



# Neural networks reveal emergent properties of collective learning in democratic but not despotic groups



Joe Morford <sup>a,\*</sup>, Patrick Lewin <sup>a,ib</sup>, Dora Biro <sup>a,b</sup>, Tim Guilford <sup>a</sup>, Oliver Padgett <sup>a</sup>, Julien Collet <sup>a,\*</sup>

<sup>a</sup> Department of Biology, University of Oxford, Oxford, U.K.

<sup>b</sup> Department of Brain and Cognitive Sciences, University of Rochester, Rochester, NY, U.S.A.

## ARTICLE INFO

### Article history:

Received 23 March 2022

Initial acceptance 23 June 2022

Final acceptance 15 August 2022

Available online 26 October 2022

MS. number: 22-00155

### Keywords:

artificial neural network  
collective learning  
consensus decision making  
emergent property  
leadership  
passenger–driver effect

Collective learning, the improvement of behaviours through experience of collective actions, is an area of animal learning that has received little attention. We investigated how individual learning during collective actions could produce improvements in collective performance, and how collective decision-making processes, including leadership dynamics, could impact upon learning. We trained artificial neural networks, either solo or paired, at an orientation task, based upon collective navigation in animals. In pairs, we implemented two rules of collective decision making: ‘democratic’ (weighted average of individual propositions) or ‘despotic’ (one individual’s proposition, determined randomly with weighted probabilities in each trial). Decision-making weightings were varied between pairs, but fixed for a given pair, with asymmetric weightings generating ‘leaders’ and ‘followers’. We found nearly all pairs improved their orientation, but more slowly than solo learners. Within pairs, leaders learnt more quickly than followers (‘the passenger–driver effect’). In democratic pairs, collective performance improved through individuals learning to compensate for partner error. This emergent process was not observed in pairs with despotic decision making, in which individuals learnt similarly to solo learners. Our model helps to clarify the links between individual learning, collective decision making and collective performance, in the context of collective navigation, and collective behaviour, more generally.

© 2022 The Authors. Published by Elsevier Ltd on behalf of The Association for the Study of Animal Behaviour. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

Many animal species are capable of associative learning (Bouton, 2007; Dukas, 1998; Pearce, 2008), a process through which individuals can improve their performance over repeated executions of a task. In associative learning, rewards or costs, received as a consequence of an individual’s behaviour, feed back to improve its future decisions and thus increase its expected net reward. Yet many animals live in groups and take part in collective actions (for example, collective movements), which result from the integration of several and sometimes divergent individual propositions (Conradt & List, 2009; Couzin et al., 2005). Surprisingly though, we still know little about how associative learning processes at the individual level interact with collective decision-making processes in group-living animals (Biro et al., 2016; Kao et al., 2014).

First, it is unclear how collective contexts affect individual learning, with most empirical studies on animal learning focusing

on individuals learning alone rather than while contributing to a collective decision (Biro et al., 2016; Kao et al., 2014). Even in social learning studies (focused upon individual learning in group contexts), tests are generally individually based, investigating whether an animal can successfully copy the behaviour of conspecifics to solve the same task individually (Biro et al., 2016; Heyes, 1994; Hoppitt & Laland, 2013). Yet it has been suggested that collective joint-action processes may fundamentally alter how and what individuals learn (Biro et al., 2016; Kao et al., 2014): reaching a consensus collective decision dilutes the relationship between an individual’s behavioural preference and the action taken (Conradt & Roper, 2005). The outcome of this collective action may be rewarding or costly, and act to reinforce learning; hence, an individual’s behavioural preference and the outcome of its behaviour are decoupled by the collective context, with potential implications for learning. Kao et al. (2014) demonstrated this, showing that identical learning algorithms learned different solutions depending on whether they acted alone or within a cohesive group making collective decisions. However, it remains unclear how well their

\* Corresponding authors.

E-mail addresses: [joe.morford@biology.ox.ac.uk](mailto:joe.morford@biology.ox.ac.uk) (J. Morford), [julien.collet@univ-lr.fr](mailto:julien.collet@univ-lr.fr) (J. Collet).

conclusions generalize to various contexts of animal collective learning, in particular contexts in which the optimal solution does not depend on whether individuals act alone or in groups.

Second, even if individuals do learn effectively from experience during collective actions, we know little about how or whether this leads to improvements in collective performance. This is partly because most theoretical models of collective decision making in animal groups have so far assumed memory-less group members (Biro et al., 2016; Kao et al., 2014). For a number of reasons, simple associative processes at the individual level may not necessarily be sufficient to produce improvements in collective performance when a group repeats a collective task. For instance, within a group, individuals may not learn at the same rate or learn the same things (de Perera & Guilford, 1999; Pettit et al., 2015); this may produce divergent actions between members of a group and so may not necessarily lead to improvements in collective performance (Conradt & List, 2009; Conradt & Roper, 2005; Couzin et al., 2005). Moreover, even if individuals do learn, this will not improve group performance unless these individuals also happen to be influential in the group decision-making process (Stroeymeyt et al., 2011). As a result, we may expect that in some cases the groups will perform less well than the most efficient member(s) would alone. In that regard, the rules by which individual preferences are combined to generate a collective decision (Biro et al., 2006; Conradt & Roper, 2005; Couzin et al., 2005) may be highly influential in whether and how individual learning generates improvements in collective performance.

On the other hand, we know empirically that at least some animal groups are capable of improving collective performance over repetitions of a collective task (homing pigeons, *Columba livia*, Flack et al., 2013; predator avoidance by fish, Hansen et al., 2021; nest emigration by ants, Langridge et al., 2004; migratory green surfing in ungulates, Jesmer et al., 2018), but the mechanisms underlying such collective learning have seldom been explored (Langridge et al., 2008). In theory, improvements in collective performance may be driven not merely by the effects of each group member learning to complete the task individually, but perhaps also by emergent properties arising from increasingly efficient interactions between group members. For instance, if a group member learns to perform a task only partially, or with some systematic error, other members of the group may simultaneously learn to compensate for it. Forms of organizational learning like this have been demonstrated experimentally in humans playing a simple cooperative game without direct communication: some dyads spontaneously divided their contribution to the task in a systematic way to reach higher performance (Andrade-Lotero & Goldstone, 2021). In some cases, the group may thus perform better than any of its members would on their own, especially if members learn more about compensating for others than about completing the task itself (Kao et al., 2014). While a few theoretical simulations suggest that this is possible (Andrade-Lotero & Goldstone, 2021; Kao et al., 2014), we still know little about the conditions that will lead (or not) to such emergent collective learning processes.

