Supplemental information

Visual sensory networks and effective information transfer in animal groups

Supplemental Data

Table S1: Means and standard deviations of best parameter fits for each model.

Means and standard deviations are calculated over 10 runs, each using 10,000 samples from parameter space. Best fit parameters are shown for the models in Figure 1C (see also Figures S2C and D).

Mean Best Fit Parameter Values

Informed

Model	*a*	*d*	*e*	*f*	*l*	*r_met*	*r_top*	*t_vis*
constant	-4.542							
target distance	-4.545	0.121	-5.228					
target visibility	-5.150				1.034			
global	-4.948			1.993				
metric	-4.898			2.196		103.401		
topological	-4.946			2.572			10.446	
Voronoi	-5.020			2.546				
visual	-5.014			2.445				0.153

Uninformed

Model	*a*	*d*	*e*	*f*	*l*	*r_met*	*r_top*	*t_vis*
constant	-5.227							
target distance	-5.343	3.107	-0.510					
target visibility	-5.865				1.281			
global	-6.553			3.660				
metric	-6.145			3.683		102.994		
topological	-6.185			3.807			13.423	
Voronoi	-6.249			3.804				
visual	-6.815			4.650				0.012

Standard Deviations of Best Fit Parameter Values

Informed

Model	*a*	*d*	*e*	*f*	*l*	*r_met*	*r_top*	*t_vis*
constant	0.004							
target distance	0.027	0.803	0.198					
target visibility	0.034				0.024			
global	0.023			0.082				
metric	0.088			0.200		15.901		
topological	0.064			0.154			2.235	
Voronoi	0.013			0.027				
visual	0.090			0.240				0.048

Uninformed

Model	*a*	*d*	*e*	*f*	*l*	*r_met*	*r_top*	*t_vis*
constant	0.002							
target distance	2.499	3.251	0.834					
target visibility	0.040				0.073			
global	0.038			0.052				
metric	0.189			0.265		7.343		
topological	0.170			0.289			2.175	
Voronoi	0.024			0.042				
visual	0.194			0.366				0.011

