

Pedestrian Movement Simulation through Augmented Potential Field Model

Jesse Leijdekker

Master of Science Thesis

Pedestrian Movement Simulation through Augmented Potential Field Model

Jesse Leijdekker

September 16, 2025

MASTER OF SCIENCE THESIS

For the degree of Master of Science in Systems and Control &
Mechanical Engineering at Delft University of Technology

Master of Science

in **Systems and Control** and **Mechanical Engineering** at the
departments of **Delft Center for Systems and Control** and
Maritime and Transport Technology

Student Number	4707745
Supervisors	Bilge Atasoy Ton van den Boom
ME MSc Track	MME
Report Number	2025.MME.9112

Faculty of Mechanical Engineering (ME) · Delft University of Technology

Summary

This thesis examines the use of a potential field model for simulating pedestrian dynamics in complex environments. The study first reviewed the different types of pedestrian dynamics models, highlighting their strengths and weaknesses. From this study, a research gap emerged regarding hybrid pedestrian movement models. As a base for such a model, a microscopic pedestrian dynamics model has been developed combining potential fields and gradient descent optimization as the drivers for trajectory selection. In this formulation, agents follow trajectories along the gradient of the potential field, naturally balancing goal seeking behavior with obstacle and inter-agent avoidance. The potential field approach was also discussed as a foundation for hybrid models, in which microscopic and macroscopic modeling strategies are combined to exploit the advantages of both.

To evaluate and calibrate the model, real-world trajectory data from a bidirectional corridor experiment and a bottleneck experiment were used. A surrogate model was constructed to accelerate the optimization process, given the high computational cost of the original simulation model. The surrogate model enabled systematic parameter calibration and sensitivity analysis, focusing on three key parameters: the goal potential function weight (KG), the wall potential function weight (KW), and the obstacle potential function weight (KO). The results demonstrated that the optimized model is capable of reproducing key crowd phenomena observed in the empirical datasets.

A sensitivity analysis further showed the relative importance of the potential function weights across different key performance indicators. Moreover, predictive uncertainty analysis confirmed that the model exhibited relatively high confidence around the optimum and avoided regions of overfitting.

Despite these contributions, the research was constrained by the computational cost of the simulation model. The reliance on a surrogate model limited the optimization to a small subset of parameters, assuming that other model parameters were already sufficiently calibrated. This assumption likely introduced some biases, such as underestimation of obstacle repulsion, leading to overly frequent close inter-agent encounters. In conclusion, this thesis has demonstrated that potential field models, when combined with real-world data and surrogate-based optimization, provide a valid and powerful framework for simulating pedestrian dynamics

in complex environments. Their ability to model pedestrian trajectories through potential functions and gradient descent makes them conceptually simple yet effective, while their extensibility offers a pathway toward hybrid models. Nevertheless, computational burden and limited parameter coverage remain key challenges, highlighting the need for more efficient implementations and broader parameter optimization in future research.

Table of Contents

1	Introduction	1
2	Previous Work	5
2-1	Literature Study	5
2-1-1	Macroscopic Models	5
2-1-2	Microscopic Models	6
2-1-3	Hybrid Models	6
2-1-4	Intermediate Conclusion	7
2-2	Proposed Model	7
2-2-1	Base Model	7
3	Methodology	9
3-1	Model Improvements	9
3-1-1	Computational burden	9
3-1-2	Pedestrian behavior	10
3-2	Model Optimization	11
3-2-1	Optimization parameters	12
3-2-2	Optimization strategy	13
3-2-3	Algorithm and search space	16
3-2-4	Simulated training data	17
3-2-5	Key performance indicators	18
3-2-6	Hardware setup and implementation	19
4	Optimization Results	21
5	Surrogate Model Validation	25
5-1	Coefficient of determination analysis	25
5-2	Local Error Analysis	27
5-3	Predictive Uncertainty Analysis	29

6	Simulation Model Testing and Validation	31
6-1	Testing	31
6-2	Validation	34
6-2-1	Bidirectional Corridor Experiment	34
6-2-2	Bottleneck experiment	37
6-3	Intermediate Conclusion	40
7	Sensitivity Analysis	41
7-1	Parameter Sensitivity	41
7-2	Local Sensitivity	43
8	Experiments and Results	45
8-1	Experiment 1: Normal Conditions	45
8-2	Experiment 2: Delayed Opening	48
8-3	Experiment 3: Busy Conditions	50
8-4	Computation Time	51
9	Conclusion and Discussion	53
9-1	Research Findings	53
9-2	Limitations and Recommendations	55
A	Overview of Uncertainty Analysis	57
B	Overview of Parameter Sensitivity Analysis	61
C	Scientific Paper	65

Chapter 1

Introduction

Understanding pedestrian dynamics is essential for the effective planning, design, and operation of public spaces. As urban environments grow more complex and densely populated, the need to model and simulate pedestrian movement has become increasingly important across a range of domains including transportation engineering, crowd safety, emergency evacuation planning, and smart infrastructure development. Accurate modeling of pedestrian behavior supports the optimization of circulation spaces, enhances user safety and comfort, and enables data-driven decisions in both public and private infrastructure projects.

This research focuses on pedestrian dynamics within a university environment, and more specifically, within the Mechanical Engineering faculty building at Delft University of Technology. The university is currently undergoing a major expansion, with plans to increase its student population from approximately 25,000 to 40,000 [3]. To accommodate this growth, new campus locations are being considered in The Hague and Rotterdam, alongside modifications to existing infrastructure in Delft. As student numbers rise, gaining a deeper understanding of pedestrian dynamics within educational environments becomes increasingly important. In particular, modeling pedestrian movement through corridors, lecture halls, and communal spaces can inform the design of more efficient circulation routes and help prevent congestion during peak transition periods. These insights are essential for designing future campus infrastructure and improving the overall user experience.

However, modeling pedestrian behavior is inherently complex. Despite the development of numerous pedestrian dynamics models [9, 19, 33, 16], accurately capturing the dynamics of human movement remains a challenge due to its context-dependent, adaptive, and often non-linear nature. Pedestrians continuously adapt to both static environmental features and dynamic interactions with other individuals. These interactions often result in emergent collective behaviors such as lane formation, bottlenecks, or stop-and-go waves that are difficult to reproduce using simple rules or flow-based models.

To address these complexities, this study introduces a microscopic pedestrian dynamics model based on agent-based potential fields. In this framework, each pedestrian is modeled as an individual agent navigating a scalar potential landscape that encodes goals, obstacles, and

interactions with other pedestrians. The agents determine their paths using a gradient descent optimization approach, allowing them to adaptively choose directions leading towards their goals while avoiding collisions. This approach draws on principles from classical optimization and control theory and offers a flexible, computationally tractable method for simulating realistic pedestrian behavior in complex indoor environments.

The aim of this research is therefore captured by the following primary research question:

To what extent can pedestrian movement in complex indoor environments be effectively modeled using a potential field model and calibrated using real-life trajectory data?

To address this main question, the following subquestions are formulated:

1. *What is the state-of-the-art for modeling pedestrian dynamics within a multi-agent system?*
2. *What makes a potential field model suitable for simulating pedestrian dynamics in complex environments?*
3. *How can a potential field model for pedestrian simulation be augmented and optimized using real-life trajectory data?*
4. *What are the strengths and limitations of using a potential field model for pedestrian simulation?*

An overview of the research is depicted in Figure 1-1. To answer subquestion one, regarding the state-of-the-art for modeling pedestrian dynamics within a multi-agent system, a literature study has been conducted as described in Section 2-1. From this literature study, a research gap has been identified and a certain method of modeling pedestrian dynamics has been elected. This method has been explained further in Section 2-2, answering subquestion two. Thereby, the first two subquestions from the basis of this study and scope the research.

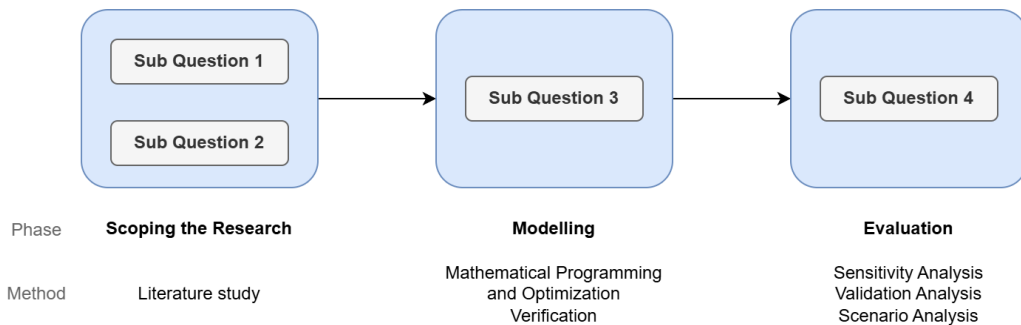


Figure 1-1: Overview of the research.

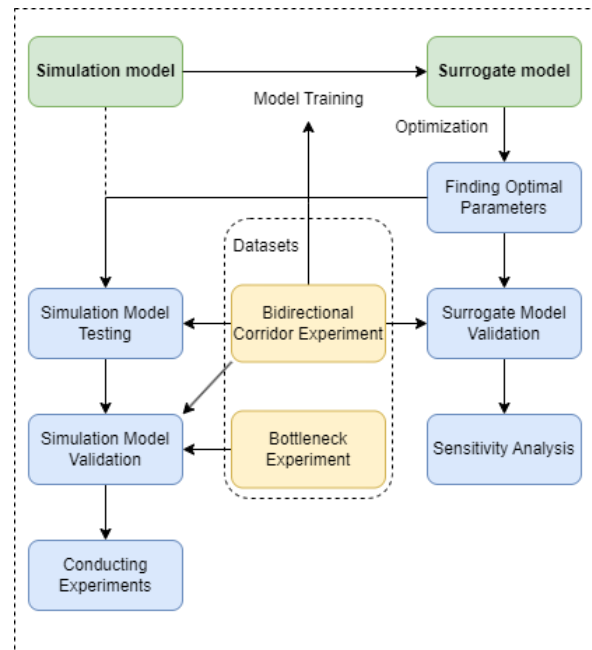


Figure 1-2: Overview of the research structure.

