

2015

An optimization-based model of collective motion

<https://hdl.handle.net/2144/14070>

"Downloaded from OpenBU. Boston University's institutional repository."

BOSTON UNIVERSITY
GRADUATE SCHOOL OF ARTS AND SCIENCES

Dissertation

**AN OPTIMIZATION-BASED MODEL OF
COLLECTIVE MOTION**

by

DIANE H. THERIAULT

B.A., Boston University, 2004
M.S., Massachusetts Institute of Technology, 2006

Submitted in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

2015

© Copyright by
DIANE H. THERIAULT
2015

Approved by

First Reader

Margrit Betke
Professor of Computer Science

Second Reader

Stan Sclaroff
Professor of Computer Science

Third Reader

Tyson L. Hedrick
Associate Professor of Biology
University of North Carolina at Chapel Hill

Acknowledgments

I would like to thank my committee, Margrit Betke, Stan Sclaroff, Tyson Hedrick, Evi-maria Terzi, and Mark Crovella, as well as other faculty who have provided guidance along the way, including Thomas Kunz, Steve Homer, and Peter Gacs. None of this work would have been possible without my collaborators, Nathan Fuller and Zheng Wu. Many of my lab mates also helped along the way, especially Danna Gurari and Mikhail Breslav. Lab assistants Gordon Towne, Brian Borucki, and Sridevi Suresh contributed important software components that helped bring the technical work to fruition. Numerous other field and lab assistants were instrumental in gathering video data of bats. I also thank Dewayne Davis and David Bamberger for property access, and the Department of Texas Parks and Wildlife for permitting assistance.

AN OPTIMIZATION-BASED MODEL OF COLLECTIVE MOTION

(Order No.)

DIANE H. THERIAULT

Boston University, Graduate School of Arts and Sciences, 2015

Major Professor: Margrit Betke, Professor of Computer Science

ABSTRACT

Computational models of collective motion have yielded many insights about the way that groups of animals or simulated particles may move together and self-organize. Recent literature has compared predictions of models with large datasets of detailed observations of animal behavior, and found that there are important discrepancies, leading researchers to reexamine some of the most widely used assumptions. We introduce FlockOpt, an optimization-based, variable-speed, self-propelled particle model of collective motion that addresses important shortcomings of earlier models. In our model, each particle adjusts its velocity by performing a constrained optimization of a locally-defined objective function, which is computed at each time step over the kinematics of the particle and the relative position of neighboring particles. Our model explains how ordered motion can arise in the absence of an explicitly prescribed alignment term and simulations performed with our model exhibit a wide variety of patterns of motion, including several not possible with popular constant-speed models. Our model predicts that variations in speed and heading of particles are coupled due to costs associated with changes in relative position. We have found that a simi-

lar coupling effect may also be present in the flight of groups of gregarious bats. The Mexican Free-tailed bat (*Tadarida brasiliensis*) is a gregarious bat that forms large maternity colonies, containing hundreds of thousands to millions of individuals, in the southwestern United States in the summer. We have developed a protocol for calibrating cameras used in stereo videography and developed guidelines for data collection. Our field protocol can be deployed in a single afternoon, requiring only short video segments of light, portable calibration objects. These protocols have allowed us to reconstruct the three-dimensional flight trajectories of hundreds of thousands of bats in order to use their flight as a biological study system for our model.

Contents

1	Introduction	1
1.1	A brief history of collective motion	2
1.2	Recent developments in collective motion	5
1.3	Stereo Videography	6
2	Data Collection	8
2.1	Results and discussion	8
2.1.1	Stereoscopic reconstructions of bat field flight	9
2.1.2	Discussion of accuracy of results and ease of experimental setup	10
2.2	Materials and methods	11
2.2.1	Experiment planning	12
2.2.2	Protocol for field experiment	16
2.2.3	Post-experiment camera calibration	18
2.3	Observations of Bats	19
2.4	Occlusion reasoning	20
3	The FlockOpt model of collective motion	22
3.1	Introduction	22
3.2	Contribution	24
3.3	The FlockOpt Model	25
3.4	How ordered motion arises	29
3.5	Patterns of motion from simulations	35
3.5.1	Objective Values	41

3.6	Coupling among spatial variation, heading, and speed	43
3.6.1	Local Heading Alignment and Spatial Variation	43
3.7	Comparison with earlier models	48
3.8	Strategies for Model Selection and Parameter Fitting	51
3.8.1	Related Work	52
3.8.2	Parameter estimation via histogram matching	53
3.8.3	Parameter estimation and model selection via model prediction	56
3.8.4	Generic models	59
3.8.5	Matching simulation parameters with natural data	60
3.8.6	Parameter estimation via objective values	61
3.9	Discussion	62
3.10	Conclusion	64
4	Analysis of behavior of bats	66
4.1	Introduction	66
4.2	Bat kinematics	66
4.2.1	Acceleration decomposition	67
4.2.2	Analysis of bat kinematics	70
4.3	Bat spatial distributions	74
4.3.1	Related Work	75
4.3.2	Coordinate systems for studying bat behavior	76
4.3.3	Computing spatial histograms	80
4.3.4	Spatial patterns of turning and pitching acceleration	84
4.4	Conclusion	85
5	Conclusion	90
	References	94

List of Figures

2.1	3D flight trajectories of 28 Brazilian Free-tailed Bats during a 1-s interval (A), estimated from video captured with three synchronized, high-speed thermal infrared cameras (B). The colored pyramids represent the three cameras while the colored lines each mark the trajectory of an individual bat. The observation distance between cameras and bats was approximately 10 m, chosen so that the nose-to-tail span of a bat in an image was at least 10 pixels. The baseline distance between the outermost cameras was approximately 6 m, chosen so that the expected uncertainty in reconstructed 3D positions at the observation distance due to image quantization and image localization ambiguity was less than 10 cm, the length of a bat. The RMS reconstruction uncertainty for the 1,656 estimated 3D positions shown was 7.8 cm. The trajectories were smoothed with Kalman filtering.	9
-----	--	---

2.2 Reconstruction uncertainty due to quantization effects is shown for six hypothetical camera configurations. The cameras were simulated to have a pixel width of $18 \mu\text{m}$ and field-of-view angle 40.5° and be positioned at a fixed height Z and aimed at a common, equidistant fixation point $F = (0, 0, Z)$. Horizontal cuts of the 3D view frustums of the cameras at height Z and lines at $D_{max}=20$ m are shown from above. Placing the cameras further apart reduces reconstruction uncertainty (A vs. B). If the cameras are placed too far apart (C), however, the view volume is "closed," and there are unobservable regions of space where the cameras will be looking past each other. If the distance between the outermost cameras is held constant, adding additional cameras may not decrease the uncertainty due to image quantization in the common observable region (D vs. E). When the image planes of the cameras are parallel (F), the common view volume is smaller and further away from the cameras than in the other configurations. -Note that these 2D cuts of the 3D view frustrums are at the level and elevation angle of the cameras, cuts at a different level or angle would show slightly greater reconstruction uncertainty but similar trends. . 15

2·3	Reconstruction uncertainty due to quantization and resolution issues. In a video frame obtained for a flight study (A), the automatically detected locations of the animals may not be at their centers (colored dots in B). When estimating reconstruction uncertainty (C, D), we include this effect by corrupting the image projections of simulated world points, generated throughout the whole space, with Gaussian noise where the standard deviation is one-sixth of the calculated apparent size of an animal at that location (circles in B). When estimating the reconstruction uncertainty, including image location ambiguity (D) increases the estimated uncertainty more than three-fold over image quantization alone (C) (note the change in color scale).	16
2·4	LoF entry	17
3·1	The objective function Φ is constructed based on the difference between two translated and horizontally scaled sigmoid functions. Appropriate constants are chosen to adjust the width of the central, flat portion, z_o , (left) and the width of the descending leg, z_a (right).	26

3·2 The objective function \mathcal{O}_i , which reflects the desirability of the anticipated future neighborhood of particle i (\mathbf{v}_i in dark blue), is visualized in a 2D simulation with one (left), two (middle), and six neighbors (right). Function values are based on the desirability $\Phi(\|\mathbf{x}_{i,k}\|)$ of the anticipated future distance $\|\mathbf{x}_{i,k}\|$ (dashed line) to its k th neighbor (\mathbf{v}_k in light blue). In 2D, high values of \mathcal{O}_i form rings surrounding the anticipated positions of the neighbors (left). Regions of intersecting rings are most desirable for particle i to move to (middle). A constrained optimization on \mathcal{O}_i is performed (right) to find the best change $\Delta\mathbf{v}_i^*$ in the velocity of particle i (red) that satisfies additional constraints, such as maintaining a preferred speed (light blue annulus) or limiting its turning angle (pink annular sector). 28

3·3 Representative spatial arrangements of particles, are shown from simulations where z_o was varied between $0.0 \Delta u$ and $2.5 \Delta u$ while other parameters were held constant ($\lambda_s = 0.0\Delta u$ and $z_a = 3.0\Delta u$). We found that the width of the flat part (z_o) of the objective function is the primary factor affecting the spatial arrangement among particles in a simulation, whereas the length of the descending leg (z_a) has less impact. When z_o is small (left, $z_o = 0.0\Delta u$), the lattice is tight and the spacing is very regular. When z_o has an intermediate value (middle, $z_o = 1.0\Delta u$) the spacing between particles is still very regular, but there may be some small perturbations in the lattice. When z_o is large (right, $z_o = 2.5\Delta u$), the lattice is loose with holes and irregular spacing between particles. 37

- 3·4 Different types of motion arise due to the rigidity of the lattice induced by the objective function. When z_o is small (top, $z_o = 0.0\Delta u$), the lattice is very tight and small perturbations in the relative position of neighboring particles can lead to large system-wide changes in motion, usually manifested as rotation. When z_o has an intermediate value (middle, $z_o = 1.0\Delta u$), the lattice spacing remains regular, but there is enough flexibility so that small differences in the heading of neighbor particles diffuse across the system without causing major disruptions, leading to a greater prevalence of purely translational motion. because the rotational motion that occurs with the tighter lattice disrupts the local heading alignment. When z_o is large (bottom, $z_o = 2.5\Delta u$), the particles are loosely held together and groups of particles will contain a heterogeneous mix of headings. 38
- 3·5 A group of particles undergoing simultaneous translational and rotational motion, where the group rolls like a ball, is shown. Red and blue are used to identify the two distinct subgroups that have emerged during this portion of this simulation. This behavior emerges when z_o and λ_s take on small values. The center of rotation is on one side of the group, not in the middle, and there is an outer shell that rotates very quickly around the group of slower-moving particles. Distinct subgroups alternate as the stationary portion and outer portion, and the overall movement of the group is much faster than the preferred speed, s^* , due to a sling-shot effect. 39

3·6 Examples of the diverse patterns of motion found in 2D and 3D simulations performed with our FlockOpt model resulted from five different settings of z_o , z_a , and λ_s . The dots indicate the current position of the particles, and the tails show the previous 10 time steps of motion. Formations A, B, and E are not possible if particles are modeled to maintain constant speed. A stronger speed penalty λ_s is associated with a greater prevalence of pure translational motion (C, D). When width z_o is larger (D), the spacing between particles becomes looser. With a large z_o and a low λ_s (E), the headings of the particles become also somewhat heterogeneous, even as the entire group slowly translates together. Stationary rotational mills (B) are more common when z_o is larger and z_a smaller, and are most often seen as a transitional rather than stable configuration. 40

3·7 The average local polarization computed for each frame for one set of initial conditions from five different parameter settings are shown. In our simulations, the objective values typically saturated within 20 time steps, but it took longer for local heading alignment to become established. Even in simulations exhibiting periodic decreases in local polarization due to rotational motion (A and B), the objective values remained nearly saturated throughout the simulations. 42

- 3.8 Joint distributions of local polarization and spatial variation are shown for five simulation parameter sets (top row) and five video segments of bat flight (bottom row). In both simulated data and natural data, spatial variation decreases as polarization increases. When the speeds of particles in simulation is more variable (A and E), the distribution spatial variance conditioned on the polarization is wider. In recordings of bats taken later in the evening (H and J), there is larger spatial variation. 46
- 3.9 The distributions of spatial variation, as observed (ψ , light gray) and after speed and heading had been forcibly decoupled ($\tilde{\psi}$, dark gray), are shown for five FlockOpt simulation parameter sets leading to coherent translation (top), and for five video segments of the flight of *T. brasiliensis* (bottom). In all cases, the distributions $P(\psi)$ and $P(\tilde{\psi})$ were not similar, according to Kolmogorov Smirnov tests ($p \ll 0.01$), and the average spatial variation was larger after speed had been decoupled from heading. In some simulations (A and E), the increase in spatial variation was very large. The increases $\Delta\psi$ observed when decoupling the speed and heading of the bats (3.5%–4.6%) were slightly smaller than the increases observed in simulations with similar average local polarization and variability in speed (3.8%–8.6%). 47

3.10 We carried out a simulation study to examine the effect of the prediction step in our algorithm. For two parameter sets, with $z_o = 1.0\Delta u$ and $z_a = 3.0\Delta u$, and $\lambda_s = 0.0\Delta u$, (top) and $\lambda_s = 0.1$ (bottom) we show the evolution of the average objective values (left) and the local polarization (right) experienced by all particles in simulations carried out with one set of initial conditions. We varied the effective step size by reducing the maximum speed s_{\max} (all panels) and preferred speed s^* (C and D). Under the fully predictive FlockOpt model (red lines) the objective values quickly converge in both circumstances. In simulations without prediction, the objective values are highest in simulations with the smallest step sizes (blue lines). The objective values from simulations without a speed penalty (A) do not converge when the step size is large (light gray lines); when the step size is small, the objective values converge but it takes many more steps (dark gray and blue lines). Even when the objective values converge, the emergence of locally aligned motion takes longer (B and D). 50

3.11 We investigated the impact of computing the responses to each neighbor individually and averaging the results, as in many other models of collective motion. For three different parameter sets known to lead to three different types of motion patterns under a pure FlockOpt model, we show, for the pure model (solid lines) and hybrid model (dashed lines), the average objective values (top) and the average local polarization, ξ (bottom) experienced by all particles in one simulation with one set of initial conditions. In the hybrid model (dashed lines), the objective values converged more slowly to lower terminal values (top row) than in a pure FlockOpt model (solid lines). Despite the lower objective values, interesting patterns of ordered motion emerged that were different from the patterns of motion of the pure FlockOpt model under identical parameters and initial conditions (bottom row). . . . 51

3.12 For simulations performed with two sets of parameters, we show the goodness of fit between a subset of other simulations based on the Earth-mover's distance between histograms of distances between nearest neighbors (columns A, B) speed (C) and local polarization (D). When the speed penalty matches, minimizing the Earth-mover's distance between histograms of nearest neighbor distances will lead to the correct values for z_o and z_a (column A). Even if the assumed speed penalty is incorrect (column B), it is generally possible to correctly match the value of z_o , though, it may be necessary to make reasonable assumptions about z_a (B, bottom row) in order to correctly infer z_o . If the neighborhood parameters are incorrect, then matching histograms of speed will not lead to the correct values (column C). Local polarization histograms are not a discriminative feature for inferring model parameters, since many simulations exhibiting coherent translation lead to very similar histograms, so the distance between histograms from parameters with different parameters is small (D). 55

3.13 The Earth-mover’s distance between the distributions of nearest neighbor distances (left) and speed (right) derived from simulations can be used to infer the parameters of the speed penalty λ_s parameter. In the top row, the true parameters are $z_o = 0.5\Delta u$, $z_a = 3.0\Delta u$ and $\lambda_s = 0.1$. In the bottom row, the true parameters are $z_o = 2.0\Delta u$, $z_a = 3.0\Delta u$ and $\lambda_s = 0.05$. Using the true values of z_o and z_a , the Earth-mover’s distance between the distributions of nearest neighbor distances (left) and speed (right) are plotted for other simulations with the correct neighborhood parameters, but different speed penalty values. If the neighborhood preference parameters are correct, then the nearest neighbor distance histogram is informative with respect to the correct speed penalty (left column), though there may be multiple local minima (bottom, left). If the neighborhood parameters are correct, then inferring the correct speed penalty by matching histograms of speeds is straightforward (right column). 56

4.1 The joint distributions between acceleration magnitude, $\|\mathbf{a}_i\|$, and turning angle, θ_i (B, middle) and speed change, Δs_i (B, bottom) are shown from one of our eleven datasets, taken in the late evening. The correlation between $\|\mathbf{a}_i\|$ and θ_i is 0.94 ($p \ll 0.01$), and the correlation between $\|\mathbf{a}_i\|$ and speed change, Δs_i is 0.075 ($p \ll 0.01$), calculated using a total of 1,263,296 data points that had appropriate support that enabled us to compute acceleration based on smoothed trajectories. For reference, the marginal distributions of acceleration magnitude, $\|\mathbf{a}_i\|$ (B, top), turning angle, θ_i (A, middle), and speed change, Δs_i (A, bottom) are shown. This finding indicates that changes in heading are a more prominent aspect of bat motion than changes in speed. . . . 71

- 4.2 The joint distributions between acceleration magnitude, $\|\mathbf{a}_i\|$, and turning acceleration, $a_i^{(\perp)}$ (B, middle) and pitching acceleration, $a_i^{(\mp)}$ (B, bottom) are shown from one of our eleven datasets (the same used in Figure 4.1, consisting of 1,263,296 data points). The shape of the joint distributions is like a fan because the magnitude of the acceleration vector constrains the maximum value of the corresponding turning or pitching acceleration. We found that the average value of the turning acceleration, $a_i^{(\perp)}$, which reflects a preference for turning right vs left, varied between our two study sites, which may be due to the geometry of the cave entry. The sign of the average value of the pitching acceleration varied on different nights, even at the same study site. For reference, the marginal distributions of acceleration magnitude, $\|\mathbf{a}_i\|$ (B, top), turning acceleration, $a_i^{(\perp)}$ (A, middle), and pitching acceleration, $a_i^{(\mp)}$ (A, bottom), are shown. 73
- 4.3 The joint distribution between turning acceleration, $a_i^{(\perp)}$, and pitching acceleration, $a_i^{(\mp)}$ is shown (B, bottom) for the same dataset used in Figure 4.1. For reference, the marginal distributions of $a_i^{(\perp)}$ (B, top), and $a_i^{(\mp)}$ (A, bottom) are shown. The “lima-bean” shape observed in this distribution is typical of the distributions also obtained from all the other datasets. By examining this distribution, it seems that pitching up is less likely to be associated with turns to the left and right than pitching down. 74

4.4 Joint distributions between speed and turning angle, θ_i (A), turning acceleration, $a_i^{(\perp)}$ (B), and pitching acceleration, $a_i^{(\uparrow)}$ (C), are shown. These distributions are from the same dataset used in Figure 4.1, but the patterns are typical of all of our data sets. We observe a negative relationship between speed and changes in heading, especially large changes in heading (A). However, we observe a positive relationship between speed and pitching up (B), so other types of changes in heading must account for the negative relationship observed in (A). We observe a “lima-bean” shape in the distribution between speed and turning acceleration (C), indicating that large turns to the left or right may be associated with slower flight. 75

4.5 The local coordinate system for a bat is shown. The X axis is along the heading of the bat, the Z axis points up, and in order to make a right-handed coordinate system, the Y axis points over the left wing. 81

4.6 The spatial histogram, \mathbf{W} (Equation 4.15), of relative position of first nearest neighbors is shown from two of our datasets, taken on the same night at different times, before and after sunset, which occurred at 8:42 PM. In the top row, cartoon diagrams show the pose of the bat, relative to the position of the viewer, and in the bottom two rows, slices through the spatial histograms, based on those viewpoints are shown. Note that the scale of the diagrams is different from the scale of the slices through the spatial histogram. The first column depicts the overhead view, cutting perpendicular to the Z axis. The center column shows a side view, cutting perpendicular to the Y axis (over the left wing). The right column shows a view from the front, cutting perpendicular to the X axis (parallel to the heading). The gray regions depict volumes where there is no data. The lack of data in the center of the spatial histogram is due to occlusion. The direction of this region is due to the location of the cameras, relative to the location of the bats. The bright halo effect occurs because most of the first nearest neighbors of the bats are within 1 meter of each other in the dense column depicted in this video segment. If more neighbors were included in these histograms, the halo would be less pronounced. Earlier in the evening, the bats fly closer together.

4.7 A spatial histogram, \mathbf{W} (Equation 4.16), depicting the motion of bats as a function of the relative position of their neighbors is shown, where the weights were given by the turning acceleration, $a_i^{(\perp)}$ (top), and pitching acceleration, $a_i^{(\parallel)}$ (bottom). In the top row, a dark color represents turning right, and a light color represents turning left. In the bottom row, a dark color represents pitching down and a light color represents pitching up. Diagrams depicting the pose of the bat relative to the viewer are shown in the first column. We observe that bats tend to turn away from things that are too close, and towards things that are too far away. For example, when neighboring bats are within 0.5 meters, and directly to the left of the focal bat, the focal bat will tend to turn to the right. However, if the first nearest neighbor is more than 1.5 meters away on the left side, the focal bat will tend to turn to the left. Note that the diagrams and spatial histograms have different scales. 86

List of Abbreviations

DLT	Direct Linear Transformation
DOG	Difference Of Gaussian (distributions)
KS	Kolmogov-Smirnov (statistical test)
PDF	Probability Distribution Function
RMSE	(Square) Root Mean Squared Error
\mathbb{R}^2	the Real (2D) plane
\mathbb{R}^3	the Real (3D) volume

Chapter 1

Introduction

Aggregations of animals, such as flocks of birds and schools of fish, have been a subject of study since the 19th century (Pryer, 1884; Breder, 1954; Hamilton, 1971) and continue to be a very active area of research (Buhl et al., 2006; Krause et al., 2009; Nagy et al., 2010; Bode et al., 2011; Bialek et al., 2012; Attanasi et al., 2014; Gerum et al., 2013; Mann et al., 2013; Calovi et al., 2014; Portugal et al., 2014). During this time, scientists have worked to understand the underlying principles of collective motion (Vicsek et al., 1995; Bialek et al., 2012; Gerum et al., 2013; Attanasi et al., 2014) and the evolution of collective behavior (Katz et al., 2011; Lukeman et al., 2010; Mann et al., 2013) by developing computational models and analyzing the observed behavior of animals and other organisms.

The goal of this work is to present a computational model of collective motion that will govern the behavior of individuals so that the motion of groups of simulated entities will be similar to the motion of groups of animals. Analysis of our model will improve our understanding of the mechanisms underlying collective motion and provide scaffolding for interpreting the behavior of animals. Computational models may be judged by richness, parsimony, and fidelity. They should be capable of reflecting a wide variety of behaviors with only a small number of parameters. The assumptions and predictions of the models should be consistent with observations of natural systems. As a study system, we have chosen *Tadarida brasiliensis*, a gregarious bat that forms large summer maternity colonies containing hundreds of thousands to millions

of individuals (Betke et al., 2008) in the southwestern United States in the summer.

1.1 A brief history of collective motion

Detailed study of schools of fish and flocks of birds in the 1950's through early 1980's (Breder, 1954; Pitcher, 1973; Radakov, 1973; Major and Dill, 1978; Partridge et al., 1980; Partridge and Pitcher, 1980; Partridge, 1981) led to simulation modeling paradigms in the late 1980's and 1990's (Aoki, 1982; Reynolds, 1987; Huth and Wissel, 1994; Vicsek et al., 1995; Czirók et al., 1999; Couzin et al., 2002) that are studied in many domains, including computer graphics (Reynolds, 1987), social networks (Krause et al., 2009), optimization (Eberhart and S., 2001; Karaboga and Akay, 2009), statistical mechanics (Bialek et al., 2012; Nagy et al., 2010; Vicsek and Zafeiris, 2012), ecology and computational biology (Okubo, 1986; Parrish and Edelstein-Keshet, 1999; Eriksson et al., 2010; Gautrais et al., 2012), vehicle and pedestrian simulation (Helbing et al., 2001, 2005), and autonomous vehicle control (Cao et al., 2008; Gazi and Passino, 2004; Jadbabaie et al., 2003; Reif and Wang, 1999). The details of these models are very diverse, but they have in common that they govern the behavior of individuals in a way so that the motion groups of individuals is similar to that observed in natural systems.

In some of the earliest work, Breder (Breder, 1954) proposed that animal behavior could be explained in terms of only attraction and repulsion. Partridge et al. (1980) and Partridge and Pitcher (Partridge and Pitcher, 1980), carried out studies in groups of schooling fish and concluded that fish also sensed and responded to the motion of conspecifics, by determining that fish schooled differently when their lateral line, responsible for sensing motion, was cut, and that fish were still able to school when blinded.

Through simulation studies, the scientific community has discovered that large

groups of entities can move coherently without leaders when group members exhibit simple behaviors in response to the relative position and heading of other individuals in the group (Aoki, 1982; Reynolds, 1987; Vicsek et al., 1995; Czirók et al., 1999; Couzin et al., 2002). Researchers have used these models to understand universal principles of collective motion by simulating groups of entities and analyzing how properties of the behavior of the system, such as distances between neighbors, spatial coherence, coherence of motion (polarization), are affected by the inputs to the simulation, such as the density of particles, the strengths and types of interactions, and the amount of noise (Czirók et al., 1999; Couzin et al., 2002; Strömbom, 2011; Calovi et al., 2014). Researchers have studied these models by simulating groups of entities and analyzing how parameters of the simulation, such as the number and density of the particles, the amount of noise, and the strengths of different interactions, affect properties of the group behavior, such as distances between neighbors, spatial coherence, coherence of motion (polarization), and the formations adopted by the groups (Czirók et al., 1999; Couzin et al., 2002; Strömbom, 2011; Calovi et al., 2014). In computational biology and behavioral ecology, researchers have worked to understand the evolution of collective behavior by using models of collective motion to discover the individual behaviors that lead to the observed collective behavior (Katz et al., 2011; Buhl et al., 2006; Lukeman et al., 2010; Herbert-Read et al., 2011).