Supplemental Figures

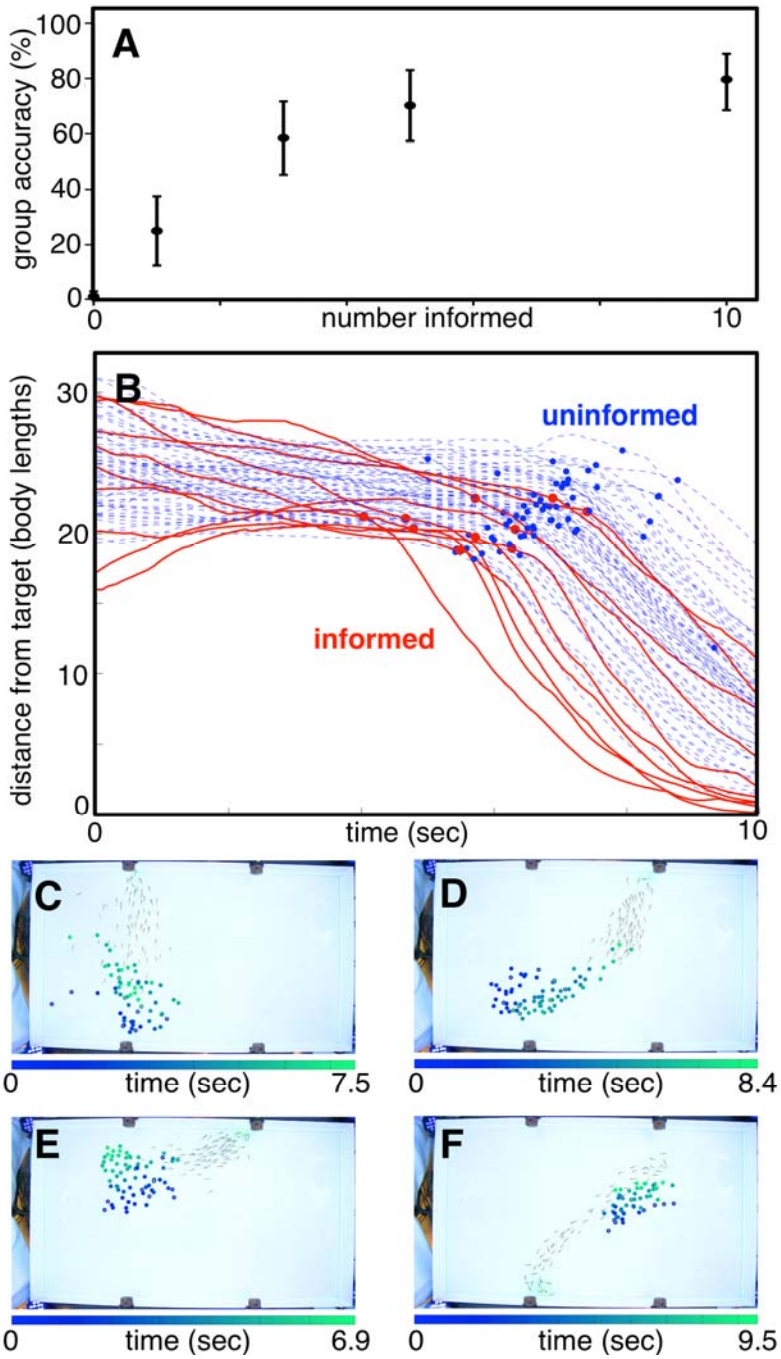

Figure S1: Effectiveness and dynamics of information transfer. (A) Group-level accuracy (percentage of the time that the group reached the target) as a function of the number of informed individuals within the group. Data points represent the aggregate probability over all trials. Error bars represent 95% confidence intervals generated via bootstrapping. (B) Extracting individual responses from trajectory data. Distance form the target over time is shown for each informed (thick red lines) and uninformed (thin blue dotted lines) individual during the same leadership event as is shown in Figure 1B. Individual behavioral responses are marked with dots. (C) The propagation of behavioral responses during four different leadership events. Circles represent responses of individual fish. Locations of circles indicate the location of each fish when it responded, while color represents the time at which the response occurred.

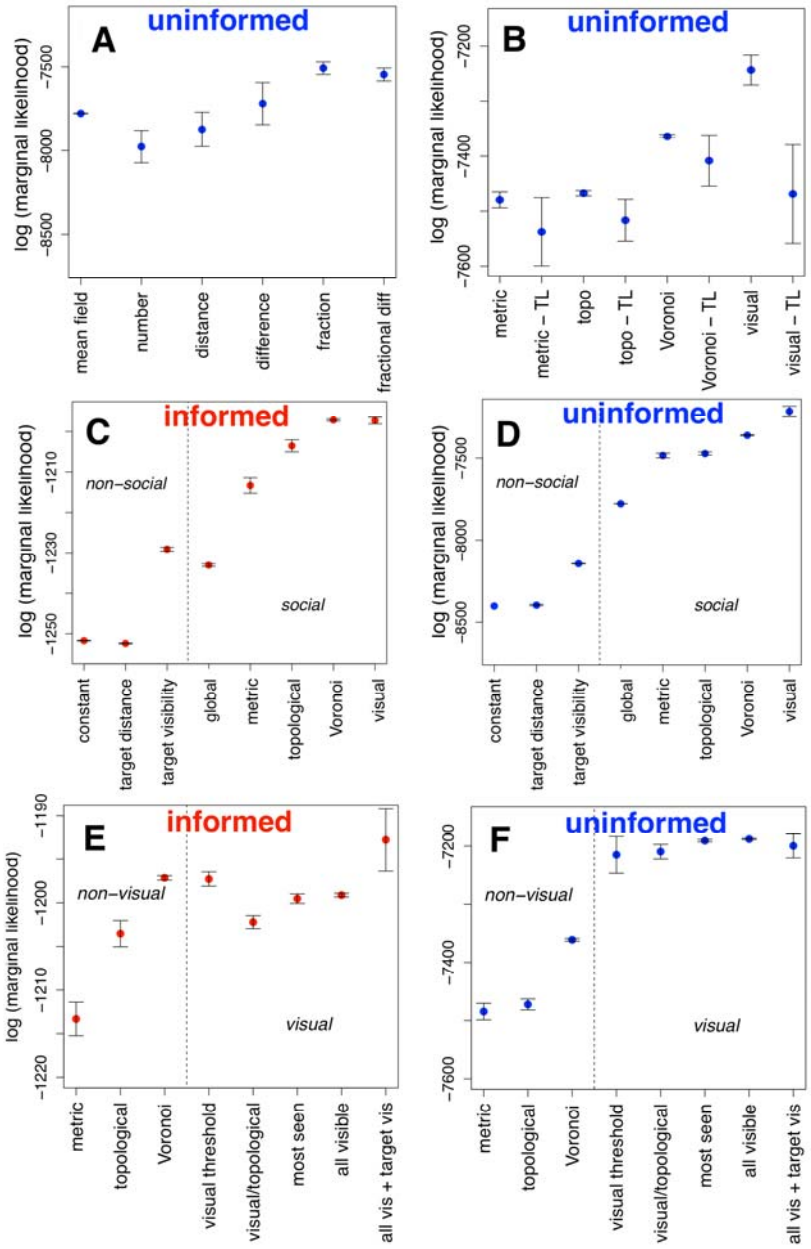

Figure S2: Empirical support for different models of information transfer in fish schools. Log marginal likelihoods are given for each model (higher values indicate more support). Values represent means over 10 runs of numerical integration, each using 5,000 (A) or 10,000 (B-F) samples from parameter space. Error bars indicate standard deviation. (A) shows the support for models assuming different functional forms for social influence, as described in the "Alternative mechanisms for information transfer" section. Here the interaction range is assumed to be metric for all models. A global model (in which a given individual is influenced by all others within the group) is also shown for comparison. (B) shows the support for social models with and without a time lag (denoted "TL"), as described in the "Models with a time lag" section. (C) and (D) show the empirical support for different models for informed and uninformed individuals respectively (same models as in main text). (E) and (F) show the support for alternative visual models, described in the "Alternative visual models" section, as compared to metric, topological, and Voronoi models.

