FishEye: Collaborative Driving by Consensus Decision Making

Pratik Gujjar[1] and Richard Vaughan[1]

***Abstract*— Behavioral psychology characterizes every individual with a set of preferences. Groups exhibit emergent capabilities like that of schooling in fish or flocking in birds by integrating disparate preferences. The paper presents our investigation of swarm intelligence to evolve an autonomous vehicle convoying behaviour. Our results demonstrate the consensus achieved by vehicle convoys to manoeuvre traffic lanes on highways. Convoys are shown to cruise at the maximum system speed, enhancing highway throughput and delivering optimal performance per vehicle.**

I. Introduction

Consensus in a multi-agent environment is a critical component since it enables coherent execution of a set of tasks. Such a decision making process is embodied in our democratic society leading to the inference that a majority vote leads to favorable decisions [1] [2]. Decisions made with mutual consensus are more reliable as they factor in information gathered by all individual agents. Isolated information when cascaded through the agents renders more knowledge than individually achieved with discounted effort. Animal groups such as birds and fish are known to employ these mechanisms in their emergent behaviours. Perhaps, a more pertinent example would be of the foraging bees. Worker bees *dance out* information about the produce of a source and a direction for its reach to other bees at the beehive [3]. The consensus by this information exchange benefits the hive in the form of efficient task allocation to the worker bees.

About one-quarter of fish species are known to *school* their entire lives. Fishes derive many benefits from schooling including defence from predators, enhanced foraging successes and increased chances of finding a potential mate. By positive feedback resulting from copying others, fish communicate and achieve a schooling consensus [4]. Schooling in multi-robot applications affords many advantages such as autonomous response to external variations in order to preserve the formation, resultant parallel velocity vectors and dynamic election of a leader. Also, converging with a successful school endows the rewards of the school to a *naïve* robot.

In this paper we emulate the *copy-feedback* mechanism of schooling to enable heterogeneous robots to form separate schools. The decision rule for distinct schooling relates a personal preference of flock to social information derived from flock size. Robots coloured red and blue employ this decision rule to achieve school consensus. We then show

[1]Pratik Gujjar and Richard Vaughan are with the Department of Computing Science, Simon Fraser University, Canada pgujjar@sfu.ca vaughan@sfu.ca

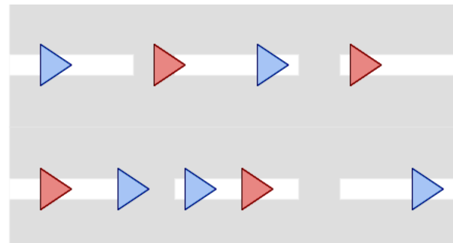

Fig. 1: Random positioning of vehicles on a two-lane highway

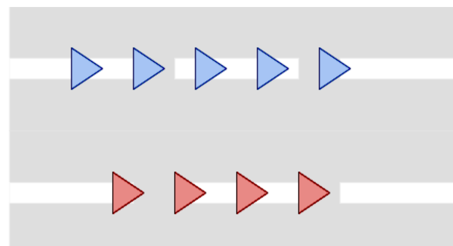

Fig. 2: Visual depiction of convoying in highways

that associating rewards for flocks minimizes the risk of a flock breakdown, thereby hastening the separation process. Our treatment of flocking and schooling is the same and the words would be used interchangeably. Furthermore, we extend the idea of a flock to convoys or platoons of vehicles which choose a lane on a highway. In doing so, the convoy of like robots cruises at its maximum speed. We present this as a comparison of current system speed - calculated as sum of current speed of all vehicles in the system to the maximum system speed - calculated as the sum of maximum speed of all vehicles. Over a finite amount of time vehicle platoons are autonomously seen to reach this maximum system speed.

The rest of the paper is organized as follows. Section II describes the related work in the field of autonomous vehicle convoying and relevant literature from modeling animal group behaviours. Section III presents the decision rule used to emerge with consensus in our robotic applications. Section IV dwells on experiments with rewarding robotic swarms, leading the way to our work on vehicle convoying in section V. Section VI elaborates our results.

II. Related Work

Vehicle platooning in an Automated Highway System (AHS) to increase highway capacity and safety was researched as early as 1994 [5]. Hedrick et al., in this paper describe a platoon as one or more vehicles travelling together as a group with relatively small spacing. Decision and control algorithms for cooperative autonomous manoeuvres have been discussed in [6]. The paper also identifies the need to develop algorithms to enable intelligent driving manoeuvres for vehicles as they go. As described in section I, flocking imparts the intelligence required for such collective steering response.

