

Mapping the Visual Coupling between Neighbors in Real and Virtual Crowds

By

Kevin W. Rio

B.S. Physics, Rensselaer Polytechnic Institute, 2009

B.S. Psychology, Rensselaer Polytechnic Institute, 2009

Submitted in partial fulfillment of the requirements

for the Degree of Doctor of Philosophy in the

Department of Cognitive, Linguistic, and Psychological Sciences at Brown University

Providence, Rhode Island

May 2015

© Copyright 2014 by Kevin W. Rio

This dissertation by Kevin W. Rio is accepted in its present form
by the Department of Cognitive, Linguistic, and Psychological Sciences
as satisfying the requirement for the degree of Doctor of Philosophy

Date _____

William H. Warren, Advisor

Recommended to the Graduate Council

Date _____

Thomas Serre, Reader

Date _____

Joo-Hyun Song, Reader

Approved by the Graduate Council

Date _____

Peter Weber, Dean of the Graduate School

Curriculum Vitae for Kevin W. Rio

kevin_rio@brown.edu
www.kevinrio.com

Dept. of Cognitive, Linguistic, & Psychological Sciences, Box 1821
Brown University, Providence RI 02912

Education

PhD, Cognitive Science, Brown University, 2015

Advisor: William H. Warren

Thesis: Mapping the local coupling between neighbors in real and virtual crowds

MS, Cognitive Science, Brown University, 2012

Advisor: William H. Warren

Thesis: Pedestrian following: Behavioral dynamics and visual information

BS, Psychology, Rensselaer Polytechnic Institute, 2009

Advisor: Mark A. Changizi

Thesis: Color vision in clinical medicine

BS, Physics, Rensselaer Polytechnic Institute, 2009

Research Interests

Crowd dynamics; perception and action; virtual reality; visual control of locomotion

Research Experience

Graduate RA, Virtual Environment Navigation Lab, Brown University, 2009-2014

Undergraduate RA, Changizi Lab, Rensselaer Polytechnic Institute, 2007-2009

Undergraduate RA, Center for THz Research, Rensselaer Polytechnic Institute, 2007

Publications

Rio KW, Rhea CK, & Warren WH (2014). Follow the leader: Visual control of speed in pedestrian following. *Journal of Vision* 14(2):4, 1-16.

Changizi MA, Brucksch M, Kotecha R, McDonald K, & **Rio KW** (2014). Ecological warnings. *Safety Science* 61, 36-42.

Rio KW & Warren WH (2012). A data-driven model of pedestrian following and emergent crowd behavior. In Weidmann, Kirsch, & Shreckenberg (Eds.), *Pedestrian and Evacuation Dynamics*, 561-574. New York: Springer.

Bonneaud S, **Rio KW**, Chevaillier P, & Warren WH (2012). Accounting for patterns of collective behavior in crowd locomotor dynamics for realistic simulations. *Lecture Notes in Computer Science 7145: Transactions on Edutainment VII*, 1-11.

Changizi MA & **Rio KW** (2010). Harnessing color vision for visual oximetry in central cyanosis. *Medical Hypotheses* 74, 87-91.

Presentations

- Rio KW** & Warren WH (2014). Visual coupling to multiple neighbors in a crowd influences walking speed and direction. Poster, Vision Sciences Society 14th Annual Meeting, St. Pete Beach FL.
- Rio KW** & Warren WH (2013). Visually-guided collective behavior in human swarms. Poster, Vision Sciences Society 13th Annual Meeting, Naples FL.
- Kiefer AW, Bonneaud S, **Rio KW**, & Warren WH (2013). Quantifying the coherence of pedestrian groups. Paper & Poster, Cognitive Science Society 35th Annual Meeting, Berlin DE.
- Rio KW**, Bonneaud S, & Warren WH (2012). Speed coordination in pedestrian groups: Linking individual locomotion with crowd behavior. Poster, Vision Sciences Society 12th Annual Meeting, Naples FL.
- Rio KW** & Warren WH (2011). A speed control law for pedestrian following based on visual angle. Poster, Vision Sciences Society 11th Annual Meeting, Naples FL.
- Rio KW**, Bonneaud S, & Warren WH (2010). Locomotor dynamics of small crowds: Characterizing individual and collective behavior. Poster, Psychonomics Society 51st Annual Meeting, St. Louis MO.
- Rio KW**, Rhea CK, & Warren WH (2010). Follow the leader: Behavioral dynamics of following. Poster, Vision Sciences Society 10th Annual Meeting, Naples FL.

Fellowships and Grants

- Link Foundation Fellowship in Advanced Training and Simulation, 2013-2014
- Brown University International Affairs Travel Grant, 2012
- EU Cog II Travel Grant, 2011
- NASA RI Space Grant Fellowship, 2011
- Honorable Mention, NSF Graduate Research Fellowship, 2010 & 2011
- Brown University, First Year Fellowship, 2009-2010

Teaching Experience

- Graduate TA, CLPS 0510 Perception, Illusion, & Visual Arts, Brown University, 2013
- Graduate TA, CLPS 0500 Perception and Mind, Brown University, 2011
- Graduate TA, CLPS 0900 Quantitative Methods in Psychology, Brown University, 2010
- Undergraduate TA, MATH 1010 Calculus I, Rensselaer Polytechnic Institute, 2008
- Undergraduate TA, PHYS 1100 Physics I, Rensselaer Polytechnic Institute, 2006

Academic Service

- Coordinator, Perception-Action Seminar, CLPS Dept., Brown University, 2012-2013
- Graduate Rep., Library Advisory Board, Brown University, 2011-2013

Professional Development

- Teaching Cert. III: Professional Development Seminar, Brown University, 2013
- Teaching Cert. I: The Sheridan Teaching Seminar, Brown University, 2011
- International Summer School: "Dynamic Field Theory: Neural Dynamics Approaches to Cognitive Robotics," University of Minho, Guimarães PT, 2011

Acknowledgments

Completing this dissertation has been a fascinating, challenging, and ultimately transformative stage in my academic career. First and foremost, I am grateful to have had Bill Warren as my advisor along the way. He has been knowledgeable, motivating, and above all, encouraging throughout. Vague ideas have an uncanny way of leaving Bill's office as concrete, actionable paths forward, and this guidance has indelibly shaped the scientist and researcher that I am today. My committee members, Thomas Serre and Joo-Hyun Song, provided helpful guidance in the planning stages and useful feedback on early versions of my dissertation. I would also like to acknowledge two of my undergraduate mentors, Mark Changizi and Larry Reid; their influence drove me to graduate school in psychology, and this dissertation would likely not exist were it not for them.

From the administrative staff to the faculty, the newly-unified CLPS department provided an ideal environment in which to learn, grow, and conduct research. I am especially thankful to my graduate cohort: Ali Arslan, Deanna Macris, Jeff Cockburn, Jie Ren, Jill Thorson, Huaiyong Zhao, who provided comradery, encouragement, and support over the years; I count them all as lifelong friends, and I wish them well. I would also like to thank my funding sources, who graciously supported the research presented here. These include the Link Foundation, the National Aeronautics and Space Administration / Rhode Island Space Grant Consortium, and the National Institutes of Health (NIH R01 EY010923).

I am grateful to all of my colleagues in the VENLab, past and present. Hugo Bruggeman, Jon Cohen, Liz Chrastil, Joost DeNijst, Jon Ericson, Chaz Firestone, Henry Harrison, and Chris Rhea were the elders when I arrived, and they served as models that I looked up, was inspired by, and learned from. Stéphane Bonneaud, Michael Fitzgerald, Adam Kiefer, Huaiyong Zhao, and Mintao Zhao joined the lab the same year as I did, and our time together was abuzz with the enthusiasm, energy, and naiveté that come with embarking on difficult, unsolved challenges. I am especially grateful to Stéphane; our many conversations and whiteboard sessions, countless journal clubs, and the occasional basket of French fries profoundly shaped the way I think about crowd behavior, experimental psychology, and science, and they helped make this dissertation what it is. Many undergraduate RA's helped in designing and conducting my experiments over the years; I am especially grateful to programmers Kurt Spindler, Reese Kuppig, and Ashwin Karumbunathan who contributed to early versions of the experiments presented here, and to Arturo Cardenas, Sara Palasits, Alex Sepolen, who served as confederates in Experiments 1 and 2. Lastly, I am thankful for the new generation in the VENLab, including Arturo Cardenas, Greg Dachner, Zach Page, and Trent Wirth, who have been outstanding colleagues and collaborators, and I wish them all many productive and transformative years in the VENLab.

I am indebted to the wild and scenic places throughout New England that have rejuvenated my spirits, given my mind the freedom to wander, and sustained my energy throughout graduate school; these include Baxter State Park, Easton's Beach, Mount Sunapee, Norman Bird Sanctuary, and Shawnee Peak.

Lastly, I must acknowledge my family and friends, whose encouragement, patience, and support made this all possible. My parents have always gone above and beyond to support my education, and nurtured a love of learning that will never go away; this dissertation is a testament to everything they have done over the years. Last, but most certainly not least, my wife Danielle has been an unending source of encouragement, friendship, patience, and warmth; without her, I am fairly certain I would have withered away (it can be difficult to remember to eat when you're writing), spontaneously combusted, or suffered some other equally unceremonious end.

Dedicated in memory of
Dr. Francis J. Rio (1912-1997)

Table of Contents

List of Tables	iii
List of Figures	iv
Chapter 1	1
Introduction: Collective Locomotion and Crowd Dynamics	
Chapter 2	5
Computational Models of Collective Locomotion	
Chapter 3	35
Empirical Studies of Collective Locomotion	
Chapter 4	54
Experiment 1: Walking in a Small Crowd of Confederates	
Chapter 5	75
Experiment 2: Validating a Virtual Crowd	
Chapter 6	85
Experiment 3: Walking in a Virtual Crowd	
Chapter 7	111
Experiment 4: Walking in a Real Crowd	
Chapter 8	122
General Discussion	
References	128

List of Tables

Table 4-1	62
Full list of conditions in Experiment 1.	

List of Figures

Figure 2-1	17
Discretization of simulation space.	
Figure 2-2	19
Schematic diagrams of the components of Reynolds' (1987) 'boids' model.	
Figure 2-3	22
Schematic diagram of an individual agent from Huth & Wissel's (1992) model of fish schooling.	
Figure 2-4	30
Four different types of aggregations, formed by different sets of parameter values for a single model.	
Figure 2-5	32
Variables in the social interaction model of Moussaid, et al. (2010).	
Figure 2-6	34
Results of Moussaid, et al. (2010).	
Figure 3-1	40
Mean walking speed as a function of log-scaled city size for 15 cities and towns.	
Figure 3-2	43
The shape of personal space.	
Figure 3-3	45
Probability density function of velocities from pedestrians at a crosswalk.	
Figure 4-1	60
Initial conditions of participant and three confederates in Experiment 1.	
Figure 4-2	61
Confederate trajectories for a sample turn left and turn right trial.	
Figure 4-3	65
Participant trajectories for all trials in which at least one confederate turns left or right.	

Figure 4-4	65
Participant trajectories for a single representative participant, for all trials in which a subgroup of confederates turned left.	
Figure 4-5	67
Mean changes in participants' final position in Experiment 1.	
Figure 4-6	67
Mean changes in participants' final position in Experiment 1, broken down by confederate number (left) and eccentricity (right).	
Figure 4-7	68
Participant speed time series for all trials in which a subgroup of confederates slowed down or sped up.	
Figure 4-8	71
Mean changes in participants' final speed in Experiment 1.	
Figure 4-9	71
Mean changes in participants' final speed in Experiment 1, broken down by confederate number and eccentricity.	
Figure 4-10	73
Mean changes in participants' final speed in Experiment 1, as observed in human data and as predicted by the model adapted from Rio, Rhea, & Warren (2014).	
Figure 4-11	73
Mean changes in participants' final speed in Experiment 1, broken down by confederate number and eccentricity, for data and model.	
Figure 5-1	86
Mean changes in participants' final lateral position in Experiment 2A and Experiment 2B.	
Figure 5-2	86
Mean changes in participants' final position in Experiment 2A and Experiment 2B, separating the effects of number and eccentricity.	
Figure 5-3	87
Mean changes in participants' final speed in Experiment 2A and Experiment 2B.	

Figure 5-4	87
Mean changes in participants' final speed in Experiment 2A and Experiment 2B, separating the effects of number and eccentricity.	
Figure 6-1	93
Metric and topological neighborhood structure.	
Figure 6-2	94
Decaying local coupling strength as a function of metric distance and topological distance.	
Figure 6-3	94
Decaying local coupling strength as a function of eccentricity, for metric and topological neighborhoods.	
Figure 6-4	97
Virtual displays from Experiment 3.	
Figure 6-5	103
Mean changes in participants' final lateral position in Experiment 3A.	
Figure 6-6	103
Mean changes in participants' collapsed final position in Experiment 3A.	
Figure 6-7	104
Mean changes in participants' final speed in Experiment 3A.	
Figure 6-8	104
Mean changes in participants' collapsed final speed in Experiment 3A.	
Figure 6-9	109
Mean changes in participants' final lateral position in Experiment 3B.	
Figure 6-10	109
Mean changes in participants' collapsed final position in Experiment 3B.	
Figure 6-11	110
Mean changes in participants' final speed in Experiment 3B.	

Figure 6-12	110
Mean changes in participants' collapsed final speed in Experiment 3B.	
Figure 6-13	115
Mean changes in participants' final lateral position in Experiment 3C, for "turn left" and "turn right" trials.	
Figure 6-14	115
Mean changes in participants' collapsed final position in Experiment 3C.	
Figure 6-15	116
Mean changes in participants' collapsed and combined final position in Experiment 3C.	
Figure 6-16	116
Mean changes in participants' final speed in Experiment 3C, for "slower" and "faster" trials.	
Figure 6-17	117
Mean changes in participants' collapsed final speed in Experiment 3C.	
Figure 6-18	117
Mean changes in participants' collapsed and combined final speed in Experiment 3C.	
Figure 7-1	122
Two helmets worn by participants in Experiment 4.	
Figure 7-2	124
A sample crowd from Experiment 4.	
Figure 7-3	124
Sample motion capture data from Experiment 4.	
Figure 7-4	126
Schematic diagrams of group polarization and angular momentum.	
Figure 7-5	127
Polarization and angular momentum for a representative trial [N=16, low initial density] in Experiment 4.	

Figure 7-6	129
Occupancy grid (percentage of frames in which at least one neighbor is found) over 6 min of data in Experiment 4.	
Figure 7-7	130
Mean absolute heading difference over 6 min of data in Experiment 4.	

CHAPTER 1

Introduction: Collective Locomotion and Crowd Dynamics

Collective locomotion is a ubiquitous feature of the natural world, found in living systems great and small—from migrating skin cells and metastasizing tumors to flocks of starlings hundreds of birds abreast and swarms of locusts so thick that they blot out the sun. In humans, collective locomotion takes the form of *crowd dynamics*. Spend a few moments gazing at the sea of people flooding through Grand Central Terminal’s Main Concourse at rush hour, and most observers cannot avoid being mesmerized by the seemingly effortless grace with which people move together, hurrying toward their trains as if tugged by invisible strings, all the while bobbing and weaving through a veritable gauntlet of potential collisions.

Unraveling the complexity of human crowd dynamics, and collective locomotion more generally, has attracted attention from a highly interdisciplinary community of scientists worldwide, with seminal contributions coming from fields as wide-ranging as animation, architecture, biology, computer science, engineering, physics, psychology, and sociology. Since the 1970’s, computational modeling has served as a common foundation that unites these efforts. Whether studying aggregations of mathematical point particles, crowds of people, or flocks of birds, there is often a similar arc: researchers propose a set of local rules governing individual locomotion, simulate interactions between individuals and their environment under those rules, and observe the resulting global patterns of collective behavior. While it has been a primarily theoretical endeavor, this modeling paradigm has made the challenge of understanding collective locomotion into a tractable problem, and it continues to be a fruitful, powerful, and well-utilized methodology.

Recently, there has also been a rising tide of data-oriented approaches, which aim to supplement idealized simulations with observations of animal collectives in the wild or

controlled experiments on individuals and small groups in laboratory settings. On the one hand, technological advances (such as improved computer vision algorithms, increasingly precise GPS, and stereoscopic videography) have facilitated the automatic extraction of locomotor trajectories from crowds, flocks, and schools containing hundreds or even thousands of members. On the other, researchers have devised clever ways to systematically probe the behavior of individual animals, such as measuring how birds, fish, insects, and even rats respond to biomimetic robots that resemble their conspecifics. Taken together, these data bolster the computational modeling approach by informing, calibrating, and validating the fundamental models of locomotor behavior that serve as building blocks for large-scale simulations. Now, rather than making *ad hoc* assumptions about the local interactions between neighbors, researchers can infer essential properties from observational data, or target them directly with experimental manipulations. This is especially important for biologists and cognitive scientists, whose primary goal is not to produce satisfactory visual reproductions of collective behavior (as in animation), but rather to understand the principle mechanisms underlying it, including the perception, action, and cognition of individual agents. While simulations can shed light on these mechanisms, hypotheses about them must ultimately be tested by measuring how organisms behave.

My dissertation is part of a continuing effort in this vein to understand how people in a crowd are visually coupled to their neighbors, and how this local coupling gives rise to global phenomena at the crowd level. I will present a series of experiments that quantify changes in participants' locomotor trajectories in response to systematic manipulations of both real and virtual crowds; I then use the results to critically test hypotheses about the

local interactions between neighbors. These experiments are some of the first of their kind, and have the potential not only to improve the fidelity of crowd behavior models, but also to reveal fundamental properties of the perceptual and cognitive mechanisms that give rise to collective behavior.

The remaining chapters of my dissertation are organized as follows: Chapter 2 reviews the existing body of literature on crowd dynamics and collective behavior, and provides a summary of theoretical and computational approaches to collective behavior. Chapter 3 draws attention to the lack of empirical data, and describes recent attempts to rectify this shortcoming with naturalistic observations and experiments. Chapters 4-7 describe a series of four experiments that are designed to test hypotheses about the visual coupling between neighbors in a crowd. Chapter 4 presents the first experiment, in which a participant walked with three confederates, whose speed and direction were systematically manipulated in order to determine the coupling between them and the participant. Chapter 5 introduces the concept and methodology of generating a virtual crowd, and validates this approach by comparing participant behavior with matched real and virtual neighbors. Chapter 6 describes a multi-part experiment with virtual crowds, in which a subset of virtual neighbors changed their speed or direction, while the density of the crowd and the number, distance, and eccentricity of the subset were systematically manipulated. Chapter 7 describes an experiment with small human crowds in a quasi-naturalistic setting, to see whether the findings from highly-controlled laboratory experiments will generalize to the real world. Finally, Chapter 8 concludes the dissertation by discussing the results and relating them to broader themes in collective behavior.

CHAPTER 2

Computational Models of Collective Locomotion

Collective behavior is a broad concept, encompassing everything from the division of labor in ant colonies to traffic patterns on the World Wide Web (Goldstone & Janssen, 2005; Gordon, 2014). The term's origins can be traced back at least as far as 1921, when Robert Park concluded his *Introduction to the Science of Sociology* with a chapter entitled "Collective Behavior," which he defined as "the behavior of individuals under the influence of an impulse that is common and collective, an impulse, in other words, that is the result of social interaction." I will focus specifically on a subset of collective behavior, namely *collective locomotion* and human crowd behavior, along with its counterparts in the animal kingdom such as flocks, schools, and swarms. For the present purposes, a *crowd* is defined as a group of people who occupy and move together in a shared space and time. Note that 'crowd' here is distinct from the more abstract, high-level notion referenced in words like 'crowdsourcing' (Estellés-Arolas & González-Ladrón-de-Guevara, 2012) and phrases like 'the wisdom of crowds' (Galton, 1907; Surowiecki, 2005). Note also that this definition excludes the interesting phenomena of human trail formation (Helbing, Keltsch & Molnár, 1997), because the people who create these trails, generally speaking, overlap in space but not in time.

Levels of Analysis

Collective locomotion can be studied at many levels, from the aerodynamics of a single bird to the 3D geometry of entire flocks. Clearly, these levels are related; after all, a flock is made up of individual birds flying together. And yet, the flock itself seems to exhibit emergent properties; at times, we lose sight of the birds, and the flock feels like a living, breathing organism all its own. The same is true in humans. Watch the crowd of

pilgrims circling the Kabaa in Mecca, and it is nearly impossible to focus on any one person; and yet, they are what give rise to the global counterclockwise flow that we so readily detect. How can these levels be related to one another? How can we connect birds and flocks, people and crowds, individuals and collectives?

The most cogent proposal for a solution comes from Sumpter and colleagues (2012), who published a paper entitled “The modelling cycle for collective animal behaviour.” In it, they distinguish between studies at two levels, local and global, and describe how these levels can be linked in two directions, from local to global global to local. Making sense of these distinctions is crucial to formulating research strategies to understand collective behavior, so it will be useful to review them in depth.

At the local level of analysis, researchers focus on the behavior of individual agents—be they birds, fish, particles, or people. The goal is to understand and ultimately predict how an individual moves in response to its immediate environment, which may include steering to goals, avoiding static and moving obstacles, and interacting with other nearby agents. Studies at this level can range from deciphering the fundamental perceptual mechanisms that underlie certain locomotor behaviors, such as how optic flow can be used to guide walking (Warren, Kay, Zosh, Duchon, & Sahuc, 2001) or flying (Srinivasan, Zhang, Lehrer, & Collett, 1996), to understanding the social factors that influence “when, where, what, and how to eat” in foraging species (Galef & Giraldeau, 2001). To draw an analogy with physical systems, analysis at the local level is equivalent to the study of classical mechanics, where the aim is to work out the kinematic equations of motion for bodies acted upon by a system of forces. Thus, individual agents in a collective occupy the same role as, say, particles in a gas.

The global level, on the other hand, is concerned with the large-scale properties of the collective as a whole, and is comparable to the study of classical thermodynamics, which deals with quantities such as temperature, pressure, entropy, and energy that are defined over entire systems. The goal is to understand and ultimately predict how large-scale properties of the collective as a whole will vary over time and react to changes in the environment. The detailed properties of individual agents are not considered. Global analyzes often consist of analyzing the overall motion pattern of the collective, which can range from disorderly chaos in swarms of insects (Kelley & Ouellette, 2013) to neatly organized translational and rotational flows in schools of fish (for an excellent overview, see Couzin, 2009).

While many studies consist of analysis at only one of these levels, there is also interest in connecting them. This can be done in two main ways, according to Sumpter. Most common are local-to-global analyzes, which seek to understand how simple interactions between agents at the local level gives rise to interesting collective phenomena at the global level. Continuing the physical analogy, local-to-global analyzes resemble statistical mechanics, which connect large-scale thermodynamic properties (e.g. heat) with the microscopic properties of interacting particles (e.g. velocity). Local-to-global analyzes are intimately connected with the notion of *self-organization* (Haken & Wunderlin, 1990; Couzin & Krause, 2003; Haken, 2006), which Camazine et al. (2001) describe as “a process in which the pattern at the global level of a system emerges solely from numerous interactions among the lower-level components of a system. Moreover, the rules specifying interactions among the system’s components are executed using only local information, without reference to the global pattern.” Large-scale agent-based simulations of collective

behavior reflect a local-to-global approach; their goal is to reproduce characteristic global patterns by modeling the behavior of simple agents and their interactions. This effort provides insight into the fundamental processes of collective locomotion and self-organization, and can help explain how complex behavior can arise from seemingly simple agents.

Less common, but equally important, are global-to-local approaches. Here, the goal is to observe patterns at the global level and use them to infer properties of agents and their interactions at the local level. Regularities in the global configuration or dynamics of the collective as a whole place constraints on the kinds of rules and models that can be proposed to characterize behaviors at the local level. For example, observing an overall measure of coordination between birds in a flock can provide clues about the local interactions between individuals, such as how many neighbors each bird responds to (Ballerini et al., 2008) or how information about a predator propagates throughout the flock (Cavagna et al., 2010), although there is a limit to what can be learned, because different local rules can give rise to identical global patterns (Vicsek & Zafeiris, 2012).

In order to truly understand collective locomotion, then, research must be conducted at both the local and global levels, and the relations between them. In the next two sections, I will review the major contributions that have been made to the study of collective locomotion, beginning with computational models and followed by the collection of empirical data. Later, at the conclusion of this chapter, I will return to these themes of levels and methods of analysis, to consider how the study of human crowd behavior in particular can be improved.

Computational Models

Over the last few decades, computational models have proliferated and become the dominant tool to study collective and crowd behavior. There are now literally hundreds of models on offer. In order to understand them, it will be useful to sketch the space of possible models by noting some important divisions that characterize it, most importantly between macroscopic and microscopic models, and between rule-based and force-based models. To borrow terminology from Mar (1982), we are chiefly interested in the theoretical or algorithmic differences between these models, rather than their computational implementation.

Macroscopic models describe the overall behavior of the entire collective and neglect individuals; that is, they focus on behavior at the global level. They make use of aggregate quantities like densities, velocity distributions, and flows (Hughes 2003). Macroscopic models are often based on analogies with canonical physical systems, such as fluids (Henderson 1974, Helbing 1992) or magnetically-charged particles (Okazaki & Matsushita 1993). These physical systems are understood in great detail, and offer a crucial advantage of yielding analytical solutions in some cases.

Microscopic models, in contrast, take the individual to be paramount. They model the local behavior of individuals and their interactions. Taken together, these interactions can lead to complex global patterns through the process of self-organization (Camazine, Deneubourg, Franks, Sneyd, Theraulaz, & Bonabeau, 2001). This has a number of benefits. In addition to the broad heterogeneity that is possible with some macroscopic models (e.g. those with multiple interacting fluids), microscopic models can also introduce heterogeneity in individual behavior—in state quantities like speed, in access to perceptual

information, and in the rules or equations governing behavior. Microscopic models also allow domain experts, such as cognitive scientists or ecologists, to use knowledge about the specific perceptual or cognitive capacities of the given species (Hildenbrandt, Carere, & Hemelrijk, 2010; Rio, Rhea, & Warren, 2014), which could result in more cognitively-plausible models.

Microscopic models can be further broken down into rule-based and force-based models. In rule-based models, agents behave by following a set of rules, which often take the form of a series of *if-then* statements: *if* some condition is met, *then* the agent should take some action. These models are often discrete, with behavioral changes occurring only when a certain set of criteria are satisfied. Force-based models, by contrast, are continuous, with behavior changing constantly in response to changes in the environment. They are often defined in the language of differential equations, and involve ‘behavioral’ (non-physical) forces like attraction, repulsion, and synchronization (Schellinck & White, 2011).

Macroscopic models

In macroscopic modeling, the goal is to describe the global behavior of the entire crowd, rather than the local behavior of its pedestrians. Typically, quantities like velocity distributions, densities, and flows are used to describe the crowd’s motion. The classic example is fluid dynamics models, which are based on analogies with the motion of particles in a fluid (liquid or gaseous). The first such model was proposed by Henderson (1974). In previous work, which will be described in Chapter 3, Henderson (1971; 1972) had observed that the velocity distribution of pedestrians fit almost exactly the Maxwell-Boltzmann distribution, which governs the motion of particles in a gas. He used this

observation to justify modeling the entire crowd as a gas, with pedestrians acting as individual particles.

Henderson's model (1971) began with a number of assumptions that have been questioned by other researchers (e.g. Helbing, 1992; Schadschneider, Klingsch, Klüpfel, Rogsch, & Seyfried, 2009). These assumptions are necessary to take advantage of the analogy with a fluid system. For example, Henderson (1971) assumed the conservation of momentum and energy, which applies to gas particles. But pedestrians in a crowd do not seem to obey these laws; for example, a pedestrian can decide to turn around and walk the way she came (if, for example, she suddenly realizes she forgot her keys). This violates the conservation of momentum because her momentum begins at some value p , then drops to zero as she comes to a stop, and then reaches a value of $-p$ when she walks in the opposite direction. This is not possible for inanimate physical systems like gas particles or billiard balls, which can only change direction (and, more generally, momentum) if they are acted on by some other force, such as gravity or interaction with another object. A pedestrian in Henderson's (1971) model can only change direction by colliding with another pedestrian—this is unrealistic, especially in low-density crowds. Moreover, Henderson (1971) assumed that each particle is statistically independent from every other; this precludes the effect of groups (which, as we will see later, are a major component of crowds) and other inter-pedestrian interactions.

Still, Henderson's (1971) model is a useful approximation in many situations. Although it may fail to capture some of the details of individual pedestrians (e.g. suddenly reversing direction), it captures the overall behavior of the crowd remarkably well. One of the most interesting properties of his model is that it allows for phase transitions between

gaseous phases (essentially random motion) and liquid phases (which flow in a particular direction with a particular speed). These phase transitions can result from conditions in the environment, such as a constriction or bottleneck that tends to compress the gaseous state into the liquid state, or vice versa. A class of students exiting a lecture hall may begin as a liquid flow, but when they reach the Main Green, they may fan out into a more gaseous state. Perhaps most importantly, Henderson (1974) makes several quantitative predictions, such as the change in speed that should result from a crowd moving through a constriction of a given width. Unfortunately, to date no one has tested these predictions against published data.

Subsequent work on fluid dynamics models has improved on Henderson's (1971) original formulation. Henderson (1974) allowed for different modes of locomotion in his model, such as standing still, walking, or running, but more recent models (Helbing, 1992; Hughes, 2003) have allowed for pedestrians to have different desired directions of motion—something that would be impossible in a classical fluid. Thus, a crowd can be thought of as made up of a number of sub-fluids, each of which has a desired position and flow through and around one another to reach it. The details of these sub-fluids can become quite complex, as in Rose and Gamble (1994), where pedestrians tire (and hence slow down) as they walk.