Supplemental Methods

Experimental procedures

Experiments with golden shiners (*Notemigonus crysoleucas*) took place at Princeton University from March 20 to April 20, 2010, with initial training running from March 20 - 27 and experimental trials running from March 28 – April 20. The first three days of experimental trials were excluded from the analysis due to a change of personnel within that time period. Following that point, A.B. Kao conducted all experiments. A total of 100 trained and 470 uninformed fish (both males and females, though their sex was not determined) were used in the experiments. Fish (approximately 5 cm in length at the time of experiments) were purchased from I.F. Anderson Farms (www.andersonminnows.com), and were housed in 27-gallon black plastic tanks for at least two months prior to the experiment. Tanks were continuously filtered and oxygenated, and water temperature was maintained at 16°C. Each tank housed approximately 100 fish. Prior to experiments, they were fed twice daily with crushed TetraMin flakes.

Fish were tagged before experiments began. We used visible implant elastomer tags of different colors to enable us to visually differentiate between informed and uninformed individuals in the videos. These tags were injected under the skin of each individual, on either side of the spine. All fish were allowed at least 2 days to recover from tagging prior to the start of experiments.

The experimental arena consisted of a 2.1 m × 1.2 m white fiberglass tank containing 5 cm of water at the same temperature as the holding tanks. Four automatic feeders were mounted, equally-spaced, on the two longer sides of the arena (as can be seen in Figure S1C-F). A green laser pointer mounted above the arena could be programmed to project the light onto the water near any of the four feeders at a given time. During training trials, the laser was projected, and food (also TetraMin flakes) was delivered at the corresponding feeder. During experimental trials, the laser was projected but no food was delivered. A Sony EX1 video camera (1280 x 720 pixels) mounted above the tank recorded the entire arena at a frame rate of 60 Hz throughout each experimental trial. (The effective frame rate was subsequently reduced to 30 Hz in post-processing to reduce computational runtime for all analyses.)

Training

Fish were trained prior to the beginning of experimental trials in groups of approximately 50. Before each set of training trials, the fish were placed into the arena and allowed to acclimate to the environment for at least 15 minutes. During each training trial, the laser was shone, and food was released from the corresponding feeder. After the fish had consumed all of the food, the laser was switched off, and remained off until the fish resumed their normal schooling activity. The feeder at which the food was presented was randomized between trials. Groups underwent 10 training trials during each training session, with at least 5 minutes between trials. One training session was performed each day for two groups of individuals, resulting in a total of 100 trained individuals. The fish were trained in this manner for 5 days prior to beginning experimental trials. Informed fish were also retrained at the end of each day after experimental trials (in one group of 70 and one group of 30) to ensure that they maintained their association of the light with food. On the two days during which experimental trials did not take place (April 10, and 11) during the three-week-long period of the experiment, informed fish still underwent training trials.

Experimental Trials

Experimental trials mimicked training trials, except that no food was presented at the feeder, and the majority of the group was replaced with uninformed (untrained) individuals. Groups were allowed approximately three minutes to reach the target once the light was turned on, after which the group was declared to have not reached the target and the light was switched off. Each group contained 70 fish, including 0, 1, 3, 5, or 10 informed (trained) individuals (with the rest uninformed). Groups with 1, 3, 5, and 10 informed individuals completed one session of five trials each day, with the order randomized. Groups with 0 informed individuals completed two sessions of five trials, at the beginning and end of the experimental session each day. Adjacent trials were separated by at least 30 seconds. Roughly the same group of uninformed individuals was used throughout the day, with informed and uninformed fish swapped out between sessions to ensure the correct number of informed individuals in each group. Uninformed fish were all drawn each day from the same stock of fish, thus individuals were reused in haphazard order on different days. Informed fish were only used during a maximum of one session on each day. The full dataset represents experimental trials from 19 days (from March 31 – April 20, but skipping April 11 and 12), resulting in a total of 95 trials per experimental condition, and 190 trials for the control (0 informed) condition. We observed with some frequency that the group exhibited an initial startle reaction when the stimulus light turned on. Because we were concerned that startles could bias results, trials in which these startle events occurred were not included in the analysis. The final dataset therefore consists of 107, 48, 53, 47, and 54 trials for experimental groups containing 0, 1, 3, 5, and 10 informed individuals respectively.

All experiments were approved by the Institutional Animal Care and Use Committee at Princeton University.

Determining how many leaders are needed to lead a group

Though our main analysis (Figure 1) focuses on the situation in which ten individuals were informed (trained), we also tested groups with 0, 1, 3, and 5 informed individuals to investigate how the probability of successful leadership scales with the number of "leaders," as described above. Theoretical predictions based on flocking simulations [S1] suggest that leadership accuracy should increase monotonically with the number of leaders, and that a relatively small number of leaders are needed to achieve a high level of accuracy. Our experimental results qualitatively confirm these predictions (Figure S1A). Here, accuracy is defined as the probability that the group was led to the target during the allotted time.