The third subquestion addresses how the developed potential field simulation model can be improved and efficiently calibrated. The computational expense of the full-scale, agent-based potential field model motivates the construction of a surrogate model that emulates the simulation’s KPI outputs, as can be seen in Figure 1-2, in which the structure of the research is depicted. The development of the optimization framework that leverages this surrogate is detailed in Chapter 3 and Chapter 4.

The surrogate is trained on KPI data generated by the simulation model for the bidirectional corridor experiment, and its predictive accuracy is quantitatively assessed and discussed in Chapter 5. In that chapter the surrogate predictions are also compared with outputs from the full-scale model at the parameter sets of interest for the bidirectional corridor experiment to establish whether the surrogate can be trusted.

Following surrogate development and validation, the optimized parameter sets and the simulation model itself are tested and validated against real-life experiments. In Chapter 6 it is described how the optimized potential field model is compared to the bidirectional corridor and bottleneck datasets both in terms of the calibration KPIs and for additional emergent phenomena not explicitly used during optimization. These tests quantify the model’s ability to reproduce observed pedestrian behavior and identify remaining discrepancies.

To probe robustness and parameter influence, a sensitivity analysis is performed using the surrogate model to limit computational cost; the methodology and results of the local and parameter sensitivity analyses are reported in Chapter 7. By exploiting the surrogate’s efficiency, these analyses explore uncertainty and parameter interactions that would be infeasible to cover exhaustively with the full-scale simulation model.

Finally, the validated and optimized potential field model is applied to a corridor scenario with lecture halls to demonstrate applicability in a more complex setting; results and interpretation

of these application experiments are presented in Chapter 8.

The complete workflow of how the simulation model, surrogate model, optimization, validation, sensitivity analysis, and final applications interrelate is visualized in Figure 1-2, which serves as a roadmap for the thesis.

Chapter 2

Previous Work

In this chapter, an overview of previously developed simulation models for pedestrian dynamics is given. In Section 2-1 models from literature are highlighted and in Section 2-2 an earlier developed potential field model is described.

2-1 Literature Study

This section provides a concise overview of previously developed models for pedestrian dynamics. In general, three categories of models are employed to simulate pedestrian behavior: macroscopic models, microscopic models, and hybrid models. For a more in depth overview of these types of models, the reader is referred to [25].

2-1-1 Macroscopic Models

Macroscopic models conceptualize pedestrian dynamics using continuous flows. Some of the popular examples of macroscopic models are Network Models and Continuum Models [7].

In Network Models, Pedestrians are represented as agents navigating a specified network consisting of nodes and links corresponding to physical spaces such as corridors, streets, and intersections. Pedestrian movement is described as flow through the links, and route choices are made at the nodes. These decisions are typically based on factors such as levels of congestion, shortest path, or individual route preference [9]. One of the main shortcomings of this approach is that decision is restricted to these single nodes and potentially may not reflect pedestrian movement between decision points most accurately. Although this may be improved by introducing additional decision nodes, this would lead to exponential growth in possible routing options, significantly increasing computational cost. Conversely, when fast approximate simulations suffice, the reduction of resolution by utilizing a macroscopic model is beneficial since it reduces computational complexity.

Continuum Models depict pedestrian crowds as a continuous density field, analogous to fluid or traffic flow models. Pedestrian motion is typically guided to accomplish tasks by a global

potential field in Continuum Models [20, 21, 35]. Like other macroscopic models, Continuum Models excel at simulating large-scale crowd behavior [20, 21, 35]. However, they lack the capacity to model decision-making and anticipation at the individual level because pedestrians react only to local and instantaneous information.

2-1-2 Microscopic Models

Microscopic models represent individual pedestrians as autonomous agents in order to represent intricate interaction and behavior. Widely-used microscopic modeling methods include Cellular Automata, Social Force Models, and Obstacle Avoidance Models [7].

Cellular Automata (CA) are discrete mathematical models consisting of a grid of cells, with each cell being able to assume one of a finite possible number of states [13, 22]. The condition of each cell evolves at discrete time steps according to uniform, predefined rules, applied simultaneously over the grid. Pedestrian movement decisions are generally regulated by transition probabilities dependent on the values of a floor field [24]. Floor fields are spatial overlays that guide pedestrian movement, and they function in a similar manner to potential fields in Continuum Models. CA models have demonstrated the ability to replicate various emergent behaviors, such as jamming, lane formation, and oscillatory patterns [5, 24]. Moreover, they provide insight into individual behavior as an agent-based framework is used for these models. Nonetheless, incorporating agent-specific destinations is challenging, as floor fields generally affect all pedestrians uniformly.

The Social Force Model describes pedestrian motion as a result of generalized forces. Pedestrians are driven to achieve a desired speed in the direction of their destination, and they respond to environmental hindrances and interpersonal distance through repulsive forces [16]. It is a suitable model for modeling individual pedestrian behavior [16, 18]. However, it assumes homogeneous response from all agents and lacks explicit route planning because the movement is a result of force dynamics and not goal-oriented decision-making.

In Obstacle Avoidance Models, pedestrians select velocities to minimize the likelihood of collisions based on real-time perception of the environment [8]. These models lack predictive power, as decisions by pedestrians are made based on an instantaneous information input based on a limited portion of the environment only. Furthermore, although Obstacle Avoidance Models capture individual-level dynamics, their computational cost can be substantial due to the agent-based representation.

2-1-3 Hybrid Models

As stated, macroscopic models offer high computational efficiency at the cost of behavioral precision, and microscopic models offer high precision at the expense of increased computational effort. Hybrid models aim to merge the benefits of both paradigms, compromising simulation speed and behavioral realism.

An example is provided in [37], where a multi-resolution approach is proposed that employs a macroscopic density-based model under stable conditions and switches to a microscopic model under dynamic or unstable conditions. Another example is provided in [38], where the space is split spatially into areas, and each area is simulated using a microscopic or macroscopic

model depending on the suitability. In [2], the pedestrian dynamics are presented macroscopically through speed-density relations, and the route choice for pedestrians is addressed microscopically.

2-1-4 Intermediate Conclusion

To conclude, pedestrian simulation models can be categorized as macroscopic, microscopic, and hybrid. Macroscopic models are computationally efficient but sacrifice individual behavior realism as pedestrians are modeled as continuous flows. Microscopic models achieve higher fidelity by simulating individual agents at the cost of increased computational demands. Hybrid models attempt to achieve a middle ground between these two extremes by exploiting the strengths of both. However, most hybrid models are tailored to specific applications, or they only partially combine the merits of the base approaches. This reveals a research gap in the development of general-purpose hybrid pedestrian simulation models.

2-2 Proposed Model

To bridge the research gap addressed in Section 2-1, this research attempts to lay the groundwork of a generalized hybrid pedestrian dynamics model. This is done by developing and optimizing a microscopic model which is easily adaptable to a hybrid approach and can therefore serve as its base. Specifically, this study offers a potential field-based approach where pedestrians are modeled as individual agents to enhance behavioral realism. The adaptability towards hybrid modeling will be covered more broadly in Chapter 9.

An earlier study has laid the foundation of the proposed model, focusing exclusively on the microscopic component [26]. In this section, the previously developed model is explained briefly. The present research aims to enhance this model such that it is prepared for integration into the hybrid modeling framework outlined earlier. However, the objective of the present research is not to further extend the model towards such a hybrid model.

2-2-1 Base Model

As stated, this study concerns a microscopic potential field-based approach for pedestrian simulation. The agents are modeled as a single point in space and find their goals by optimizing a potential function using gradient descent optimization. The agent specific goals are defined as the global minimum of the potential function. For every time step, the agents take a step in the direction of the computed negative gradient. Furthermore, to prevent agents from getting stuck in the environment, intermediate goals are defined, aiding the agents along their trajectories.

The potential function consists of three separate potential functions: the obstacle potential function VO , the wall potential function VW and the goal potential function VG . The obstacle potential function is designed to prevent collisions with both agents and dynamic or static obstacles. This is achieved by defining a potential function that exhibits a maximum at the location of the obstacle. The wall potential function operates in a similar manner, but employs a more elongated peak shape for each wall of the environment to represent the

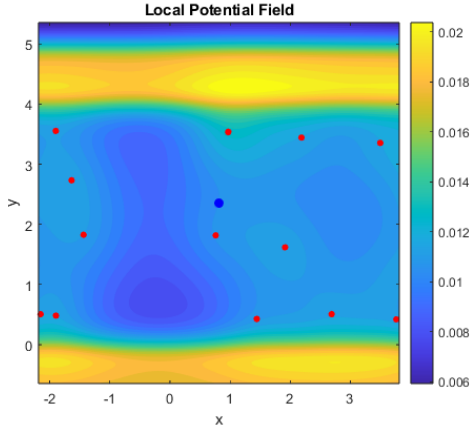


Figure 2-1: Example of a local potential field. The agents for which the potential field is drawn is depicted in blue and other agents in the environment are depicted in red.

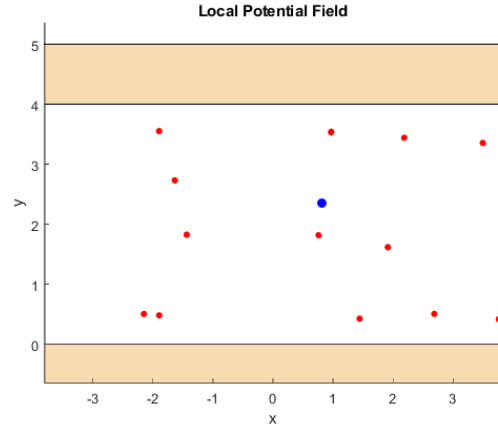


Figure 2-2: The corresponding environment for the local potential field. The walls in the environment are depicted in orange.

repulsive influence of the boundaries. In this way, agents are guided away from walls while still maintaining realistic movement patterns. The summation of all these individual potentials corresponding to a certain wall results in the complete potential function. The goal potential function is used to attract the agents towards their goal location. For a more extensive description of the individual potential functions, the reader is referred to [26]. An example of an agent based potential function for a certain time step and the corresponding environment is depicted in Figure 2-1 with its corresponding environment depicted in Figure 2-2.