With these improvements, macroscopic models possess a number of attractive features. First, they are often analytically tractable in ways that microscopic models are not. This is of particular interest to physicists and mathematicians, who are more interested in the collective phenomena exhibited by crowd motion than by accurately describing the behavior of any individual pedestrian. As Hughes (2003) writes, “[Microscopic] simulation

forms a very useful tool for practical applications, but as a research tool, it suffers from a lack of analytical tractability that makes deriving general results difficult.” A second, and closely related, advantage is that there have been centuries of work in analogous physical systems, which means that much of the difficult analytical work has already been carried out. Rather than having to work out the solution to a difficult set of equations, the challenge for the macroscopic modeler is to choose an appropriate model of fluid dynamics and check whether and under what circumstances its assumptions hold.

Despite these positive attributes, macroscopic models have several important limitations. By neglecting individual pedestrians, they fail to capture the specifics of crowd motion. Moreover, because they deal with aggregate and statistical quantities, they are difficult to interpret when applied to low-density scenarios (Hughes 2003). Large numbers of pedestrians are needed to ‘smear out’ individual variability and generate stable statistical quantities, just as large numbers of particles are needed in order to do statistical thermodynamics (Gibbs 1960). For this reason, macroscopic models have tended to focus on high-density scenarios, such as emergency evacuations (Colombo, Goatin, & Rosini, 2010), where they still widely used, but are less common in models of everyday locomotion, which typically involve low to medium densities (Moussaïd, Perozo, Garnier, Helbing, & Theraulaz, 2010). Moreover, they cannot be used to study isolated pedestrians or pedestrians in small groups (Moussaïd et al., 2010; Bonneaud, Rio, Chevaillier, & Warren, 2010), which makes them difficult to test experimentally.

Mesoscopic hybrids

The distinction between purely macroscopic and purely microscopic models is a rough one. There is a middle ground, the mesoscopic level, which Hoogendoorn and Bovy (2000) define as models that “do not distinguish each pedestrian individually but rather describe the pedestrians and their characteristics in more aggregate terms (e.g. velocity distributions). However, pedestrian behavior is described at a microscopic level, albeit only in terms of probabilities.” Although there are some models that completely ignore individual pedestrians (i.e. Henderson’s (1971)), and many that do not explicitly model behavior above the level of the individual pedestrian, there are numerous models in between, which either use macroscopic models to derive the microscopic behavior of individual pedestrians, or use the microscopic behavior of individual pedestrians to drive the macroscopic behavior of the crowd as a whole. Let us examine an example of each of these approaches in turn.

Helbing (1992) presents a fluid dynamics model that incorporates microscopic effects at the individual level. In addition to defining multiple sub-fluids, which contain pedestrians with different intended directions and speeds, he also includes several terms that are specific to individual pedestrians. In Henderson’s (1971) model, pedestrians are approximated with inanimate particles, so the only possible interaction is collision. In Helbing’s (1992) model, by contrast, individuals are modeled as perceiving and acting agents, and their interactions—both physical and intentional—are explicitly modeled. Each pedestrian has a desired velocity that they tend toward, but this tendency is not absolute, as in Henderson’s (1971) model; rather than the conservation of momentum and energy, there is a *tendency to conserve* momentum and energy. A pedestrian might not be able to

reach his desired velocity if, for example, other slower-moving pedestrians are nearby. Instead of colliding with these pedestrians, as in Henderson's (1971) model, a pedestrian in Helbing's (1992) model will try to pass them. If passing is not possible, as in an extremely dense crowd, the pedestrian will slow down. Moreover, this passing behavior is anisotropic, meaning it can be biased to the left or the right, reflecting cultural factors. It is also probabilistic, meaning there is a certain probability that a pedestrian will pass on the right, as well as a non-zero probability of passing on the left. Still, despite these microscopic effects, Helbing's (1992) model can still be considered largely macroscopic. The principal quantities of this model are the velocity distributions and densities of the various sub-fluids that make up the crowd. Individual pedestrians are only distinguished when they are engaged in an interaction with pedestrians from another sub-fluid; after the interaction is over, they are no longer distinguished as individuals, and the focus is once again on densities. Although this model does not describe individual pedestrians at every time step, it is able to capture a much wider range of phenomena than a purely macroscopic model like Henderson's (1971), such as the development of lanes, stop-and-go waves, and pedestrian jams.

Hoogendoorn and Bovy (2000) take the opposite approach. Rather than using microscopic behavior to derive macroscopic quantities, they use macroscopic behavior to infer microscopic behavior. They begin with a gas-kinetic model that is similar in most respects to Henderson's, but they use the macroscopic behavior of the system as a whole to drive the microscopic behavior of individuals, using what they call a "particle discretization method." As shown in Figure 2-1, they divide the simulation space into discrete cells, compute the mean velocity vector of all agents in each cell, and use this mean

vector as input to the macroscopic gas-kinetic equations, which then drive the movements of individual agents.

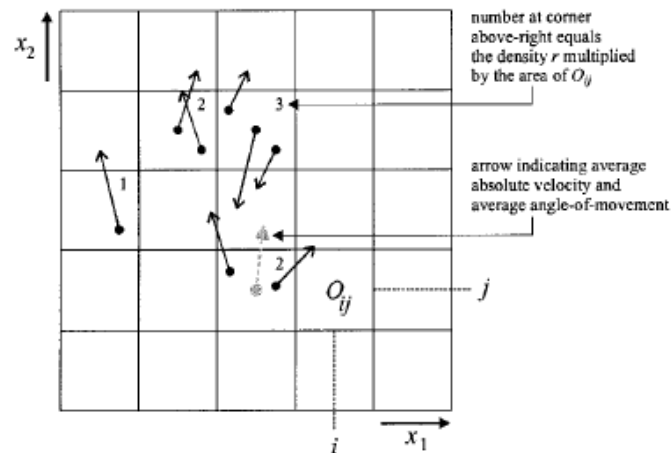


Figure 2-1: Discretization of simulation space. For each cell, the mean velocity is computed, which is then used as input to the macroscopic gas-kinetic equations. From Hoogendorn & Boovy (2000).

This method thus “upgrades” a mesoscopic model to a microscopic model. What is the advantage of a model that has been upgraded in this way? Hoogendorn and Boovy (2000) note that their approach is more computationally efficient when simulating large crowds. However, this is less of an issue now than it was when their paper was first published 14 years ago; the difficulty of animating large crowds today comes from graphical constraints on generating sufficiently detailed pedestrians, not from computing their motions. And this slight edge in computational efficiency is far outweighed by the disadvantages that come with such an ‘upgrade.’ On the one hand, mesoscopic models lack the analytical tractability and analogy with physical systems that makes macroscopic models so appealing. On the other, they lose the plausibility of microscopic models. It is easy to imagine how a set of rules or social forces could be cognitively implemented; it is

nearly impossible to imagine how that would work for a mesoscopic model. Do individuals somehow compute the mean velocity of groups of neighbors? And even if they could, how could this information make its way to some large-scale gas-kinetic equation, and then back down to drive their behavior? Given these difficulties, and without a compelling reason to prefer them over microscopic or macroscopic models, explicitly mesoscopic models like that of Hoogendorn and Boovy (2000) remain relatively uncommon. Helbing, for example, abandoned his mesoscopic model (Helbing 1992) in favor of a microscopic model, the social force model (Helbing & Molnar, 1995; Helbing, Farkas, & Vicsek, 2000; Moussaïd, Helbing, Garnier, Johansson, Combe, & Theraulaz, 2009).

Microscopic models

Rule-based models.

In rule-based models, an agent perceives its environment, uses that information to determine whether certain conditions have been met, and then acts according to a set of rules. A classical way to formulate these rules is in the form of *if-then* statements; *if* condition X is met, *then* perform action Y. The rules can range from extremely simple to highly complex, incorporating high-level cognitive behaviors like planning and navigation. One of the first, and most influential, of these models was introduced in 1987 by Craig Reynolds.

Reynolds' (1987, 1999) model contains three basic rules that each agent ('boid') follows. These generate what he calls "happy aimless flocking," and by providing a global direction or position that all birds flock to, scripted routines can be generated. Reynolds' three basic rules have come to influence virtually all microscopic models that followed, and they are: (1) collision avoidance—agents steer away from nearby flockmates; (2)

velocity matching—agents attempt to match the velocity of nearby flockmates, and (3) flock centering—agents steer toward the center of nearby flockmates. There are several details regarding each of these components, their combination, and the definition of ‘nearby’ that are important and warrant further attention.

The first rule is collision avoidance (Reynolds 1987) or separation (Reynolds 1999): agents avoid colliding with nearby flockmates, by steering away from them. This direction is a vector sum, computed from the vectors which point from nearby flockmates to the agent. It is weighted, such that closer agents have a greater influence on direction; Reynolds (1999) uses a factor of $1/r$, where r is the distance from each flockmate to the agent. A schematic diagram in Figure 2-2 illustrates this model of collision avoidance, showing the vectors (green lines) pointing from nearby flockmates (blue triangles) to the agent (green triangle) as well as the resulting steering direction (red arrow).




Figure 2-2: Schematic diagrams of the components of Reynolds' (1987) 'boids' model. (a) Collision avoidance: agents steer away from nearby flockmate; (b) velocity matching: agents match the velocity of nearby flockmates; (c) flock centering: Agents steer toward the center of nearby flockmates. From Reynolds (1999).

A second, related rule is velocity matching (Reynolds 1987) or alignment (Reynolds 1999): agents steer toward the average heading of nearby flockmates, and adjust their speed so that it matches the average speed of nearby flockmates. Reynolds (1987) writes that velocity matching complements collision avoidance. “It is a predictive version

of collision avoidance: if the boid does a good job of matching velocity with its neighbors, it is unlikely that it will collide with any of them any time soon. ... Static collision avoidance serves to establish the minimum required separation distance; velocity matching tends to maintain it.” Figure 2-2 depicts a schematic diagram of velocity matching. It shows the heading (blue lines) of nearby flockmates (blue triangles) as well as the current heading (green line) of the agent (green triangle). The blue line attached to the agent represents its desired heading, while the red line represents the vector difference of the agent’s current heading and its desired heading.

The third component of Reynolds’s model is a rule for flock centering (Reynolds 1987) or cohesion (Reynolds 1999): agents steer toward the average position of nearby flockmates; namely, toward their center of mass. This component will be particularly active at the edges of the flock, where there will be a strong anisotropy in the positions of nearby flockmates; deep inside the flock, flockmates will be distributed roughly homogeneously (Reynolds 1987). It also gives the flock a degree of flexibility, such as the ability to split up and rejoin after obstacles. Figure 2-2 depicts a schematic diagram of flock centering. It shows the position of nearby flockmates (blue triangles) as well as their center of mass (green dot), which the agent (green triangle) moves in the direction of (red arrow).

In order to generate behavior, however, these three rules must be combined with some mechanism of integrating their influences. Each rule effectively results in a desired velocity, which Reynolds (1987) terms acceleration requests. What if different components ‘request’ accelerations in different directions? Reynolds’ model (1987) relies on a hierarchy, which he terms prioritized acceleration allocation, whereby the three rules are arranged in order of descending priority. Collision avoidance is first, and highest-priority,

followed by alignment, and then cohesion. In addition to a desired velocity, each component also produces a magnitude ranging from 0 to 1, which specifies the intensity of its request. If a nearby flockmate is dangerously close, collision avoidance would have a magnitude of 1; if the nearby flockmate was still quite far away, the magnitude would be close to 0. The total magnitude of all three forces cannot exceed 1. Thus, according to Reynolds (1999), “by adjusting their magnitude, higher priority behaviors could decide whether or not to leave any steering force for use by lower priority behaviors.” In essence, if a higher priority behavior has a high magnitude, there will be no more acceleration left over to be allocated.

Finally, it must be noted that while the ‘boids’ model is simple in principle, its actual implementation (Reynolds 1987) contains additional constraints that help produce realistic behavior. These include assumptions such as requiring a goal; setting a maximum speed and acceleration; choosing a spherically, exponentially-decaying neighborhood of influence; and including a ‘blind spot’ behind each agent.

Even with these additional features, however, Reynolds’ (1987) model does not require its agents to have detailed (and unrealistic) knowledge about the entire flock or its behavior, and it was the first individual-based model of collective locomotion. His algorithm has served as the basis for over two decades of advances in computer animation. It also inspired work in the animal literature, which embraced the rule-based approach. Huth and Wissel (1992) modeled fish schools using four components. Fish swim away from neighbors if they come within a small, fixed radius ($r < r1$), swim parallel to neighbors if they come within some fixed distance interval ($r1 < r < r2$), swim toward neighbors if they are within another fixed distance interval ($r2 < r < r3$), and swim randomly (‘search’)

if there are no fish within a maximum detectable distance ($r > r_3$). These radii are shown in Figure 2-3. The first three components of Huth and Wissel's (1992) model are analogous to Reynolds' collision avoidance, velocity matching, and flock centering, respectively. Once again, these simple rules are able to generate realistic displays of schooling fish.




Figure 2-3: Schematic diagram of an individual agent from Huth & Wissel's (1992) model of fish schooling. Dots ($r < r_1$) represent the 'repulsion' region; vertical bars ($r_1 \leq r < r_2$) represent the 'parallel' region; slanted bars ($r > r_3$) represent the 'attraction' region. From Huth & Wissel (1992).

While simple models like Reynolds' (1987) or Huth and Wissel's (1992) might be able to produce realistic displays for computer animation, some researchers have proposed rule structures that incorporate higher-level cognitive factors like planning, navigation, and decision-making. Paris and Donikian's (2009) activity-driven populace is one such model. It is hierarchical, and contains both reactive and predictive components. There is a model of decision-making, so that agents can decide which of multiple goals they should move toward. There is a model of memory, so that agents can learn the layout of their environment. There is a path-planning model, which computes the optimal path for each agent given its current goals and known obstacles. The agent's position is then compared

against this planned path, to make sure that it stays on track. This model can reproduce complex series of actions, such as a pedestrian at a train station walking toward the ticket window, waiting in line, purchasing the ticket, and walking to the train to be boarded. Similar hierarchical models exist for other specific situations, like escaping a predator (Funge, Tu, & Terzopoulos, 1999) or dealing with uncertain environments (Yu & Terzopoulos 2007).

In recent work, Moussaïd, Helbing, and Theraulaz (2011) have proposed a new model based on simple rules that they call heuristics. It also features a synthetic vision component, which uses first-person information about each agent's environment to drive behavior. Each pedestrian (hereafter, 'the agent') has a gaze direction \vec{H}_i , and a field of view which extends $\pm \phi$ degrees to the left and right of \vec{H}_i . For all possible directions α within this field of view, the model computes the distance to the first collision $f(\alpha)$ if the agent moved in that direction at their preferred speed, taking into account the other pedestrians' walking speed and direction (i.e. predicting their trajectories). If no collision is expected in direction α , $f(\alpha)$ is set to a default value d_{max} .

Along with this vision model, Moussaïd et al. (2011) use two heuristics to drive behavior. The first concerns steering to goals while avoiding obstacles, and they state it in words as follows: "[The agent] chooses the direction α_{des} that allows the most direct path to destination point O_i , taking into account the presence of obstacles." This chosen direction, α_{des} , is computed by minimizing the distance to the destination $d(\alpha)$ over all angles α in the field of view, where $d(\alpha)$ is given by:

$$d(\alpha) = d_{max}^2 + f(\alpha)^2 - 2d_{max}f(\alpha)\cos(\alpha_0 - \alpha) \quad (1)$$

where α_0 is the direction of the goal.

The second heuristic can be stated as: “A pedestrian maintains a distance from the first obstacle in the chosen walking direction that ensures a time to collision of at least τ .”

This determines the desired walking speed of the agent, v_{des} , which is defined by:

$$v_{des} = \min\left(v_i^0, \frac{d_h}{\tau}\right) \quad (2)$$

where v_i^0 is the agent’s comfortable walking speed (a parameter of the model), d_h is the distance between the agent and the first obstacle in the desired direction α_{des} , and τ is a parameter which represents the minimum time to contact that is allowed. Thus, if a collision is imminent, the agent will slow down; if not, it will continue to move at its comfortable speed.

Force-based models.

An alternative to the rule-based approach is the force-based approach. This framework borrows concepts from physics—most often from classical mechanics—to describe pedestrian trajectories. Rather than physical forces like gravity and electromagnetism, these forces are ‘social forces,’ reflecting the various drives an individual pedestrian experiences. As Helbing and Molnar (1995) write, “The social force is not exerted by the environment on a pedestrian’s body. It is rather a quantity that describes the concrete *motivation to act*,” (emphasis in original). Typical forces include things like reaching a preferred velocity, attraction to goals, and repulsion from obstacles. These forces can be combined in various ways. The most typical is superposition, because this is most common in physical systems. Just like the net force due to gravity on a body is the sum of the individual forces of nearby bodies, the social force on a pedestrian is the sum of the individual forces from nearby pedestrians, goals, and obstacles.

The first, and most influential, of these models was introduced by Helbing and Molnar (1995). It is made up of several components which act to change the preferred (or desired) velocity, $\vec{w}_\alpha(t)$ of each model pedestrian (hereafter, ‘the agent’) α . Note that this is a vector quantity, so it involves both speed and direction, and it is time-varying. In the absence of any obstacles or other pedestrians, each agent α has a desired speed, $v_\alpha^0(t)$, and direction, $\vec{e}_\alpha(t)$. The desired direction points toward the agent’s goal, and is defined as the unit vector between the agent’s current position, $\vec{r}_\alpha(t)$, and the position of the nearest goal k , \vec{r}_α^k .

$$\vec{e}_\alpha(t) = \frac{\vec{r}_\alpha^k - \vec{r}_\alpha(t)}{\|\vec{r}_\alpha^k - \vec{r}_\alpha(t)\|} \quad (3)$$

If the agent’s current velocity differs from the preferred velocity defined by the desired speed and direction, there will be a net (social) force which causes the agent to accelerate.

This force, \vec{F}_α^0 , takes the form:

$$\vec{F}_\alpha^0(\vec{v}_\alpha, \vec{v}_\alpha^0 \vec{e}_\alpha) = \frac{1}{\tau_\alpha} (\vec{v}_\alpha^0 \vec{e}_\alpha - \vec{v}_\alpha) \quad (4)$$

where \vec{v}_α is the agent’s current velocity, $\vec{v}_\alpha^0 \vec{e}_\alpha$ is the agent’s desired velocity, and τ_α is a relaxation time term, which modulates how differences in velocity produce acceleration. If τ_α is very large, small differences in velocity have a negligible effect on acceleration; by contrast, if τ_α is very small, small differences in velocity will be amplified.

Of course, there are few situations in which a pedestrian is in a totally open environment that is devoid of other pedestrians. Thus, the social force model (Helbing & Molnar 1995) adds repulsive forces, which reflect the fact that pedestrians avoid colliding with other pedestrians, as well as environmental barriers like walls and. The first of these concerns other pedestrians, and takes the form:

$$\vec{F}_{\alpha\beta}(\vec{r}_{\alpha\beta}) = -\nabla_{\vec{r}_{\alpha\beta}} V_{\alpha\beta}[b(\vec{r}_{\alpha\beta})] \quad (5)$$

where $\vec{r}_{\alpha\beta}$ is the vector from agent α to nearby pedestrians β , $\nabla_{\vec{r}_{\alpha\beta}}$ is the gradient of that vector, and $V_{\alpha\beta}$ is a repulsive potential, which is monotonic decreasing and elliptical (b denotes the semiminor axis of the ellipse). Thus, nearby pedestrians have a larger influence than more distant ones, and the shape of this ‘neighborhood of influence’ is elliptical; in particular, it is elongated along the direction of motion, such that pedestrians directly in front of the agent are more repulsive than those on the side.

Interestingly, in Helbing & Molnar’s (1995) original formulation of the social force model, there is no explicit mention of point-like environmental obstacles, such as trash cans or signposts. Presumably, these might be handled in the same way as pedestrians, perhaps with a different potential function. Helbing & Molnar (1995) do describe a force concerning interaction with extended obstacles, such as walls. This takes the form:

$$\vec{F}_{\alpha B}(\vec{r}_{\alpha B}) = -\nabla_{\vec{r}_{\alpha B}} U_{\alpha B}(\|\vec{r}_{\alpha B}\|) \quad (6)$$

where $\vec{r}_{\alpha B}$ is the unit vector from the agent to the nearest point on the border, B , and $U_{\alpha B}$ is a monotonic and decreasing potential which depends on the distance from the agent to the point B , $\|\vec{r}_{\alpha B}\|$.

In addition to these repulsive effects, Helbing & Molnar (1995) also include an attractive force, which can involve either pedestrians or inanimate objects. Its form is similar to the repulsive forces, and is defined by:

$$\vec{F}_{\alpha i}(\vec{r}_{\alpha i}, t) = -\nabla_{\vec{r}_{\alpha i}} W_{\alpha i}(\|\vec{r}_{\alpha i}\|, t) \quad (7)$$

where $\vec{r}_{\alpha i}$ is the unit vector between the agent and the object i , and $W_{\alpha i}$ is a monotonic increasing attractive potential. Note that the attractiveness of an object, $\|\vec{f}_{\alpha i}\|$, can be a function of time, which can reflect decreasing interest, as in a street musician.

The net social force (\vec{F}_{α}) on an agent is the sum of these various forces: maintaining a preferred velocity (\vec{F}_{α}^0), attraction ($\vec{F}_{\alpha i}$), and repulsion (pedestrians: $\vec{F}_{\alpha\beta}$; barriers: $\vec{F}_{\alpha B}$), and can be written as:

$$\vec{F}_{\alpha}(t) = \vec{F}_{\alpha}^0 + \sum_{\beta} \vec{F}_{\alpha\beta} + \sum_B \vec{F}_{\alpha B} + \sum_i \vec{F}_{\alpha i} \quad (8)$$

In addition, there is a noise term, to account for random variations in the behavior of pedestrians, and a limit for the maximum possible acceleration.

Helbing and Molnar (1995) use their model to demonstrate a variety of qualitative and quantitative phenomena. These include the formation of lanes and counterflow at bottlenecks. Lanes form in a corridor when pedestrian density exceeds a critical density, and the number of lanes scales linearly with the width of the corridor. Counterflow at bottlenecks arises when two groups of pedestrians try to pass through a constriction in opposite directions. Because of variability in preferred velocity and attraction to goals, pedestrians will be more or less motivated to move through the constriction. Several pedestrians from one side will pass through at a time, until the supply of motivated pedestrians is exhausted; then, more highly-motivated pedestrians from the other side will begin moving in the opposite direction, and the cycle will continue until all pedestrians have cleared the bottleneck.

Social forces models, such as the one introduced by Helbing & Molnar (1995), have been criticized because while they seem to do a good job of capturing the overall behavior of crowds, “they tend to create simulations that look more like particle animation than

human movement,” (Pechano, Allbeck, & Badler, 2007); in other words, behavior may seem realistic at the global level, but unrealistic at the local level. In addition, they cannot natively deal with high-density crowds, where social forces like attraction to goals and repulsion from obstacles are overwhelmed by physical forces between very tightly-packed agents. Helbing, Farkas, and Vicsek (2000) added physical forces to the social force model of Helbing and Molnar (1995) to reflect these high-density effects. In order to make sure that two pedestrians could not occupy the same space, they added a body force, which resists compression from nearby agents, and a sliding friction force, which reflects the change in velocity due to the direct pressure applied by a pedestrian making contact. Both of these forces only apply when pedestrians are actually in contact with one another; otherwise, they have no effect.

The social force model, and models like it, have become so ubiquitous that they have been formalized as the *attraction-repulsion framework*. Schellink and White (2011) define the three essential properties of this kind of individual-based modeling: the movement of an individual at a particular time t is defined by:

1. the attraction and repulsion equations (along with equation parameters) associated with individuals in the aggregate.
2. the inputs to the attraction and repulsion equations at that time, t , which themselves depend on the state of the environment at that time—in particular the position and velocity of other individuals in the environment.

3. the outputs of the attraction and repulsion equations, possibly combined with additional reaction equations, which determine how the position and velocity of the individual will change from time t to time $t + 1$.

Within this framework, two quite different approaches have followed from the initial work of Helbing and Molnar (1995) and others like them. One push has come from theoretical physicists, who look at abstract, highly general models to determine things like the minimum set of rules needed to generate stable flocking and how transitions between states can be produced by changing parameters or environmental conditions. The other push has come from animal behavior, and more recently cognitive science, to develop more specialized models, which deal with particulars of a given animal species—ants, starlings, or humans.

The first seminal contribution from the theoretical physics camp came from Vicsek and colleagues (Vicsek, 1995; Czirok, Stanley, & Vicsek, 1997; Czirok & Vicsek 2000), who introduced the concept of self-propelled particles. These are massless particles which move with a constant velocity and align themselves with the velocity of their nearest neighbors. They showed that even this remarkably simple model (which, in essence, contains the velocity matching component of Reynolds (1987) without attraction or collision avoidance) could produce stable aggregations for some values of particle density and noise parameters. Thus, they showed that alignment is sufficient for aggregation, under certain conditions. Couzin et al. (2002) used three rules: maintaining a minimum distance, attraction, and alignment, and showed how this could produce four distinct forms of aggregation, which they labeled as swarm (high cohesion, low polarization), torus (individuals rotating around an empty center), dynamic parallel group (more mobile than

swarm or torus), and highly parallel group (highly aligned arrangement, rectilinear movement). An example of each type is shown in Figure 2-4.




Figure 2-4: Four different types of aggregations, formed by different sets of parameter values in a single model. Clockwise from top right: (A) swarm, (B) torus, (C) dynamic parallel group, and (D) highly parallel group. From Couzin et al. (2002).

More interesting, Couzin et al. (2002) also show how changes between these forms result from changes in model parameters, namely the size of the zones corresponding to repulsion and attraction. This result highlights how different large-scale behaviors can result from relatively small changes to the individual rules governing each agent. This could explain rapid changes in animal behavior, such as a sudden rearrangement in response to the detection of a predator. For example, Ebeling & Schimansky-Geier (2008) show how transitions between a translational (corresponding to the ‘parallel groups’ in Couzin et al., 2002) and rotational (‘torus’) mode can occur as a result of changes in

parameters. Both Couzin et al. (2002) and Ebeling and Schimansky-Geier (2008) also show hysteresis effects; that is, the threshold of the parameter value for changing from one mode to another is different depending on which mode the system is currently in. This is a characteristic property of nonlinear systems (Haken, Kelson, & Bunz, 1985).

In contrast to these highly general models, work in animal behavior has begun to focus on specialized, species-specific models. One of the most prominent examples comes from recent studies of starlings. Starlings form massive and highly-coordinated flocks with tens of thousands of individuals, but until recently it has been difficult to track individual birds. Recent advances in stereo photography (Ballerini et al., 2008; Cavagna, Cimorelli, Giardina, Orlandi, Parisi, Procaccini, Santagati, & Stefanini, 2008) have made it possible to do this, renewing interest in starlings as a model species. Hildenbrandt, Carere, and Hemelrijk (2010) developed a highly specialized model for starlings. In addition to having biologically-derived parameter values (e.g. mass, wing span, average and maximum speed, etc.), the model also includes components which reflect flight characteristics specific to starlings. There are components for lift and drag, as well as banking in and out of turns. These components, far from being superfluous, help to explain the dynamics of turning in starlings—starlings do not form a rigid body, but rather exchange positions within the flock as they turn. This falls naturally out of a model which includes banking, because banking reduces speed and hence changes a bird's relative position in the flock.

This kind of highly-specialized approach brings us to cognitive science. There is a newfound interest in using findings from cognitive science to inspire and constrain models of crowd behavior that reflect human capabilities (Goldstone & Gureckis, 2009). Perhaps the most important work has come from Moussaid et al. (2010), who include these kinds

of human-specific components in a version of Helbing's social force model. Moussaïd et al. (2010) extend the social force model by including social interactions among people walking in groups. They introduce a new force (\vec{f}_i^{group}) that represents these interactions, and is combined with the forces for attraction, repulsion, and maintaining a preferred velocity defined by Helbing & Molnar (1995). They also add a synthetic vision model, which purports to reflect something about the perception-action capabilities of real pedestrians.

Each model pedestrian i (hereafter, 'the agent') has a gazing direction vector, \vec{H}_i , and a field of view which extends $\pm \phi$ degrees to the left and right of the gazing direction. The agent turns its head by an angle α_i so that the center of mass of the group, c_i , is contained within their field of view, as shown in Figure 2-5. It is assumed that greater values of α_i are less comfortable, so the agent adjusts its position to reduce head rotation; this is the first term, \vec{f}_i^{vis} , that contributes to \vec{f}_i^{group} , and can be written:

$$\vec{f}_i^{vis} = -\beta_1 \alpha_i \vec{V}_i \quad (9)$$

where β_1 is a parameter describing the strength of interaction between the agent and other group members, and \vec{V}_i is the agent's velocity vector.




Figure 2-5: Variables in the social interaction model of Moussaïd et al. (2010). Each agent has a gaze direction, H_i , and a field of view which is defined as $\pm \phi$ to the left and right of H_i . c_i represents the center of mass of the group, and agents rotate their gaze direction by an angle α such that c_i is contained in their field of view. From Moussaïd et al. (2010).

Agents also maintain a certain distance from the group's center of mass, which increases with group size. This results in a second force, \vec{f}_i^{att} , which can be written:

$$\vec{f}_i^{att} = -q_A \beta_2 \vec{U}_i \quad (10)$$

where β_2 is a parameter describing the strength of attraction, \vec{U}_i is the unit vector pointing from the agent to the center of mass, and q_A is a binary variable with $q_A = 1$ if the distance between the agent and the center of mass exceeds a threshold value $\left(\frac{[N-1]}{2}\right)$, where N is the number of pedestrians in the group), and 0 otherwise.