Waves of information transfer: Obtaining individual trajectories and extracting "responses"

We extracted 25 representative leadership events from our videos for the trajectory-based analysis. Leadership events were selected from the experimental condition with 10 informed individuals because of the high repeatability of leadership and strong response of individuals in that condition. Video clips were extracted such that they started before any visible response from the group was observed, and ended when the first fish reached the target. Ending the clips at this point avoided the confounding effects of negative feedback from leading fish once they realized that no food was present in experimental trials, as well as any boundary effects as the group approached the edge of the arena. The selected clips ranged from 10 to 15 seconds in length. Trajectories were obtained using a novel, fully-automated tracking software developed by H. Wu (SchoolTracker). This software extracts the position and orientation of each individual in each

frame, and links these detections together into full trajectories. For our data, the software detected 69-71 fish (i.e. up to one false positive or one false negative) in 98% of frames.

Individual behavioral responses (transitions from typical ambient movement to directed movement towards the target) were extracted from trajectory data using a combination of velocity and acceleration thresholds. When an individual exceeded both of these thresholds, it was considered to have responded. The values for the two thresholds were set manually such that they corresponded to the visually-apparent response points (Figure S1B, Movie S2). Once set, the thresholds were kept the same for all videos. We checked each video manually to ensure that individual responses were accurately captured by this thresholding process. Because typically not all fish had responded before the first fish reached the target, the number of observed responses in each leadership event ranged between 29 and 66.

These responses can be seen to propagate in a wave across the group (Figure S1B, Movie S2). While reminiscent of the "waves of agitation" observed in schools of fish responding to a fright stimulus [S2], waves of information transfer during leadership events are relatively slow (approximately: 7.6 +/- 2.8 individuals/sec; 4.87 +/- 0.08 cm/sec; 0.94 +/- 0.02 body lengths/sec). In addition, leadership events involve maximum speeds similar to those seen in free schooling behavior, and are not nearly as fast as those seen in startle events: The maximum speed reached by a responding fish near (i.e. within a half-second of) its response point was 20.8 ± 5.2 cm/sec (mean ± standard deviation over all individuals), a value comparable to the speeds of up to approximately 20-30 cm/sec seen in freely schooling fish [S3], but far removed from the speeds typically seen in startle events in the same species (77.6 ± 18.2 cm/sec; unpublished data). Furthermore we note that transitions in group behavior, similar to those investigated here, characterize natural schooling behavior in golden shiners [S3].

Reconstructing the visual information available to each individual

The field of view for each individual in the school was estimated using a ray-casting algorithm from the vantage point of each eye of each individual, after determining the position and pose of each individual's body in space, for a given moment in time. Individual head positions were determined using the SchoolTracker novel tracking software (H. Wu), developed to deal with many partially occluded and overlapping bodies in space. Next, body estimation was achieved in two stages. First, the midline of each individual was estimated, allowing for multiply overlapping bodies, using a curve-tracing algorithm similar to [S4]. A uniform, two-knot cubic basis spline was fit to each midline estimation, and integrated with prior estimations using an Unscented Kalman Filter. Taking inspiration from [S5] but simplifying the approach, the flanks of each individual were estimated using a simple linear fit to the displacements of the body-background boundary from the midline b-spline (fits were done using the Theil-Sen estimator for robust linear regression). Flank slope and offset parameter estimates were improved over time using simple constant scalar Kalman filters. Eye positions were approximated at fixed positions near the estimated position and orientation of the head. After body and eye estimations were completed, bodies were encoded in a simple, regular grid spatial data structure along with boundary information, in preparation for efficient ray-casting. Ray-casting was used to determine, for each eye of each individual, what bodies or boundaries are visible to that eye (Figure 1A). Two thousand equally-spaced rays were cast from each eye of each individual, and represented the fish's line of sight in every direction. Each of these rays terminated, ultimately at a surface: the body of another fish, the fish's own body, the target (i.e. stimulus light), or the tank wall. From the number of rays hitting various surfaces, we estimated the angular area of each object subtended on the retina of the focal individual. For all subsequent analysis, the maximum

angular area between the two eyes was used, thus avoiding any complications arising from binocular vision. Visual field and body pose estimation methods were developed by C.R. Twomey.

Statistical Analyses

Spread of responses within a group
We used Linear Mixed Models (LMMs) to determine whether responses spread out from the target or from the first fish to respond (the "source"), and whether informed individuals tended to respond before uninformed individuals (Table S1). The order in which responses occurred was determined for each leadership event, and this was used as the dependent variable. The time when the "source" individual responded was considered to be the beginning of the leadership event, t_0. The fixed effects included in the LMM were: an individual's distance from the target at t_0, the individual's distance from the "source" individual at t_0, and whether the individual was informed or uninformed. The event number and the group from which the tracks were obtained were included as nested random effects. The first individual to respond in each event was removed from the dataset, as this individual would always have a distance from the "source" of 0.