An influential simulation model (Aoki, 1982) incorporated repulsion, attraction, and heading alignment influences applied to neighbors within different concentric neighborhood zones. Entities moved away from neighbors that were too close (repulsion zone), towards neighbors that were too far away (attraction zone) and matched headings otherwise (alignment zone) by averaging the headings of neighboring entities. This model was adopted and expanded in the computer graphics community (Reynolds, 1987), and in the theoretical biology community (Huth and Wissel, 1992).

Couzin et al. (2002) carried out a detailed study of the behavior of a similar model under different parameter settings. In the Vicsek model, popular in the physics community, heading alignment is the only influence (Vicsek et al., 1995; Czirók et al., 1999).

In Aoki’s model (Aoki, 1982), speed and direction are treated independently; direction is computed based on entities in the local neighborhood and speed is drawn from a gamma distribution based on the previous speed. In Reynold’s boids model (Reynolds, 1987), the speed of entities is clamped to some desired range. In many models, it is assumed that entities maintain a constant speed (Vicsek and Zafeiris, 2012). Mishra et al. (2012) proposed a model where the speed of motion was set as a function of some derived property of the local neighborhood.

To encourage system stability, the motion of entities is often damped in some way, often by performing a weighted average of the current and desired velocity, or by simply applying a simple threshold on the maximum turning angle (Couzin et al., 2002; Vicsek et al., 1995). The local neighborhood is frequently defined in terms of a metric distance (Aoki, 1982; Reynolds, 1987; Vicsek et al., 1995), and recently, other models of proximity have been developed, such as a topological distance (i.e. “the first K nearest neighbors”) (Ballerini et al., 2008), or the first shell of a Voronoi diagram (Gautrais et al., 2012).

In control theory (Reif and Wang, 1999; Gazi and Passino, 2004), researchers are concerned with developing control laws to govern the behavior of robots undertaking tasks such as social foraging, which require robots to disperse into the environment while maintaining some desirable configuration and avoiding collisions. To accomplish this, influences of attraction and repulsion are understood in terms of the gradients of potential functions, inspired by physical forces. The total potential experienced by an entity is the sum of the potentials induced by each neighbor. Performing gradient

descent of the total potential function leads to a weighted averaging of the attraction and repulsion influences due to each neighbor, and action is taken subject to some step size that may either be constant or based on the magnitude of the gradient. These models often omit a heading alignment influence and work in this area often focuses on proving the convergence of the system to some desired state.

1.2 Recent developments in collective motion

In recent years, as it has become computationally feasible to collect larger datasets of observations of animals, researchers have begun to make more detailed comparisons between the predictions of behavioral models and observations of natural behavior, including locusts (Buhl et al., 2006), starlings (Ballerini et al., 2008; Bialek et al., 2012; Cavagna et al., 2010; Attanasi et al., 2014), surf scoters (ducks) (Lukeman et al., 2010), mosquitofish (Herbert-Read et al., 2011), golden shiners (fish) (Katz et al., 2011), barred flagtails (fish) (Gautrais et al., 2012), and glass prawns (Mann et al., 2013). Based on these studies, some of the fundamental, widely used design assumptions of models of collective motion have been called into question (Calovi et al., 2014).

To combine the influences of multiple neighbors, typically the responses to each neighbor are averaged (possibly with weights) (Aoki, 1982; Vicsek et al., 1995; Couzin et al., 2002), but Katz et al. (2011) performed a detailed study of the behavior of fish in groups of two fish and three fish and they found that the response of fish in groups of three was not an average of the responses of fish in groups of two; there is some unexplained residual effect. Another important issue is speed regulation. Although there are a very small number of variable speed models in the literature (Aoki, 1982; Reynolds, 1987; Mishra et al., 2012; Reif and Wang, 1999), the vast majority of models assume that entities move with a constant speed (Vicsek and Zafeiris, 2012),

and this assumption has been used in more recent literature as well (Gautrais et al., 2012). Challenging this long-held assumption, Katz et al. (2011) and Herbert-Read et al. (2011) studied the behavior of fish in detail and found that speed regulation is actually a very important aspect of fish behavior.

1.3 Stereo Videography

Researchers have studied many diverse topics, including animal behavior and plant seed dispersal, by recording and quantifying the movement of animals and plants using video (Hayashi et al., 2009; Olaniran et al., 2013; Bhandiwad et al., 2013). For example, in 2012, seventy papers published in the *Journal of Experimental Biology* reported using video to measure kinematics, representing 11% of the papers published in the journal that year. Many of these studies used two or more cameras to measure the 3D locations of points of interest in the scene. In order to use image observations from multiple cameras to reconstruct three-dimensional (3D) world positions via triangulation, the relative position and orientation of the cameras (extrinsic parameters) and their focal lengths and principal points (intrinsic parameters) must be given. The process by which these parameters are estimated is known as “camera calibration” and typically involves matching points on a calibration object across camera views.

In the recent biology literature, the most commonly mentioned method for calibrating cameras is Direct Linear Transformation (DLT) (Abdel-Aziz and Karara, 1971). When using DLT, it is important to obtain calibration points throughout the volume of interest; otherwise reconstruction accuracy may be reduced (Hedrick, 2008). Previous authors who have used DLT in a field setting have constructed a large physical calibration object at the field site (Clark, 2009; Munk, 2011), limiting the size of the calibration volume. Others have carefully measured the extrinsic parameters by hand (Cavagna et al., 2008), relying on a semi-permanent placement of

their cameras in a sheltered location. We propose a different calibration approach that is particularly useful for field settings where the volume of interest may be tens of thousands of cubic meters or where cameras cannot be left in place for multiple sessions. Our approach uses the sparse bundle adjustment (SBA) calibration algorithm (Lourakis and Argyros, 2009) which minimizes the difference between the observed and ideal locations of the calibration points in each camera view. Bundle adjustment has been chosen by biologists for its 3D reconstruction accuracy (Walker et al., 2009) but lacks wide use because of the absence of well-documented and freely available software implementations that are directly applicable to analysis of field data.

Chapter 2

Data Collection

Stereo videography is a powerful technique for quantifying the kinematics and behavior of animals, but it can be challenging to use in an outdoor field setting. We here present a workflow and associated software for performing calibration of cameras placed in a field setting and estimating the accuracy of the resulting stereoscopic reconstructions. We demonstrate the workflow through example stereoscopic reconstructions of bat and bird flight. We provide software tools for planning experiments and processing the resulting calibrations that other researchers may use to calibrate their own cameras. Our field protocol can be deployed in a single afternoon, requiring only short video clips of light, portable calibration objects.

The protocols used in this study are consistent with the American Society of Mammalogists (Sikes and Gannon, 2011), and were approved by Boston University’s Animal Care and Use Committee and Texas Parks and Wildlife Department Permit Number: SPR-0610-100.

2.1 Results and discussion

When developing our stereo videography workflow and software, we focused on accuracy and ease. It is important to estimate the level of calibration and reconstruction accuracy when the goal of 3D videography is to quantify the kinematics of airborne animals and facilitate the study of their behavior. Any uncertainty in the estimation of their 3D position affects the uncertainty in derived calculations like velocity and

acceleration which may be of direct biological interest. We here first describe the results of our experiments that yielded stereoscopic reconstructions of bat and bird flight and then discuss their accuracy and the ease in obtaining them.

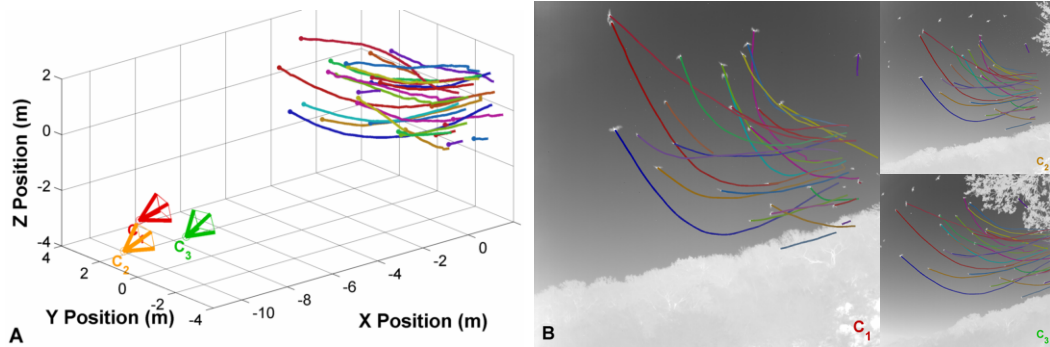


Figure 2-1: 3D flight trajectories of 28 Brazilian Free-tailed Bats during a 1-s interval (A), estimated from video captured with three synchronized, high-speed thermal infrared cameras (B). The colored pyramids represent the three cameras while the colored lines each mark the trajectory of an individual bat. The observation distance between cameras and bats was approximately 10 m, chosen so that the nose-to-tail span of a bat in an image was at least 10 pixels. The baseline distance between the outermost cameras was approximately 6 m, chosen so that the expected uncertainty in reconstructed 3D positions at the observation distance due to image quantization and image localization ambiguity was less than 10 cm, the length of a bat. The RMS reconstruction uncertainty for the 1,656 estimated 3D positions shown was 7.8 cm. The trajectories were smoothed with Kalman filtering.

2.1.1 Stereoscopic reconstructions of bat field flight

We recorded video of Brazilian Free-tailed Bats during their evening emergence from Davis Blowout Cave in Blanco County, Texas (Fig. 2-1) and used the proposed stereo videography workflow to support the estimation of the flight paths of 28 bats flying through a $1,400\text{-m}^3$ volume during a 1-s interval with a 7.8-cm root mean squared (RMS) uncertainty in their 3D positions. We used easyWand and easySBA to estimate the extrinsic and intrinsic camera parameters, using, from each view, 226 points manually digitized from the ends of a 1.56-m calibration wand, 2,010

points manually digitized from hot packs thrown in the air, and 7,135 points on the bats flying through the volume of interest, identified using automated methods (Wu et al., 2009). Manufacturer-provided values served as initial estimates of the intrinsic parameters. The calibrated space was aligned to gravity by calculating the acceleration of hot packs thrown in the volume of interest. The standard deviation of the estimated wand length divided by its mean was 0.046 (4.6%). The respective reprojection errors were 0.63, 0.74, and 0.59 pixels for the three views. The protocols used in bat observation are consistent with the American Society of Mammalogists (Sikes and Gannon, 2011), and were approved by the Institutional Animal Care and Use Committee (IACUC) of Boston University and the Texas Parks and Wildlife Department (Permit Number: SPR-0610-100).

2.1.2 Discussion of accuracy of results and ease of experimental setup

As our results demonstrate, the proposed field videography workflow and software yield accurate calibration of multi-camera systems and enable accurate reconstruction of 3D flight paths of bats and birds in field settings. The values of our two measures of calibration inaccuracy were sufficiently small to indicate accurate calibrations without any errors in determining corresponding calibration points.

We balanced our two conflicting observational objectives of recording sufficiently long flight paths in a large volume of interest and recording the animals so that they appeared sufficiently large in all camera views. The resulting camera placement yielded a level of uncertainty in estimated 3D locations of the animals that was less than the length of a bat and half the length of a bird. A different camera placement may have resulted in a different level of accuracy. The Methods section below describes how researchers can determine this level in pre-experiment planning using the proposed easyCamera software tool. In post-experiment processing, easyWand can

be used to estimate the accuracy of the calibration.

Our recordings were part of extensive multi-day experiments. Camera placement and calibration recordings required 45 minutes at the beginning of each daily recording bout. This short setup time may be critical when the study organism or group is on the move and must be followed, or when site location, daily weather patterns, tides, or safety considerations dictate.

We suggest that the methods and implementations provided here will substantially aid biologists seeking to make quantitative measures of animal movements and behavior in field settings, as they have done for our own work on bat and bird flight. While previous studies have gathered similar data, they required heroic attempts at calibration frame construction in the field or carefully controlled field environments. We believe that the ease of setup and accuracy of calibration afforded by our methods opens up a wide range of previously unachievable studies and plan to continue refining the publicly available software implementations to fit a variety of needs.

2.2 Materials and methods

The proposed three-step workflow for performing stereo videography provides guidance on field-experiment planning, capture, and post-processing. During planning, appropriate camera equipment is chosen and the placement of the cameras is determined so that the captured imagery will satisfy observational requirements. In the field, a protocol is followed for moving calibration objects through the volume of interest. During post-processing, the image locations of the calibration points are digitized in each camera view. Then, easyWand uses corresponding points to estimate the relative position and orientation of the cameras.

2.2.1 Experiment planning

Camera systems

When selecting a multi-camera system, scientists should consider whether the frame rate, spatial resolution, field-of-view, and synchronization ability of the cameras are appropriate for the size and speed of the study organisms. Camera synchronization, in particular, is a critical requirement for successful stereoscopy. In our field work, we used hardware synchronization to ensure accurate temporal alignment of frames across cameras. Multi-camera systems without precise frame synchronization could be calibrated using these methods if wand and animal pixel motion per frame is small, a large number of calibration points are used and some motionless background points are visible in all cameras.

For capturing video of Brazilian Free-tailed Bats, we used three thermal infrared cameras (FLIR SC8000, FLIR Systems, Inc., Wilsonville, OR) with variable-focus 25-mm lenses, and a pixel width of 18 μm , providing a 40.5° field of view. The 14-bit per pixel video has a frame size of 1024 \times 1024 pixels and frame rate of 131.5 Hz.

Camera placement

When designing camera placement, scientists should consider their observational objectives, the amount of 3D reconstruction uncertainty they may tolerate, and potential additional requirements introduced by the manual or automatic post-experiment video analysis.

Two observational objectives that commonly conflict are the size of the volume of space in which the animals are observed and the spatial resolution at which they are observed. When designing an experiment for such studies, we suggest to impose a lower bound on the size of animals in the image, so that they are not recorded at sizes that will make post-experiment analysis difficult. Based on a pinhole camera model,

the bound x_{max} on the pixel-span of the animal in the image can only be guaranteed if the observation distance between animals and each camera is at most $D_{max} = \frac{fX}{x_{max}p}$, where X is the length of the animal, f is the focal length of the camera, and p is the physical width of a pixel. For our studies, in the interest of observing the flight paths of the animals over a large distance, we chose to allow a small image size. A 10-pixel nose-to-tail span of a 10-cm-long bat in an image was ensured for animals that flew at distances smaller than $D_{max}=(25 \text{ mm} \times 10 \text{ cm})/(10 \text{ pixels} \times 18 \text{ } \mu\text{m}/\text{pixel})=13.8$ m from our thermal cameras.

An inescapable source of uncertainty in the stereoscopic reconstruction of a 3D point is the quantization of intensity measurements (light or thermal radiation) into an array of discrete pixels. Each pixel, projected into space, defines a pyramidal frustum expanding outward from the camera. The location of the 3D point resides somewhere in the intersection of the frustums defined by pixels in each image. For any camera configuration, we can estimate this uncertainty for every 3D point observed by at least two cameras via simulation using the easyCamera software. Our procedure first projects the 3D point onto the image plane of each camera and quantizes the location of each projection according to the pixel grid of each camera. The discrete pixel coordinates of the image points in each camera are then used to reconstruct a 3D position via triangulation. The reconstruction uncertainty is finally computed as the difference between the original and reconstructed positions of the point.

Our simulation results, shown in Fig. 2-2, indicate that the size and shape of the observation volumes and the uncertainty due to quantization within these volumes can differ significantly depending on the number of cameras and their placement. Ensuring that the angle between the optical axes is not wider than the field-of-view angle of the cameras leads to "open" intersection volumes that extend infinitely far away from the cameras (all examples in Fig. 2-2 except C), which is desirable because

it facilitates recording even if the animals appear in an unexpected location. The level of uncertainty increases with the distance from the cameras because the volume of the intersection of the pixel frustums also increases with this distance.

In addition to the reconstruction uncertainty created by quantization, we also consider the reconstruction uncertainty arising from the difficulty in identifying the location of an animal in an image (Fig. 2·3). The location of an animal is often thought of as a single point, e.g., at the center of its body. Localization accuracy of this ill-defined point depends on the resolution of the animal in the image (Fig. 2·3B). To estimate uncertainty in localization, we included a stochastic element to the simulation procedure described above by adding noise to the 2D projections before quantization. Over 100 trials, the RMS distance between the original and reconstructed scene points gives an estimate of the reconstruction uncertainty at the original point (Fig. 2·3D). The camera placements we selected for our bat field experiments (Fig. 2·1A), were similar to the configuration shown in Fig. 2·2B. With our simulation, we were able to determine, prior to any field work, that the levels of uncertainty due to quantization and localization issues would be acceptable for us (Fig. S1).

Reconstruction error occurs when image locations corresponding to different animals are mistakenly used to reconstruct 3D positions. These "data association" errors are commonly made by automated tracking methods, especially if the animals appear similar and small in the images. Camera selection and placement can reduce the potential occurrence of data association errors by imposing appropriate geometric constraints on the triangulation (Fig. 2·4). We recommend use of three or more cameras and a non-collinear camera placement that ensures that the image planes are not parallel (avoiding the configuration in Fig. 2·2F).




Figure 2-2: Reconstruction uncertainty due to quantization effects is shown for six hypothetical camera configurations. The cameras were simulated to have a pixel width of $18 \mu\text{m}$ and field-of-view angle 40.5° and be positioned at a fixed height Z and aimed at a common, equidistant fixation point $F = (0, 0, Z)$. Horizontal cuts of the 3D view frustums of the cameras at height Z and lines at $D_{max}=20$ m are shown from above. Placing the cameras further apart reduces reconstruction uncertainty (A vs. B). If the cameras are placed too far apart (C), however, the view volume is "closed," and there are unobservable regions of space where the cameras will be looking past each other. If the distance between the outermost cameras is held constant, adding additional cameras may not decrease the uncertainty due to image quantization in the common observable region (D vs. E). When the image planes of the cameras are parallel (F), the common view volume is smaller and further away from the cameras than in the other configurations. -Note that these 2D cuts of the 3D view frustums are at the level and elevation angle of the cameras, cuts at a different level or angle would show slightly greater reconstruction uncertainty but similar trends.




Figure 2-3: Reconstruction uncertainty due to quantization and resolution issues. In a video frame obtained for a flight study (A), the automatically detected locations of the animals may not be at their centers (colored dots in B). When estimating reconstruction uncertainty (C, D), we include this effect by corrupting the image projections of simulated world points, generated throughout the whole space, with Gaussian noise where the standard deviation is one-sixth of the calculated apparent size of an animal at that location (circles in B). When estimating the reconstruction uncertainty, including image location ambiguity (D) increases the estimated uncertainty more than three-fold over image quantization alone (C) (note the change in color scale).

2.2.2 Protocol for field experiment

The protocol we recommend for field work includes two phases. In the first phase, prior to any camera setup and recordings, a preliminary plan for the location of the recording space and placement of the cameras is made. The easyCamera software is then used to estimate the uncertainty in localizing the study organism in this space, and adjustments to the plan can be made by experimenting with other hypothetical camera configurations. In the second phase, the actual camera setup in the field can be done easily because no field measurements of camera pose or distances to the animals of study are needed. The only measurements required are references for scene scale and orientation. In our experiments, the known length of a calibration wand moved through the scene provided a scale reference, and gravitational acceleration, estimated from the ballistic trajectories of thrown objects, provided a reference for

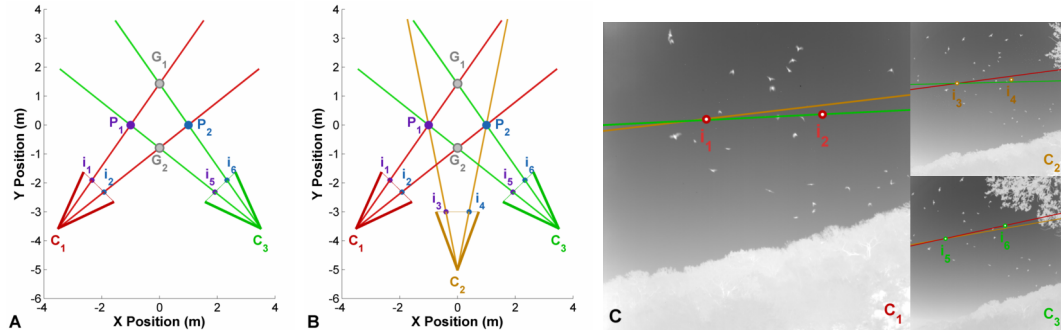


Figure 2.4: An illustration of how, using three cameras, instead of two, in an appropriate configuration, leads to reduced data association errors. A data association error may occur with a two camera system (A) if image points i_1 and i_6 are used to reconstruct "ghost" point G_1 , and i_2 and i_5 to reconstruct "ghost" point G_2 . The correct association would be to match i_1 and i_5 to reconstruct P_1 , and i_2 and i_6 to reconstruct P_2 . If a third camera is introduced (B), this error cannot be made because the image points i_3 and i_4 can only be projections of P_1 and P_2 . A similar three-camera configuration was used for the bat recordings. It facilitated the process of matching corresponding image points across camera views: The ray through i_3 and P_1 (yellow) and the ray through i_5 and P_1 (green) appear as "epipolar" lines in the image of camera C_1 (large image in panel C), which intersect in image point i_1 (red). Similarly, the projections into camera C_2 of the (red) ray through i_1 and P_1 and the (green) ray through i_5 and P_1 at image point i_3 in camera C_2 (yellow in panel C); the projections into camera C_3 intersect at i_5 (green). These "epipolar constraints" are only satisfied if the image points i_1 , i_3 , and i_5 are matched to reconstruct the position P_1 . A collinear camera placement (Fig. 2.2F) would result in parallel epipolar lines, making the matching of corresponding image points difficult.

scene orientation.

Our calibration method generally produces more accurate results when more sets of corresponding image points are used (Table S5, column 2). Thus, we used recordings of the animals, digitized automatically using a preliminary wand-only calibration, to augment the number of calibration points and volume encompassed by them. This augmentation is a feature of our SBA based calibration pipeline not possible with calibration frame based DLT methods. See Table S5 for an exploration of the effects of using animal points in the calibration. We typically recorded our study videos of bats and birds after obtaining videos of calibration objects, but this order can be

reversed.

2.2.3 Post-experiment camera calibration

Our easyWand calibration software bundles a modular pipeline of algorithms that can be used to estimate the relative positions and orientations of the cameras and their intrinsic parameters. The first, most time-consuming step of the calibration procedure is to manually or automatically digitize the image locations of objects recorded in all views. In our post-experiment analysis, we identified thousands of sets of matching image points.

Using the focal lengths and principal points obtained directly from the lenses and image sensors as preliminary estimates of the intrinsic camera parameters and the 8-point algorithm (Hartley and Zisserman, 2004), our software computes preliminary estimates of the camera pose and 3D positions of the calibration objects. Our software then applies the sparse bundle adjustment algorithm (Lourakis and Argyros, 2009) to obtain refined estimates for all calibration parameters. Finally, it converts to a representation of the camera calibration parameters in the form of the DLT coefficients in order to easily integrate into previously existing workflows. None of the 8-point, SBA, and DLT algorithms explicitly requires use of a wand, and other sources of matched camera points could be used as input. Wands are convenient for their mobility and as means to measure scene scale and conduct additional error checking.

The easyWand software tool computes three measures of calibration inaccuracy: The "reprojection error," measured in pixels for each camera, is the RMS distance between the original and reprojected image points of each calibration point, where the "reprojected" image points are computed using the estimated 3D position of the calibration point and the estimated camera parameters. The second measure of inaccuracy is the ratio of the standard deviation of wand-length estimates to their mean.

A large ratio may indicate problems with the calibration, for example, unidentified lens distortion. The third measure is the average uncertainty in the position of each wand tip, estimated from the distance between the two tips. The easyCamera tool can be used with the estimated extrinsic parameters to compute the uncertainty of the reconstructed 3D positions of the study animals (Figs. S2 and S3).

2.3 Observations of Bats

The nightly emergences of bats from two large maternity colonies. Davis cave in Blanco county, Texas and Bamberger Chiroptorium in Johnson City, Texas, were imaged against vegetation and clouded sky. Thermal infrared video was recorded at 131.5 Hz using three hardware-synchronized FLIR SC8000 cameras with 25mm lenses. The field of view was approximately 40 degrees and the image resolution was 1024 by 1024 pixels. The apparent nose-to-tail length of the bats was between 8 and 20 pixels. We made observations on one night at Davis cave, and three nights at the Bamberger Chiroptorium. The duration of recording one each night varied between 12 and 40 minutes.

For capturing video of Brazilian Free-tailed Bats, we used three thermal infrared cameras (FLIR SC8000, FLIR Systems, Inc., Wilsonville, OR) with variable-focus 25-mm lenses, and a pixel width of 18 μm , providing a 40.5° field of view. The 14-bit per pixel video has a frame size of 1024×1024 pixels and frame rate of 131.5 Hz. The three cameras were spatially calibrated following the protocol described by Theriault et al. (2014).

Image locations of bats were estimated using an image processing pipeline similar to Betke et al. (2007). The background was modeled using a running average with exponential decay. Foreground regions were identified using background subtraction with empirically determined thresholds. Multiple detected objects per foreground

region were identified using brightness peaks, defined as local maxima that were sufficiently bright (thresholds determined heuristically) and sufficiently far away from other local maxima (based on apparent target size). We followed the reconstruction-tracking approach described by Wu et al. (2009); image locations of bats from each of the three cameras were used to reconstruct three dimensional object detections, using epipolar geometry and the DLT algorithm. Tracks were constructed using multiple hypothesis tracking with a sliding window. Spurious tracks shorter than 100 ms in duration or 1 meter in length were discarded. Up to 295 bats were simultaneously observable via 3D reconstruction and tracking. Over 20 million individual data points were obtained from over 200,000 trajectories from the dataset taken at the Davis site. Between 2 million and 4.9 million total data points were obtained from between 44,000 and 87,000 trajectories in each of the three nights at the Bamberger site. From 70 minutes of video showing 11 bursts of activity, over 392,000 flight trajectories with over 31,000,000 data points were reconstructed using image processing, epipolar geometry, and multi-target tracking. The object locations in each track were filtered with a Kalman filter to mitigate the effects of reconstruction uncertainty (Theriault et al., 2014).

