness benefit b to the recipient at a cost c to the actor. Given a measure r of the genetic relatedness between the two, the rule states that this gene should increase in frequency if $rb > c$. A key consequence of Hamilton’s rule is that an altruistic behaviour that targets the actor’s kin will undergo more positive selection than an equivalent behaviour which is indiscriminate or favours non-kin. This is referred to as ‘kin selection’ [11], and can be seen as a more narrow form of inclusive fitness. Targeting of behaviour towards kin can be facilitated either by mechanisms of kin recognition, or more simply by population viscosity. Since the publication of Hamilton’s rule, empirical studies across a range of different species [1] have confirmed its predictions for social behaviours such as guarding among female Allodapine bees and defence of dominant male Wild turkeys by non-dominant males [12,10]. Kin selection theory has also been used to explain insect eusociality [7,9] and allomothering in certain species of monkey [3], and has even been invoked in accounts of language evolution [4,5].

In this paper, we demonstrate an emergent kin selection mechanism within a population of continuously evolving artificial agents situated in an ecologically plausible environment. We place thousands of agents within a large gridworld filled with resources that regenerate through time (Figure 1). Each agent is controlled by its own artificial neural network, mapping its local observation of the environment to its possible actions: moving in space, consuming a resource, feeding the agent next to it, or reproducing. We depart from the standard evolutionary computation paradigm: our simulations include no explicit fitness measure, instead relying on a ‘minimal criterion’ for individual agents’ survival and reproduction, governed by a simple physiological model. In addition, we do not reset the agent population or environment state at any point during simulation (a framework referred to as non-episodic neuroevolution [8]). Importantly, the only kin-specific mechanism we introduce is the ability of agents to distinguish between adults, infants and their own offspring. The answer to our main research question is therefore far from trivial: in these simple, ecologically plausible conditions, will agents evolve the altruistic behaviour of feeding their own offspring? It is important to note that doing so will not favour their own survival and reproduction in any way. It will only favour the propagation of their own (mutated) genome, but without any explicit incentive to do so. To study this research question, we manipulate certain features of the environment, indirectly varying the fitness tradeoff $rb - c$. Our results show that altruistic feeding behaviour emerges as this differential increases. We also show that an increase in the prevalence of the altruistic feeding behaviour is associated with an increase in its selectivity towards agents’ own offspring. To our knowledge, this is the first demonstration that kin selection naturally evolves in realistic ecological conditions where artificial agents continuously evolve in a non-episodic environment.

2 Environment and agents

2.1 Environment dynamics

The environment we use is based on the non-episodic neuroevolution framework described in [8]. Simulations are run in a 100x100 gridworld which contains both food resources (plants) and agents, with each tile containing at most one agent, up to a maximum population size of 2500. Resources are generated and regenerated via simple spontaneous growth: at the start of simulation, each tile has probability $p = 0.1$ of containing a plant; at every subsequent timestep, each non-plant-containing tile has constant probability $p = 0.003$ of spontaneously spawning one. The environment is initialised with a starting population of 2000 agents, all with random neural network weights. Agents’ survival and reproduction is governed entirely by simple physiological rules, to enable minimal-criterion neuroevolution [2]. Agents are born with an initial energy of 65. They lose 0.1 energy passively at each timestep, with a further loss of 1.0 for every non-idle action they perform. An agent dies when its energy falls below 0. Agents can increase their energy (up to a maximum of 200) by consuming plants, with each plant worth 20 units of energy. If an agent’s energy is at or above a threshold of 85, and the population is less than the maximum of 2500, it can use the reproduce action to produce a single offspring (at an energy cost of 30). Offspring spawn into the nearest empty square and receive a copy of their parent’s neural network weights randomly mutated by the addition of zero-mean Gaussian noise, i.e. $\theta_{\text{offspring}} := \theta_{\text{parent}} + \mathcal{N}(0, \sigma)$ (for all simulation runs reported in this paper we use $\sigma = 0.02$). Agents younger than 100 timesteps are considered to be ‘infants’—they are unable to reproduce, and may also experience slightly different environment dynamics or action mechanics (see Section 3 for more detail).