The third component is a repulsion force, \vec{f}_i^{rep} , which prevents group members from overlapping. It is defined by:

$$\vec{f}_i^{rep} = \sum_k q_R \beta_3 \vec{W}_{ik} \quad (11)$$

where \vec{W}_{ik} is the unit vector from the agent to each group member k , β_3 is a parameter describing the strength of repulsion, and q_R is a binary variable with $q_R = 1$ if the agent and group member k overlap one another, and 0 otherwise.

The overall social interaction, then, is simply the sum of these three forces:

$$\vec{f}_i^{group} = \vec{f}_i^{vis} + \vec{f}_i^{att} + \vec{f}_i^{rep} \quad (12)$$

Moussaid et al. (2010) show that their model can reproduce a pattern observed from observational data; namely that the angle between group members increases with increasing density, as shown in Figure 2-6. This is not particularly surprising, however, because it is a relatively low-dimensional pattern, and their model is quite complicated: it has 3 components and 5 free parameters, in addition to the components and parameters of Helbing & Molnar's (1995) original model.




Figure 2-6: Results of Moussaid et al. (2010). Their social interaction model reproduces the pattern observed in data, that the angle between agents in a group increases with increasing density (top: low density, bottom: high density). From Moussaid et al. (2010).

The heuristics-based model of Moussaid et al. (2011) is simpler than the social force model, but it can reproduce many of the same patterns of behavior, such as lane formation and stop-and-go waves. By taking a first-person view, it also deals with situations where environmental obstacles, including other pedestrians, are hidden or out of view, without requiring to additional components or assumptions. Still, it is difficult to call this model ‘cognitive.’ It uses simple heuristics, which can be easily stated in a sentence, but each contains a number of assumptions that may or may not be plausible in humans. For example, the model assumes that agents are able to estimate the current position and velocity of their neighbors accurately, and use them to estimate their future positions. These

are empirical claims that can and must be tested before the model of Moussaïd et al. (2011) can be properly referred to as a cognitive one.

The next chapter will deal with precisely this problem: How can empirical data be used to inform, calibrate, and validate the theoretical models that have just been presented? What kinds of data have been collected from human crowds? What can each type of data contribute, and how can they be combined—with each other and with theoretical models—to yield a broad and comprehensive understanding of crowd behavior?

CHAPTER 3

Empirical Studies of Collective Locomotion

Empirical Data

As Chapter 2 demonstrates, there is a tremendous variety of models and theoretical approaches on offer for the study of human crowds. Despite this heterogeneity, the field seems to agree that, eventually, models will need to be tested against empirical data. This has been noted since the earliest days of crowd modeling; Henderson (1974) writes in one of the first papers on the subject, “[The validity of a model] can only be established empirically by further and more sophisticated comparisons [with data].” Since then, modelers have continued to call for comparison with empirical data; unfortunately, such efforts remain the exception, not the rule. In large part, this can be attributed to the difficulties involved with collecting data: compared to running numerical simulations, making measurements from real crowds is costly and time-consuming. It is also difficult to devise controlled experiments or record naturalistic observations that reflect the ‘general’ scenarios that models are designed to simulate. Moreover, there is an unfortunate tendency to prefer slick graphics and demonstrations over rigorous, grounded, and painstaking empirical work. All of these factors conspire to make empirical tests an uncommon enterprise. Nevertheless, such work is being carried out, and my dissertation is part of that effort.

Before reviewing the data-driven studies that have been conducted thus far, it will be useful to consider more generally the kinds of data that can be collected, and what they can contribute to our understanding of human crowd behavior.

First, data can be qualitative or quantitative. Qualitative data are non-numerical, and typically take the form of observations. These observations can be as simple as noticing regularities in the walking behavior of pedestrians in cities, or they may involve dozens of cameras placed in key locations on a single street corner. Qualitative data are useful

because they are easily produced, and they can offer insight into the characteristic phenomena that any model should be able to capture. Quantitative observations, by contrast, are more formal. They can consist of global measures aggregated over the entire crowd, such as densities or velocity distributions, or local variables like the heading directions or walking trajectories of individual pedestrians. While quantitative data is most often collected continuously over a period of time, one can also take discrete ‘snapshots’ of a crowd to estimate quantities like interpersonal distance or group structure. A common difficulty with both qualitative and quantitative approaches is determining which variables or features are relevant; in the former, this means the phenomena that are interesting and worth capturing in a model, while in the latter, this means the variables that should be modeled.

A second fundamental distinction can be drawn between observational and experimental data. Observational data allows one to learn about crowds in a naturalistic setting. It can also be collected in ways that would be unethical experimentally, such as using video surveillance from real evacuations. Experimental data, on the other hand, allow for better control of the environment and the goals of each individual. It allows for repeated trials with (more or less) the same initial conditions, which allows one to average across many trials. This helps to reduce the variability that is inherent in any crowd, or any behavioral experiment for that matter. Most important, experiments permit tests of causal relations; for instance, one can test whether the behavior of neighbors in a crowd is coupled, or merely correlated. But the control provided by experimental studies is often their greatest weakness; if the task or environment is too constraining, generalizability is limited. On the other hand, observational studies are limited to the specific circumstances they

observe. It is therefore necessary to use both approaches in conjunction with one another, across a variety of scenarios and situations.

A third and final aspect to consider is the underlying motivation for collecting data. This can range from testing, validating, or calibrating a particular model ('model-dependent data') to the more general goal of learning facts about crowds ('model-independent data'). These two enterprises are quite different, and again, both are necessary to fully understand crowd behavior. Models are the most powerful tool in the study of crowd behavior; they reify theoretical approaches from mere abstractions into something tangible. They also generate testable predictions, which is the basic process by which science progresses. But there is also a danger, that by focusing too narrowly on the kind of data that is needed to validate a given model, researchers may be blinded to fundamental aspects of crowd behavior that their model cannot explain. This is where model-independent data can help. Not bound by any particular theoretical construct, researchers can simply report general properties of crowd behavior, and provide concrete details and phenomena from the real world that candidate models should attempt to reproduce through simulations.

Chapter 3 will present a summary of the empirical studies that have been conducted to investigate crowd behavior, from all of these perspectives: qualitative and quantitative, observational and experimental, model-dependent and model-independent.

Model-Independent Data

Observational.

Some of the earliest empirical data on crowds comes from observations of pedestrian walking speed in naturalistic environments. One of the first, and most well-known, of the studies was conducted by Bornstein and Bornstein (1976). They collected data from 300 pedestrians in 15 cities and towns in Europe, Asia, and North America and found—remarkably—that walking speed scaled linearly with the logarithm of city size, as shown in Figure 3-1.




Figure 3-1: Mean walking speed as a function of log-scaled city size for 15 cities and towns. There is a strong ($r = 0.91$) and significant ($p < 0.001$) linear correlation. From Bornstein & Bornstein (1976).

Subsequent studies have elucidated this relationship, showing that it is at least partly due to the different demographics of cities and towns. Wirtz and Reis (1992) show that large cities tend to have higher proportions of young males and lower proportions of people over 60 years old, both of which serve to increase mean speed. Not surprisingly, other factors like density (Boles 1981, Walmsley & Lewis 1989), time of day (Hoel 1968), and even weather (Wirtz & Reis 1992) have been shown to affect walking speed.

Walking speed has also been studied in small groups, instead of isolated individuals. Boles (1981) found that individuals walking alone are slower than individuals walking in same-sex dyads. Costa (2010) extended this finding by showing that walking speed is even slower in mixed dyads. There is conflicting evidence about whether larger group sizes affect walking speed. Moussaïd et al. (2010) collected data from 1353 groups in two samples, and found a significant negative correlation between group size and walking speed. Costa (2010), however, found no relationship using data from 2,544 groups in five samples. This difference could be due to cultural factors (Moussaïd et al. collected data in France, Costa in Italy). Bonneaud et al. (2010) found that individuals adopt a common walking speed when walking in a group, such that slower individuals speed up and faster individuals slow down.

Moussaïd et al. (2010) and Costa (2010) also studied the composition and shape of small groups. Moussaïd et al. (2010) found that a majority of pedestrians (between 55% and 70%) walk in groups of two or more. Costa (2010) found that larger group sizes are less common than smaller ones. In his data, 61.8% of groups were dyads, 24.9% triads, 8.8% four-person groups, 2.8% five-person groups, 1.4% six-person groups, 0.1% seven-person groups, and 0.1% eight-person groups. The larger groups tended to split into smaller subgroups of dyads, triads, and single individuals. Costa (2010) further divides these groups based on their spatial configuration, and gives a detailed description and proportion for each of them.

While Costa's (2010) data consisted of static snapshots, Moussaïd et al. (2010) used video data to show how group shape changes dynamically with environmental condition. As described in Chapter 2, Moussaïd et al. (2010) studied the angle between pedestrians

walking abreast, and found that they increased with density, going from a 'line' to a 'V,' as shown in Figure 2-6. The authors note that these shapes do *not* have an aerodynamic shape, like the inverted V of migrating birds. Instead, they hypothesize that these shapes arise because group members want to communicate with each other, as agents in their model do.

Experimental.

Some of the earliest relevant experimental work came from the study of personal space. This was an area of intense interest from the late 1960s to the early 1980s, spawning no less than seven review articles (Aiello & Thompson 1980, Altman & Vinsel 1977, Hayduk 1978, Hayduk 1983, Mallenby & Roberts 1973, Obudho 1974, Sundstrom & Altman 1976) and several books (most notably Hall, 1969) on the topic. One of the most rigorous experimental studies was conducted by Hayduk (1981), and was designed to determine not only the distances but also the shape of personal space. The experimenter walked toward participants from one of eight angles arranged in a complete circle around them. Participants were instructed to report when they "just began to feel uncomfortable about [the experimenter's] closeness." They were also given one of several instructions which specified how they could move their eyes or heads. In some conditions, participants could look freely, while in others they were instructed to keep their head still while moving their eyes, and in still others they were instructed to keep their head still and fixate a point directly in front of them. The basic finding is that personal space is roughly circular, but it is compressed longitudinally, meaning participants allowed the experimenter to get closer when he started from the back than from the front. The results from Hayduk (1981) are shown in Figure 3-2.




Figure 3-2: The shape of personal space. Numerals 1-8 represent the starting positions of the experimenter. Lines represent the mean uncomfortable distance. Condition 3: eyes and head free to rotate; condition 4: head position fixed; condition 9: repeat of condition 3; condition 10: same as condition 3, but experimenter looked at subjects' eyes instead of neckline (as he did in all other conditions). From Hayduk (1981).

Dyer, Ioannou, Morrell, Croft, Couzin, & Waters (2008) conducted experiments to understand how groups of people reach a consensus about where to travel. They found that it took only a small minority of informed individuals (i.e. individuals who knew the destination) to effectively guide the group. When informed individuals were given conflicting goals, the group went to the goal where a higher proportion of individuals were informed to go. Interestingly, this conflict condition did not take longer than the condition where all informed individuals went to the same goal; Dyer takes this as a sign that “consensus decision making in conflict situations is possible, and highly efficient.”

In follow-up studies, Dyer and his colleagues (Dyer, Johansson, Helbing, Couzin, & Krause 2009, Faria, Dyer, Tosh, & Krause 2010) also found that the spatial distribution

of informed individuals affected consensus decision making. In particular, they found that having informed individuals in the center and at the edge of the group (rather than in between) increased the speed and accuracy with which the group reached its target. They also studied the effects of different instructions, and found that performance improved when uninformed individuals were told to follow an (unknown) leader, rather than simply to stay with the group, as they had been in previous experiments. Performance was even further improved when uninformed individuals knew some distinguishing, but not unique, property of the (still unknown) leader—in this case, the color of a sash he or she was wearing, which was also shared by other uninformed individuals.

Finally, Dyer et al. (2009) also extended these results to study much larger groups ($N = 100$ and 200). This is the only rigorous experimental study conducted on a large crowd to date. They replicated their findings from small groups, showing that the number and spatial distribution affected consensus decision making, and demonstrated that only a small minority of informed individuals (approximately 10%) could effectively steer the group.

Model-Dependent Data

Observational.

One of the first studies to use observational data to inform modeling was Henderson's (1971) article. After describing certain parallels between particles in a fluid and pedestrians in a crowd, he proposed that the Maxwell-Boltzmann distribution would be a reasonable approximation (as described in Chapter 2). He then tested this hypothesis by using video data collected from students on a footpath ($N = 693$), pedestrians at a crosswalk ($N = 628$), and children on a playground ($N = 779$). After obtaining the velocity for each pedestrian, Henderson compared them to the Maxwell-Boltzmann distribution,

and found remarkable agreement. Figure 3-3 shows the probability density function for velocity in the crosswalk data, as well as the Maxwell-Boltzmann distribution with best fit parameters. The fit is nearly perfect, except for a conspicuous bimodal peak near the maximum of the distribution.




Figure 3-3: Probability density function of velocities from pedestrians at a crosswalk. Solid lines represent human data; dashed lines represent the Maxwell-Boltzmann distribution with best-fit parameters. The left panel shows all pedestrians, while the right panel subdivides the data into males (curve a) and females (curve b). From Henderson (1971) and Henderson and Lyons (1971).

In a follow-up study, Henderson and Lyons (1972) hypothesized that this bimodality may have been caused by a sex difference, and separated males and females in their analysis. They found that males have a higher mean speed (1.59 m/s) than females (1.49 m/s), and that using two Maxwell-Boltzmann distributions produces a better fit than using just one, as shown in Figure 3-3. As a result, Henderson and Lyons conclude, “the sexes behave roughly like isotopes of a crowd species.”

Henderson’s work is important because it is one of the earliest attempts to fit a theoretical model to data collected from human crowds. But it only captures one phenomenon—the distribution of velocities—and fails to capture many other phenomena

that characterize crowd behavior. More critically, Henderson fails to test alternative models that might account for this pattern. Perhaps a simple Gaussian distribution of velocity, rather than the more complicated Maxwell-Boltzmann distribution, would have also been sufficient to account for the data in Figure 3-3.

Experimental.

A common approach to using model-dependent experimental data is to calibrate the parameters of a specific model. Moussaïd et al. (2009) performed a simple experiment to parameterize the social force model of Helbing and Molnar (1995). He had pedestrians walk alone down a narrow hallway, and used their trajectories to calibrate the component for accelerating to a preferred velocity. He also had two pedestrians walk in the same hallway—in one scenario, pedestrians walked in opposite directions, while in another, one pedestrian stood stationary at the center of the corridor—and used these trajectories to calibrate the parameters of the obstacle avoidance component in the social force model.

More recently, Olivier, Marin, Crétual, & Pettré (2012) performed an experiment in which pairs of pedestrians walked on a collision course, and used the data to test a model of obstacle avoidance. They proposed a new metric, the minimum predicted distance (MPD), which is the minimum separation between the two pedestrians at the point when they pass one another. They provide evidence that participants make adaptations only when the MPD is below a threshold value of about 1 m, but their model is never fully specified and they provide no evidence that other models could not also account for this behavior. Nevertheless, their effort is one of the few attempts to systematically isolate and experimentally test a single local interaction (obstacle avoidance) in human pedestrians.

Pétré's group has also investigated one-dimensional following in pedestrians walking along a circular track (Lemercier et al. 2012). They derived a dynamical systems model that accounts for the acceleration of the follower as a function of the speed of the person in front of them and the local density of the crowd. They calibrated this model to experimental data collected from 28 participants, and showed that it could reproduce the local motion of individual pedestrians as well as global patterns such as stop-and-go waves.

Shortcomings of Current Empirical Approaches

The data collected from the studies described above is essential, and has recently received greater attention from the pedestrian and crowd modeling communities. There is still much to be done, however. The empirical work suggests a pre-paradigmatic science, one which lacks a coherent set of background principles to guide research. In this section, I will examine some of these shortcomings, and in the next, I will suggest ways to address them.

The first problem with current approaches is that there is no consistent paradigm for data collection. There is little agreement about what kinds of data should be collected, or how it should be collected. Should we start by examining microlevel behavior, of individuals and perhaps small groups of individuals? Or would our efforts be better spent examining macrolevel behavior of entire crowds? How can we link these two levels? (For a prescient attempt to do this, see Yamori 1998). How should experimental and observational studies be used in conjunction with one another? What kinds of phenomena should we focus on first, and what order should we tackle them in? These are questions that too few researchers are asking. Typically, a researcher chooses some phenomenon or

another because he or she has a model that can reproduce it. Occasionally, they will validate it with empirical data, but there is almost no discussion about *why* one phenomenon is more important than any other. Nor do most experimenters explain *why* they favor experimental or observational data of a particular kind. And indeed, without a coherent framework on which to scaffold disparate empirical findings, there is no good reason to prefer any one approach to any other. Clearly, there is room for many kinds of projects—there is a lot to learn about crowd behavior. But it seems equally clear that some efforts should be prioritized over others. Should we really be spending time studying the specific effects of wheelchairs (Shimada & Naoi 2006) or crawling (Nagai, Fukamachi, & Nagatani 2005) before we know the basics of crowd behavior? Perhaps, but we should be careful not to put too much emphasis on these highly special cases.

A related problem is *how* to link data with theory. How can we reduce the variability that is inherent in crowd behavior? Should we do this at all? Is it enough for a model to reproduce behavior, or should it be able to predict behavior as well? Is it enough to capture the macrolevel behavior of the whole crowd, or do we also need to get the microlevel right? What kinds of variables are most relevant, and how do we choose them? Again, there does not seem to be much discussion about these issues.

Yet another key problem with existing models of human crowds is that they are *ad hoc*. Researchers begin with a set of rules or equations, and observe the behavior that is produced. By tweaking these rules, and their parameters, they bring their models closer to the desired behavior—either qualitatively or quantitatively. But the local rules themselves are almost never tested against data on behavioral interactions. As an example, many (if not most) individual-based models include some provision that limits the number of nearby

agents that influence a given agent. This reflects the reasonable assumption that a person in a crowd cannot gather information about *all* of the other people in the crowd. Most modelers assume that this is based on a metric measure, typically a fixed radius (e.g. Reynolds, 1987; Helbing & Molnar, 1995; Moussaid et al., 2010). A minority assumes a topological measure, generally a fixed number of nearest neighbors (e.g. Cavagna et al. 2008; Ma, Song, Zhag, Lo, & Liao, 2010). But no studies have been conducted in humans to determine which type of rule comes closest to describing pedestrian behavior. Two studies of birds (starlings: Nagy, Akos, Biro, & Vicsek, 2008; surf scoters: Lukeman, Li, & Edelstein-Keshet, 2010) provided evidence in favor of a topological strategy. By observing large numbers of birds, Ballerini et al. (2008) and Lukeman et al. (2010) could compute the correlation in heading between birds at different metric and topological distances, and found that birds were coupled to approximately seven neighbors (a constant topological distance), regardless of metric distance. This evidence does not seem to have had an impact on the field, however; modelers choose the rules that produce the best results, rather than the ones that are most biologically accurate or plausible.

Finally, there are no agreed-upon benchmarks for a “successful” model. Is a model successful if it reproduces some phenomenon, or does it need to be able to capture a range of phenomena? What does it mean to ‘reproduce’ or ‘capture’ a phenomenon at all—is a qualitative ‘eyeball’ observation sufficient, or is some quantitative measure required? How can we compare across models? If a model captures one phenomenon very well, but another model captures two or three less well, which should be preferred?

Singh et al. (2009) have made a valiant attempt to address this problem with their SteerBench framework. SteerBench is a benchmark suite that tests models on a wide range

of scenarios—ranging from simple ones like obstacle avoidance and binary pedestrian interactions to complex ones like bidirectional traffic in a hallway and simulations with 5000 agents. For each scenario, a number of performance measures are computed, such as the number of collisions, distance travelled, average speed change, and so on. These measures are combined to yield a score for each model on each scenario, which can then be used to objectively compare models. This allows for easy comparison, both for highly specific needs (e.g. one wants the best model for simulating a particular scenario) and to see which models generally outperform others across a wide range of tasks. However, the tasks and measures adopted themselves must be justified and may be open to criticism.

Efforts like those of Singh et al. (2009) are a good start, and provide a basis for evaluating and comparing different models. But what the study of human crowds most needs is a coherent framework around which to build a base of knowledge, rather than the hodge-podge, grab-and-dash approach that has been taken so far. In the next section, I will outline the approach that our lab has taken to introduce such a framework, and describe how the experiments in my dissertation will further our goal of producing a cognitively-plausible, data-driven model of human crowd behavior.

A Behavioral Dynamics Approach to Crowd Behavior

In this section, I will describe the efforts that my colleagues and I have made in the VENLab over the past several years to develop a cognitively-plausible, data-driven approach to studying human crowd behavior. In this section I will summarize this work so far and set the stage for the experiments that I will present in the succeeding chapters.

Motivated by Gibson's (1950; 1979) proposals that optic flow could be used for the visual control of locomotion, Warren and colleagues conducted some of the first work to show that humans can perceive heading direction from optic flow (Warren, Morris, & Kalish, 1988; Warren & Hannon, 1988) and utilize this information to regulate steering during walking (Warren, Kay, Zosh, Duchon, & Sahuc, 2001). The creation of the Virtual Environment Navigation Laboratory, one of the world's largest ambulatory virtual environments (Tarr & Warren, 2002), enabled the experimental study of such questions. Virtual reality made it possible to systematically manipulate optic flow separately from other sources of information about heading (e.g. proprioception), and a series of experiments provided further evidence to support the hypothesis that optic flow is used to control steering in an on-line fashion (Bruggeman, Zosh, & Warren, 2007; Bruggeman & Warren, 2010).

In an effort to integrate Gibson's concepts about perception with the emerging dynamical systems approach to action (Kelso, 1995; Kugler & Turvey, 1987), Fajen and Warren introduced the behavioral dynamics framework, which offers an account of how locomotor behavior can be regulated based on on-line information about an observer's environment. This work began with a dynamical model of steering and obstacle avoidance (Fajen & Warren, 2003), which could account for route selection in novel environments without the need for explicit path planning. Subsequent research modeled new behaviors within the same framework, such as intercepting a moving target (Fajen & Warren, 2006, 2007) and avoiding a moving obstacle (Cohen, Bruggeman, & Warren, 2005; Cohen, Bruggeman, & Warren, 2006). The steering dynamics model was also extended to more complex scenarios, such as integrating moving targets with stationary and moving

obstacles (Bruggeman & Warren, 2005), avoiding extended barriers (Gerin-Lajoie & Warren, 2008), and navigating through fields of up to 12 stationary obstacles (Fajen & Warren, 2010).

The next step was to consider whether pedestrian interactions could be explained using the models that were developed for inanimate objects. Cohen (2009) found that pursuit and evasion could be accounted for with pairs of pedestrians in some scenarios, using the moving target and moving obstacle components. Specifically, the steering dynamics model could account for mutual evasion (where each pedestrian treats the other as a moving obstacle), pursuit-evasion (where a ‘pursuing’ pedestrian treats an ‘evading’ pedestrian as a moving target, and the evader treats the pursuer as a moving obstacle), although the model was less successful at explaining mutual pursuit (where each pedestrian treated the other as a moving target to be intercepted). Cohen also showed that participants could perceive whether individual pedestrians in a virtual crowd were pursuing or evading them (in effect, treating the participant as a moving target or a moving obstacle, respectively).

The behavioral dynamics approach was then applied to crowd behavior in several parallel lines of research. An essential aspect of collective locomotion is that neighbors in a crowd coordinate their movements with one another. We thus conducted experiments to investigate coordination among pedestrians, building up from dyads walking together. This included work on speed control in following (Rio, Rhea, & Warren, 2014) and side-by-side walking (Page & Warren, 2012; Page & Warren, 2013), as well as heading alignment in both cases (Dachner & Warren, in press). In Rio, Rhea, & Warren (2014), I investigated both the visual information and behavioral strategies that were used to coordinate speed in

following, and found evidence to support the hypothesis that a follower walks behind a leader by nulling changes in the leader's optical expansion. Based on these results, I developed a dynamical model of speed control in which the follower's acceleration is a function of the difference in speed with the leader. In Chapter 4, I will use this model to predict acceleration in participants walking in a small group, to test whether it generalizes to coordination with multiple neighbors.

We also began to extend the framework to larger groups of pedestrians. Stephane Bonneaud and I began with a pilot experiment using small groups of 4 pedestrians, in an experimental design that closely resembled that of the original Fajen & Warren (2003) study. Participants walked toward a common goal, and both goal location (left, center, right) and the initial separation of the group (0.5, 1.0, 1.5, 2.5 m) were manipulated. Bonneaud, Rio, Chevallier, & Warren (2012) reported several global patterns in these data, including the maintenance of an initial group configuration, the adoption of a common walking speed, and anisotropy in interpersonal distance, with participants maintaining a larger anterior-posterior (front-back) than lateral (side-side) distance. Kiefer, Bonneaud, Rio, & Warren (2014) extended these analyzes by conducting a principal components analysis and cross-recurrence quantification on the four walking trajectories, and found evidence for dimensional compression in real groups when compared to a random shuffling of participants in matched conditions, indicating a local coupling between participants when walking in small groups.

To enable simulations of the model with large numbers of agents, Bonneaud developed a multiagent simulation platform (Bonneaud & Warren, 2014). This platform makes it possible to test whether different model components are necessary to account for

behavior in different scenarios, and potentially to compare alternative models of each component. For example, it appears that neighbors are treated as moving obstacles in some scenarios, but applying this component in other scenarios results in unrealistic movement patterns (Warren & Bonneaud, 2014). In collaboration with Thomas Serre, Bonneaud and others have also integrated a “front end” model of the human visual system to analyze the information that is used to guide behavior in complex environments and promote the neural plausibility of the steering dynamics model (Bonneaud, Warren, Olfers, Irwin, & Serre, 2013; Barhomi, Yanke, Bonneaud, Warren, & Serre, 2014).

This work provides the foundation that I will build upon in my dissertation. My goals are to scale up this framework from the study of binary pedestrian interactions to the study of human crowds, and to conduct experiments that will help further our attempts to model crowd dynamics. In particular, my dissertation is an attempt to understand the local coupling between neighbors in a crowd. The steering dynamics model, like most models of collective behavior (Moussaid et al., 2012), is built out of binary interactions between pairs of pedestrians. Thus, in order to extend the model to large crowds, we must answer specific questions about the functional connections between neighbors in a crowd. Is there a limit to the number of neighbors that influence an individual pedestrian? How are the influences from multiple neighbors combined? How does the strength of the interaction between neighbors change as a function of their distance, their lateral position (eccentricity), or as a function of the density of the crowd? Is the neighborhood metric or topological? These questions must be addressed, either explicitly or implicitly, in order not only to simulate crowd behavior, but to understand the mechanisms that generate it. One of the main goals of my dissertation is to provide answers to these questions and support

the continued development of the behavioral dynamics framework. Answers to these fundamental questions will not be limited to our own model, however. Rather, they should inform the community as a whole, by providing constraints on the form and parameterization of any model of crowd behavior. Finally, they will shed light on the basic mechanisms that underlie collective locomotion, and in conjunction with studies in other animal species, will help us better understand how local interactions can give rise to global phenomena through the process of self-organization.

CHAPTER 4

Experiment 1: Walking in a Small Crowd of Confederates

Collective locomotion requires that agents perceive and respond to features of their environment, including the movements of those around them (Vicsek, 2008); in humans, this is primarily accomplished using the visual system, which is the principal sensory modality implicated in the control and guidance of locomotion (Patla, 1997; Warren, 2009). Pedestrians in a crowd use visual information about the people around them in order to regulate walking speed and direction and move together as a collective entity. This perceptual-informational link, which I will call the *local visual coupling* between neighbors, is at the core of human crowd dynamics. Despite this, fundamental questions about the local coupling remain unanswered. For example: How does the local coupling strength between a pedestrian and a neighbor vary as a function of factors like the neighbor's distance or their eccentricity relative to the direction of travel? When pedestrians are coupled to more than one neighbor, how are these influences combined?

In order to address these questions, and to investigate the visual coupling between a pedestrian and a few neighbors, Experiment 1 was conducted with small crowds (N=4). A participant walked together with three confederates, a subgroup of whom changed their walking speed (slower/faster) or direction (left/right) on each trial. Varying parameters of this subgroup made it possible to observe how these factors affected the local coupling. In particular, the number of confederates in the manipulated subgroup (0, 1, 2, 3), the eccentricity of the single confederate (center, side), and the initial distance between pedestrians (1, 2 m) were all varied, in order to test specific hypotheses about the local coupling.

One common hypothesis (e.g. Helbing & Molnar, 1995) is that a pedestrian should be more strongly coupled to neighbors located in front (that is, aligned with the direction

of travel) than those on the sides. This hypothesis would be supported if participants showed a greater response to changes in the speed or direction of a confederate with small eccentricity (center) than to similar changes in a confederate with large eccentricity (side). On the other hand, it could be the case that eccentricity has no effect, or the opposite might be true—participants might respond more to confederates with small eccentricities. Another common hypothesis is that the coupling strength should decrease as neighbors get farther away; this hypothesis predicts that participants should show a greater response when confederates are close (1 m) rather than far away (2 m). On the other hand, some accounts of collective locomotion (e.g. Ballerini et al. 2008) hold that what matters is topological distance (i.e. the number of neighbors between the participant and a given neighbor), rather than the absolute metric distance. This would predict no difference between the two conditions, because the topological structure remains unchanged. Lastly, there are several ways in which the influences from multiple neighbors (0, 1, 2, 3) could be combined. The influences could be linearly combined, with each additional confederate eliciting a greater response from participants; or, there might be a ‘threshold’ or ‘saturation’ effect, whereby a participant’s response to a single confederate’s speed change, say, is quantitatively similar to their response to similar speed changes in 2 or 3 confederates.