We utilized a navigational paradigm to model collective learning. This provides a simple and easily quantifiable model of a task solvable both solo and collectively, to test the implications of the collective context on learning. Our model is based upon navigational learning at a single site, whereby, over repeated visits, a solo animal or group of animals improves its orientation towards a fixed target from the site, such that it navigates more efficiently (Fig. 1). In the model, each neural network proposes a direction (from 0 to 360°). Neural networks either learnt this task alone (solo learners, Fig. 1a), or within pairs (collective learners, Fig. 1b, c). Solo neural network learners generated an output direction and ‘paid a

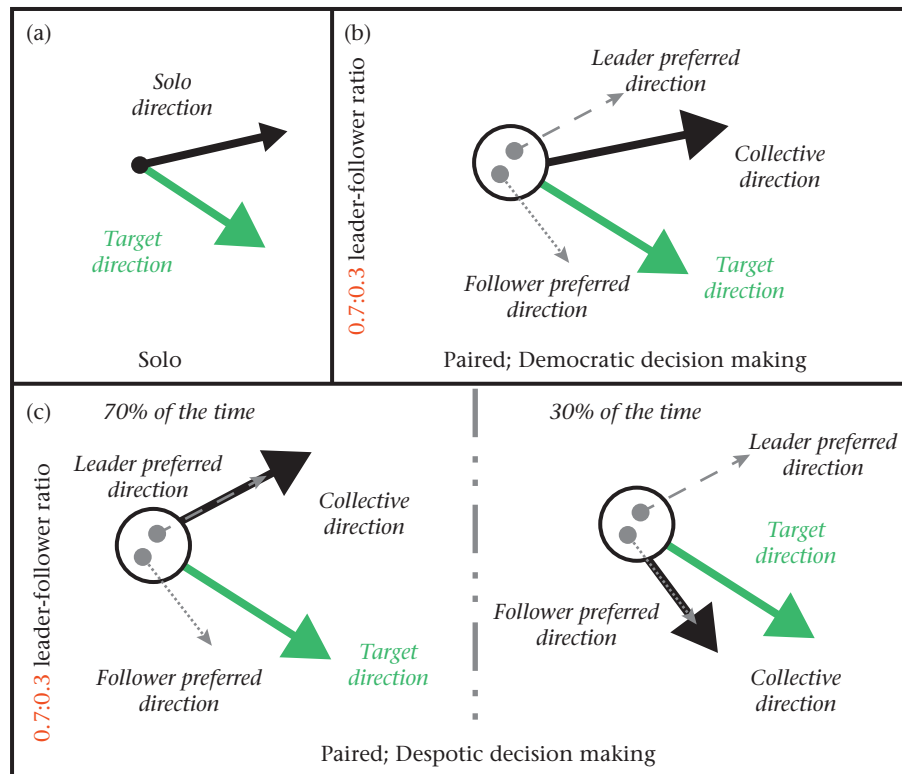
cost’ equal to the squared difference between its output and the ‘correct’ direction. This cost then fed into the learning of the network through backpropagation (Gulli & Pal, 2017). Within pairs, each neural network proposed a direction, and a single consensus direction of the pair was then determined through a collective decision-making function. The difference between the group (postconsensus) direction and the correct direction determined the cost for each individual member. As in solo learners, this cost fed into the learning of each member’s (preconsensus) proposition. This coupled individual learning with the collective outcome, a crucial element of collective learning.

For pairs, we implemented two types of collective decision-making function: either a weighted average of the two proposed directions (‘democratic’; Fig. 1b), or one of the individually proposed directions, determined randomly in accordance with weighted probabilities (‘despotic’; Fig. 1c). We explored the effects of different dynamics of leadership/followership by varying between pairs the weightings of members’ contributions to the consensus decision making. These two functions therefore allowed us to capture a wide range of consensus decision-making processes observed in animal groups and to vary both the extent to which decisions are shared between group members and the consistency of leadership within groups (Biro et al., 2006; Conradt & Roper, 2005, 2007). We expected that leaders might learn more quickly than followers, a phenomenon that has been termed the ‘passenger–driver’ effect in relation to empirical research (de Perera & Guilford, 1999). This is expected as leaders receive the most consistently appropriate reinforcement (cost) from the consequences of the collective action.

We assessed and compared the learning performance of pairs as a collective with different consensus rules and leader/follower weightings and made comparisons with the performance of solo learners. This allowed us to observe whether there was any ‘collective intelligence’ effect, with groups performing better than solo learners, and how this was affected by different decision-making processes. Within pairs, we examined what each individual member learnt to propose (‘preconsensus propositions’), to investigate how well individual members learnt the task within a collective, whether this depended on the leader–follower weighting (to test for the passenger–driver effect) and/or the type of collective decision rule (despotic or democratic). To test for emergent organizational forms of collective learning across different collective decision-making rules, we compared the performance of each individual’s preconsensus proposition (individual-preferred direction) with its performance after consensus (collective direction): ‘collective membership gain’ was observed if the individual performed better through the consensus than it would have without. This is highly related to the idea of ‘consensus costs’ (Conradt & Roper, 2005) that individuals pay by forgoing an optimal individual action to comply with the collective consensus. We qualitatively examined how our findings resemble empirical results on collective navigation, and collective behaviour more generally, focusing particularly on the collective navigation of homing pigeons, the best studied model species in relation to collective navigation.