2.2 Agents

At each timestep, every agent in the population samples an action from a policy parameterised by an individual neural network (NN) controller, implemented as simple 3-layer convolutional neural network (CNN). The weights of each agent’s controller remain constant throughout their lifetime—i.e. there is no learning, and behavioural adaptation can happen only through neuroevolution. Agents observe the environment within an 11x11 window centred at their current location, within which they see the locations of both plants and other agents. They can also see, for each other agent within this window, two binary values indicating whether that agent is A) an infant, and B) their own offspring. The action space is given by {forward, turn left, turn right, eat, reproduce, feed, idle}. The ‘eat’ action has the effect of consuming a single plant resource, and is only effective if the agent is occupying a tile containing a food resource. The ‘feed’ action transfers one food resource’s worth of energy to a single other agent, and is only effective if the two agents are in adjacent tiles, and the feeder is directly oriented towards the feedee (facilitating some degree of selectivity in feeding behaviour). The ‘reproduce’ action is as described above (Section 2.1). Fig 1 gives an illustration of the environment, agent architecture, and energy dynamics.

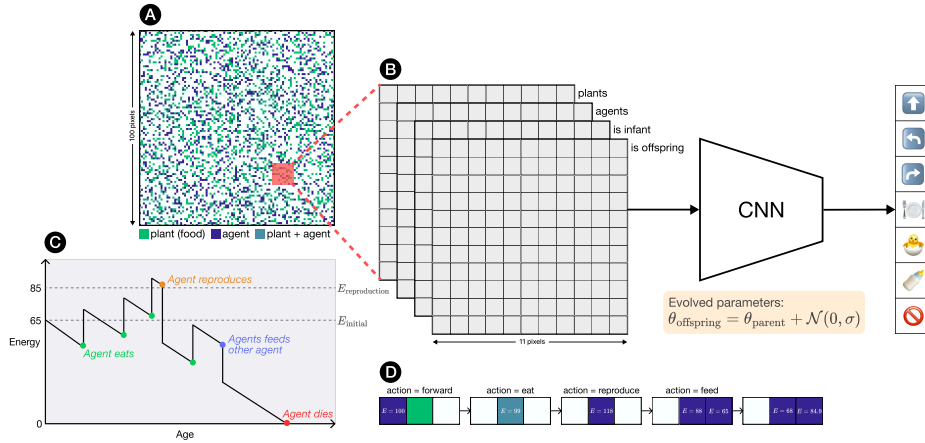
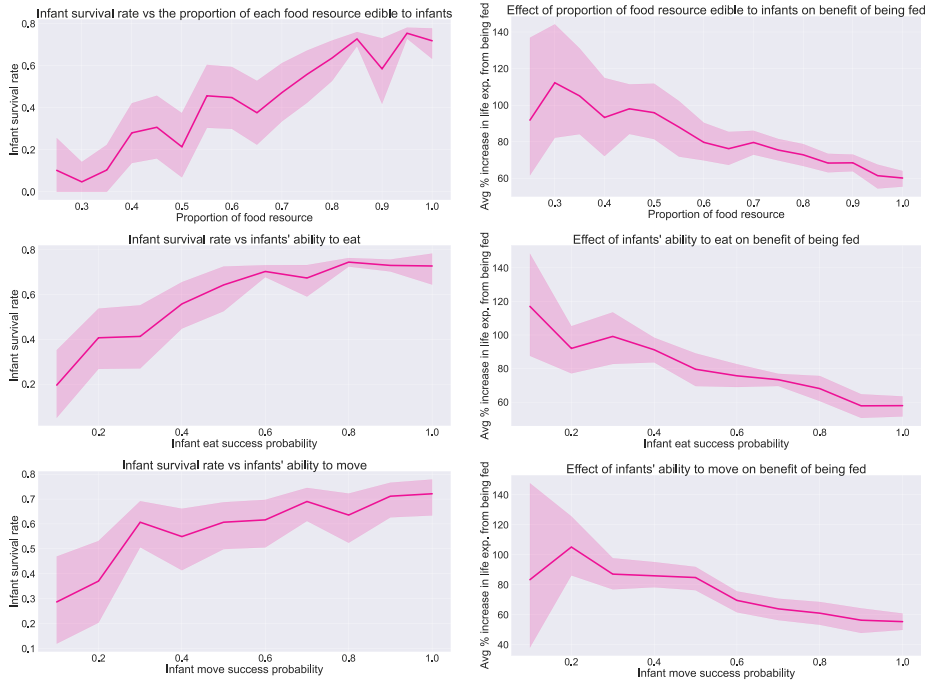


