Walking Through a Crowd: Modeling Pedestrian Collision Avoidance Behavior

By

Kyra Veprek

B.A., University of Massachusetts Amherst, 2021

Thesis submitted in partial fulfillment of the requirements for the Degree of Master of Science in the Department of Cognitive, Linguistic, & Psychological Sciences at Brown University

Providence, RI

May 2023

This thesis by Kyra Elizabeth Veprek is accepted in its present form
by the Department of Cognitive, Linguistic and Psychological Sciences as satisfying the thesis
requirements for the degree of Masters of Sciences.

Date ______________________ ____________________________________
Dr. William H. Warren Jr., Advisor

Approved by the Graduate Council

Date ______________________ ____________________________________
Dr. Thomas A. Lewis
Dean of Graduate the School

Table of Contents

1. Introduction

1.1 Avoiding Collision with a Single Moving Obstacle

Walking on our daily commutes, we often navigate through bustling crowds of people on commutes of their own, but rarely do we end up colliding with any of them. Although this might seem like a trivial issue, successful collision avoidance requires an ability to perceive potential collisions with other individuals and control our own movements to avoid accidents. This problem is further complicated when successful navigation requires avoiding collisions with *multiple* moving people. To understand how an individual can avoid collisions in a crowd, we take a bottom-up approach by first asking how people avoid collisions with a single moving person.

Previous research on collision avoidance in traffic science, robotics, and computer animation has qualitatively described avoidance strategies and developed models that produce collision-free paths. However, few of these models reflect plausible human behavior. Humans exhibit consistent patterns of path geometry and velocity control when engaged in goal-oriented locomotion and stationary obstacle avoidance (Fajen & Warren, 2003; Hicheur et al., 2007). However, when a moving obstacle is introduced, agents must engage in a collision avoidance strategy. Two approaches to path planning can be applied to the problem of path planning: global path-planning approaches, in which an agent predicts the paths of all moving objects and computes an optimal trajectory, and local path-planning approaches, in which the agent plans a trajectory based on the motions of nearby objects. Thus, three potential strategies have been previously identified which might qualitatively inform how information is processed to generate a collision-free path. A fourth proposed strategy, however, best accounts for human locomotor behavior.

One strategy assumes a global, or feedforward, approach in which an agent computes an optimal trajectory in advance. Such global path planning assumes an agent's ability to predict the paths of all obstacles in the scene. In this case, online adjustments are secondary to a pre-planned trajectory in determining the ultimate path taken by an individual (Todorov, 2004). Paris, Pettre, and Donikian (2007) adopted this predictive global approach in a model that sought to address agent collision avoidance in simulated pedestrian navigation for animated films and television. For each navigation problem, the model explored the reachable space of the pedestrian in any given direction within a range of speed values, subsequently searching for possible collisions with neighboring obstacles. The final trajectory was then computed based on the optimal speed and orientation according to the environment state and goal.

The model produced by Paris et al. (2007) successfully solved many problems with local approaches, such as a continuous oscillation between agents or agents halting due to environmental constraints. Compared to human behavior, however, the simulated pedestrian demonstrated a lack of realism. One interpretation of this finding is that humans do not follow the pre-planning strategy employed in the model. Interestingly, human locomotion in the presence of static obstacles has been shown to follow a maximum smoothness constraint (Pham et al., 2007; Pham & Hicheur, 2009). Consequently, research conducted by Basili et al. (2012) involved several experiments in which chosen trajectories in a collision-avoidance paradigm were measured based on smoothness maximization. This set of experiments and simulations demonstrated that pedestrians do not select a path that is optimally smooth when avoiding a *moving* object, in contrast to static objects, suggesting humans use present feedback to make online adjustments in a locomotor trajectory.

A second possible collision avoidance strategy is a local approach in which the agent treats the obstacle as static at each time step. Researchers and engineers have employed this approach to collision avoidance as a solution to fundamental motion planning problems in robotics (Cai et al., 2007) and cellular automata models (Canny, 1998; Latombe, 2012). Such models have successfully guided the trajectories of robot systems in goal-oriented path planning and collision avoidance. While suitable for controlling robots, these algorithms do not reflect human locomotion and produce trajectories and speed variations that are not predictive of human behavior.

A third possible strategy for collision avoidance involves local planning coupled with global trajectory prediction to continuously estimate the path of an interfering obstacle (Todorov, 2004). In their 2012 walking experiments, Basili et al. (2013) found that participants decided whether to pass in front or behind a moving obstacle when located approximately 0.8 meters before passing the objects, suggesting the participant chose a trajectory that prevented the need for halting. This research further demonstrated, however, that the maximally smooth path did not predict the behavior of human pedestrians in response to a moving obstacle, unlike in the cases of a static obstacle. Maximum trajectory smoothness, therefore, does not underlie trajectory formation in the case of avoiding a moving obstacle. Humans, then, engage in online planning while walking. To develop a scientifically adequate model of human behavior, however, an understanding of the information humans use to update their approach must be clarified.

The proposed fourth approach assumes online control of locomotion informed by present perceptual information. Previous research provides robust evidence that sensory information guides trajectory formation (Warren et al., 2001; Wilkie & Wann, 2006). Further, vision largely guides human locomotion (Gibson, 1958; Rieser et al., 1995; Warren, 1998). Renault et al. (1990) made an early attempt to integrate visual input into a collision avoidance model in a project aimed at animating realistic virtual pedestrians. In Renault et al.'s (1990) simulations, an actor was designed to walk down a corridor in which it was required to avoid obstacles and moving pedestrians. The actors' movement was based on a combination of synthetic vision and Displacement Local Automata (DLA), which informs an agent's motion within the context of a specific environment. For example, this simulation included a DLA for "displacement in a corridor," and "obstacle avoidance," which could be combined to inform an agent's trajectory within a unique situation. The researchers implemented vision as a 30 × 30 pixels screen with known depth information for each pixel. This is an example of the two-step approach to solve collision avoidance, which assumes knowledge not available to humans based on perception, such as veridical distance. Therefore, this model does not accurately explain human collision avoidance behavior.

Ondrej et al. (2010) developed a synthetic-vision-based approach to collision avoidance, which relied on visual information to specify both future collisions and the level of danger that a collision presented. Trajectories were based on a steering strategy to avoid future collisions and a deceleration strategy to prevent imminent collisions. At each time point, the risk of future collision is computed based on the time derivative of the angle between the agent's heading and an obstacle's direction, and the risk of collision is computed based on the time-to-interaction. The results demonstrated the ability of their approach to generate self-organized patterns of moving agents; however, their model does not rely entirely on visual input. The computation of time-to-interaction requires omniscient knowledge of the velocity and distance of obstacles and, therefore, cannot accurately predict human behavior.

The model developed by Ondrej et al. (2010) was elaborated upon to replicate human behavior patterns more closely. Moussaid et al. (2011) proposed the cognitive heuristic model. The input to the model includes the computed distance to the first possible collision if the pedestrian moves at any given speed and direction, considering other pedestrians' walking speeds and body sizes. Two heuristics influence the subsequent path: First, pedestrians choose a direction of motion that allows the most direct path to a destination point, and second, a pedestrian maintains a distance from the first obstacle in the chosen walking direction that ensures a time to collision of at least τ, which is the period required for a pedestrian to stop in the case of an unexpected obstacle. While this model presents a meaningful shift away from the physics-inspired interaction models, it still depends on variables that do not reflect the visual information available to an agent. To develop a behaviorally-informed model, we must necessarily abandon omniscient models and strictly incorporate variables that reflect the visual information available to humans.

1.2 The Steering Dynamics Model

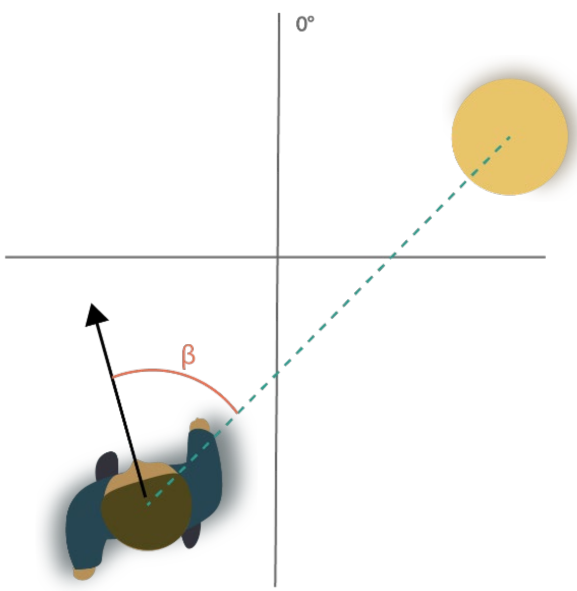

Figure 1.2.1. A diagram of the obstacle-heading angle (β) between a moving agent and an obstacle.