## METHODS

Neural networks learnt the navigational task of returning an arbitrarily correct bearing either individually or collectively, in pairs. The neural network comprised multilayer perceptrons of only six neurons and a single dense hidden layer (four neurons) with a rectified linear unit (ReLU) activation function (Chollet, 2015; Gulli & Pal, 2017). In each task, neural networks were given a constant input of 1 and outputted a proposed direction between 0 and 360°, so that the task comprised simply the honing of the output, without



**Figure 1.** Model outline. Our model is based upon an orientation decision towards a target from a single site in solo or collectively navigating animals. (a) Solo learners, with the output direction of the learner (Solo direction) and the correct direction (Target direction) shown. (b, c) Paired learners with 0.7:0.3 decision-making weightings, defining a leader and a follower. Their collective behavioural output is determined by a decision-making rule: either (b) democratic or (c) despotic. The individual output direction of each learner is shown (Leader preferred direction and Follower preferred direction), as well as the collective output (Collective direction). In all cases, learning relies on the 'cost paid', which feeds back into neural network learning and is determined by the difference between the overall output direction (Solo direction or Collective direction) and the correct direction (Target direction). Over repeated trials, the solo and paired learners could improve their output orientations, getting closer to the correct target direction, through learning.

changing the input. The output was generated with a linear activation function (units: degrees/360) and, in pairs, the two outputs were combined using a collective decision rule to generate a single consensus bearing. In each trial, the cost paid (the loss function) was equal to the squared error in orientation ( $0-180^\circ$ ) of the solo or consensus orientation. In pairs, this cost was applied equally to each member and in all neural networks, the cost was used to optimize the orientation in subsequent trials through standard gradient descent (learning rate = 0.05; momentum = 0; decay = 0). Training involved a single learning trial (training datapoint) at a time (batch size = 1, epochs = 1). The networks were implemented in Python (Van Rossum & Drake Jr, 1995), with libraries keras (Chollet, 2015), sys, numpy (Harris et al., 2020), scipy (Virtanen et al., 2020), matplotlib (Hunter, 2007) and tensorflow (Abadi et al., 2016).

During collective learning, two consensus decision-making rules were used: a democratic (averaging) decision-making rule and a despotic (probabilistic) decision-making rule. The democratic decision-making rule comprised a weighted circular mean of the two individual output directions. The despotic decision-making rule comprised randomly selecting one of the individual output directions using weighted probabilities. In the despotic instance, whether an individual led or followed on a given trial was not input into their learning process. The relative weightings in both democratic and despotic instances were set to 0.5:0.5, 0.7:0.3 or 0.9:0.1 for a given pair, and each individual retained its relative weighting throughout the learning process. Individuals in a pair would therefore either input equally into consensus decision-making (0.5:0.5) or would act as a leader and follower (0.7:0.3 and

0.9:0.1). Henceforth, 'leader' is used to mean an individual with a weighting of greater than 0.5 and 'follower' to mean an individual with a weighting less than 0.5.


Two hundred and fifty neural networks were trained to complete the task solo, each learning over 25 learning trials (training datapoints). Similarly, for each consensus decision rule and each leadership ratio, 250 pairs of neural networks were trained over 25 trials. The networks were tested first after model initialization but before any training, and then after every learning trial. Comparative tests of performance were made using pairwise Mann–Whitney *U* tests after five learning trials (to assess learning rates) and at the conclusion of learning, after 25 learning trials (to assess final asymptotic learning performance). Statistical analysis and graphical output were produced using R (R Core Team, 2018; RStudio Team, 2022), including package scales (Wickham, 2018).

## RESULTS

### Overall Solo and Collective Performance

Over 25 learning trials, all the solo learners ( $N = 250$ ) improved at the navigational task, with better performance after the final learning trial than before the first. Median performance is shown in Fig. 2a.

Similarly, almost all (741 of 750) of the pairs of neural networks using the democratic decision-making rule improved in collective performance during training (performed better in the last trial than the first). Learning was quicker and final performance was better when pair members contributed less equally to decision making,



**Figure 2.** Overall performance of solo and paired learners. Average learning curves are shown for (a) solo learners and (b, c) pairs of networks learning collectively through (b) a democratic rule or (c) a despotic rule. (a) Median absolute angular error of 250 learners across trials (black curve). (b, c) Three median absolute angular errors of pairs after their consensus (collective performance), each of 250 pairs of neural networks. The average learning curve of solo learners (a) is also shown in black in (b) and (c) for comparison.

such that one member was a clear leader and the other a clear follower (Fig. 2b). At the conclusion of training, performance was significantly better in pairs with a 0.9:0.1 decision-making weighting than in pairs with 0.7:0.3 and equal (0.5:0.5) decision-making weightings (Mann–Whitney  $U$  tests:  $P < 0.0001$  in both cases); additionally, performance was significantly better in pairs with a 0.7:0.3 decision-making weighting than pairs with an equal decision-making weighting (Mann–Whitney  $U$  test:  $P < 0.001$ ). These differences in performance between pairs with different leadership/followership ratios were already detectable after only five learning trials. Performance of solo learners was significantly better than the collective performance of paired learners with decision-making weightings of 0.5:0.5, 0.7:0.3 and 0.9:0.1 both after five learning trials (Mann–Whitney  $U$  tests:  $P < 0.0001$ ,  $P < 0.0001$  and  $P = 0.024$ , respectively) and at the conclusion of training (Mann–Whitney  $U$  tests:  $P < 0.0001$  in all cases).

Of 750 pairs of learners using the despotic decision-making rule, 730 improved in collective performance during training (performed better in the last trial than the first). Again, learning was quicker and final performance was better when pairs contributed less equally, with stronger leaders and followers. After five learning trials and at the conclusion of training, performance was significantly better in pairs with a 0.9:0.1 decision-making weighting than pairs with a 0.7:0.3 or 0.5:0.5 weighting (Mann–Whitney  $U$  tests:  $P < 0.0001$ , in all cases); however, there was no significant difference between pairs with 0.5:0.5 and 0.7:0.3 weightings (Mann–Whitney  $U$  tests:  $P = 0.942$  after five trials;  $P = 0.216$  at conclusion of training; Fig. 2c). Both after five trials and at the conclusion of training, performance by solo learners was significantly better than the collective performance of paired learners with all decision-making weightings (Mann–Whitney  $U$  tests:  $P < 0.02$  in all cases).