Fig. 1: (A) A snapshot of the simulation environment. (B) an illustration of the agent architecture, observation and action spaces. Each agent receives an 11x11x4 observation of its local environment—this is fed through a convolutional neural network (CNN) to produce a vector of action probabilities. (C) an example illustration of the energy level over time for a single agent. (D) an illustration of a short action sequence in which an agent moves onto a food resource, eats the food resource, reproduces and then feeds their offspring.

3 Experiments

In our experiments, we seek to determine whether altruistic feeding behaviour (i.e. the zero-sum transfer of resources from one agent to another) can emerge in our neuroevolution framework as a result of kin selection. We focus on a narrow version of kin selection that considers only the relationship between parents and their direct offspring. Since in this case the genetic relatedness r is close to 1 (specifically, $1 - \sigma$), Hamilton’s rule [6] tells us that in our environment, resource transfer from parent to child should be selected for when the fitness benefit b to the child of gaining the resource exceeds the fitness cost c to the parent of losing the resource. Our hypothesis is thus that in simulation environments where $b - c$ is greater, we should observe both more feeding behaviour *and* feeding behaviour that is more selectively directed towards agents’ own offspring. If on the other hand there is no kin selection process,

While controlling this tradeoff directly is challenging, we can take an indirect approach by varying the relative effect of a marginal food resource on the continued survival of parent and child. We do this by manipulating certain environment conditions that make it harder for *infant* agents to survive to adulthood (and thus to reproductive age), while keeping the survival difficulty for adult agents unchanged. This should have the effect of increasing b while keeping c constant, and thus increasing the differential $b - c$. The parameters we manipulate for this are:



(a) The proportion of agents that survive to adulthood, as a function of our three experimental parameters, measured over the final 10k timesteps of each 100k-timestep simulation, in environments where agents do **not** have the ability to feed one another. (b) The estimated benefit (to life expectancy) of being fed as an infant, as a function of our three experimental parameters, measured over the final 50k timesteps of each 500k-timestep simulation

Fig. 2: The effect of varying the three chosen simulation parameters on (a) the infant survival rate (ISR) with feeding disabled, and (b) the estimated benefit of being fed as an infant (\hat{b}), given by the average percentage increase in lifespan for infants that are fed relative to those that aren't. Error bars in all plots represent bootstrapped 95% confidence intervals over 20 seeds per parameter value.

- **infant food energy proportion**: the proportion of available energy that infant agents gain from each food resource (where adults get 100%), defaulting to 1.0. For example, if this value is 0.5, then an infant’s energy level will increase by 10 upon consuming a plant, whereas an adult’s would increase by 20.
- **infant eat success probability**: the probability with which an infant’s ‘eat’ action succeeds, assuming they are located at a food resource (where adults’ ‘eat’ actions are deterministically successful), defaulting to 1.0. For example, if this value is 0.5, an infant would have to use the ‘eat’ action twice as often as an adult to get the same expected energy gain.
- **infant move success probability**: equivalent to **infant eat success probability** but for the ‘move’ action, defaulting to 1.0. For example, if this value is 0.5, an infant would move around the environment twice as slowly as an adult following the same action policy.