Most recent work is focused on developing motion control algorithms for collaborative driving [7]. With many deriving from visual perception of obstacles [8] [9] [10], efficient communication technologies with high reliability and low latency have also been supporting the control algorithms [11]. In the background of existing work on motion control algorithms, our work delves into the autonomy of vehicles to join, switch or if permitted create traffic lanes.

Information from personal preferences and social cohesion driving collective decisions in golden shiner fish has been discussed and modeled by Miller et al [12]. We adapt their interpretation of self-information and social cohesion factor s as a sigmoidal decision function [13]. Autonomous platooning is achieved by supplementing this decision rule with another rule to defer from flocking behaviour when expected rewards are not met.

III. Schooling Behaviour

Humans and animals balance their personal information (from past experiences) and social information (from interaction with other individuals) to arrive at a single consensus decision. A group can influence the choice of behaviour of an individual by its sheer size. It is also imperative that the individual's personal preferences are respected. This relation is quantified as a sigmoidal function yielding a score for all choices of groups presented to an agent. Personal preferences are modeled as a_X, a value $0 \leq a_X \leq 1$ representing probability of choosing a group X. Social cohesion is instrumented as an exponent over parameter s to number of agents n_X executing the behaviour X. The score, $P(X)$ is pairwise evaluated against a_m, the agent's preference to move independently of any other group. A decision is made in favour of the alternative if $P(X)$ is greater than a_m. Where $a_{mX} = a_m/a_X$, the score $P(X)$ is encoded as:

$$P(X) = \frac{1}{1 + a_{mX} s^{-n_X}}$$

The equation demonstrates how multiple informational dimensions are integrated within groups to achieve consensus, even though no individual has an explicit preference for the consensus option. Having close preference values for two options also makes it possible for an agent to opt for the less favoured choice when size of such a group is dominant enough. The resulting behaviour is that of an aggregation of agents executing a behaviour X favoured by "early followers", as for the rest, the term n_X would dominate.

Figures 3 and 4 plot the variation of the score function for a behaviour with the number of agents executing that bahaviour. Figure 3 additionally has $a_X = 0.3999$, $a_m = 0.6$ and $s = 1.55$. An agent with these settings will follow the group conduct if it sees more than 2 other agents with that conduct. Figure 4 has $a_X = 0.0001$, $a_m = 0.6$ and $s = 1.55$. The low preference for this behaviour sets the threshold for its election to a high number of agents executing that behaviour.

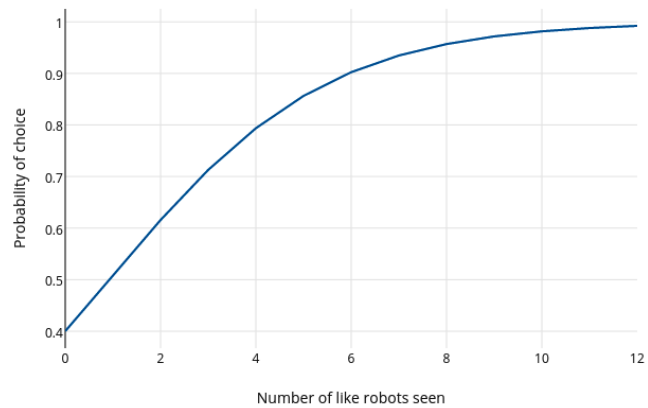

Fig. 3: Probability of electing a choice with $a_X = 03999$ against $a_m = 0.6$ and $s = 1.55$

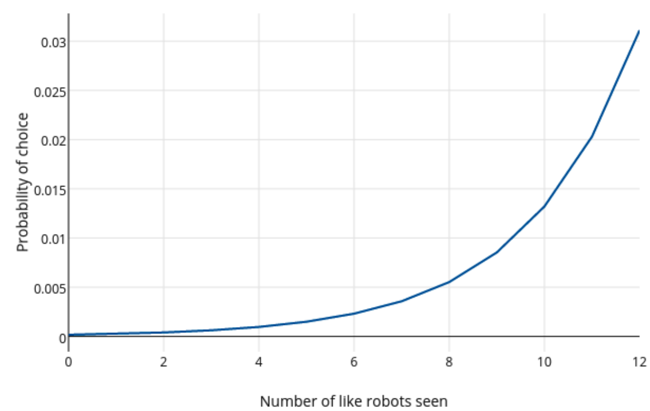

Fig. 4: Probability of electing a choice with $a_X = 0.0001$ against $a_m = 0.6$ and $s = 1.55$