The remainder of my dissertation will continue to explore these questions in a variety of contexts, including with larger crowds and in virtual reality, but Experiment 1 will set the stage by studying the local coupling between neighbors in a small crowd under controlled conditions in order to isolate the specific effects of neighbor number, eccentricity, and distance.

Methods

Participants

Ten participants, 4 female and 6 male, were recruited for Experiment 1. None reported having any visual or motor impairment, and they were paid for their participation. A separate group of three trained experimenters, 2 female and 1 male, acted as confederates in the experiment.

Apparatus

The research was conducted in the Virtual Environment Navigation Laboratory (VENLab) at Brown University. The participant and three confederates walked together in a 12x14 m room, and each wore a set of wireless headphones connected to a PC that delivered pre-recorded instructions. Head position and orientation were recorded at a sampling rate of 60 Hz using a hybrid ultrasonic-inertial tracking system (IS-900, Intersense, Billerica MA) affixed to the headphones. (Note that virtual displays were not used in this experiment.)

Procedure

Participants in Experiment 1 were not informed of the confederates' complicity in the experiment. All four pedestrians (the participant and three confederates) received the same verbal instructions at the beginning of the experiment; they were told simply to "walk together across the room." On each trial, the group was directed via wireless headphones to go to one of 8 starting locations, which were marked on the floor with pieces of colored tape. The starting locations were arranged in two concentric diamond configurations, with initial distances of 1 or 2 m on a side (see Figure 4-1). Only trials in which the participant was located in the back position (i.e. the rear vertex of the diamond) were analyzed further.

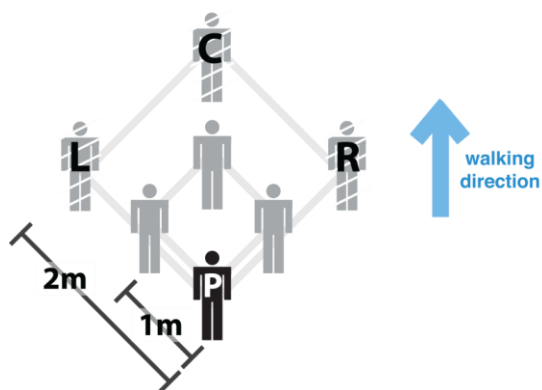


Figure 4-1: Initial conditions of participant (P) and three confederates (L, C, R) in Experiment 1.

In addition, the confederates could be covertly instructed to turn left, turn right, walk faster, or walk slower on each trial. The same instruction was given to a subgroup of 1, 2, or 3 of the confederates, depending on the condition; to avoid collisions, in the subgroup = 1 condition the Right confederate was instructed to “turn right” (and vice versa), and in the group = 2 condition the Right and Center confederates were instructed to “turn right” (and vice versa). The remaining confederates, who were not manipulated, walked straight ahead at their preferred speed. To determine whether the unmanipulated confederates were successfully following these instructions—or whether their behavior was inadvertently affected by the manipulated confederates—I computed their change in final lateral position (see Results) as a function of whether their neighboring confederates were manipulated. A one-way ANOVA on these lateral positions showed that there was an effect of the manipulated confederate on the unmanipulated confederates (turn left: $F(3,298) = 3.65$, $p < .05$; turn right: $F(3,302) = 3.38$, $p < .05$); however, Bonferroni-corrected follow-up tests showed that only one difference (of eight) was significant at $p < .05$; namely, manipulated confederates drifted right ($M = 0.27$ m, $SD = 0.20$ m) when a single confederate in the center was manipulated (“R1c”). Overall, these results suggest that unmanipulated

confederates were unaffected by the manipulated confederates in the majority of conditions, and thus were largely able to successfully follow the instructions they were given.

Each trial commenced when a verbal command (“Begin”) was played to all four pedestrians, which instructed them to start walking across the room. When the first confederate crossed an invisible line located 4 m in front of the participant’s starting location, the subgroup of confederates who received manipulation instructions heard a pure tone indicating they should initiate the speed or heading change. They maintained this new speed or heading direction for the duration of the trial. Once the first confederate reached the opposite side of the room (after walking about 13 m), all four pedestrians received a verbal command (“End”), indicating the end of the current trial, and instructions for the next trial were presented. Figure 4-2 shows the confederate trajectories for two sample trials.

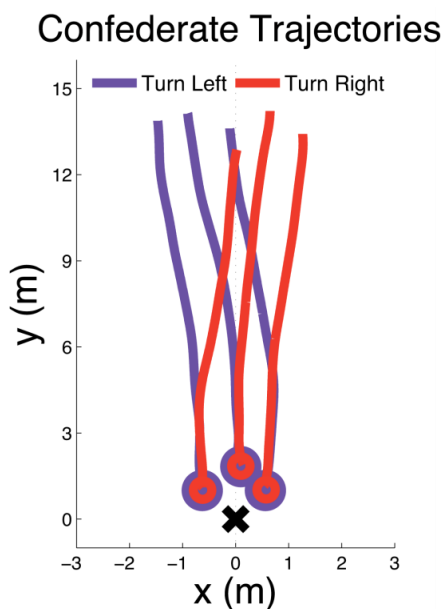


Figure 4-2: Confederate trajectories for a sample turn left (purple) and turn right (red) trial. Initial positions marked with an “O” for confederates and an “X” for the participant.

Design

Because some heading changes were not tested to avoid collisions, there were 8 heading changes with 4 trials of each (for 32 heading-change trials), and 12 speed changes with 2 trials of each (for 24 speed change trials), plus 8 no-change control trials, yielding a total of 64 trials per participant. The 12 speed changes were collapsed into 8 conditions to correspond with heading changes – refer to Table 4-1. This yielded a 2 (initial distance) x 2 (heading or speed change) x 9 (confederate manipulations) design.

Condition	Confederate Manipulation	Type	Repetitions
L3	L/L/L	Heading	4
L2	L/L/-	Heading	4
L1c	-/L/-	Heading	4
L1s	L/-/-	Heading	4
Ctrl	-/-/-	Heading	8
R1s	-/-/R	Heading	4
R1c	-/R/-	Heading	4
R2	-/R/R	Heading	4
R3	R/R/R	Heading	4
S3	S/S/S	Speed	2
S2	S/S/-; -/S/S	Speed	4
S1c	-/S/-	Speed	2
S1s	S/-/-; S/-/-	Speed	4
Ctrl	-/-/-	Speed	8
F1s	-/-/F; -/-/F	Speed	4
F1c	-/F/-	Speed	2
F2	-/F/F; F/F/-	Speed	4
F3	F/F/F	Speed	2

Table 4-1: Full list of conditions in Experiment 1. Confederate Manipulation lists the instructions given to the left, center, and right confederates, respectively; "L"=turn left, "R"=turn right, "S"=slower, "F"=faster, "-"=no change. Note that there are fewer Heading Change conditions, because manipulations that resulted in confederates crossing the participants' path were eliminated (e.g. L/-/- is included, but not -/-/L).

Data analysis

The time series of head position for all four pedestrians were recorded in three dimensions, but only data in the horizontal xy -plane were analyzed. Each time series was filtered, using a forward and backward 4th-order low-pass Butterworth filter with a cutoff frequency of either 1.0 Hz or 0.6 Hz, to reduce occasional tracker error and attenuate both anterior-posterior and lateral accelerations due to the step cycle. The 1.0 Hz cutoff was used for analyses of speed, as in Rio, Rhea, & Warren (2014), because this eliminates anterior-posterior accelerations on each full step associated with the head shifting forward and backward relative to the center of mass, while the 0.6 Hz cutoff was used for analyses of heading direction, as in Fajen & Warren (2003), because this also eliminates lateral oscillations associated with side-to-side head movements on each half-step. To eliminate edge effects from filtering at the end of the trial ('endpoint error'), the position time series were extended by 2 s using linear extrapolation based on the last 0.5 s of data (Vint & Hinrichs, 1996; Howarth & Callaghan, 2009); the extrapolated portions were not included in any subsequent analysis. The filtered position data were used to compute time series of speed, heading, and acceleration. Due to tracking errors, the participant's trajectory could not be recovered on 78 trials (8.3%), which were excluded from further analysis.

Results

Heading

Descriptive. Figure 4-3 shows an overhead view of participants' trajectories for trials in which a subgroup of one or more confederate turned left (purple) or right (red),

across all participants and trials. In general, participants turned left (or right) on trials in which a subgroup of confederates did. Furthermore, participants were sensitive to the number of confederates in the subgroup. Figure 4-3 shows a single participant's trajectories for all trials in which a subgroup turned left. Shade represents the number of confederates in the manipulated subgroup; going from light (1 confederate) to medium (2 confederates) to dark (3 confederates). Again, it can be observed that participants generally turn more when more confederates are manipulated.

To quantify these observations, participants' mean changes in final lateral position were computed in each condition. The change in lateral position on each trial was computed by subtracting the participant's x -position at $t=1$ s before the trial ended (to minimize any edge effects, such as prematurely slowing down or turning around to begin the next trial) from the participant's mean final x -position at the same time for all control trials. A positive value indicates a rightward change, while a negative value indicates a leftward change. The mean final lateral position on control trials was +0.04 m, suggesting a weak tendency to drift to the right. A one-sample t -test confirmed that this drift in final lateral position was significantly different from zero (i.e. straight ahead), $t(9) = 2.35, p < .05$.

Measuring the participant's final lateral position in these ways bolsters the intuitions that can be gleaned from Figures 4-3 and 4-4. Figure 4-5 shows the mean change in lateral position for each of the 9 speed change conditions listed in Table 4-1, averaged across participants, separated by the initial distance. A two-way ANOVA on final lateral position, with the confederate manipulation and initial distance as factors, confirmed there was a significant main effect of the manipulation ($F(8,162) = 22.40, p < .001$), no effect of initial distance ($F(1,162) = 0.26, p > .05$), and no interaction ($F(8,162) = 0.18, p > .05$).




Figure 4-3: Participant trajectories for all trials in which at least one confederate turns left (purple) or right (red).

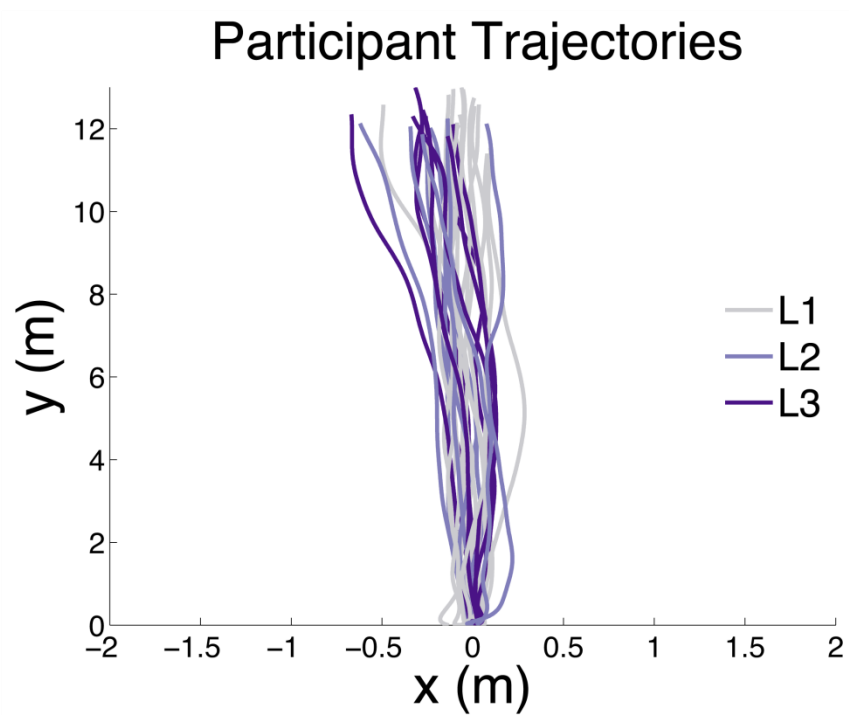


Figure 4-4: Participant trajectories for a single representative participant, for all trials in which a subgroup of confederates turned left. Shade represents the number in the subgroup (L1 = one confederate, L2 = two confederates, L3 = three confederates).

Participants were not only sensitive to the direction of change in confederate motion, but also to the number of confederates in the subgroup as well as their eccentricity. Figure 4-6 shows the same data as Figure 4-5, but it has been reshaped in two ways. First, in order to provide a better estimate of the effects of number and eccentricity, rather than direction, “left” conditions have been combined with “right” by multiplying the former by negative one, which yields the unsigned magnitude of the change in final lateral position. For example, data from “L3” (in which three confederates receive instructions to turn left) and “R3” trials are both grouped as “Number=3” or “N=3” trials. Second, number and eccentricity have been factored out by combining data from the single confederate conditions in Figure 4-5. Thus, “L1s” and “L1c” are both treated as “N=1,” but “L1s” is treated as “Eccentricity=Side” and “L1c” is treated as “Eccentricity=Center.” Analyzing the data in this way makes the effect of these two factors more readily interpretable, and also reduces the variability in any given condition, which now contains more observations.

Separate two-way ANOVAs were conducted for number and eccentricity. For the former, there was a significant main effect of number ($F(3,172) = 12.75, p < .001$), no effect of initial separation ($F(1,172) = 0.41, p > .05$) and no interaction ($F(3,172) = 0.03, p > .05$). Participants’ change in final lateral position increased with the number of confederates in the manipulated subgroup, from one ($M = 0.07$ m, $SD = 0.13$ m) to two ($M = 0.15$ m, $SD = 0.17$ m) to three ($M = 0.21$ m, $SD = 0.17$ m). Post-hoc tests using the Bonferroni correction revealed that there were significant differences between all number conditions, $p < .05$, except between the control condition ($N=0$) and $N=1$, and between $N=2$ and $N=3$, which were not significantly different from one another, $p > .05$. A similar

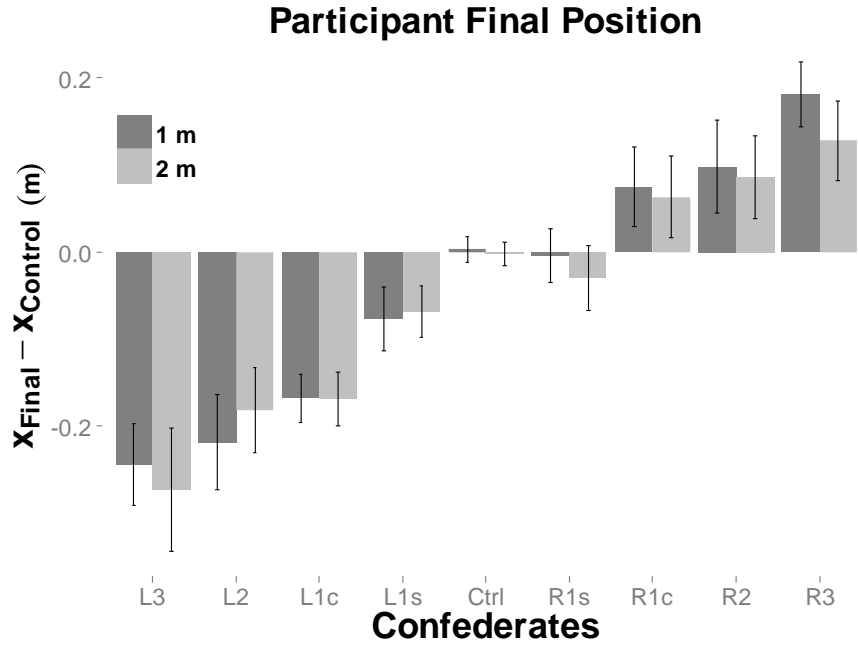


Figure 4-5: Mean changes in participants' final position in Experiment 1.




Figure 4-6: Mean changes in participants' final position in Experiment 1, broken down by confederate number (left) and eccentricity (right). "Left" conditions are combined with "right" by reversing the sign of the former.

eccentricity ($F(2,94) = 10.08, p < .001$), no effect of initial distance ($F(1,94) = 0.20, p > .05$) and no interaction ($F(2,94) = 0.03, p > .05$). Participants' change in final lateral position was greater for a single confederate in the center ($M = 0.12$ m, $SD = 0.13$ m) than a single confederate on the side ($M = 0.03$ m, $SD = 0.11$ m).

To further quantify how the visual coupling depends on the number and eccentricity of neighbors, I performed a multiple linear regression of participant final positions on the 3 confederates' final positions. This analysis showed that nearly half of the variance in participant final position ($R^2 = 0.43, p < .001$) could be accounted for by the confederate manipulations, and the weight of the center confederate was greater than that of the left and right confederates ($w_L = 0.27, w_C = 0.47, w_R = 0.26$).

Speed

Descriptive. The results obtained for speed follow a similar pattern as the results obtained for heading. Figure 4-7 shows participants' time series of speed for all trials in which a subgroup of confederates walked slower (blue) or faster (orange). In general, participants walked slower (or faster) when a subgroup did.

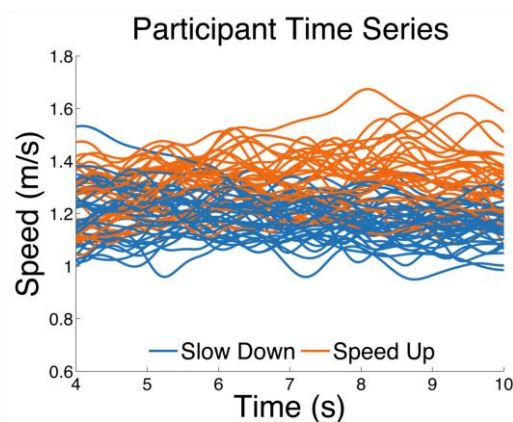


Figure 4-7: Participant speed time series for all trials in which a subgroup of confederates slowed down (blue) or sped up (orange). Data is shown from $t=4$ to $t=10$ s, to highlight the main 'steady-state' portion of the trial and eliminate initial acceleration from standstill.

Similar to the change in lateral position, the participants' change in speed on each trial was computed by subtracting the participant's mean final speed during a 1 s interval centered on $t=1$ s before the trial ended, (i.e. from $t=1.5$ s to $t=0.5$ s before the end of the trial) from the participant's mean final speed over the same interval on control trials. Positive values indicate that the final speed was faster, and negative values that final speed was slower, than in control trials. The mean final speed on control trials was significantly faster in the 1 m initial distance conditions ($M = 1.24$ m/s, $SD = 0.04$ m/s) than in the 2 m conditions ($M = 1.19$, $SD = 0.05$ m/s), $t(9) = 2.33$, $p < .05$, so these two different values were subtracted. Note that these walking speeds of around 1.2 m/s are typical of human walking in similar experimental contexts (Fajen & Warren, 2003; Rio, Rhea, & Warren, 2014).

As before, these measurements of participants' final speed bolster the intuition from Figure 4-7. Figure 4-8 shows the mean change in lateral position for each of the 9 speed change conditions listed in Table 4-1, averaged across participants, separated by the initial distance. A two-way ANOVA on final speed, with the confederate manipulation and initial distance as factors, confirmed there was a significant main effect of the manipulation ($F(8,160) = 48.77$, $p < .001$), a significant main effect of initial distance ($F(1,160) = 7.88$, $p < .01$), and no interaction ($F(8,160) = 0.20$, $p > .05$).

To better separate the effects of number and eccentricity, the data were reshaped as before for heading changes, and are shown in Figure 4-9. "Slow" trials were combined with "fast" trials by multiplying the former by negative one, yielding the unsigned magnitude of the change in final speed, and conditions were combined as before to isolate number and eccentricity. Separate two-way ANOVAs for number and eccentricity

produced a similar pattern of results to those obtained for heading changes. There was a significant main effect of number ($F(3,170) = 29.37, p < .001$), no main effect of initial distance ($F(1,170) = 0.01, p > .05$), and a marginal interaction ($F(3,170) = 2.17, p = 0.09$). Participants' change in final speed increased with the number of confederates in the manipulated subgroup, from one ($M = 0.05$ m/s, $SD = 0.07$ m/s) to two ($M = 0.13$ m/s, $SD = 0.08$ m/s) to three ($M = 0.15$ m/s, $SD = 0.09$ m/s). Post-hoc tests using the Bonferroni correction revealed that there were significant differences between all number conditions, $p < .05$, except between the $N=2$ and $N=3$ conditions, which were not significantly different from one another, $p > .05$. A similar pattern of results was observed for eccentricity, where there was a main effect of eccentricity ($F(2,93) = 16.87, p < .001$), no effect of initial distance ($F(1,93) = 0.48, p > .05$) and no interaction ($F(2,93) = 0.14, p > .05$). Participants' change in final speed was greater for a single confederate in the center ($M = 0.08$ m/s, $SD = 0.06$ m/s) than a single confederate on the side ($M = 0.03$ m/s, $SD = 0.07$ m/s).

As with the heading change conditions, I performed a multiple linear regression of participant final positions on the 3 confederates' final positions. This analysis showed that nearly half of the variance in participant final speed ($R^2 = 0.46, p < .001$) was accounted for by the confederate manipulations, and the weight of the center confederate was greater than that of the left and right confederates ($w_L = 0.19, w_C = 0.47, w_R = 0.34$).

Modeling and simulation. To better understand speed coordination in this experiment, I simulated the participants' walking trajectories using Rio, Rhea, & Warren's (2014) model of speed control developed to account for following behavior in dyads. In this model, the acceleration of the follower is a function of the difference in speed between

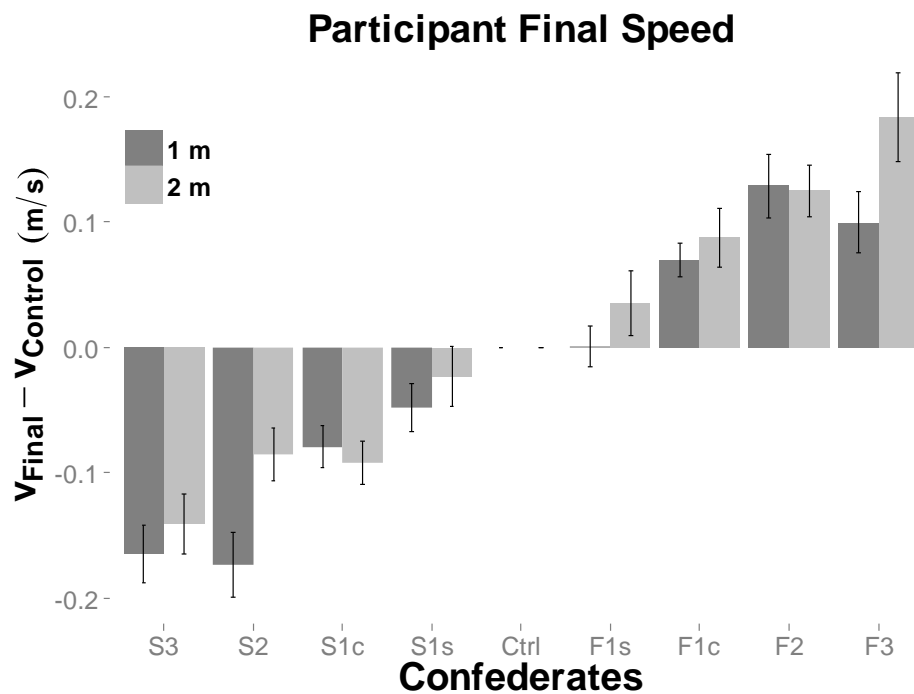


Figure 4-8: Mean changes in participants' final speed in Experiment 1.

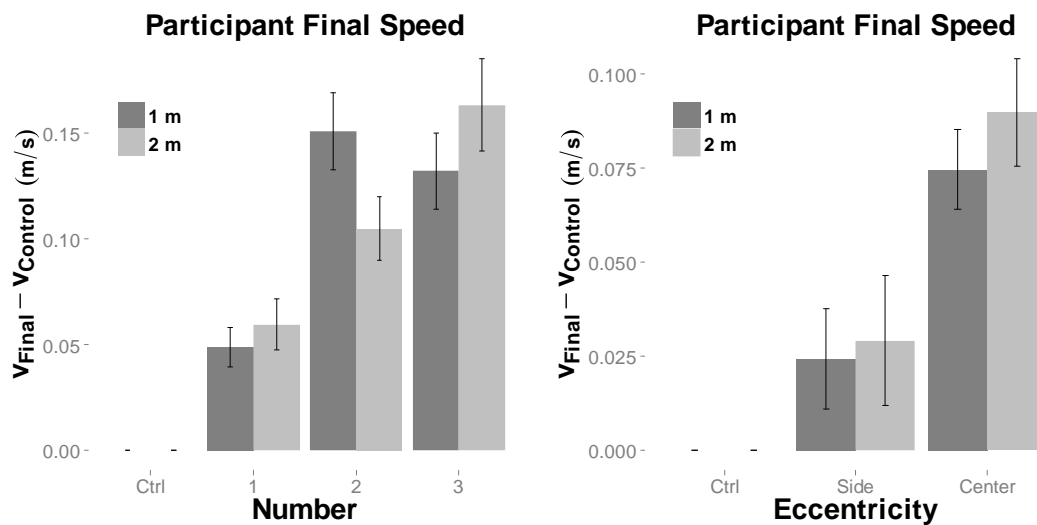


Figure 4-9: Mean changes in participants' final speed in Experiment 1, broken down by confederate number (left) and eccentricity (right). "Slow" conditions are combined with "fast" by reversing the sign of the former.

the leader and the follower, with a gain k :

$$\ddot{x}_{follower} = k(\dot{x}_{leader} - \dot{x}_{follower}) \quad (1)$$

A straightforward way to generalize this model to multiple neighbors would be to predict the participant's acceleration using a weighted linear combination of the difference in speed with each of the 3 confederates:

$$\ddot{x}_{participant} = k[w_L(\dot{x}_L - \dot{x}_p) + w_C(\dot{x}_C - \dot{x}_p) + w_R(\dot{x}_R - \dot{x}_p)] \quad (2)$$

Adopting the gain k from our previous dyad data and the weights w for eccentricity from the multiple regression on the present speed data, I simulated the participant's time series of acceleration on each trial using the 3 confederates' time series of speed as input. Consistent with the method used in Rio, Rhea, & Warren (2014), the human data were truncated to remove the first 3 s of each trial, to eliminate acceleration transients. Not surprisingly, the trial-by-trial model fits based on the group input are lower (mean $r = 0.28$) than the previous fits with a single leader (mean $r = 0.67$). But as shown in Figure 4-10 and 4-11, the model closely reproduces the pattern of mean final speeds in Figures 4-8 and 4-9, respectively.

As before, two-way ANOVAs were conducted on the final speeds, broken down by number and eccentricity, using data type (either observed from the human data or predicted by the model) as the second factor. There was a significant main effect of number ($F(3,223) = 65.46, p < .001$), a marginal effect of data type ($F(3,223) = 3.13, p = 0.08$), and no interaction ($F(3,223) = 1.45, p > .05$). Participants' change in final speed increased with the number of confederates in the manipulated subgroup, from one ($M = 0.04$ m/s, $SD = 0.07$ m/s) to two ($M = 0.13$ m/s, $SD = 0.06$ m/s) to three ($M = 0.17$ m/s, $SD = 0.08$ m/s). Post-hoc tests using the Bonferroni correction revealed that there were significant

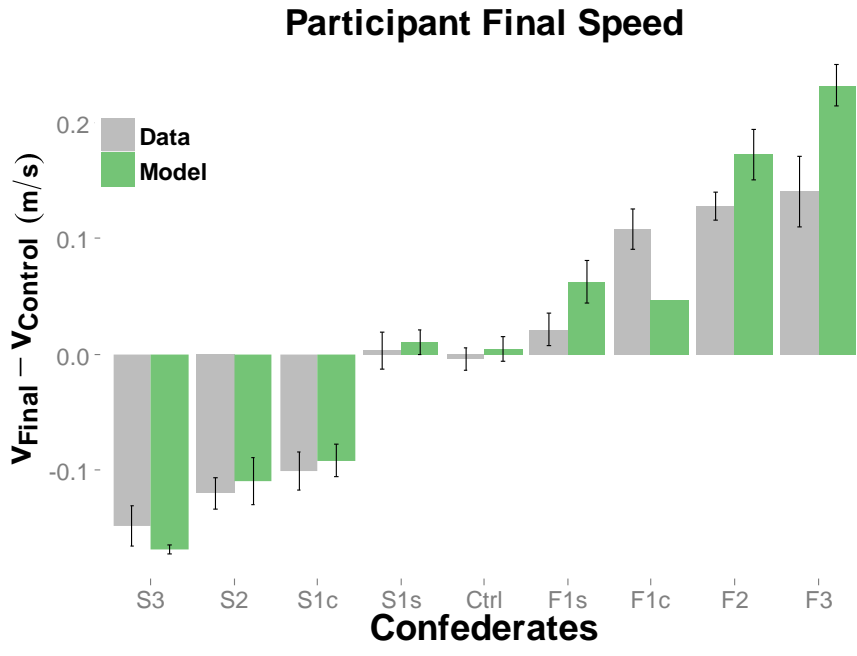


Figure 4-10: Mean changes in participants' final speed in Experiment 1, as observed in human data (grey) and as predicted by the model adapted from Rio, Rhea, & Warren (2014) (green).