To establish the effect of each of these parameters on the environment survivability, we first run a set of simulations with **feeding behaviour disabled** (i.e. agents’ action spaces do not contain the ‘feed’ action), and track the ‘infant survival rate’ (ISR) over time. For a window $[t_1, t_2]$, we compute this as

$$\text{ISR} = 1 - \frac{\text{num infant deaths in } [t_1, t_2]}{\text{total num deaths in } [t_1, t_2]} \quad (1)$$

Figure 2a shows the ISR recorded over the final 10k steps of 100k-step simulations as a function of each of the three parameters listed above, with 20 seeds run per parameter value (and the other parameters held constant at their default value of 1.0). For these preliminary experiments, we chose to end simulations at 100k timesteps after observing minimal changes in the overall action distribution past this point (indicating that agents’ policies had evolved to a stable equilibrium). While the trends are not monotonic, and there is a fair amount of variance over the 20 seeds, we can see that the recorded ISR increases with all of the three chosen parameters—as we would expect.

Having established that our chosen parameters do affect infants’ ability to survive to adulthood (and thus in theory the excess fitness $b - c$), we re-enable agents’ ability to feed others, and run a new set of simulations over 20 seeds x 500k timesteps per parameter value.

4 Results

4.1 The benefit of being fed is inversely correlated with infants’ ability to survive

While we use the three simulation parameters listed above as experimental proxies for $b - c$, we would like to obtain a more direct estimate \hat{b} of the fitness benefit to agents of being fed. Since agents that live longer will in expectation produce more offspring, we can do this by computing, over the final 50k steps of each

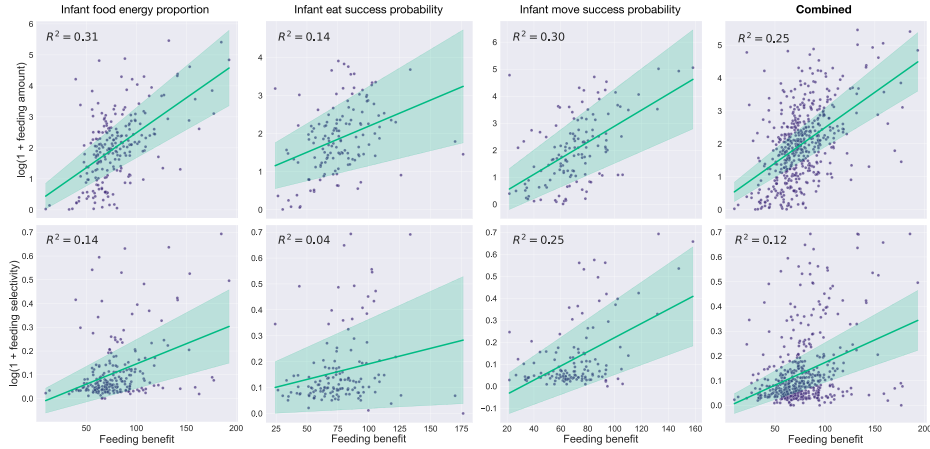


Fig. 3: The relationship between the estimated benefit to infants of being fed and both the amount and selectivity of feeding observed, shown separately for each of the three experimental parameters we varied (and combined in the rightmost column). Each scatterplot point represents a single 500k-timestep simulation run (with values averaged over the final 50k timesteps); regression lines (with 95% confidence intervals) are shown in green. Note that the y -axis shows $\log(\text{measure})$ for both amount and selectivity.

simulation run, the average difference in lifespan between agents that do and do not receive feeding as infants. Fig 2b shows how \hat{b} changes as a function of each of the three experimental parameters. Across all three parameters, we see that the benefit of being fed during infancy on agents’ life expectancy (and thus ability to reproduce) is higher in environments that are more hostile to infants’ survival.