The entire potential function is defined as follows:

$$V_{tot,i}(\mathbf{x}_{i,t}, \mathbf{x}_{j,t}, t) = \sum_{j=1}^{N(t)} VO_j(\mathbf{x}_{j,t}) + VW(\mathbf{x}_{i,t}) + VG_i(\mathbf{x}_{i,t}, t) \quad (2-1)$$

for $j \neq i$ with $\mathbf{x}_{i,t} = [x_{i,t} \ y_{i,t}]^T \in \mathbb{R}^2$ being the coordinate vector for agent i at time t and $N(t)$ being the number of agents present at time t . The movement of the agent then results from the following gradient descent algorithm:

$$\mathbf{x}_{i,t+1} = \mathbf{x}_{i,t} - \alpha \nabla V_{tot,i}(\mathbf{x}_{i,t}, \mathbf{x}_{j,t}, t) \quad (2-2)$$

Here, α is a parameter to control the step size and $\nabla V_{tot,i}(\mathbf{x}_{i,t})$ represents the gradient of the potential function at the agent's current position.

Chapter 3

Methodology

Given that this study focuses on researching to what extent a potential field model for pedestrian movement simulation can yield realistic results, it is crucial to find what areas the previously developed model can be improved on. This chapter explains what improvements have been applied to the model and how it is optimized such that improved realism can be achieved.

3-1 Model Improvements

In this section, the improvements applied to the earlier developed model from [26] are explained.

3-1-1 Computational burden

A significant performance enhancement to the simulation model was achieved through the parameterization of symbolic expressions. Initially, symbolic variables representing the potential fields were redefined and differentiated at every simulation timestep, leading to substantial computational inefficiency due to repeated symbolic processing and function generation. To resolve this, all static symbolic expressions such as wall potentials and their gradients were computed once outside the main simulation loop using MATLAB's Symbolic Math Toolbox and converted into anonymous function handles via `matlabFunction`. This eliminated redundant symbolic differentiation and allowed for fast numerical evaluation during runtime. For time-varying components, such as agent-specific obstacle potential functions, the symbolic form was parameterized with agent-dependent coefficients (e.g., position and gain parameters), allowing the simulation to reuse a generic symbolic template and update only the numerical values. This restructuring reduced simulation runtime by an order of magnitude in larger scenarios and enabled efficient scaling with respect to both agent count and simulation duration, without compromising the accuracy of the gradient-based movement dynamics.

The performance was improved further by simplifying the wall potential function. As stated in [26], the wall potential function used to be defined as:

$$VW_k = \left(\frac{1}{1 + e^{-c_1 f_1(x,y)}} + \frac{1}{1 + e^{-c_2 f_2(x,y)}} - 1 \right) \left(\frac{1}{1 + e^{-c_3 f_3(x,y)}} + \frac{1}{1 + e^{-c_4 f_4(x,y)}} - 1 \right) \quad (3-1)$$

Here, VW_k is the wall potential for wall k , c_i is the gain controlling the steepness of the wall for the appurtenant side, $f_1(x, y)$ and $f_2(x, y)$ are the functions defining the location of the borders of the wall and $f_3(x, y)$ and $f_4(x, y)$ are the functions defining the location of the borders of the wall, generally perpendicular to $f_1(x, y)$ and $f_2(x, y)$. In the previous version of the model, $f_1(x, y), \dots, f_4(x, y)$ used to be sigmoid functions. However, building an intricate environment using these sigmoid functions results in a complex wall function with a lot of terms. This increases the computational burden as the derivative of this wall function has to be computed for the gradient descent algorithm. Therefore, in the new model, for straight walls, the wall potential function is defined as:

$$VW_k = \left(\frac{1}{1 + e^{-c_1 f_1(x,y)}} + \frac{1}{1 + e^{-c_2 f_2(x,y)}} - 1 \right) \left(\frac{1}{1 + c_3 f_3(x, y)} \right) \quad (3-2)$$

Here, $f_3(x, y)$ generally is a linear function with very few terms. This significantly reduces the complexity of the derivative operation because of which the computational speed of the model is improved.

3-1-2 Pedestrian behavior

In pedestrian dynamics, the speed-density relationship is well established. It has been shown that free-flowing walking speeds of pedestrians decrease significantly as the density of people walking around them increases [29]. This was not implemented in the base model, because of which simulating pedestrian behavior in crowded areas was difficult. Furthermore, studies have evinced the fact that interpersonal distance shrinks in crowded spaces [15]. This means that pedestrians keep less distance to one another at higher densities. Again, this had not been implemented in the original model. Therefore, two adaptations have been made to mimic these phenomena.

First of all, the speed-density relationship has been altered to establish decreasing walking speeds in high density environments. A commonly used formulation for modeling the relationship between pedestrian speed and local density is the exponentially decreasing function, such as $v(\rho) = v_0 e^{-\alpha \rho}$, where $v(\rho)$ is the pedestrian speed at density ρ , v_0 is the desired (free-flow) speed, and α is a sensitivity parameter. This formulation captures the empirically observed behavior that pedestrian speed decreases non-linearly as density increases. The exponential function e^{-x} is strictly decreasing, continuous, and smooth, making it suitable for gradient-based modeling approaches. Its asymptotic behavior ensures that speed never becomes negative, and its curvature naturally models the diminishing marginal effect of density on speed at higher congestion levels. The speed-density relationship in this study is based on the above formulation and was determined using an iterative process:

$$v_i = \alpha_i e^{-0.015c_i} \quad (3-3)$$

Here, c_i is a measure of density, namely the number of agents within a certain vicinity of agent i . The distance threshold establishing this vicinity is set to three meters. Thus, c_i is defined as the number of agents within a circle with a three meter radius and with the centerpoint being agent i . Furthermore, $\alpha_i \sim \mathcal{N}(\mu_\alpha, \sigma_\alpha)$ is the agent-based desired free-flow speed scaling factor drawn from a Gaussian distribution with mean $\mu_\alpha = 40$ and standard deviation $\sigma_\alpha = 10$.

Such exponentially decreasing speed-density relationships have been used and validated in various studies. A notable example is Weidmann's fundamental diagram [36], which, although originally presented as a tabulated relationship, is often approximated by exponential forms in simulation models. Similarly, Helbing's Social Force Model [17] includes exponential functions to describe interpersonal repulsive forces and has inspired exponential decay models in velocity-density contexts as well.

The use of an exponential speed-density model thus offers a balance between empirical realism and mathematical tractability, making it a suitable choice for simulating and optimizing pedestrian behavior in crowd dynamics scenarios in regards to real-life data.

Also, the obstacle potential function has been augmented to effectuate the decrease of interpersonal distances in high density environments, thereby improving the realism of the model under crowded conditions. As before, c_i denotes the number of agents in the vicinity of agent i and is utilized here to account for local crowding effects. The obstacle potential function is now defined as follows:

$$VO_{i,j} = \frac{KO}{0.05(1 + 0.1(5 + c_j))((x_j - x_i)^2 + (y_j - y_i)^2) + 0.25} \quad (3-4)$$

This formulation reduces the spatial reach of the repulsive influence from obstacles as local crowd density increases. In effect, the potential peak becomes narrower in high-density scenarios, as illustrated in Figure 3-1 and Figure 3-2. As a result, the agents are allowed to navigate closer to obstacles when surrounded by many other pedestrians, reflecting more natural behavior in congested environments.

3-2 Model Optimization

In this section, the optimization process of the previously developed potential field model, including all earlier explained adaptations and augmentations, is unfolded. The optimization is tailored towards improving the realism of the simulation put forward by the potential field model. The optimization process is mapped by explaining the optimization parameters and the optimization method, and describing the real-life data, the KPIs, and the hardware setup and implementation.




Figure 3-1: Obstacle potential function for low density ($c_j = 1$)




Figure 3-2: Obstacle potential function for low density ($c_j = 50$)

3-2-1 Optimization parameters

As a first step of the optimization process, the optimization parameters need to be exploited. Possible optimization parameters include:

- Function weights K_w , K_o and K_g .
- Time step size δt
- Speed distribution for the pedestrians
- Type of potential function for a static obstacle
- Type of potential function for a dynamic obstacle

The weights assigned to the different potential field components such as static obstacles, dynamic obstacles, and goal attraction play a crucial role in shaping pedestrian behavior in the simulation. These weights influence how strongly a pedestrian responds to environmental cues like walls, other pedestrians, or destinations. Calibrating these weights is essential to ensure that the simulated movement patterns align with real-life pedestrian behavior.

The time step size determines how frequently the pedestrian positions are updated during the simulation. If the time step is too large, pedestrians may make unrealistically large jumps, potentially skipping over obstacles or interacting unnaturally with the environment. If the time step is too small, the simulation becomes computationally expensive without significant gains in accuracy. Optimizing the time step size ensures a balance between numerical stability, realism, and computational efficiency.

The speed distribution reflects the natural variability in how fast different pedestrians move. In real-life scenarios, pedestrian speeds are not uniform. Factors like age, intent, or congestion

affect walking speeds. By tuning the speed distribution used in the simulation, the model can better reproduce realistic crowd dynamics, such as lane formation in bidirectional flows or slower movement near bottlenecks. Accurate modeling of speed variability is particularly important when comparing temporal KPIs like time-to-target or dwell times.

Static obstacles like walls and furniture shape how pedestrians navigate the environment. The mathematical form of the potential function used to represent these obstacles affects how sharply or gradually pedestrians respond to their presence. For example, an exponential decay may lead to smoother avoidance behavior, while strong repulsion may result in abrupt changes in trajectory. Optimizing this function type helps capture realistic pedestrian-wall interaction patterns, such as the tendency to avoid walking too close to walls in open spaces.

Dynamic obstacles, mainly other pedestrians, require special treatment to simulate crowd behavior effectively. The potential function used for these interactions determines how pedestrians anticipate and avoid collisions with others. Choosing an appropriate function type (e.g., isotropic vs. velocity-based repulsion) is crucial for reproducing emergent behaviors such as self-organized lanes or synchronized movement patterns in crowded environments. This parameter directly affects the realism of collective behavior in the simulation.