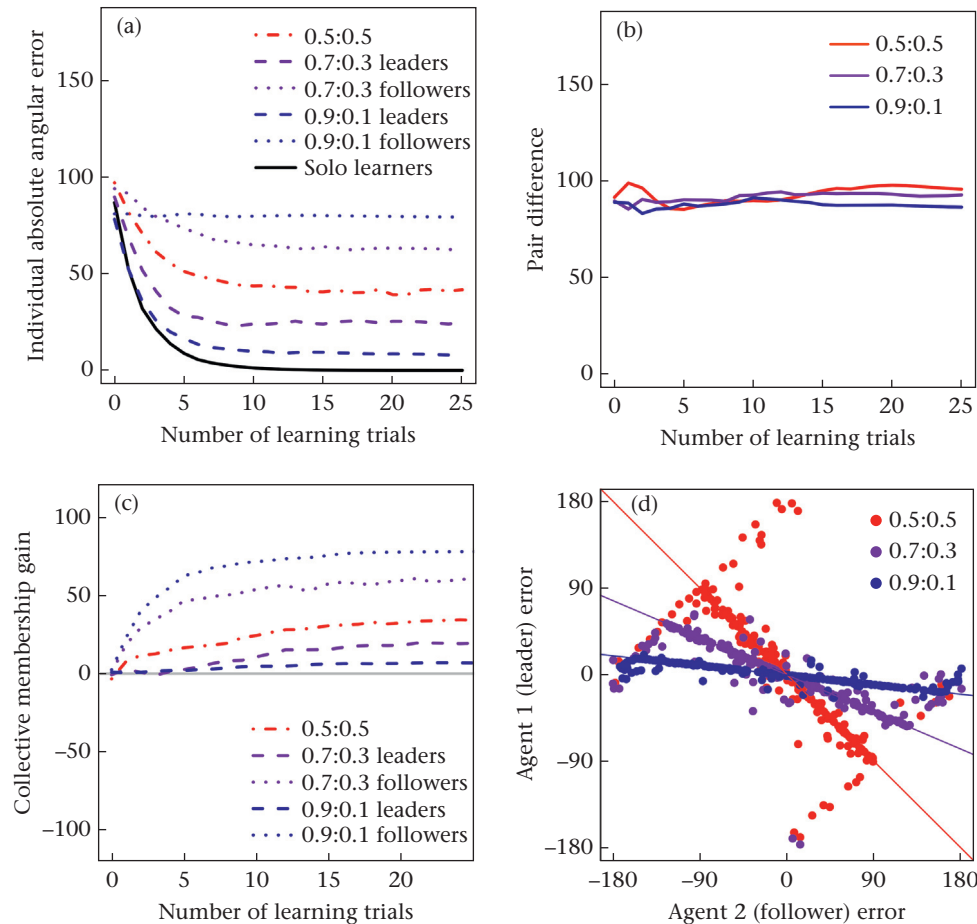
### Democratic Decision Making

In pairs of learners using a democratic decision-making rule, the error of each individual progressively reduced, but the median error appeared to quickly plateau at levels well above zero (Fig. 3a). Individual error at the conclusion of learning was significantly greater in networks learning in pairs with democratic decision-making rules than in solo learners, irrespective of decision-making weighting (Mann–Whitney  $U$  tests:  $P < 0.0001$  in all cases). Errors plateaued at a significantly higher level in learners more prone to follow than to lead (Mann–Whitney  $U$  tests: at final trial of learning,  $P < 0.01$  for all pairwise

combinations) but in neither leaders nor followers did the error appear to be approaching zero (Fig. 3a). On average, members of a pair did not learn to reduce the difference between their respective output angles ('pair difference'; Fig. 3b), with average pair difference remaining approximately level across trials, and no significant difference between the pair difference in the first and last trials, irrespective of the decision-making weightings of the pair (Mann–Whitney  $U$  tests:  $P = 0.289$ ,  $P = 0.164$  and  $P = 0.707$  for pairs with 0.5:0.5, 0.7:0.3 and 0.9:0.1 leadership weightings, respectively). Hence, the increase in collective performance was not achieved by both members of a pair converging on the correct solution. This indicates that the collective context changed the nature of the solution that individual learners were reaching.

The difference between an individual's error considered alone and the collective performance of its group is termed 'collective membership gain' if collective error is smaller than individual error, or 'collective membership loss' otherwise. During learning, there was an average increase in collective membership gain for democratic paired learners (Fig. 3c), with significantly greater collective membership gain in the final learning trial than in the first, whether they were leaders (Mann–Whitney  $U$  tests:  $P < 0.0001$  for both pairs with 0.7:0.3 and 0.9:0.1 leadership weightings), followers (Mann–Whitney  $U$  tests:  $P < 0.0001$  for both pairs with 0.7:0.3 and 0.9:0.1 leadership weightings), or individuals contributing equally to decision making (Mann–Whitney  $U$  test:  $P < 0.0001$ ). Collective membership gain at the conclusion of learning was, on average, greater in individuals that contributed the least to decision making (Mann–Whitney  $U$  tests:  $P < 0.0001$  for all pairwise comparisons). To understand how this collective membership gain emerged within pairs we quantified the correlation between the errors of the two individuals within a pair. We predicted that if individuals were learning to compensate for the error of their partner, there would be negative correlations between the errors of the two individuals within a pair (e.g. a large anticlockwise error by one member of a 0.5:0.5 pair could be compensated by an equally large clockwise error by its partner). We therefore regressed the errors of the leaders on the errors of the followers at the conclusion of training (or, for the pairs with equally contributing members, split each pair randomly and performed the regression). We found negative relationships in each instance (linear regressions:  $P < 0.0001$  in all cases). These correlations appeared to show excellent qualitative fit to the theoretical gradient expectations of  $-1$ ,  $-3/7$  and  $-1/9$  in pairs with decision-making weightings of 0.5:0.5, 0.7:0.3 and 0.9:0.1, respectively, as shown in Fig. 3d.





**Figure 3.** Democratic decision making: performance and contribution of each individual member. (a) Median absolute angular error of each pair member's output before consensus (coloured curves), compared to the median error observed in solo learners (black curve), across training trials. (b) Median within-pair difference between members' output directions before consensus, across training trials. (c) Collective membership gain across trials: the difference between the collective performance of a pair and the performance of each of its members considered individually (individual error – collective error). (d) Points show the relationships between directions proposed by members of the same pair at the last (25th) learning trial; lines show the theoretical gradient expectations of -1, -3/7 and -1/9 in pairs with decision-making weightings of 0.5:0.5, 0.7:0.3 and 0.9:0.1, respectively.