2.4 Occlusion reasoning

Since we model the bats as points, it is important for the image detection to be located on the body as consistently as possible. Our image processing pipeline works for this goal by localizing the brightest (warmest) part of the thorax. When bats occlude or partially occlude each other in the original video, the image detection for one or more of the bats may be displaced. This is a problem because an inaccurate image location can lead to an inaccurate 3D reconstructed position. This inaccuracy can propagate to computation of higher order moments, including velocity and acceleration. The net

effect is that bats that are in shadow tend to have larger apparent heading variability, speed, and acceleration magnitude than bats that are not occluded. Therefore, it is important to identify bats that are participating in an occlusion event so that they can be excluded from computations that aggregate information (e.g. the distribution of bat flight speed). Some properties we compute depend on the velocities of neighboring bats, and so if a neighbor is in shadow, then that computation should either exclude the neighbor, or the property for the focal bat should be excluded.

The potential displacement of the image detection will be more severe when the bats are closer together. For example, the image of two bats that partially occlude each other because the image of their wings are touching is much less likely to result in a displaced image detection than two bats that occlude because the wing of one bat is covering part of the body of the other bat. The potential image detection displacement affects both bats in the interaction. The reconstruction accuracy will be further diminished if the bats are occluding in more than one camera. When bats are far apart in 3D space, this doesn't reduce the effect of occlusion on image detection accuracy, though it does reduce the chance that the two bats will occlude each other in more than one camera.

Occlusion occurs when two or more objects are located along the same ray from the camera center and out into the scene. Given a representative 3D point on an object, \mathbf{x}_i , and the camera center \mathbf{c}_c , the view vector from the camera center through \mathbf{x}_i is given by $l_{i,c}[t] = \frac{\mathbf{x}_i - \mathbf{c}_c}{\|\mathbf{x}_i - \mathbf{c}_c\|}$. The distance between some other point, \mathbf{x}_j and this line is given by $d(l_{i,c}[t], \mathbf{x}_j) = \|\mathbf{x}_j - (\mathbf{x}_j \cdot l_{i,c}[t])l_{i,c}[t]\|$. Given the wingspan of the bats, $d_{\text{wing}}=30$ cm (for *Tadarida brasiliensis* (Farney and Fleharty, 1969)), we conclude that a bat i at time t is occluded by bat j at time t if $d(l_{i,c}[t], \mathbf{x}_j) < \frac{1}{2}d_{\text{wing}}$ for any of the three cameras ($c \in \{1, 2, 3\}$), or if the sum of the distances between all of the lines of sight $\sum_{c=1}^3 d(l_{i,c}[t], \mathbf{x}_j) < 3 d_{\text{wing}}$.

Chapter 3

The FlockOpt model of collective motion

3.1 Introduction

The study of the collective behavior of animals has fascinated scientists for over 100 years, (Pryer, 1884; Breder, 1954; Hamilton, 1971) and continues to be a very active area of research (Calovi et al., 2014; Buhl et al., 2006; Krause et al., 2009; Nagy et al., 2010; Bode et al., 2011; Bialek et al., 2012; Attanasi et al., 2014; Gerum et al., 2013; Mann et al., 2013; Portugal et al., 2014).

Computational models of group motion are algorithms designed to control simulated agents so that the behavior of the group will be similar to natural systems. Computational models can be judged by fidelity, richness, and parsimony. Simulations performed with computational models should have properties in common with natural systems, and the model should be able to generate a wide variety of behaviors by varying a small number of parameters.

Through simulation studies, the scientific community has discovered that large groups of entities, such as simulated particles or living organisms, can move coherently without leaders when group members exhibit simple behaviors in response to the relative position and heading of neighboring individuals (Aoki, 1982; Reynolds, 1987; Vicsek et al., 1995; Czirók et al., 1999; Couzin et al., 2002). In an influential simulation model by Aoki (1982), repulsion, attraction, and heading alignment influences are applied to neighbors within concentric zones in order to govern the behavior

of individual particles. In models by Vicsek et al. (1995) and Czirók et al. (1999), heading alignment is the only influence. In social potential field models, introduced by Reif and Wang (1999), influences of attraction and repulsion are understood in terms of the gradients of potential functions, inspired by physical forces.

The notion of similarity used to compare models and natural data has evolved over time. Early models were judged based on emergent characteristics, such as the formation adopted by the models (Aoki, 1982; Couzin et al., 2002). There is a growing body of work studying the challenges of matching a model of collective motion and associated parameters with a biological study system. Authors have observed that many different models may lead to similar macroscopic behavior (Strömbom, 2011; Lopez et al., 2012; Mann et al., 2013). This observation led to increased interest in new methodologies for building models in a data-driven way (Gautrais et al., 2012; Mann et al., 2013) and measuring the goodness of fit between a model and observed behavior of animals (Eriksson et al., 2010; Mann, 2011; Bode et al., 2011). Mann et al. (2013) suggested that for a model to be considered a good fit to a biological system, it must predict aspects of the behavior at both individual and system-wide scales.

In recent years, as it has become feasible to collect large datasets of thousands of observations of animals, detailed studies have been carried out to study the behavior of locusts (Buhl et al., 2006), starlings (Ballerini et al., 2008; Bialek et al., 2012; Cavagna et al., 2010; Attanasi et al., 2014), surf scoters (ducks) (Lukeman et al., 2010), mosquitofish (Herbert-Read et al., 2011), golden shiners (fish) (Katz et al., 2011), barred flagtails (fish) (Gautrais et al., 2012), and glass prawns (Mann et al., 2013). Some of these studies have provided an opportunity to reexamine some of the widespread assumptions underpinning many models of collective motion.

Early models defined neighborhoods in terms of metric, absolute distances (Aoki,

1982; Couzin et al., 2002). Ballerini et al. (2008) provided evidence that interactions in flocks of starlings are dependent on a topological distance instead. In most models, the behavioral response to a set of neighbors is computed as an average of the responses to each neighbor individually (possibly with weights) (Aoki, 1982; Vicsek et al., 1995; Couzin et al., 2002; Gazi and Passino, 2004). In a detailed study of golden shiners, Katz et al. (2011) discovered that the assumption that the response of an animal to a set of neighbors could be understood as an average of the responses to each neighbor individually was not consistent with data.

The vast majority of existing models assume that simulated particles move with a constant (Vicsek and Zafeiris, 2012) or random (Aoki, 1982; Huth and Wissel, 1992) speed (exceptions include Mishra et al. (2012), Reif and Wang (1999), Reynolds (1987), and Bode et al. (2010)) because early authors (Aoki, 1982) concluded that the heading of animals was the most important aspect of their motion. In two detailed studies of fish behavior, Herbert-Read et al. (2011) and Katz et al. (2011) found that speed regulation is actually a very important aspect of interactions between fish. In response to findings revealing the importance of speed regulation, variable speed models have recently been introduced (Bode et al., 2010; Mishra et al., 2012). In these models, the heading of the entities is determined in the same way as in older, constant-speed models, and the speed is determined separately, either as a function of a property of the local neighborhood (Mishra et al., 2012) or based on a behavioral rule (Bode et al., 2010).

3.2 Contribution

We introduce FlockOpt, an optimization-based, variable-speed, self-propelled particle model of collective motion where the speeds and headings of simulated particles are determined holistically. Groups of particles simulated with our model spontaneously

exhibit ordered motion, even though our algorithm does not explicitly prescribe heading alignment. Analysis of our model explains how velocity matching follows from the impetus of particles to maintain their relative positions, revealing that the heading alignment influence as formulated in earlier models may be redundant for understanding ordered motion in natural systems. Our model also predicts that the speed and heading of entities should be coupled by a causal dependence on maintaining relative position. We contribute to the growing body of work comparing the predictions of computational models with observations of animals by studying the behavior of *Tadarida brasiliensis*, a gregarious bat that forms large summer maternity colonies containing hundreds of thousands to millions of individuals in the southwestern United States (Betke et al., 2008). By measuring how spatial relationships change over time, we found that individual variations in speed and heading are coupled, consistent with our model.

3.3 The FlockOpt Model

FlockOpt is a self-propelled particle model of collective motion where individual particles make decisions about how to move by performing a constrained optimization of a local objective function, \mathcal{O}_i , computed for each particle at each moment in time based on the kinematics of the particle and the relative position and kinematics of its neighbors. Repulsion and attraction occur when moving away from or closer to neighboring particles improves the objective value that a particle experiences. Alignment of heading is not explicitly prescribed, but is an emergent effect (explained in the next section).

To reinterpret the design of classic zonal models under our paradigm, we design a continuous function that encodes the preference over the distance between two particles, Φ (Figure 3-1). To represent the range of distances that are equally desirable




Figure 3.1: The objective function Φ is constructed based on the difference between two translated and horizontally scaled sigmoid functions. Appropriate constants are chosen to adjust the width of the central, flat portion, z_o , (left) and the width of the descending leg, z_a (right).

(z_o), there is an elevated flat part in the middle. To represent the less desirable distances that are too close (z_r) and too far (z_a) there are flanking ascending and descending legs. Consistent with the maximum sight distance imposed by many models, the function flattens beyond a certain distance, because all further distances are equally bad. An appropriate shape for this neighbor distance preference function, Φ (Figure 3.1), is given by the difference of two translated and horizontally scaled sigmoid functions as follows

$$\Phi(\gamma; \theta) = \phi(\gamma; \theta_1, \theta_2) - \phi(\gamma; \theta_3, \theta_4), \quad (3.1)$$

where the sigmoid function with slope $\frac{\alpha}{4}$ centered at β is given by $\phi(\gamma; \alpha, \beta) = (1 + e^{-\alpha(\gamma+\beta)})^{-1}$ and constants, θ , are chosen to reflect the desired shape.

The state of particle i at time t consists of position, \mathbf{x}_i , and velocity, \mathbf{v}_i , (heading and speed) and its neighbors are $\mathbf{N}_i = \{p_k\}$. Time is discretized into steps, Δt . The arbitrary spatial units of the simulation are denoted by Δu instead of an absolute unit, such as meters. The current relative position of particle i and particle k at time t is given by $\mathbf{x}_{i,k} = \mathbf{x}_k - \mathbf{x}_i$. The current relative velocity between the two particles is given by $\mathbf{v}_{i,k} = \mathbf{v}_k - \mathbf{v}_i$.

If, at time t , particle p_i adopts a hypothetical change in velocity $\Delta \tilde{\mathbf{v}}_i$ while the neighbor p_k maintains its current velocity, then the anticipated relative velocity be-

tween a particle and its neighbor will be $\widehat{\mathbf{v}}_{i,k} = \mathbf{v}_{i,k} - \Delta\widetilde{\mathbf{v}}_i$. The anticipated relative position, $\widehat{\mathbf{x}}_{i,k}$, between entities i and k in the future will be

$$\begin{aligned}\widehat{\mathbf{x}}_{i,k} &= (\mathbf{x}_k + \mathbf{v}_k\Delta t) - (\mathbf{x}_i + \mathbf{v}_i\Delta t + \Delta\widetilde{\mathbf{v}}_i\Delta t) \\ &= \mathbf{x}_{i,k} + (\mathbf{v}_{i,k} - \Delta\widetilde{\mathbf{v}}_i)\Delta t \\ &= \mathbf{x}_{i,k} + \widehat{\mathbf{v}}_{i,k}\Delta t\end{aligned}\quad (3.2)$$

The desirability of the anticipated future position of a single neighbor (Figure 3·2, left) is given by

$$\Phi(\|\widehat{\mathbf{x}}_{i,k}\|; \theta). \quad (3.3)$$

To express the desirability of the predicted future configuration of the entire local neighborhood, given a candidate change that may be adopted, $\Delta\widetilde{\mathbf{v}}_i$, the desirability of all of the neighbors are summed (Figure 3·2, middle)). The joint objective function over the whole neighborhood is given by

$$\mathcal{O}_i(\Delta\widetilde{\mathbf{v}}_i; \theta) = \sum_{k \in \mathbf{N}_i} \Phi(\|\widehat{\mathbf{x}}_{i,k}\|; \theta). \quad (3.4)$$

Flying animals must maintain sufficiently fast flight in order to avoid aerodynamic stall, but excessively fast flight may be energetically unsustainable. Under our model, such costs and constraints may either be folded into the objective function via penalty terms, or expressed as a constrained optimization (Figure 3·2, right). The hypothetical vector velocity of a particle due to $\Delta\widetilde{\mathbf{v}}_i$ is given by $\widehat{\mathbf{v}}_i = \mathbf{v}_i + \Delta\widetilde{\mathbf{v}}_i$. We include a preference for some particular speed s^* using a penalty term, λ_s , leading to

$$\begin{aligned}\mathcal{O}_i(\Delta\widetilde{\mathbf{v}}_i; \theta) &= \\ &[\sum_{k \in \mathbf{N}_i} \Phi(\|\widehat{\mathbf{x}}_{i,k}\|; \theta)] - \lambda_s(\|\widehat{\mathbf{v}}_i\| - s^*)^2.\end{aligned}\quad (3.5)$$

To restrict the range of allowable speeds, we define constraints on the minimum and maximum speeds s_{\min} and s_{\max} . To encourage system stability, we also define a constraint on the maximum change in velocity, v_{\max} . Then, we perform the following




Figure 3-2: The objective function \mathcal{O}_i , which reflects the desirability of the anticipated future neighborhood of particle i (\mathbf{v}_i in dark blue), is visualized in a 2D simulation with one (left), two (middle), and six neighbors (right). Function values are based on the desirability $\Phi(\|\mathbf{x}_{i,k}\|)$ of the anticipated future distance $\|\mathbf{x}_{i,k}\|$ (dashed line) to its k th neighbor (\mathbf{v}_k in light blue). In 2D, high values of \mathcal{O}_i form rings surrounding the anticipated positions of the neighbors (left). Regions of intersecting rings are most desirable for particle i to move to (middle). A constrained optimization on \mathcal{O}_i is performed (right) to find the best change $\Delta \mathbf{v}_i^*$ in the velocity of particle i (red) that satisfies additional constraints, such as maintaining a preferred speed (light blue annulus) or limiting its turning angle (pink annular sector).

constrained optimization (Figure 3-2, right), using sequential quadratic programming (Nocedal and Wright, 2006)

$$\begin{aligned} \Delta \mathbf{v}_i^* &= \arg \max_{\Delta \tilde{\mathbf{v}}_i \in \mathcal{R}^d} \mathcal{O}_i(\Delta \tilde{\mathbf{v}}_i; \theta) \\ \text{s.t. } &\begin{cases} \|\hat{\mathbf{v}}_i\| \leq s_{\max} \\ \|\hat{\mathbf{v}}_i\| \geq s_{\min} \\ \|\Delta \tilde{\mathbf{v}}_i\| \leq v_{\max}. \end{cases} \end{aligned} \quad (3.6)$$

The FlockOpt model extends trivially to three or more dimensions with the neighbor preference function and penalties we have described, since the functions are defined only over the distances between points or the lengths of velocity vectors.

In order to incorporate modeling of noisy perception in the FlockOpt model, without defining a new objective function, random noise could be added to the predicted positions of the neighbors before evaluating the objective function, as $\mathcal{O} = \sum_{k \in \mathbf{N}_i} \Phi(\|\hat{\mathbf{x}}_{i,k} + \mathcal{N}(0, \sigma)\|)$. As an implementation detail, the noise added to

the perceived neighbor position should be held constant through all iterations of the optimization to enable convergence.

3.4 How ordered motion arises

In our paradigm, each particle always takes action to optimize the objective function. There are two distinct sets of circumstances: when a particle can improve its objective value, and when it cannot. Most importantly, as an emergent effect that is not explicitly prescribed, if no improvement is possible, or if a decrease is inevitable due to the aggregate motion of its neighbors, in order to minimize the decrease in the objective value that a particle experiences, it must minimize the change in relative position between it and its neighbors. This effect leads to velocity matching (below). We use the notation as defined in Section 3.6.

The current relative position of entity i and entity k at time t is given by $\mathbf{x}_{i,k} = \mathbf{x}_k - \mathbf{x}_i$. If, at time t , particle i adopts a hypothetical change in velocity $\Delta\tilde{\mathbf{v}}_i$ while the neighbor, particle k , maintains its current velocity, then the predicted relative position between entities i and k in the future will be $\hat{\mathbf{x}}_{i,k} = (\mathbf{x}_k + \mathbf{v}_k\Delta t) - (\mathbf{x}_i + \mathbf{v}_i\Delta t + \Delta\tilde{\mathbf{v}}_i\Delta t)$.

Observe that $\hat{\mathbf{x}}_{i,k} = \mathbf{x}_{i,k} + \hat{\mathbf{v}}_{i,k}$, where $\hat{\mathbf{v}}_{i,k} = \mathbf{v}_{i,k} - \Delta\tilde{\mathbf{v}}_i$ represents the total relative motion between two entities due to both the relative motion due to the current relative velocities, $\mathbf{v}_{i,k} = \mathbf{v}_k - \mathbf{v}_i$, and the hypothetical change to be adopted, $\Delta\tilde{\mathbf{v}}_i$. We will use the Φ symbol from our description of our model, but the exact shape is not important for the purpose of this explanation.

We consider an additive joint objective function over all neighbors in the local neighborhood, such as

$$\mathcal{O}_i(\Delta\tilde{\mathbf{v}}_i; \theta) = \sum_{k \in \mathbf{N}_i} \Phi(\|\hat{\mathbf{x}}_{i,k}\|; \theta). \quad (3.7)$$

Due to the motion of entities in the system, the objective value experienced by

each entity will vary over time, as follows

$$\begin{aligned}
& \Delta \mathcal{O}_i(\Delta \tilde{\mathbf{v}}_i; \theta) \\
&= \sum_{k \in \mathbf{N}_i} \Delta \Phi(\|\hat{\mathbf{x}}_{i,k}\|; \theta) \\
&= \sum_{k \in \mathbf{N}_i} \Phi(\|\hat{\mathbf{x}}_{i,k}\|; \theta) - \Phi(\|\mathbf{x}_{i,k}\|; \theta).
\end{aligned} \tag{3.8}$$

Recalling that $\hat{\mathbf{x}}_{i,k} = \mathbf{x}_{i,k} + \hat{\mathbf{v}}_{i,k}$, the change in the objective value can be approximated using $\Delta \Phi_i(\|\hat{\mathbf{x}}_{i,k}\|; \theta) \approx \nabla \Phi(\|\mathbf{x}_{i,k}\|; \theta)^T \hat{\mathbf{v}}_{i,k}$ due to a first order Taylor series expansion as follows:

$$\begin{aligned}
& \Delta \Phi(\|\hat{\mathbf{x}}_{i,k}\|; \theta) \\
&= \{\Phi(\|\hat{\mathbf{x}}_{i,k}\|; \theta)\} - \Phi(\|\mathbf{x}_{i,k}\|; \theta) \\
&\approx \{\Phi(\|\mathbf{x}_{i,k}\|; \theta) + \nabla \Phi(\|\mathbf{x}_{i,k}\|; \theta)^T \hat{\mathbf{v}}_{i,k}\} - \Phi(\|\mathbf{x}_{i,k}\|; \theta) \\
&= \nabla \Phi(\|\mathbf{x}_{i,k}\|; \theta)^T \hat{\mathbf{v}}_{i,k}.
\end{aligned} \tag{3.9}$$

This formalism captures the intuition that, for each neighbor, the anticipated change in the objective value due to the anticipated relative motion, $\hat{\mathbf{v}}_{i,k}$, will be larger when the particles move directly towards or away from each other, and smallest when the direction of the relative velocity is perpendicular to the vector between the two neighbors. The change in the objective value will also be larger when the slope of $\Phi(\|\mathbf{x}_{i,k}\|; \theta)$ is larger, e.g. when the current distance between the two neighbors falls on the ascending or descending legs of the distance preference function as opposed to a flat part of the function.

Finally, the change will be proportional to the magnitude of the anticipated relative velocity between the two particles, $\|\hat{\mathbf{v}}_{i,k}\|$ (particles moving towards each other faster will generally experience a larger change in the objective value). This last point is very important because in a system containing many particles, each particle will be surrounded by neighbors, so, generally, it will not be possible for a particle to minimize the change in the objective value by moving perpendicularly to all neighbors

simultaneously (exceptions may occur at the edges of the group). Minimizing the magnitude of the anticipated relative velocity, $\|\widehat{\mathbf{v}}_{i,k}\|$, between two particles is equivalent to minimizing the magnitude of the change in the relative position, $\|\widehat{\mathbf{x}}_{i,k} - \mathbf{x}_{i,k}\|$, between two particles.

By expanding $\widehat{\mathbf{v}}_{i,k}$ as

$$\begin{aligned}\widehat{\mathbf{v}}_{i,k} &= \mathbf{v}_k - (\mathbf{v}_i + \Delta\widetilde{\mathbf{v}}_i) \\ \widehat{\mathbf{v}}_{i,k} &= (\mathbf{v}_k - \mathbf{v}_i) - \Delta\widetilde{\mathbf{v}}_i \\ \widehat{\mathbf{v}}_{i,k} &= \mathbf{v}_{i,k} - \Delta\widetilde{\mathbf{v}}_i,\end{aligned}\tag{3.10}$$

we obtain the following illustrative equation:

$$\Delta\mathcal{O}_i(\Delta\widetilde{\mathbf{v}}_i; \theta) \approx \left\{ \sum_{k \in \mathbf{N}_i} \nabla\Phi(\|\mathbf{x}_{i,k}\|; \theta)^T \mathbf{v}_{i,k} \right\} - \left\{ \sum_{k \in \mathbf{N}_i} \nabla\Phi(\|\mathbf{x}_{i,k}\|; \theta)^T \Delta\widetilde{\mathbf{v}}_i \right\}.\tag{3.11}$$

From this last equation, we can understand the change in the objective value as composed of two parts, one due to the current relative motion of the neighbors, $\mathbf{v}_{i,k}$, and another due to the motion that the focal entity is trying to choose, $\Delta\widetilde{\mathbf{v}}_i$.

In circumstances when an improvement in the objective value is possible due to either or both terms, then the entity will derive the most benefit from maximizing $|\Delta\mathcal{O}_i|^2$ and we cannot predict $\Delta\widetilde{\mathbf{v}}_i$ without knowing the exact objective function to be optimized.

On the other hand, when an entity will experience a decrease in the objective value due to the aggregate motion of its neighbors ($\mathbf{v}_{i,k}$), and it is not possible to completely compensate by choosing an appropriate $\Delta\widetilde{\mathbf{v}}_i$, then the entity will suffer the least harm by minimizing $|\Delta\mathcal{O}_i|^2$ and the behavior can be predicted as follows.

Due to the chain rule,

$$\nabla\Phi(\|\mathbf{x}_{i,k}\|; \theta) = \alpha_{i,k} \frac{\mathbf{x}_{i,k}}{\|\mathbf{x}_{i,k}\|},\tag{3.12}$$

where α_k is the slope of Φ at $\|\mathbf{x}_{i,k}\|$, given θ . Then

$$\Delta\mathcal{O}_i(\Delta\tilde{\mathbf{v}}_i; \theta) \approx \sum_{k \in \mathbf{N}_i} \alpha_{i,k} \frac{\mathbf{x}_{i,k}}{\|\mathbf{x}_{i,k}\|}{}^T \widehat{\mathbf{v}}_{i,k}. \quad (3.13)$$

By examining each term of this expression, we can consider three ways to minimize the harm due to $\Delta\tilde{\mathbf{v}}_i$: 1.) if $\widehat{\mathbf{v}}_{i,k} \perp \mathbf{x}_{i,k}$, 2.) if α_k are small or 3.) if $\|\widehat{\mathbf{v}}_{i,k}\|$, the magnitudes of the anticipated relative velocities, are small.

In a system with many particles, most of the particles will be surrounded by neighbors, so there will be a variety of directions for $\mathbf{x}_{i,k}$ and it will not be possible for a particle to move perpendicularly to all of the neighbors simultaneously (exceptions may exist on the boundary of the group).

If the distance between neighbors falls along a part of Φ with small slope, then $|\alpha_k|$ will be small and $\alpha_{i,k} \frac{\mathbf{x}_{i,k}}{\|\mathbf{x}_{i,k}\|}{}^T \widehat{\mathbf{v}}_{i,k}$ will be small, even if $\|\widehat{\mathbf{v}}_{i,k}\|$ is large. If $|\alpha_k|$ are not small, and $\|\widehat{\mathbf{v}}_{i,k}\|$ is large, then any improvement in the objective value possible due to a large value of α_k for some neighbor must be at least offset by other values of α_k of opposite sign, otherwise an improvement in the objective value would be possible, and our analysis would not apply.

In general, in a situation in a steady state system where particles are surrounded by neighbors and the objective value cannot be improved, choosing $\Delta\tilde{\mathbf{v}}_i$ to minimize $\|\widehat{\mathbf{v}}_{i,k}\| = \|\mathbf{v}_{i,k} - \Delta\tilde{\mathbf{v}}_i\|$ will lead to the smallest decrease in the objective value.