Although each of the aforementioned parameters plays a meaningful role in pedestrian simulation, optimizing all of them simultaneously would be computationally infeasible due to the long simulation runtime and the high-dimensional search space. For this reason, the optimization in this project was limited to adjusting the weights of the potential field components. These weights have a strong and direct influence on pedestrian behavior and could be calibrated using surrogate modeling techniques to reduce computational cost. By focusing on the function weights, the study balances realism with tractability while still capturing the core dynamics necessary for validation against real-life data.

3-2-2 Optimization strategy

After determining the optimization parameters, the optimization strategy has been elected. Several optimization strategies were considered to calibrate the parameters of the pedestrian simulation model. Each method differs in terms of how it interacts with the simulation and balances computational cost, flexibility, and effectiveness. The optimization methods scrutinized for this project are brute force optimization, curve fitting and surrogate modeling.

One of the most direct approaches is brute force optimization, sometimes referred to as black-box or direct optimization. In this method, the pedestrian simulation itself is treated as an opaque function: the optimizer does not need to know how the simulation works internally, only what the input parameters are and what output it produces. The optimization algorithm proposes new parameter sets in each iteration. For each proposal, the entire simulation is run, and a set of performance indicators are computed. The optimizer then measures how far these indicators deviate from real-life data and uses this error to guide the next iteration. While conceptually simple and requiring minimal setup, this method is highly computationally expensive. Each function evaluation may require minutes of simulation time, and the optimization often requires a significantly large number of evaluations. Furthermore, if the simulation includes randomness, the optimization landscape becomes noisy and potentially misleading for gradient-based algorithms. This method is only viable if the number of parameters is very limited and if simulation time is not a constraint.

Another method is curve fitting, which attempts to represent the relationship between input parameters and output indicators with a known mathematical form, such as a polynomial or exponential function. By running the simulation for various parameter combinations and recording the outcomes, a regression model is fitted to the data. This fitted function can then be used to predict outcomes for new parameter values. The advantage of curve fitting is that it is fast and simple to implement. However, its performance heavily depends on the choice of functional form. If the true relationship between parameters and performance indicators is highly nonlinear or exhibits complex interactions, a simple curve-fitting model may fail to capture these dynamics. This can result in significant prediction errors, especially outside the range of the training data. Additionally, curve fitting typically does not offer any way to estimate the uncertainty of its predictions, making it difficult to judge the reliability of the results.

To overcome the limitations of both brute force optimization and curve fitting, surrogate modeling can be used as a more advanced and efficient alternative. A surrogate model is a data-driven approximation of the simulation that can predict outputs based on input parameters with far less computational cost than the actual model. The process begins by selecting a diverse set of input parameters using a design of experiments method. The simulation is then run once for each of these parameter sets to build a training dataset. A regression model is then trained on this dataset to learn the mapping from inputs to outputs. Once the surrogate model is trained, optimization can be performed quickly using standard algorithms without running the full simulation. In addition, areas of high uncertainty can be identified and used to guide additional simulations to iteratively improve the surrogate's accuracy in critical regions. This approach greatly reduces the number of full simulation runs needed, making it highly suitable for expensive or noisy simulations.

In the context of this pedestrian simulation project, surrogate modeling offers the best balance between computational efficiency and predictive power. The original simulation is complex, stochastic, and time-consuming to run, making brute force optimization impractical. Meanwhile, as there is no apriori knowledge regarding the type of curve describing the relationships between parameters and performance indicators, simple curve fitting cannot capture these relationships reliably. Surrogate modeling enables the use of powerful optimization techniques while keeping simulation cost under control and allowing for uncertainty-aware decision making. As such, it forms a core component of the methodology used in this thesis.

To construct an efficient and accurate surrogate model for the pedestrian simulation, several types of regression models and sampling strategies were considered. Regression models serve as the core of surrogate modeling by learning the relationship between input parameters and simulation outputs based on a limited set of training data. Common types include linear regression, polynomial regression, radial basis function (RBF) models, neural networks, and Gaussian Process Regression (GPR). Linear and polynomial regressions are simple and interpretable but often struggle to capture nonlinearities present in complex simulation outputs. RBF models are more flexible and can adapt to local variations, yet they lack built-in uncertainty quantification [4]. Neural networks are powerful for modeling highly nonlinear relationships but require a large amount of data and do not inherently provide uncertainty estimates [14]. In contrast, Gaussian Process Regression models are particularly well-suited for surrogate modeling, especially when training data is limited and uncertainty quantification is important. GPR models not only provide predictions but also quantify the confidence of each prediction, which is valuable in guiding optimization and understanding model reliability [30].

Mathematically, Gaussian Process Regression assumes that the function to be learned is a sample from a multivariate Gaussian process defined by a mean function $m(x)$ and a covariance (kernel) function $k(x, x')$. Given training inputs $X = \{x_1, x_2, \dots, x_n\}$ and corresponding outputs $y = \{y_1, y_2, \dots, y_n\}$, the GPR model predicts the output y_* at a new input x_* by computing the conditional distribution of the joint Gaussian distribution of training and test outputs. The prediction is given by a Gaussian distribution with mean

$$\mu(x_*) = k(x_*, X)[K + \sigma_n^2 I]^{-1}y \quad (3-5)$$

and variance

$$\sigma^2(x_*) = k(x_*, x_*) - k(x_*, X)[K + \sigma_n^2 I]^{-1}k(X, x_*) \quad (3-6)$$

where K is the covariance matrix constructed from the kernel function over the training inputs, and σ_n^2 represents the noise variance. The kernel function encodes assumptions about the smoothness and behavior of the underlying function, with its parameters learned from the data [34].

The GPR model has been trained on data created by the pedestrian simulation model using a number of parameter sets. The experiment which was used to create this data is explained in Section 3-2-4. For effective surrogate modeling, the choice of the sampling for these parameter sets is of large importance. A naïve approach would be to randomly sample input parameters, but this often leads to clusters and gaps in the parameter space, especially in higher dimensions. Latin Hypercube Sampling (LHS) divides the input parameter space into equally probable intervals along each dimension and ensures that each interval is sampled exactly once, resulting in a more uniform and efficient sampling distribution, especially in higher dimensions [27]. For these reasons, LHS was selected as the sampling method when developing the surrogate model.

In this project, a Gaussian Process Regression model was trained on data generated from Latin Hypercube Sampling of the input parameter space. Each sample point represented a unique configuration of the simulation parameters, and the associated outputs were the Key Performance Indicators (KPIs) computed from the pedestrian simulation. This combination of GPR and LHS enabled the construction of a surrogate model that was both computationally efficient and capable of generalizing across a broad parameter space with quantified prediction uncertainty. This surrogate model was subsequently used for optimization, allowing for rapid evaluation of model performance and efficient parameter calibration without requiring repeated full simulation runs.

A critical component of GPR is the choice of kernel function, also known as the covariance function, which defines the similarity between input points and encodes prior assumptions about the smoothness and structure of the target function. Several types of kernel functions are commonly used in GPR, each with distinct characteristics and suitability for different modeling tasks [30].

One widely used kernel is the Squared Exponential (SE) kernel, also known as the Radial Basis Function (RBF) kernel. It assumes the function being modeled is infinitely differentiable and thus very smooth. Its covariance between two input points x and x' is given by:

$$k_{\text{SE}}(x, x') = \sigma_f^2 \exp\left(-\frac{\|x - x'\|^2}{2\ell^2}\right) \quad (3-7)$$

where σ_f^2 is the signal variance and ℓ is the length-scale parameter controlling how quickly the correlation decays with distance.

Another popular family of kernels is the Matern class, which introduces an additional smoothness parameter ν . The Matern kernel provides more flexibility in modeling functions with varying degrees of differentiability. In particular, the Matern-3/2 kernel, defined as:

$$k_{\text{M32}}(x, x') = \sigma_f^2 \left(1 + \frac{\sqrt{3}\|x - x'\|}{\ell}\right) \exp\left(-\frac{\sqrt{3}\|x - x'\|}{\ell}\right) \quad (3-8)$$

assumes that the target function is once differentiable and is less smooth than the SE kernel, while it still is assumed continuous [34].

Other commonly used kernels include the Matern-5/2 kernel, which assumes the function is twice differentiable and offers a balance between flexibility and smoothness, and the Rational Quadratic kernel, which can be seen as a scale mixture of SE kernels and is useful for modeling functions with varying length scales [30].

In this study, the Matern-3/2 kernel was chosen for the GPR model. Compared to the SE kernel, the Matern-3/2 kernel avoids over-smoothing and better captures localized variations in the simulation response surface, which is particularly important given the noisy and nonlinear nature of pedestrian dynamics.

3-2-3 Algorithm and search space

To identify the optimal parameters for the simulation model, the surrogate model was optimized using the particle swarm optimization (PSO) algorithm. PSO is a population-based, stochastic optimization method inspired by the social behavior of flocks of birds and schools of fish [23]. In this approach, a set of candidate solutions (particles) explore the search space, adjusting their positions based on their own best-found solutions and the global best solution identified by the swarm. This mechanism enables a balance between exploration of the search space and exploitation of promising regions, which is particularly beneficial for non-convex, multimodal objective functions where gradient information is unavailable or unreliable.

A swarm size of 30 was chosen to provide a sufficient level of diversity in the search process while keeping the computational cost of the surrogate-based optimization manageable. This choice reflects a trade-off between the algorithm's convergence speed and the likelihood of escaping local optima.

PSO was preferred over other global optimization methods such as genetic algorithms or Bayesian optimization. Bayesian optimization is widely recognized for its efficiency in low-dimensional, smooth problems. However, its performance can degrade in high-dimensional or multimodal settings, which increases the sample complexity required to model the objective function accurately [12]. Additionally, as bayesian optimization with GPR often assumes some degree of smoothness, the GPR surrogate model can struggle to represent potential

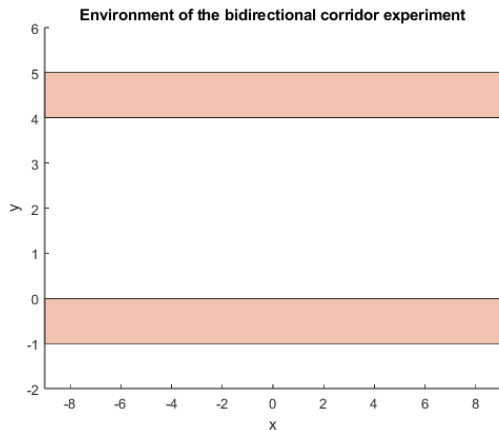


Figure 3-3: The environment of the bidirectional corridor experiment. The walls are depicted in red.