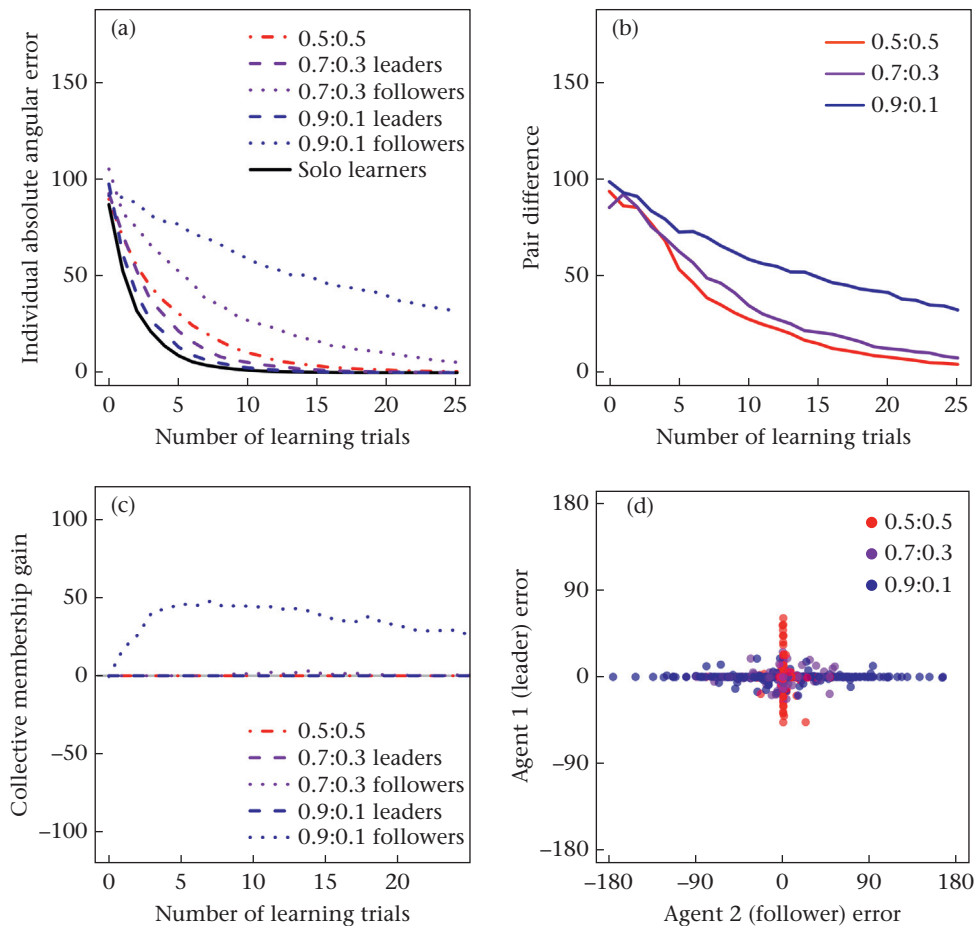
### Despotic Decision Making

Conversely, in pairs using a despotic decision-making rule, the error of each individual appeared to approach zero across trials in all cases (Fig. 4a). Individual error at the conclusion of learning was significantly greater than in solo learners for followers (Mann–Whitney  $U$  tests:  $P < 0.0001$  in both cases), equally contributing individuals (Mann–Whitney  $U$  test:  $P < 0.0001$ ) and leaders (Mann–Whitney  $U$  tests:  $P < 0.0001$  and  $P = 0.027$  for pairs with 0.7:0.3 and 0.9:0.1 decision-making weightings, respectively). Individual error fell more quickly in leaders than in individuals contributing equally to decision making, and more slowly still in followers (Mann–Whitney  $U$  tests:  $P < 0.02$  for all pairwise combinations after five learning trials). The pair difference (the difference in output between the members of a pair before consensus) decreased during learning (Fig. 4b), with significantly lower pair difference in the last trial than the first for pairs with all three decision-making weightings (Mann–Whitney  $U$  tests:  $P < 0.0001$  in all three cases). During training, collective membership gain increased significantly in followers with 0.7:0.3 and 0.9:0.1 decision-making weightings (Mann–Whitney  $U$  tests:  $P < 0.02$  and  $P < 0.0001$ , respectively), and decreased in leaders (Mann–Whitney  $U$  tests:  $P < 0.01$  and  $P < 0.0001$  for the two decision-making weightings, respectively). However, collective gain after training in the majority of individuals was zero, and hence these

changes were not reflected by the median collective membership gain (shown in Fig. 4c). There was no change in the collective membership gain during training for individuals in equally contributing pairs (Mann–Whitney  $U$  test:  $P = 0.940$ ). Finally, no significant relationships were found between the errors of leaders and followers (linear regressions:  $P = 0.501$  and  $P = 0.990$  for pairs with 0.7:0.3 and 0.9:0.1 weightings, respectively), or randomly split pairs in equally contributing pairs (linear regression:  $P = 0.566$ ), providing no evidence for organizational learning in pairs with probabilistic decision-making rules (Fig. 4d). Hence, learning in this collective context, with a despotic decision-making rule, seemed to be qualitatively similar to learning in a solo context. In particular, the final solution reached by the learners converged upon the correct solution, as in a solo learning context, with the rate of learning determined by the decision-making weightings.

### DISCUSSION

Our artificial simulations of associative learning processes showed that most pairs of neural networks increased their performance through a repeated orientation task. However, the collective context altered the rate of learning and final performance of both the group itself and its individual members. Overall, we found that pairs and their members learnt more slowly than solo learners. The type of consensus decision-making rule of a pair affected the



**Figure 4.** Despotic decision making: performance and contribution of each individual member. (a) Median absolute angular error of each pair member's output before consensus (coloured dashed curves), compared to the median error observed in solo learners (black curve), across training trials. (b) Median within-pair difference between members' output directions before consensus, across training trials. (c) Collective membership gain across trials: the difference between the collective performance of a pair, and the performance of each of its members considered individually (individual error – collective error). (d) Points show the relationships between directions proposed by members of the same pair at the last (25th) learning trial.

processes through which they improved their collective performance. Democratic groups increased performance by improving the complementarity of their member's contributions, giving rise to an emergent form of collective learning; however, despotic groups improved entirely through individual improvements. The degree of leadership in pairs affected both the individual learning rate, with leaders learning faster under most circumstances, and the rate of collective learning (pairs with greater asymmetry in decision-making contribution tended to learn faster).