If we assume that each $\Delta\Phi(\|\widehat{\mathbf{x}}_{i,k}\|; \theta) \propto \|\mathbf{v}_{i,k} - \Delta\tilde{\mathbf{v}}_i\|$ individually, then by a least squares derivation, we see that the optimal $\Delta\mathbf{v}_i^*$ is the average of the relative velocity of the neighbors, as follows:

$$\begin{aligned}
\Delta \mathbf{v}_i^* &= \arg \min_{\Delta \tilde{\mathbf{v}}_i} \sum_{k \in \mathbf{N}_i} (\Delta \Phi(\hat{\mathbf{x}}_{i,k}; \theta))^2 \\
\Delta \mathbf{v}_i^* &\propto \arg \min_{\Delta \tilde{\mathbf{v}}_i} \sum_{k \in \mathbf{N}_i} (\|\mathbf{v}_{i,k} - \Delta \tilde{\mathbf{v}}_i\|)^2 \\
\frac{d}{d\Delta \tilde{\mathbf{v}}} \sum_{k \in \mathbf{N}_i} (\|\mathbf{v}_{i,k} - \Delta \tilde{\mathbf{v}}_i\|)^2 &= 0 \\
2 \sum_{k \in \mathbf{N}_i} (\mathbf{v}_{i,k} - \Delta \tilde{\mathbf{v}}_i) &= 0 \\
\Delta \mathbf{v}_i^* &= \frac{1}{|\mathbf{N}_i|} \sum_{k \in \mathbf{N}_i} \mathbf{v}_{i,k}.
\end{aligned} \tag{3.14}$$

It may be too strong to say that the change in the objective value due to each neighbor individually is proportional to $\|\mathbf{v}_{i,k} - \Delta \tilde{\mathbf{v}}_i\|$, since the changes with respect to some neighbors might be negative while others are positive. If we instead assume that $\sum_{k \in \mathbf{N}_i} \Delta \Phi(\|\hat{\mathbf{x}}_{i,k}\|; \theta) \propto \|\sum_{k \in \mathbf{N}_i} (\mathbf{v}_{i,k} - \Delta \tilde{\mathbf{v}}_i)\|$ in the aggregate, then by setting up the desired minimization and simply rearranging terms, we see again that $\Delta \mathbf{v}_i^*$ is the average of the relative velocity of the neighbors as follows:

$$\begin{aligned}
&\sum_{k \in \mathbf{N}_i} \Delta \Phi(\|\hat{\mathbf{x}}_{i,k}\|; \theta) \\
&\propto \left\| \sum_{k \in \mathbf{N}_i} (\mathbf{v}_{i,k} - \Delta \tilde{\mathbf{v}}_i) \right\| \\
&= \left\| \sum_{k \in \mathbf{N}_i} (\mathbf{v}_{i,k}) - |\mathbf{N}_i| \Delta \tilde{\mathbf{v}}_i \right\| \\
&= \left\| \left(\frac{1}{|\mathbf{N}_i|} \sum_{k \in \mathbf{N}_i} \mathbf{v}_{i,k} \right) - \Delta \tilde{\mathbf{v}}_i \right\| \\
\Delta \mathbf{v}_i^* &= \arg \min_{\Delta \tilde{\mathbf{v}}_i} \left\| \left(\frac{1}{|\mathbf{N}_i|} \sum_{k \in \mathbf{N}_i} \mathbf{v}_{i,k} \right) - \Delta \tilde{\mathbf{v}}_i \right\| \\
\Delta \mathbf{v}_i^* &= \frac{1}{|\mathbf{N}_i|} \sum_{k \in \mathbf{N}_i} (\mathbf{v}_{i,k}).
\end{aligned} \tag{3.15}$$

From this line of reasoning, we conclude that ordered motion can arise in groups of entities when each individual entity is unable to improve its relationship to its neighbors and instead takes actions to avoid worsening its relationship to its neighbors. The alignment influence, prescribed as velocity averaging in earlier models, arises as an emergent effect if the preference over relative position is expressed as we prescribe.

We can understand the evolution of a system in two stages. From an initial configurations, entities in the system will initially work to increase their objective

values by changing their relative position. Once the objective values of the entities in the system saturate, and the entities are in a configuration such that most of the entities are unable to increase their objective values, then the entities will take action to avoid decreasing their objective value by maintaining their relative position. These theoretical insights are borne out in our simulation studies; We found that, across a wide range of parameters and initial conditions, the objective values experienced by the particles quickly saturated and remained nearly saturated, despite the occurrence of diverse patterns of motion (Figure 3·7).

In both our simulation studies and observations of bats in flight, we observed that neighborhoods where local headings were more polarized had lower spatial variability (Figure 3·8). In the natural data, this is not the inevitable conclusion: headings could be aligned, but the bats could be moving at different speeds, which would allow the relative positions to change. The natural behavior that we observe is consistent with our model and our explanation of how ordered motion can arise in groups of moving entities.

Our analysis predicts that particles will attempt to maintain their relative positions in order to minimize the decrease in objective values. Due to the first order Taylor series expansion, this is not completely accurate. With an objective function defined over the distances between particles, to maintain objective values, particles will need to minimize the change in *distances* between themselves and their neighbors. This is approximated in the first order by maintaining relative position.

Understanding that ordered motion can arise in order to preserve objective values by maintaining distances leads to an important insight about the dominant modes of group movement, such as translating flocks and rotational mills: euclidean transformations that maintain such distances are composed of rotations and translations.

We can also understand why some objective function shapes are more conducive

to the emergence of stable, coherent translational motion than others. According to our theory, there are two phenomena that can disrupt ordered motion: the extent to which improvements are generally possible and the extent to which many positions are equally good. Given the objective function shape we have used throughout this paper, when z_o is small, unless the motion of the entities is perfectly parallel, the particles will constantly suffer decreases in their objective values due to the motion of the neighbors and take action to recover, frequently causing large system-wide destabilizations. When z_o is large, then many relative positions are equally good and the organization of the motion is much looser. We have chosen a particular shape for the simulation studies in this paper, but many other shapes are possible. We have used objective functions based on both Gaussian and gamma distributions, which lead to behavior that is most like our simulations with the double-sigmoid function with a short z_o .

3.5 Patterns of motion from simulations

We generated both 2D and 3D simulations that were closed systems (no particles entered or left during the simulation) on an open, unbounded volume. We held $z_r = 1\Delta u$ and varied z_o between $0 \Delta u$ and $2.5 \Delta u$. and z_a between $0.5 \Delta u$ and $3.5 \Delta u$. For each setting of z_o and z_a , we varied λ_s between 0.0 and 0.1. Speed was constrained to be less than $s_{\max} = 0.4\Delta u/\Delta t$. We used a topological definition of proximity; each particle interacted with the first six nearest neighbors in 2D or the first twelve nearest neighbors in 3D, subject to a maximum sight distance of $20 \Delta u$. In each simulation, 50 particles were simulated for 1000 steps. In each of 12 initial conditions, positions were uniformly distributed in a square or box around the origin with identical speeds of $0.1 \Delta u$ and headings uniformly distributed on a circle or sphere. At each time step t , the vector velocity of each particle was updated by

calculating $\Delta \mathbf{v}_i^*$ to optimize the objective function by adding the vector components of position and velocity.

We observe that the type of motion exhibited by groups of particles under our simulation model depends most strongly on the width of the flat part of the objective function, z_o , and the speed penalty, λ_s . The value of z_o determines the spacing and tightness of the particles in the group (Figure 3-3). As z_o increases, both the average and variance of the distance between first nearest neighbors increases, independent of other parameters.

In simulations with a small value for z_o , such as 0.0 or 0.5 Δu , the distances between neighbors in a group undergoing stable translational motion will be very regular and small perturbations in the relative position between particles may lead to large disruptions in the overall group motion (Figure 3-4, top row). For example, if the heading of one particle on the edge of the group is not perfectly parallel to the other particles, that particle may move slightly closer to another particle, and this will precipitate a wide destabilization, causing many neighboring particles to change their heading and speed, initiating large rotational motion. With a medium value of z_o , such as 1.0 or 1.5 Δu , even though the spacing between neighbors remains very regular, such perturbations would lead to a smaller disruption involving fewer particles, or merely a change in the direction of translational motion (Figure 3-4, middle row). In groups with a large value of z_o , such as 2.0 or 2.5 Δu , there is much more variability in the distances between neighbors, and much more heterogeneity in headings and speeds within the group (Figure 3-4, bottom row). Even when the overall motion of the group is translational, there will be a variety of headings and speeds, due to a looser spatial organization with holes and gaps. Changes in the spatial relationships between particles are common and individual interactions generally will have little observable effect on the overall motion of the group. Simulations in 2D and 3D were

qualitatively similar (Figure 3·6), though they were not the same as measured by statistical tests of distributions of nearest neighbor distances, kinematics, or other properties.




Figure 3·3: Representative spatial arrangements of particles, are shown from simulations where z_o was varied between $0.0 \Delta u$ and $2.5 \Delta u$ while other parameters were held constant ($\lambda_s = 0.0\Delta u$ and $z_a = 3.0\Delta u$). We found that the width of the flat part (z_o) of the objective function is the primary factor affecting the spatial arrangement among particles in a simulation, whereas the length of the descending leg (z_a) has less impact. When z_o is small (left, $z_o = 0.0\Delta u$), the lattice is tight and the spacing is very regular. When z_o has an intermediate value (middle, $z_o = 1.0\Delta u$) the spacing between particles is still very regular, but there may be some small perturbations in the lattice. When z_o is large (right, $z_o = 2.5\Delta u$), the lattice is loose with holes and irregular spacing between particles.

The speed penalty plays an important role in the type of motion. More rotational structures require a large difference in speeds among particles. A stronger penalty encourages translational motion. With a large speed penalty, rotational disruptions are less common, and events that might have precipitated a large rotational disruption with a lower speed penalty will instead lead to a brief period of disorder. Or, if the rotational disruption occurs, it will quickly dissipate into translational motion, rather than remaining as a fixture of the group motion.

There are different types of rotational motion. In our simulations, when z_o is short and the lattice is tight, the most common type of motion is a translating rotational ball (Figure 3·4, top row. Figure 3·5). The center of rotation is actually on one side of the group, not the middle, and there is an outer shell that rotates around

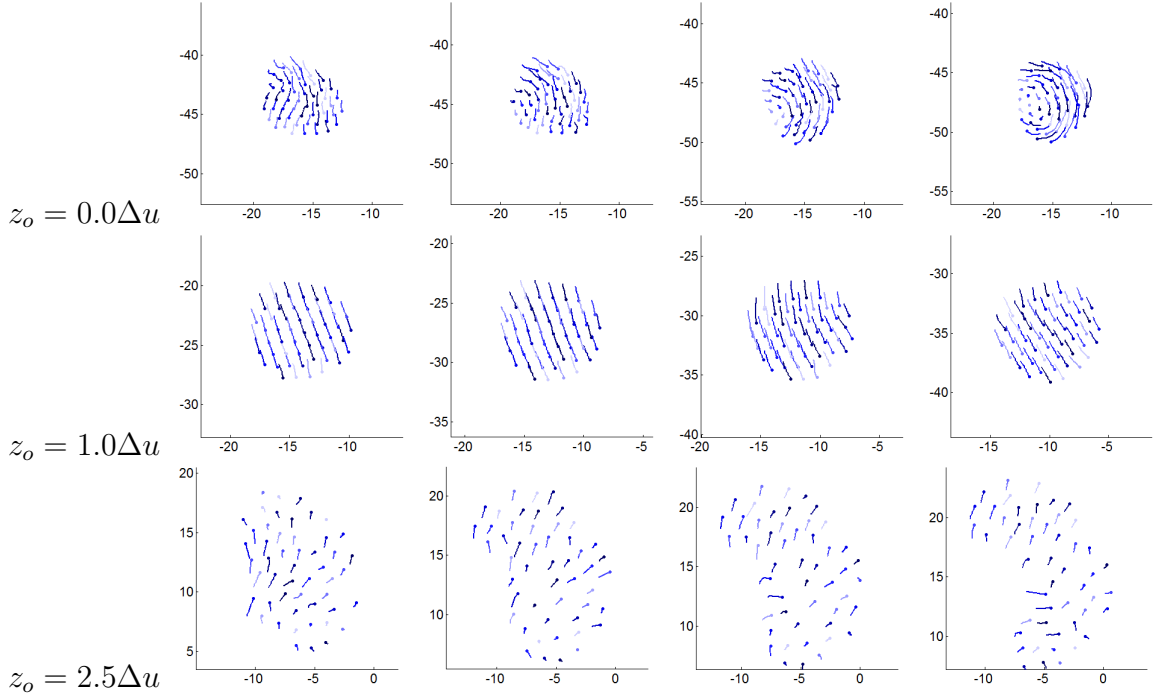


Figure 3-4: Different types of motion arise due to the rigidity of the lattice induced by the objective function. When z_o is small (top, $z_o = 0.0\Delta u$), the lattice is very tight and small perturbations in the relative position of neighboring particles can lead to large system-wide changes in motion, usually manifested as rotation. When z_o has an intermediate value (middle, $z_o = 1.0\Delta u$), the lattice spacing remains regular, but there is enough flexibility so that small differences in the heading of neighbor particles diffuse across the system without causing major disruptions, leading to a greater prevalence of purely translational motion. Because the rotational motion that occurs with the tighter lattice disrupts the local heading alignment. When z_o is large (bottom, $z_o = 2.5\Delta u$), the particles are loosely held together and groups of particles will contain a heterogeneous mix of headings.

that center very quickly. Distinct subgroups alternate as the stationary portion and outer portion, and the overall movement of the group is much faster than s^* due to a sling-shot effect. Simulations with these parameter settings often fragment in the middle of an alternation.

Stationary rotational mills, where the overall group remains stationary and the center of rotation is the center of the group usually occur in simulations with an intermediate value of z_o , and the group will transition between stationary rotation and




Figure 3.5: A group of particles undergoing simultaneous translational and rotational motion, where the group rolls like a ball, is shown. Red and blue are used to identify the two distinct subgroups that have emerged during this portion of this simulation. This behavior emerges when z_o and λ_s take on small values. The center of rotation is on one side of the group, not in the middle, and there is an outer shell that rotates very quickly around the group of slower-moving particles. Distinct subgroups alternate as the stationary portion and outer portion, and the overall movement of the group is much faster than the preferred speed, s^* , due to a sling-shot effect.

pure translation if the speed penalty is low. With a larger speed penalty, these types of rotational mills often occur early in simulations that later converge to stable translational motion. Finally, these formations also sometimes occur in simulations with a small z_o as a transitional phase between purely translational motion and combined translation and rotation.

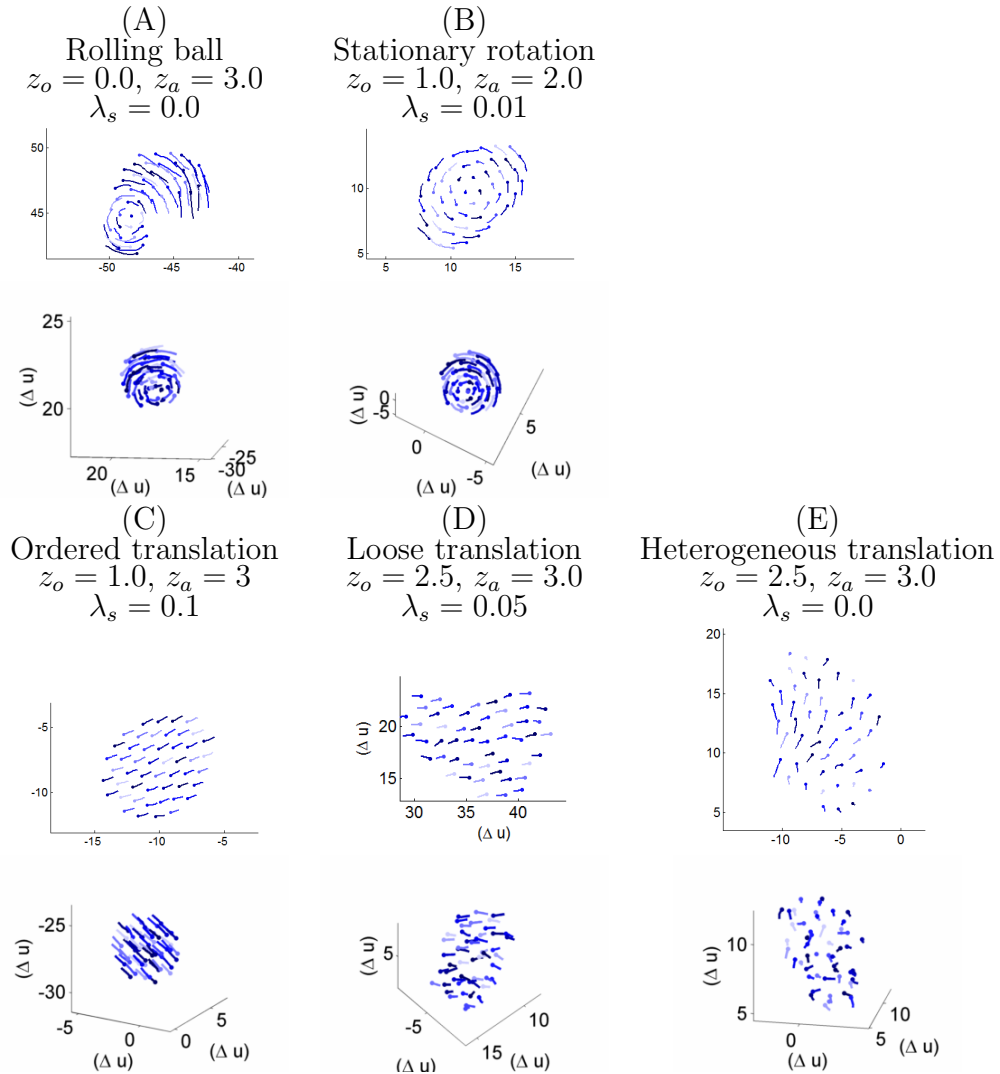


Figure 3-6: Examples of the diverse patterns of motion found in 2D and 3D simulations performed with our FlockOpt model resulted from five different settings of z_o , z_a , and λ_s . The dots indicate the current position of the particles, and the tails show the previous 10 time steps of motion. Formations A, B, and E are not possible if particles are modeled to maintain constant speed. A stronger speed penalty λ_s is associated with a greater prevalence of pure translational motion (C, D). When width z_o is larger (D), the spacing between particles becomes looser. With a large z_o and a low λ_s (E), the headings of the particles become also somewhat heterogeneous, even as the entire group slowly translates together. Stationary rotational mills (B) are more common when z_o is larger and z_a smaller, and are most often seen as a transitional rather than stable configuration.

3.5.1 Objective Values

In our simulation studies, we computed the objective value that each particle experienced at each time step. We found that the objective values in our simulations quickly saturated near the maximum value and remained nearly saturated throughout the entire simulation (Figure 3-7), even though our simulations exhibited very diverse patterns of motion (Figure 3-6). Polarization only emerged after the objective values had saturated. This is consistent with our theoretical result, which predicts that velocity matching only arises when no improvement in the objective value is possible. Our experimental results demonstrate that this is the most common case.

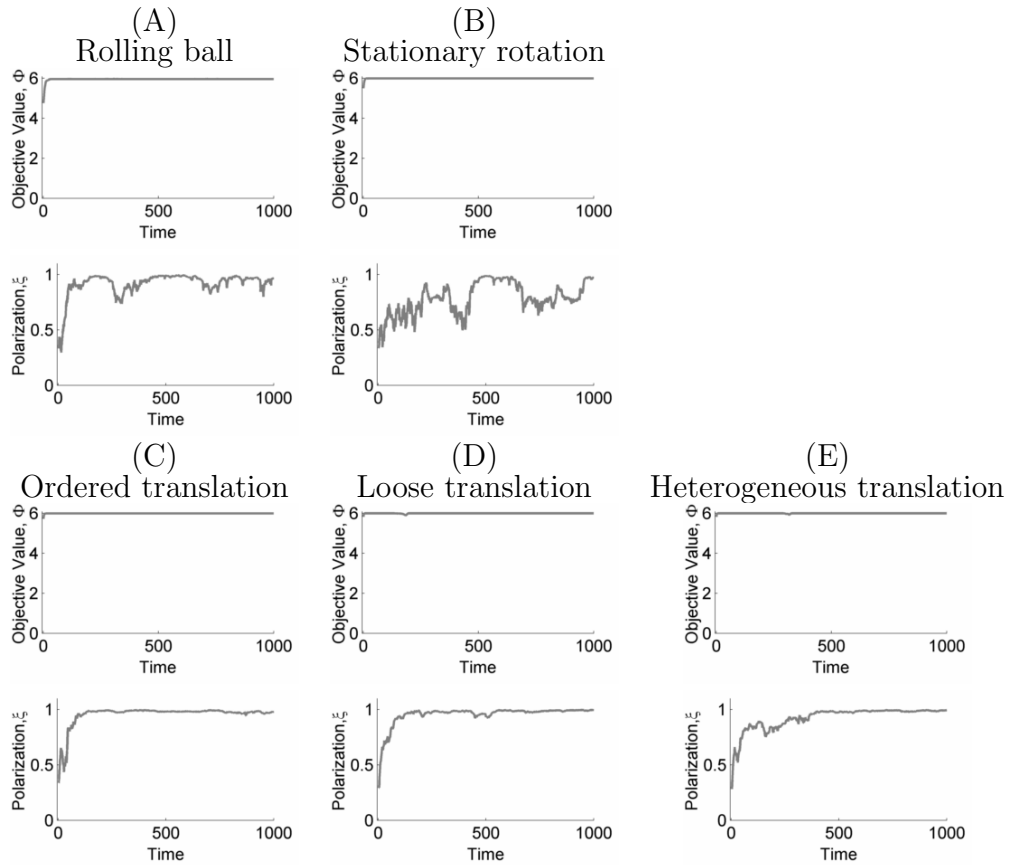


Figure 3.7: The average local polarization computed for each frame for one set of initial conditions from five different parameter settings are shown. In our simulations, the objective values typically saturated within 20 time steps, but it took longer for local heading alignment to become established. Even in simulations exhibiting periodic decreases in local polarization due to rotational motion (A and B), the objective values remained nearly saturated throughout the simulations.

3.6 Coupling among spatial variation, heading, and speed

3.6.1 Local Heading Alignment and Spatial Variation

One of the remarkable things about natural systems that has been reflected in models of collective motion is the way that entities move together. Previous scientists have focused on the extent to which the headings of entities in a system are aligned. To describe the local alignment of heading vectors, we use local polarization as defined by Mishra et al. (2012)

$$\xi_i = \frac{1}{|\mathbf{N}_i|} \left\| \sum_{k \in \mathbf{N}_i} \frac{\mathbf{v}_k}{\|\mathbf{v}_k\|} \right\|. \quad (3.16)$$

Due to our analysis demonstrating that velocity matching follows from maintaining relative position, we suggest also considering how relative position changes over time. We define the scalar quantity *spatial variation*, which is computed over a small time window τ

$$\psi_i = \sqrt{\frac{1}{|\mathbf{N}_i|} \sum_{k \in \mathbf{N}_i} \|\Delta \mathbf{x}_k - \Delta \mathbf{x}_i\|^2}, \quad (3.17)$$

where $\Delta \mathbf{x}_i = (\mathbf{x}_i[t + \tau] - \mathbf{x}_i[t])$.

For our simulations, local polarization, ξ_i , and spatial variation, ψ_i , were computed over a neighborhood of radius $2.25 \Delta u$, approximately twice the average distance between nearest neighbors across all of our simulations. Spatial variation was computed over a window of $\tau = 1 \Delta t$. For our analysis of observations of bats, local polarization and spatial variation were computed over a neighborhood with a radius of 2 meters, twice the average distance between first nearest neighbors over all of our data. Spatial variation was computed over a time window of 100 ms, approximately one wingbeat (Norberg and Norberg, 2012).

Under our model, we are able to simulate motion patterns with a variety of local polarization signatures. Even though the bats seem to fly in organized, highly coordinated formations, the headings of the bats and the distances between neighbors are

actually somewhat heterogeneous, and are similar to systems from our simulations exhibiting looser formations (Figure 3·6, D and E). In our simulation study, we find that systems with higher polarization generally exhibit less spatial variation, although this effect is less strong when the variance of speed is larger, compared with the mean (Figure 3·6 A – E). Similar trends are present in our observations of bats (Figure 3·6 F–J). Emergence bursts from late evening exhibit a somewhat different signature than others (Figure 3·6, H and J).

For any particular set of headings, speed differences can either lead to increased spatial variation, if the speeds are random, or decreased spatial variation, as in a group of particles that are rotating. In our simulation model, we know that the speed and heading of particles are coupled through a dependence on maintaining relative position. In other variable-speed simulation models, heading and speed are decoupled because they are determined by separate procedures.

To understand which mechanism might be at work in the behavior of *T. brasiliensis*, we developed a non-parametric, stochastic procedure to forcibly decouple the speed from the heading by permuting the speeds among bats present at each moment in time while maintaining their headings. This strategy obviated the need to choose a parametric distribution for the speed, or to consider how the parameters of the speed distribution changed over time. We measured spatial variation in the data as observed, ψ , and under this permutation, $\tilde{\psi}$. For every entity in the system, we replaced $\Delta\mathbf{x}_i$ with $\Delta\tilde{\mathbf{x}}_i = \|\Delta\mathbf{x}_j\| \frac{\Delta\mathbf{x}_i}{\|\Delta\mathbf{x}_i\|}$ where j was chosen with uniform probability, without replacement, from any of the entities present at time t . Then,

$$\tilde{\psi}_i = \sqrt{\frac{1}{|\mathbf{N}_i|} \sum_{k \in \mathbf{N}_i} \|\Delta\tilde{\mathbf{x}}_k - \Delta\tilde{\mathbf{x}}_i\|^2}. \quad (3.18)$$

If the speed and heading of individual particles were separable, then the spatial variation under both conditions would be similar. To demonstrate this, we gener-

ated sets of random particles where headings and speeds were drawn independently from normal distributions $\theta_i \sim \mathcal{N}(0, \sigma_\theta)$ and $s_i \sim \mathcal{N}(0, \sigma_s)$. The displacement was computed as $\Delta \mathbf{x}_i = s_i [\cos(\theta_i), \sin(\theta_i)]$. We used a variety of values for σ_θ and σ_s . Then, we computed spatial variation as initially generated, ψ_i , and after permuting the speeds, $\tilde{\psi}_i$. As anticipated, across a variety of parameter values for the variability of heading and speed, the distributions of ψ_i and $\tilde{\psi}_i$ were similar ($p \gg 0.05$, measured by Kolmogorov-Smirnov (KS) tests).