Models were built up in the following stepwise procedure. The significance of main effects was first determined by comparing a model including terms for all possible main effects to a model with the relevant term removed. If there were two or more significant main effects, the significance of interaction effects was then determined by comparing a model including the interaction and main effects to a model with only the main effects. Degrees of freedom were determined by the difference between the numbers of parameters in the two models under comparison. Parameter estimates are reported from the full model (including all possible main effects) or from a model including interactions when they were found to be significant. Results are shown below:

LMM; Response variable: Rank (order of responses)

Effect	*Estimate*	*Std. Error*	χ^2	*d.f.*	*P*
Intercept	24.95	3.63	NA	NA	NA
Informed	**-9.92**	**2.46**	**150.56**	**1**	**<0.0001**
Distance to source individual	**2.89**	**0.19**	**225.56**	**1**	**<0.0001**
Distance to target	-0.28	0.14	2.91	1	0.09
Informed : Distance to source individual	**-0.96**	**0.42**	**5.07**	**1**	**0.02**

Relative positioning of informed individuals within the group
We defined the "frontness score" of an individual fish to be the number of fish that that individual is in front of at any given time. If an imaginary line is drawn through the current position of a focal fish, perpendicular to its orientation vector, any fish that falls above that line is considered to be in front of the focal fish. Note that, by this definition, two fish can both be in front of each other if they are facing one another.

We investigated the frontness scores of informed and uninformed fish at both the beginning (when the first fish responded, t_0) and the end (when the last fish responded, t_f) of leadership events. We computed the average of these frontness scores for all informed individuals (over all trials), and compared this average frontness test statistic to a null model that assumes informed individuals were randomly positioned within the group. To determine statistical significance, we

used a permutation test. This method involved randomly permuting the labels of informed and uninformed fish within each group, then computing the resulting average frontness statistic. By repeating this process 1000 times, we constructed a null distribution of average frontness statistics. We then compared the observed average frontness to this distribution to determine a (two-tailed) P-value.

Permutation tests on average frontness score test statistic. P-values are two-tailed.

Time	*Informed average frontness score*	*Uninformed average frontness score*	*P*
Beginning	**34.48**	**29.97**	**<0.001**
End	**50.97**	**32.01**	**<0.001**

Comparing the support for different models

Following the methodology of [S6], we employed Bayesian model selection to compare the empirical support for different models of information spread. Bayesian model selection is a technique that allows one to directly compare the level of support for different models of varying complexity. In contrast to traditional model fitting, which finds the parameters that maximize the likelihood of the data given a particular model, Bayesian model selection involves integrating the marginal likelihood over a specified prior distribution of parameters. The resulting marginal likelihood describes the overall support for a given model (assuming the prior distribution of parameters), and the marginal likelihoods of different models can then be compared. Because prior distributions are normalized, this method automatically penalizes models with more parameters [S6].

Here, we are interested in modeling individual responses. In each frame (1/30 sec) we classify each individual as "pre-response" (before the response) or "post-response" (after the response). During each frame, an individual in the pre-response state either does or does not transition to the post-response state. Each of our models predicts the probability that such a transition occurs in a given frame based on both environmental and social factors. The models all assume that once an individual has transitioned to the post-response state (i.e. responded), its probability of responding becomes 0. Hence, after individuals have responded, they no longer contribute to the likelihood calculations.

Using these predicted probabilities, we then compute the likelihood of the actual responses given a particular model and parameter set. The likelihood is defined as the probability of the data given the model, $P(D \mid M,\{\alpha_k\})$, where D is the data, M is the model, and $\{\alpha_k\}$ is the set of parameters. If we define X_{itj} as 1 if individual i responds at time step t during leadership event j, and 0 otherwise, then the likelihood can be computed as follows:

$$P(D \mid M,\{\alpha_k\}) = \prod_j \prod_t \prod_i P_{trans}^{itj}(s_{itj} \mid \{\alpha_k\}) X_{itj} + \left(1 - P_{trans}^{itj}(s_{itj} \mid \{\alpha_k\})\right)\left(1 - X_{ijt}\right), \quad \text{(S1)}$$

where $P_{trans}^{ijt}(s_{itj} \mid \{\alpha_k\})$ is the model-predicted probability that individual i responds in time t during leadership event j, given the parameter set, and the products run over all trials, time steps, and individuals.

The marginal likelihood of the model given the data can then be computed by integrating over the entire (possibly multivariate) parameter space, A:

$$P(D \mid M) = \int_A P(D \mid M, \{\alpha_k\}) P(\{\alpha_k\} \mid M) dA, \tag{S2}$$

where $P(\{\alpha_k\} \mid M)$ is the prior distribution of parameters. Following [S6], we randomly sample parameter space according to specified prior parameter distributions (defined below), then perform numerical integration to compute an estimate of the marginal likelihood of a given model. To obtain an estimate of the error in this integration, we perform 10 calculations of the marginal likelihood using 10,000 samples from parameter space in each calculation for each model and then compute the standard deviation of the resulting distribution.