4.2 Agents engage in more feeding and more selective feeding when the benefit of being fed is higher

Having established an approximate measure of the benefit to reproductive fitness of being fed, we now test the relationship between this benefit and the actual feeding behaviour observed. If feeding behaviour is being evolved via kin selection, then we should expect two effects as we increase the importance of being fed to young agents’ reproductive fitness:

- The amount of feeding behaviour observed in the population should increase,
- The selectivity of feeding behaviour towards agents’ own offspring should increase.

We measure the amount of feeding as simply the number of times the ‘feed’ action was successfully used over the final 50k timesteps of each simulation run, divided by the average population size over the same period. To measure selectivity, we

first take the proportion of these feeding events that were directed from parent to offspring. But this proportion is not necessarily a reliable measure of selectivity—it could be higher, for example, if infants stay closer to their parents on account of being less able to move. To account for this, we track all instances where one agent (agent A) is directly facing another agent (agent B), and compute the proportion of times where B is A’s offspring. We can treat this proportion as a baseline for the relative level of feeding towards offspring that we would expect if there were no ‘real’ selectivity; i.e. if agents were equally likely to use the feed action when facing any other agent. Our final measure of selectivity is then obtained by subtracting this baseline from the original proportion.

$$\text{selectivity} = \frac{\# \text{ times offspring fed}}{\# \text{ all feeding events}} - \frac{\# \text{ times agent facing offspring}}{\# \text{ times agent facing any agent}} \quad (2)$$

	log(amount) \sim benefit			log(selectivity) \sim benefit		
	β	r^2	p -val	β	r^2	p -val
varying infant food energy proportion	0.0225	0.312	<0.001	0.00170	0.138	<0.001
varying infant eat success probability	0.0137	0.138	<0.001	0.00120	0.043	0.014
varying infant move success probability	0.0297	0.297	<0.001	0.00320	0.247	<0.001
combined	0.0215	0.254	<0.001	0.00180	0.117	<0.001

Table 1: Results of OLS linear regression analysis for feeding amount and selectivity against feeding benefit (to infants)

To test for the two trends listed above, we carried out a linear regression analysis (by the OLS method). We found significant positive correlations between feeding benefit and both log(feeding amount) and log(feeding selectivity), indicating that the two trends are indeed observed, and have an exponential shape. The regression results are given in Table 1, and visualised in Figure 3.

4.3 The amount of feeding observed increases with the selectivity of feeding towards offspring

If a behaviour evolves through kin selection, then we should expect to see a positive correlation between the overall prevalence of the behaviour and the extent to which it is used selectively towards agents’ own kin. As an additional check, we therefore test the relationship between the amount and selectivity of feeding observed in our simulation runs. As is shown in Table 2 and Figure 4, we find a very strong positive correlation between the two measures, across all three experimental parameters.

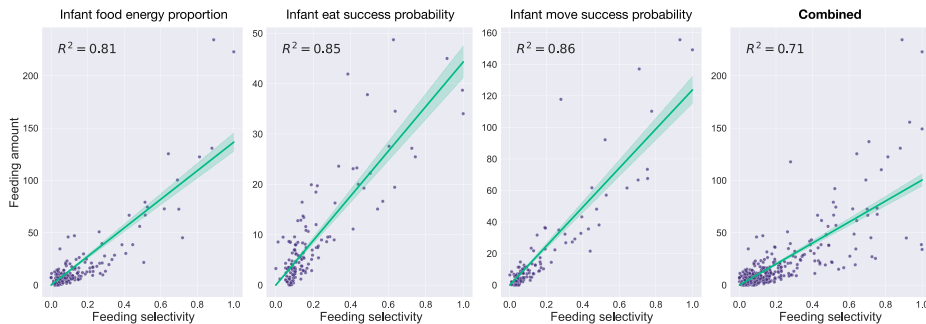


Fig. 4: The relationship between the amount of feeding behaviour observed and its selectivity towards agents’ own offspring, shown separately for each of the three experimental parameters we varied (and combined in the rightmost column). Each scatterplot point represents a single 500k-timestep simulation run (with values averaged over the final 50k timesteps); regression lines (with 95% confidence intervals) are shown in green.