In data from our simulations, where we know that headings and speeds are coupled, distributions of ψ and $\tilde{\psi}$ were not similar across all parameter sets ($p \ll 0.01$), and the mean of $\tilde{\psi}_i$ was over 40% larger than ψ_i for some patterns of motion (Figure 3-9 A). In simulations with average local polarization and speed variability (coefficient of variation, σ_s/μ_s) similar to the bats, the increase was between 3.8% and 8.6% (e.g. Figure 3-9, B).

In natural data of bats in flight, the distributions of $\tilde{\psi}_i$ and ψ_i were not similar ($p \ll 0.01$ for all video segments). The increase in the average spatial variation was between 3.4% and 6.0% (Figure 3-9). Based on these results, we conclude that individual speed and heading are coupled in the motion of groups of bats.




Figure 3-8: Joint distributions of local polarization and spatial variation are shown for five simulation parameter sets (top row) and five video segments of bat flight (bottom row). In both simulated data and natural data, spatial variation decreases as polarization increases. When the speeds of particles in simulation is more variable (A and E), the distribution spatial variance conditioned on the polarization is wider. In recordings of bats taken later in the evening (H and J), there is larger spatial variation.

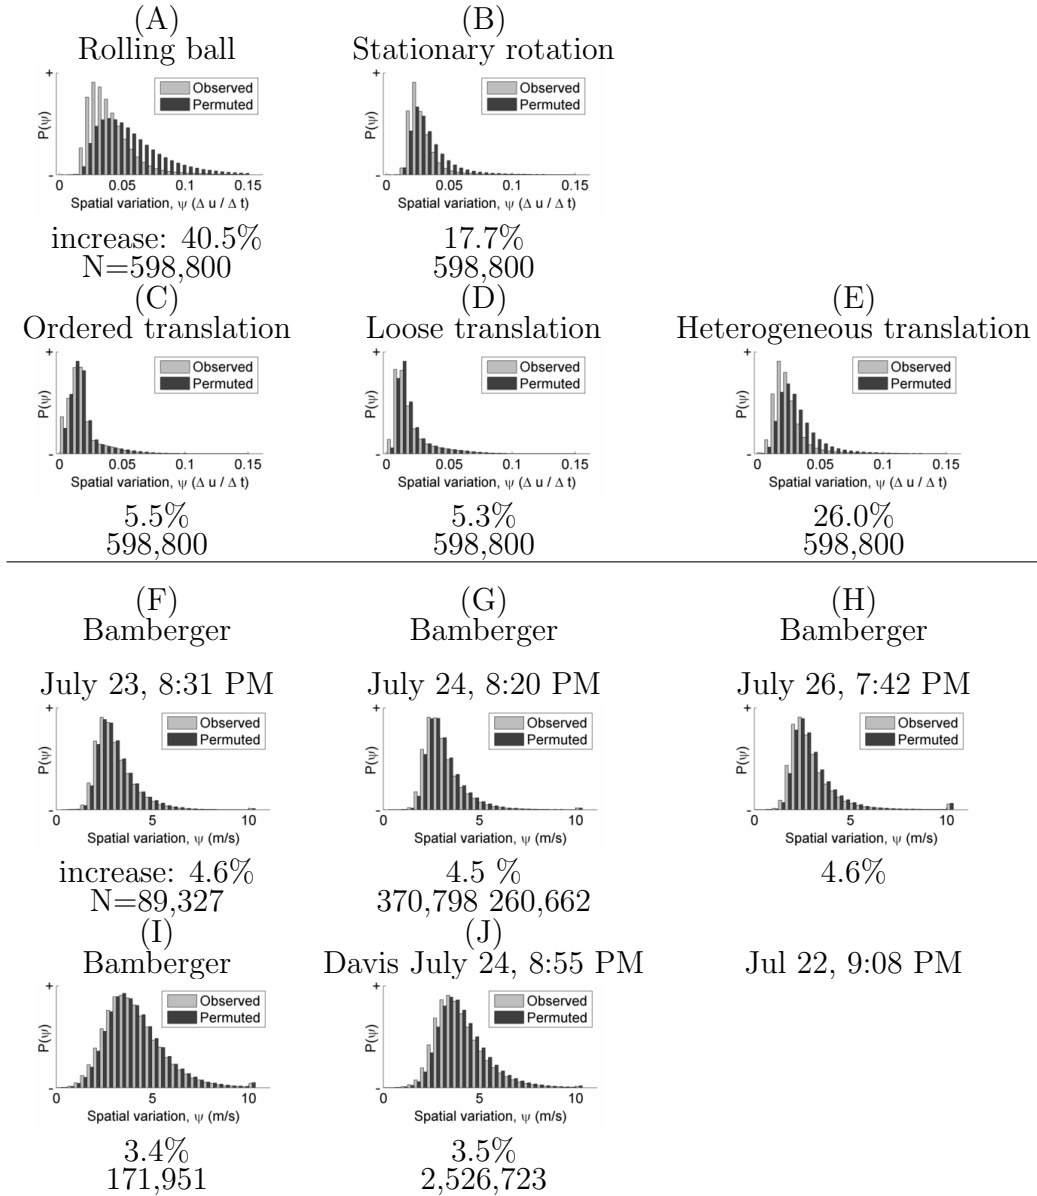


Figure 3-9: The distributions of spatial variation, as observed (ψ , light gray) and after speed and heading had been forcibly decoupled ($\tilde{\psi}$, dark gray), are shown for five FlockOpt simulation parameter sets leading to coherent translation (top), and for five video segments of the flight of *T. brasiliensis* (bottom). In all cases, the distributions $P(\psi)$ and $P(\tilde{\psi})$ were not similar, according to Kolmogorov Smirnov tests ($p \ll 0.01$), and the average spatial variation was larger after speed had been decoupled from heading. In some simulations (A and E), the increase in spatial variation was very large. The increases $\Delta\psi$ observed when decoupling the speed and heading of the bats (3.5%–4.6%) were slightly smaller than the increases observed in simulations with similar average local polarization and variability in speed (3.8%–8.6%).

3.7 Comparison with earlier models

We were inspired by zonal models (Couzin et al., 2002), and our optimization-based methodology has conceptual similarities to potential field models (Reif and Wang, 1999).

A critical distinguishing feature of our algorithm is choosing the course of action for each particle by optimizing the objective function with respect to the anticipated *future* relative position, $\hat{\mathbf{x}}_{i,k}$. With a large step size, this is crucial for convergence of the objective values and the emergence of ordered motion (s^* and s_{\max} were set to $0.1 \Delta u / \Delta t$ with $z_r = 1 \Delta u$ for the experiments described earlier in this chapter). To illustrate this, we carried out a simulation study without the predictive aspect. We chose two parameter sets using the same neighbor preference parameters ($z_o = 1.0 \Delta u$ and $z_a = 3.0 \Delta u$) and two speed penalties ($\lambda_s = 0.0$ and $\lambda_s = 0.1$). Under the fully predictive model, when $\lambda_s = 0.0$, the simulations tended to exhibit a mixture of rotational and translational motion, whereas when $\lambda_s = 0.1$, the motion was purely translational, following the initial stabilization period.

We varied the effective step size over five orders of magnitude by reducing the maximum speed s_{\max} and preferred speed s^* . Simulations were run for 5,000 steps by calculating $\Delta \mathbf{v}_i^*$ to optimize $\sum_{k \in \mathbf{N}_i} \Phi(\mathbf{x}_{i,k} - \Delta \tilde{\mathbf{v}} \Delta t)$ instead of $\sum_{k \in \mathbf{N}_i} \Phi(\hat{\mathbf{x}}_{i,k}) = \sum_{k \in \mathbf{N}_i} (\mathbf{x}_{i,k} + \mathbf{v}_{i,k} \Delta t - \Delta \tilde{\mathbf{v}} \Delta t)$. We found that without a speed penalty, with large step sizes, such as $s_{\max} = 0.1 \Delta u$, the objective values failed to stabilize (Figure 3-10, A), which precluded the emergence of ordered motion. With small step sizes, the objective values did eventually converge, but the terminal values were lower than with the fully predictive model, and the polarization of particles within the simulation remained low (Figure 3-10, B). With the damping effect of a penalty on deviation from a preferred speed, the objective values converged slowly for all step sizes (Figure 3-10, C), although the steady state objective values were slightly lower than with the fully

predictive model. Even though the objective values converged and this parameter setting was found to lead to purely translational motion in the fully predictive case, ordered translation did not stabilize in the first 5,000 steps of the simulation, versus less than 200 steps with a fully predictive model with a large step size (Figure 3·10, D).

In potential field models, the potential experienced by each particle is optimized by performing gradient descent of the potential field, which leads to averaging the responses to each neighbor individually in order to determine the final course of action. We made the important design decision to use hard constraints instead of very large potentials to express undesirable behavior, and so we used constrained optimization, implemented by sequential quadratic programming, which is somewhat less susceptible to local maxima than gradient ascent.

Our choice has led us away from averaging the responses to each neighbor individually and towards optimizing the joint objective function holistically. This design decision provides an alternative to the averaging approach used by most other models. To explore the importance of this decision, we carried out a simulation study with a hybrid model where the objective function was optimized for each neighbor individually to obtain $\Delta \mathbf{v}_{i,k}^* = \arg \max \Phi(\hat{\mathbf{x}}_{i,k}; \theta)$, then, to determine the final course of action, the responses to each neighbor individually were averaged as $\Delta \mathbf{v}_i^* = \sum_{k \in \mathbf{N}_i} \Delta \mathbf{v}_{i,k}^*$. Even though simulations with this hybrid model developed heading alignment, we found that the objective values converged more slowly and the terminal values were lower with the averaging approach (Figure 3·11). This was manifested in simulations by less ordered spatial arrangements of particles than with our original model, given identical parameters and initial conditions.

As we have demonstrated, speed and heading are coupled by a dependency of relative position. Our model controls variations in speed in two ways: by using

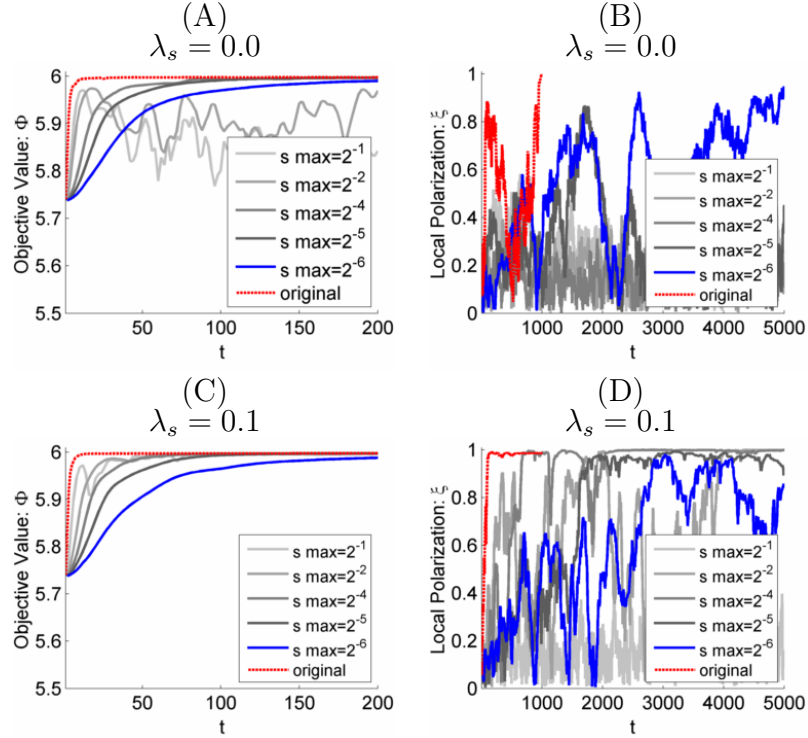


Figure 3-10: We carried out a simulation study to examine the effect of the prediction step in our algorithm. For two parameter sets, with $z_o = 1.0\Delta u$ and $z_a = 3.0\Delta u$, and $\lambda_s = 0.0\Delta u$, (top) and $\lambda_s = 0.1$ (bottom) we show the evolution of the average objective values (left) and the local polarization (right) experienced by all particles in simulations carried out with one set of initial conditions. We varied the effective step size by reducing the maximum speed s_{\max} (all panels) and preferred speed s^* (C and D). Under the fully predictive FlockOpt model (red lines) the objective values quickly converge in both circumstances. In simulations without prediction, the objective values are highest in simulations with the smallest step sizes (blue lines). The objective values from simulations without a speed penalty (A) do not converge when the step size is large (light gray lines); when the step size is small, the objective values converge but it takes many more steps (dark gray and blue lines). Even when the objective values converge, the emergence of locally aligned motion takes longer (B and D).

penalty terms in the objective function and by defining a constrained optimization.

We have found that relying on penalties produces more Gaussian-like distributions of speed, and that constraints can lead to edge effects at the boundaries.




Figure 3.11: We investigated the impact of computing the responses to each neighbor individually and averaging the results, as in many other models of collective motion. For three different parameter sets known to lead to three different types of motion patterns under a pure FlockOpt model, we show, for the pure model (solid lines) and hybrid model (dashed lines), the average objective values (top) and the average local polarization, ξ (bottom) experienced by all particles in one simulation with one set of initial conditions. In the hybrid model (dashed lines), the objective values converged more slowly to lower terminal values (top row) than in a pure FlockOpt model (solid lines). Despite the lower objective values, interesting patterns of ordered motion emerged that were different from the patterns of motion of the pure FlockOpt model under identical parameters and initial conditions (bottom row).

3.8 Strategies for Model Selection and Parameter Fitting

In order to understand the behavior and evolution of animals through computational modeling based on the FlockOpt paradigm, it is important to be able to choose a parametric form for the objective function and constraints, and to estimate the associated parameters that best predicts the behavior of real animals. In real animals, the objective function might be computed over variables we are unable to observe, such as neighbor visibility or aerodynamic considerations. We had mixed success achieving a match between a simulation model and natural data of *T. brasiliensis* (Section 3.8.5). Here, we discuss parameter fitting approaches that have worked in

our simulation studies (Section 3.8.2 and Section 3.8.3), which may point the way towards future techniques for performing model selection and parameter estimation based on natural data, using the FlockOpt paradigm.

3.8.1 Related Work

The fit between a model and data of animals may be considered in different ways. One point of view is to evaluate simulations performed with a model based on emergent characteristics, such as group cohesion and the type of motion (disordered, rotation, translation), as compared to the emergent characteristics of the natural system. We observe that in natural systems, such as large groups of bats, it is normal for pieces of the group of emerging bats to break off, so maintaining perfect group cohesion is not necessarily a prerequisite for a model that may reflect the natural behavior of animals. A qualitative match between the behavior of particles in simulations and animals in a group is necessary but not sufficient. It has been shown (Gautrais et al., 2012) that many different types of models and parameter settings can lead to emergent behavior that is similar.

A second, more detailed point of view, is to examine quantitative, aggregate properties of the simulated and natural systems, as done by Ballerini et al. (2008), Katz et al. (2011), Lukeman et al. (2010) and Herbert-Read et al. (2011), although Mann et al. (2013) observed that many different models and parameters may lead to similar statistics.

As a third point of view, some authors (Eriksson et al., 2010; Mann, 2011; Mann et al., 2013; Bode et al., 2011) have suggested that computational models should also be able to accurately predict the detailed behavior of individual animals. Eriksson et al. (2010) proposed using a “force matching” approach where the behavior of each animal at each moment in time is predicted by a behavioral model, and then the

acceleration actually experienced by each animal is compared to the acceleration that would be required to accomplish the course of action prescribed by the model. Using this approach, they were able to correctly discern between two models using different proximity definitions, even though the simulated study systems had similar aggregate statistics. Mann et al. (2013) suggested that in order to determine the consistency of a model with data, it is necessary for the model to match in different ways simultaneously.

3.8.2 Parameter estimation via histogram matching

We evaluated whether or not it was possible to infer simulation parameters by comparing distributions between simulations of first nearest neighbor distances, speed, and local polarization. Our experiment only concerned parameter estimation, not model selection, and so we used the parametric form of the objective function used to generate the simulations. Using the simulation data from our study in Section 3.5, we computed distributions of first nearest neighbor distances, speed, and local polarization. We computed the similarity between each simulation and all other simulations by computing the Earth-mover’s distance (EMD, a measure of similarity between histograms) between the histograms of each quantity.

We found that by comparing the distributions of the distances between first nearest neighbors, it is generally possible to determine z_o accurately (Figure 3·12, column A). This is true even when the assumed speed penalty is incorrect (Figure 3·12, column B), given the reasonable assumption that z_a is larger than z_o (Figure 3·12, column B, bottom). Without this assumption, incorrect neighbor preference parameters may lead to a better match than the correct parameters (Figure 3·12, B, bottom, $z_o = 2.0$ is correct, but the best match is $z_o = 1.5$, with an incorrect speed penalty).

If the presumed speed penalty is incorrect, comparing histograms of speeds will

not lead to the true neighborhood preference parameters, z_o and z_a (Figure 3-12, C, top, $z_o = 1.0$ smallest distance, but $z_o = 0.5$ is the correct value). We found that local polarization was not a useful feature for matching because a wide variety of parameter sets that lead to coherent translation all lead to very similar distributions, so the distances among the distributions are very small (Figure 3-12, D).

Once the parameters of the neighbor preference function, z_o and z_a , have been matched, the distribution of nearest neighbor distances is surprisingly discriminative with respect to the speed penalty (Figure 3-13, left), however, there may be multiple local minima (Figure 3-13 left, bottom). We found that some neighbor preference parameters will encourage rotational structures, while others will not, even when λ_s is the same. The distributions of speeds from simulations with rotational structures tend to have larger variance, therefore, the distribution of speeds is not necessarily independent of the neighborhood preferences, so a reasonably good estimate of the neighborhood parameters is necessary before attempting to match the speed penalty. Given correct neighbor preference parameters, inferring the speed penalty from a distribution of the speeds is straightforward (Figure 3-13, right).

Based on our findings, we suggest that parameter estimation via histogram matching might proceed effectively using an iterative approach, alternating between inferring the neighborhood preference parameters and the kinematic parameters. We suggest that evaluating the match between a study system and a bank of simulated systems based on distributions of speed or neighbor distances will help to restrict the parameter space to search for the best fit parameters in natural data, and possibly provide multiple starting points for optimizing parameters via a detailed fit, as described below.

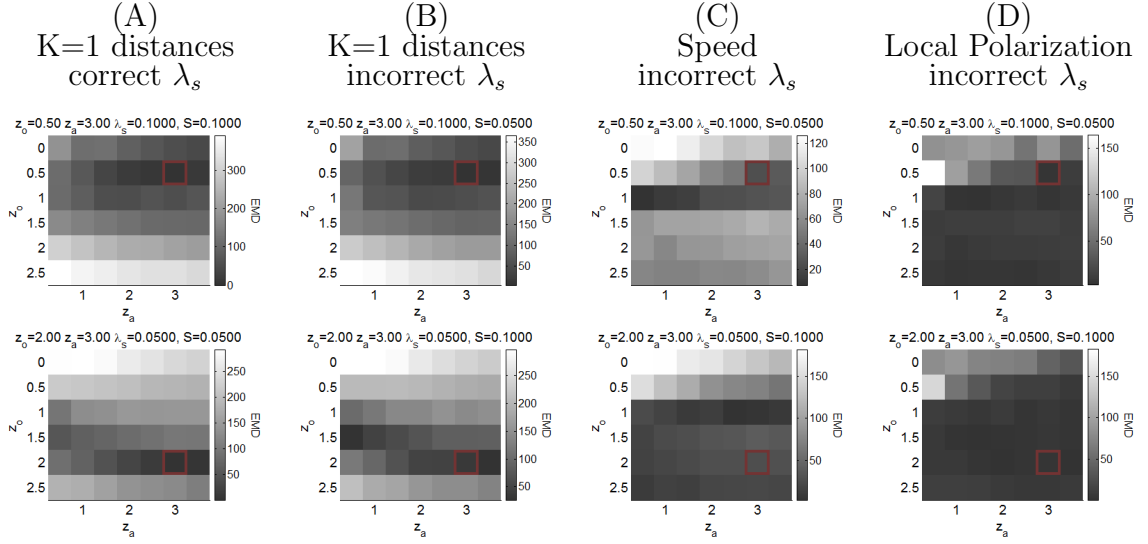


Figure 3-12: For simulations performed with two sets of parameters, we show the goodness of fit between a subset of other simulations based on the Earth-mover's distance between histograms of distances between nearest neighbors (columns A, B) speed (C) and local polarization (D). Low values (dark) are desirable and high values (bright) are undesirable. In each grid, each box depicts the Earth-mover's distance between the histograms derived from simulations with true parameters compared with simulations with other parameters. Values of z_o vary along the rows and values of z_a vary along the columns, so that each grid is able to display the goodness of fit between simulations with the true parameters and all other parameters with a given speed penalty. The true values of z_o and z_a are indicated with a dark red box. In the top row, the true parameters are $z_o = 0.5\Delta u$, $z_a = 3.0\Delta u$ and $\lambda_s = 0.1$. In the bottom row, the true parameters are $z_o = 2.0\Delta u$, $z_a = 3.0\Delta u$ and $\lambda_s = 0.05$. In column A, the histograms of nearest neighbor distances are compared with other simulations with a speed penalty matching the true value. In column B, histograms of nearest neighbor distances are compared with simulations with a different speed penalty. In columns C and D, the histograms of speed and local polarization respectively are compared with other simulations using a speed penalty that is different from the true value. When the speed penalty matches, minimizing the Earth-mover's distance between histograms of nearest neighbor distances will lead to the correct values for z_o and z_a (column A). Even if the assumed speed penalty is incorrect (column B), it is generally possible to correctly match the value of z_o , though, it may be necessary to make reasonable assumptions about z_a (B, bottom row) in order to correctly infer z_o . If the neighborhood parameters are incorrect, then matching histograms of speed will not lead to the correct values (column C). Local polarization histograms are not a discriminative feature for inferring model parameters, since many simulations exhibiting coherent translation lead to very similar histograms, so the distance between histograms from parameters with different parameters is small (D).




Figure 3-13: The Earth-mover’s distance between the distributions of nearest neighbor distances (left) and speed (right) derived from simulations can be used to infer the parameters of the speed penalty λ_s parameter. In the top row, the true parameters are $z_o = 0.5\Delta u$, $z_a = 3.0\Delta u$ and $\lambda_s = 0.1$. In the bottom row, the true parameters are $z_o = 2.0\Delta u$, $z_a = 3.0\Delta u$ and $\lambda_s = 0.05$. Using the true values of z_o and z_a , the Earth-mover’s distance between the distributions of nearest neighbor distances (left) and speed (right) are plotted for other simulations with the correct neighborhood parameters, but different speed penalty values. If the neighborhood preference parameters are correct, then the nearest neighbor distance histogram is informative with respect to the correct speed penalty (left column), though there may be multiple local minima (bottom, left). If the neighborhood parameters are correct, then inferring the correct speed penalty by matching histograms of speeds is straightforward (right column).

3.8.3 Parameter estimation and model selection via model prediction

Instead of estimating parameters by comparing histograms between simulations, a detailed fit may be performed by using a computational model and associated parameters to predict the motion of individual animals. We found that the force-matching approach espoused by Eriksson et al. (2010) contained an unnecessary level of indirection that was sensitive to the exact details of how the acceleration was computed, and

so we adopted a *position matching* approach. Given a (simulated or natural) study system and a behavioral model with associated parameters to compare, we used the current state of the system in order to predict a change in velocity $\Delta \mathbf{v}_i^*$ for each entity. Then, the predicted position of each particle was computed as $\mathbf{x}_i + \mathbf{v}_i + \Delta \mathbf{v}_i^*$. This prediction was compared with the actual position of the entity in the study system at the next time step. The goodness-of-fit measure is the RMSE of the differences between the predicted and observed positions. This approach can also be used to evaluate the goodness of fit with other types of models, such as a constant velocity model (Section 3.8.4). Performing the required predictions in order to evaluate this score for one model and one set of parameters requires a similar amount of computation as that required to run a simulation. In the case of a FlockOpt model, the behavioral model consists of the objective function and constraints.

To better understand the way the position matching goodness-of-fit measure behaves, as study systems, we selected 12 sets of simulated systems from our simulation study that led to archetypal patterns of motion, using a small selection of parameter values. Holding $z_a = 3\Delta u$, we varied z_o between 0 and $2.5 \Delta u$. Holding $z_o = 1\Delta u$, we varied z_a between 1 and $3 \Delta u$. For two sets of z_o and z_a , we varied λ_s between 0.0 and 0.1. For each of these simulation sets, we evaluated the goodness of fit between the original simulation data and the same model with different parameters using the position matching approach.

Due to the computational expense, we restricted the explored range of parameters for comparison to the parameter space immediately surrounding the true parameters that generated the simulation. We found that, similar to the histogram matching approach, an incorrect value of z_o has the largest impact in the prediction error, often on the same order as the maximum allowable magnitude of the change in velocity, v_{\max} . A value that is too small is worse than a value that is too large, especially when

the true value of z_o is large. An incorrect value of z_a is less detrimental, typically on the order of the value of the preferred speed, s^* . The speed preference penalty λ_s has the least impact of the three parameters, but the impact is larger when the spacing between particles is looser due to larger values of z_o .