We assume that the probability that an individual responds during a given time step, $P_{trans}(s)$ (where we now drop the subscripts and superscripts for convenience), is given by the logistic expression

$$P_{trans}(s) = \frac{1}{1+e^{-s}} \tag{S3}$$

where *s* may depend on environmental and/or social factors present in the environment of the individual at that particular time step. Our different models represent different ways of determining *s* at a given time step. The model definitions and descriptions are given below.

Models:

Non-social models:

In non-social models, *s* of each focal individual is determined by factors unrelated to the current state of other individuals. The models compared are:

constant – constant response probability over time:

$$s = a \tag{S4}$$

target distance – response probability depends on an individual's distance from the target, δ_{tar}:

$$s(d_{tar}) = a + d\delta_{tar}^{\ e} \tag{S5}$$

target visibility – response probability depends on the angular area of the target on the retina of a focal individual, α_{tar}:

$$s(\alpha_{tar}) = a + lH(\alpha_{tar}) \tag{S6}$$

where $H(x,y)$ is the Heaviside step function, defined as:

$$H(x) = \begin{cases} 0, & x < 0 \\ 1, & x > 0 \end{cases} \tag{S7}$$

Social models:

In all social models, s is determined by the fraction of "neighboring" individuals that have previously responded, n_{resp}, according to:

$$s = a + fn_{resp} \tag{S8}$$

To investigate which "neighbors" actually influence an individual's response probability, we compared models using different neighborhood definitions:

global – all individuals in the group

metric – only individuals within a radius r_{met} of the focal individual

topological – only the r_{top} closest individuals to the focal individual

Voronoi – only individuals that share an edge with the focal individual in a Voronoi tessellation of the group

visual threshold – only those individuals that take up an angular area on the eye of the focal individual that is greater than or equal to a threshold, τ_{vis}.

Prior parameter distributions are given below, where $U_c(x,y)$ and $U_d(x,y)$ indicate continuous and discrete (integer) uniform distributions on the interval [x,y] respectively.

Parameters	Prior distribution
a	$U_c(-10,0)$*unitless
d	$U_c(-10,10)$*units are 1/pixels, where 1 fish body length ~40 px
e	$U_c(-10,10)$*unitless
f	$U_c(-10,10)$*unitless
l	$U_c(-10,10)$*unitless
r_{met}	$U_c(1,519)$*units are pixels
r_{top}	$U_d(1,70)$*units are individuals
τ_{vis}	$U_c(0,0.76)$*units are angular area

To introduce the least amount of bias, the upper limit of r_{met} was chosen to be the average distance between the focal individual and the farthest-away neighbor, averaged over all individuals and all trials. Similarly, the upper limit of τ_{vis} was chosen to be the angular area threshold at which individuals had an average of 1 neighbor. These priors were chosen so as to make models using different neighborhood definitions as comparable as possible. All best fit parameter values (Table S1) fell well within these parameter ranges for all models, giving us confidence that priors were reasonable.

While previous simulations of collective motion have typically adopted metric, topological, or Voronoi assumptions, two notable exceptions are the recent theoretical work of Lemasson et al. [S7,S8] and Kunz & Hemelrijk [S9], both of whom develop models of collective motion based explicitly on the visual information available to individuals. These models differ from the visual models described here in that they use a self-propelled particle framework to describe velocity changes of individuals over time, whereas our analysis employs experimental data to test among alternative models of how a behavioral change is propagated within groups. While the substantial differences make it difficult to compare these models directly, commonalities include the

incorporation of occlusion and a threshold on angular area to determine which neighbors are influential.

Alternative mechanisms for information transfer

For the social models shown in the main text (Figure 1C), we assume that individuals conditioned their probability of responding on the fraction of their neighbors that had already responded. In a preliminary analysis, we also tested several other functional forms for predicting responses. From this comparison, we determined that the fraction of neighbors that have already responded was the most likely model, and we therefore focused on this measure in our main comparison of different neighborhood definitions. Below we list the alternative mechanisms, and their marginal likelihoods are shown in Figure S2A. All models described below use a metric neighborhood definition (with a globally-connected model also shown for comparison).