The unequal learning rates of leaders and followers can be termed the 'passenger–driver effect' (de Perera & Guilford, 1999), in which individuals contributing more strongly to the collective decision learn more quickly. The passenger–driver effect has been observed in the individual navigational learning of pigeons within flocks (Pettit et al., 2015). Our model suggests that this can arise because individuals contributing most (leaders) to the collective decision receive the most consistently appropriate reinforcement (cost) from the consequences of the collective action. Here, this was true in both democratic and despotic groups through slightly different mechanisms. In democratic groups, the consensus collective action is determined through a weighted averaging of the propositions of the leader and follower. This generates noise in the relationship between an individual's proposition and the reinforcement it receives, which depends on the collective action. The noise in this relationship slows learning and slows the learning of

followers to a greater extent than that of leaders, as the leader's proposed direction is always closer to the collective action than the follower's proposition. Conversely, in despotic groups, each individual determines the collective action with a given probability. In trials in which they lead, they learn as if performing the task solo, whereas, when following, no learning takes place as there is no relationship between the individual proposition and the collective action. This effectively reduces the number of trials in which learning can take place, and leaders (individuals that lead more often) therefore learn more quickly than followers as they learn on a greater proportion of trials. Somewhat similarly, in the collective learning model of Kao et al. (2014), individuals only learn about cues when they indicate the same discrete option as chosen by the group. This might be expected to slow learning in collective contexts relative to solo contexts, and to slow the learning of followers relative to leaders, by reducing the number of trials in which learning occurs.

The slower learning of individuals in both democratic and despotic groups in this model precluded the possibility of any 'collective intelligence' effect. Initially both solo learners and pairs performed equally well (orienting randomly), and solo learners outperformed pairs after learning, although would be expected to plateau at the same level, after sufficient trials. This contrasts with the modelling results of Kao et al. (2014), in which groups could outperform individuals in various simulated scenarios by

successfully exploiting cues with a low reliability for individuals, by averaging out the errors of group members. This is a manifestation of the ‘many-wrongs principle’ (Simons, 2004), a driver of collective intelligence. Further modelling work (Falcón-Cortés et al., 2019) shows how collective intelligence can emerge in foraging tasks through information transfer between individuals. Additionally, in empirical research, animal groups have often been observed to outperform solo individuals (Conradt & Roper, 2005; Simons, 2004), including in navigational contexts (Sasaki & Biro, 2017; Tamm, 1980), although not in all cases (Guilford & Chappell, 1996; Keeton, 1970). These collective intelligence effects may derive from perceptual errors or execution errors that are independent between individuals in a group and are not captured in our simple model.

The mechanism by which collective learning occurred was highly dependent on the collective decision-making rule. An organizational form of learning emerged in pairs adopting a democratic rule of weighted average between member propositions, but not in pairs adopting a despotic rule to determine which partner had total control in each trial. In democratic pairs, each member could learn to compensate for the error of its partner, proposing directions with error in the opposite direction to the error of its partner. With the rare exception of pairs not improving their collective performance, each democratic pair thus found its own idiosyncratic equilibrium between members (a form of ‘convention’; Stephens & Heinen, 2018), such that the average error of the collective approached zero. As a result, collective accuracy was, on average, better than the propositions of either individual member. Hence, if an individual that had learned as part of a group subsequently had to complete the task alone, their decision would, in almost all cases, be less accurate than the group’s consensus decision. This is collective membership gain, with individuals performing better as members of a collective than alone. This represents an emergent property of collective learning: the collective context altered not only the learning rate of individual members, but also the nature of the solution upon which individuals converged (as in Kao et al., 2014, but in a context where the optimal solution was independent of the social context). These results highlight the potentially complex relationship between individual and collective learning processes.

Our model includes a number of assumptions that may appear unrealistic of collective dynamics in navigating animals. First, for simplicity, we forced pairs to remain cohesive, even if they had highly divergent directional preferences. However, in animal groups, if there is a large conflict of interest and therefore ‘consensus cost’ (Conradt & Roper, 2005), then groups are less likely to come to a consensus decision and may split. For instance, in pigeons, paired individuals often split, especially on first releases when partners are unfamiliar with the homing route (Flack et al., 2013; Guilford & Chappell, 1996), or when they have very divergent directional preferences (Biro et al., 2006). Nevertheless, once each bird, flying solo, converges towards the most direct route, cohesion is often restored and so a phase of solo learning might be followed by collective learning in these instances. While such complications will have some effect upon the outcomes and processes of learning, they are not central to the interface between collective decision making and learning and hence were not included in our model. Second, we imposed a fixed ratio of leadership between paired individuals. In reality, this might itself change through learning as groups of individuals gain experience: leadership could vary randomly between trials, or some individuals may learn to lead more than others, and this could be contingent on the accuracy of each individual’s proposition. In the latter case, there might be feedback between learning and leadership, whereby individuals that learn fastest become leaders and,

once leaders, are able to learn faster still. In our model, learning and the accuracy of past decisions did not affect leadership and faster learning was only a consequence not a cause of leadership. However, a stable leadership ratio appears relatively consistent with various examples of leadership in animal groups, in which stable individual attributes such as age, size and boldness (Beauchamp, 2000; Fischhoff et al., 2007; Pettit et al., 2015; Sasaki et al., 2018) have been shown to influence the level of contribution to collective decisions.

Third, one potential criticism of our model implementation is that the cost paid, feeding back into learning, is unrealistic: in an orientation task, an animal, or group of animals, could not possibly know its precise level of error without knowing the correct orientation. Furthermore, navigational tasks (and other behaviours) are more complex than a single orientation decision, so the reinforcement generated from completing the task (i.e. homing) will not relate perfectly to the error of the initial orientation. However, any reinforcement that an animal, or group of animals, could use to learn to improve an orientation task would likely show a strong relationship with the orientation error. For instance, an animal could use the length of time it took to reach a goal as a measure of its orientation performance. Providing that there is a relationship between this measure of performance and its actual orientation error, learning will operate similarly in our model as in this more realistic potential scenario. We expect that the process of learning reinforcement will be strongly related to task performance in solo and collective behavioural tasks more generally, and hence this simplifying assumption is unlikely to generate unrealistic modelling results.

A final limitation of our model may be that the neural networks are relatively constrained in their learning in our modelling environment, and cannot, for instance, learn about or remember the collective action in previous training trials. In contrast, it might seem possible for individuals to simply remember and recapitulate behaviours previously executed within groups, and this is potentially true in homing pigeons (Pettit et al., 2013; Sasaki & Biro, 2017). Similarly, individuals could learn within groups through social learning mechanisms (Heyes, 1994; Hoppitt & Laland, 2013), for instance the follower within a pair observing and imitating the behaviour of the leader. Considering social learning within groups may complicate the predictions of our model. For instance, the passenger–driver effect could be lessened or disappear through the followers learning from leaders. Additionally, given that social learning can occur within groups and not in solo learners, social learning might contribute towards a collective intelligence effect, facilitating better performance in groups than in solo learners, contrary to the results of our model.