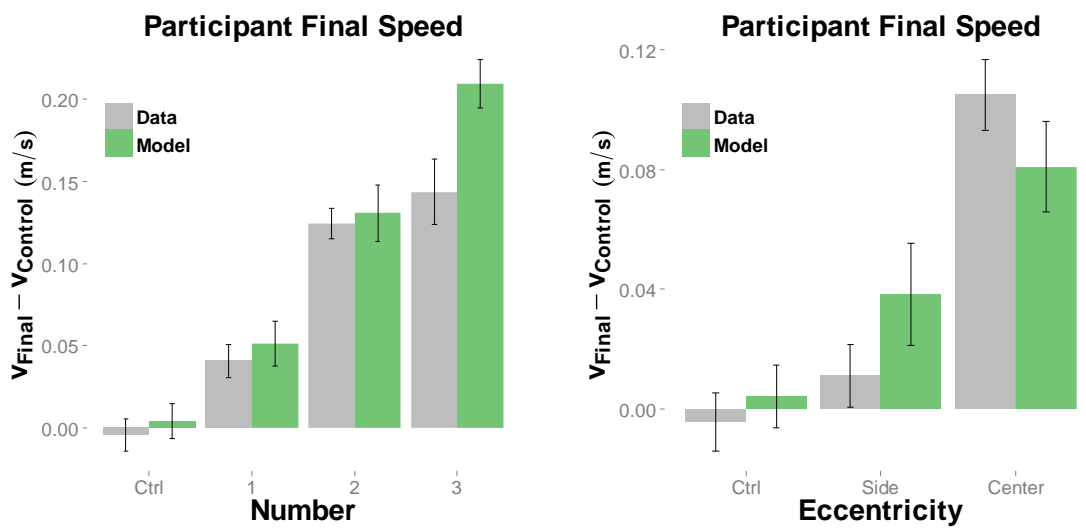


Figure 4-11: Mean changes in participants' final speed in Experiment 1, broken down by confederate number (left) and eccentricity (right), for data (grey) and model (green).

differences between all number conditions, $p < .05$. A similar pattern of results was observed for eccentricity, where there was a main effect of eccentricity ($F(2,137) = 23.68$, $p < .001$), no effect of data type ($F(1,137) = 0.66$, $p > .05$) and no interaction ($F(2,137) = 0.80$, $p > .05$). Participant's change in final speed was greater for a single confederate in the center ($M = 0.10$ m/s, $SD = 0.05$ m/s) than a single confederate on the side ($M = 0.02$ m/s, $SD = 0.06$ m/s). Overall, the pattern of results obtained is the same from both the human data and the model predicts, and the differences between them were either not significant (for eccentricity) or marginally significant (for number), which indicates that the following model from Rio, Rhea, & Warren (2014) can be extended to predict acceleration in small crowds.

Discussion

Overall, the results of Experiment 1 are broadly consistent with the hypothesis that the influences from multiple neighbors are linearly combined. For both speed and heading changes, participant response increased with each additional neighbor that was manipulated. Moreover, a model based on a linear combination of influences, weighted by eccentricity, was able to reproduce the general pattern of results for speed (compare Figures 4-10 and 4-11 to Figures 4-8 and 4-9). One caveat to note is that the difference between $N=2$ and $N=3$ confederates was not significant for either speed or heading changes, although they did trend in the direction predicted by a pure linear combination. There are several possible explanations for this effect. One, of course, is that the relatively small number of participants may have resulted in insufficient power to detect a difference between $N=2$ and $N=3$ confederates. But there are other, more interesting alternatives: for

example, participants may have already been responding maximally when two confederates were manipulated, such that adding a third confederate did not produce any further effect on their behavior. Experiment 3A will provide a similar test of the effect of number in a larger crowd, to see whether the linear trend observed in Experiment 1 is a general phenomenon.

There was a significant effect of eccentricity; participant response was greater to a single confederate on the side than a single confederate in the center (Figures 4-6 and 4-9). This is consistent with the hypothesis that the local coupling is greater for pedestrians with small eccentricities (along the direction of travel) than those with greater eccentricities (alongside). It is difficult to draw definitive conclusions from only two data points (center and side), however. It could be the case that a pedestrian aligned with the direction of travel (i.e. the center confederate) is treated qualitatively differently from other pedestrians, or there could be a more gradual decay of influence as a function of eccentricity. Experiment 3C will investigate eccentricity in finer detail, by manipulating three levels (center, side, and far side) to better distinguish between these alternatives.

Finally, the effects of initial distance were mixed. There was no main effect of initial distance for heading changes (Figure 4-5), although there was for speed changes (Figure 4-8). When data were reshaped to isolate number and eccentricity, however, there was no main effect of initial distance for either variable on speed change conditions (Figures 4-9 and 4-10). These results suggest that participants respond similarly when confederates are located at 1 m or 2 m. This could be a general effect, which is consistent with the topological hypothesis—the local coupling depends on a neighbor's topological distance (nearest neighbor number), rather than their absolute metric distance. However,

because the range of distances tested was quite small, it is difficult to make such a conclusion from the data in Experiment 1. Experiment 3 will examine this metric-topological distinction in greater detail.

Overall, the results from Experiment 1 suggest that pedestrians are locally coupled to multiple neighbors and that their influence is linearly combined, weighted by eccentricity, with neighbors to the side having a weaker influence than neighbors in the direction of travel. There was no evidence for an effect of distance, although as noted this may reflect the small range of distances that were tested. Experiments 2-4 will continue to investigate the local coupling in a broader range of contexts, beginning with Experiment 2, which serves as a partial replication of Experiment 1 and a validation of the methodology used in Experiment 3.

CHAPTER 5

Experiment 2: Validating a Virtual Crowd

The purpose of the second experiment was to validate a virtual crowd paradigm by comparing participants' trajectories in matched real-world and virtual reality environments. Experiment 2A was a partial replication of Experiment 1, with participants walking together with three confederates. A subgroup of confederates received instructions to change their heading direction (turn left/right) or speed (walk faster/slower), and the number of confederates in the subgroup (0, 1, 2, 3) and the eccentricity of the single confederate (center, side) were varied. Experiment 2B was conducted with a new group of participants, who walked in virtual reality with avatars that were programmed to follow the same trajectories as the confederates in Experiment 2A. Thus, any significant differences in participant responses between Experiment 2A and 2B may be attributable to the virtual environment.

In addition, two head-mounted displays (HMDs) with different specifications (resolution, field of view, etc.) were also compared in Experiment 2B, with half the trials being performed in each HMD. This made it possible to test whether any differences with the real environment were due to a general effect of virtual reality, or to specific properties of the HMD.

To my knowledge, this is the first experiment to compare human walking trajectories in matched real and virtual environments. Pelechano, Stocker, Allbeck, & Badler (2008) published a paper entitled "Being a Part of the Crowd: Towards Validating VR Crowds Using Presence." They immersed participants in a virtual environment populated with crowds that were generated using four different pedestrian models (including a social force model and a rule-based model). Pelechano et al. (2008) made behavioral observations of the participants (i.e. that participants stepped to the side to avoid

walking into a virtual pedestrian, or turned their head to watch as virtual pedestrians passed by) and reported qualitative findings from post-experiment questionnaires and free response comments. They concluded that “some people do think about the interaction with virtual agents in a similar way as when they interact with real people.” This work differed from the present experiment in significant ways: (1) avatars in Pelechano et al. (2008) were driven by models of human crowd behavior rather than human trajectory data; (2) the virtual environment was not ambulatory; instead, participants stood in place while viewing the world through a head-mounted display; (3) participants’ behavior was not directly compared across conditions, and (4) Pelechano et al. (2008) collected only qualitative data, which makes detailed comparisons between conditions difficult or impossible. Cohen (2009) conducted a series of virtual reality experiments with avatars that behaved according to the constant bearing strategy in the steering dynamics model. He found that participants could perceive the intention of the avatars (to intercept or avoid the participant), but the avatars were not driven by human trajectory data. The present experiment compares the behavior of a participant following human confederates with virtual avatars that are driven by the confederates’ data.

Experiment 2A: Human Confederates

Methods

Participants. Ten participants, 5 female and 5 male, were recruited for Experiment 2A. None reported having any visual or motor impairment, and they were paid for their participation. Two separate groups of three trained experimenters (Group 1: 2 female and

1 male; Group 2: 1 female and 2 male) acted as confederates in the experiment. 5 participants were tested with Group 1, and 5 participants were tested with Group 2.

Apparatus. The research was conducted in the Virtual Environment Navigation Laboratory (VENLab) at Brown University. The participant and confederates walked together in a 12x14 m room, and each wore a set of wireless headphones connected to a PC that delivered pre-recorded instructions. Head position and orientation were recorded at a sampling rate of 60 Hz using a hybrid ultrasonic-inertial tracking system (IS-900, Intersense, Billerica MA) affixed to the headphones. (Note that virtual displays were not used in this experiment.)

Design and Procedure. The design and procedure of Experiment 2A were similar to those of Experiment 1, although changes were made to accommodate the smaller field-of-view of the HMDs used in Experiment 2B. In particular, only one initial distance (2 m) was tested, and the confederates' initial positions were located along a 45° arc in front of the participant with a radius of 2 m. Unlike Experiment 1, participants always began in the back position; there were no 'filler trials.' The 8 heading-change, 12 speed-change, and control conditions were the same as in Experiment 1, and there were 2 repetitions of each heading change condition (16 trials), 1 repetition of each speed change condition (12 trials), and 4 repetitions of the control condition (4 trials), for a total of 32 trials in each session. The speed change conditions were collapsed as in Experiment 1, yielding a 2 (heading or speed change) x 9 (confederate manipulations) design.

At the beginning of each trial, the four pedestrians were directed via wireless headphones to go to their starting locations, which were marked on the floor with pieces of colored tape. The participant always began in the back position; each confederates'

initial position was randomly chosen on each trial. In addition, the confederates were covertly instructed to “turn left,” “turn right,” “speed up,” or “slow down,” on each trial. The same instruction was given to a subset 1, 2, or 3 of the confederates, depending on the condition.

Each trial began when a verbal command (“Begin”) was played to all four pedestrians, which instructed them to begin walking across the room. When the first confederate crossed an invisible line located 4 m in front of the participant’s starting location, the subset of confederates heard a pure tone indicating they should initiate the speed or heading change. They maintained this speed or direction for the duration of the trial. Once the first confederate reached the opposite side of the room (about 13m), all four pedestrians received a verbal command (“End”), indicating the end of the current trial, and instructions for the next trial were presented.

Experiment 2B: Virtual Avatars

Methods

Participants. Ten participants, 5 female and 5 male, were recruited for Experiment 2B. None reported having any visual or motor impairment, and they were paid for their participation. There were no confederates in this experiment. Rather, the trajectories of confederates recorded in Exp. 2A were used to drive virtual avatars in Exp. 2B.

Apparatus and Displays. The experiment was conducted in the Virtual Environment Navigation Laboratory (VENLab) at Brown University. The participant walked in a 12 x 14 m room while viewing a virtual environment through one of two stereoscopic head-mounted displays. HMD 1 (SR-80A, Rockwell Collins, Cedar Rapids

IA) provided a $63^\circ \times 53^\circ$ (horizontal x vertical) field of view with a resolution of 1280 x 1024 pixels in each eye. HMD 2 (Rift DK1, Oculus, Irvine CA) provided a 110° diagonal field with a resolution of 640 x 800 pixels in each eye. Displays were generated on a Dell XPS workstation (Round Rock TX) at a frame rate of 60 fps, using the Vizard software package (WorldViz, Santa Monica CA). Head position and orientation were recorded at a sampling rate of 60 Hz using a hybrid ultrasonic-inertial tracking system (IS-900, Intersense, Billerica MA). Head coordinates from the tracker were used to update the display with a latency of approximately 50 ms (three frames).

The virtual environment included a visual surround consisting of a ground plane mapped with a grayscale granite texture and a blue sky. A green home pole (radius 0.2 m, height 3 m) and a red orienting pole (radius 0.2 m, height 3 m) appeared at either extreme of the room's diagonal, 12.73 m apart. The virtual avatars were generated using high-polygon ($M = 8099$, $SD = 813$), textured (2048 x 2048 pixels) 3D models and were animated with a 60 frames per second walking gait.

Design and Procedure. The design of Experiment 2A was replicated twice in Experiment 2B, once with each HMD, in separate sessions. In one session, participants wore HMD 1 (Rockwell-Collins) and viewed avatars driven by confederate data from a corresponding participant in Experiment 2A, including all trials in their original order. In the other session, the same participants wore HMD 2 (Oculus Rift) and viewed avatars driven by data from a different participant in Experiment 2A. Sessions were conducted immediately after one another, and session order was counterbalanced. There were 32 trials per session, for a total of 64 trials.

To begin each trial, participants stood at the green home pole and faced the red orientating pole. After 2 s, the three virtual pedestrians appeared in front of the participant. After an additional 1 s, a verbal command ("Begin") was delivered to the participant via wireless headphones, and the virtual pedestrians began walking. Their positions at each time step were determined by the recorded trajectories of the human confederates from the corresponding trial in Experiment 2B, at a rate of 60 Hz. Participant data were analyzed as in Experiment 2A.

Results

As in Experiment 1, we computed the mean change in lateral position or change in speed for all three display conditions: the real world ("RW") condition in Experiment 2A, the Oculus Rift ("OR") condition and the Rockwell-Collins ("RC") condition in Experiment 2B. The mean final lateral position on control trials was close to zero in each conditions (RW: $M = 0.04$ m, $SD = 0.09$ m; OR: $M = -0.02$ m, $SD = 0.25$ m; RC: $M = 0.01$ m, $SD = 0.23$ m), with no significant differences among them, $F(2,114) = 0.74$, $p > .05$. The mean final speed on control trials was slightly lower than the 1.22 m/s observed in Experiment 1 (RW: $M = 1.09$ m/s, $SD = 0.12$ m/s; OR = 1.16 m/s, $SD = 0.10$ m/s; RC = 1.17 m/s, $SD = 0.09$ m/s). There was a significant effect of display conditions, $F(2,114) = 5.97$, $p < .01$, and Bonferroni post-hoc corrections revealed that walking speed in the real-world was actually slower than in the two virtual reality conditions ($p < .05$), which did not differ from one another ($p > .05$).

As in Experiment 1, participants responded to changes in their neighbors' behavior in all three display conditions. Overall, the participant turned left (or right) when one or more confederates did so, as shown in Figure 5-1. A two-way ANOVA on final lateral

position, with manipulation condition and display condition (RW, OR, RC) as factors, confirmed a significant main effect of manipulation ($F(8,243) = 80.34, p < .001$), no main effect of display condition ($F(2,243) = 0.32, p > .05$), and a significant manipulation-by-display interaction ($F(16,243) = 3.51, p < .001$). To isolate the effects of number and eccentricity, the final position data were collapsed as in Experiment 1. A two-way ANOVA on the number conditions revealed that there was a main effect of number ($F(3,258) = 66.34, p < .001$), a main effect of display ($F(2,258) = 20.67, p < .001$), and no interaction ($F(6,258) = 1.74, p > .05$). Bonferroni post-hoc tests revealed that the differences between all number conditions were significant ($p < .05$); participants' responses increased with each additional neighbor that was manipulated, from the control condition ($M = 0.01$ m, $SD = 0.12$ m) to 1 ($M = 0.16$ m, $SD = 0.22$ m), 2 ($M = 0.39$ m, $SD = 0.34$ m), and 3 ($M = 0.61$ m, $SD = 0.27$ m) neighbors. Participants responded significantly less ($p < .001$) in the real-world condition ($M = 0.16$ m, $SD = 0.32$ m) than in either of the head-mounted display conditions (OR: $M = 0.35$ m, $SD = 0.30$ m; RC: $M = 0.37$ m, $SD = 0.33$ m), which did not differ from one another ($p > .05$). A similar two-way ANOVA on the eccentricity conditions revealed that there was a main effect of eccentricity ($F(2,141) = 7.01, p < .01$), a main effect of display condition ($F(2,141) = 5.63, p < .01$), and no interaction ($F(4,141) = 0.64, p > .05$). Bonferroni post-hoc tests showed that participants responded significantly more to a single confederate in the center ($M = 0.15$ m/s, $SD = 0.22$ m/s) than a single confederate on the side ($M = 0.17$ m/s, $SD = 0.23$ m/s), $p < .01$. The response to these manipulations did not differ from one another ($p > .05$), but they were both significantly greater than the control condition ($M = 0.01$ m/s, $SD = 0.12$ m/s), $p < .01$. Bonferroni post-

hoc tests showed that the differences between display conditions were not significant, $p > .05$.

Participants' final speeds are shown in Figure 5-2. A two-way ANOVA on final speed confirmed a significant main effect of manipulation ($F(8,232) = 20.70, p < .001$), a significant main effect of display ($F(2,232) = 6.38, p < .01$), and no interaction ($F(16,232) = 0.88, p > .05$). Again, to isolate the effects of number and eccentricity, the final speed data were collapsed. A two-way ANOVA on the number conditions revealed that there was a significant main effect of number ($F(3,218) = 2.76, p < .05$), no main effect of display condition ($F(2,218) = 2.11, p > .05$), and no interaction ($F(6,218) = 0.29, p > .05$). Bonferroni post-hoc tests revealed that participants responded significantly less when 2 neighbors were manipulated ($M = -0.09$ m/s, $SD = 0.07$ m/s) than when 1 neighbor was manipulated ($M = 0.01$ m/s, $SD = 0.14$ m/s), $p < .05$; all other differences were not significant, $p > .05$. A similar two-way ANOVA on the eccentricity conditions found no main effect of eccentricity ($F(2,135) = 0.37, p > .05$), no main effect of display condition ($F(2,135) = 1.31, p > .05$), and no interaction ($F(4,135) = 0.32, p > 0.05$).

Taken together, these results indicate that changes in final position were significantly more pronounced in the two HMD conditions in Experiment 2B (OR, RC) than in the real-world condition in Experiment 2A, as well as in Experiment 1. In contrast, changes in final speed were significantly greater in the RW condition than in the HMD conditions (OR, RC), but only for 'speed up' trials.

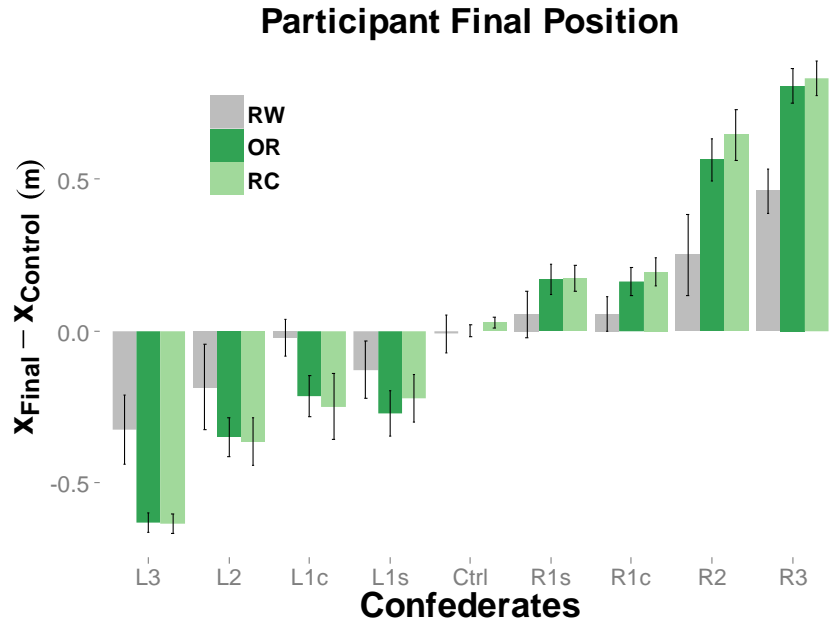


Figure 5-1: Mean changes in participants' final lateral position in Experiment 2A (grey) and Experiment 2B (green).

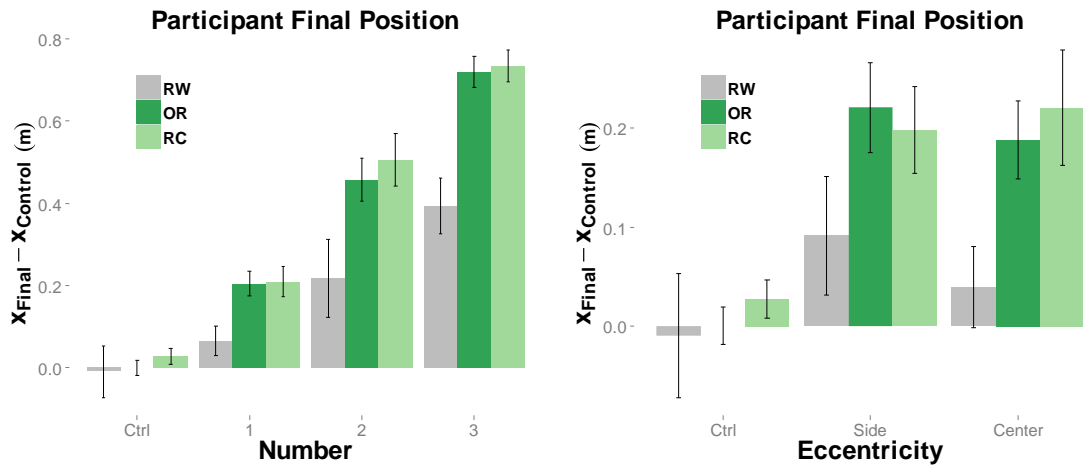


Figure 5-2: Mean changes in participants' final position in Experiment 2A (grey) and Experiment 2B (green), separating the effects of number and eccentricity.

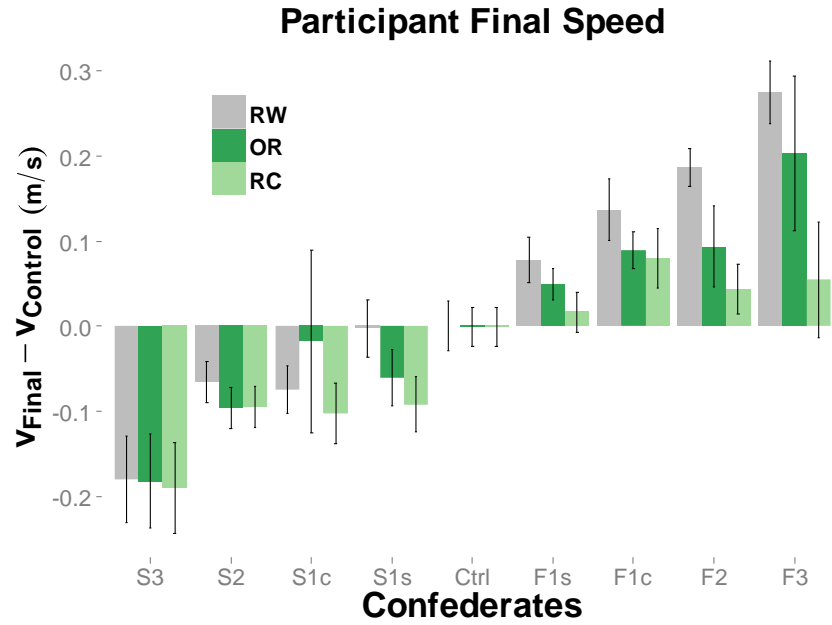


Figure 5-3: Mean changes in participants' final speed in Experiment 2A (grey) and Experiment 2B (green).

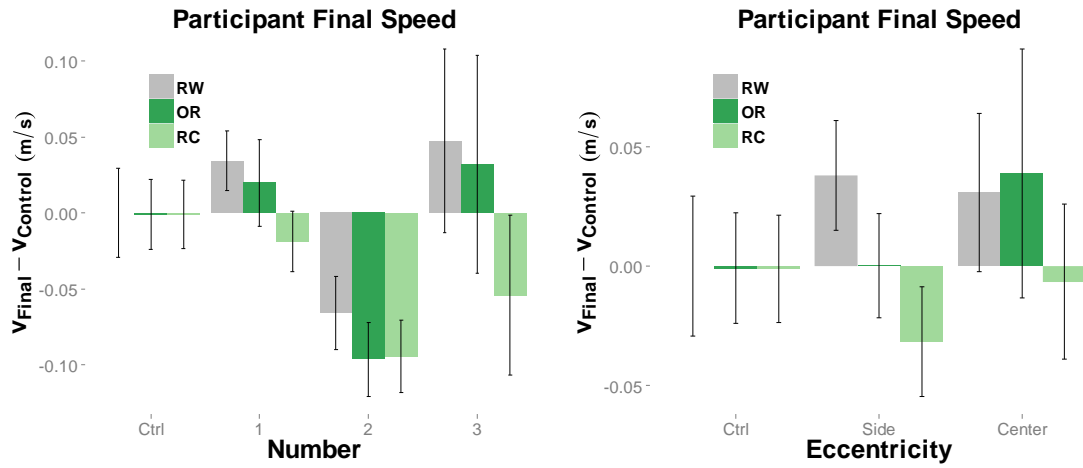


Figure 5-4: Mean changes in participants' final speed in Experiment 2A (grey) and Experiment 2B (green), separating the effects of number and eccentricity.

Discussion

Overall, the pattern of results was qualitatively similar whether participants were responding to movements of real humans or virtual avatars. When a subgroup of neighbors changed their heading or speed, participants followed suit, and the magnitude of the response was again proportional to the number of neighbors in the subgroup. The results clearly imply that such manipulations of virtual pedestrians will yield findings similar to experiments using real pedestrians. The potential to deploy virtual crowds will enable more flexible and powerful experiments on crowd behavior.

There were, naturally, a few differences between the display conditions that should be discussed. First, participants' responses in the turn left and turn right conditions were greater in the HMDs (Experiment 2B) than in the real world (Experiment 2A). This indicates that participants were either more sensitive to changes in the virtual pedestrians' heading direction(s), more likely to respond when a change in the virtual pedestrians' heading direction(s) was detected, or both. It is possible that participants were concerned about colliding with one of the confederates in the RW condition, which may have inhibited or reduced their response to changes in the confederates' heading direction(s). Collisions may have been a lesser concern in virtual reality, although participants were instructed to "treat the virtual people as if they were real people" and to avoid collisions with them. Another possibility is that participants in the RW condition used peripheral vision of the physical room as a fixed visual reference frame, which stabilized their walking trajectories and reduced their response. By contrast, in the two virtual reality conditions, peripheral vision was limited by the relatively small field of the views of the two head-mounted

displays, and the physical boundaries of the room such as walls and corners were not present, only an endless ground plane extending to the horizon.

In contrast, participants were more likely to speed up in the real world (Experiment 2A) than in the Oculus Rift HMD, which in turn was greater than the Rockwell Collins HMD (Experiment 2B). We have generally observed a reluctance to accelerate in response to a speed-up manipulation in other experiments (Rio, Rhea, & Warren, 2014), whereas a slow-down manipulation increases the chances of collision and elicits a greater response. Participants may have felt more social pressure to walk with the group in Experiment 2A, where the confederates were real people in the room, than in Experiment 2B, where the avatars were computer-generated virtual pedestrians. This ‘pull’ to be with the group may help explain why participants were less likely to accelerate in response to a speed up manipulation in Experiment 2B.

Participants behaved in largely the same way across the two virtual reality conditions (OR and RC), despite differences in the specifications of the two head-mounted displays. One exception is that participants showed a reduced speed response in the RC condition (see Figure 5-3) compared to the OR and RW conditions. This may be due to the reduced field of view of the Rockwell-Collins HMD compared to the Oculus Rift; participants may have been less able to detect when avatars, especially those in the periphery, increased their speed, thus reducing their overall response in the RC condition. Resolution did not seem to play an important role; even though the Oculus Rift had a lower resolution than the Rockwell-Collins HMD, performance was comparable in both conditions.

Overall, despite these minor differences between display conditions, participants' responses were qualitatively similar across both experiments, indicating that participants treated the avatars in the same manner as they treated the real confederates. This serves to validate the use of virtual pedestrians, which will serve a central role in Experiment 3.

CHAPTER 6

Experiment 3: Walking in a Virtual Crowd

In Experiment 1, participants in a small crowd responded to systematic manipulations of their confederate neighbors, providing evidence for a local visual coupling. Experiment 2 validated the use of virtual reality as a tool to investigate crowd behavior, by showing that participants responded in largely the same way to matched human confederates and virtual avatars. Experiment 3 combines these methodologies, by immersing participants in a large crowd (N=30) of virtual pedestrians and systematically manipulating their behavior. This novel approach makes it possible to map out the local visual coupling in greater detail, and further test hypotheses regarding the interactions between neighbors.

Experiment 3 will focus on two fundamental properties of the local coupling between neighbors; namely, the neighborhood structure and the spatial pooling of multiple neighbors. There are two main hypotheses regarding the neighborhood structure, and they hold that it is defined by metric or topological distance, respectively. The metric view (e.g. Huth & Wissel, 1994) holds that people in a crowd are coupled to all neighbors within a fixed radius, defined in absolute distance units like meters or relative units like body length or stride length. The topological account (e.g. Ballerini et al., 2008), on the other hand, holds that people in a crowd are coupled to a fixed number of their nearest neighbors, regardless of their absolute distance. The “topological distance” is an ordinal measure that refers to the number of neighbors between two pedestrians, such as the first nearest neighbor or the seventh nearest neighbor. Figure 6-1 shows a schematic diagram of these two views.

These two hypotheses make different predictions about the effect of density. In a metric neighborhood, as density increases, more neighbors will fall within the fixed radius;

thus, the number of influential neighbors increases with density. In a topological neighborhood, by contrast, the number of influential neighbors is fixed, and does not change with density.




Figure 6-1: Metric (left) and topological (right) neighborhood structure. The focal pedestrian (black) is locally coupled to some neighbors (blue), but not others (grey). Pedestrians outside the field of view (light grey with slashes) are not considered.