Figure 3-4: A frame from the real-life bidirectional corridor experiment. Taken from [11].

sharp discontinuities ingrained in the objective function, limiting its effectiveness in such scenarios. As the objective function is unknown in this case, it is preferred to not presuppose any smoothness assumptions.

Furthermore, genetic algorithms offer a comparable global search capability to PSO, but often require more complex parameter tuning and can converge more slowly. In contrast, PSO combines a relatively simple implementation with good empirical performance on high-dimensional, non-convex optimization problems [32], making it a suitable choice for this application. It requires relatively few design parameters and avoids gradient dependence, which makes it attractive for optimization tasks where function smoothness or convexity is not guaranteed.

Based on the initial parameter values of the original model from [26], the search space for the PSO algorithm has been determined to be as follows:

$$\begin{aligned} 0.001 &\leq KW \leq 0.15 \\ 0.001 &\leq KG \leq 0.2 \\ 0.0001 &\leq KO \leq 0.015 \end{aligned}$$

3-2-4 Simulated training data

As stated, the surrogate model is trained by a dataset created by running an experiment with the simulation model over a number of parameter sets. This experiment is based on a study consisting of an experiment conducted in a bidirectional corridor environment [11]. The environment consists of two perpendicular walls, as depicted in Figure 3-3. The agents walk into the environment either from the left or the right side to eventually reach the opposite side of the corridor. A camera was mounted at the ceiling which recognizes the pedestrian by a certain helmet that they are wearing, as can be seen in Figure 3-4. The agents are recognized by the camera between approximately $x = -5$ and $x = 5$.

3-2-5 Key performance indicators

To calibrate the pedestrian simulation model and ensure its outputs align with empirically observed human behavior, a set of twelve key performance indicators (KPIs) has been defined. These KPIs have been chosen to comprehensively capture both low-level interaction dynamics and high-level behavioral patterns, ensuring that the model replicates real-life pedestrian movement under a variety of conditions to a high degree.

The first four KPIs quantify the frequency of inter-agent proximity events within thresholds of 0.1 m, 0.3 m, 0.5 m, and 1.0 m, respectively. These metrics are crucial for modeling social spacing behavior, as humans tend to maintain a personal buffer zone to avoid discomfort or collisions. By incorporating multiple distance thresholds, the model can be tuned to capture not only hard collisions but also near misses and general crowding effects.

In addition to inter agent interactions, physical interactions with the environment are considered through the wall collision KPI, which measures the frequency with which agents collide with walls and other static obstacles. This helps penalize unrealistic behavior such as direct contact with impenetrable boundaries.

Beyond collision-related KPIs, the model's ability to reproduce realistic movement patterns is assessed via the distributional properties of agent speeds and lateral (y-axis) trajectories. Specifically, the speed distribution is summarized by fitting a beta distribution to the observed speed samples and extracting its α and β parameters. These parameters serve as compact representations of the underlying statistical behavior and allow for easy comparison between simulated and experimental data. A similar approach is taken for the lateral distribution of paths: since the experimental setup involves straight-line motion in a corridor along the x-axis, only the y-coordinates are informative for analyzing how agents deviate from the central path. Fitting a beta distribution to the y-positions across all trajectories again yields α and β parameters that serve as robust KPIs.

To further capture deviations in path smoothness and irregularities, the variance in the y-direction of each trajectory is computed. This KPI quantifies the "squigginess" of agent paths, with higher variance indicating erratic or unstable behavior, which is generally not observed in controlled pedestrian movement.

To obtain a more general and systematic assessment of the simulation model's performance relative to real-life data, a step-by-step comparison was conducted at the level of individual agents and timesteps. Specifically, at each timestep in the empirical dataset, the complete state of the system, that is, the positions of all agents and other obstacles, was provided as input to the simulation model. The model then computed the subsequent position (or "next step") for each agent based on this input. This predicted position was then compared to the actual next position of the corresponding agent in the empirical data. The discrepancy between the predicted and observed steps offers a direct validation of the model's decision-making process, allowing for a detailed evaluation of its ability to replicate pedestrian behavior under realistic conditions.

Lastly, the goal reached percentage is included as a sanity check to ensure that the simulated agents successfully complete their navigation tasks. This metric captures failure cases due to poor potential field design, such as agents getting stuck in local minima or being blocked by high-gradient regions.

Since the sampling frequency of the empirical data (25 Hz) differs from that of the simulation model (4 Hz), all KPIs that are influenced by the sampling rate, which are all collision-related KPIs, the variance in y-direction KPI, and the goal reached percentage KPI, were normalized to be invariant with respect to the sampling rate. This has been done by multiplying each KPI by the sampling rate corresponding to its respective dataset. Furthermore, because the number of agents spawning in the simulation is governed by a stochastic process and may not match the number of agents observed in the real-life experiment, these KPIs were also normalized by dividing by the total number of agents present in the corresponding dataset. These adjustments ensure that the KPIs reflect per-agent and per-second rates, allowing for meaningful comparisons between simulated and empirical data.

Together, these KPIs provide a multi-scale evaluation framework that enables robust optimization of the model parameters and ensures both physical plausibility and behavioral realism.

The values for all KPIs cannot be less than zero. However, for some KPIs, e.g. the wall collision KPI, the values for the training data will be close to zero. Therefore, to ensure that the surrogate model does not predict any negative values for any of the KPIs, the surrogate model has been trained on the square root of the training data values. The final predicted KPI results from squaring the surrogate model prediction value. In this manner, no negative values will be predicted by the surrogate model.

3-2-6 Hardware setup and implementation

For all computations in this study, the software used was MATLAB R2021b and it was performed on a machine with a Ryzen 3700x CPU and 16GB of RAM.

Chapter 4

Optimization Results

The optimization process was established by assigning an objective function of the form:

$$\min_{\tilde{\mathbf{y}} \in \mathcal{R}^{12}} f(\tilde{\mathbf{y}}) = [\mathbf{y} - \tilde{\mathbf{y}}]^T \Sigma [\mathbf{y} - \tilde{\mathbf{y}}] \quad (4-1)$$

Here, \mathbf{y} is a vector containing the values of the KPIs for the real-life bidirectional corridor experiment, $\tilde{\mathbf{y}}$ is a vector containing the KPI values predicted by the surrogate model and Σ is a weight matrix.

The optimization has been performed in three steps. First of all, the optimization was performed by assigning unity weights to all KPIs by using an identity matrix as the weight matrix. This optimization yielded the following results:

$$KW^* = 0.0251, \quad (4-2)$$

$$KG^* = 0.0192, \quad (4-3)$$

$$KO^* = 0.0045. \quad (4-4)$$

Subsequently, a local linearization was performed around this optimum to estimate the covariance matrix. Each parameter was perturbed five times with a maximum absolute perturbation of 10^{-4} . For each perturbation, a vector was constructed containing the residuals between the predicted KPI values at the optimum and those at the perturbed parameter values. These residual vectors were assembled into a matrix, where each row corresponds to the perturbations of a single parameter. The covariance matrix was then computed from this residual matrix and the resulting covariance matrix is shown below:

$$\begin{bmatrix}
1 & 0.99 & 0.99 & 0.97 & 0.91 & -0.49 & 0.63 & 0.95 & 0.84 & 0.96 & 0.06 & -0.45 \\
0.99 & 1 & 0.99 & 0.98 & 0.93 & -0.54 & 0.62 & 0.95 & 0.87 & 0.97 & 0.15 & -0.54 \\
0.99 & 0.99 & 1 & 0.98 & 0.92 & -0.54 & 0.63 & 0.96 & 0.88 & 0.96 & 0.16 & -0.55 \\
0.97 & 0.98 & 0.98 & 1 & 0.98 & -0.67 & 0.48 & 0.89 & 0.79 & 0.98 & 0.26 & -0.58 \\
0.91 & 0.93 & 0.92 & 0.98 & 1 & -0.74 & 0.30 & 0.80 & 0.68 & 0.97 & 0.30 & -0.53 \\
-0.49 & -0.54 & -0.54 & -0.67 & -0.74 & 1 & 0.18 & -0.33 & -0.21 & -0.64 & -0.56 & 0.58 \\
0.63 & 0.62 & 0.63 & 0.48 & 0.30 & 0.18 & 1 & 0.80 & 0.85 & 0.49 & -0.17 & -0.34 \\
0.95 & 0.95 & 0.96 & 0.89 & 0.80 & -0.33 & 0.80 & 1 & 0.97 & 0.90 & 0.06 & -0.52 \\
0.84 & 0.87 & 0.88 & 0.79 & 0.68 & -0.21 & 0.85 & 0.97 & 1 & 0.80 & 0.11 & -0.58 \\
0.96 & 0.97 & 0.96 & 0.98 & 0.97 & -0.64 & 0.49 & 0.90 & 0.80 & 1 & 0.28 & -0.57 \\
0.06 & 0.15 & 0.16 & 0.26 & 0.31 & -0.56 & -0.17 & 0.06 & 0.11 & 0.28 & 1 & -0.82 \\
-0.45 & -0.54 & -0.55 & -0.58 & -0.53 & 0.59 & -0.34 & -0.52 & -0.58 & -0.57 & -0.82 & 1
\end{bmatrix}
\quad (4-5)$$