IV. Robot Schooling

The decision rule described in the previous section is applied to 30 robotic agents in an OpenGl based simulator for a toroidal world. Each agent has a 270 degree field of view, their colors being used are flock identifiers. Robots are equally distributed in red and blue colors. Motion control for the agents is an elementary pose copier. When a flock is favoured, the pose of the closest robot from the flock is adapted by the deciding robot. The robot favours its own kind with a preference value of 0.39999. This is done to absolutely avoid flocking with the opposite kind which is modeled with a preference of 0.00001. Preference for independent movement is at $a_m = 0.6$. The social cohesion parameter is $s = 1.55$. Collision avoidance is not embodied.

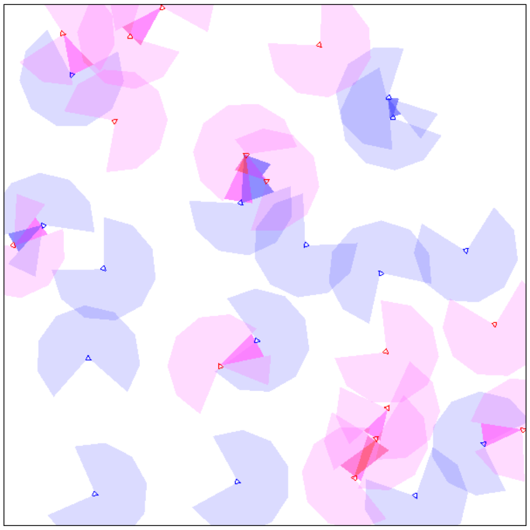

Fig. 5: Dispersed starting postions for robots

A. Naïve Schooling

Experiments run with the above mentioned preset constitute naïve schooling. An agent is found to be very fickle in choosing its group. This is because multiple small groups formed locally, influence peripheral robots from other groups to abandon their initial choice. Through this dynamic leader election process, heuristically all small flocks coalesce to emerge with a single school of like robots. Figure 5 shows a snapshot of the dispersed positions all robots start with.

B. Associating Rewards with Schools

The capricious demeanour of robots seen with naïve schooling is restrictive in the amount of time taken to reach perfect flocking. In terms of information of flock existence, we need inertia in this existence knowledge, so that a robot will not be influenced by another moving in solitude. By associating a boolean reward to a robot in a flock, we observe that separation is achieved at an accelerated pace. Figure 6 is a snapshot showing separate schools for rewarded robots. Rewarded robots are drawn to be encircled with the same colour as the robot. The graph from figure 7 is a plot of the number of robots rewarded as a function of time. As theorized, information about a flock's existence is accelerated through all agents as more and more robots become rewarded. This results in an increased pace of separation of the robots.

V. Autonomous Convoying

Although schooling in the form described in section IV renders benefits to scenarios like foraging, autonomous driving is one avenue that can tremendously benefit from like-vehicle convoying. We conceive a highway of vehicles located in random lanes. A sports car would be restricted in speed by the minivan that is ahead of it in its lane, unless it is able to switch to a faster one. Schooling of vehicles driving at similar speeds and allocating them to a lane delivers high throughput, optimal performance per vehicle, less frequent