With this method for computing a goodness-of-fit score, parameters of a model can be estimated by repeatedly comparing the study system with a behavioral model using different parameters. We suggest that for this approach to work well and be computationally feasible, the initialization needs to be near the correct starting point, with bounds to restrict the search space to some reasonable values. Because our model is not linear, we are unable to use simple linear regression to infer the parameters, and so we must use an iterative optimization algorithm. One possibility is to use gradient ascent. We do not know if the error surface over the parameter space is convex, or if there are many local minima, so if it is not possible to restrict the search space, we favor an algorithm like simulated annealing. We have found that sometimes evaluating a set of parameters with more than one incorrect value can sometimes match better than a parameter set with a single incorrect value, due to the existence of critical points in the parameter space, where the behavior of simulated systems changes dramatically with a small change in parameters, so we suggest using multiple starting points with any optimization algorithm.

In this chapter, we have suggested an objective function based on the difference of two sigmoids, with appropriate constants, but other forms are possible e.g. a Gaussian. If the assumed parametric form of the objective function does not match the study system, then estimating the best fit parameters may not lead to useful insights about animal behavior, and so it may be necessary to fit parameters of multiple models in order to gain insight into animal behavior.

3.8.4 Generic models

Since the double-sigmoid neighbor preference function is very specific, we tried to develop a small set of more generic models that would enable us to learn basic things about a system without trying to guess the parametric form of the objective function or constraints. There were four generic models. 1.) a constant-velocity model, where particles simply continued moving with their current velocity. 2.) a velocity averaging model, where particles adopted a velocity equal to the average velocity of their neighbors. 3.) a position-maintenance model, where particles tried to minimize the change in relative position between themselves and their neighbors, formulated in the FlockOpt framework by defining the objective function $\mathcal{O} = -\sum_{k \in \mathbf{N}_i} (\|\widehat{\mathbf{x}}_{i,k} - \mathbf{x}_{i,k}\|^2)$, and 4.) a distance-maintenance model where particles tried to minimize the changes experienced in the distances between themselves and their neighbors, also formulated in the FlockOpt framework by defining the objective function $\mathcal{O} = -\sum_{k \in \mathbf{N}_i} (\|\widehat{\mathbf{x}}_{i,k}\| - \|\mathbf{x}_{i,k}\|)^2$.

We found that when z_o was fairly small and the spacing between particles in simulations was very tight, then a constant velocity model performed the worst of all of the generic models; the RMS position matching error (as described in Section 3.8.3) was as much as twice as large as the error associated with the position-maintenance or distance-maintenance models. In simulations with rotational motion, the predictions of the velocity averaging model were similar to the constant velocity model, and much less accurate than the distance-maintenance and position-maintenance models. In simulations with stronger tendencies towards translational motion, the position matching error of the velocity averaging model was more similar to the position and distance maintenance models than the constant velocity model. For values of z_o and z_a that led to loose spacing between particles, the RMS position matching error of the constant velocity model was substantially *lower* than the error of any of the other

generic models, which were three times larger.

When we evaluated the goodness of fit between simulated study systems and the generic models described in this section, we used the neighborhood definition that was used to generate the simulation. For example, simulations from our simulation study performed with the FlockOpt model were carried out in 2D using the six nearest neighbors as the neighborhood definition, so when we compared those study systems with the velocity averaging model, we averaged the velocities of the first six nearest neighbors. For the position-maintenance and position-maintenance models formulated in the FlockOpt framework, we also used the same kinematic constraints (restricting $\|\Delta\tilde{\mathbf{v}}\|$ to $0.4 \Delta u/\Delta t$), so we do not know if any of the latter three models would be useful for inferring the correct neighborhood definition or kinematic constraints, in the absence of better knowledge about the true model and associated parameters.

3.8.5 Matching simulation parameters with natural data

To fit a behavioral model to natural data, several aspects of the system must be considered. The first thing to consider is the prediction interval. Mann (2011) observed that when simulation data were interpolated or decimated, recovering the true parameters of the simulation became more difficult. It is not known whether to use an interval based on the reaction time of an animal under other conditions, or how to incorporate knowledge of an animal’s locomotion. We speculate that different types of interactions may happen at different time scales e.g. moving away from imminent collisions happens very quickly, but moving towards the group after drifting away may happen more slowly.

We tried fitting models to bats by using prediction intervals of one or two wing beats (100 ms or 200 ms). When we tried to fit a FlockOpt model to the flight of

bats, using a double-sigmoid neighbor distance preference function, we found that the optimization typically led to very large values of z_o , indicating a large range of equally preferable distances. The practical implication of this is that the “best fit” models simply degenerated into constant velocity models.

Several factors may have contributed to our difficulties fitting FlockOpt model parameters to observations of bats. Most obviously, our assumed parametric form may be a mismatch with what the bats are actually doing. More subtly, our prediction interval may have been incorrect or our assumption that there is a single prediction interval may be inconsistent with animal behavior. The correct neighborhood definition is unknown. The bats may pay more attention to other bats that are directly in front of them, for example. It may also be that the column where we observed it, very close to the entrance of the cave, is not yet in a steady state and so our assumption that bats might try to maintain neighbor distances or relative positions may be invalid. Finally, occlusion may play an important role, since bats may be reacting to the actions of neighbors that are not observable with our image processing and tracking pipeline.

We also tried fitting all of the generic models described above to bat behavior, and we found that the constant velocity model was the best predictor. Of the other models, the prediction error of the position-maintenance model was more than 25% better than the velocity averaging model on one of our datasets. The position maintenance model was slightly more accurate than the distance maintenance model, suggesting that the neighbor position function for bats might be anisotropic.

3.8.6 Parameter estimation via objective values

From our simulation studies, we have come to understand that even when the behavior of entities in the context of a system is observable, that does not necessarily give

good insight into what the entities “want” (the circumstances that maximize their individual objective value) or what they would do outside of that context. There is a complex interaction among the desires of individual particles, which leads to group behavior that may be very different from desired individual behavior.

For example, in our simulations, we used objective functions specified over neighbor distances, with long flat parts in the middle, but the neighbor distributions that we observe in the simulations generated with those objective functions are much more concentrated than the input, having modes very close to the edge of the ascending leg of the function and small variances. Consequently, directly comparing the shape of Φ and the distribution of first nearest neighbors is not informative about the parameters that drove the simulation.

Another approach that is ineffective is to select model parameters to maximize the objective values achieved by all entities in the system. Because the distances between nearest neighbors are generally concentrated near the edge of the ascending leg of the Φ function, the tail of the distribution will contribute to higher objective values when z_o is large, and so objective values computed with larger values of z_o will almost always be larger than smaller values.

3.9 Discussion

We chose to use six neighbors in 2D and twelve neighbors in 3D because these are the numbers suggested by theory of sphere packing. In other experiments, we have found that in general, when the neighborhood definition is expanded to include more neighbors, generally, the spacing between particles decreases, whereas when the neighborhood definition is restricted to include fewer neighbors, groups of particles become more prone to fragmentation.

In our simulations, we have chosen not to use a blind angle in the definition of

proximity because real animals are able to move their heads in order to see things all around them. Furthermore, even if neighboring entities cannot be seen, it may be possible to hear them, and so the range of an animal’s perception may be larger than that suggested by their field of view. Varied interactions between particles due to ease of perception in different relative locations could be incorporated either by reintroducing a blind angle into the neighborhood definition, thus removing imperceptible neighbors from consideration. An alternative approach would be to define an objective function that weighted the influence of neighbors in different relative locations, for example by constructing an objective function such as $\mathcal{O} = \sum_{k \in \mathbf{N}_i} \gamma_{i,k} \Phi(\|\hat{\mathbf{x}}_{i,k}\|)$ where the weights, $\gamma_{i,k}$ are calculated based on the viewing direction between the two particles.

When a particle is not in a stable state, the quadratic optimization may consider multiple equally good solutions. The solution that is ultimately chosen will depend on the implementation of the algorithm that is used. Any criteria that makes one change in velocity more favorable than another should be encoded with appropriate penalty terms in the optimization. We used a penalty against the deviation of the particle’s movement speed from a preferred speed. Another penalty could be added to penalize the magnitude of the change in velocity, which would favor smaller changes. We considered but did not use such a penalty because it did not alter the overall patterns of motion and led to increased group fragmentation when particles were initialized randomly, even with small values of λ . If particles were initialized in a different way that was closer to a configuration where the objective values were saturated, a penalty against $\|\Delta\tilde{\mathbf{v}}\|$ might not have as much of an adverse effect on group cohesion.

The possibilities for implementing variations of our model are endless. In addition to objective function terms or constraints based on kinematics of individual particles or their neighbors, terms could be incorporated to reflect interactions in heterogeneous

systems, goals, walls, and obstacles. To represent interactions in a heterogeneous system, the objective function would require terms with weights based on the type of the neighbor. Environmental obstacles, such as walls are probably best modeled as a combination of both constraints, to prevent particles from moving through the wall, and penalties, to encourage obstacle avoidance.

3.10 Conclusion

The FlockOpt model is rich paradigm for modeling collective motion. We have selected one important, previously unexplored way that the motion observed in our simulations is similar to the observed behavior of the bats. The purpose of this work is to introduce a new class of model for researchers to choose from as they work to understand animal behavior, and to illustrate the insights into the fundamental nature of collective motion that this model affords us. Simulations performed with our model reflected a great diversity of patterns of motion, even when only three parameters (z_o , z_a , and λ_s) were varied. Our model is promising because it addresses important shortcomings that have been discovered in existing models, such as regulating speed and combining the influences of neighbors.

Because our model is able to express variability in speed, simulations with our model exhibit patterns of motion not possible with earlier, constant-speed models. We presented evidence that our model captures a potentially important aspect of natural behavior, the coupling between speed and heading through a tendency to maintain relative positions, and our analysis of observations of bats confirms that a similar behavior may exist in the flight of gregarious bats. Our explanation leads to an important insight about the dominant modes of group movement in a variety of natural systems, such as translating flocks and rotational mills: Euclidean transformations that maintain relative position are composed of rotations and translations.

Rule-based models of collective motion became popular in the early 1980's and spawned three decades of research in diverse areas, including computer graphics, biology, physics, robotics, and civil engineering. Just as the classic zonal models gave rise to over 30 years of insights and countless variations that reflected different aspects of natural behavior, we expect that the FlockOpt paradigm will provide a new avenue for exploring collective motion which will be applicable in all of the domains where models of collective motion are currently employed.

Chapter 4

Analysis of behavior of bats

4.1 Introduction

We have studied the behavior of *Tadarida brasiliensis* in the emergence column, based on the hundreds of thousands of trajectories that we have reconstructed, containing millions of data points obtained by analyzing 70 minutes of video taken over four nights at two large maternity colonies. In particular, we have focused on characterizing the most frequently observed patterns of motion and confirming or refuting assumptions typically found in models of collective motion, such as attraction and repulsion.

We captured eleven video segments over four days at two colonies. The bats emerged in dense columns with multiple bursts of activity separated by long periods of quiet, consistent with the observations of Betke et al. (2008). Generally, each video segment captured a different emergence burst. Each video segment resulted in a dataset of reconstructed trajectories.

4.2 Bat kinematics

The object locations in each track were filtered with a Kalman filter to mitigate the effects of reconstruction uncertainty. We observed periodic motion with a dominant frequency of approximately 10 Hz. Since this frequency is consistent with the flapping frequency of *Tadarida brasiliensis*, we further applied a box filter of 100 ms to mitigate

this effect when computing higher moments of the tracks. A longer filter may have damped this periodic motion more effectively, but we needed to balance the length of our filter with the percentage of data lost due to inadequate support when computing acceleration. The average displacement between points in the original tracks and tracks smoothed for flapping was less than 1 cm. We used the smoothed tracks when computing kinematics and the original tracks when computing spatial relationships.

Velocity and acceleration were computed using a window length, τ , of approximately one wing-beat, using the central differences as follows:

$$\mathbf{v}_i [t] = \frac{1}{\tau} (\mathbf{x}_i \left[t + \frac{\tau + 1}{2} \right] - \mathbf{x}_i \left[t - \frac{\tau - 1}{2} \right]), \quad (4.1)$$

$$\mathbf{a}_i [t] = \frac{1}{\tau^2} (\mathbf{x}_i [t + \tau] + \mathbf{x}_i [t - \tau] - 2\mathbf{x}_i [t]). \quad (4.2)$$

In order for the computation of the acceleration vector to be valid (computed with appropriate support from the smoothed track) a data point at $t \pm 2\tau$ must be present in the track, so that the track positions used to compute the acceleration vector can be filtered appropriately. Data points where the acceleration could not be computed with the required support were discarded in our analysis. We used over 15 million valid data points from the Davis site and between 1 million and 2.5 million valid data points from each of the three nights at the Bamberger site.

4.2.1 Acceleration decomposition

We examined the motion of bats in a number of different ways. As described above, velocity, \mathbf{v}_i and acceleration, \mathbf{a}_i , were computed using a window length, τ , of approximately one wing-beat, using central differences (Equations 4.1 and 4.2). The heading of the bat is the velocity of the bat normalized to unit length, $\mathbf{h}_i [t] = \frac{\mathbf{v}_i}{\|\mathbf{v}_i\|}$.

The instantaneous angular change in velocity, or turning angle, was computed as

$$\theta_i = \frac{1}{\tau} \arccos \left(\frac{\mathbf{x}_i[t] - \mathbf{x}_i[t - \tau]}{\|\mathbf{x}_i[t] - \mathbf{x}_i[t - \tau]\|} \cdot \frac{\mathbf{x}_i[t + \tau] - \mathbf{x}_i[t]}{\|\mathbf{x}_i[t + \tau] - \mathbf{x}_i[t]\|} \right). \quad (4.3)$$

Note that this definition conflates changes in heading to the right and left with changes in heading up and down. We describe our remedy for this below.

The instantaneous speed change was computed as

$$\Delta s_i = \frac{1}{\tau} (\|\mathbf{x}_i[t + \tau] - \mathbf{x}_i[t]\| - \|\mathbf{x}_i[t] - \mathbf{x}_i[t - \tau]\|). \quad (4.4)$$

Note that the definitions of the turning angle and speed change given here are centered at time t , whereas computing the change between the velocity vectors $\mathbf{v}_i[t]$ and $\mathbf{v}_i[t - 1]$ would lead to a value that was delayed in time by $\frac{\tau}{2}$ since they are each computed using central differences. For notational convenience, we now drop the t subscripts and quantities such as \mathbf{v}_i will be assumed to occur at time t unless otherwise denoted.

As written here, the acceleration is a vector quantity, and so we decompose the acceleration into three components. This is in the same spirit as Katz et al. (2011), who performed a similar decomposition when studying the 2D motion of fish. We have adapted their thinking for our 3D analysis. We use a coordinate system defined where the X axis is parallel to the heading, the Y axis is a vector over the left wing, given by the cross product between the heading and the unit vector opposite of gravity, \mathbf{u}_g , and the Z axis is the cross product of the X and Y axes (similar to, but not necessarily parallel to \mathbf{u}_g , if \mathbf{u}_g and \mathbf{h}_i are not perpendicular).

The component of the acceleration parallel to the heading is given by

$$a_i^{(\parallel)} = \mathbf{a}_i \cdot \frac{\mathbf{v}_i}{\|\mathbf{v}_i\|} = \mathbf{a}_i \cdot \mathbf{h}_i. \quad (4.5)$$

Since the bats move in 3D, there are two components of the acceleration perpendicular to the heading. We separate these into a turning acceleration, $a_i^{(\perp)}$, corresponding to movement left and right, and a pitching acceleration, $a_i^{(\uparrow)}$, corresponding to movement up and down. Due to the right-handed coordinate system, the turning acceleration $a_i^{(\perp)}$ is positive to the *left*, and negative to the right. The axis $\hat{\mathbf{u}}_i^{(Y)}$ points roughly over the left wing, and $\hat{\mathbf{u}}_i^{(Z)}$ points roughly up. The two axes perpendicular to the heading are given by

$$\begin{aligned}\hat{\mathbf{u}}_i^{(Y)} &= \frac{\mathbf{u}_g \times \mathbf{h}_i}{\|\mathbf{u}_g \times \mathbf{h}_i\|} \\ \hat{\mathbf{u}}_i^{(Z)} &= \frac{\mathbf{h}_i \times \hat{\mathbf{u}}_i^{(Y)}}{\|\mathbf{h}_i \times \hat{\mathbf{u}}_i^{(Y)}\|}.\end{aligned}\tag{4.6}$$

The turning acceleration and pitching acceleration are the projections of the acceleration vector onto the respective axis, computed as

$$a_i^{(\perp)} = \mathbf{a}_i \cdot \hat{\mathbf{u}}_i^{(Y)},\tag{4.7}$$

and

$$a_i^{(\uparrow)} = \mathbf{a}_i \cdot \hat{\mathbf{u}}_i^{(Z)}.\tag{4.8}$$

Their units are m/s², just like the units of the full acceleration vector.

We chose to allow $\hat{\mathbf{u}}_i^{(Z)}$ to deviate from gravity so that the acceleration components would correspond to the deviations from the current heading. (If the heading of the bat is already pointed up, continuing in exactly that direction will yield $\hat{\mathbf{u}}_i^{(Z)} = 0$.) Katz et al. (2011) described $a_i^{(\parallel)}$ as a “speeding force,” but this nomenclature is slightly misleading, since this component of the acceleration may be negative if the bat turns while maintaining a constant speed. For this reason, we generally prefer to look directly at the speed change Δs_i instead of $a_i^{(\parallel)}$.

4.2.2 Analysis of bat kinematics

In order to better understand the way that bats move in the emergence column, we computed basic statistics of all kinematics quantities, such as mean, variance, and skew. To understand the relationships among kinematic quantities, we computed correlation coefficients and also examined joint distributions, which allowed us to better understand relationships that might not be captured well by simple correlation analysis.

In summary, we found that, across all of our datasets, the mean of the magnitude of the acceleration, $\|\mathbf{a}_i\|$ was between 8.8 and 12.0 m/s². A non-zero mean of the acceleration magnitude suggests that the mean of the acceleration vectors might not be centered at zero. We also found that the tendency to turn right or left, reflected by the mean of the turning acceleration $a_i^{(\perp)}$, seemed to vary based on the geometry of the cave entrance (more data would be helpful for establishing this trend). The tendency to pitch up or down, reflected by the mean of the pitching acceleration $a_i^{(\uparrow)}$ was different on different nights, even at the same site, so more data would be needed to establish a pattern.

By examining the joint distribution between total acceleration magnitude, $\|\mathbf{a}_i\|$ and turning angle θ_i (Equation 4.3, Figure 4-1, A), we found that changes in heading were a much more prominent aspect of the motion of bats (correlation ρ between 0.93 and 0.97) than changes in speed Δs_i (Equation 4.4) (correlation ρ less than 0.25 for nine out of eleven of our video clips) (Figure 4-1 shows the distributions for one example dataset).

Since turning angle θ_i represents total changes in heading and conflates turning left and right with pitching up and down, we further investigated the way that bats change their heading by examining the joint histograms between the acceleration magnitude $\|\mathbf{a}_i\|$ and turning acceleration $a_i^{(\perp)}$ (Equation 4.7), which describes turns

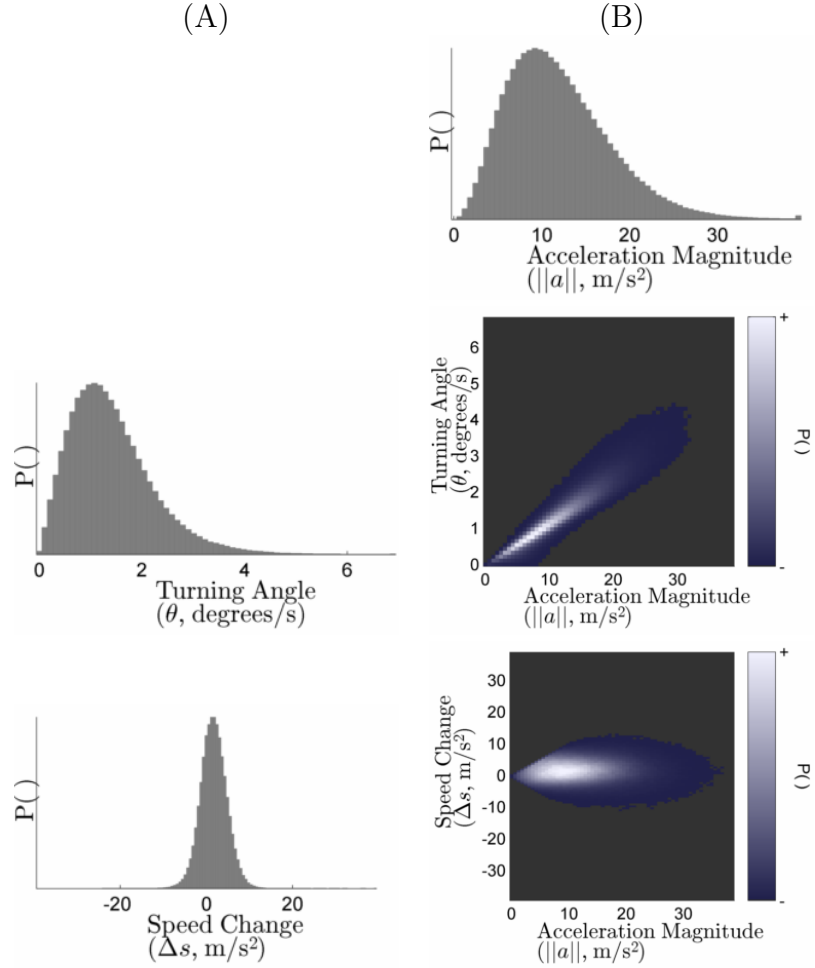


Figure 4-1: The joint distributions between acceleration magnitude, $\|\mathbf{a}_i\|$, and turning angle, θ_i (B, middle) and speed change, Δs_i (B, bottom) are shown from one of our eleven datasets, taken in the late evening. The correlation between $\|\mathbf{a}_i\|$ and θ_i is 0.94 ($p \ll 0.01$), and the correlation between $\|\mathbf{a}_i\|$ and speed change, Δs_i is 0.075 ($p \ll 0.01$), calculated using a total of 1,263,296 data points that had appropriate support that enabled us to compute acceleration based on smoothed trajectories. For reference, the marginal distributions of acceleration magnitude, $\|\mathbf{a}_i\|$ (B, top), turning angle, θ_i (A, middle), and speed change, Δs_i (A, bottom) are shown. This finding indicates that changes in heading are a more prominent aspect of bat motion than changes in speed.

left and right, and pitching acceleration $a_i^{(-)}$ (Equation 4.8) which describes motion up and down. The shapes of the joint distributions are like a fan (Figure 4.2) because large changes in heading can be accounted for by turns in up, down, left, or right, and the total acceleration magnitude restricts the possible range of values of both pitching and turning acceleration. Stated another way, both the turning acceleration and pitching acceleration are essentially lengths of the projections of the acceleration vector, \mathbf{a}_i onto axes defining a local coordinate frame of the bat, namely $\hat{\mathbf{u}}_i^{(Y)}$ and $\hat{\mathbf{u}}_i^{(Z)}$. The length of the projection of a vector cannot be larger than the length of the vector itself. Due to the shape of the joint distribution between acceleration magnitude and turning or pitching acceleration, the relationships are not amenable to correlation analysis. By examining the joint distributions among acceleration magnitude, pitching acceleration, and turning acceleration (Figure 4.2), we found that in this dataset, bats engage in turns to both the left and right, but pitching up is more likely than pitching down.

We also investigated the relationship between turning and pitching acceleration and found that there was a characteristic “lima-bean” shape to the joint distribution (Figure 4.3), across all of our datasets, suggesting that pitching up is less likely than pitching down to be associated with turning either left or right.

Finally, we investigated the relationship between heading changes and flight speed (Figure 4.4). We found that turning angle, θ_i (Equation 4.3), was negatively correlated with speed (correlation ρ between -0.13 and -0.33). From the joint distribution, we can see that this is especially true of large turning angles (Figure 4.4, A shows an example for one dataset). Investigating changes in heading more closely by examining speed as related to pitching acceleration, we found that pitching up was positively associated with faster speed (correlation ρ between 0.21 and 0.55) (Figure 4.4, B). When examining the speed as related to turning acceleration, we observe the “lima-




Figure 4.2: The joint distributions between acceleration magnitude, $\|\mathbf{a}_i\|$, and turning acceleration, $a_i^{(\perp)}$ (B, middle) and pitching acceleration, $a_i^{(\top)}$ (B, bottom) are shown from one of our eleven datasets (the same used in Figure 4.1, consisting of 1,263,296 data points). The shape of the joint distributions is like a fan because the magnitude of the acceleration vector constrains the maximum value of the corresponding turning or pitching acceleration. We found that the average value of the turning acceleration, $a_i^{(\perp)}$, which reflects a preference for turning right vs left, varied between our two study sites, which may be due to the geometry of the cave entry. The sign of the average value of the pitching acceleration varied on different nights, even at the same study site. For reference, the marginal distributions of acceleration magnitude, $\|\mathbf{a}_i\|$ (B, top), turning acceleration, $a_i^{(\perp)}$ (A, middle), and pitching acceleration, $a_i^{(\top)}$ (A, bottom), are shown.