The inverse of this covariance matrix was used as the weight matrix in the second optimization step, which yielded:

$$KW^* = 0.0117, \quad (4-6)$$

$$KG^* = 0.0045, \quad (4-7)$$

$$KO^* = 0.0016. \quad (4-8)$$

Since this optimum differed substantially from the first, another linearization was performed around the new optimum to obtain an updated covariance matrix:

$$\begin{bmatrix}
1 & 0.97 & 0.99 & 0.94 & 0.45 & -0.77 & 0.43 & 0.97 & 0.95 & 0.68 & -0.11 & -0.54 \\
0.97 & 1 & 0.99 & 1 & 0.63 & -0.86 & 0.20 & 0.91 & 0.95 & 0.83 & -0.06 & -0.43 \\
0.99 & 0.99 & 1 & 0.98 & 0.59 & -0.82 & 0.29 & 0.95 & 0.97 & 0.79 & -0.06 & -0.46 \\
0.94 & 1 & 0.98 & 1 & 0.69 & -0.87 & 0.13 & 0.88 & 0.94 & 0.87 & -0.04 & -0.39 \\
0.45 & 0.63 & 0.59 & 0.69 & 1 & -0.65 & -0.52 & 0.46 & 0.62 & 0.95 & 0.31 & 0.24 \\
-0.77 & -0.86 & -0.82 & -0.87 & -0.65 & 1 & 0.12 & -0.63 & -0.70 & -0.77 & 0.11 & 0.34 \\
0.43 & 0.20 & 0.29 & 0.13 & -0.52 & 0.12 & 1 & 0.47 & 0.31 & -0.31 & -0.47 & -0.61 \\
0.97 & 0.91 & 0.95 & 0.88 & 0.46 & -0.63 & 0.47 & 1 & 0.98 & 0.67 & -0.00 & -0.43 \\
0.95 & 0.95 & 0.97 & 0.94 & 0.62 & -0.70 & 0.31 & 0.98 & 1 & 0.81 & 0.02 & -0.33 \\
0.68 & 0.83 & 0.79 & 0.87 & 0.95 & -0.77 & -0.31 & 0.67 & 0.81 & 1 & 0.26 & 0.04 \\
-0.11 & -0.06 & -0.06 & -0.04 & 0.31 & 0.11 & -0.47 & -0.00 & 0.02 & 0.26 & 1 & 0.52 \\
-0.54 & -0.43 & -0.46 & -0.39 & 0.24 & 0.34 & -0.61 & -0.43 & -0.33 & 0.04 & 0.52 & 1
\end{bmatrix}
\quad (4-9)$$

Using this covariance matrix as the weight matrix in the third optimization step yielded the same optimum as in the second step. Consequently, this solution was adopted as the final result of the surrogate-model-based optimization.

Since the optimization was performed on the surrogate model, minor adjustments were necessary to obtain the corresponding optimum for the full simulation model. An iterative refinement process was therefore applied, leading to the following best-performing parameters:

$$KW = 0.0128, \quad (4-10)$$

$$KG = 0.0060, \quad (4-11)$$

$$KO = 0.0014. \quad (4-12)$$

Surrogate Model Validation

In this chapter, the accuracy of the surrogate model is checked by means of a validation analysis. This is done to validate its use for the optimization of the simulation model and the sensitivity analysis.

As described, the surrogate model is a mathematical representation of the potential field model. The surrogate model is used only for optimization purposes, not for direct simulation or prediction. Therefore, the main goal of the validation analysis is to ensure that the surrogate model is accurately mimicking the potential field model, especially around the found optimum.

Three validation analysis methods are used to validate the surrogate model: coefficient of determination analysis, local error analysis and predictive uncertainty analysis.

5-1 Coefficient of determination analysis

A coefficient of determination, R^2 , is a statistical measure used to assess how well a mathematical model explains the variability in a set of observed data [31]. It is defined as follows:

$$R^2 = 1 - \frac{\sum_{i=1}^N (y_{i,true} - y_{i,pred})^2}{\sum_{i=1}^N (y_{i,true} - \bar{y}_{true})^2} \quad (5-1)$$

Here, N represents the number of data points, y_{true} is the true data and y_{pred} is the output predicted by the surrogate model. Generally, a surrogate model with an R^2 value above 0.8 is considered a ‘good’ model [6]. The R^2 values for the surrogate models for each KPI are depicted in Table 5-1.

Among all surrogate models, the one predicting the β parameter of the path distribution in the y-axis direction exhibits the lowest predictive accuracy, with an R^2 value falling below the threshold of 0.8. A scatter plot comparing predicted and true values for this model is shown in Figure 5-1, where the ten least accurate predictions are highlighted in red. Notably,

the majority of these poorly predicted points correspond to scenarios with $y_{\text{true}} \geq 1.3$, a region that is relatively sparsely populated in the dataset. This suggests that relatively few simulations captured pedestrian dynamics resulting in such skewed path distributions. As a consequence, the surrogate model has had limited training exposure to this region of the parameter space, which likely contributes to the reduced accuracy.

Table 5-1: R^2 values for all surrogate models

KPI	R^2
Agent collisions (<0.1 m)	0.9298
Agent collisions (<0.3 m)	0.9917
Agent collisions (<0.5 m)	0.9883
Agent collisions (<1.0 m)	0.9867
Wall collisions	0.8129
Speed distribution alpha	0.9184
Speed distribution beta	0.9284
Paths y-axis distribution alpha	0.9035
Paths y-axis distribution beta	0.7891
Variance y-direction	0.9840
Discrepancies	1.00
Goal reached percentage	0.9957




Figure 5-1: Scatter plot of surrogate fit of beta parameter paths distribution y-axis model. The ten worst fitted points are depicted in red and the line visualizing a perfect fit is depicted in pink.

Despite the lower model fit, this limitation is unlikely to undermine the validity of the simulation outcomes. Since high β values in the y-axis path distribution are relatively rare, the model's predictive performance remains reliable in the regions most relevant to practical scenarios. This conclusion is supported by the true and predicted distributions for the three worst-fitting cases, shown in Figure 5-2. Even in these cases, the surrogate model reproduces

the general shape of the true distributions reasonably well, suggesting that the predictive limitations are not critical.

Furthermore, Figure 5-3 displays the three best-fitting distribution cases, which illustrate the surrogate model's capacity to closely match the true distributions when sufficient training data is available. These results underscore that the model is robust and accurate in the parameter regions most frequently encountered during simulation, supporting its overall suitability for the sensitivity and optimization tasks at hand.

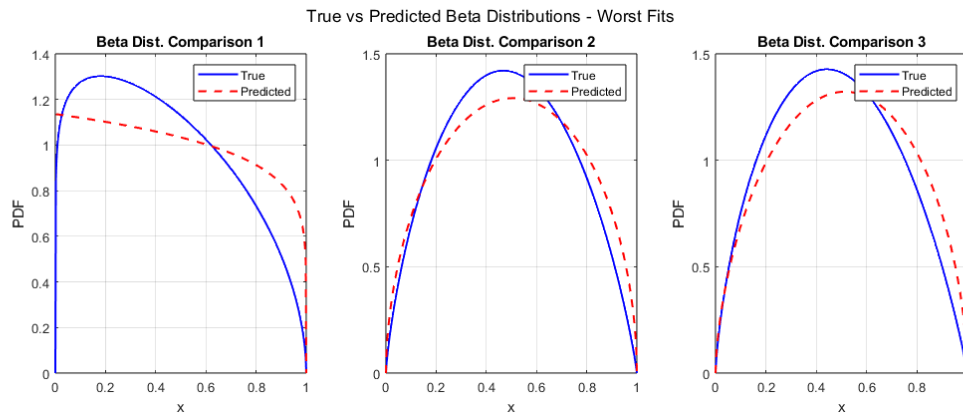


Figure 5-2: The three worst fitted distributions for the beta parameter paths distribution y-axis model with appurtenant alpha values.

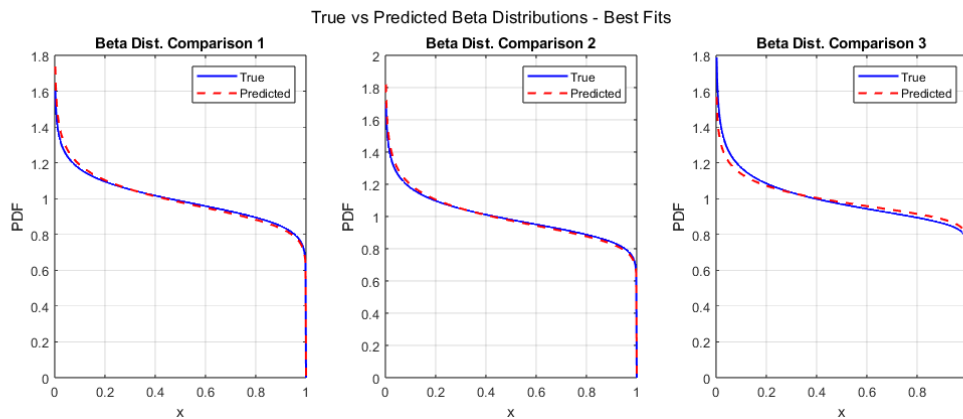


Figure 5-3: The three best fitted distributions for the beta parameter paths distribution y-axis model with appurtenant alpha values.

5-2 Local Error Analysis

As an additional validation check, the values of the KPIs predicted by the surrogate model are compared to the values of the KPIs produced by simulating the simulation model for the bidirectional corridor experiment, as depicted in Table 5-2. The values are an average of three separate simulations.

Table 5-2: KPI values for simulation model and surrogate model at the optimum

KPI	Simulation Model	Surrogate Model
Agent collisions (<0.1 m)	0.0823	0.0850
Agent collisions (<0.3 m)	1.7215	1.7335
Agent collisions (<0.5 m)	4.3509	4.7346
Agent collisions (<1.0 m)	15.6158	19.5395
Wall collisions	0	7.5813e-04
Speed distribution alpha	0.8199	0.7955
Speed distribution beta	20.3820	20.2791
Paths y-axis distribution alpha	1.1252	1.2179
Paths y-axis distribution beta	1.1736	1.2724
Variance y-direction	1.3246	1.5661
Discrepancies	2.2124e+05	2.1420e+05
Goal reached percentage	0.9576	0.9086

The comparison between the KPI values obtained from the surrogate model and the simulation model indicates that the surrogate model reproduces the majority of the key behavioral and interactional features of the pedestrian dynamics with high fidelity. For short-range inter-agent proximity events (distances below 0.1 m and 0.3 m), the surrogate model shows excellent agreement with the simulation model, suggesting that its learned interaction dynamics are well calibrated for collision avoidance and close-range spacing. Similarly, the speed distribution parameters (α and β) and the lateral trajectory distribution parameters are reproduced with only minor deviations, indicating that both movement speeds and spatial positioning patterns are faithfully captured.