On the other hand, an improvement in the proposition of a group member at the individual level through recapitulation of a previous collective output or imitation of another individual could counterintuitively worsen performance at the group level, through a failure to compensate for the errors of other group members. Furthermore, in many cases it may not be simple or even possible for an individual within a group to be able to remember and execute a collective behavioural output. This may be because collective behavioural outputs have many inputs from the propositions of group members themselves responding to various cues, making the behaviour difficult to perceive and replicate for an individual member of the group. Alternatively, complex cooperative collective behaviours such as group foraging behaviours (Stander, 1992) or the biparental care of offspring may involve a spatiotemporal separation of cooperative individuals. Nevertheless, the collective performance (the number of prey caught by the group, or the condition of the offspring) may feed back into the learning of the individual group members.

Overall, our model helps clarify the complex links between individual learning and collective decision making. Our results highlight that individual associative processes can lead to improvements in group level performance through experience, both through an individual increase in performance and through organizational learning, where group members here learnt to compensate for the errors of others, even without explicit rules for them to learn about each other. Future research could explore in more depth how the collective and individual learning properties exposed in our study are affected by biologically realistic changes in model parameters, model complexity and/or model tasks. Additionally, empirical focus on the interaction between collective decision-making processes and individual learning would allow the assumptions and predictions of our model to be tested and would develop understanding of animal learning in collective contexts.

## Author Contributions

**Joe Morford:** Conceptualization, Methodology, Software, Formal analysis, Investigation, Writing – Original Draft, Writing – Review and Editing, Visualization. **Patrick Lewin:** Methodology, Writing – Original Draft, Writing – Review and Editing, Visualization. **Dora Biro:** Writing – Review and Editing, Supervision. **Tim Guilford:** Writing – Review and Editing, Supervision. **Oliver Padgett:** Methodology, Supervision. **Julien Collet:** Conceptualization, Investigation, Writing – Original Draft, Writing – Review and Editing, Supervision.

## Data Availability

The code and data used in this research is available on GitHub: <https://github.com/Morfordjoe/Collective-learning-neural-network-model>, and at request from the lead author.

## Declaration of Interest

The authors have no conflicts of interest.

## Acknowledgments

J.M. was supported by funding from the Biotechnology and Biological Sciences Research Council (BBSRC) [grant number BB/M011224/1]. Financial support for D.B. and J.C. was provided by the Templeton World Charity Foundation's 'Diverse Intelligences' scheme [grant number TWCF0316]. O.P. was funded by a Junior Research Fellowship from St John's College, Oxford. T.G.'s research was supported by Merton College, University of Oxford, and by the Mary Griffiths award. We thank Theresa Burt de Perera and Alex Kacelnik for feedback on the ideas and modelling covered in this manuscript, Louis Morford for proof reading, and two referees and the editor for their feedback and comments on the manuscript.