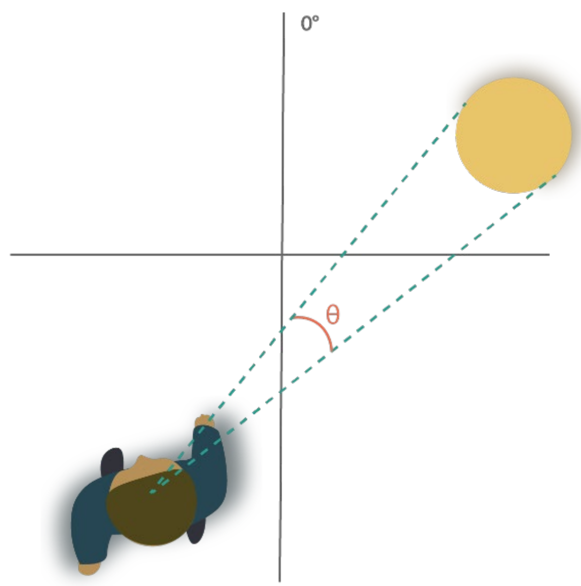

Figure 1.2.2. A diagram of the visual angle between an agent and an obstacle from the agent's perspective. The visual angle increases as the distance between the agent and obstacle decreases.

Notation	Meaning
ϕ	Agent heading direction
ψ_o	Bearing angle of obstacle
r_o	Distance to goal and obstacle
β_o,	Obstacle-agent angle
θ	Visual angle of obstacle
k, c	Model parameters

Table 1.2 Notation definitions

To develop such a model, previous researchers in the Virtual Environment and Navigation Lab at Brown University have identified the relevant visual variables necessary to control collision avoidance. At any given time point, β represents the angle between an agent's heading and the heading of a moving obstacle (See Figure 1.2.1). Similarly, the visual angle, θ, represents the size of an obstacle from the agent's perspective, which provides information to the agent about their distance from the obstacle (See Figure 1.2.2). These optical variables are the basis for the visual information necessary to inform a collision avoidance model. Optical Expansion ($\dot{\theta}$), or the first derivative of visual angle, is the increase of the size of an obstacle within an agent's visual field. Previous research from the Virtual Environment Navigation Lab has demonstrated that optical expansion is informative in pedestrian following (Rio et al., 2014; Bai & Warren, 2019). The Bearing angle is the angle between a reference direction (in our case, 0 degrees), and the line of sight between an agent and an obstacle. Prior work has also shown that bearing angle can be used in pedestrian models intended for both intercepting a moving target (Fajen & Warren, 2004, 2007) and avoiding a moving obstacle (Cohen et al., 2005; Ondrej et al., 2010; Warren & Fajen, 2008).

$$\ddot{\phi} = -k_9\dot{\psi}_o\left(e^{-c_{13}|\dot{\psi}_o|}\right)\left(1 - e^{-c_{14}max(0,\dot{\theta})}\right)\mathbb{1}_{|\beta_o|<\frac{\pi}{2}} \quad \textbf{(1.2a)}$$

$$\ddot{s} = k_{10}sign(\beta_o)\dot{\psi}_o\left(e^{-c_{15}|\dot{\psi}_o|}\right)\left(1 - e^{-c_{16}max(0,\dot{\theta})}\right)\mathbb{1}_{|\beta_o|<\frac{\pi}{2}} \quad \textbf{(1.2b)}$$

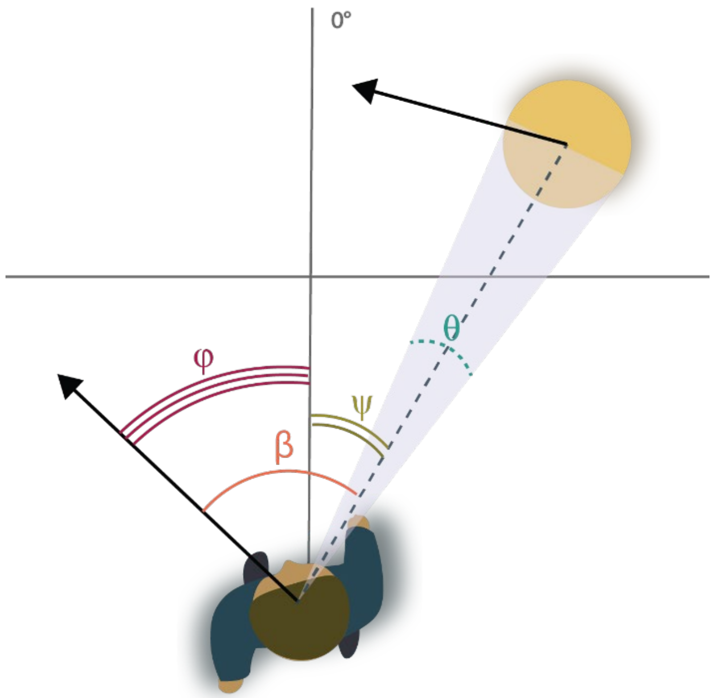

Figure 1.2.3. A diagram of the variables corresponding to visual information available to an agent, which inform the Steering Dynamics Model.

The collision avoidance model applied in the analysis in this paper is based on the model originally developed by John Cohen and William H. Warren, and later revised by Jiuyang Bai and William H. Warren. The foundation of the model is based on the observation that the rate of change of the bearing angle ($\dot{\psi}$) indicates the risk of collision with an obstacle (See Figure 1.2.3). For example, a constant bearing angle ($\dot{\psi}_o = 0$) indicates an agent is on course for a collision, while a fast-changing bearing angle indicates the obstacle and agent are moving away from a potential collision point. The model thus incorporates a repulsion function, $e^{-c_{17}|\dot{\psi}_o|}$, which forces the agent to move such that the bearing angle increases in a positive direction when an agent passes in front of an obstacle, and in a negative direction when passing behind. The magnitude of this response decays exponentially as the rate of change in the bearing angle increases. The agent thus prevents collision by avoiding a constant bearing angle in relation to an obstacle.

Both the heading and trajectory taken around an obstacle are determined by the repulsion function in the collision avoidance model (1.2a & 1.2b). If ($\dot{\psi}_o > 0$) the agent turns in the positive direction to increase the absolute value of ($\dot{\psi}$), while if ($\dot{\psi}_o < 0$), the agent turns in the negative direction to increase the absolute value of ($\dot{\psi}$). This response is modulated by the imminence of collision, which is defined by the equation $1 - e^{-c_8 max(0,\dot{\theta})}$. This term, which is a function of optical expansion, amplifies the angular acceleration as the distance between the agent and the obstacle decreases. When the rate of optical expansion is equal to zero, indicating a constant distance between the agent and obstacle, the function is also equal to zero. Meanwhile, when the rate of optical expansion is large, close to or reaching one, full power is extended to steering. While $1 - e^{-c_8 max(0,\dot{\theta})}$ is correlated with distance, it allows the model to reflect better the perceptual information available to an agent than previous models which relied on distance as an omniscient variable.

Speed control (1.2b) is similarly determined by the repulsion function. The term $k_{12} sign(\beta_o) sign(\dot{\psi}_o)$ determines the passing order of the agent and the obstacle. When the signs of β and $\dot{\psi}$ differ, and the net value is therefore negative, the agent decelerates to pass behind the obstacle comfortably. Alternatively, when the signs of β and $\dot{\psi}$ are the same, the agent accelerates to pass in front of the obstacle. So, when $\beta < 0$ and $\dot{\psi} > 0$, the obstacle is on the left side of the agent, and the model will decelerate and steer to the left of an object, causing rear passing. When, $\beta < 0$ and $\dot{\psi} < 0$, the obstacle is on the left side of the agent, and the agent accelerates to pass in front. When $\beta > 0$ and $\dot{\psi} > 0$, the obstacle is on the right of the agent and the agent accelerates to pass in front of the obstacle. Finally, when $\beta > 0$ and $\dot{\psi} < 0$, the obstacle is on the right of the agent and agent decelerates to pass behind the obstacle.