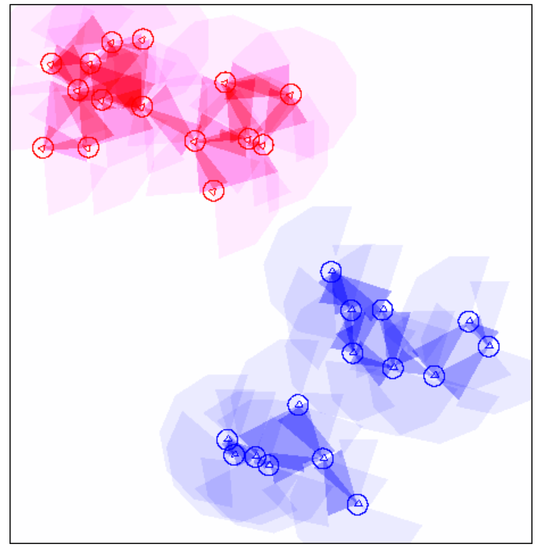

Fig. 6: Schooling of like coloured robots with associated rewards

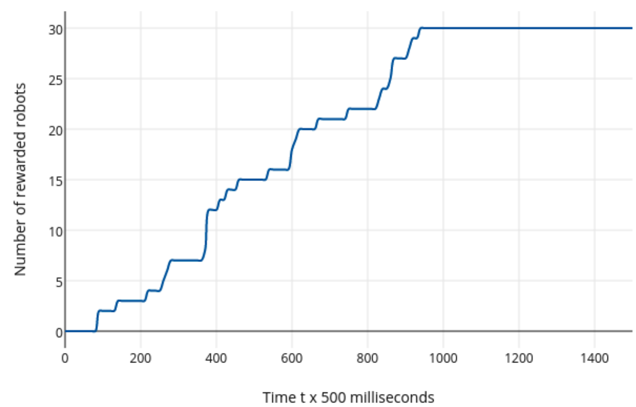

Fig. 7: Rate of communication of rewarded-ness of a robot

lane-changing events and enhanced safety as described in **??**. On a highway of cruising vehicles, adding all individual speeds can be defined as the system speed. Performance of our robot schooling implementation is measured as current system speed against the maximum system over time.

Lane driving is implemented by adapting the robot schooling manoeuvres to form convoys. A more sophisticated controller is added to our robots to follow the closest robot from its elected flock. The goal of this motion controller is to keep the velocity vector of every robot in the resulting convoy in parallel to one-another. Robots dispersed in various lanes trace a smooth curve to each other, forming linear traffic. A boolean reward is activated when a robot conforms to the lane discipline by not deviating from the convoy vector beyond a threshold. The reward is an indicator of convoy adherence and is depicted as a circle around the robot described in section IV-B. Lanes themselves are defined as a relatively closely spaced platoons of robot-vehicles. No physical boundaries are defined for such lanes and collision avoidance is also neglected. The focus of our work is to demonstrate an empirically intelligent convoying behaviour for robot-vehicles. We argue that our work can be extended to respect the extensively researched motion control algorithms.

VI. Results

Experiments are run on robots starting at randomly dispersed positions on two lanes. Decision rule parameters are as described in section IV. Like robots are represented in the same color. Red robots are configured to be slower than the blue. Seven trials are run with different starting positions for a system with 18 robots. The toroidal world restricts the number of robots in the system because of the interference of head of the convoy with its tail. However, more robots can be added to the system by increasing the simulation world size. System speed is encoded as a percentage of the maximum system speed. All robots are assumed to start with their maximum speeds at time $t = 0$, and then adapt their speeds to subsequent formation. Speeds of individual robots are in conformance of the lowest speed of the lane they are in. Robots race at their maximum speed when no other robot is close-by and could potentially cross-paths. Figure 8 is a graph of the percentage system speed at a time t.

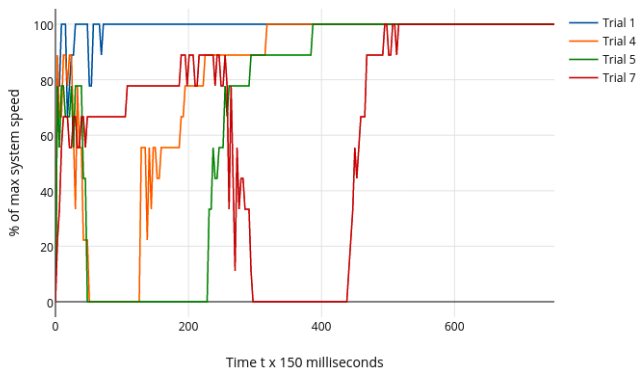

Fig. 8: Current system speed at time t

Trials 1-3 show similar behaviour. For ease of illustration, only trial 1 has been depicted on the graph. System speed oscillates before saturating to the maximum speed because of the lane changing manoeuvre. Robots turn to join the convoy and momentarily have no robots in the vicinity, allowing them to switch to their top speed. Speed however is readjusted once they are back in the convoy.