## References

- Abadi, M., Barham, P., Chen, J. M., Chen, Z. F., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., Kudlur, M., Levenberg, J., Monga, R., Moore, S., Murray, D. G., Steiner, B., Tucker, P., Vasudevan, V., Warden, P., ... Zheng, X. (2016). TensorFlow: A system for large-scale machine learning. In *Proceedings of OSDI '16: 12th USENIX Symposium on Operating Systems Design and Implementation* (pp. 265–283).
- Andrade-Lotero, E., & Goldstone, R. L. (2021). Self-organized division of cognitive labor. *PLoS One*, 16(7), Article e0254532. <https://doi.org/10.1371/journal.pone.0254532>
- Beauchamp, G. (2000). Individual differences in activity and exploration influence leadership in pairs of foraging zebra finches. *Behaviour*, 137, 301–314. <https://doi.org/10.1163/156853900502097>
- Biro, D., Sasaki, T., & Portugal, S. J. (2016). Bringing a time-depth perspective to collective animal behaviour. *Trends in Ecology & Evolution*, 31(7), 550–562. <https://doi.org/10.1016/j.tree.2016.03.018>
- Biro, D., Sumpter, D. J., Meade, J., & Guilford, T. (2006). From compromise to leadership in pigeon homing. *Current Biology*, 16(21), 2123–2128. <https://doi.org/10.1016/j.cub.2006.08.087>
- Bouton, M. E. (2007). *Learning and behavior: A contemporary synthesis* (2nd ed.). Sinauer Associates.
- Chollet, F. (2015). keras, GitHub. <https://github.com/fchollet/keras>.
- Conradt, L., & List, C. (2009). Group decisions in humans and animals: A survey. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 364(1518), 719–742. <https://doi.org/10.1098/rstb.2008.0276>
- Conradt, L., & Roper, T. J. (2005). Consensus decision making in animals. *Trends in Ecology & Evolution*, 20(8), 449–456. <https://doi.org/10.1016/j.tree.2005.05.008>
- Conradt, L., & Roper, T. J. (2007). Democracy in animals: The evolution of shared group decisions. *Proceedings of the Royal Society B: Biological Sciences*, 274(1623), 2317–2326. <https://doi.org/10.1098/rspb.2007.0186>
- Couzin, I. D., Krause, J., Franks, N. R., & Levin, S. A. (2005). Effective leadership and decision-making in animal groups on the move. *Nature*, 433(7025), 513–516. <https://doi.org/10.1038/nature03236>
- de Perera, T. B., & Guilford, T. (1999). The social transmission of spatial information in homing pigeons. *Animal Behaviour*, 57, 715–719.
- Dukas, R. (1998). *Cognitive ecology: The evolutionary ecology of information processing and decision making*. University of Chicago Press.
- Falcón-Cortés, A., Boyer, D., & Ramos-Fernández, G. (2019). Collective learning from individual experiences and information transfer during group foraging. *Journal of the Royal Society Interface*, 16(151), Article 20180803.
- Fischhoff, I. R., Sundaresan, S. R., Cordingley, J., Larkin, H. M., Sellier, M. J., & Rubenstein, D. I. (2007). Social relationships and reproductive state influence leadership roles in movements of plains zebra, *Equus burchellii*. *Animal Behaviour*, 73, 825–831. <https://doi.org/10.1016/j.anbehav.2006.10.012>
- Flack, A., Freeman, R., Guilford, T., & Biro, D. (2013). Pairs of pigeons act as behavioural units during route learning and co-navigational leadership conflicts. *Journal of Experimental Biology*, 216(Pt 8), 1434–1438. <https://doi.org/10.1242/jeb.082800>
- Guilford, T., & Chappell, J. (1996). When pigeons home alone: Does flocking have a navigational function? *Proceedings of the Royal Society B: Biological Sciences*, 263(1367), 153–156. <https://doi.org/10.1098/rspb.1996.0024>
- Gulli, A., & Pal, S. (2017). *Deep learning with Keras*. Packt Publishing.
- Hansen, M. J., Burns, A. L., Monk, C. T., Schutz, C., Lizier, J. T., Ramnarine, I., Ward, A. J. W., & Krause, J. (2021). The effect of predation risk on group behaviour and information flow during repeated collective decisions. *Animal Behaviour*, 173, 215–239. <https://doi.org/10.1016/j.anbehav.2021.01.005>
- Harris, C. R., Millman, K. J., van der Walt, S. J., Gommers, R., Virtanen, P., Cournapeau, D., Wieser, E., Taylor, J., Berg, S., Smith, N. J., Kern, R., Picus, M., Hoyer, S., van Kerkwijk, M. H., Brett, M., Haldane, A., Fernández del Río, J., Wiebe, M., Peterson, P., & Oliphant, T. E. (2020). Array programming with NumPy. *Nature*, 585(7825), 357–362. <https://doi.org/10.1038/s41586-020-2649-2>
- Heyes, C. M. (1994). Social learning in animals: Categories and mechanisms. *Biological Reviews of the Cambridge Biological Society*, 69(2), 207–231. <https://doi.org/10.1111/j.1469-185x.1994.tb01506.x>
- Hoppitt, W., & Laland, K. N. (2013). *Social learning: An introduction to mechanisms, methods, and models*. Princeton University Press.
- Hunter, J. D. (2007). Matplotlib: A 2D graphics environment. *Computing in Science & Engineering*, 9(3), 90–95. <https://doi.org/10.1109/mcse.2007.55>
- Jesmer, B. R., Merkle, J. A., Goheen, J. R., Aikens, E. O., Beck, J. L., Courtemanch, A. B., Hurley, M. A., McWhirter, D. E., Miyasaki, H. M., Monteith, K. L., & Kauffman, M. J. (2018). Is ungulate migration culturally transmitted? Evidence of social learning from translocated animals. *Science*, 361(6406), 1023–1025. <https://doi.org/10.1126/science.aat0985>
- Kao, A. B., Miller, N., Torney, C., Hartnett, A., & Couzin, I. D. (2014). Collective learning and optimal consensus decisions in social animal groups. *PLoS Computational Biology*, 10(8), Article e1003762. <https://doi.org/10.1371/journal.pcbi.1003762>
- Keeton, W. T. (1970). Comparative orientational and homing performances of single pigeons and small flocks. *Auk*, 87(4), 797–799. <https://doi.org/10.2307/4083715>
- Langridge, E. A., Franks, N. R., & Sendova-Franks, A. B. (2004). Improvement in collective performance with experience in ants. *Behavioral Ecology and Sociobiology*, 56(6), 523–529. <https://doi.org/10.1007/s00265-004-0824-3>
- Langridge, E. A., Sendova-Franks, A. B., & Franks, N. R. (2008). How experienced individuals contribute to an improvement in collective performance in ants. *Behavioral Ecology and Sociobiology*, 62(3), 447–456. <https://doi.org/10.1007/s00265-007-0472-5>
- Pearce, J. M. (2008). *Animal learning & cognition: An introduction* (3rd ed.). Psychology Press.
- Pettit, B., Akos, Z., Vicsek, T., & Biro, D. (2015). Speed determines leadership and leadership determines learning during pigeon flocking. *Current Biology*, 25(23), 3132–3137. <https://doi.org/10.1016/j.cub.2015.10.044>
- Pettit, B., Perna, A., Biro, D., & Sumpter, D. J. (2013). Interaction rules underlying group decisions in homing pigeons. *Journal of the Royal Society Interface*, 10(89), Article 20130529. <https://doi.org/10.1098/rsif.2013.0529>
- R Core Team. (2018). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing. <https://www.R-project.org/>.
- RStudio Team. (2022). *RStudio: Integrated Development for R*. RStudio. <http://www.rstudio.com/>.



- Sasaki, T., & Biro, D. (2017). Cumulative culture can emerge from collective intelligence in animal groups. *Nature Communications*, 8, Article 15049. <https://doi.org/10.1038/ncomms15049>
- Sasaki, T., Mann, R. P., Warren, K. N., Herbert, T., Wilson, T., & Biro, D. (2018). Personality and the collective: Bold homing pigeons occupy higher leadership ranks in flocks. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 373(1746). <https://doi.org/10.1098/rstb.2017.0038>
- Simons, A. M. (2004). Many wrongs: The advantage of group navigation. *Trends in Ecology & Evolution*, 19(9), 453–455. <https://doi.org/10.1016/j.tree.2004.07.001>
- Stander, P. E. (1992). Cooperative hunting in lions—The role of the individual. *Behavioral Ecology and Sociobiology*, 29(6), 445–454.
- Stephens, D. W., & Heinen, V. K. (2018). Modeling nonhuman conventions: The behavioral ecology of arbitrary action. *Behavioral Ecology*, 29(3), 598–608. <https://doi.org/10.1093/beheco/ary011>
- Stroeymeyt, N., Franks, N. R., & Giurfa, M. (2011). Knowledgeable individuals lead collective decisions in ants. *Journal of Experimental Biology*, 214(Pt 18), 3046–3054. <https://doi.org/10.1242/jeb.059188>
- Tamm, S. (1980). Bird orientation — Single homing pigeons compared with small flocks. *Behavioral Ecology and Sociobiology*, 7(4), 319–322. <https://doi.org/10.1007/bf00300672>
- Van Rossum, G., & Drake, F. L., Jr. (1995). *Python reference manual*. Centrum voor Wiskunde en Informatica Amsterdam.
- Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S. J., Brett, M., Wilson, J., Millman, K. J., Mayorov, N., Nelson, A. R. J., Jones, E., Kern, R., Larson, E., & SciPy 1.0 Contributors. (2020). SciPy 1.0: Fundamental algorithms for scientific computing in Python. *Nature Methods*, 17(3), 261–272. <https://doi.org/10.1038/s41592-019-0686-2>
- Wickham, H. (2018). *scales: Scale Functions for Visualization*. <https://CRAN.R-project.org/package=scales>.