The Steering Dynamics model demonstrated robust results in comparison to human behavior in collision avoidance with moving non-human obstacles. In experiments conducted by Bai for his dissertation, the model produced small distance errors and high passing order accuracy. Across three experiments in which participants were tasked with avoiding moving obstacles while walking towards a goal, the distance errors remained below 30 cm, and ranged between 14-16 cm with the introduction of a visual threshold. One limitation of Bai's work, however, is that, across all three avoidance experiments, participants only encountered a moving non-human pole. There may be qualities unique to avoiding humans that would require parameter adjustments for the model (Zhou et al., 2019); Perhaps human behavior in response to humans either a) requires a reweighting of the model parameters or b) cannot be properly modeled based solely on higher order visual information. The first experiment addressed in this paper outlines our approach to this problem.

1.3 Navigation Through and Around a Crowd

The problem of avoiding multiple moving obstacles has been described in previous research as well, primarily through investigations of pedestrian and crowd interactions. One primary qualitative question about human interaction with a moving crowd is whether an agent treats each interaction with another pedestrian sequentially, or chooses a path based on the simultaneous processing of multiple pedestrians. Researchers Meerhoff et al. (2018) sought to answer this question by comparing avoidance strategies between real world pedestrians in dyadic (1 person vs. 1 person) interactions versus triadic interactions (1 person vs. group of 2) in terms of the evolution over time of the Minimal Predicted Distance between pedestrians and the Dynamics of the Gap, which includes passing order and perceived possibility of a gap. The results from this experiment demonstrated that it is difficult to draw broad conclusions about the adoption of collision avoidance strategies; Some triadic interactions invited sequential interactions, producing similar avoidance strategies as dyadic interactions, while others resulted in simultaneous interactions.

Another point of investigation in the avoidance of multiple moving obstacles is the conditions under which a pedestrian elects to pass through or deviate around a crowd. Researchers Bruneau et al. (2015) conducted an experiment in which participants were asked to go to a visible target in a virtual environment, which included a virtual crowd. The crowd was modified considering three independent conditions: Interpersonal Distance ({1.1, 1.4, 1.7, 2, 2.3} m), Group Relative Movement (approaching from head-on towards the participant or from either side) and Group Type (ordinary people, zombies, or soldiers). Trajectories were evaluated based on energetic cost. Results demonstrated that users elected to go around a crowd when the group motion was orthogonal to their own motion, and that Predicted Minimal Expenditure can be a useful predictor of pedestrian behavior. Further, subjects make a smooth transition between passing through the crowd at low densities and around the crowd at high densities, with the 50% point occurring at an interpersonal distance of 1.7m.

Approaches similar to those used in modeling single object obstacle avoidance have also been taken in efforts to model multiple-obstacle collision avoidance. To realistically model human avoidance behavior, physicist Francesco Zanlungo (2007) developed a model based on Theory of Mind. To explain human behavior, Zanlugo and colleagues represented each agent as a rigid disc that undergoes elastic collisions. Each agent has a goal and a fitness, which is defined as a positive term given by the inverse of time needed to reach a goal plus a negative one determined by the momentum exchanged in bounces, which is meant to stand in for the experience of "pain." The evolvable parameters include the visual experience (radius and angle of sight) and the detection mechanism (perception of its own size, attraction to the goal, and Theory of Mind level that determines the prediction of the other agents' motion).

Zanlungo's model accurately reproduced some of the basic features of the organized motion of pedestrians, and he used his findings to later elaborate upon the Social Force Model developed by Helbing in 1999 (Zanlungo et al., 2011). Within the Social Force Model Framework, Zanlungo introduces an explicit computation for collision prediction. The key concept of the model is that each pedestrian will try to compute, for each approaching pedestrian in her environment, the time in the future at which their relative distance will be minimum. Based on this, circular interaction forces are applied based on the minimal predicted distance in the imagined future situations. The simulations of this model demonstrated that introducing an explicit computation of collisions improves the capability of the Social Force Model to reproduce pedestrian trajectories. While the Social Force Model has been successful in reproducing crowd behavior at the macroscale, it fails to generate human-like trajectories at the local level. Further, modeling human behavior using physical forces is not biologically motivated, as they are merely a metaphor for the visual information actually used to control locomotion.

An outstanding question, then, is whether the steering dynamics model can be generalized to a pedestrian navigating through a crowd of moving pedestrians. Experiment 1 tested whether the collision avoidance model generalizes from avoiding a single moving object (a pole) to avoiding a walking person (a virtual avatar). The ultimate goal is to simulate crowd behavior in which agents act and react in response to one another, producing collective behavior such as that seen in crossing flows (Naka, 1977). As a step towards this goal, Experiment 2 tested whether the collision avoidance model generalizes from avoiding a single avatar to avoiding collisions with a crowd of avatars.

2. Experiment 1: Collision avoidance with a pole or a human

The goal of Experiment 1 was to test and evaluate whether there are differences in the walking trajectories taken by participants when avoiding a moving pole and a moving avatar. Previous research has modeled human pedestrian collision avoidance behavior in response to a single moving pole (Bai, 2022), and this experiment seeks to determine whether that model can be generalized to the avoidance of collisions with walking pedestrians. Participant trajectories in the pole and avatar conditions were compared using several quantitative measures. The performance of the Steering Dynamics Model was evaluated across both conditions.

2.1 Method

Participants

Separate groups of 12 participants were recruited for Experiment 1 (7F, 5M), using an online bulletin board through Brown University email. No participants reported any visual or motor impairments, or any nausea due to the virtual environment. The research protocol was approved by Brown University's Institutional Review Board, in accordance with the principles expressed in the Declaration of Helsinki. Informed consent was obtained from all participants, who were paid for their participation.

Equipment

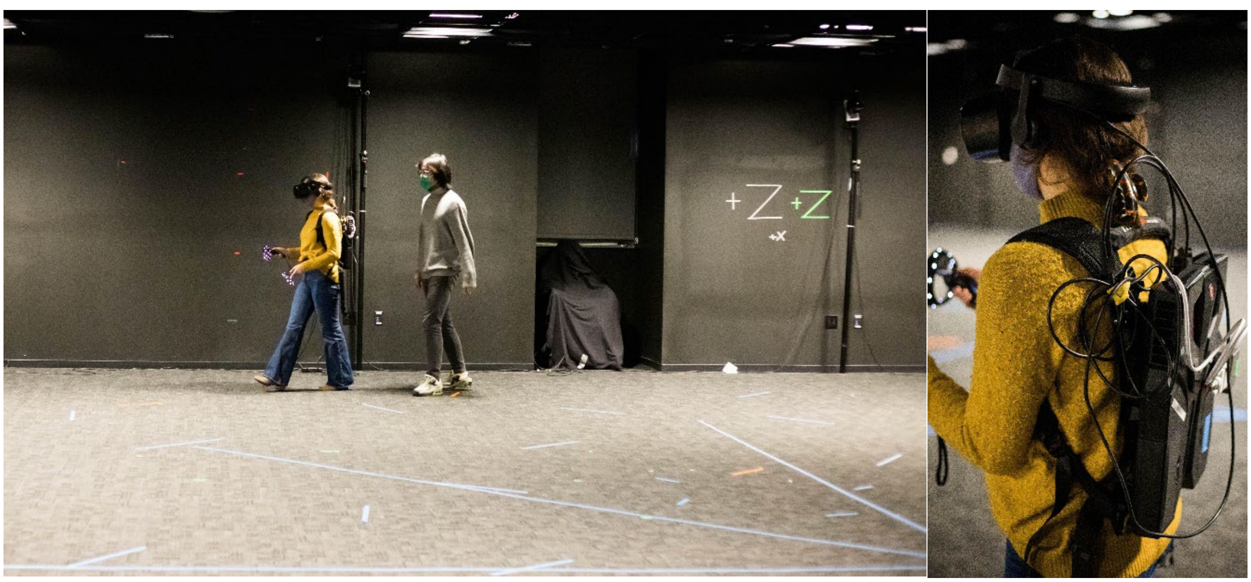

Figure 2.1.1. A participant wearing the Odyssey head-mounted display device (a) and MSI backpack (b) computer while walking in the 12x14 m free walking space.

All experiments were conducted in Brown University's Virtual Environment Navigation Lab (VENLab). Participants freely walked in the 12x12 m lab space while viewing a virtual environment in a wireless head-mounted display (HMD) device. The HMD was the Samsung

Odyssey (Seoul, S. Korea), 101°H x 105°V, 1600 x 1440 pixels per eye, 90 Hz refresh rate. Head position and orientation were recorded using the HMD's inside-out tracking system, which consisted of two onboard cameras and an inertial measurement unit (Accelerometer (6 Axis), Gyrometer (6 Axis), Compass (3 Axis)). Stereoscopic displays were generated on an MSI VR-One backpack computer (New Taipei City, Taiwan, weight 3.3 kg, Intel Core i7-7820HK processor, NVidia GTX 1070 graphics) at a frame rate of 90 fps, with a total latency of ~11ms.