Trial 4 brings to the fore a latent consequence of our decision rule. If enough number of robots from both colour groups start at the same lane, new convoys for each group will be formed overlapping the other group (fig 9b). This results in a single convoy of robots at the lowest of their speed, or as seen from the graph, at 0% of the maximum system velocity. The robot works around this inadequacy by temporarily turning off the convoying behaviour when it learns that it is not cruising at its best speed. It is intuitively satisfying for the robot to abandon the convoy that is not according the robot its full potential. The robot then elects to switch to a new lane and continues its quest

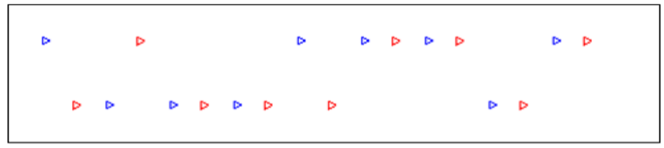

(a)

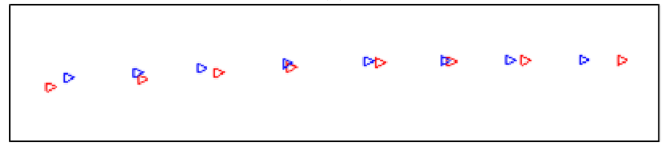

(b)

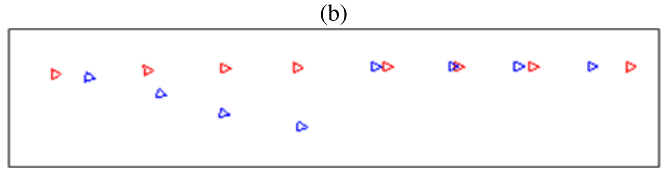

(c)

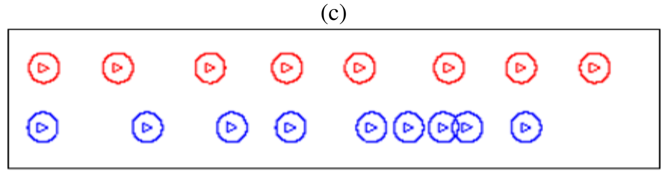

(d)

Fig. 9: System snapshots for experiment trial 4: (a) Initial dispersion (b) Convoying manoeuvre makes both red and green choose top lane (c) Robots leaving the devolving lane (d) System achieves maximum speed

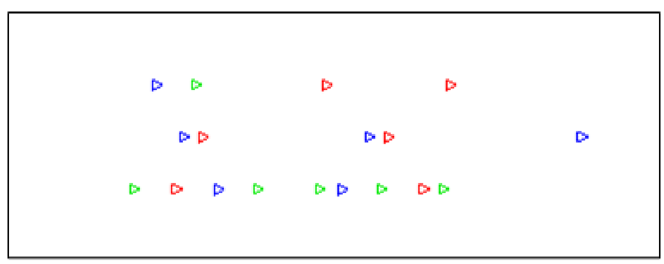

(a)

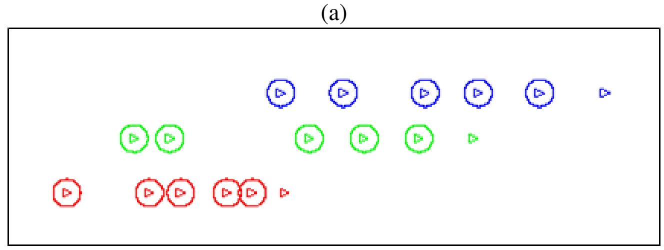

(b)

Fig. 10: System snapshots for three categories of robots: (a) Initial dispersion (b) Three lanes formed; one each for a group of like colours

for a group best suited to its interests. Although we assume that a lane is always available, global information pertaining to lane occupancy is a trivial addition. In applications like moving items, it is indeed supplementary to throughput of the system that a new lane of transporter robots will be formed autonomously.

Trials 5-7 are when the early responders to the local minima in speed (as with trial 4) are towards the tail end of the convoy. Such robots leave the devolving convoy and cruise at their maximum speed alongside it. Other like robots influenced with the flair of the speeding robot will elect to school with it. However when the convoy-behaviour turn-off time is insufficient for the early responder, it will drive the newly formed convoy to rejoin the devolving convoy. This explains the delayed plunge of the percentage system speed seen with trial 7. The system however recovers in the same way as described for trials 4 and 5.