A second set of questions can be asked regarding the spatial pooling of influences from multiple neighbors. Are all neighbors within the metric or topological neighborhood equally influential, or does the local coupling strength decay with their distance or eccentricity? Is the local coupling stronger for nearby neighbors, or for neighbors with low eccentricity? The diagrams on the following page show what this would look like for distance (Figure 6-2) and eccentricity (Figure 6-3), for both the metric and topological accounts of neighborhood structure, respectively.




Figure 6-2: Decaying local coupling strength as a function of metric distance (left) and topological distance (right). Shade represents the strength of the local coupling; nearby neighbors are darker, reflecting a stronger coupling.




Figure 6-3: Decaying local coupling strength as a function of eccentricity, for metric (left) and topological (right) neighborhoods. Shade represents the strength of the local coupling; neighbors with small eccentricities are darker, reflecting a stronger coupling.

With these hypotheses in mind, Experiment 3 consisted of three sub-experiments. Experiment 3A varied the number of neighbors in the manipulated subset (0, 3, 6, 9, 12) and the distance ‘zone’ that those neighbors were in (near zone: 1.5 m, far zone: 3.5 m). This tests the hypothesis that coupling strength decreases with distance (Figure 6-2). If participants respond more to manipulations in near zone than in the far zone, this suggests that neighbors in the near zone have a greater influence and therefore a stronger coupling. Alternatively, if there is no difference between zones, this suggests that neighbors are treated equivalently, regardless of their distance.

Experiment also varied the number of neighbors in the manipulated subset (0, 3, 6, 9, 12) as well as the density of the crowd (mean initial distance = 2.5 m or 5.5 m). This sub-experiment tests the metric and topological hypotheses regarding neighborhood structure. The topological hypothesis predicts that response should be the same in the two density conditions, because the coupling depends on topological distance (number of nearest neighbors), which does not change with crowd density. On the other hand, the metric hypothesis predicts that response should decrease in the lower density condition, because more neighbors will fall outside the cutoff distance (Figure 6-1).

Finally, Experiment 3C varied the eccentricity of the manipulated subset in 30° ‘wedges’ of the visual field at 15° intervals, as well as the distance zone of the manipulated subset (near zone: 1.5 m, far zone: 3.5 m). This sub-experiment tests the hypothesis that coupling strength varies with neighbors’ eccentricity. This hypothesis predicts that manipulations in ‘wedges’ containing neighbors at greater eccentricities should elicit weaker participant responses than manipulations of neighbors at lower eccentricities.

General Methods

Apparatus

The participant walked in the VENLab while viewing a virtual environment stereoscopically in an Oculus Rift HMD (Rift DK1, Oculus, Irvine CA) that provided a 110° diagonal field of view with a resolution of 640 x 800 pixels in each eye. The rest of the apparatus was the same as in the preceding experiment (Experiment 2B).

Displays

The virtual environment was the same as in the preceding experiment (Experiment 2B), except the three virtual avatars were replaced with a virtual crowd. The crowd consisted of 30 virtual pedestrians that were generated using high-polygon ($M = 8073$, $SD = 780$), textured (2048 x 2048 pixels) 3D models; 12 of them were rendered within a participant's typical horizontal field of view in the HMD (90°), while the other 18 were placed outside the field of view to enhance the sense of immersion. The latter could be seen if participants turned their head (which was infrequent), but they were not manipulated and will be ignored in subsequent analyses. The virtual pedestrian models were animated with a walking gait, with an update rate of 60 frames per second. The beginning of the gait cycle of each virtual pedestrian was randomly offset, to reduce the appearance of synchronization. The virtual crowd contained equal numbers of men and women and diverse races and ethnicities. Figure 6-4 shows a sample of the displays.



Figure 6-4: Virtual displays from Experiment 3.

The initial positions of the virtual pedestrians were distributed in two ‘zones,’ which were defined by circles with radii 1.5 m and 3.5 m (high density) or 3.5 m and 7.5 m (low density). Of the 12 manipulated virtual pedestrians, 5 were initially positioned on the near circle (1.5 m in the high density condition) and 7 were positioned on the far circle (3.5 m in the high density condition) at equal intervals along a 90° arc centered about the line connecting the home pole to the orienting pole (the participants’ direction of travel if they were to walk straight ahead). The other 18 virtual pedestrians were equally distributed along the two remaining 270° arcs. These initial positions were randomly ‘jittered’ in polar coordinates, independently for distance (r) and eccentricity (θ). For each virtual pedestrian, the jitter was randomly chosen from two Gaussian distributions, with $M = 0$ m, $SD = 0.15$ m for distance r and $M = 0^\circ$, $SD = 8^\circ$ for eccentricity θ .

Virtual pedestrians accelerated from a standstill (0 m/s) to a speed of 1.0 m/s on a straight path over a period of 3 s using an ogive function (normal cumulative distribution function with $\mu = 0$ and $\sigma = 0.5$ s) fitted to data from the confederates in Experiment 1. After this initial acceleration, the virtual pedestrians continued to travel with the same heading and speed for 2 s. At this point, on some trials, a subset of the virtual pedestrians

gradually changed their heading or speed over a period of 0.5 s, using an ogive function similar to that for the initial acceleration (normal cumulative distribution function with $\mu = 0$ and $\sigma = 0.083$ s). Heading changes consisted of a 10° change in heading direction; speed changes consisted of a ± 0.3 m/s change in speed, from the baseline speed (1.0 m/s) to 0.7 m/s or 1.3 m/s for slower and faster trials, respectively. The trajectories of all virtual pedestrians were generated in advance. Trials in which the separation between any two virtual pedestrians became less than 0.4 m were removed to eliminate collisions, and were regenerated until no such collisions occurred.

Procedure

Participants were instructed to walk as naturally as possible, to treat the virtual pedestrians as if they were real people, and to stay together with the crowd. To begin each trial, participants stood at the green home pole and faced the red orientating pole. After 2 s, the 12 virtual pedestrians appeared in front of the participant. 1 s later, a verbal command ("Begin") was heard over a loudspeaker, and the virtual pedestrians began walking. When the participant had walked about 12 m, a second verbal command ("End") signaled the end of the current trial. The participant then began the next trial by walking to another green pole, which appeared near their stopping point on the previous trial.

Data Analysis

Data were processed and analyzed as in the previous experiments.

Experiment 3A

Experiment 3A was designed to test hypotheses about spatial pooling, by manipulating the number of neighbors in the subgroup and their distance from the

participant. If neighbor influence is simply averaged, as in the SPP model, then participant response should increase proportionally to the number in the subgroup, regardless of their distance. If neighbor influence is a weighted linear combination that depends on distance, then the participant response should be greater for manipulations in the near zone than the same manipulations in the far zone. Finally, if the neighborhood is topological, neighbor influence may be linearly combined up to some maximum number.

Methods

Participants. Ten participants, 6 female and 4 male, were recruited for Experiment 3A, with the same criteria as in Experiments 1 and 2.

Design. Experiment 3A had a factorial design: 2 (heading or speed change) x 2 direction of change (left/right or fast/slow) x 2 distance (near or far zone) x 5 number (0, 3, 6, 9, 12 in subset), for a total of 40 conditions. Only the high-density condition was tested in this sub-experiment, so the near zone was centered at 1.5 m and the far zone at 3.5 m. There were 4 repetitions per condition, for a total of 160 trials per participant. These trials were run across two sessions on consecutive days, in a randomized order.

Procedure. The manipulated subset (0, 3, 6, 9, or 12 virtual neighbors) were randomly chosen from one of the two distance ‘zones’ (near: 1.5 m, far: 3.5 m). Because there were 5 neighbors in the near zone and 7 in the far zone, there was some overlap between zones for subsets of size 6, 9, and 12. In these cases, the extra neighbors were randomly selected from the other zone, so the subset was ‘biased’ toward the first zone. In the N=12 condition, the subset included all virtual pedestrians in both rows, so it was equivalent for the ‘near’ and ‘far’ condition.

Results

As before, participants' changes in final lateral position and final speed were computed relative to control trials. The mean final lateral position in control trials across all participants was +0.07 m (SD = 0.08 m), suggesting a tendency to drift to the right, as in the previous experiments. A one-sample *t*-test showed that this drift was significant, $t(9) = 2.87, p < .05$. The mean final speed in control trials across all participants was 1.17 m/s (SD = 0.02 m/s), which is similar to that observed in Experiment 1 (1.22 m/s).

Figure 6-5 shows the mean change in position for each of the 9 neighbor conditions, across the two distance zone conditions. A two-way ANOVA on change in position shows a main effect of number ($F(8,162) = 300.59, p < .001$), no effect of distance zone ($F(1,162) = 1.79, p > .05$), but a significant interaction ($F(8,162) = 25.85, p < .001$). To better interpret these results, the turn left and turn right manipulations were collapsed by multiplying the final lateral position on turn left trials by negative one. This yielded the unsigned magnitude of the change in speed, regardless of direction, which provides a better estimate of the effect of the variables being investigated, namely number and distance zone. Figure 6-6 shows these collapsed data.

A two-way ANOVA on the collapsed final lateral position data revealed a main effect of number ($F(4,170) = 151.89, p < .001$), a main effect of distance ($F(1,170) = 96.72, p < .001$), and a significant interaction ($F(4,170) = 22.95, p < .001$). Change in lateral position increased with the number of neighbors in the subgroup, from 0 in the control condition (M = 0.02 m, SD = 0.11 m) to 3 virtual pedestrians (M = 0.22 m, SD = 0.22 m) through 6 (M = 0.44 m, SD = 0.38 m), 9 (M = 0.78 m, SD = 0.22 m), and 12 virtual pedestrians (M = 0.92 m, SD = 0.16 m). Bonferroni post-hoc corrections confirmed that

these differences were all significant, $p < .05$, except for between 9 and 12 neighbors, which were not significantly different from one another, $p > .05$. Change in lateral position was greater when the subgroup was closer to the near zone ($M = 0.65$ m, $SD = 0.35$ m) than the far zone ($M = 0.40$ m, $SD = 0.41$ m). A simple effects analysis was performed to further understand the number-by-zone interaction. For a given number of neighbors, participants responded more to manipulations biased toward the near zone than the far zone for 3, 6, and 9 neighbors ($p < .05$), but not for 12 neighbors or the control condition ($p > .05$). Note that in these latter two conditions, the near and far subgroups were identical, so no difference was expected. These differences account for the significant interaction that was observed.

Similar results were obtained for changes in speed. Figure 6-7 shows the mean change in speed for each of the 9 neighbor conditions, across the two distance zone conditions. There was a main effect of number ($F(8,162) = 79.57$, $p < .001$), no effect of distance zone ($F(1,162) = 1.11$, $p > .05$), but a significant interaction ($F(8,162) = 7.48$, $p < .001$). As before, slow down and speed up trials were collapsed by multiplying the final speed on slow down trials by negative one, yielding the unsigned magnitude of the change in speed. Figure 6-8 shows these collapsed data. A two-way ANOVA on the collapsed final speed data revealed a main effect of number ($F(4,170) = 34.28$, $p < .001$), a main effect of distance zone ($F(1,170) = 24.33$, $p < .001$), and a significant interaction ($F(4,170) = 0.18$, $p < .01$). Change in speed increased with the number of neighbors in the subgroup, from 0 in the control condition ($M = -0.01$ m/s, $SD = 0.03$ m/s) to 3 neighbors ($M = 0.06$ m/s, $SD = 0.07$ m/s) through 6 ($M = 0.13$ m/s, $SD = 0.13$ m/s), 9 ($M = 0.20$ m/s, $SD = 0.13$ m/s), and 12 ($M = 0.26$ m/s, $SD = 0.12$ m/s). Bonferroni post-hoc corrections confirmed that these

differences were all significant, $p < .05$, except for three of them: between 0 and 3 neighbors, between 6 and 9, and between 9 and 12. Change in lateral position was greater when the subgroup was closer to the near zone ($M = 0.18$ m/s, $SD = 0.14$ m/s) than the far zone ($M = 0.11$ m/s, $SD = 0.13$ m/s). Again, a simple effects analysis revealed that for a given number of neighbors, participants responded more to manipulations biased toward the near zone than the far zone for 3, 6, and 9 neighbors ($p < .05$), but not for 12 neighbors or the control condition ($p > .05$), which accounts for the interaction.

Overall, the results of Experiment 3A indicate that, for both speed and heading changes, participant response increases with the size of the manipulated subset, up to at least 12 neighbors. For both dependent variables, the effect of subset number was greater for the near zone than the far zone. That is, a given subset of neighbors in the near zone elicited a larger participant response than an equally-sized subset in the far zone (compare $N=3$ and $N=6$ in Figures 6-6 and 6-8). This is consistent with the hypothesis that the local coupling can be explained by a linear combination of neighbor influences, weighted by their distance.

There are alternative explanations, however. If the neighborhood structure is purely topological, with a maximum limit of less than 12, some of the neighbors in the far row might fall outside of it. Participants might respond equally to the 5 neighbors in the near row and, say, the 2 or 3 neighbors in the far row who happen to be closest to them. This alternative could be tested by (a) systematically manipulating either the nearest and farthest members of the 'far' row, or (b) splitting the 'far' row in half, with some neighbors at the original 'far' distance and some at an 'even farther' distance, and then

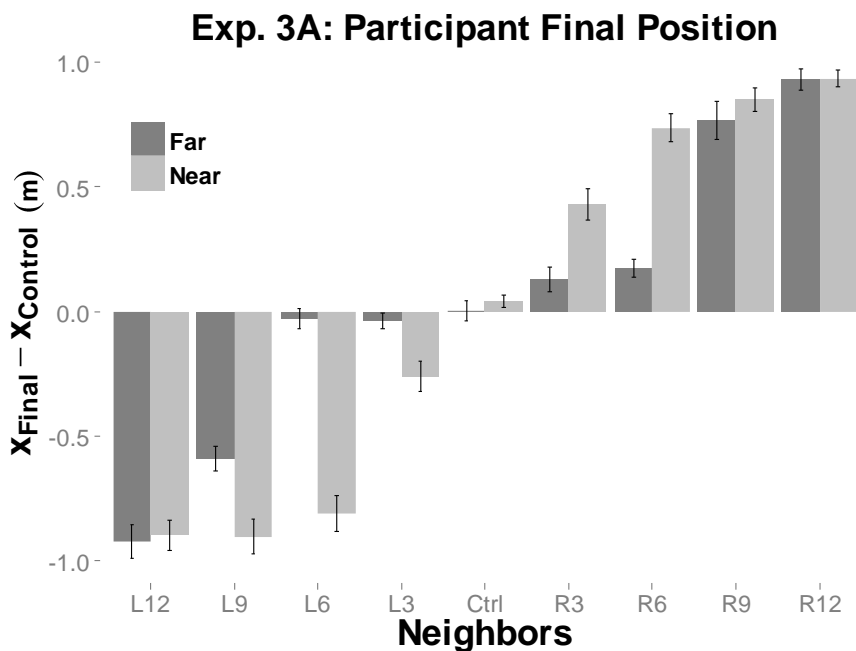


Figure 6-5: Mean changes in participants' final lateral position in Experiment 3A.

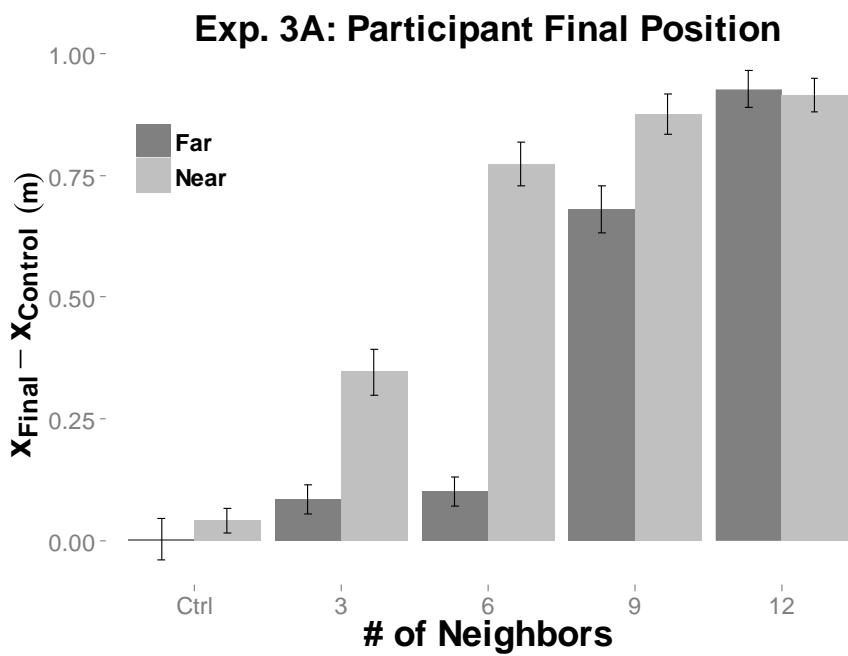


Figure 6-6: Mean changes in participants' collapsed final position in Experiment 3A. Note that the near and far zones converge for $N=12$ because these conditions are identical (see Procedure).

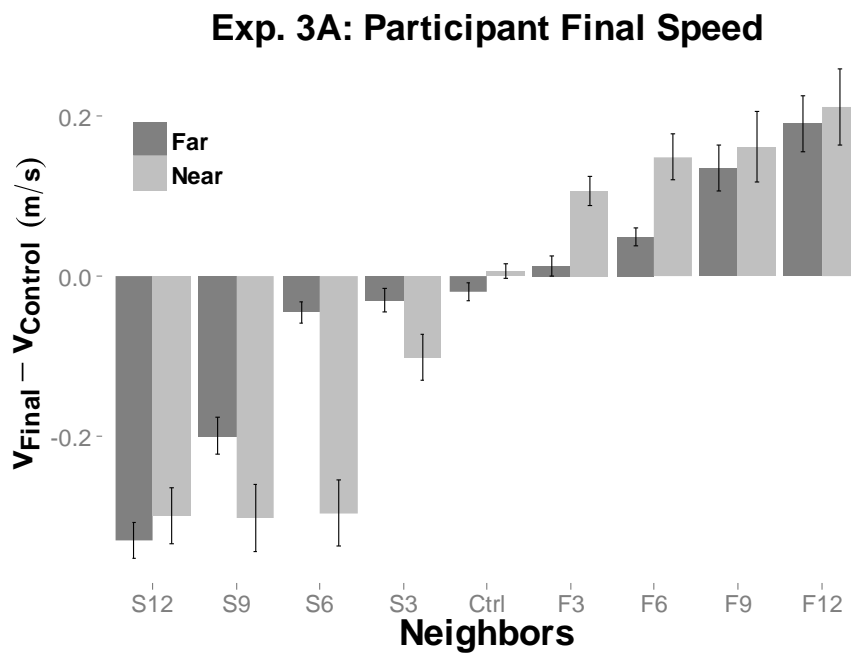


Figure 6-7: Mean changes in participants' final speed in Experiment 3A.

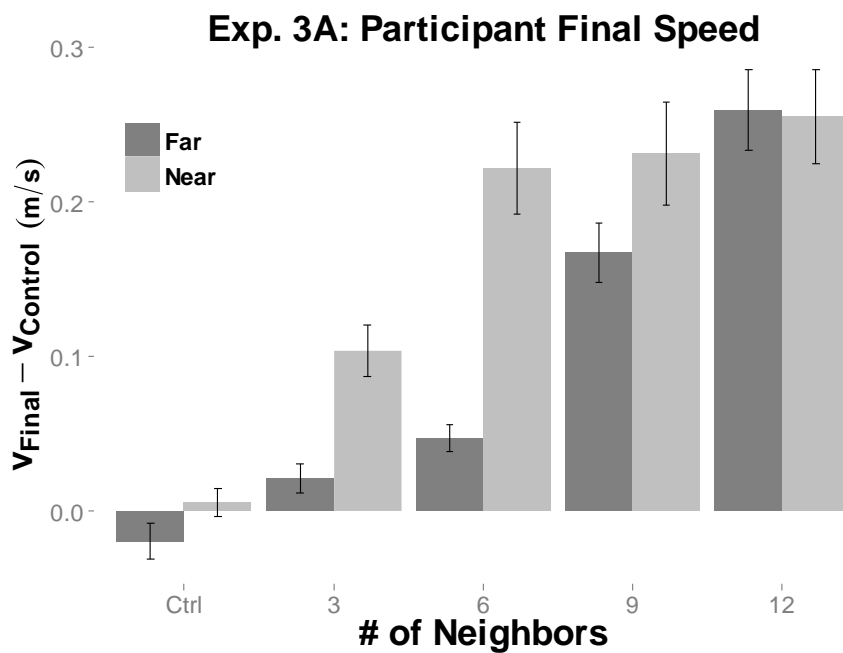


Figure 6-8: Mean changes in participants' collapsed final speed in Experiment 3A. Note that the near and far zones converge for $N=12$ because these conditions are identical (see Procedure).

systematically manipulating either the ‘far’ row or the ‘even farther’ row. Similarly, another alternative account is that there is a purely metric neighborhood, where the maximum limit falls exactly at 3.5 m; because the initial positions of the virtual pedestrians were jittered, some neighbors in the far row would fall inside this limit and some would fall outside. This alternative could be tested by systematically varying the distance to the far row; in part, this was the purpose of Experiment 3C.

Finally, the results of Experiment 3A do not provide strong evidence for a maximum limit to the number of influential neighbors. Participant response increased up to 12 neighbors (although the difference between $N=9$ and $N=12$ was not significant for both speed and heading). This provides some evidence against a purely topological neighborhood with a maximum limit less than 12 neighbors. Experiment 3B will more precisely test the metric and topological hypotheses, by varying the overall density of the crowd.

Experiment 3B

Experiment 3B was designed to test the hypothesis of a metric or topological neighborhood of interaction, by manipulating the overall crowd density, as well as the number of manipulated neighbors in the subgroup. A purely metric neighborhood implies that people interact with all neighbors who fall within a fixed metric distance; for example, 3 m. Pedestrians outside this radius should not have an effect. A purely topological neighborhood, on the other hand, implies that people interact with a fixed number of their nearest neighbors, regardless of their metric distance; the number 7 is often assumed, because there is evidence from starling flocks that birds are topologically coupled to 7 of

their nearest neighbors (Ballerini et al., 2008). Thus, these two accounts make different predictions. On the purely metric account, participant response should decrease with decreasing density, because neighbors who fall outside the fixed radius no longer have an influence. On the purely topological account, density should have no effect on participant response, because the nearest neighbors have the same influence whether or not they fall outside the metric radius. By varying density, we aimed to provide experimental evidence regarding these hypotheses.

Methods

Participants. Ten participants, 5 female and 5 male, were recruited for Experiment 3B, with the same criteria as in Experiments 1 and 2.

Design. Experiment 3B had a 2 (heading or speed change) x 2 direction of change (left/right or fast/slow) x 2 initial density (2.5 m or 5.5 m initial distance) x 5 subset (0, 3, 6, 9, 12 neighbors) factorial design, yielding 40 conditions. There were 4 repetitions per condition, for a total of 160 trials per participant. These trials were run across two sessions on consecutive days, in a randomized order.

Procedure. The procedure was the same as in the preceding experiment, with two exceptions. First, the virtual pedestrians were presented in two initial density conditions. The high density condition was identical to Experiment 3A, with 5 pedestrians at a distance of 1.5 m (near zone) and 7 pedestrians at a distance of 2.5 m (far zone). The low density condition consisted of 5 pedestrians at a distance of 3.5 m (near zone) and 7 pedestrians at a distance of 7.5 m (far zone). The 18 additional virtual pedestrians outside the field of view were distributed at the same distances, respectively. Second, the manipulated subset

(0, 3, 6, 9, 12) was randomly chosen from all 12 neighbors across both the near and far zone.

Results

Participants' changes in final lateral position and final speed on each trial were computed relative to control trials as in Experiment 3A. The mean final lateral position on control trials for all participants was +0.06 m (SD = 0.11 m), again suggesting a tendency to drift to the right. A one-sample *t*-test showed that this drift was not significant, $t(9) = 1.81, p > .05$. The mean final speed on control trials for all participants was 1.20 m/s (SD = 0.04 m/s).

Figure 6-9 shows the mean change in lateral position for each of the 9 neighbor conditions, and the two crowd density conditions. A two-way ANOVA on change in position shows a main effect of number ($F(8,162) = 131.01, p < .001$) and a main effect of density ($F(1,162) = 1.62, p < .05$), but no interaction ($F(8,162) = 1.62, p > .05$). It appears that the higher density produced greater heading changes in the turn left conditions, but not in the turn right conditions. Nevertheless, to better isolate the effects of number and density, turn left and turn right trials were collapsed, as in the preceding experiment. The collapsed data appear in Figure 6-10. A two-way ANOVA on these collapsed lateral positions finds a main effect of number ($F(4,170) = 71.28, p < .001$), no effect of density ($F(1,170) = 1.75, p > .05$), and a significant interaction ($F(4,170) = 2.33, p < .01$). Change in lateral position increased with the number of neighbors in the subset, from 0 in the control condition (M = 0.02 m, SD = 0.10 m) to 3 (M = 0.20 m, SD = 0.17 m) through 6 (M = 0.36 m, SD = 0.23 m), 9 (M = 0.61 m, SD = 0.23 m), and 12 neighbors (M = 0.80 m, SD = 0.24 m). Bonferroni post-hoc corrections confirmed that these differences were all significant, $p < .05$, except

between the control condition and 3 neighbors, $p > .05$. A simple effects analysis showed that participants responded more to manipulations at high density for 12 neighbors ($p < .05$), but there was no difference for other manipulations ($p > .05$); this, the effect of density appears to increase with the number of manipulated neighbors, with greater heading response at high density.

A similar pattern of results was obtained for change in final speed. Data are shown in Figure 6-11. There was a main effect of number ($F(8,162) = 79.29, p < .001$), no effect of density ($F(1,162) = 0.01, p > .05$), and a significant interaction ($F(8,162) = 2.70, p < .01$). In this case, the higher density produced greater speed changes in both the slower and faster conditions. Slow down and speed up trials were collapsed, as before, as shown in Figure 6-12. A two-way ANOVA on the collapsed final speed data revealed main effects of number ($F(4,170) = 36.05, p < .001$) and density ($F(1,170) = 11.59, p < .001$), with no interaction ($F(4,170) = 0.81, p > 0.5$). Change in speed increased with the number of neighbors in the subset, from 0 in the control condition ($M = 0.00$ m/s, $SD = 0.04$ m/s) to 3 ($M = 0.04$ m/s, $SD = 0.08$ m/s) through 6 ($M = 0.14$ m/s, $SD = 0.12$ m/s), 9 ($M = 0.19$ m/s, $SD = 0.10$ m/s), and 12 neighbors ($M = 0.25$ m/s, $SD = 0.11$ m/s). Bonferroni post-hoc corrections confirmed that these differences were significant, $p < .05$, except between 6 and 9, and between 9 and 12, $p > .05$. These results show that participants have a greater speed response overall in the high density condition.

The results of Experiment 3B thus confirm that the influences from multiple neighbors are linearly combined. In addition, they demonstrate that participants exhibit a greater response to both speed changes and heading changes in a higher density crowd. This implies a stronger effect when more neighbors are present within a given radius,




Figure 6-9: Mean changes in participants' final lateral position in Experiment 3B.

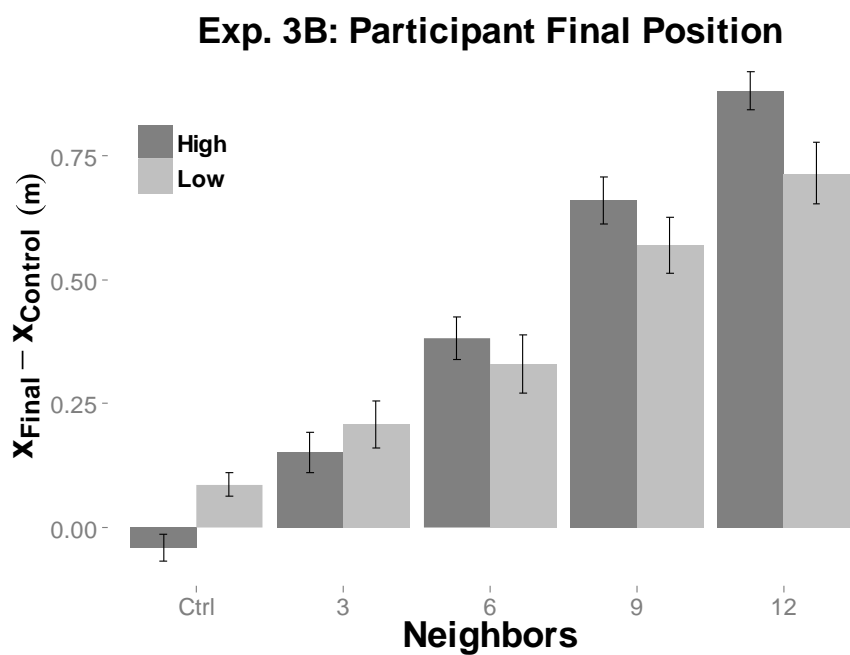


Figure 6-10: Mean changes in participants' collapsed final position in Experiment 3B.

Exp. 3B: Participant Final Speed




Figure 6-11: Mean changes in participants' final speed in Experiment 3B.

Exp. 3B: Participant Final Speed




Figure 6-12: Mean changes in participants' collapsed final speed in Experiment 3B.

consistent with the metric neighborhood hypothesis. On the other hand, it is contrary to the topological neighborhood hypothesis, which predicts that a fixed number of nearest neighbors should have a similar influence regardless of their density. Thus, these results provide evidence against a purely topological account, in which there is a sharp cutoff after a fixed number of nearest neighbors.