	amount \sim selectivity		
	β	r^2	p -val
varying infant food energy proportion	136.5	0.806	<0.001
varying infant eat success probability	44.31	0.848	<0.001
varying infant move success probability	123.8	0.862	<0.001
combined	100.4	0.708	<0.001

Table 2: Results of OLS linear regression analysis for feeding amount against feeding selectivity

4.4 Kin selection operates through a combination of kin recognition and population viscosity

As mentioned briefly in Section 1, kin selection is understood to operate via two possible mechanisms: kin recognition, and population viscosity [6,7]. In the case of kin recognition, individuals have some faculty for recognising their own kin, which allows them to target certain behaviours specifically towards them. In the absence of such a faculty, kin selection may still operate if the average movement of individuals from their birthplace is sufficiently slow. In such populations, the local social environment of an individual likely consists in large part of their relatives, simply by default—and so mere targeting of behaviour towards neighbours is sufficient for kin selection to occur.

As a final experiment, we seek to understand which of these two mechanisms is responsible for the kin selection we observe in our virtual agent populations. To do this, we perform two independent ablations to our existing setup, and re-run the full set of simulations with each. To test the importance of kin recognition, we simply disable the observation feature that tells an agent whether each other

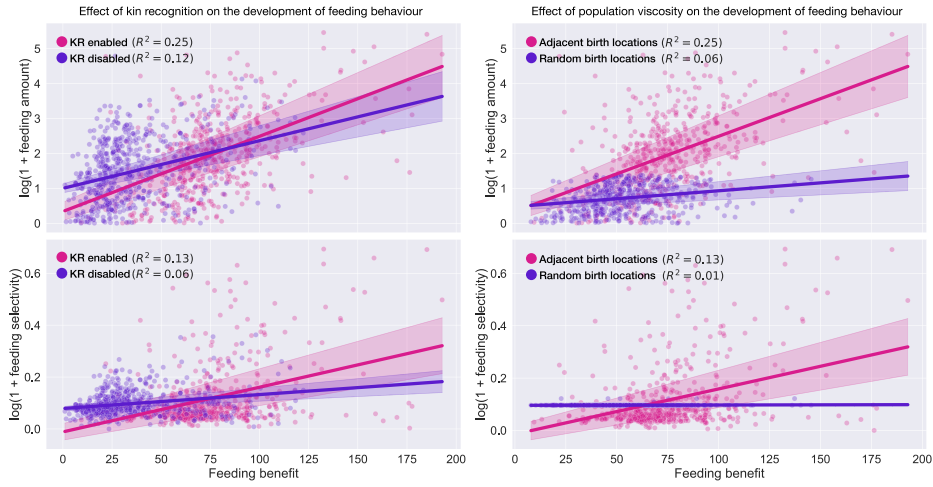


Fig. 5: The effect of disabling kin recognition (left) and reducing population viscosity (right) on the relationships between feeding benefit and feeding amount and feeding selectivity. Each dot represents a single 500k-timestep simulation run. Simulations were run over the same values of the three experimental parameters as before, and as before we ran 20 seeds per parameter value. 95% confidence intervals are shown for all regression lines.

agent they can see is one of their own offspring (see Figure 1; we retain the ‘is infant’ feature). Population viscosity is a little trickier to intervene on directly. Our approach was to modify the reproduction dynamics. Previously, each new offspring was born into the tile their parent was facing at the time of reproduction (meaning parent and child were initially directly adjacent to one another). Instead, we now randomly select the birth location of all new offspring agents, meaning that a newborn is no more likely to be near to its parent as to any other agent in the population.