Displays

The virtual environment consisted of a ground plane with gray granite texture and a light blue sky. Several vertical poles rested on the ground, all with granite textures: a blue start pole (radius 0.12 m, height 1.35 m), a yellow orientation pole (radius 0.2 m, height 3 m), a green goal pole (radius 0.2 m, height 3 m), and a moving red obstacle pole (radius 0.1 m, height 1.8 m). The goal pole appeared at a distance of 26.3m from the participant. The initial obstacle positions fell on a circle with a radius of 5.1 m, with the participant at the origin. The initial positions of the obstacle relative to the position of participant were (0m, 10.2m), (±2m, 9.8m), (±3.6m, 8.7m), (±4.7m, 7.1m), and (±5.1m, 5.1m). The obstacle pole moved across the participant's path at 9 different angles (180, ±157.5, ±135, ±112.5, ±90°), with 2 different speeds (1.0, 1.2 m/s). Obstacle trajectories are illustrated in Figure 2.1.2.

Procedure

Figure 2.1.2 Experimental Design. The participant was asked to walk towards the goal pole while avoiding a collision with a moving obstacle.

Before each trial, only the start and orientation poles were visible. The participant was instructed to stand at the start pole and face the orientation pole. After 2 seconds, the start and orientation pole disappeared, and the goal pole appeared behind the orientation pole at a distance of 26.3 m; the participant was instructed to walk toward the goal while avoiding the moving obstacle. After the participant had walked 2 m, either a moving obstacle pole or walking avatar appeared. The participant walked past the obstacle and toward the goal until they were <2 m from the location of the next start pole, at which point the start and orientation poles for the next trial appeared. Participants were informed that they could take a break or discontinue the experiment if they felt uncomfortable or experienced symptoms of simulator sickness or felt uncomfortable. All participants completed the experiments without incident.

Design

Experiment 1 had a factorial design with 5 moving obstacle angles (±180, ±157.5, ±135, ±112.5, ±90°) × 2 obstacle speeds (1.0, 1.2 m/s) × 2 obstacle types (pole vs. avatar), with six trials per condition. 10 free-walk trials, in which a participant walked freely to the goal, for a total of 130 trials. All variables were within-subject, and trials were presented in a random order.

Data Processing

The basic data from each trial were the time series of head position in the horizontal X-Y plane. Each time series was filtered using a forward and backward fourth-order low-pass Butterworth filter to reduce the effect of gait oscillations due to the step cycle. A 1.0 Hz cut-off was used prior to computing speed to reduce anterior-posterior oscillations on each step, whereas a 0.6 Hz cut-off was used prior to computing heading to reduce lateral oscillations on each stride. The first two seconds of data corresponding to standing were removed. Time series analyses of walking speed and heading direction were computed from the filtered position data. Right and left turn trials were then collapsed by multiplying the lateral deviation on left trials by -1. A mean time series of heading and speed was computed for each subject in each condition, and the mean final heading and mean final speed were computed over the last 2s of each. The collapsed data were analyzed using Jasp statistical software. All trials with an obstacle angle of ±180° were excluded from model comparisons because head-on collisions produced a high uncertainty regarding the side of avoidance (left vs. right).

Data Analysis: Passing Order, Maximum heading, and Maximum deviation

When avoiding collision with a moving obstacle, participants have two options, which we call the *passing order*: to pass behind the moving obstacle or in front of the moving obstacle. The collision avoidance model relies on the sign of the obstacle-heading angle (β, where positive is to the right of the heading direction) and the sign of the change in bearing angle ($\dot{\psi}$, where positive means the obstacle is moving clockwise) to control the order of passing.

A second measure used in analyses was the *maximum lateral deviation* of the participant's position from the center line. Maximum deviation is meant to give us insight as to how far participants moved from the most direct path to the goal pole. Maximum distance was computed as the maximum distance from the center line in the perpendicular direction.

A third measure used in analyses was the *maximum heading* angle reached while the participant was in motion. Maximum heading angle is meant to give us insight as to the degree to which participants steered away from a direct heading to the goal. Maximum heading angle was computed as the angular deviation from the line extending from the start pole to the goal pole; a participant walking directly from the start to the goal would have a maximum heading angle of zero.

Data Analysis: Model Comparison

k_{11}=6.19, c_{17}=5.27, c_{18}=16.70, k_{12}=2.20, c_{19}=8.10, c_{20}=21.27

Table 2.1.1 Fitted parameters with the visual threshold.

The parameters used in model simulations (see Table 2.1.1) were taken from Bai (2022), who fit them to experimental data. Bai also discovered the importance of a visual threshold for the onset of collision avoidance behavior, which is based on the observation that humans only initiate collision avoidance once needed. He found that best-fit to the human data was a threshold on the combined rate of optical expansion ($\dot{\theta}$) and bearing change ($|\dot{\psi}|$). We used $\dot{\theta}|\dot{\psi}|$ = 0.03 °/s as the threshold for the first experiment (see Bai, 2022) by initializing the simulated trials once the participant had reached the threshold.

The primary goodness-of-fit measure for the model was the Root Mean Squared Error. For each individual trial for each subject, the model was initialized with the initial position, heading, and speed of the participant, and the optical variables for the moving obstacle were input to the model. At each time step (90 Hz) the root mean squared error (RMSE) between the participant's position and the model's predicted position was computed to determine the Euclidean distance between the two positions. The mean RMS Error was then calculated for each trial.

2.2 Results

We first analyze the human data in the experimental conditions, with particular interest in comparing trajectories around the pole and the avatar, and then report model performance in those two conditions.

Passing Order: Binomial Logistic Regression

Predictor	**χ2**	***df***	**p**
Obstacle Type	2.53214	1	0.488
Trajectory	2.14716	3	0.544
Speed	6.75212	2	0.009
Trajectory * Obstacle	4.49412	3	0.213
Trajectory * Speed	2.89894	3	0.407
Obstacle * Speed	1.91186	1	0.167
Trajectory * Obstacle * Speed	1.81567	3	0.612

Table 2.2.1 Omnibus Likelihood Ratio Tests

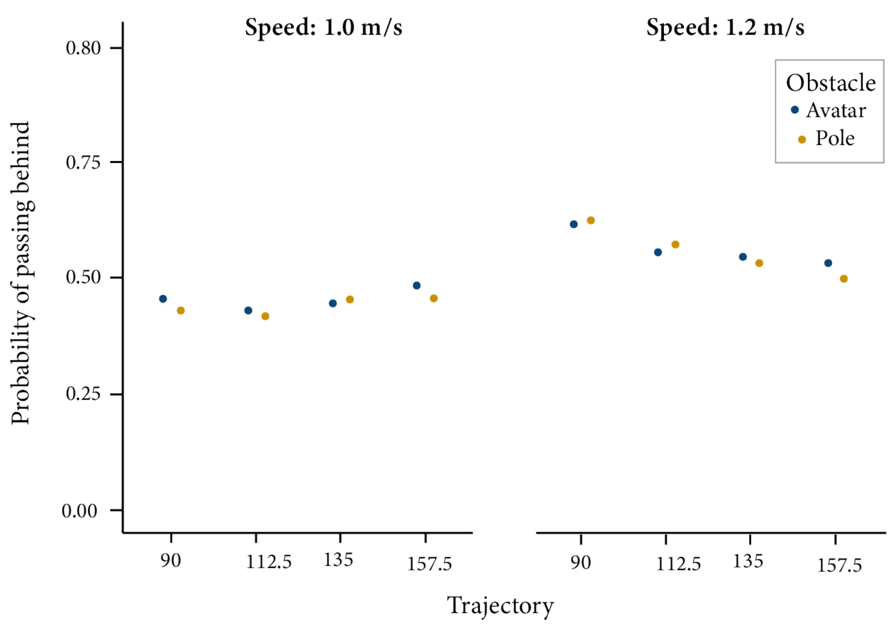

Figure 2.2.1 Binomial Logistic Regression. The probability of the participant passing behind dependent on each the speed and trajectory of the obstacle was compared

To determine whether there are differences in passing order between conditions, we performed a binomial logistic regression, with explicit interest in whether passing order differed for the pole and the avatar (the obstacle type condition). Statistical results appear in Table 2.2.1.

The Omnibus Likelihood Ratio Test revealed that the passing order is dependent on the speed of the obstacle ($\chi^2 = 6.75212, p = 0.009$): the probability of the participant passing behind the obstacle was greater in the high speed condition (Figure 2.2.1, right panel) than in the low speed condition (left panel). There was no evidence of a significant effect of trajectory angle or obstacle type, or any interactions, on the passing order. This finding suggested that participants did not adopt different strategies when avoiding pedestrians versus avatars.