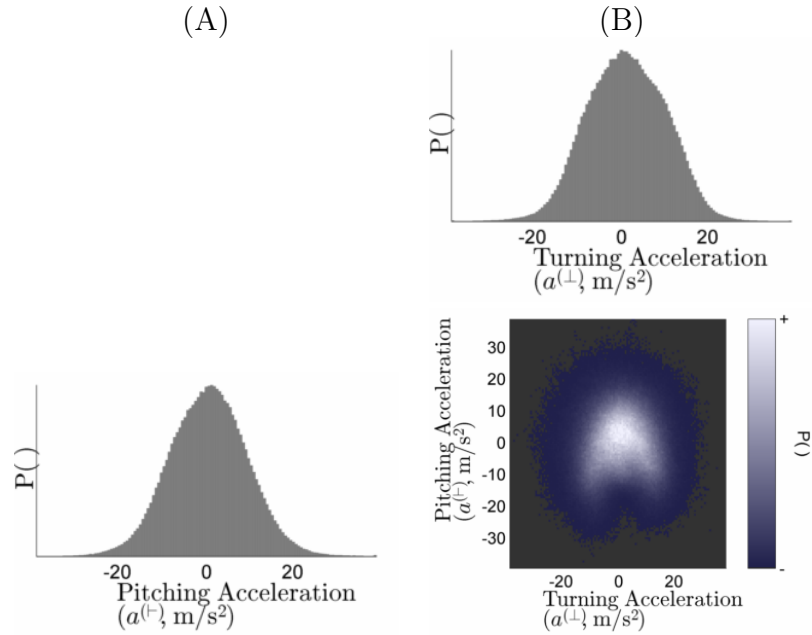


Figure 4-3: The joint distribution between turning acceleration, $a_i^{(\perp)}$, and pitching acceleration, $a_i^{(+)}$ is shown (B, bottom) for the same dataset used in Figure 4-1. For reference, the marginal distributions of $a_i^{(\perp)}$ (B, top), and $a_i^{(+)}$ (A, bottom) are shown. The “lima-bean” shape observed in this distribution is typical of the distributions also obtained from all the other datasets. By examining this distribution, it seems that pitching up is less likely to be associated with turns to the left and right than pitching down.

bean” shape across all of our datasets, suggesting that turns to the left and right are more likely to be associated with slower movement speed, possibly because it may be necessary to slow down in order to maneuver successfully. Due to the shape of the distribution, it is not amenable to correlation analysis.

4.3 Bat spatial distributions

Understanding bat behavior by examining aggregate statistics, such as joint and marginal distributions, can give us insights into some aspects of bat behavior, but the behaviors of animals as a function of the relative position of their neighbors may




Figure 4.4: Joint distributions between speed and turning angle, θ_i (A), turning acceleration, $a_i^{(\perp)}$ (B), and pitching acceleration, $a_i^{(\perp)}$ (C), are shown. These distributions are from the same dataset used in Figure 4.1, but the patterns are typical of all of our data sets. We observe a negative relationship between speed and changes in heading, especially large changes in heading (A). However, we observe a positive relationship between speed and pitching up (B), so other types of changes in heading must account for the negative relationship observed in (A). We observe a “lima-bean” shape in the distribution between speed and turning acceleration (C), indicating that large turns to the left or right may be associated with slower flight.

not be isotropic, and so further analysis of the spatial patterns of various kinematic quantities is warranted.

4.3.1 Related Work

Ballerini et al. (2008) studied flocks of starlings and found that the typical arrangement between birds was anisotropic – it was more likely for neighboring birds to occur along certain directions than others, relative to the direction of flight. Earlier work (Partridge et al., 1980) found similar effects in three species of commercially important fish. Lukeman et al. (2010) and Katz et al. (2011) studied the spatial arrangements of ducks and golden shiners respectively by making spatial histograms of the relative position of neighboring animals. For each animal, they defined a local coordinate frame, based on the animal’s heading, and computed the relative position between it and its neighbors, rotated into that coordinate frame (described in more detail for bats in Section 4.3.2). Then, they created a two dimensional grid, with one

bin per squared spatial unit of some appropriate size. Then, they created the spatial histogram by counting the number of neighboring animals that occurred in each bin (more details to follow in Section 4.3.3).

By understanding the circumstances when animals engage in certain behaviors, such as turning right or left, we can gain insight into how behaviors such as attraction and repulsion are realized. By creating spatial histograms with appropriate weights, such analysis is possible. For example, Katz et al. (2011) examined the spatial pattern of both turning and changes in speed in order to learn that changes in speed are an important aspect of the behavior of fish, in contrast with many models of collective motion that assume that changes in heading are the most important factor.

4.3.2 Coordinate systems for studying bat behavior

We have examined the behavior of bats as a function of the spatial relationships between all bats and their neighbors by computing the spatial histograms in three different coordinate systems. Previous authors studied systems in 2D, so in order to study the bats, we needed an extra step in order to choose an appropriate “up” vector, needed to construct an appropriate 3D coordinate system. A global coordinate system is a single frame of reference based on the geometry of the cameras. A column-based coordinate system is a time-varying coordinate system, computed once for each frame, that captures the current location and orientation of the column in space. A local coordinate system is computed for each moment in time for each bat, and depends on the heading of individual animals, as done by previous authors. We have tended to favor the local coordinate system (Subsection 4.3.2), but describe all three below for completeness.

Global

In order to align the scene with gravity, we calculated the direction and magnitude of gravity, \mathbf{g} , by observing the ballistic trajectory of several hot packs thrown in the air. A unit vector in the direction opposite of gravity, $\mathbf{u}_{\mathbf{g}}$, was used to define the Z axis, $\mathbf{u}^{(G,Z)}$, of the world coordinate system. The X axis of the world coordinate system, $\mathbf{u}^{(G,X)}$, is defined parallel to the direction of the optical axis of the center camera, projected onto a plane perpendicular to gravity. We place our cameras in the field so that their view direction is roughly perpendicular to the orientation of the bat column, so that the column progresses roughly along the Y axis. The arbitrary origin of the system is defined by the center of the set of calibration points.

Column / Group

During emergence, the group of *T. brasiliensis* bats move in a column, which is a long, undulating structure. We denote all of the bats observable at any moment in time as $\mathbf{B}[t]$. Over a short distance, such as our field of view, the shape of the column can reasonably be modeled as an oriented elliptic cylinder whose center $\mathbf{b}[t]$, orientation $\mathbf{u}^{(\mathbf{B},X)}[t]$, and cross-section shape (size, aspect ratio, and orientation) may change over time. To compute these quantities so that they vary smoothly throughout the sequence, we use the position and velocity of bats from frames both before and after the frame t .

The orientation of the column, $\mathbf{u}^{(\mathbf{B},X)}[t]$, is defined as the unit vector in the direction of the average velocity of all bats present. We define the column-relative orthonormal coordinate system based on $\mathbf{u}^{(\mathbf{B},X)}[t]$ and $\mathbf{u}_{\mathbf{g}}$, a unit vector in the direction of gravity, \mathbf{g} . The Z axis of the column coordinate system is not held parallel to gravity because this coordinate frame is used to represent the position and orientation of the column as an object, as opposed to representing relative direction, as in the

local coordinate systems below:

$$\begin{aligned}
\mathbf{u}^{(\mathbf{B},X)} [t] &= \sum_{i=1}^{N_t} \mathbf{v}_i \\
\mathbf{u}^{(\mathbf{B},Y)} [t] &= \mathbf{u}_g \times \mathbf{u}^{(\mathbf{B},X)} [t] \\
\mathbf{u}^{(\mathbf{B},Z)} [t] &= \mathbf{u}^{(\mathbf{B},X)} [t] \times \mathbf{u}^{(\mathbf{B},Y)} [t] \\
\mathbf{U}^{(\mathbf{B})} [t] &= \begin{bmatrix} \frac{\mathbf{u}^{(\mathbf{B},X)} [i]t}{\|\mathbf{u}^{(\mathbf{B},X)} [i]t\|} \\ \frac{\mathbf{u}^{(\mathbf{B},Y)} [i]t}{\|\mathbf{u}^{(\mathbf{B},Y)} [i]t\|} \\ \frac{\mathbf{u}^{(\mathbf{B},Z)} [i]t}{\|\mathbf{u}^{(\mathbf{B},Z)} [i]t\|} \end{bmatrix}.
\end{aligned} \tag{4.9}$$

The ‘‘location’’ of the column is defined with respect to the world coordinate system. The centroid of the positions of the bats present at time t is $\frac{1}{N} \text{sum}_{i=1}^N \mathbf{x}_i$, where N is the total number of bats. To facilitate using the first column-relative coordinate to represent the progression of a bat along the length of the column, we stabilize the column ‘‘location,’’ $\mathbf{b} [t]$, to lie in a plane parallel to gravity, coincident with the optical axis of the center camera by projecting the centroid of positions of the animals along $\mathbf{u}^{(\mathbf{B},X)} [t]$ until it intersects the plane, as follows:

$$\mathbf{b} [t] = \frac{1}{N} \sum_{i=1}^N \mathbf{x}_i + (c_{2,Y} - x_i^{(Y)} [t]) \mathbf{u}^{(\mathbf{B},X)} [t] \tag{4.10}$$

The center line of the column is defined as a line with the orientation of $\mathbf{u}^{(\mathbf{B},X)} [t]$ (parallel to the average velocity of all bats present at time t), passing through $\mathbf{b} [t]$. The coordinates of the bats, relative to the column are obtained by subtracting the location of the column center from the absolute position of the bats, then rotating into the column coordinate frame by multiplying by $\mathbf{U}^{(\mathbf{B})} [t]$, as follows:

$$\mathbf{x}_{i,\mathbf{B}} [t] = \mathbf{U}^{(\mathbf{B})} [t] (\mathbf{x}_i - \mathbf{b} [t]). \tag{4.11}$$

Note that the X coordinate now measures the progression along the column, whereas the Y and Z coordinates represent the position of the bat with respect to

the column center line. To compute the orientation, aspect ratio, and length of the elliptical column cross section, we perform principle components analysis on the Y and Z components of the column-relative coordinates, $\mathbf{x}_{i,\mathbf{B}} [t]$ of the bats. This is done by computing an outer product of the Y and Z components of the points, and then computing the eigenvectors and eigenvalues of that matrix. The eigenvectors of the covariance matrix give the orientation of the elliptical cross-section of the column and the square root of the eigenvalues of the covariance matrix give the aspect ratio. We use three times the square root of the eigenvalues as the length of the major and minor axes of the column cross section, $v_1^{\mathbf{B}} [t]$ and $v_2^{\mathbf{B}} [t]$. When performing subsequent analysis of the properties of the flight of the bats, we exclude any bat whose distance from the column center line is larger than $v_1^{\mathbf{B}} [t]$.

Local

The local coordinate frame (visualized in Figure 4.5) helps to express the relative positions of animals, capturing spatial relationships such as “behind”, “in front”, “to the right” and “to the left.” At each moment in time, the local coordinate frame, $\mathbf{U}_i [t]$, is defined based on gravity and the heading, \mathbf{h}_i of the animal. This local coordinate frame is represented as a 3×3 orthonormal matrix, where the rows are the local coordinate axes.

Since all animals are affected by gravity, we use \mathbf{u}_g as the local Z axis, $\mathbf{u}_i^{(Z)} [t]$ for all animals. The X axis, $\mathbf{u}_i^{(X)} [t]$, is the heading vector, projected into a plane perpendicular to \mathbf{u}_g . The Y axis, $\mathbf{u}_i^{(Y)} [t]$ is the cross product of these two vectors. When constructing the local coordinate frames, pitching and banking are purposefully not modeled.

$$\begin{aligned}
\mathbf{u}_i^{(X)} [t] &= \mathbf{u}_i^{(Y)} [t] \times \mathbf{u}_g \\
\mathbf{u}_i^{(Y)} [t] &= \mathbf{u}_g \times \mathbf{h}_i \\
\mathbf{u}_i^{(Z)} [t] &= \mathbf{u}_g
\end{aligned}
\tag{4.12}$$

$$\mathbf{U}_i [t] = \begin{bmatrix} \frac{\mathbf{u}_i^{(X)} [t]}{\|\mathbf{u}_i^{(X)} [t]\|} \\ \frac{\mathbf{u}_i^{(Y)} [t]}{\|\mathbf{u}_i^{(Y)} [t]\|} \\ \frac{\mathbf{u}_i^{(Z)} [t]}{\|\mathbf{u}_i^{(Z)} [t]\|} \end{bmatrix}$$

The relative position, $\mathbf{x}_{i,k}$ of bat b_k with respect to bat b_i is computed by computing the absolute relative position between the two animals, and then rotating it into the local coordinate frame, $\mathbf{U}_i [t]$, as follows.

$$\mathbf{x}_{i,k} = \mathbf{U}_i [t] (\mathbf{x}_k - \mathbf{x}_i).
\tag{4.13}$$

The local coordinate frame, $\mathbf{U}_i [t]$ may differ for each bat. For a pair of animals, i and k , the relative positions between the two animals, computed in their respective local coordinate frames, may not be symmetric, unless their coordinate frames are the same. In other words, $\mathbf{x}_{i,k} \neq -\mathbf{x}_{k,i}$, unless $\mathbf{U}_i [t]$ is equal to $\mathbf{U}_k [t]$

4.3.3 Computing spatial histograms

To tabulate the relative positions of bats, we established a 3D grid of voxels (3D pixels). The grid was four meters wide by 4 meters long (X and Y), and two meters tall (Z). The length of each side of the voxels was $d = 5$ cm. The relative positions of the bats were tabulated in one of the coordinate systems described above, usually the local coordinate system (Section 4.3.2). For this section, we will denote the relative position of two animals as $\mathbf{x}_{i,k}$, regardless of the coordinate system used. For each voxel, we counted the number of animals whose relative position was within a d cm³ cube centered at coordinates \mathbf{q} . Formally, given the set of all relative positions, \mathbf{X} ,




Figure 4.5: The local coordinate system for a bat is shown. The X axis is along the heading of the bat, the Z axis points up, and in order to make a right-handed coordinate system, the Y axis points over the left wing.

we can define a histogram function, \mathbf{W} , that takes as arguments the the center of a voxel, \mathbf{q} . It returns the number of neighboring positions, $\mathbf{x}_{i,k}$, that are contained within that voxel, as follows

$$\mathbf{W}(\mathbf{q}; \mathbf{X}) = \sum_{i=1}^N \sum_{k=1}^K \sum_{\Delta=-d/2}^{d/2} \delta(\mathbf{q} - \mathbf{x}_{i,k} + \Delta), \quad (4.14)$$

where K is the number of neighbors to include and δ is the Kronecker function. The spatial histogram may be computed with the following pseudocode:

```

for all q:
    W[q]=0;

for all q:
    for all x_ik:
        if q - d/ 2 < x_ik <= q + d/2:
            W[q]++;

```

A problem with this definition is that it does not account for the size of the bats. Bats are larger than 5 cm^3 because their wingspan is approximately 30 cm and their body length is approximately 10 cm. It would be undesirable to have voxels that were the same size as a bat, so we use kernel density estimation in order to create a histogram where each data point is convolved with a Gaussian, $\mathcal{N}(\mathbf{x}_{i,k} - \mathbf{q}, \mathbf{G})$, so that each single data point can be counted in multiple voxels, with a soft assignment, based on the distance between the voxel center, \mathbf{q} , and the point $\mathbf{x}_{i,k}$. \mathcal{N} is the normal distribution, and \mathbf{G} is a diagonal matrix with diagonal values that are $(1/6)^2$ of the wingspan of a bat. The histogram computed with kernel density estimation is described by

$$\mathbf{W}(\mathbf{q}; \mathbf{X}, \mathbf{G}) = \sum_{i=1}^N \sum_{k=1}^K \mathcal{N}(\mathbf{q} - \mathbf{x}_{i,k}, \mathbf{G}). \quad (4.15)$$

and computed with the following pseudocode,

```

for all q:
    W[q]=0;

for all q:
    for all x_ik:
        W[q] += mvnpdf(q - x_ik , G);

```

where $\text{mvnpdf}(\mathbf{q}-\mathbf{x}_{ik})$ returns $\frac{1}{(2\pi)^{\frac{3}{2}}|\mathbf{G}|^{\frac{1}{2}}}e^{-(\mathbf{q}-\mathbf{x}_{i,k})^T\mathbf{G}^{-1}(\mathbf{q}-\mathbf{x}_{i,k})}$.

In order to visualize the spatial histogram, it is necessary to examine 2D slices. We typically take slices through the center of the histogram, from each of the three major axes (Figure 4.6).




Figure 4-6: The spatial histogram, \mathbf{W} (Equation 4.15), of relative position of first nearest neighbors is shown from two of our datasets, taken on the same night at different times, before and after sunset, which occurred at 8:42 PM. In the top row, cartoon diagrams show the pose of the bat, relative to the position of the viewer, and in the bottom two rows, slices through the spatial histograms, based on those viewpoints are shown. Note that the scale of the diagrams is different from the scale of the slices through the spatial histogram. The first column depicts the overhead view, cutting perpendicular to the Z axis. The center column shows a side view, cutting perpendicular to the Y axis (over the left wing). The right column shows a view from the front, cutting perpendicular to the X axis (parallel to the heading). The gray regions depict volumes where there is no data. The lack of data in the center of the spatial histogram is due to occlusion. The direction of this region is due to the location of the cameras, relative to the location of the bats. The bright halo effect occurs because most of the first nearest neighbors of the bats are within 1 meter of each other in the dense column depicted in this video segment. If more neighbors were included in these histograms, the halo would be less pronounced. Earlier in the evening, the bats fly closer together.

4.3.4 Spatial patterns of turning and pitching acceleration

In order to visualize the spatial patterns of kinematic quantities, such as turning or pitching acceleration, we define weights associated for each data point, and compute a weighted spatial histogram as

$$\mathbf{W}(\mathbf{q}; \mathbf{X}, \mathbf{w}, \mathbf{G}) = \sum_{i=1}^N \sum_{k=1}^K w_i \mathcal{N}(\mathbf{q} - \mathbf{x}_{i,k}, \mathbf{G}), \quad (4.16)$$

computed with the following pseudocode:

```

for all q:
    W[q]=0;

for all q:
    for all x_ik:
        W[q] += w_i mvnpdf(q - x_ik , G);

```

To visualize the spatial pattern of *average* turning or pitching acceleration, we display $\mathbf{W}(\mathbf{q}; \mathbf{X}, \mathbf{w}, \mathbf{G}) / \mathbf{W}(\mathbf{q}; \mathbf{X}, \mathbf{1}, \mathbf{G})$. In other words, the weighted histogram normalized by the total number of points occurring in each bin.

By examining spatial patterns of changes in motion, specifically turning acceleration and pitching acceleration, we have found evidence supporting the repulsion and attraction assumptions of many behavioral models. By examining the spatial distribution of turning acceleration, (Figure 4.7, top row), we found that when the first nearest neighbors of bats are within 0.3 and 0.75 meters directly to the right or left, the focal bat will tend to turn to the opposite direction. Whereas at longer distances, greater than 1 m, the bats will tend to turn towards their neighbors. Our findings confirm that at very short distances, bats will tend to turn away from each other, whereas when bats are far apart, they will tend to turn towards each other. This pattern depends on the time of night when the bats are observed. Both distances are

longer later in the evening, after sunset.

By examining the spatial distribution of pitching acceleration (Figure 4.7, bottom row), we found that when the first nearest neighbors of bats are above their current position, the bats will tend to pitch up, even when neighbors are as little as 0.3 meters above. Bats will tend to pitch downwards only when their first nearest neighbors are further away, at least 0.5 meters. This may reflect a larger global trend, namely that as bats fly out of their caves, the entire stream of bats tends to pitch upwards. We did not find evidence of repulsion in the spatial distribution of pitching. Turning interactions may happen at slightly longer spatial scales than pitching because the wingspan of *T. brasiliensis* is 30cm (Farney and Fleharty, 1969), whereas their body is only a few centimeters tall.

The spatial distribution of speed changes did not reveal evidence of repulsion and attraction. This may be because the bats prefer to resolve such conflicts through changes in heading. It might also be the case that bats use speed to regulate their spatial relationships with their neighbors, but the way they do it is highly variable, so there is not a consistent spatial pattern. Finally, we observe that the mean of the distribution of speed changes was always larger than zero, suggesting that the bats are generally increasing their speed as they fly up and out of the cave.

4.4 Conclusion

The findings presented in this chapter give new insight into the behavior of bats as they fly in the emergence column and engage in complex social interactions. We conclude that the principal way that bats move within the column is by changing their heading. Across all of our datasets, we observed that the relative position between bats differs, depending on the time of night (Figure 4.6); early in the evening the bats

Davis, July 22

Bamberger, July 24

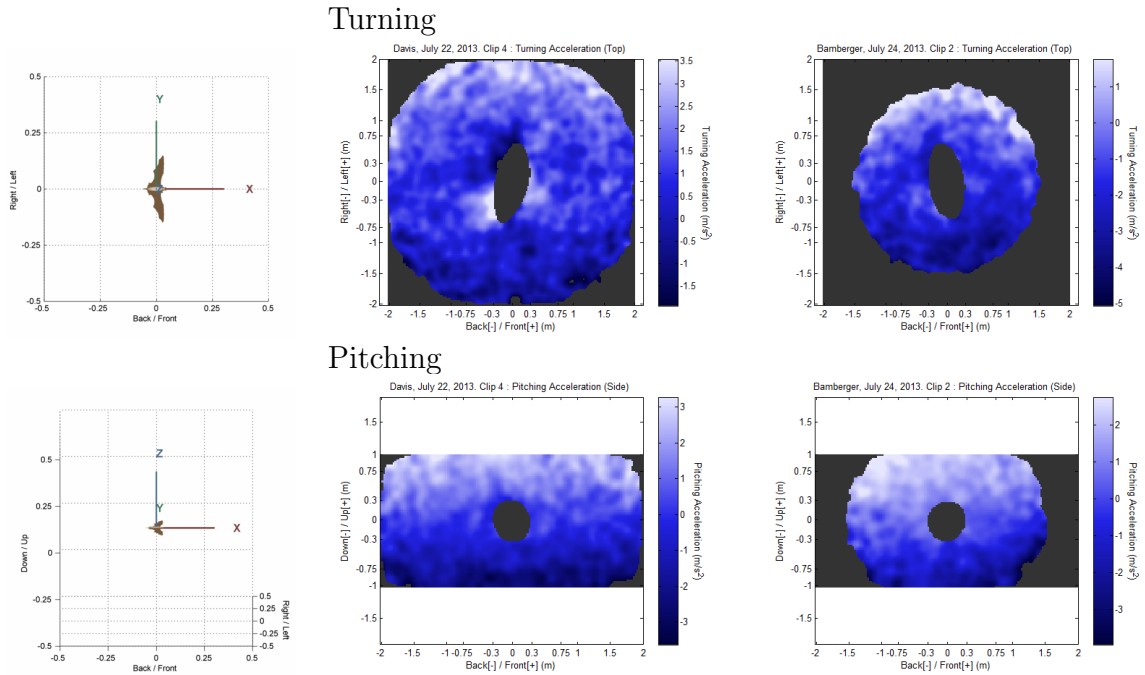


Figure 4-7: A spatial histogram, \mathbf{W} (Equation 4.16), depicting the motion of bats as a function of the relative position of their neighbors is shown, where the weights were given by the turning acceleration, $a_i^{(\perp)}$ (top), and pitching acceleration, $a_i^{(\parallel)}$ (bottom). In the top row, a dark color represents turning right, and a light color represents turning left. In the bottom row, a dark color represents pitching down and a light color represents pitching up. Diagrams depicting the pose of the bat relative to the viewer are shown in the first column. We observe that bats tend to turn away from things that are too close, and towards things that are too far away. For example, when neighboring bats are within 0.5 meters, and directly to the left of the focal bat, the focal bat will tend to turn to the right. However, if the first nearest neighbor is more than 1.5 meters away on the left side, the focal bat will tend to turn to the left. Note that the diagrams and spatial histograms have different scales.

tend to fly closer together, whereas later in the evening, their formations are looser.

We find that there is a spatial pattern to the way bats adjust their heading (Figure 4.7), turning away from other bats that are less than half a meter away and towards other bats when their first nearest neighbors are more than a meter away. Although the datasets we have collected form the most comprehensive set of observations of bats in their natural habitat that has been collected to date, even larger datasets, collected over more nights, would enable the understanding of trends in behavior related to wind, weather, cave entrance geometry, or even the reproductive cycle.

Katz et al. (2011) performed analysis of fish in groups of two fish and three fish, and found that the interactions of fish in groups of three was not an average of the responses to fish in groups of two. For our work, there are two possible lines of inquiry that are suggested by this finding. It might be very interesting to try to repeat their experiment by focusing on the behavior of bats only during moments in time when only two or three bats were present (or groups of two and three bats were physically well separated from other bats). This could lead to a better understanding of how bat behavior changes as the number of interacting conspecifics increases.

A straightforward extension of the work presented in this chapter would be to investigate the motion of bats by creating higher-dimensionality histograms. For example, using a six-dimensional histogram, it would be possible to study the behavior of bats as a function of both their first *and* second neighbors (as opposed to creating a three dimensional histogram with both the first and second neighbors as data points). The work by Katz et al. (2011) suggests that analysis of these two different ways of aggregating the responses of bats to multiple neighbors might lead to different results.

Data-driven simulation is a possible application of detailed histograms of the behavior of bats. If we tabulated the 3D acceleration of the bats as a function of the

relative position of their neighbors, it might be possible to use this information to drive a simulation. For each particle at each moment in time, the relative position of other simulated particles would be determined, then used to look up the motion that bats tended to adopt in a similar situation. The particle under consideration could then be assigned a motion “typical” of a bat. In this way, without understanding why bats do what they do, we might be able to capture *what* they do well enough to be able to generate simulations with similar properties, which might be of interest in computer graphics.