Figure 5 shows the same regression analysis as performed previously for the two independent ablation runs. We can see that for both ablations, the relationships between feeding benefit (b) and the amount and selectivity of feeding are weakened. Interestingly, the effect is considerably stronger for randomised birth locations than for disabled kin recognition—suggesting that while both mechanisms play a role in the kin selection we observe, population viscosity is more important in facilitating the evolution of kin-selective feeding behaviour.

5 Discussion

In this paper, we have presented what is to our knowledge the first evidence of kin selection emerging from interactions between simulated agents undergoing neuroevolution in a simple but ecologically plausible environment. Focusing specifically on the phenomenon of agents feeding (donating energy to) their own

infant offspring, we showed that this feeding behaviour saw greater adoption as a function of the fitness differential $rb - c$, as predicted by Hamilton’s rule. We also found a very strong relationship between the prevalence of feeding behaviour observed and its selectivity towards agents’ own offspring, supporting the idea that the use of the ‘feed’ action emerged principally through a mechanism that reinforced its use towards kin but not non-kin. Finally, we investigated the mechanisms by which this kin selection process occurred, and found that while both kin recognition and population viscosity played a role, the latter had a more significant impact.

As well as serving as a methodological demonstration for how kin selection and inclusive fitness can be studied using neuroevolution of simulated agents, we believe that our results may have some broader theoretical implications. The commonly accepted account of kin selection is made in terms of fitness maximisation from the perspective of individual genes shared by genetically related organisms. However, our simulations involve neither a proper notion of ‘gene’ (or any discrete genetic subunit), nor an explicit notion of fitness maximisation. Agents are governed by, and pass to their offspring, a set of continuous neural network weights with no particular structure; agents survive and reproduce as long as they can maintain their internal energy above a certain level. The fact that we were still able to observe a kin selection phenomenon emerging suggests the possibility of a more general framing of the theory for such systems. One possible interpretation of our results is that by increasing the difficulty for infants to survive without being fed (i.e. increasing $rb - c$), we increased the extent to which the feeding of offspring was necessary for an agent to propagate a lineage over several generations. That is, as $rb - c$ increased, the probability increased for the genome of any agent *not* engaging in offspring-feeding to disappear from the population, increasing the overall prevalence of the behaviour in the phenotype of the population.

While we do believe our results to be of some interest, we close by noting a number of limitations that could be addressed in future work that builds upon our framework. First, our agents’ behaviour is governed by a very simple reactive architecture, in which actions are selected based purely on the current state observation, and there is no within-lifetime adaptation. This limits the complexity of the behavioural policies that agents can implement, and it would be interesting to explore a variant using (for example) recurrent architectures. For instance, an agent with memory could in principle ‘remember’ which other agents are its own offspring without needing access to an explicit marker; or could adapt their behaviour to focus more on caregiving in later life. Secondly, our kin recognition mechanism is also very simplistic, only allowing agents to distinguish their own direct offspring. For a more complete test of the quantitative predictions made by Hamilton’s rule, we would need to provide a continuous measure of the genetic relatedness between agents, in terms of their distance in neural parameter space. Finally, while not specifically a limitation of the current work, we believe that another fruitful research direction could be in exploring the evolution of more complex social behaviours within the same non-episodic neuroevolution

framework. For example, could environment dynamics that change over time (instead of being stationary) foster the emergence of teaching or demonstration by parents for the benefit of their newborn offspring?

A Appendix

All code is available at this anonymised Github repo.

Acknowledgments: This work was supported by the United Kingdom Research and Innovation (grant EP/S023208/1), EPSRC Centre for Doctoral Training in Robotics and Autonomous Systems (RAS). This work was also partially funded by the French National Research Agency (<https://anr.fr/>, project ECOCURL, Grant ANR-20-CE23-0006). This work also benefited from access to the HPC resources of IDRIS under the allocation 2023-[A0151011996] made by GENCI, using the Jean Zay supercomputer.