Maximum Deviation: Bayesian ANOVA

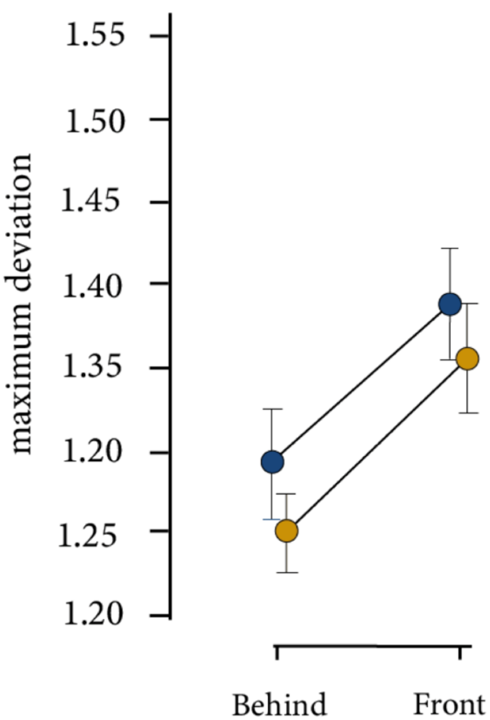

Cases	P(incl)	P(incl|D)	BF_{incl}
Front/Behind	.6	1.00	5.904E+7
Obstacle Type	.6	0.413	0.469
*Front/Behind * Obstacle Type*	.2	0.062	0.281

Figure 2.2.2. Bayesian Analysis of Variance on the maximum lateral deviation taken by the participant. The yellow dots correspond with the mean of pole trials and the blue dots correspond with the mean of the avatar trials.

Table 2.2.1 Bayesian analysis of Effects on Maximum Heading

To determine effects on the participant's trajectory, we first performed a Bayesian Analysis of Variance on the maximum lateral deviation. We hypothesized that there would be no difference between avoiding avatars and avoiding poles, and thus we did not expect to find statistical effects of Obstacle Type on maximum deviation. Indeed, the ANOVA yielded moderate evidence against an effect of obstacle type ($\boldsymbol{BF_{incl}} = 0.469$) and anecdotal evidence against an interaction between obstacle type and passing order ($\boldsymbol{BF_{incl}} = 0.281$). In contrast, there was decisive evidence for a main effect of passing order on maximum deviation ($\boldsymbol{BF_{incl}} = 5.9047e^{7}$), indicating that participants had larger heading deviations when going in front of a moving object than when passing behind it.

Maximum Heading: Bayesian ANOVA

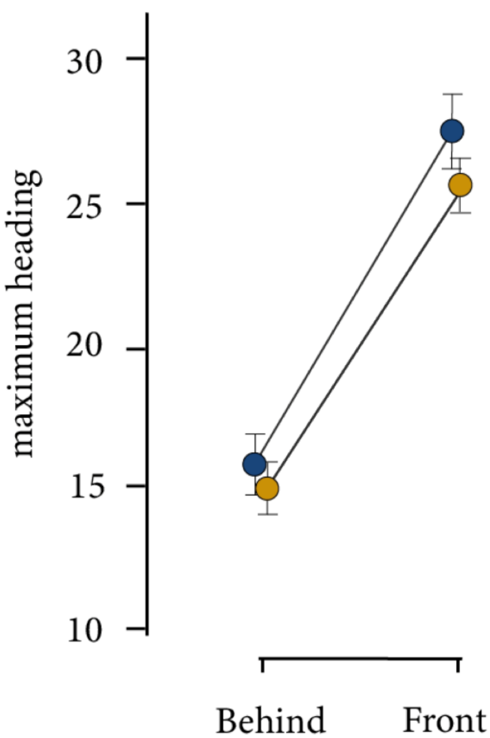

Cases	P(incl)	P(incl|D)	$\boldsymbol{BF_{incl}}$
Front/Behind	.6	1.0	6.672E+14
Obstacle Type	.6	0.052	0.222
*Front/Behind * Obstacle Type*	.2	0.164	0.131

Figure 2.2.3. Bayesian Analysis of Variance on the maximum heading taken by the participant. The yellow dots correspond with the mean of pole trials and the blue dots correspond with the mean of the avatar trials.

Table 2.2.2 Bayesian Analysis of Effects on Maximum Heading

Next, we performed a Bayesian Analysis of Variance on the maximum heading to further determine the effects on the participant's trajectory. Once again, we did not expect to see a difference between avatars and poles. Consistent with this hypothesis, the ANOVA yielded moderate evidence against an effect of obstacle type ($\boldsymbol{BF_{incl}} = 0.222$) and moderate evidence against an interaction between obstacle type and passing order ($\boldsymbol{BF_{incl}} = 0.131$). On the other hand, there was decisive evidence for a main effect of passing order ($\boldsymbol{BF_{incl}} = 6.672e^{14}$), with a larger maximum heading when passing in front of than behind the obstacle.

Model Performance

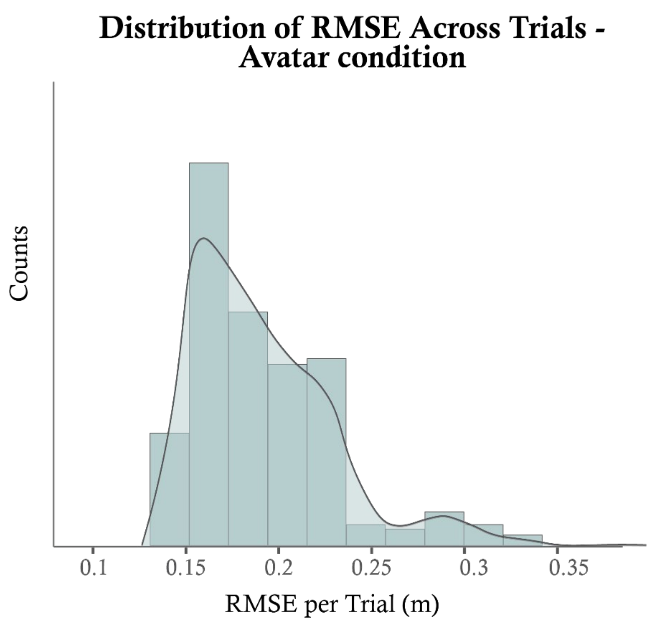

Mean RMS Error (m)	0.178
Standard deviation	0.38
Passing Order Accuracy	0.96

Figure 2.2.4. A histogram displaying the average RMS Error per trial, in meters, in the trials in which participants avoided an avatar.

Table 2.2.3 Measures of model performance for trials in which participants avoided an avatar

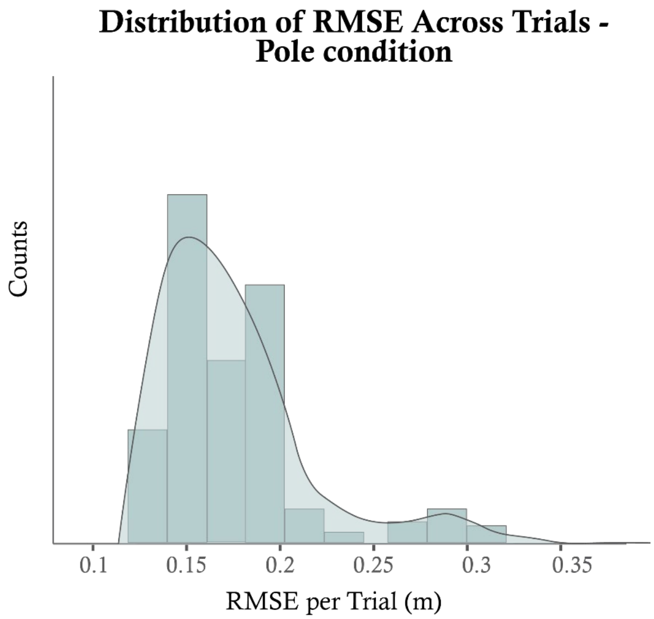

Mean RMS Error (m)	0.173
Standard deviation	0.38
Passing Order Accuracy	0.98

Figure 2.2.5. A histogram displaying the average RMS Error per trial in the trials in which participants avoided a pole.

Table 2.2.3 Measures of model performance for trials in which participants avoided a pole.