Experiment 3C

Experiment 3C was designed to test the hypothesis that spatial pooling depends on the eccentricity of neighbors, as well as their distance. The results of Experiment 1 suggest that neighbors aligned with the direction of travel have a larger influence than pedestrians with larger eccentricities. A number of factors limit the conclusions that can be drawn from Experiment 1: (1) there was a relatively small number of neighbors ($N=3$), (2) only two eccentricities were tested (0° and $\pm 22.5^\circ$), which leaves open the possibility that the 0° may be qualitatively different than other eccentricities, and (3) only one distance (2 m) was tested. Thus, additional data are necessary to determine whether spatial pooling depends on eccentricity in the general case. Experiment 3C was designed to provide such evidence, by manipulated the eccentricity of a subgroup of neighbors (centered on 0° , $\pm 15^\circ$, or $\pm 30^\circ$), as well as the distance zone of the subgroup (1.5 m or 3.5 m).

Methods

Participants. Ten participants, 6 female and 4 male, were recruited for Experiment 3C, with the same criteria as in the previous experiments.

Design. Experiment 3C had a 2 (heading or speed change) x 2 direction of change (left/right or fast/slow) x 5 eccentricity (-30° , -15° , 0° , $+15^\circ$, $+30^\circ$) x 2 distance zone (near

or far) factorial design, yielding 40 conditions. There were also 4 control conditions, in which no virtual pedestrians were manipulated. Thus, there were a total of 44 conditions, with 4 repetitions per condition, for a total of 176 trials per participant. These trials were run across two sessions on consecutive days, in a randomized order.

Procedure. As explained in General Methods, the 12 neighbors were distributed in two distance zones (near zone: 5 pedestrians centered at 1.5 m; far zone: 7 pedestrians at 3.5 m), with jittered positions. To vary eccentricity, the manipulated subgroup lay within five 30° “wedges” of the visual field, at 15° intervals, from left to right: -30° (spanning -45° to -15°), -15° (-30° to 0°), 0° (-15° to +15°), +15° (0° to +30°), and +30° (+15° to +45°). In the “near” condition, only virtual neighbors in the near zone (1.5 m) were manipulated; in the “far” condition, only virtual pedestrians in the far zone (3.5 m) were manipulated. The manipulated subgroup consisted of about one or two neighbors on each trial; on average, more neighbors were manipulated in the “far” conditions ($M = 2.18$, $SD = 0.81$) than the “near” conditions ($M = 1.23$, $SD = 0.73$).

Results

We computed the change in final lateral position or final speed on each trial relative to control trials, as in Experiment 3A and 3B. The mean final lateral position on control trials for all participants was +0.04 m ($SD = 0.07$ m), suggesting a tendency to drift to the right. A one-sample t -test showed that this drift was not significant, $t(9) = 1.60$, $p < .05$. The mean of final speed on control trials for all participants was 1.20 m/s ($SD = 0.02$ m/s).

The mean changes in lateral position for the 10 left/right eccentricity conditions, and the two distance conditions, appear in Figure 6-13. A two-way ANOVA on change in position shows a main effect of eccentricity ($F(9,180) = 5.07$, $p < .001$), no main effect of

distance zone ($F(1,180) = 0.01, p > .05$), and a significant interaction ($F(9,180) = 3.14, p < .01$). As before, to better isolate the effects of eccentricity and distance, turn left and turn right trials were collapsed. The collapsed data appear in Figure 6-14. A two-way ANOVA, with eccentricity and distance zone as factors, reveals there was no main effect of eccentricity ($F(4,190) = 0.65, p > .05$), a main effect of distance ($F(1,190) = 11.19, p < .001$), and a significant interaction ($F(4,190) = 2.66, p < .05$). Participants responded more to manipulations in the near zone ($M = 0.11$ m, $SD = 0.16$ m) than in the far zone ($M = 0.03$ m, $SD = 0.18$ m). Response was greatest to manipulations at 0° eccentricity ($M = 0.08$ m, $SD = 0.17$ m) than other eccentricities, but these differences were not significant. A simple effects analysis showed that participants responded significantly more to manipulations in the near zone than in the far zone for the -30° and $+30^\circ$ eccentricity conditions ($p < .01$); there were no significant differences for the other eccentricity conditions, which accounts for the significant interaction. Thus, the position data from Experiment 3C confirm that response is greater for neighbors in the near zone than in the far zone, and they provide weak evidence for an eccentricity effect—namely, that participants respond less to manipulations at higher eccentricities, but only when those manipulations occur at relatively large distances (i.e. in the far zone).

To further simplify the data, I combined the -30° and $+30^\circ$ conditions, and the -15° and $+15^\circ$ conditions. This yielded Figure 6-15, which can be more easily compared to Figure 4-6. Again, there was no main effect of eccentricity ($F(2,194) = 0.61, p > .05$), a main effect of distance zone ($F(1,194) = 11.27, p < .001$), and a significant interaction ($F(2,194) = 4.67, p < .05$). Participants responded more to manipulations in the near zone ($M = 0.11$ m, $SD = 0.16$ m) than in the far zone ($M = 0.03$ m, $SD = 0.18$ m). A simple

effects revealed that participants responded significantly more for near manipulations in the 30° eccentricity condition ($M = 0.14$ m, $SD = 0.18$ m) than far manipulations ($M = -0.02$ m, $SD = 0.12$ m) ($p < .001$); the differences for other eccentricity conditions were not significant.

The mean changes in final speed are shown in Figure 6-16. A two-way ANOVA on change in speed shows a significant main effect of eccentricity ($F(9,180) = 5.35$, $p < .001$), a significant main effect of zone ($F(1,180) = 22.42$, $p < .001$), but the eccentricity-by-zone interaction was not significant ($F(9,180) = 1.45$, $p > .05$). Again, the slow and fast conditions were collapsed, and the collapsed data are shown in Figure 6-17. A two-way ANOVA showed there was no main effect of eccentricity ($F(4,190) = 0.72$, $p > .05$), a main effect of distance ($F(1,190) = 7.84$, $p < .01$), and no significant interaction ($F(4,190) = 0.08$, $p > .05$). Participants responded more to manipulations in the near zone ($M = 0.04$ m/s, $SD = 0.07$ m/s) than in the far zone ($M = 0.01$ m/s, $SD = 0.05$ m/s). Again, these data provide evidence that neighbors in the near zone have a greater influence than neighbors in the far zone. However, as with the results for lateral position, the results for speed do not provide evidence for an eccentricity effect. Thus, overall, the data from Experiment 3C fail to support the hypothesis that spatial pooling depends on the eccentricity of neighbors.

Again, to further simplify the data, I combined the -30° and +30° conditions, and the -15° and +15° conditions. This yielded Figure 6-18, which can be more easily compared to Figure 4-9. As before, there was no main effect of eccentricity ($F(2,194) = 0.30$, $p > .05$), a main effect of distance zone ($F(1,194) = 7.90$, $p < .01$), and no interaction ($F(2,194) = 0.21$, $p > .05$). Participants responded more to manipulations in the near zone ($M = 0.04$ m/s, $SD = 0.07$ m/s) than in the far zone ($M = 0.01$ m/s, $SD = 0.05$ m/s).

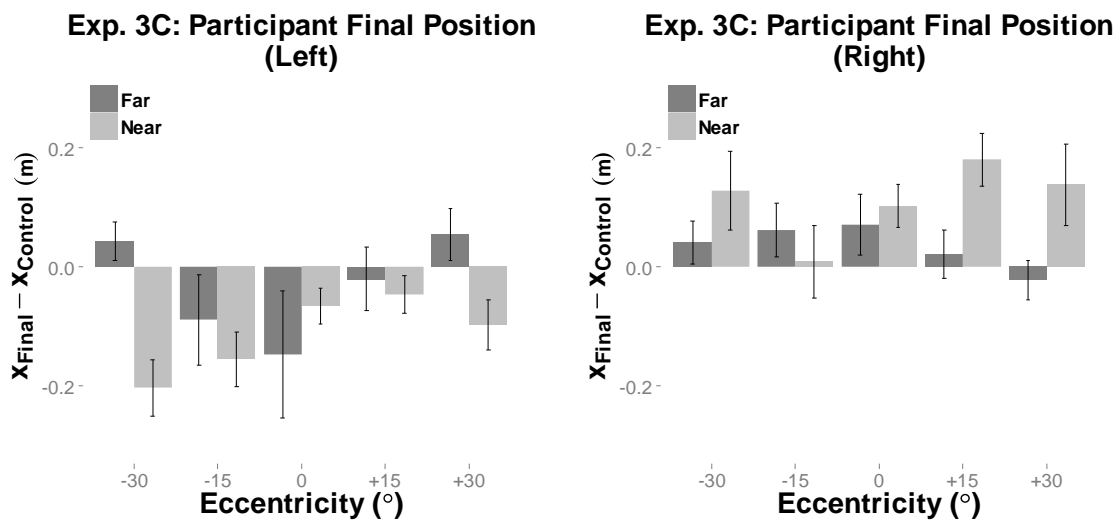


Figure 6-13: Mean changes in participants' final lateral position in Experiment 3C, for "turn left" trials (left) and "turn right" trials (right).

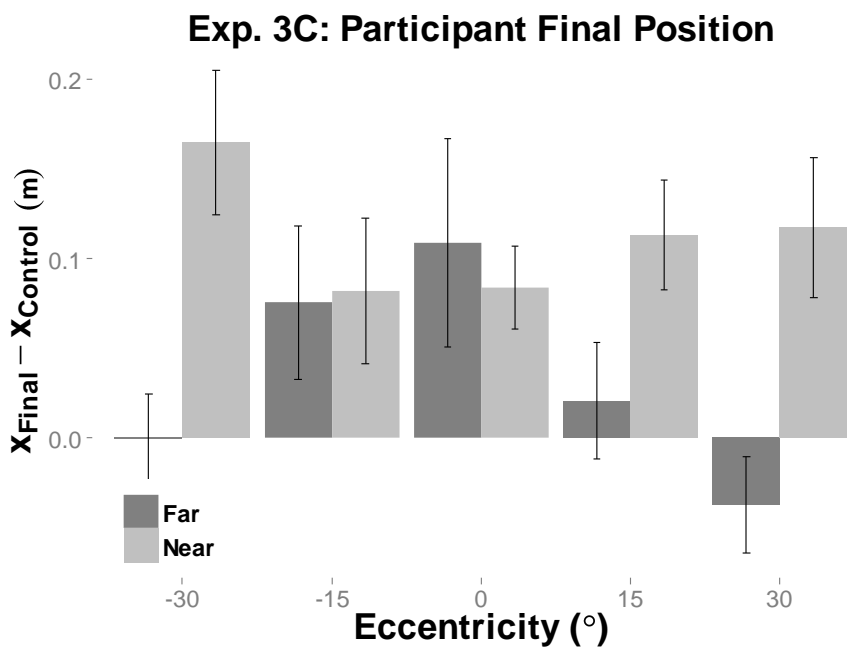


Figure 6-14: Mean changes in participants' collapsed final position in Experiment 3C.




Figure 6-15: Mean changes in participants' collapsed and combined final position in Experiment 3C.




Figure 6-16: Mean changes in participants' final speed in Experiment 3C, for "slower" trials (left) and "faster" trials (right).




Figure 6-17: Mean changes in participants' collapsed final speed in Experiment 3C.




Figure 6-18: Mean changes in participants' collapsed and combined final speed in Experiment 3C.

Discussion

The results from Experiment 3A and 3B suggest that the influences from multiple neighbors are combined, consistent with the results from Experiment 1 (compare Figures 4-6 and 4-9 with Figures 6-10 and 6-12 for position and speed changes, respectively). Moreover, these influences appear to be linearly combined up to as many as 12 neighbors (Figures 6-10 and 6-12). Thus, the linear combination of neighbor influences appears to generalize from the small crowds ($N=3$) of Experiment 1 to the larger crowds ($N=30$) of Experiment 3.

Regarding spatial pooling, the local coupling strength appeared to decay with distance (Experiment 3A) but not eccentricity (Experiment 3C). Participants' responses were greater to neighbors in the near zone than in the far zone (Figures 6-6 and 6-8), and for crowds at high density than crowds at low density (Figures 6-10 and 6-12). While the former results would be predicted by both a metric and topological neighborhood structure, the latter would only be predicted by a metric structure. The topological account predicts that response should stay the same across changes in overall crowd density, because the topological distance to each neighbor (defined in terms of nearest neighbor number, i.e. the first or seventh nearest neighbor) does not change. Thus, Experiment 3B provides evidence against a purely topological neighborhood structure.

There was no consistent influence of neighbor eccentricity in Experiment 3C (see Figures 6-14 and 6-16). This differs from the results in Experiment 1, where greater response was observed for confederates with low eccentricity (Figures 4-6 and 4-9). There are several differences that may account for this discrepancy. First, Experiment 1 was conducted with human confederates; although they received training and practice to

execute their instructions reliably, it was impossible to generate the consistency that is possible with virtual pedestrians. Thus, the human confederates with high eccentricities might have behaved differently from those with low eccentricities (for example, they may have made smaller behavioral changes when instructed), thereby biasing participants' responses in that experiment. Alternatively, participants in Experiment 1 may have treated the confederate in front of them in a qualitatively different manner than the confederates on the side; in Experiment 3, which had larger crowds and 'jittered' initial positions, this difference may have been less pronounced. Finally, note that the eccentricity manipulations in Experiment 1 affected a larger fraction of the crowd (1 of 3 confederates, i.e. 33% of the crowd) than in Experiment 3 (1 or 2 of 12 virtual pedestrians, i.e. 8% or 17% of the crowd). It is possible that the manipulations in Experiment 3 were too small to produce changes in participants' behavior; however, there was a limit to the number of pedestrians that could be manipulated, given the participants' field of view and the resolution of the 'wedges' needed to create 3 levels of eccentricity. Further experiments with larger manipulations could be useful to further explore these differences.

Taken together, the results of Experiment 3 are consistent with a local coupling that is a linear combination of the influences from multiple neighbors and that decays with metric distance, but not eccentricity.

CHAPTER 7

Experiment 4: Walking in a Real Crowd

The purpose of Experiment 4 was to study coordination between neighbors in a real crowd in a more naturalistic setting than the laboratory conditions of Experiments 1-3. In order to measure the coupling between neighbors in a large group, participants were given general instructions, so Experiment 4 is somewhere between observation and experiment. Specifically, participants were instructed to walk together as a group and to veer left and right at will, with no designated leader and no signal about when or where to turn. The resulting “swarm” behavior made it possible to adapt some of the measures that other researchers have developed to quantify coordination in animal collectives. More generally, Experiment 4 makes it possible to question whether the findings from the controlled conditions of Experiments 1-4 will generalize to more naturalistic settings.

Methods

Participants

A total of 36 participants were recruited to create two crowds (N=16 and N=20) for Experiment 3, with the same criteria as before. Each crowd had a roughly equal distribution of genders (crowd 1: 6 females, 10 males; crowd 2: 10 females, 10 males). In addition to being paid, participants received pizza, soft drinks, and a T-shirt.

Apparatus

The experiment was conducted in a large hall (Sayles Hall at Brown University, 14.5 x 22.3 m). Sixteen infrared motion-capture cameras (Qualisys Oqus, Deerfield IL) were set up at regular intervals around the perimeter of the hall on tripods to record head positions (60 Hz sampling rate). The 12 x 20 m tracking area was marked on the floor with red tape; several starting boxes were also marked at each end of the tracking area in

different colored tape. Each participant wore a lightweight bicycle helmet with five reflective markers on protruding stalks (30 to 40 cm long) in a unique constellation. Figure 7-1 shows two sample helmets.



Figure 7-1: Two helmets worn by participants in Experiment 4. Each helmet was outfitted with five reflective markers on protruding stalks to create a unique 'constellation' that could be detected and tracked.

Design

Each group completed four trials (2 low density, 2 high density), each lasting 2 min. This yielded 8 min of crowd data at each density, for a total of approximately 1,152,000 data points.

Procedure

The two groups were tested in separate sessions. Participants were instructed to “walk about the room on a random path at a normal speed, while staying together as a group, veering left or right, for a couple of minutes.” They were also told to stay within the tracking area. Participants began each trial by randomly arranging themselves within one of two start boxes (4x4 m or 7x7 m), which defined two initial densities (high and low,

respectively). A trial was initiated by the experimenter saying, “OK, ready, begin,” and terminated by saying, “OK, stop.” Figure 7-2 shows a sample crowd, and Figure 7-3 shows sample motion capture data.

Data analysis

Head position in three dimensions was recovered on each frame from the motion capture data. The five motion capture markers on each participant’s helmet created a unique “constellation,” whose 3D geometry uniquely identified that helmet and differed from the geometry of every other helmet. On each frame, these constellations were used to identify the visible markers corresponding to each participants’ helmet. Due to occlusion and other tracking errors, all five markers were not always visible on every frame. Three visible markers were used to define a coordinate system (the ‘marker-centered’ coordinate system), find the participant’s head position in that coordinate system, and then perform a coordinate transformation to recover the head position in global coordinates (the ‘world’ coordinate system). When more than three markers were visible, an independent estimate of head position was computed for each set of three markers, and the mean of these estimates was used; when four markers were present, there were three such estimates, and when all five markers were present, there were ten such estimates. Thus, using five markers on each helmet introduced redundancy and improved accuracy. Across all trials and participants, head position was recovered on 88% of frames. Once head position was recovered in this way, the data were processed as before.



Figure 7-2: A sample crowd from Experiment 4. Participants' faces have been obscured to preserve anonymity.




Figure 7-3: Sample motion capture data from Experiment 4. Circles represent the current xy-position of each helmet marker; trails represent the previous 10 s of data. The 12x20 m tracking area is shown, with each grid representing 1x1 m².

Results

Global coordination

Group polarization and group angular momentum are two measures that purportedly capture the overall linear and circular alignment of a crowd's heading direction, respectively (Couzin et al., 2002). Both measures range from 0 to 1, with 0 representing either no coordination or coordination in opposing directions, and 1 representing perfect coordination. They are defined for every time step, producing a time series for each trial. Group polarization, P , is defined at every time step t by the following equation:

$$P(t) = \frac{1}{N} \left| \sum_{i=1}^N v_i(t) \right|$$

where N is the number of pedestrians in the crowd and $v_i(t)$ is the unit velocity vector for each individual at time t (with direction defined by heading and magnitude defined by speed). The quantity inside the vertical bars represents the vector sum of these unit vectors; taking the magnitude of this vector sum and dividing by N yields a quantity that ranges from 0 to 1. If everyone in the crowd is perfectly aligned, each unit velocity vector will point in the same direction, so the result will be a vector pointing in the same direction with magnitude N ; hence, multiplying N by $1/N$ yields a value of 1, the maximum possible group polarization. A minimal value of 0 can be produced in several ways. If the velocities are randomly distributed, they will cancel each other out (because, on average, there will always be another velocity vector pointing in the opposite direction). Alternatively, if there are two streams of pedestrians, each of them perfectly polarized with one another, the overall group polarization will again be 0, because the velocity vectors from these two streams point in opposite directions, and will again cancel out.

Group angular momentum, M , measures the degree to which pedestrians rotate about the center of mass of the group. It is defined by:

$$M(t) = \frac{1}{N} \left| \sum_{i=1}^N r_{ic}(t) \times v_i(t) \right|$$

where N and $v_i(t)$ are defined as before, and $r_{ic}(t)$ is the unit vector pointing from each pedestrian to the crowd's center of mass. Taking the cross product of $r_{ic}(t)$ and $v_i(t)$ yields a vector that is perpendicular to both. If all pedestrians are moving along a circle equidistant from the crowd's center, their velocity vectors $v_i(t)$ will be tangent to that circle, and perpendicular to the vectors pointing to the center of the crowd, $r_{ic}(t)$. Thus, the cross product for each pedestrian will be perfectly aligned, and taking the vector sum's magnitude and dividing by N will yield the maximal value of 1. Again, the minimal value of 0 can be produced in several ways, including random motion, rotation about the center in opposite directions, or linear flow.

Figure 7-4 shows a schematic diagram illustrating these two measures, while Figure 7-5 shows the time series of each measure for a sample trial in Experiment 4.




Figure 7-4: Schematic diagrams of group polarization (blue) and angular momentum (red). For group polarization, the vector sum (large grey arrow) of the individual velocity vectors (blue arrows) is shown below.

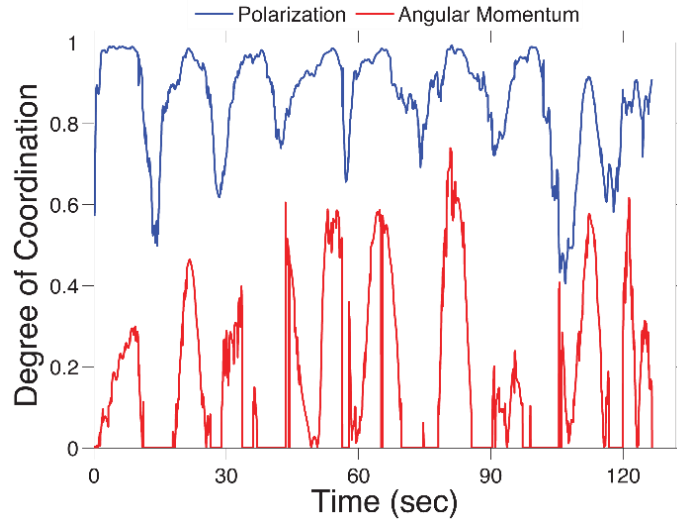


Figure 7-5: Polarization (blue) and angular momentum (red) for a representative trial [$N=16$, low initial density] in Experiment 4.

As illustrated in Figure 7-5 for a representative trial [$N=16$, low initial density], polarization was generally higher ($M = 0.86$, $SD = 0.13$) than angular momentum ($M = 0.17$, $SD = 0.20$). There was no reliable trade-off between the two measures; although angular momentum sometimes increased when polarization decreased (and vice versa), this was not always the case. This may reflect the observation that there were gradual transitions between linear and curved alignments from the front to rear of the swarm, such that both were often present simultaneously. This pattern can be observed in Figure 7-3, where there is a period of relatively linear motion in the middle of the room preceded and followed by a period of rotational motion as the crowd turns to stay within the tracking area. Note that the angular momentum measure assumes a circular configuration of pedestrians, in order for the center of mass to be properly defined; for example, if pedestrians form a half-circle, the center of mass would be defined closer to the perimeter of that half-circle than it would be for an identical full circle. A measure of angular momentum that does not assume a fully

circular configuration would be more fitting for these data, given the relatively small number of individuals being observed. Nevertheless, these measures do provide insight into the global coordination between members of the crowd, and further analyses that could be performed based on them will be discussed at the conclusion of this chapter.

Local coordination

To investigate the local coordination between pedestrians in the swarm, pairwise statistics were computed between individual pedestrians on each frame and aggregated across frames and trials. Many of these statistics were computed between the ‘focal’ pedestrian (the participant nearest to the center of the swarm on each frame) and their neighbors. This approach minimizes edge effects that can occur with pedestrians on the edge of the swarm (Cavagna, 2008).

An occupancy grid for 6 min of swarm data from one $N=16$ and two $N=20$ trials, all with low initial density, appears in Figure 7-6. Data is plotted in the focal pedestrians’ coordinate frame; the origin (0,0) is the focal pedestrian’s location, the positive y-axis (up) is aligned with their heading direction, and color temperature corresponds to the percentage of frames that neighbors are found in each 0.2 m^2 cell. There appears to be an exclusion zone with radius 0.5 m, in which neighbors almost never appear, surrounded by a ‘preferred’ zone with radius 1-2 m where neighbors are most often found. It is possible that these zones indicate a preferred interpersonal distance in human crowds, consistent with distance-based attraction/repulsion models (Schellinck & White, 2011). However, they may also reflect a tendency to maintain the initial density of the swarm in the present experiment. Further manipulation of initial density, and further analysis of the high initial density trials, will test the invariance of these zones.




Figure 7-6: Occupancy grid (percentage of frames in which at least one neighbor is found) over 6 min of data in Experiment 4.




Figure 7-7: Heat map of mean absolute heading difference over 6 min of data in Experiment 4.

The heat map in Figure 7-7 represents the mean absolute difference in heading between the focal pedestrian and neighbors in positions for the same 6 min of data, plotted in the focal pedestrian's coordinate frame. It vividly depicts the degree to which their headings are aligned, providing an estimate of the coupling strength in the focal participant's neighborhood. Notably, the heading difference is small ($<20^\circ$) within a radius of about 2 m, and nearly doubles (35-40°) by 4 m, indicating a strong local coupling that decays rapidly with distance. This is quite consistent with the distance weighting observed in Experiments 3A and 3C, and the circularly symmetric pattern offers little suggestion of an eccentricity effect, consistent with the results of Experiment 3C.

Discussion

The purpose of Experiment 4 was to generalize from the highly-controlled laboratory conditions of Experiments 1-3 to a more naturalistic setting. Overall, the results appear to be broadly consistent across these two contexts. Participants in Experiment 4 exhibit behavior that is consistent with combining influences from multiple neighbors, weighted by their distance and relatively unaffected by their eccentricity (Figure 7-7). Moreover, participants show signs of broad global coordination, especially during translation, as quantified by group polarization (Figure 7-5). Thus, Experiment 4 provides little evidence that the fundamental aspects of collective locomotion observed in Experiments 1-3 fail to generalize to larger crowds or more naturalistic settings.

Naturally, the rich data collected in Experiment 4 offer a number of avenues for future work. Additional pairwise measures, such as the time delay between participants, can be computed in order to capture additional details about the local visual coupling. This could ultimately lead to a representation of the network structure of the crowd—which

neighbors are coupled to which others, and how does influence propagate through these links? Ultimately, one goal is to simulate the behavior of an individual within one of these crowds, using their neighbors as input. Indeed, Warren & Bonneaud (2014) presented preliminary evidence that individual trajectories in this experiment could be predicted if pedestrians aligned their heading with one another, but not if they treated one another as moving obstacles—the latter approach led to a divergent solution, with the crowd scattering radially outward in all directions. Further testing of different combinations of model components, informed by behavioral experiments such as those presented earlier, should help to facilitate this modeling effort.

CHAPTER 8

General Discussion

Collective locomotion is inherently complex. It is a phenomenon that subsumes behavior across spatiotemporal scales, from the perception, cognition, and locomotion of individual animals to the global patterns of behavior that emerge from myriad interactions between them. This dissertation represents the first step in a long road toward an empirically-grounded, cognitively-plausible understanding of collective behavior of human crowds. The results presented thus far shed new light onto the fundamental properties of the local coupling between neighbors, and provide some of the first experimental evidence to quantify these core building blocks of crowd behavior.

There are three major themes that were consistently observed throughout Experiments 1-4, which can be expressed as properties of the local coupling. First, the influence from multiple neighbors *is additive*. Manipulating more neighbors elicited greater changes in participants' behavior (speed or heading) in virtually every experiment; this can be clearly observed in Figures 4-6, 4-9, 6-6, 6-8, 6-10, and 6-12. Although this finding may not be surprising, it is important; additivity or superposition of neighbor influences is assumed in virtually every model of collective locomotion, but it has not been so clearly demonstrated in a well-controlled experimental context. Second, neighbor influence *decays strongly with distance*. This can be observed in Experiments 3A and 3C, where manipulating neighbors in the far zone (3.5 m) had a weaker effect than manipulating neighbors in the near zone (1.5 m), as well as in Experiment 3B, where participant response was weaker at low density (mean distance: 4.5 m) than at high density (mean distance: 2.5 m). Experiment 4 confirmed this result; mean absolute heading difference decreased significantly with distance, as shown in the heat map in Figure 7-7. Experiment 1 seems to contradict this finding, because there was no effect of initial

distance; however, as discussed earlier, this was likely due to the relatively small range of distances that were tested (1 and 2 m), which may not have produced detectable changes in behavior. The third and final theme is that neighbor influence decays *weakly with eccentricity*. In Experiment 1, participants responded significantly more to manipulations of a single confederate located in the center (low eccentricity) than on the side (high eccentricity). This finding was not replicated in Experiment 3C, but again, this was lack of an effect is likely attributable to experimental design. Only one or two virtual pedestrians were manipulated on each trial, which did not elicit detectable changes in participants' behavior. Nevertheless, simple effects analyses did reveal some slight differences as a function of eccentricity—namely, that participants responded less to manipulations in the far zone when those manipulations occurred at high eccentricity. This, along with the results of Experiment 1, do provide some evidence that neighbor influence decays, albeit weakly, with eccentricity.

More generally, the results of Experiments 1-3 provide quantitative, empirical evidence about properties of the local coupling between neighbors. In Experiment 4, these results were confirmed, both at the local level—where there is evidence that neighbor influence decays with distance (Figure 7-7), as predicted from earlier results—and the global level. Through the propagation of local interactions, crowds exhibit global patterns of coordination, which can be quantified using measures such as group polarization (Figure 7-5). My dissertation also clearly demonstrates that virtual reality is a powerful tool in the study of collective locomotion. Virtual reality makes it possible for experiments to precisely manipulate any property of a crowd's composition or trajectory. Humans are not the only animals for which virtual reality systems have been developed

(Abramson, Buckbee, Edwards, Bowe, 1996; Hölscher, Schnee, Dahmen, Setia, & Mallot, 2005; Youngstrom & Strowbridge, 2012), but the ambulatory virtual environment that is produced in the VENLab has not been replicated in other species.

Taken together, the results of my dissertation provide some of the first experimental evidence that can be brought to bear on fundamental questions regarding the local interactions that give rise to human crowd dynamics. Moreover, my dissertation also represents an important first pass at completing the modeling cycle that was laid out by Sumpter, Mann, & Pernea (2012). By linking the local and global levels of analysis, and relying on data collected from real and virtual crowds, this work represents a significant contribution to the field and a major step forward on the path toward an empirically-grounded, cognitive-plausible understanding of human crowd dynamics.