However, discrepancies become more noticeable for longer-range proximity events (0.5 m and 1.0 m), where the surrogate model slightly overestimates the frequency of such events, potentially reflecting a reduced sensitivity to mid-range personal space. A similar overestimation is observed in the variance of the y-direction trajectories, suggesting a tendency towards marginally more erratic lateral movement in the surrogate model compared to the simulation model.

The wall collision rate in the surrogate model is negligible but not zero, whereas the simulation model reported no wall collisions. Although the absolute magnitude of this difference is small, it could indicate a subtle divergence in obstacle avoidance behavior. The “discrepancies” KPI, which directly measures the difference between predicted and actual next-step positions, shows close agreement, further confirming that the surrogate model’s decision-making is consistent with the simulation model at the local, per-timestep level.

Finally, the goal completion percentage is slightly lower in the surrogate model, which may reflect occasional inefficiencies in navigation or increased susceptibility to local minima in the potential field. Overall, the surrogate model demonstrates strong predictive capability across all KPIs, with only modest deviations in mid-range proximity handling and lateral stability, making it a reliable computational proxy for the full simulation model in the studied scenario.

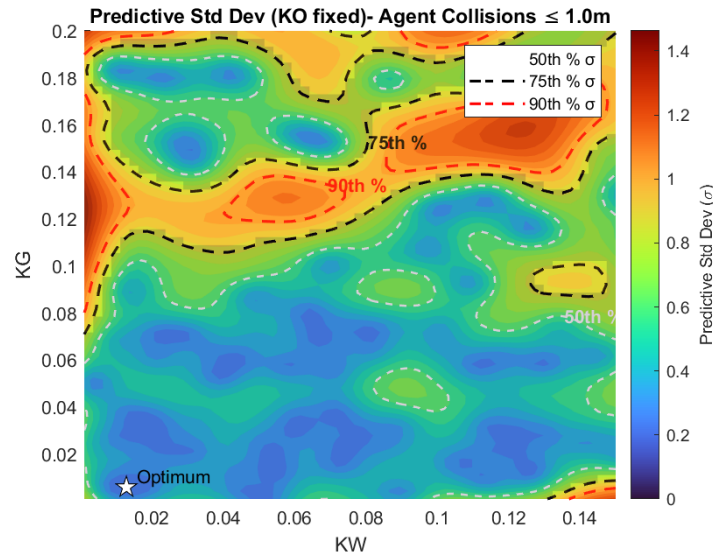


Figure 5-4: Standard deviations for GPR model for agent collisions < 0.1 m KPI. Here, KO has been kept fixed.

5-3 Predictive Uncertainty Analysis

Since the GPR model inherently provides estimates of the standard deviation, these were utilized to perform a predictive uncertainty analysis in order to assess whether the model exhibits sufficient certainty in the vicinity of the optimum. In this analysis, one parameter was fixed (KO in this case), while the other two parameters were perturbed across the search space defined for the optimization. The corresponding standard deviation values were then computed and mapped for all KPIs. An illustrative example of the resulting standard deviation distribution for one of the KPI models is presented in Figure 5-4.

The results indicate that the optimum is located well within the 50th percentile of the predicted standard deviation values. This finding suggests that the model demonstrates relatively high confidence in its predictions near the optimum. Furthermore, since multiple regions of the parameter space exhibit comparable levels of standard deviation, the optimum does not appear to be located in an overtrained region but rather in a representative and stable portion of the parameter space.

Since this behavior is consistently observed across all KPIs, as demonstrated in Appendix A, which provides an overview of the complete uncertainty analysis, the model outputs in the vicinity of the optimum can be regarded as reliable, neither excessively uncertain nor indicative of overtraining.

Simulation Model Testing and Validation

In this chapter, the simulation model is subjected to testing and validation. Testing is conducted by replicating the bidirectional experiment used for training the surrogate model and comparing the simulation outputs with the real-life data. Validation is performed using both additional metrics from the bidirectional corridor experiment and a distinct experiment involving a bottleneck scenario.

6-1 Testing

Since the surrogate model was trained on data from a bidirectional corridor experiment, the simulation model as explained in Chapter 4 is tested against this dataset as well. This dataset concerns the experiment as explained in Section 3-2-4. For each agent, the captured data consists of the agent's ID, their position and the corresponding timestep of the experiment. A twosome of these experiments has been used for the validation data, consisting of 151 and 176 agents, respectively.

As the simulation model is computationally heavy and it has to be ran for a lot of different parameter sets to create a reliable surrogate model, the simulation time for the simulations in the dataset cannot be too large. Therefore, for the real-life dataset to be suitable, it has been filtered to only contain agents that spawn within the simulation time of the simulation used when creating the training data for the surrogate model.

For testing, first of all, the KPIs of the simulation model have been compared to the real-world dataset of the bidirectional experiments as presented in Table 6-1. As the discrepancies KPI is only valuable for optimizing the surrogate model, this KPI has been left out for this comparison.

First of all, the collision-related KPIs (<0.1 m, <0.3 m, <0.5 m, and <1.0 m) reveal several important patterns. At the <0.1 m threshold, the real-life data reports no such close encounters, whereas the simulation yields a non-zero value (0.0823). Although this value is still low,

Table 6-1: KPI values for simulation model and real-life data for bidirectional corridor experiment

KPI	Simulation Model	Real-Life Data
Agent collisions (<0.1 m)	0.0823	0
Agent collisions (<0.3 m)	1.7215	0.3273
Agent collisions (<0.5 m)	4.3509	4.138
Agent collisions (<1.0 m)	15.6158	33.7264
Wall collisions	0	0
Speed distribution alpha	0.8199	0.7048
Speed distribution beta	20.3820	20.2518
Paths y-axis distribution alpha	1.1950	1.3777
Paths y-axis distribution beta	1.2696	1.7309
Variance y-direction	1.3246	0.4993
Goal reached percentage	0.9576	0.9432

it indicates that the model occasionally permits agents to approach distances that, in the experiment, were never observed. Given that separations below 0.1 m represent near-physical contact, this difference is notable. Moreover, for <0.3 m encounters, the simulation records substantially more occurrences than the real-life data, pointing to a tendency for agents to accept closer interpersonal distances than observed experimentally. Overall, it seems that the model struggles somewhat to represent pedestrian dynamics at low proximities.

Second, in the intermediate range around <0.5 m, the simulation aligns closely with empirical observations. This suggests that for medium-range avoidance, where pedestrians must navigate but do not necessarily take evasive action, the model reproduces human behavior with high fidelity. Furthermore, at the largest threshold (<1.0 m), the simulation produces fewer proximity events than the empirical dataset, indicating that it may underestimate the frequency of mild crowding situations or larger-scale interactions.

The wall collision KPI remains zero in both datasets, confirming that the simulation reliably prevents agents from making contact with static boundaries.

Regarding movement-related KPIs, the fitted beta distribution parameters for both speed, which are depicted in Figure 6-1 and Figure 6-2, are in close agreement with those from the real-life data, indicating that the overall speed profiles. For the lateral (y-axis) positions, which are depicted in Figure 6-3 and Figure 6-4, there is some difference as the trajectories seem to be more spread out towards the edges of the corridor for the simulation experiment while being slightly more concentrated towards the center of the corridor for the real-life data. However, as can be seen in the plots, the overall shape of the distribution is similar which indicates that the general spatial usage of the corridor is captured well.

However, a more pronounced difference emerges in the variance of the y-direction: the real-life trajectories exhibit a variance of 0.4993, whereas the simulation yields a higher value of 1.3246. This indicates that the simulated agents follow more laterally oscillatory or “squiggly” paths than their real-life counterparts. This could be a byproduct of the gradient-descent navigation method producing small oscillations in direction.

Finally, the goal-reached percentage is nearly identical for simulation and reality, confirming that almost all agents successfully reach their destination in both scenarios.

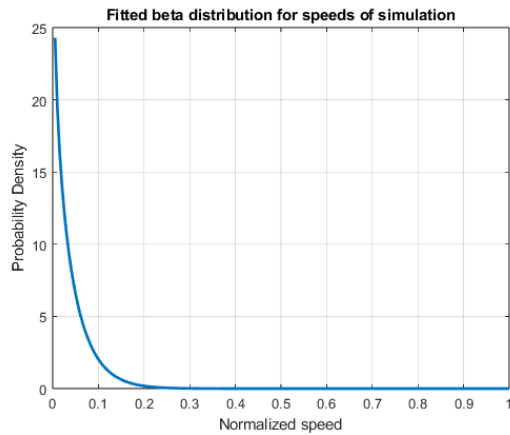


Figure 6-1: Beta distribution of normalized speeds for simulated experiment.

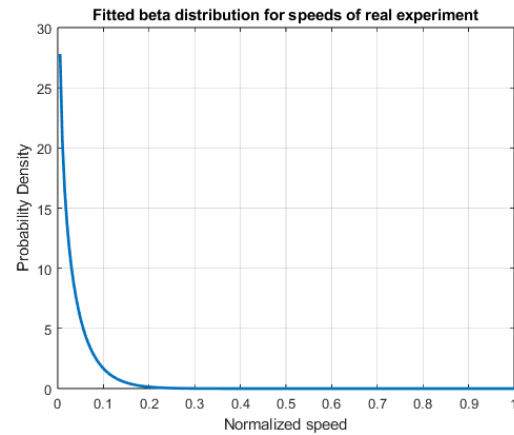


Figure 6-2: Beta distribution of normalized speeds for real-life experiment.

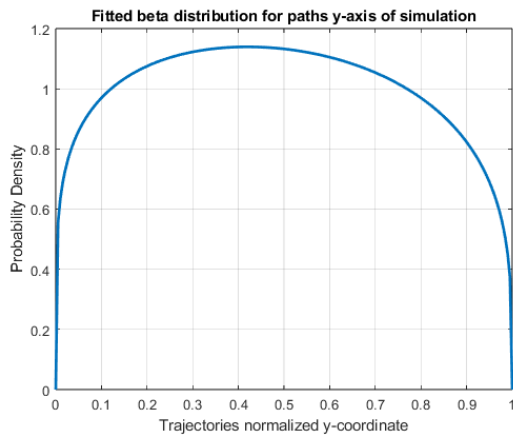


Figure 6-3: Beta distribution of normalized y-coordinates of trajectories for simulated experiment.