The mean distance error per trial was used as the measure of model performance across trials. The results from the experiment in Bai's dissertation were replicated, with the model producing a mean error of 0.176m across all trials, with a passing order accuracy of 97 percent. The model performed equally as well for the trials in which participants avoided an avatar (mean error = 0.178, SD = 0.38) and the trials in which participants avoided a pole (mean error = 0.173, SD = 0.38). The high performance of the model in the avatar condition suggested to us that the model would not need to be refit to account for differences in behavior.

2.3 Discussion

The descriptive results, including passing order, maximum deviation, and maximum heading, support the claim that participants did not use different strategies when avoiding a pole versus an avatar. This is further supported by the similar performance of the model in the two obstacle type conditions. While this is the result we expected, it does appear to contradict models of a front-back asymmetry in personal space, such as that developed by Zhou et al.(2019). There are two explanations which could account for the apparent lack of effect of personal space when comparing the trajectories of people around avatars versus poles. First, because the pole moves with a specific velocity, there is an asymmetry between 'front' and 'back' trajectories based on its direction of motion, with a larger deviation when passing in front. Under this explanation, the moving non-human obstacle assumes the qualities of personal space, causing the collision avoidance behavior to be the same as when avoiding a human. A second explanation is that, when avoiding a person in motion, the effect of an individual's personal space is nulled or reduced.

Based on our results from Experiment 1, we concluded that there is no need to refit the model to accommodate differences in the locomotor behavior of avoiding moving humans versus moving non-human obstacles. The next question, then, is whether the model can be generalized to the avoidance of multiple moving humans when walking through a crowd.

3. Experiment 2: Collision avoidance in a crowd

The goal of Experiment 2 was to model collision avoidance as an agent walks through a crowd, when avoiding multiple moving pedestrians. As a first step, Experiment 2 tests a participant avoiding multiple avatars on prescribed trajectories, which stayed on a preplanned course and did not respond to the participant. We then simulated the data with the collision avoidance model, to determine whether it would generalize to multiple moving objects. Based on the results, we modified the model and used it to investigate human collision avoidance strategies.

If the steering dynamics model generalizes from a single avatar to multiple avatars, this would imply that all obstacles within the visual threshold are processed simultaneously. If the model requires modifications, their nature would reveal that humans may combine obstacles within a given distance (Cohen, Bruggeman, & Warren, 2005), or process avatars sequentially (Meerhoff, et al., 2018). On the other hand, a failure to fit the data would imply that a different approach is needed, perhaps based on other optical variables, shooting gaps, or predicting obstacle trajectories.

3.1 Method

Participants

Twelve participants were recruited for Experiment 2 (7F, 5M), using an online bulletin board through Brown University email. No participants reported any visual or motor impairments, or any nausea due to the virtual environment.

Displays

Equipment and displays were the same as in Experiment 1, except for the following exceptions. In Experiment 2, sixteen randomized avatars from the Vizard library rested on the ground in a pseudo-random pre-determined configuration. The crowd moved across the participant's path from 9 different angles (180, ±157.5, ±135, ±112.5, ±90°). The members of the crowd always walked at 1.0 m/s.

Procedure

The procedure was the same as in Experiment 1, except that the participant was tasked with walking through a crowd of sixteen avatars.

Design

Experiment 2 had a factorial design with 5 crowd angles (±180, ±157.5, ±135, ±112.5, ±90°) × 2 crowd densities × 2 crowd configurations, with five trials per condition. 8 free-walk trials were added, in which a participant walked freely to the goal, for a total of 108 trials. All variables were within-subject, and trials were presented in a random order.

3.2 Experiment 2 Results

Visual Model Performance

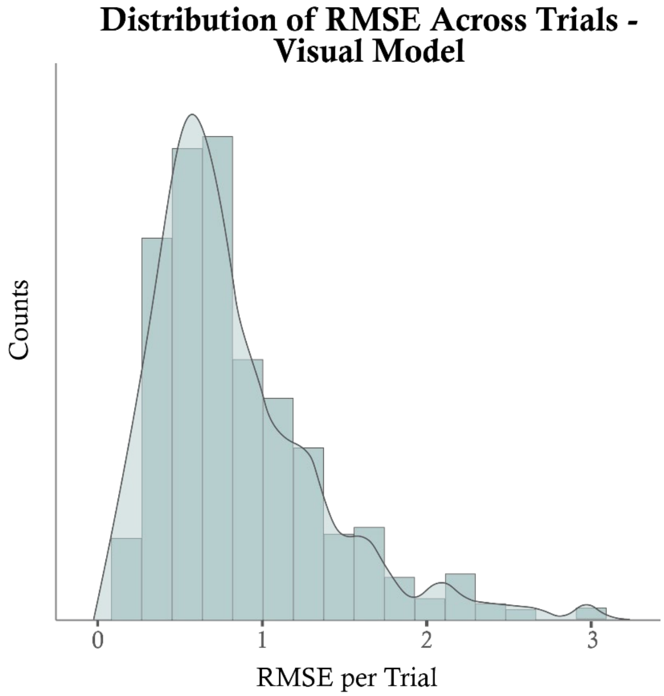

Mean Error (m) ***RMSE***	0.848
Standard Deviation	0.512

Figure 3.2.1. A histogram displaying the average RMS Error per trial between the behavioral data and the Steering Dynamics Model.

Table 3.2.1 Measures of the Steering Dynamics model performance across all trials

For each trial, the simulation was initialized with the initial position, heading, and speed of the participant at the time at which the crowd appeared. The mean positional error was computed at each time point, and averaged across each trial. The mean error across trials was high (Mean Error = 0.848m, SD = 0.512) compared to the model performance in Experiment I (Mean Error = 0.176 m, SD = 0.38).

To gain insight on the shortcomings of the model in this scenario, we observed the behavior of the simulated agent at the trial level. What we found was that the simulated agent appeared to greatly slow down at the beginning of each trial, ultimately passing around the periphery of the group. This suggested to us that, as the model stood, the speed control component was too sensitive to account for real human behavior; the sum of the collective optical expansion of the crowd members was inhibiting realistic performance of the model.

Visual Model Performance without Speed Control

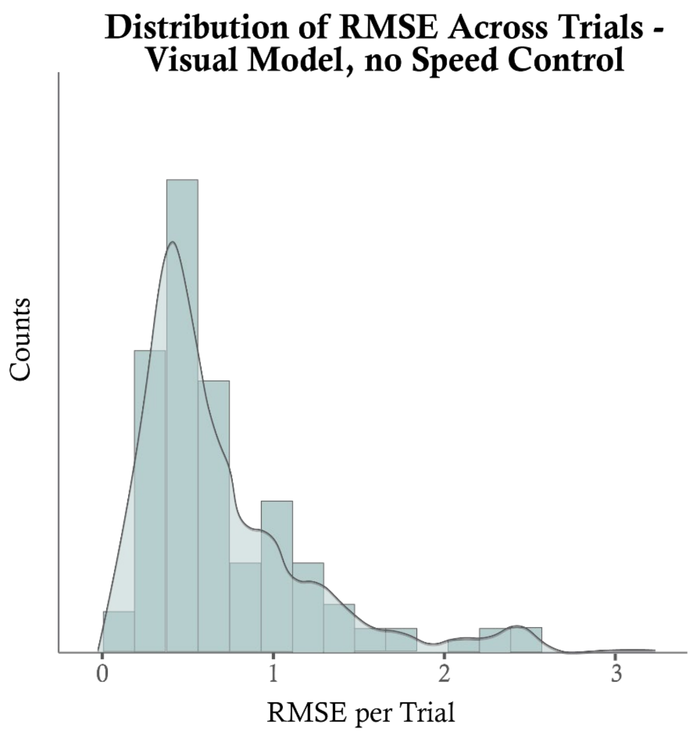

Mean Error (m) *RMSE*	0.779
Standard Deviation	0.468

Figure 3.2.2. A histogram displaying the average RMS Error per trial between the behavioral data and the Visual Model without speed control

Table 3.2.2 Measures of the Visual model performance without speed control across all trials.

Our first step in identifying whether speed control could account for the large mean error was to simulate trials without the speed control component, such that the simulated agent only adjusts their heading as a means of collision avoidance. As with the previous comparison, the model was initialized with the initial position, heading, and speed of the participant at the point in time at which the crowd of avatars appeared.

Removing speed control greatly improved the model's performance, with the mean error reducing to 0.779m (*SD = 0.468*). This confirmed for us that one primary problem with the model was the sensitivity of the speed component. However, this does not agree with our intuition that humans do engage in some speed control in order to avoid collisions. While this adjustment is an improvement, it does not reproduce real human behavior.