Limitations

Despite these successes, there are also a number of limitations that must be taken into account when drawing conclusions from the research presented thus far. Designing experiments to tackle an exceedingly complex problem such as crowd behavior requires making deliberate choices about which variables to manipulate, which measurements to make, and which effects to test for. These choices, while they are necessary to ensure tractability, are not easy to make, and require compromises that are not always readily apparent until the data are in.

First, there are a number of limitations related to the virtual reality methodology used in Experiments 2 and 3. Some of these are specific to the current technologies tested, and the state of the art of in virtual reality hardware. The Oculus Rift DK1 used in

Experiments 2 and 3, for example, provides a relatively wide field of view, but it is still smaller than typical human vision; this places artificial limits on the number of virtual pedestrians that can be presented, as well as limits on their maximum eccentricity. Other causes for concern include the relatively low resolution displays (which may affect the ability to detect distant pedestrians) and relatively high latencies (which may affect how participants interact with their virtual neighbors) associated with current display technologies. If virtual reality continues its current rate of progression, many of these technical issues will be resolved within a few years. Other issues will remain, however. Although Experiment 2 showed that participants behave similarly when interacting with real and virtual pedestrians, the design of this experiment was somewhat limited. Virtual pedestrians in Experiments 2 and 3 followed pre-specified trajectories, and were not influenced by participants' behavior. This was done by design, to isolate participants' responses to particular stimuli. In the future, however, it will be desirable to create virtual crowds that are bidirectionally coupled to participants; that is, the virtual crowd should not only influence participants, but also *be influenced* by them. This will require considerable artificial intelligence efforts to ensure that virtual pedestrians behave and react appropriately, and will require further validation studies to ensure that participants continue to treat them similarly.

A second limitation concerned the number and range of variables that could be tested in the present experiments. Even with a crowd of just 12 virtual pedestrians, the state space of possible manipulations is virtually infinite—in practice, it is impossible to test a sizeable fraction of the possible combinations of densities, initial distances, preferred walking speeds, changes in speed and heading direction, subsets of manipulated and non-

manipulated neighbors, and countless other variables that define the virtual crowd's trajectories. The conditions that *were* tested were carefully chosen to optimize the predicted return of investment given a finite number of participants and participant hours. Nevertheless, some of these choices made it difficult to detect effects or test hypotheses. For example, the goal of Experiment 3C was to see whether the eccentricity effect from Experiment 1 would scale to larger crowds, where more fine-grained manipulations of eccentricity were possible. Rather than simply detecting a difference between “high” and “low” eccentricity (more precisely, “high” and “no,” because confederates in Experiment 1 were located either on the side or directly in front), I wanted to see how this effect varied. Did neighbor influence linearly decay with eccentricity, or was it exponential? Did influence level off at a certain point, or did it steadily decay to zero? I chose to test 3 levels in Experiment 3C, rather than 2 in Experiment 1, to address these more nuanced questions. But slicing up an already constrained visual field into 5 non-overlapping regions meant that only one or two virtual pedestrians could be manipulated on any given trial. There are several possible solutions—using non-overlapping slices, increasing the number or density of the crowd to ensure more pedestrians were manipulated, and so on. The general point to be made is that conducting experiments in such a large space is difficult, and that while it may not be possible to anticipate all such difficulties *a priori*, careful designs are essential in order to test hypotheses with a reasonable number of experiments, conditions, and participant hours. Care should be taken to disentangle variables that covary; for example, reducing crowd density has the effect of increasing the distance to neighbors, on average.

A third limitation concerns a subtle distinction between the absolute number of people in a crowd, and the relative proportion or degree of consensus within the crowd.

Experiments 1-3 each manipulated the number of real or virtual pedestrians who changed their speed or direction, and it was observed that participants' responses increase as a function of this number. It is unclear whether this effect is due to the absolute number, or the degree of consensus, because these factors are confounded. In Experiment 3, participants responded more when 9 neighbors changed their speed or direction than when 3 did; but this was equivalent to saying that participants responded more when 75% of the crowd changed their speed or direction than when 25% did, because the size of the crowd was always fixed at 12. Which is truly important? Is it the absolute number of neighbors that matters, or is it their percentage of the crowd? The answer to this question has important theoretical implications—for example, it may provide evidence to support the hypothesis that people perceive summary statistics about the overall speed or direction of the crowd as a whole, rather than perceiving the speed and direction of each neighbor individually. Fortunately, these questions can be addressed with a relatively straightforward follow-up experiment. In addition to the manipulations in Experiment 3 (i.e. manipulating 0, 3, 6, 9, or 12 pedestrians in a crowd of 12, or 0, 25, 50, 75, or 100% of the crowd), an additional factor should be added to the design—the size of the crowd. To keep the proportions simple, this could be done in crowds of 12 (a replication of Experiment 3) and 18. Manipulating 50% of the crowd would now mean manipulating a different absolute number of neighbors—6 for a crowd of 12, and 12 for a crowd of 24. Likewise, manipulating the same absolute number would correspond to different proportions—25% of a crowd of size 12 is 3, while 25% of a crowd of size 24 is 6. Comparing participant responses across these different conditions would reveal whether absolute number or degree of consensus has the largest effect on behavior.

Future Directions

The experiments presented here were inspired, informed, and are ultimately indebted to over a decade's worth of research on the visual control of locomotion in the VENLab. They represent an important step forward—by scaling up from individuals to large crowds, and by validating the use of virtual reality as an experimental methodology to study collective locomotion, they pave the way for future explorations of the fundamental processes underlying human crowd dynamics. It will be informative to consider a few directions that could be pursued to that end.

In earlier work (Rio, Rhea, & Warren, 2014), I studied the behavioral dynamics of pedestrian following, and also worked out the visual information that was used to regulate behavior, in order to derive a visual control law for this simple behavior. This strategy becomes increasingly difficult as locomotor behaviors grow in complexity, but there are ways in which the study of visual perception can inform studies of pedestrian and crowd behavior. One of the most obvious of these relates to occlusion. People in a crowd tend to occlude one another, and this effect varies systematically with certain spatial properties, such as the distance to a given pedestrian and the density of the crowd as a whole. How does this affect the interactions between neighbors? Occlusion offers a plausible explanation for why neighbor influence may decay with distance, at least in some cases. It can be experimentally manipulated in several interesting ways. Changing the relative difference in height between the participant and the crowd is one example. When participants are taller relative to their neighbors, they can see farther (and hence potentially interact with more of them). This could be accomplished by changing the eye-

height in virtual reality or, perhaps more preferably, by reducing the height of the virtual pedestrians. Another method—whose ecological validity is open to question—would be to render the virtual pedestrians as translucent, enabling participants to see potentially all members of the crowd simultaneously.

This latter manipulation opens up another series of questions, regarding the effects of attention. Visual attention is often cited as one reason for the limited neighborhood with which animals interact (Ballerini et al., 2008), and studies in the VENLab have shown that attention plays a key role in the visual control of locomotion (Cohen, 2009; Harrison & Warren, 2011). It would be interesting to combine the virtual reality methodology from Experiment 3 with eye-tracking, to determine which neighbors participants spend the most time looking at, and what strategies they use to pick up information from multiple neighbors simultaneously (or nearly simultaneously). Other experiments along these lines could include systematically varying the rate at which participants can pick up information from their neighbors; rendering virtual pedestrians on intermittent frames or using shutter glasses in the real world provide two ways of doing this. Again, evidence from these kinds of studies would be useful in determining whether people rely on perceiving the speed and direction of individual neighbors, or whether they instead perceive summary statistics over the crowd as a whole.

Further work is also needed to connect the data presented here with modeling and simulation. Although Experiment 1 showed that a binary model could account for behavior in a small crowd, this approach must be further scaled up to larger crowds, as in Warren & Bonneaud (2014). This is a difficult task. Naturalistic experiments, which constrain behavior in certain ways but are otherwise uncontrolled, such as Experiment 4

in this dissertation, are a promising strategy for such efforts. During the “Sayles Swarm,” my colleagues and I conducted over a dozen experiments in addition to those presented here as Experiment 4. Some of these have already been analyzed and presented in various forms (Warren, Bonneaud, & Kiefer, 2013; Warren & Bonneaud, 2014; Dachner & Warren, in press), but some have not. These data make it possible to focus modeling efforts on specific behaviors—such as obstacle avoidance in the Grand Central scenario (Bonneaud & Warren, 2014)—in the hopes of reproducing specific phenomena such as lane formation or hysteresis. These ‘constrained, naturalistic’ experiments provide a useful compliment to the unconstrained observations that can be recorded from crowds ‘in the wild,’ where the aims, goals, and objectives of individual people are difficult to decipher.

Ultimately, one of the most difficult challenges facing the field of collective locomotion is how to combine, synthesize, and reconcile the disparate findings from theory and experiment, across a wide range of contexts and scenarios (Sumpter, Mann, & Pernea, 2012). We are unlikely to ever find a ‘grand unified theory’ of collective locomotion; by its very nature, it is a complex phenomenon. But it is important to continue searching for one, because theoretical approaches provide the scaffolding upon which empirical data can come to rest. As the field continues to produce mountains of data on crowds, flocks, schools, and swarms, theory is needed to tie them all together, to produce a coherent framework, to build a deeper and more fundamental understanding rather than simply a collection of facts. I hope that my dissertation has contributed toward this goal, and that further work continues to pull back the layers of mystery surrounding human crowd dynamics.

References

- Abramson, C. I., Buckbee, D. A., Edwards, S., & Bowe, K. (1996). A demonstration of virtual reality in free-flying honeybees: *Apis mellifera*. *Physiology & Behavior*, *59*(1), 39-43.
- Aiello, J.R., & Thompson, D.E. (1980). Personal space, crowding, and spatial behavior in a cultural context. In Altman, I., Rapoport, A., & Wohwill, J. (Eds.), *Human Behavior and Environment: Advances in Theory and Research* (Vol. 4). New York: Plenum Press.
- Altman, I., & Vinsel, A.M. (1977). Personal space: An analysis of E.T. Hall's Proxemics framework. In Altman, I., & Wohwill, J. (Eds.), *Human Behavior and Environment: Advances in Theory and Research* (Vol. 2). New York: Plenum, 1977.
- Ballerini, M., Cabibbo, N., Candelier, R., Cavagna, A., Cisbani, E., Giardina, I., Lecomte, V., Orlandi, A., Parisi, G., Procaccini, A., Viale, M., & Zdravkovic, V. (2008). 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.
- Barhomi, Y., Yanke, A., Bonneaud, S., Warren, W., & Serre, T. (2014). A data-driven approach to learning strategies for the visual control of navigation. Poster, Vision Sciences Society 14th Annual Meeting, St. Pete Beach, FL.
- Boles, W. (1981). The effect of density, sex, and group size upon pedestrian walking activity. *Man-Environment Systems*, *11*, 37-40.
- Bonneaud, S., Rio, K., Chevaillier, P., & Warren, W. H. (2012). Accounting for patterns of collective behavior in crowd locomotor dynamics for realistic simulations. In *Transactions on Edutainment VII* (pp. 1-11). Springer Berlin Heidelberg.

- Bonneaud, S., Warren, W. H., Olfers, K., Irwin, G., & Serre, T. (2013). Towards a biologically-inspired vision system for the control of locomotion in complex environments. *Journal of Vision*, *13*(9), 753-753.
- Bonneaud, S., & Warren, W. H. (2014). An empirically-grounded emergent approach to modeling pedestrian behavior. In *Pedestrian and Evacuation Dynamics 2012* (pp. 625-638). Springer International Publishing.
- Bornstein, M. H., & Bornstein, H. G. (1976). The pace of life. *Nature*, *259*, 557-559.
- Bruggeman, H., Zosh, W., & Warren, W. H. (2007). Optic flow drives human visuo-locomotor adaptation. *Current Biology*, *17*(23), 2035-2040.
- Bruggeman, H., & Warren, W. H. (2010). The direction of walking—but not throwing or kicking—is adapted by optic flow. *Psychological Science*, *21*(7), 1006-1013.
- Camazine, S., Deneubourg, J.-L., Franks, N. R., Sneyd, J., Theraulaz, G., & Bonabeau, E. (2001). *Self-Organization in Biological Systems*. Princeton University, Princeton, NJ.
- Cavagna, A., Cimorelli, A., Giardina, I., Orlandi, A., Parisi, G., Procaccini, A., Santagati, R. et al. (2008). New statistical tools for analyzing the structure of animal groups. *Mathematical Biosciences*, *214*(1-2), 32-37.
- Cavagna, A., Cimorelli, A., Giardina, I., Parisi, G., Santagati, R., Stefanini, F., & Viale, M. (2010). Scale-free correlations in starling flocks. *Proceedings of the National Academy of Sciences*, *107*(26), 11865-11870.
- Cohen, J. A., Bruggeman, H., & Warren, W. H. (2005). Switching behavior in moving obstacle avoidance. *Journal of Vision*, *5*(8), 312-312.

- Cohen, J. A., Bruggeman, H., & Warren, W. H. (2006). Combining moving targets and moving obstacles in a locomotion model. *Journal of Vision*, 6(6), 135-135.
- Cohen, J. A. (2009). Perception of pursuit and evasion by pedestrians. Unpublished PhD thesis.
- Costa, M. (2010). Interpersonal distances in group walking. *Journal of Nonverbal Behavior*, 34, 15-26.
- Couzin, I. D., Krause, J., James, R., Ruxton, G. D., & Franks, N. R. (2002). Collective memory and spatial sorting in animal groups. *Journal of Theoretical Biology*, 218, 1-11.
- Couzin, I. D., & Krause, J. (2003). Self-organization and collective behavior in vertebrates. *Advances in the Study of Behavior*, 32, 1-75.
- Couzin, I. D. (2009). Collective cognition in animal groups. *Trends in Cognitive Sciences*, 13(1), 36-43.
- Czirok, A., Stanley, H. E., & Vicsek, T. (1997). Spontaneously ordered motion of self-propelled particles. *Journal of Physics A: Mathematical and Theoretical*, 30, 1375-1385.
- Czirok, A., & Vicsek, T. (2000). Collective behavior of interacting self-propelled particles. *Physica A*, 281, 17-29.
- Dachner, G. & Warren, W.H. (in press). Behavioral dynamics of heading alignment in pedestrian following. *Transportation Research Procedia*.
- Dyer, J., Ioannou, C., Morrell, L., Croft, D., Couzin, I., Waters, D., & Krause, J. (2008). Consensus decision making in human crowds. *Animal Behaviour*, 75(2), 461-470.

- Dyer, J. R. G., Johansson, A., Helbing, D., Couzin, I. D., & Krause, J. (2009). Leadership, consensus decision making and collective behaviour in humans. *Philosophical Transactions of the Royal Society B*, 364, 781-789.
- Goldstone, R. L., & Janssen, M. A. (2005). Computational models of collective behavior. *Trends in Cognitive Sciences*, 9(9), 424-430.
- Ebeling, W., & Schimansky-Geier, L. (2008). Swarm dynamics, attractors and bifurcations of active Brownian motion. *The European Physical Journal Special Topics*, 157(1), 17-31.
- Estellés-Arolas, E., & González-Ladrón-de-Guevara, F. (2012). Towards an integrated crowdsourcing definition. *Journal of Information Science*, 38(2), 189-200.
- Fajen, B. R., & Warren, W. H. (2007). Behavioral dynamics of intercepting a moving target. *Experimental Brain Research*, 180(2), 303-19.
- Fajen, B. R., & Warren, W. H. (2003). Behavioral dynamics of steering, obstacle avoidance, and route selection. *Journal of Experimental Psychology: Human Perception and Performance*, 29(2), 343-362.
- Fajen, B.R., & Warren, W.H. (2010, May). Route selection in complex environments emerges from the dynamics of steering and obstacle avoidance. Talk presented at the Annual Meeting of the Vision Sciences Society, Naples, FL. Abstract published in *Journal of Vision* 10(7), 1046.
- Faria, J. J., Dyer, J. R. G., Tosh, C. R., & Krause, J. (2010). Leadership and social information use in human crowds. *Animal Behaviour*, 79, 895-901.

- Funge, J., & Terzopoulos, D. (1999). Cognitive modeling: Knowledge, reasoning, and planning for intelligent characters. *Proceedings of the 1999 ACM SIGGRAPH International Conference on Computer Graphics and Interactive Techniques*, 29-38.
- Galef Jr, B. G., & Giraldeau, L. A. (2001). Social influences on foraging in vertebrates: causal mechanisms and adaptive functions. *Animal Behaviour*, 61(1), 3-15.
- Galton, F. (1907). Vox populi (The Wisdom of Crowds). *Nature* 1949(75), 450-451.
- Gérin-Lajoie, M., & Warren, W. (2008). The circumvention of barriers: Extending the steering dynamics model. *Journal of Vision*, 8(6), 1158-1158.
- Gibbs, J. (1960). *Elementary Principles in Statistical Mechanics*, New York: Dover.
- Goatin, P., Colombo, R. M., & Rosini, M. D. (2010). A macroscopic model for pedestrian flows in panic situations. *GAKUTO International Series Mathematical Sciences and Applications*, 32, 255-272.
- Gibson, J. J. (1950). *The Perception of the Visual World*.
- Gibson, J. J. (2013). *The Ecological Approach to Visual Perception*. Psychology Press.
- Goldstone, R. L., & Gureckis, T. M. (2009). Collective behavior. *Topics in Cognitive Science*, 1(3), 412-438.
- Gordon, D. M. (2014). The ecology of collective behavior. *PLoS Biology*, 12(3), e1001805.
- Haken, H., Kelso, J.A.S., & Bunz, H. (1985). A theoretical model of phase transitions in human hand movements. *Biological Cybernetics*, 51, 347-356.
- Haken, H., & Wunderlin, A. (1990). Synergetics and its paradigm of self-organization in biological systems. *The Natural-Physical Approach to Movement Control*. Amsterdam: VU University Press. ISBN, 472634347.

- Haken, H. (2006). *Information and Self-organization: A Macroscopic Approach to Complex Systems*. Springer.
- Hall, E.T. (1969). *The Hidden Dimension*. Garden City: Doubleday Anchor.
- Harrison, H. S., & Warren, W. H. (2011). Inconsistent routes in moving obstacle avoidance are due to sensitivity to initial conditions, not attention. *Journal of Vision, 11*(11), 912-912.
- Hayduk, L.A. (1978). Personal space: An evaluative and orienting overview. *Psychological Bulletin, 85*, 117-134.
- Hayduk, L. A. (1981). The permeability of personal space. *Canadian Journal of Behavioural Science/Revue canadienne des sciences du comportement, 13*(3), 274.
- Hayduk, L. A. (1983). Personal space: Where we now stand. *Psychological Bulletin, 94*(2), 293-335.
- Helbing, D. (1992). A fluid-dynamic model for the movement of pedestrians. *Complex Systems, 6*, 391-415.
- Helbing, D., & Molnár, P. (1995). Social force model for pedestrian dynamics. *Physical Review E, 51*(5), 4282-4286.
- Helbing, D., Keltsch, J., & Molnar, P. (1997). Modelling the evolution of human trail systems. *Nature, 388*(6637), 47-50.
- Helbing, D., Farkas, I., & Vicsek, T. (2000). Simulating dynamical features of escape panic. *Nature, 407*(6803), 487-90.
- Henderson, L. F. (1971). The statistics of crowd fluids. *Nature, 229*, 381-383.
- Henderson, L. F., & Lyons, D. J. (1972). Sexual differences in human crowd motion. *Nature, 240*, 353-355.

- Henderson, L. F. (1974). On the fluid mechanics of human crowd motion. *Transportation Research*, 8(6), 509-515.
- Hoel, L.A. (1968). Pedestrian travel rates in central business districts. *Traffic Engineering*, 38, 10-13.
- Hölscher, C., Schnee, A., Dahmen, H., Setia, L., & Mallot, H. A. (2005). Rats are able to navigate in virtual environments. *The Journal of Experimental Biology*, 208(3), 561-569.
- Hildenbrandt, H., Carere, C., & Hemelrijk, C. K. (2010). Self-organized aerial displays of thousands of starlings: A model. *Behavioral Ecology*, 21(6), 1349-1359.
- Hoogendoorn, S., & Bovy, P. H. L. (2000). Gas-kinetic modeling and simulation of pedestrian flows. *Transportation Research Record*, 1710, 28-36.
- Howarth, S. J., & Callaghan, J. P. (2009). The rule of 1s for padding kinematic data prior to digital filtering: Influence of sampling and filter cutoff frequencies. *Journal of Electromyography and Kinesiology*, 19(5), 875-881.
- Hughes, R. L. (2003). The flow of human crowds. *Annual Review of Fluid Mechanics*, 35(1), 169-182.
- Huth, A., & Wissel, C. (1992). Simulation of fish schools. *Journal of Theoretical Biology*, 156, 365-385.
- Kelley, D. H., & Ouellette, N. T. (2013). Emergent dynamics of laboratory insect swarms. *Scientific Reports* 3(1073), 1-7.
- Kelso, J. A. S. (1995). *Dynamic Patterns: The Self-Organization of Brain and Behavior*.

- Kiefer, A. W., Bonneaud, S., Rio, K., & Warren, W. H. (2014). Quantifying the coherence of pedestrian groups. *Proceedings of the Cognitive Science Society 35th Annual Meeting*.
- Kugler, P. N., & Turvey, M. T. (1987). *Information, Natural law, and the Self-Assembly of Rhythmic Movement*.
- Lemercier, S., Jelic, A., Kulpa, R., Hua, J., Fehrenbach, J., Degond, P., ... & Pettré, J. (2012, May). Realistic following behaviors for crowd simulation. In *Computer Graphics Forum* (Vol. 31, No. 2pt2, 489-498). Blackwell Publishing Ltd.
- Lukeman, R., Li, Y.-X., Edelstein-Keshet, L. (2010) Inferring individual rules from collective behavior. *Proceedings of the National Academy of Sciences*, 107(28), 12576-12580.
- Mallenby, T., & Roberts, R.G. (1973). *A bibliography of research on spatial and social behavior*. Winnipeg: Thomas Todd.
- Ma, J., Song, W. G., Zhang, J., Lo, S. M., & Liao, G. X. (2010). k-Nearest-Neighbor interaction induced self-organized pedestrian counter flow. *Physica A: Statistical Mechanics and its Applications*, 389(10), 2101-2117.
- Marr, D. (1982). *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*.
- Moussaïd, M., Helbing, D., Garnier, S., Johansson, A., Combe, M., & Theraulaz, G. (2009). Experimental study of the behavioural mechanisms underlying self-organization in human crowds. *Proceedings of the Royal Society B*, 276(1668), 2755-62.

- Moussaïd, M., Perozo, N., Garnier, S., Helbing, D., & Theraulaz, G. (2010). The walking behaviour of pedestrian social groups and its impact on crowd dynamics. *PLoS ONE*, 5(4), 1-7.
- Moussaïd, M., Helbing, D., & Theraulaz, G. (2011). How simple rules determine pedestrian behavior and crowd disasters. *Proceedings of the National Academy of Sciences*, 108(17), 6884-6888.
- Nagai, R., Fukamachi, M., & Nagatani, T. (2006). Evacuation of crawlers and walkers from corridor through an exit. *Physica A: Statistical Mechanics and its Applications*, 367, 449-460.
- Nagy, M., Akos, Z., Biro, D., & Vicsek, T. (2010). Hierarchical group dynamics in pigeon flocks. *Nature*, 464(7290), 890-3.
- Obudho, C.E. (1974). The proxemic behavior of man and animals: An annotated bibliography. *Council of Planning Librarians Exchange Bibliography Numbers*, 646-647, 1-88.
- Okazaki, S., & Matsushita, S. (1993). A study of simulation model for pedestrian movement with evacuation and queuing. *Proceedings of the 1993 International Conference on Engineering for Crowd Safety*, 271-280.
- Olivier, A. H., Marin, A., Crétual, A., & Pettré, J. (2012). Minimal predicted distance: A common metric for collision avoidance during pairwise interactions between walkers. *Gait & Posture*, 36(3), 399-404.
- Page, Z., & Warren, W. (2012). Visual control of speed in side-by-side walking. *Journal of Vision*, 12(9), 188-188.

- Page, Z., & Warren, W. (2013). Speed-matching strategy used to regulate speed in side-by-side walking. *Journal of Vision*, 13(9), 950-950.
- Paris, S., & Donikian, S. (2009). Activity-driven populace: A cognitive approach to crowd simulation. *IEEE Computer Graphics and Applications*, 29(4), 34-43.
- Patla, A. E. (1997). Understanding the roles of vision in the control of human locomotion. *Gait & Posture*, 5(1), 54-69.
- Pechano, N., Allbeck, J.M., & Badler, N.I. (2007). Controlling individual agents in high-density crowd simulation. *Proceedings of the 2007 ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (pp. 99-108).
- Pelechano, N., Stocker, C., Allbeck, J., & Badler, N. (2008, May). Being a part of the crowd: towards validating VR crowds using presence. In *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems- Volume 1* (pp. 136-142). International Foundation for Autonomous Agents and Multiagent Systems.
- Reynolds, C. W. (1987). Flocks, herds and schools: A distributed behavioral model. *Computer Graphics*, 21(4), 25-34.
- Reynolds, C.W. (1999). Steering behaviors for autonomous characters. *Proceedings of the 1999 Game Developers Conference*, 763-782.
- Rio, K. W., Rhea, C. K., & Warren, W. H. (2014). Follow the leader: Visual control of speed in pedestrian following. *Journal of Vision*, 14(2), 4.
- Rose, J., & Gamble, J.G. (Eds.). (1994). *Human Walking*. Baltimore: William & Wilkins (2nd ed.)

- Schadschneider, A., Klingsch, W., Klüpfel, H., Rogsch, C., & Seyfried, A. (2008). Evacuation dynamics: empirical results, modeling and applications. In B. Meyers (Ed.), *Encyclopedia of Complex and Systems Science* (pp. 1-57). Berlin: Springer.
- Schellinck, J., & White, T. (2011). A review of attraction and repulsion models of aggregation: Methods, findings and a discussion of model validation. *Ecological Modeling*, 222, 1897-1911.
- Shimada, T., & Naoi, H. (2006). An experimental study on the evacuation flow of crowd including wheelchair users. *Fire Science and Technology*, 25(1), 1-14.
- Singh, S., Kapadia, M., Faloutsos, P., & Reinman, G. (2009). SteerBench: A benchmark suite for evaluating steering behaviors. *Computer Animation And Virtual Worlds*, 20, 533-548.
- Srinivasan, M., Zhang, S. W., Lehrer, M., & Collett, T. (1996). Honeybee navigation en route to the goal: Visual flight control and odometry. *Journal of Experimental Biology*, 199(1), 237-244.
- Sumpter, D. J., Mann, R. P., & Perna, A. (2012). The modelling cycle for collective animal behaviour. *Interface Focus*, rsfs20120031.
- Sundstrom, E., & Altman, I. (1976). Interpersonal relationships and personal space: Research review and theoretical model. *Human Ecology*, 4, 47-67.
- Surowiecki, J. (2005). *The Wisdom of Crowds*. Random House.
- Tarr, M. J., & Warren, W. H. (2002). Virtual reality in behavioral neuroscience and beyond. *Nature Neuroscience*, 5, 1089-1092.
- Vicsek, T. (1995). Novel type of phase transition in a system of self-driven particles. *Physical Review Letters*, 75(6), 4-7.

- Vicsek, T. (2008, October). Universal patterns of collective motion from minimal models of flocking. In *Self-Adaptive and Self-Organizing Systems, 2008. SASO'08. Second IEEE International Conference on* (pp. 3-11).
- Vicsek, T., & Zafeiris, A. (2012). Collective motion. *Physics Reports*, *517*(3), 71-140.
- Vint, P. F., & Hinrichs, R. N. (1996). Endpoint error in smoothing and differentiating raw kinematic data: an evaluation of four popular methods. *Journal of Biomechanics*, *29*(12), 1637-1642.
- Warren, W. H., & Hannon, D. J. (1988). Direction of self-motion is perceived from optical flow. *Nature*, *336*(6195), 162-163.
- Warren, W. H., Morris, M. W., & Kalish, M. (1988). Perception of translational heading from optical flow. *Journal of Experimental Psychology: Human Perception and Performance*, *14*(4), 646.
- Warren, W. H., Kay, B. A., Zosh, W. D., Duchon, A. P., & Sahuc, S. (2001). Optic flow is used to control human walking. *Nature Neuroscience*, *4*(2), 213-216.
- Warren, W. H. (2009). How do animals get about by vision? Visually controlled locomotion and orientation after 50 years. *British Journal of Psychology*, *100*(S1), 277-281.
- Warren, W.H., & Bonneaud, S. (2014). Multi-agent simulation of collective behavior in human crowds. Talk, Vision Sciences Society 14th Annual Meeting, St. Pete Beach FL.
- Wirtz, P., & Ries, G. (1992). The pace of life - reanalyzed: Why does walking speed of pedestrians correlate with city size? *Behavior*, *123*(1-2), 77-83.

- Yamori, K. (1998). Going with the flow: Micro-macro dynamics in the macrobehavioral patterns of pedestrian crowds. *Psychological Review*, *105*, 530-557.
- Youngstrom, I. A., & Strowbridge, B. W. (2012). Visual landmarks facilitate rodent spatial navigation in virtual reality environments. *Learning & Memory*, *19*(3), 84-90.
- Yu, Q., & Terzopoulos, D. (2007). A Decision Network Framework for the Behavioral Animation of Virtual Humans. *Proceedings of the 2007 ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (pp. 119-128).