All trials are seen to heuristically approach the maximum system speed and in finite time as hypothesized. Further, the convoying behaviour developed can be seen to hold when extended to three groups of robots. Figure 10 shows that three lanes, one for each group of same coloured robots are formed as a result. Green robots were configured as the slowest, blue- the fastest and red with an intermediate speed.

VII. Conclusion

We introduced multi-agent schooling as an emergent behaviour of consensus decision making. The paper theorizes and demonstrates that controlling motion in a school leads to autonomous platooning of robot-vehicles. Such platoons by electing a lane for themselves elevate their speeds to the maximum, thereby improving a highway's capacity. Scenarios like the single lane convoying are also managed by selectively modifying parameters of the decision rule. By emulating animal behaviour such as humans swarming at a sale or fish schooling to avoid predators, we realize a group consensus by aggregating individual intelligence.

References

[1] A. J. King and G. Cowlishaw, "When to use social information: the advantage of large group size in individual decision making," vol. 3, no. 2, pp. 137–139, 2007.

[2] C. List, "Democracy in animal groups: A political science perspective." *Trends Ecol. Evol.*, vol. 19, pp. 168–169, 2004.

[3] T. D. Seeley and S. C. Buhrman, "Group decision making in swarms of honey bees," *Behavioral Ecology and Sociobiology*, vol. 45, no. 1, pp. 19–31, 1999. [Online]. Available: http://dx.doi.org/10.1007/s002650050536

[4] D. J. Sumpter, J. Krause, R. James, I. D. Couzin, and A. J. Ward, "Consensus decision making by fish," *Current Biology*, vol. 18, no. 22, pp. 1773 – 1777, 2008.

[5] J. K. Hedrick, M. Tomizuka, and P. Varaiya, "Control issues in automated highway systems," *IEEE Control Systems*, vol. 14, no. 6, pp. 21–32, Dec 1994.

[6] J. Baber, J. Kolodko, T. Noel, M. Parent, and L. Vlacic, "Cooperative autonomous driving: intelligent vehicles sharing city roads," *IEEE Robotics Automation Magazine*, vol. 12, no. 1, pp. 44–49, March 2005.

[7] T. Kaneko, Y. Kuriyagawa, and I. Kageyama, "Calculation algorithm for motion control target for the platooning of autonomous heavy vehicles," in *2014 UKACC International Conference on Control (CONTROL)*, July 2014, pp. 696–701.

[8] A. Broggi, M. Bertozzi, A. Fascioli, C. G. L. Bianco, and A. Piazzi, "Visual perception of obstacles and vehicles for platooning," *IEEE Transactions on Intelligent Transportation Systems*, vol. 1, no. 3, pp. 164–176, Sep 2000.

[9] P. Fernandes and U. Nunes, "Platooning with ivc-enabled autonomous vehicles: Strategies to mitigate communication delays, improve safety and traffic flow," *IEEE Transactions on Intelligent Transportation Systems*, vol. 13, no. 1, pp. 91–106, March 2012.

[10] P. Fernandes and U. Nunes, "Platooning with dsrc-based ivc-enabled autonomous vehicles: Adding infrared communications for ivc reliability improvement," in *2012 IEEE Intelligent Vehicles Symposium*, June 2012, pp. 517–522.

[11] M. Y. Abualhoul, O. Shagdar, and F. Nashashibi, "Visible light inter-vehicle communication for platooning of autonomous vehicles," in *2016 IEEE Intelligent Vehicles Symposium (IV)*, June 2016, pp. 508–513.

[12] N. Miller, S. Garnier, A. T. Hartnett, and I. D. Couzin, "Both information and social cohesion determine collective decisions in animal groups," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 110, no. 13, pp. 5263–5268, 2013. [Online]. Available: http://www.jstor.org/stable/42583444

[13] A. Strandburg-Peshkin, C. R. Twomey, N. W. F. Bode, A. B. Kao, Y. Katz, C. C. Ioannou, S. B. Rosenthal, C. J. Torney, H. S. Wu, S. A. Levin, and I. D. Couzin, "Visual sensory networks and effective information transfer in animal groups," *Current Biology*, vol. 23, no. 17, pp. R709–R711, 2016/12/15. [Online]. Available: http://dx.doi.org/10.1016/j.cub.2013.07.059