Omniscient Model Performance with Distance Threshold

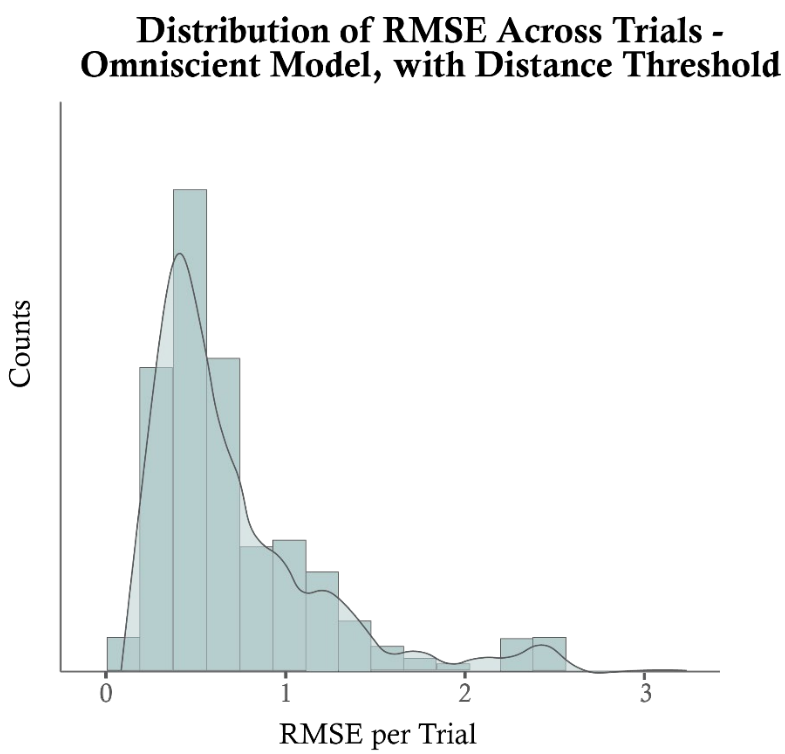

Mean Error (m) ***RMSE***	0.698
Standard Deviation	0.476

Figure 3.2.2. A histogram displaying the average RMS Error per trial between the behavioral data and the omniscient model, with a distance threshold of 2m.

Table 3.2.2 Measures of the Steering Dynamics model performance without speed control across all trials.

A second possible explanation for the errors of the full visual model is that agents are not responding to all potential collisions simultaneously. To gain a better intuition about this problem, we introduced an arbitrary distance threshold of 2m at which participants began to respond to moving obstacles. Because this threshold was based on Euclidean distance instead of optical expansion, we used the “omniscient” version of the collision avoidance model. Out of the three versions of the model tested against the behavioral data, this version performed the best (Mean error = 0.698m, *SD* = 0.476). This suggests that humans do not simultaneously avoid all upcoming collisions, but rather respond to potential collisions with some degree of sequential order.

Model Fitting

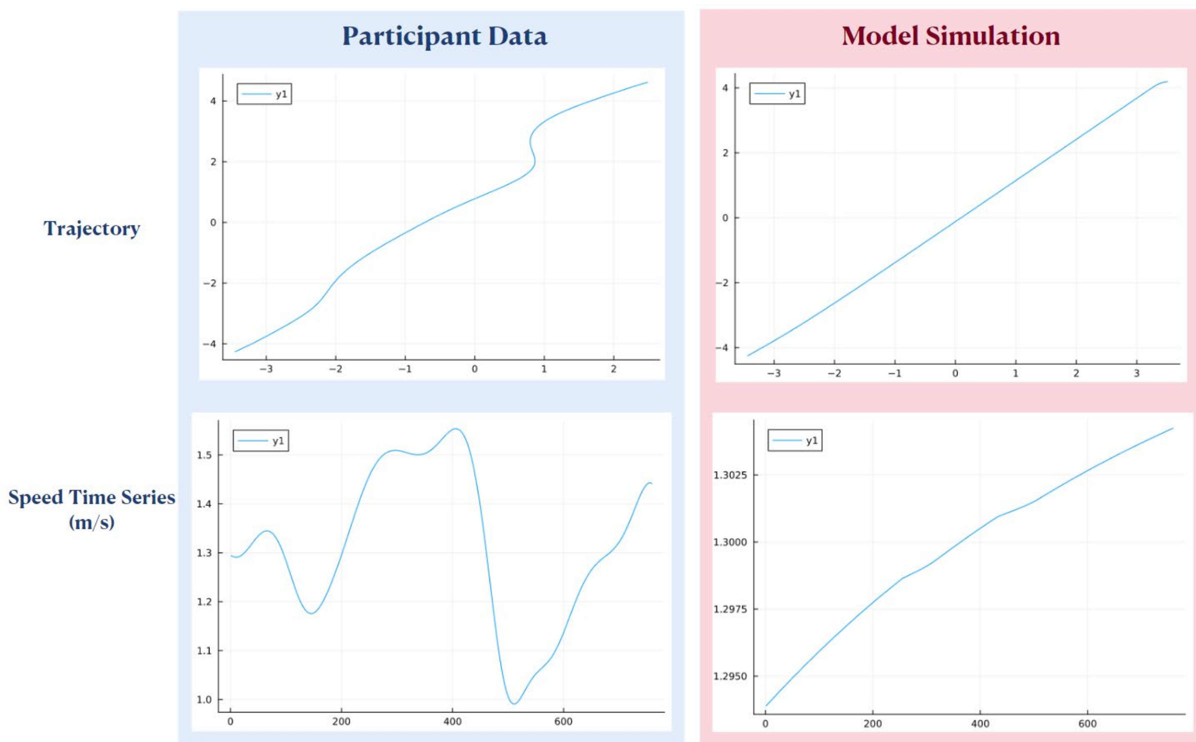

Figure 3.2.3. Graphs demonstrating model data based on the best-fit parameters, compared to real participant data. The comparison demonstrates the major limitation of the best-fit model: it favored straight trajectories, which do not precisely match human behavior.

k_{11}=1, c_{17}=9, c_{18}=1, k_{12}=1, c_{19}=1, c_{20}=1

Table 3.2.1 Best-fit parameters for the model given the participant data from Experiment 2.

The next step was to fit the model's parameter values to see the degree of precision which could be reached using the model. From a course grid search, the best-fit parameters (see Table 3.2.1) yielded an average RMSE of 0.737m. Upon comparing the participant trajectory and Speed Time Series with the simulated model data, it was found that the best-fitting model had a stiffness which produced straight trajectories and near-constant speed.

3.3 Discussion

The comparison of the behavioral data to the three different versions of the model revealed the shortcomings of the collision avoidance model when used to avoid multiple moving obstacles. The first comparison demonstrated that the visual model, which was largely successful in generating human-like behavior in the case of avoiding a single moving obstacle, only partially generalizes to avoiding multiple moving obstacles. The second comparison, to the visual model without speed control, confirmed our intuition that the shortcoming was the sensitivity of the speed component. The third comparison to the omniscient model with a 2m distance threshold revealed that participant responses to upcoming collisions do not occur simultaneously.

One problem that arises from these results is how to most effectively compute a threshold for initiating collision avoidance behavior in response to a moving obstacle. Three potential measures might inform a potential threshold. One possibility is a distance threshold, as we used in our third model comparison in Experiment 2. While the distance threshold improved model performance, it depends on the visual perception of object distance, which is known to be inaccurate. A second possibility is a threshold on optical expansion, as used in the simulations in Experiment 1. While this threshold correlates with distance, the advantage is that it is an optical variable that can be perceived. A third possible threshold would be a ‘nearest neighbor’ threshold, which would limit which potential collisions an agent responds to in accordance with a limited processing capacity. Future research will be conducted to test which type of threshold best accounts for the present data.

Furthermore, the model fitting procedure demonstrated that the best performing model was generating near straight trajectories, which does not reflect true human path formation. We hypothesize that our best performing model produced straight trajectories because it most closely fits the average trajectory. The problem of avoiding multiple sequential collisions lies in the abundance of bifurcation points for each trajectory; If an agent passes eight obstacle pedestrians within their path, there are eight sequential opportunities for the participant to make a passing order decision different from that of the model. The problem of a proliferation of bifurcation points also means that, if a single passing order decision is “incorrect” in accordance with the model early in the trial, the error will be high despite the possibility that a human pedestrian may have formed a similar trajectory, had they made the same passing order decision.

4. General Discussion

The first experiment replicated the success of Bai’s (2022) visual collision avoidance model with a single moving obstacle, and demonstrated that the model generalizes to the avoidance of moving pedestrians. This result supports our claim that human pedestrians use the visual information of optical expansion and bearing angle change to control collision avoidance behavior. Future work is needed to determine why principles of personal space do not have implications for avoiding moving obstacles.