References

1. Bourke, A.F.G.: Hamilton's rule and the causes of social evolution. *Philosophical Transactions of the Royal Society B: Biological Sciences* **369**(1642), 20130362 (May 2014). <https://doi.org/10.1098/rstb.2013.0362>, <https://royalsocietypublishing.org/doi/10.1098/rstb.2013.0362>, publisher: Royal Society
2. Brant, J.C., Stanley, K.O.: Minimal criterion coevolution: a new approach to open-ended search. In: *Proceedings of the Genetic and Evolutionary Computation Conference*. p. 67–74. GECCO '17, Association for Computing Machinery, New York, NY, USA (2017). <https://doi.org/10.1145/3071178.3071186>, <https://doi.org/10.1145/3071178.3071186>
3. Fairbanks, L.A.: Reciprocal benefits of allomothering for female vervet monkeys. *Animal Behaviour* **40**(3), 553–562 (1990). [https://doi.org/https://doi.org/10.1016/S0003-3472\(05\)80536-6](https://doi.org/https://doi.org/10.1016/S0003-3472(05)80536-6), <https://www.sciencedirect.com/science/article/pii/S0003347205805366>
4. Fitch, W.T.: Kin selection and "mother tongues": A neglected component in language evolution. In: *Evolution of Communication Systems: A Comparative Approach*. The MIT Press (07 2004). <https://doi.org/10.7551/mitpress/2879.003.0022>, <https://doi.org/10.7551/mitpress/2879.003.0022>
5. Fitch, W.T.: *Evolving Meaning: The Roles of Kin Selection, Allomothering and Paternal Care in Language Evolution*, pp. 29–51. Springer London, London (2007). https://doi.org/10.1007/978-1-84628-779-4_2, https://doi.org/10.1007/978-1-84628-779-4_2
6. Hamilton, W.D.: The genetical evolution of social behaviour. I. *Journal of Theoretical Biology* **7**(1), 1–16 (Jul 1964). [https://doi.org/10.1016/0022-5193\(64\)90038-4](https://doi.org/10.1016/0022-5193(64)90038-4), <https://www.sciencedirect.com/science/article/pii/0022519364900384>
7. Hamilton, W.: The genetical evolution of social behaviour. II. *Journal of Theoretical Biology* **7**(1), 17–52 (1964). [https://doi.org/https://doi.org/10.1016/0022-5193\(64\)90039-6](https://doi.org/https://doi.org/10.1016/0022-5193(64)90039-6), <https://www.sciencedirect.com/science/article/pii/0022519364900396>

8. Hamon, G., Nisioti, E., Moulin-Frier, C.: Eco-evolutionary dynamics of non-episodic neuroevolution in large multi-agent environments. In: Proceedings of the Companion Conference on Genetic and Evolutionary Computation. p. 143–146. GECCO '23 Companion, Association for Computing Machinery, New York, NY, USA (2023). <https://doi.org/10.1145/3583133.3590703>, <https://doi.org/10.1145/3583133.3590703>
9. Hughes, W.O.H., Oldroyd, B.P., Beekman, M., Ratnieks, F.L.W.: Ancestral monogamy shows kin selection is key to the evolution of eusociality. *Science* **320**(5880), 1213–1216 (2008). <https://doi.org/10.1126/science.1156108>, <https://www.science.org/doi/abs/10.1126/science.1156108>
10. Krakauer, A.H.: Kin selection and cooperative courtship in wild turkeys. *Nature* **434**, 69–72 (2005), <https://doi.org/10.1038/nature03325>
11. Smith, J.M.: Group selection and kin selection. *Nature* **201**, 1145–1147 (1964), <https://doi.org/10.1038/2011145a0>
12. Stark, R.E.: Cooperative nesting in the multivoltine large carpenter bee *xylocopa sulcatipes maa* (apoidea: Anthophoridae): Do helpers gain or lose to solitary females? *Ethology* **91**(4), 301–310 (1992). <https://doi.org/https://doi.org/10.1111/j.1439-0310.1992.tb00871.x>, <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1439-0310.1992.tb00871.x>