In our data, we observed periodic motion consistent with the wingbeat frequency of *T. brasiliensis*. It is remarkable that it is possible to see the bats flapping their wings in our kinematic data, without any other information about pose from the video. This is a very rich area for further inquiry. We could study how bats change the way they flap their wings as they engage in certain actions, such as turning or pitching. We could also investigate the extent to which bats synchronize their wing beats with their neighbors, and at what distances such synchronization might occur.

Moving beyond the study of the interactions between individuals, it will also be interesting to study the interaction between individuals and “the group.” For example, we could study the way that the motion of bats is different as a function of their position within the column, such as their flight speed or tendency to change heading. We could also examine the way bats tend to change their relationship to the column over time, investigating whether bats tend to maintain their position within the column or whether they tend to move back and forth between the center and the periphery. Finally, it would be interesting to study whether or not the interactions between individuals change as a function of the position of the individuals within the column. For example, when bats are too close together, they may do something

different if both bats are near the top of the column, versus the right side of the column or the center.

Unfortunately, our ability to understand interactions over very short spatial scales is at odds with our need to disregard data points where bats may occlude each other, leading to inaccurate reconstructions of 3D positions, and hence inaccurately estimated kinematics. This problem cannot be solved with higher resolution cameras, since the bats will remain the same physical size, and so any bats that are sufficiently close may occlude each other regardless of the resolution of the cameras. This problem may be somewhat remedied by using more varied camera placements. It was a coincidence that the cave entrance geometry of both colonies that we studied lent themselves to configurations where the cameras were placed so that the bats flew across the field of view from right to left. Finding different points of view might help to rule out whether or not certain effects are due the actual behavior of bats, or to occlusion effects.

Chapter 5

Conclusion

For more than thirty years, scientists have worked to understand the underlying principles of collective behavior by developing computational models and analyzing the observed behavior of animals and other organisms. Through simulation studies, the scientific community has discovered that large groups can move together without leaders when group members exhibit simple behaviors in response to the relative position and heading of other individuals in the group. Recently, as it has become feasible for researchers to collect large datasets of detailed observations of animals, the scientific community is presented with new opportunities to characterize and discover interesting aspects of animal behavior. The contributions of this thesis are 1.) an improved strategy for collecting data of airborne animals in a field setting 2.) a new computational model of collective motion, and 3.) the first large-scale dataset of 3D tracks of large groups of bats in flight in their natural habitat and subsequent analysis.

Our strategy for collecting data (Theriault et al., 2014) is an improvement over the previous state of the art because our approach can be deployed in a single afternoon with minimal extra equipment, and determining the relative pose of the cameras requires only annotations of matching points across video streams, combined with the internal parameters of the cameras. We also developed an approach for characterizing

the anticipated reconstruction uncertainty inherent due to the camera placement and determined that in the datasets we collected, led to reconstruction uncertainty smaller than the size of the body of a bat. In previous work (Cavagna et al., 2008) where 3D observations of starlings were collected, the relative positions of the cameras were meticulously measured by hand and left in place in a sheltered location and they did not consider or quantify reconstruction uncertainty. In our experience, the toe-in camera placement approach that we proposed enables reconstruction and tracking algorithms to successfully localize many more bats simultaneously than a traditional camera placement approach with parallel image planes. With this approach, we were able to collect a substantial dataset of hundreds of thousands of trajectories, containing millions of data points from 70 minutes of video taken over four nights.

The FlockOpt model proposed in this thesis is a rich paradigm for modeling collective motion. Our model is promising because it addresses important shortcomings that have been discovered in existing models. Because our model is better able to express variability in speed, simulations with our model exhibit patterns of motion not possible with earlier, constant-speed models. Our model gives us an important insight into the fundamental nature of collective motion: that ordered motion may arise in order to preserve spatial relationships, and variations in speed may be structured to achieve this goal, not simply random variations. Based on this insight, we examined our data of bats and we concluded that the spatial relationships between bats are more stable than they would be if variations in speed were random perturbations. We also found that a generic model that predicts that bats will minimize the change in relative position between themselves and their neighbors better predicted the individual behavior of bats than a model that assumed that bats averaged the velocities of their neighbors. Analysis of our model leads to a simple but powerful insight into

natural patterns of collective motion which are composed predominantly of groups that either rotate or translate coherently together: Euclidean transformations that maintain relative position include rotations and translations. The FlockOpt model espoused in this thesis can be used in any of the diverse applications where models of collective motion have been used productively, including computer graphics, biology, physics, robotics, and civil engineering.

With the observations of bats in flight that we collected, we were presented with an unprecedented opportunity to study the behavior of bats in flight. We focused on characterizing the properties of their motion, both in the aggregate and as a function of the relative position between themselves and their neighbors. Studying the changes in the motion of the bats as a function of the relative position of their neighbors allowed us to examine the validity of the repulsion and attraction assumptions of many behavioral model. We found that most of the changes in the motion of the bats that were observed were accounted for by turning, both left/right and up/down. Changes in speed did not display an obvious spatial pattern, but we know from our analysis of the changes in spatial relationships over time that speed regulation must play an important role in their behavior, so this is an opportunity for further study. Additional opportunities for study are presented by characterizing and understanding relationships between individual bats and the group, such as the way that bats move within the column, and whether or not there are column-level spatial patterns in kinematics. Although our dataset is the most comprehensive to date, it is still not large enough to fully understand the effects of wind, weather, or the variations in behavior over the course of a season. Collecting such a comprehensive dataset, using the camera calibration and placement approach proposed in this thesis, would enable the study of many different aspects of bat behavior.

In this thesis, we contributed to the state of the art by proposing an improved approach for data collection, formalizing a new model of collective motion that addresses shortcomings of earlier models and leads to new insights, and collecting and analyzing the largest dataset of observations of bats in flight that has been collected to date. Our approach for data collection promises to enable other scientists to collect comprehensive datasets by substantially reducing the technical difficulty, relative to previous approaches. Our computational model represents a new paradigm for modeling collective motion that we suggest will open new lines of inquiry in the study of collective animal behavior, as well as new approaches in other areas such as control theory and computer graphics. Our analysis of the behavior of bats in the dataset we collected provides a tantalizing glimpse into the possibilities of the understanding that the approaches presented in this thesis will enable in the future.

References

- Y. I. Abdel-Aziz and H. M. Karara. Direct linear transformation from comparator coordinates into object space coordinates in close-range photogrammetry. In *Proceedings of the Symposium on Close Range Photogrammetry*, pages 1 – 18, Falls Church, VA, 1971. American Society of Photogrammetry. Direct linear transformation into object space coordinates in close-range photogrammetry.
- I. Aoki. A simulation study on the schooling mechanism in fish. *Bulletin of the Japanese Society of Scientific Fisheries*, 48, 1982.
- A. Attanasi, A. Cavagna, L. Del Castello, I. Giardina, T. S. Grigera, A. Jelic, S. Melillo, L. Parisi, O. Pohl, E. Shen, and M. Viale. Information transfer and behavioural inertia in starling flocks. *Nature Physics*, 10(9):691–696, 2014. doi: 10.1038/nphys3035. URL <http://dx.doi.org/10.1038/nphys3035>.
- M. Ballerini, N. Cabibbo, R. Candelier, A. Cavagna, E. Cisbani, I. Giardina, V. Lecomte, A. Orlandi, G. Parisi, A. Procaccini, M. Viale, and V. Zdravkovic. Interaction ruling animal collective behavior depends on topological rather than metric distance: Evidence from a field study. *Proceedings of the National Academy of Sciences*, 105(4):1232–1237, January 2008.
- M. Betke, D. E. Hirsh, A. Bagchi, N. I. Hristov, N. C. Makris, and T. H. Kunz. Tracking large variable numbers of objects in clutter. In *Computer Vision and Pattern Recognition, 2007. CVPR '07. IEEE Conference on*, pages 1–8, June 2007. doi: 10.1109/CVPR.2007.382994.
- M. Betke, D. E. Hirsh, N. C. Makris, G. F. McCracken, M. Procopio, N. I. Hristov, S. Tang, A. Bagchi, J. D. Reichard, J. W. Horn, S. Crampton, C. J. Cleveland, and T. H. Kunz. Thermal imaging reveals significantly smaller brazilian free-tailed bat colonies than previously estimated. *Journal of Mammalogy*, 89(1):18–24, 2008. doi: 10.1644/07-MAMM-A-011.1. URL <http://www.asnjournals.org/doi/abs/10.1644/07-MAMM-A-011.1>.
- Ashwin A Bhandiwad, David G Zeddies, David W Raible, Edwin W Rubel, and Joseph A Sisneros. Auditory sensitivity of larval zebrafish (*Danio*

- rerio) measured using a behavioral prepulse inhibition assay. *The Journal of experimental biology*, 216(18):3504 – 3513, 2013.
- W. Bialek, A. Cavagna, I. Giardina, T. Mora, E. Silvestri, M. Viale, and A. M. Walczak. Statistical mechanics for natural flocks of birds. *Proceedings of the National Academy of Sciences*, 109(13):4786–4791, 2012. doi: 10.1073/pnas.1118633109. URL <http://www.pnas.org/content/109/13/4786.abstract>.
- N. W. F. Bode, J. J. Faria, D. W. Franks, J. Krause, and A. J. Wood. How perceived threat increases synchronization in collectively moving animal groups. *Proceedings of the Royal Society B: Biological Sciences*, 277(1697):3065–3070, 2010. doi: 10.1098/rspb.2010.0855. URL <http://rspb.royalsocietypublishing.org/content/277/1697/3065.abstract>.
- N. W. F. Bode, D. W. Franks, and A. J. Wood. Limited interactions in flocks: relating model simulations to empirical data. *Journal of The Royal Society Interface*, 8(55):301–304, 2011. doi: 10.1098/rsif.2010.0397. URL <http://rsif.royalsocietypublishing.org/content/8/55/301.abstract>.
- C. M. Breder. Equations descriptive of fish schools and other animal aggregations. *Ecology*, 35(3):361–370, 1954. URL <http://www.jstor.org/stable/1930099>.
- J. Buhl, D. J. T. Sumpter, I. D. Couzin, J. J. Hale, E. Despland, E. R. Miller, and S. J. Simpson. From disorder to order in marching locusts. *Science*, 312(5778):1402–1406, 2006. doi: 10.1126/science.1125142. URL <http://www.sciencemag.org/content/312/5778/1402.abstract>.
- D. S. Calovi, U. Lopez, S. Ngo, C. Sire, H. Chaté, and G. Theraulaz. Swarming, schooling, milling: phase diagram of a data-driven fish school model. *New Journal of Physics*, 16(1):015026, 2014. URL <http://stacks.iop.org/1367-2630/16/i=1/a=015026>.
- M. Cao, A. S. Morse, and B. D. O. Anderson. Reaching a consensus in a dynamically changing environment: A graphical approach. *SIAM Journal on Control and Optimization*, 47(2):575–600, 2008.
- A. Cavagna, I. Giardina, A. Orlandi, G. Parisi, A. Procaccini, M. Viale, and V. Zdravkovic. The STARFLAG handbook on collective animal behaviour: 1. Empirical methods. *Animal Behaviour*, 76(1):217 – 236, 2008.
- A. Cavagna, A. Cimorelli, I. Giardina, G. Parisi, R. Santagati, F. Stefanini, and M. Viale. Scale-free correlations in starling flocks. *Proceedings of the*

- National Academy of Sciences*, 2010. doi: 10.1073/pnas.1005766107. URL <http://www.pnas.org/content/early/2010/06/11/1005766107.abstract>.
- Christopher James Clark. Courtship dives of Anna’s Hummingbird offer insights into flight performance limits. *Proceedings of the Royal Society B: Biological Sciences*, 276(1670):3047 – 3052, 2009.
- I. D. Couzin, J. Krause, R. James, G. D. Ruxton, and N. R. Franks. Collective memory and spatial sorting in animal groups. *Journal of Theoretical Biology*, 218(1):1 – 11, 2002. ISSN 0022-5193. doi: 10.1006/jtbi.2002.3065.
- A. Czirók, M. Vicsek, and T. Vicsek. Collective motion of organisms in three dimensions. *Physica A: Statistical Mechanics and its Applications*, 264(1):299–304, 1999. ISSN 0378-4371. doi: 10.1016/S0378-4371(98)00468-3. URL <http://www.sciencedirect.com/science/article/pii/S0378437198004683>.
- R. Eberhart and Yuhui S. Particle swarm optimization: developments, applications and resources. In *Proceedings of the 2001 Congress on Evolutionary Computation*, volume 1, pages 81 –86, 2001. doi: 10.1109/CEC.2001.934374.
- A. Eriksson, M. Nilsson Jacobi, J. Nyström, and K. Tunstrøm. Determining interaction rules in animal swarms. *Behavioral Ecology*, 21(5): 1106–1111, 2010. doi: 10.1093/beheco/arq118. URL <http://beheco.oxfordjournals.org/content/21/5/1106.abstract>.
- John Farney and Eugene D. Fleharty. Aspect ratio, loading, wing span, and membrane areas of bats. *Journal of Mammalogy*, 50(2):362–367, 1969. ISSN 0022-2372. doi: 10.2307/1378361.
- J. Gautrais, F. Ginelli, R. Fournier, S. Blanco, M. Soria, H. Chaté, and G. Theraulaz. Deciphering interactions in moving animal groups. *PLoS Comput Biol*, 8(9):e1002678, 09 2012. doi: 10.1371/journal.pcbi.1002678. URL <http://dx.doi.org/10.1371%2Fjournal.pcbi.1002678>.
- Veysel Gazi and Kevin M. Passino. A class of attractions/repulsion functions for stable swarm aggregations. *International Journal of Control*, 77(18): 1567–1579, 2004. doi: 10.1080/00207170412331330021. URL <http://dx.doi.org/10.1080/00207170412331330021>.
- R. C. Gerum, B. Fabry, C. Metzner, M. Beaulieu, A. Ancel, and D. P. Zitterbart. The origin of traveling waves in an emperor penguin huddle. *New Journal of Physics*, 15(12):125022, 2013. URL <http://stacks.iop.org/1367-2630/15/i=12/a=125022>.

- W. D. Hamilton. Geometry for the selfish herd. *Journal of Theoretical Biology*, 31(2):295–311, 1971.
- R. Hartley and A. Zisserman. *Multiple view geometry in computer vision*. Cambridge University Press, 2 edition, 2004.
- Marika Hayashi, Kara L Feilich, and David J Ellerby. The mechanics of explosive seed dispersal in orange jewelweed (*Impatiens capensis*). *Journal of Experimental Botany*, 60(7):2045 – 2053, 2009.
- Tyson L Hedrick. Software techniques for two- and three-dimensional kinematic measurements of biological and biomimetic systems. *Bioinspiration & Biomimetics*, 3(3):034001, 2008. URL <http://stacks.iop.org/1748-3190/3/i=3/a=034001>.
- D. Helbing, P. Molnar, I.J. Farkas, and K. Bolay. Self-organizing pedestrian movement. *Environment and Planning B*, 28(3):361–384, 2001.
- D. Helbing, L. Buzna, A. Johansson, and T. Werner. Self-organized pedestrian crowd dynamics: experiments, simulations, and design solutions. *Transportation Science*, 39(1):1–24, 2005.
- J. E. Herbert-Read, A. Perna, R. P. Mann, T. M. Schaerf, D. J. T. Sumpter, and A. J. W. Ward. Inferring the rules of interaction of shoaling fish. *Proceedings of the National Academy of Sciences*, 108(46):18726–18731, 2011. doi: 10.1073/pnas.1109355108. URL <http://www.pnas.org/content/108/46/18726.abstract>.
- A. Huth and C. Wissel. The simulation of the movement of fish schools. *Journal of Theoretical Biology*, 156(3):365 – 385, 1992. ISSN 0022-5193. doi: 10.1016/S0022-5193(05)80681-2. URL <http://www.sciencedirect.com/science/article/pii/S0022519305806812>.
- A. Huth and C. Wissel. The simulation of fish schools in comparison with experimental data. *Ecological Modelling*, 75/76(0):135 – 146, 1994. ISSN 0304-3800. doi: 10.1016/0304-3800(94)90013-2. URL <http://www.sciencedirect.com/science/article/pii/0304380094900132>.
- A. Jadbabaie, J. Lin, and A.S. Morse. Coordination of groups of mobile autonomous agents using nearest neighbor rules. *Automatic Control, IEEE Transactions on*, 48(6):988–1001, 2003.
- D. Karaboga and B. Akay. A survey: algorithms simulating bee swarm intelligence. *Artificial Intelligence Review*, 31:61–85, 2009. ISSN 0269-2821. URL 10.1007/s10462-009-9127-4. 10.1007/s10462-009-9127-4.

- Y. Katz, K. Tunstrøm, C. C. Ioannou, C. Huepe, and I. D. Couzin. Inferring the structure and dynamics of interactions in schooling fish. *Proceedings of the National Academy of Sciences*, 108(46):18720–18725, 2011. doi: 10.1073/pnas.1107583108. URL <http://www.pnas.org/content/108/46/18720.abstract>.
- J. Krause, D. Lusseau, and R. James. Animal social networks: an introduction. *Behavioral Ecology and Sociobiology*, 63:967–973, 2009. ISSN 0340-5443. URL 10.1007/s00265-009-0747-0. 10.1007/s00265-009-0747-0.
- Ugo Lopez, Jacques Gautrais, Iain D. Couzin, and Guy Theraulaz. From behavioural analyses to models of collective motion in fish schools. *Interface Focus*, 2012. doi: 10.1098/rsfs.2012.0033.
- M.I. A. Lourakis and A.A. Argyros. SBA: A Software Package for Generic Sparse Bundle Adjustment. *ACM Trans. Math. Software*, 36(1):1 – 30, 2009. doi: <http://doi.acm.org/10.1145/1486525.1486527>.
- R. Lukeman, Y. Li, and L. Edelstein-Keshet. Inferring individual rules from collective behavior. *Proceedings of the National Academy of Sciences*, 107(28):12576–12580, 2010. doi: 10.1073/pnas.1001763107. URL <http://www.pnas.org/content/107/28/12576.abstract>.
- P. F. Major and L. M. Dill. The three-dimensional structure of airborne bird flocks. *Behavioral Ecology and Sociobiology*, 4(2):111–122, 1978. URL <http://www.jstor.org/stable/4599193>.
- R. P. Mann. Bayesian inference for identifying interaction rules in moving animal groups. *PLoS ONE*, 6(8):e22827, 08 2011. doi: 10.1371/journal.pone.0022827. URL <http://dx.doi.org/10.1371/journal.pone.0022827>.
- R. P. Mann, Andrea Perna, D. Strömbom, R. Garnett, J. E. Herbert-Read, D. J. T. Sumpter, and A. J. W. Ward. Multi-scale inference of interaction rules in animal groups using bayesian model selection. *PLoS Comput Biol*, 9(3):e1002961, 03 2013. doi: 10.1371/journal.pcbi.1002961. URL <http://dx.doi.org/10.1371/journal.pcbi.1002961>.
- S. Mishra, K. Tunstrøm, I. D. Couzin, and C. Huepe. Collective dynamics of self-propelled particles with variable speed. *Phys. Rev. E*, 86:011901, Jul 2012. doi: 10.1103/PhysRevE.86.011901. URL <http://link.aps.org/doi/10.1103/PhysRevE.86.011901>.
- Jonathan Daniel Munk. *The Descent of Ant*. PhD thesis, University of California, Berkeley, 2011.

- M. Nagy, Z. Ákos, D. Biro, and T. Vicsek. Hierarchical group dynamics in pigeon flocks. *Nature*, 464(7290):890–893, 2010. doi: 10.1038/nature08891. URL <http://www.nature.com/nature/journal/v464/n7290/abs/nature08891.html>.
- J. Nocedal and S. J. Wright. *Numerical optimization*. Springer Series in Operations Research, Springer Verlag, 2 edition, 2006.
- Ulla M. Lindhe Norberg and R. Ake Norberg. Scaling of wingbeat frequency with body mass in bats and limits to maximum bat size. *The Journal of Experimental Biology*, 215(5):711–722, 2012. doi: 10.1242/jeb.059865. URL <http://jeb.biologists.org/content/215/5/711.abstract>.
- A. Okubo. Dynamical aspects of animal grouping: Swarms, schools, flocks, and herds. *Advances in Biophysics*, 22(0):1 – 94, 1986. ISSN 0065-227X. doi: 10.1016/0065-227X(86)90003-1. URL <http://www.sciencedirect.com/science/article/pii/0065227X86900031>.
- Oladele A Olaniran, Akella VS Sudhakar, Falko P Drijfhout, Ian AN Dublon, David R Hall, James GC Hamilton, and William DJ Kirk. A male-predominant cuticular hydrocarbon, 7-methyltricosane, is used as a contact pheromone in the western flower thrips *Frankliniella occidentalis*. *Journal of chemical ecology*, pages 1 – 10, 2013.
- J. K. Parrish and L. Edelstein-Keshet. Complexity, pattern, and evolutionary trade-offs in animal aggregation. *Science*, 284(5411):99–101, 1999. doi: 10.1126/science.284.5411.99. URL <http://www.sciencemag.org/content/284/5411/99.abstract>.
- B. L. Partridge. Internal dynamics and the interrelations of fish in schools. *Journal of Comparative Physiology A: Neuroethology, Sensory, Neural, and Behavioral Physiology*, 144:313–325, 1981. ISSN 0340-7594. URL <http://dx.doi.org/10.1007/BF00612563>. 10.1007/BF00612563.
- B. L. Partridge and T. J. Pitcher. The sensory basis of fish schools: Relative roles of lateral line and vision. *Journal of Comparative Physiology A: Neuroethology, Sensory, Neural, and Behavioral Physiology*, 135:315–325, 1980. ISSN 0340-7594. URL <http://dx.doi.org/10.1007/BF00657647>. 10.1007/BF00657647.
- B. L. Partridge, T. Pitcher, J. M. Cullen, and J. Wilson. The three-dimensional structure of fish schools. *Behavioral Ecology and Sociobiology*, 6:277–288, 1980. ISSN 0340-5443. URL <http://dx.doi.org/10.1007/BF00292770>. 10.1007/BF00292770.

- T. J. Pitcher. The three-dimensional structure of schools in the minnow, *Phoxinus phoxinus* (L.). *Animal Behaviour*, 21(4):673 – 686, 1973. ISSN 0003-3472. doi: 10.1016/S0003-3472(73)80091-0. URL <http://www.sciencedirect.com/science/article/pii/S0003347273800910>.
- S. J. Portugal, T. Y. Hubel, J. Fritz, S. Heese, D. Trobe, B. Voelkl, S. Hailes, A. Wilson, and J. R. Usherwood. Upwash exploitation and downwash avoidance by flap phasing in ibis formation flight. *Nature*, 505(7483):399–402, 2014. doi: <http://dx.doi.org/10.1038/nature12939>. URL <http://www.nature.com/nature/journal/v505/n7483/abs/nature12939.html>.
- H. Pryer. An account of a visit to the birds'-nest caves of British Borneo. *Proceedings of the Zoological Society of London*, 52:532–538, 1884.
- D. V. Radakov. *Schooling in the Ecology of Fish*. John Wiley and Sons, New York NY, 1973.
- John H. Reif and Hongyan Wang. Social potential fields: A distributed behavioral control for autonomous robots. *Robotics and Autonomous Systems*, 27(3):171 – 194, 1999. doi: [http://dx.doi.org/10.1016/S0921-8890\(99\)00004-4](http://dx.doi.org/10.1016/S0921-8890(99)00004-4). URL <http://www.sciencedirect.com/science/article/pii/S0921889099000044>.
- C. W. Reynolds. Flocks, herds and schools: A distributed behavioral model. In *ACM SIGGRAPH Computer Graphics*, volume 21, pages 25–34. ACM, 1987.
- R.S. Sikes and W.L. Gannon. Animal Care and Use Committee of the American Society of Mammalogists guidelines of the american society of mammalogists for the use of wild mammals in research. *Journal of Mammalogy*, 92:235–253, 2011.
- Daniel Strömbom. Collective motion from local attraction. *Journal of Theoretical Biology*, 283(1):145 – 151, 2011. doi: 10.1016/j.jtbi.2011.05.019. URL <http://www.sciencedirect.com/science/article/pii/S002251931100261X>.
- D. H. Theriault, N. W. Fuller, B. E. Jackson, E. Bluhm, D. Evangelista, Z. Wu, M. Betke, and T. L. Hedrick. A protocol and calibration method for accurate multi-camera field videography. *The Journal of Experimental Biology*, 217(11):1843–1848, 2014. doi: 10.1242/jeb.100529. URL <http://jeb.biologists.org/content/217/11/1843.abstract>.

- T. Vicsek and A. Zafeiris. Collective motion. *Physics Reports*, 517(3-4):71 – 140, 2012. ISSN 0370-1573. doi: 10.1016/j.physrep.2012.03.004. URL <http://www.sciencedirect.com/science/article/pii/S0370157312000968>.
- T. Vicsek, A. Czirók, E. Ben-Jacob, I. Cohen, and O. Shochet. Novel type of phase transition in a system of self-driven particles. *Physical Review Letters*, 75:1226–1229, Aug 1995. doi: 10.1103/PhysRevLett.75.1226. URL <http://link.aps.org/doi/10.1103/PhysRevLett.75.1226>.
- Simon M Walker, Adrian LR Thomas, and Graham K Taylor. Photogrammetric reconstruction of high-resolution surface topographies and deformable wing kinematics of tethered locusts and free-flying hoverflies. *Journal of The Royal Society Interface*, 6(33):351 – 366, 2009.
- Z. Wu, N. I. Hristov, T. H. Kunz, and M. Betke. Tracking-reconstruction or reconstruction-tracking? comparison of two multiple hypothesis tracking approaches to interpret 3d object motion from several camera views. In *Motion and Video Computing, 2009. WMVC '09. Workshop on*, pages 1–8, Dec 2009. doi: 10.1109/WMVC.2009.5399245.

CURRICULUM VITAE