The second experiment revealed differences between avoiding a single moving obstacle and avoiding multiple moving obstacles. Primarily, the initial poor performance of the visual model suggests that humans do not respond simultaneously to all potential collisions. This is true in the case for the control of speed especially. Our comparisons suggest that a threshold on optical expansion rate, distance or number of moving obstacles may explain human collision avoidance behavior. Future work will determine which type of threshold best represents the strategy used by humans to select which moving obstacles to avoid.

The problem of the threshold also presents the question of whether our interaction with a crowd is limited by our attentional set. Research in the domain of Multiple Object Tracking (MOT) has demonstrated that humans are able to track up to 5 independently moving obstacles accurately and that more objects can be tracked at lower speeds (Pylyshyn & Storm, 1988; Alvarez & Franconeri, 2007). Humans may respond only to future collisions with moving obstacles to which they are attending. The capacity limit for object tracking could account for a bound on how many potential collisions humans respond to when avoiding multiple moving obstacles.

Another possibility, as mentioned before, is that the limit is not related to attentional set, but rather optical expansion. This conclusion most closely aligns with our intuitions: given that optical expansion provides information about the collision imminence, we may only initiate collision avoidance once an optical expansion threshold has been reached. For our second experiment, fitting the model included fitting the threshold of optical expansion, which originated in Bai's 2022 experiments. As it currently stands, the fit of this threshold yet to prove to be descriptive, as the best-fit model generates near straight trajectories.

The shortcomings in model performance for the second experiment introduce another primary question about model fitting procedures in path formation. In this study, we elected to use RMS Error to measure of model performance because that is the standard method used in path formation and locomotion research. However, this measure may not be optimal for problems of avoiding multiple sequential collisions; Each collision affords a bifurcation point in the path formation in which an agent elects to pass in front of or behind the obstacle. As mentioned before, despite electing the wrong passing order, the model still produces human-like trajectories. Multiple other measures have been considered for future model fitting, such as RMSE for heading, RMSE for speed, and passing order accuracy. While each measure of fit presents both advantages and disadvantages, further work is needed to determine how to optimally assess and fit the model in the case of multiple moving obstacle collision avoidance.

Limitations and Future Work

One potential constraint of this work is the need for more analysis in social and cultural domains. Cross-cultural comparisons have demonstrated differences in preferences for interpersonal distance and orientation toward others during interactions (Beaulieu, 2004). The best-fit parameters for the model may need to be adjusted based on cultural context, or willingness to be near others. However, we do expect that the general visual-based model will continue to be generalizable, as the visual guidance of locomotion is a species-specific adaptation.

Another limitation is the ecological validity of the scenarios posed by the experiments, especially Experiment 2. The intention of the asking a participant to walk through a uniformly moving crowd was to force a path through the crowd, which would necessitate multiple deviations from center. The configuration and uniformity of motion of the crowd may, however, have had an effect that was the opposite of what we intended; Participants may have been constrained by the crowd's density, limiting their ability to respond to any potential collisions. Further, the problem of independently pushing oneself into a large crowd may not be congruent with social constraints and a heightened risk of potential collision in the real world.

5. References

Alvarez, G. A., & Franconeri, S. L. (2007). How many objects can you track?: Evidence for a resource-limited attentive tracking mechanism. *Journal of vision*, *7*(13), 14-14.

Bai, J. (2022). *Visual Control Laws for Collision Avoidance* [Doctoral Dissertation].

Basili, P., Sağlam, M., Kruse, T., Huber, M., Kirsch, A., & Glasauer, S. (2013). Strategies of locomotor collision avoidance. *Gait & posture*, *37*(3), 385-390.

Beaulieu, C. (2004). Intercultural study of personal space: A case study. *Journal of applied social psychology*, *34*(4), 794-805.

Bruneau, J., Olivier, A. H., & Pettre, J. (2015). Going through, going around: A study on individual avoidance of groups. *IEEE transactions on visualization and computer graphics*, *21*(4), 520-528.

Cai, C., Yang, C., Zhu, Q., & Liang, Y. (2007, August). Collision avoidance in multi-robot systems. In *2007 International Conference on Mechatronics and Automation* (pp. 2795-2800). IEEE.

Canny, J. (1988). *The complexity of robot motion planning*. MIT press.

Fajen, B. R., & Warren, W. H. (2003). Behavioral dynamics of steering, obstable avoidance, and route selection. *Journal of Experimental Psychology: Human Perception and Performance*, *29*(2), 343.

Gibson, J. J. (1958). Visually controlled locomotion and visual orientation in animals. *British journal of psychology*, *49*(3), 182-194.

Hicheur, H., Pham, Q. C., Arechavaleta, G., Laumond, J. P., & Berthoz, A. (2007). The formation of trajectories during goal-oriented locomotion in humans. I. A stereotyped behavior. *European Journal of Neuroscience*, *26*(8), 2376-2390.

Hopkins, B., Churchill, A., Vogt, S., & Rönnqvist, L. (2004). Braking reaching movements: A test of the constant tau-dot strategy under different viewing conditions. *Journal of motor behavior*, *36*(1), 3-12.

Koh, W. L., & Zhou, S. (2007, October). An extensible collision avoidance model for realistic self-driven autonomous agents. In *11th IEEE International Symposium on Distributed Simulation and Real-Time Applications (DS-RT'07)* (pp. 7-14). IEEE.

Latombe, J. C. (2012). *Robot motion planning* (Vol. 124). Springer Science & Business Media.

Meerhoff, L. A., Pettré, J., Lynch, S. D., Crétual, A., & Olivier, A. H. (2018). Collision avoidance with multiple walkers: Sequential or simultaneous interactions?. *Frontiers in psychology*, *9*, 2354.

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.

Naka, Y. (1977). Mechanism of cross passenger frow–study on complicated passenger flow in railway station (part 1). *Report Collection of Journal of Architectural Institute of Japan*, *258*, 93-102.

Ondřej, J., Pettré, J., Olivier, A. H., & Donikian, S. (2010). A synthetic-vision based steering approach for crowd simulation. *ACM Transactions on Graphics (TOG)*, *29*(4), 1-9.

Paris, S., Pettré, J., & Donikian, S. (2007, September). Pedestrian reactive navigation for crowd simulation: a predictive approach. In *Computer Graphics Forum* (Vol. 26, No. 3, pp. 665-674). Oxford, UK: Blackwell Publishing Ltd.

Pham, Q. C., & Hicheur, H. (2009). On the open-loop and feedback processes that underlie the formation of trajectories during visual and nonvisual locomotion in humans. *Journal of neurophysiology*, *102*(5), 2800-2815.

Pham, Q. C., Hicheur, H., Arechavaleta, G., Laumond, J. P., & Berthoz, A. (2007). The formation of trajectories during goal-oriented locomotion in humans. II. A maximum smoothness model. *European Journal of Neuroscience*, *26*(8), 2391-2403.

Pylyshyn, Z. W., & Storm, R. W. (1988). Tracking multiple independent targets: Evidence for a parallel tracking mechanism. *Spatial vision*, *3*(3), 179-197.

Renault, O., Thalmann, N. M., & Thalmann, D. (1990). A vision-based approach to behavioural animation. *The journal of Visualization and computer animation*, *1*(1), 18-21.

Rieser, J. J., Pick, H. L., Ashmead, D. H., & Garing, A. E. (1995). Calibration of human locomotion and models of perceptual-motor organization. *Journal of Experimental Psychology: Human Perception and Performance*, *21*(3), 480.

Schrater, P. R., Knill, D. C., & Simoncelli, E. P. (2000). Mechanisms of visual motion detection. *Nature neuroscience*, *3*(1), 64-68.

Todorov E. Optimality principles in sensorimotor control. *Nature Neuroscience* 2004;7:907–15.

Warren Jr, W. H. (1998). Visually controlled locomotion: 40 years later. *Ecological Psychology*, *10*(3-4), 177-219.

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.

Wilkie, R. M., & Wann, J. P. (2006). Judgments of path, not heading, guide locomotion. *Journal of Experimental Psychology: Human Perception and Performance*, *32*(1), 88.

Zanlungo, F. (2007). A collision-avoiding mechanism based on a theory of mind. *Advances in complex systems*, *10*(supp02), 363-371.

Zanlungo, F., Ikeda, T., & Kanda, T. (2011). Social force model with explicit collision prediction. *EPL (Europhysics Letters)*, *93*(6), 68005.

Zhou, C., Miao, M. C., Hu, Y. F., & Kuai, S. G. (2019). Invisible social space alters human walking behaviours. *Journal of Vision*, *19*(10), 179b-179b.