Model	**Response probability depends on…**	**Functional Form**
Global	Fraction of all individuals that have already responded	$s(n_{resp}) = a + fn_{resp}$ where n_{resp} is the fraction of all individuals that have responded
Number	Number of neighbors that have already responded	$s(N_{resp}) = a + fN_{resp}$ where N_{resp} is the number of neighbors that have already responded
Distance	Distance to all neighbors that have already responded	$s\left(\{d_i\}\right) = a + g\sum_i d_i^{\,h}$ where d_i is the distance to a given individual and the sum runs over all individuals that have already responded
Difference	Difference between number of neighbors that have and have not responded	$s(N_{resp}, N_{no-resp}) = a + f(N_{resp} - N_{no-resp})$ where N_{resp} and $N_{no\text{-}resp}$ are the numbers of individuals that have and have not responded respectively
Fraction	Fraction of neighbors that have already responded	$s(n_{resp}) = a + fn_{resp}$ where n_{resp} is the fraction of neighbors that have responded
Fractional difference	Difference between number of neighbors that have and have not responded divided by total number of neighbors	$s(N_{resp}, N_{no-resp}) = a + f\left(\frac{N_{resp} - N_{no-resp}}{N_{resp} + N_{no-resp}}\right)$ where N_{resp} and $N_{no\text{-}resp}$ are the number of neighbors that have and have not responded respectively

We also tested an additional nonsocial model that is not shown in the main text. This model assumes that individuals' response times are drawn from a Gaussian distribution with a mean μ and standard deviation σ. The Gaussian is truncated at $t = 0$ (so that individuals do not respond before the stimulus light), and renormalized such that the total probability of responding is equal to 1. The log marginal likelihood of this model is -21157.55 ± 9.62. Due to the extremely low likelihood of this model, its marginal likelihood is not plotted in Figure S2A. We speculate that this model performs so poorly because the fish often did not notice the stimulus light immediately, and therefore the absolute response times varied widely, in contrast to the model assumptions.

Models with a time lag

In the social models discussed above, fish are assumed to respond to the characteristics of their instantaneous neighborhood. As an extension, we also considered the possibility that individuals respond to their neighborhood configurations some time in the past. In other words, there may be time lag between when an individual observes its neighborhood and when it responds to this observation. To test this idea, we created models identical to the social models shown in Figure 1C (including metric, topological, Voronoi, and visual neighborhoods), but with the addition of a time lag, *h*. Thus, in these models, individuals conditioned their response probabilities on their neighborhood configurations some time, *h*, in the past.

We found little evidence for the existence of such a time lag (Figure S2B). Each time-lag model performed worse (i.e. had a lower marginal likelihood over all data from the uninformed individuals) than its non-time-lag counterpart. In addition, the mean best-fit values for the time lag (*h*) were very small (0-1 frames), further supporting the approximation that fish respond to their instantaneous local neighborhoods.

Models for informed individuals

In our main analysis, we test models for predicting individual responses of the uninformed (untrained) individuals within the group (Figure 1C). We chose to use only the uninformed individuals because the informed individuals could have shown biased behavior due to their training. In addition, there were six times as many uninformed individuals as informed individuals in each group, therefore providing more data. However, we also performed the same analysis for the informed individuals, using the same models (Figure S2C). The relative performance of different models is very similar for both informed and uninformed individuals (compare Figure S2C and D) suggesting that, despite the difference in information, both make decisions using the same, or similar, mechanism(s). The differences between informed and uninformed individuals instead manifest themselves through different best fit parameter values (Table S1). However, due to the smaller amount of data and the potential confounding factor of prior training, we must be cautious in interpreting the results pertaining to informed individuals.

Alternative visual models

We tested four other vision-based models, in addition to the model shown in Figure 1C. The first is a combination visual / topological model which assumes that interactions are topological, but that only visible individuals are relevant. In other words, it is equivalent to a topological model where individuals that are not visible are not counted. The second assumes that individuals are ranked by the angular area they occupy on a given focal fish's retina, and that the focal individual only pays attention to a certain number of these "most seen" neighbors. This model is similar to a topological model, but it uses visual information rather than physical distance to determine which individuals count as neighbors. The third assumes that individuals interact with all visible neighbors (equivalent to a visual threshold of 0). The final model is an additive combination of this "all visible" model and the "target visibility" model. The results are shown for both informed and uninformed individuals in Figure S2E and F respectively. For uninformed individuals, all of these alternative visual models perform similarly well to the models shown in the main text. For informed individuals, the first three perform slightly worse and the last performs slightly better.

AIC scores

Bayesian model selection requires us to choose a prior distribution of parameters. To ensure that our choice of priors did not influence our results, we compared the relative performance of models based on our full analysis (Figure 1C; Figure S2C,D) to the relative support based on the models' Akaike Information Criterion (AIC) scores, a measure which is independent of the prior

distribution of parameters. As shown below, the relative performance of the various models based on AIC scores is consistent with their performance based on the full Bayesian analysis, suggesting that our results are robust to the prior distribution of parameters chosen.

Models describing informed individuals

Model	*AIC*
Voronoi	2383.08
visual threshold	2385.26
topological	2399.65
metric	2419.45
target visibility	2445.73
global	2455.45
constant	2497.17
target distance	2501.11

Models describing uninformed individuals

Model	*AIC*
visual threshold	14343.95
Voronoi	14709.10
topological	14914.29
metric	15026.86
global	15545.04
target visibility	16273.18
target distance	16790.98
constant	16796.00

Network analysis

To get a broader picture of information flow in the various models, we created hypothetical interaction networks from our data based on each model's assumptions. For metric networks, each individual was taken to be connected to all individuals within a given distance, r_{met}. For topological networks, each was connected to its r_{top} nearest neighbors. For Voronoi networks, each individual was connected to its Voronoi neighbors. For visual networks, each individual was connected to individuals that occupied an angular area on its retina greater than a threshold value, τ_{vis}. Different parameterizations of these networks were generated by varying their respective interaction ranges (r_{met}, r_{top}, and τ_{vis}). For Voronoi networks, only one parameterization exists because the interaction range is not variable.

For each network parameterization, we computed three basic network properties (shown in Figure 1D and E). All measurements represent mean values over all of our data, randomly subsampled to include 10 networks per sequence (and thus 250 networks total). The measured properties include:

Average degree
This measure gives the average number of neighbors for a given individual within the network (averaged over all individuals).

Network efficiency
This measure provides an estimate of how fast information can spread through the network. It is defined by computing the inverse of the shortest (directed) path length for each pair of individuals within the network (i.e. the inverse of the fewest number of links that must be traversed to pass from one individual to another). These values are then averaged (in an unweighted fashion) over the entire network to yield one value.

Network transitivity
This measure represents the extent to which individuals who share neighbors are also neighbors themselves. In other words, transitivity measures the extent to which triplets of individuals are mutually connected. The measure is defined on an undirected graph as the ratio between the number of connected triads (triplets of nodes) to the number of triads in which at least two of the nodes are connected. Here, for simplicity, we use the undirected version of this measure even on our directed graphs (visual and topological) by converting all directed edges to undirected edges. However, using a directed definition of transitivity does not qualitatively change results.

While the transitivity measure described above only addresses redundancy of connections at a local scale, we can generalize this measure to account for longer path lengths. One such generalized definition is: given that there exists some path of length n from a start node to an end node, what is the probability that the start node is also connected to the end node? In other words, for every pair of nodes connected by a path of length n (not necessarily a shortest path), how often are these nodes also directly connected? We call this probability the *n-transitivity* of a graph, and suggest that its value may reflect the correlation of information at a given scale (the scale becomes more global as n increases). From our data, we computed the n-transitivity of metric, topological, Voronoi, and visual networks across a range of average degrees for n up to 25. Our results (graphs not shown) indicate that the difference between visual / Voronoi and metric / topological networks persists even past the local scale, and indeed is still apparent up to $n = 10$, a relatively long path length. As n increases, the curves converge, and eventually coincide (for approximately $n > 15$).

All network properties were calculated using igraph v0.6 [S10] in R v2.15.0.

Author Contributions

Ariana Strandburg-Peshkin - all aspects of data analysis, including Bayesian model selection, network analysis, trajectory analysis, and data processing; wrote the paper
Colin R. Twomey – developed novel visual field software; network analysis; wrote the paper
Nikolai W. F. Bode –Bayesian model selection analysis
Albert B. Kao – conducted experiment
Yael Katz – designed and conducted experiment
Christos C. Ioannou –statistical analysis (non-Bayesian)
Sara B. Rosenthal –application of visual field methods
Colin J. Torney –trajectory analysis (extracting responses)
Haishan Wu –developed novel tracking software
Simon A. Levin – co-advised and directed data analysis
Iain D. Couzin – designed and conducted experiment, advised and directed data analysis; wrote the paper

Supplemental References

S1. 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*, 513-516.

S2. Radakov, D.V. (1973) *Schooling in the Ecology of Fish* (Halsted Press, New York).

S3. Tunstrøm, K., Katz, Y., Ioannou, C.C., Huepe, C., Lutz, M., Couzin, I.D. (2013). Collective states, multistability, and transitional behavior in schooling fish. PLoS Comput Biol *9(2)*, e1002915.

S4. Raghupathy, K., Parks, T.W. (2004). Improved curve tracing in images. International Conference on Acoustics, Speech, and Signal Processing.

S5. Fontaine, E., et al. (2008). Automated visual tracking for studying the ontogeny of zebrafish swimming. J Exp Biol *211(8)*, 1305-1316.

S6. Mann, R.P., et al. (2013). Multi-scale inference of interaction rules in animal groups using Bayesian model selection. PLoS Comput Biol *9(3)*, e1002961.

S7. Lemasson, B.H., Anderson, J.J., Goodwin, R.A. (2009). Collective motion in animal groups from a neurobiological perspective: The adaptive benefits of dynamic sensory load and selective attention. J Theor Biol *261*, 501-510.

S8. Lemasson, B.H., Anderson, J.J., Goodwin, R.A. (2013). Motion-guided attention promotes adaptive communications during social navigation. Proc R Soc B *280*, 20122003.

S9. Kunz, H., Hemelrijk, C.K. (2012). Simulations of the social organization of large schools of fish whose perception is obstructed. Appl Anim Behav Sci *138(3-4)*, 142-151.

S10. Csárdi, G., Nepusz, T. (2006). The igraph software package for complex network research. InterJournal Complex Systems *1695*, http://igraph.sf.net.