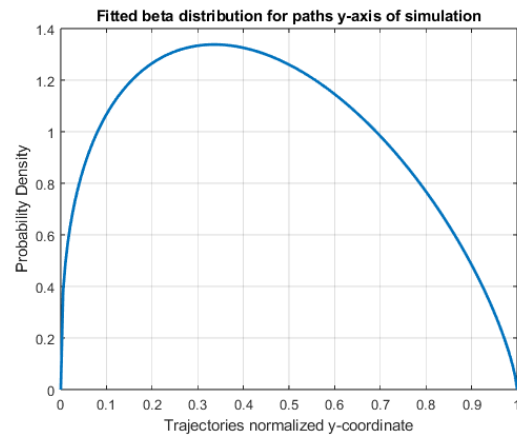


Figure 6-4: Beta distribution of normalized y-coordinates of trajectories for real-life experiment.

In summary, the simulation model reproduces medium-range collision avoidance and overall locomotion patterns with high accuracy, but diverges from the empirical data at very short interpersonal distances (<0.1 m), overestimates the frequency of close encounters (<0.3 m), and generates more laterally variable trajectories. These findings suggest that while the model successfully captures the main features of the bidirectional corridor flow, further refinement of short-range avoidance behavior and trajectory smoothness could yield even closer correspondence with observed pedestrian dynamics. However, as the general characteristics of the bidirectional corridor experiment are matched by the model, it is a reliable basis for further validation and extension.

6-2 Validation

In this section, the simulation model is validated by means of performance metrics that were not employed during the optimization of the surrogate model. Specifically, this validation is carried out by analyzing additional metrics from the bidirectional corridor experiment and evaluating metrics derived from an separate scenario involving a bottleneck configuration.

6-2-1 Bidirectional Corridor Experiment

In addition to the tests presented in the previous section, the bidirectional corridor experiment is further utilized to assess the validity of the simulation model. The focus of this assessment lies on metrics related to the phenomenon of lane formation, which describes the self-organized emergence of unidirectional paths within a crowd. Individuals entering such lanes tend to remain within them by following those ahead, thereby reinforcing the structure.

As commonly observed in densely populated environments, lane formation also emerged in the empirical bidirectional corridor experiment, despite the fact that participants were entirely unconstrained in their point of entry and their choice of route through the corridor [11]. However, since the agents in the real-life experiment had already entered the corridor before their positions were captured by the camera, formation characteristics were already present prior to the start of the recorded experiment. To enable a meaningful comparison, the distributions of the agents' initial y -coordinates for each walking direction were examined. Figure 6-7 and Figure 6-8 present histograms of these initial y -positions for both walking directions in the empirical bidirectional corridor experiment. The results show that both distributions are skewed toward the opposite side of the corridor, indicating that agents were more likely to begin their trajectories near one side of the corridor depending on their walking direction. To reproduce this effect in the simulated experiment, beta distributions were fitted to the empirical starting y -coordinates for each walking direction, and the simulated agents' initial y -positions were subsequently sampled from these fitted distributions.

Figure 6-5 illustrates a snapshot of the real experiment, displaying the direction of the target location for each participant. The snapshot corresponds to a time step 10 seconds after the start of the experiment, which reflects the delay required for lane formation to develop. It can be observed that the participants organized into two predominant groups moving in opposite directions, with only a few individuals deviating from this structure.

For the simulated experiment, a snapshot at the same timestep is depicted in Figure 6-6. A similar observation can be made as there are again two main groups, although they seem to be slightly less organized than for the real-life experiment. Furthermore, it can be observed that the agents seem to walk closer to each other, in line with the conclusion from model testing.

Furthermore, Figure 6-9 presents two overlapping histograms of the lateral distributions of agents walking from left to right and from right to left in the real-life experiment. It is clearly observable that the peak of the histogram corresponding to one walking direction generally coincides with the trough of the distribution for the opposite direction, and vice versa. This pattern provides additional evidence for the presence of lane formation. Again, the much similarity can be seen in Figure 6-10, which are the same histograms for the simulated experiment, although the inverse correlation is less strong.

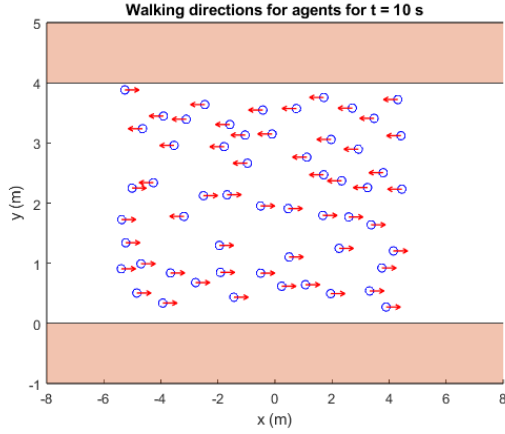


Figure 6-5: Walking directions for agents of real-life bidirectional corridor experiment at $t = 10$ s.



Figure 6-6: Walking directions for agents of simulated bidirectional corridor experiment at $t = 10$ s.

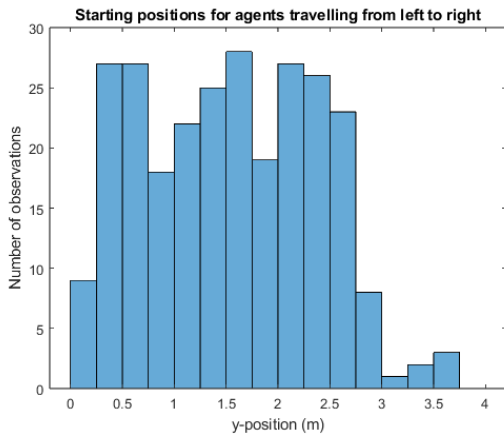


Figure 6-7: Histogram of starting y-coordinates for agents of real-life bidirectional corridor experiment walking towards the positive x-direction.

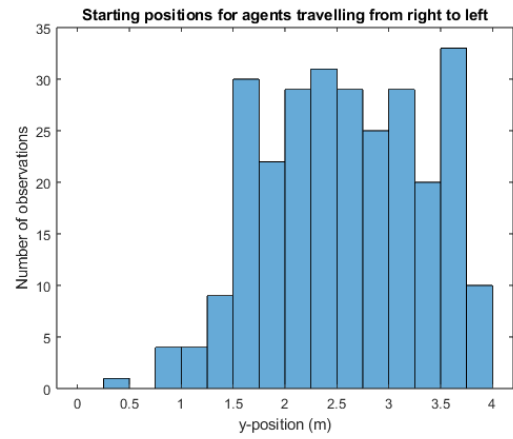


Figure 6-8: Histogram of starting y-coordinates for agents of real-life bidirectional corridor experiment walking towards the negative x-direction.

Additionally, to get a more quantified understanding of the lane formation in the empirical bidirectional corridor experiment, a number of metrics has been computed. First of all, the overlap coefficient (OVL) is used. This metric is used to determine the overlap of distributions and is defined as:

$$OVL = \int \min\{p_+(y), p_-(y)\} dy \quad (6-1)$$

Here, $p_+(y)$ is the distribution corresponding to the trajectories for the agents walking in the positive x-direction while $p_-(y)$ is the distribution for the agents walking in the negative x-direction. In this context, the OVL can be used to determine to what extent the distributions of the lateral trajectory positions coincide for both walking directions. The coefficient ranges

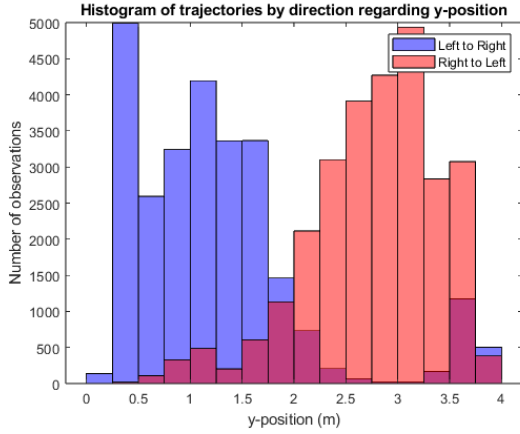


Figure 6-9: Overlapping histograms of the distribution of lateral trajectory positions for both walking directions in the real-life experiment.




Figure 6-10: Overlapping histograms of the distribution of lateral trajectory positions for both walking directions in the simulated experiment.

from 0 to 1, with a value of zero meaning there is no coincidence of the two distributions while a value of 1 means the distributions are exactly the same.

The second metric is the zero-lag coefficient which is derived from the cross-correlation between the two distributions for lag $\tau = 0$:

$$R[\tau = 0] = \sum_{n=1}^N p_+[n] \cdot p_-[n] \quad (6-2)$$

Here, N is the number of bins for the histograms. When computing the zero-lag coefficient, this number has been increased compared to the histograms in the figures, to create a more accurate representation of the distribution.

A positive zero-lag correlation means that the density peaks of the two directions occur at similar lateral positions, indicating a lack of lane segregation and more overlap between flows. A negative zero-lag correlation implies that the peaks of one distribution align with the troughs of the other. This is characteristic of interleaving lanes, where pedestrians moving in opposite directions occupy alternating lateral positions. Finally, a value close to zero indicates no clear relationship.

The third metric is average peak offset which quantifies the lateral misalignment between the peak density locations of opposing pedestrian streams. In histograms of the lateral positions, each direction tends to produce peaks corresponding to preferred walking lanes. In well-formed bidirectional lanes, these peaks are interleaved: the maxima of one distribution coincide with the minima of the other, minimizing conflicts. The peak offset is defined as the average lateral distance between these peaks, serving as a measure of how well the two flows avoid spatial overlap. A small peak offset implies weak or unstable lane segregation, whereas a larger and more consistent offset indicates effective interleaving and stronger spatial coordination. This metric complements lane spacing by directly measuring the spatial alignment between opposing streams.

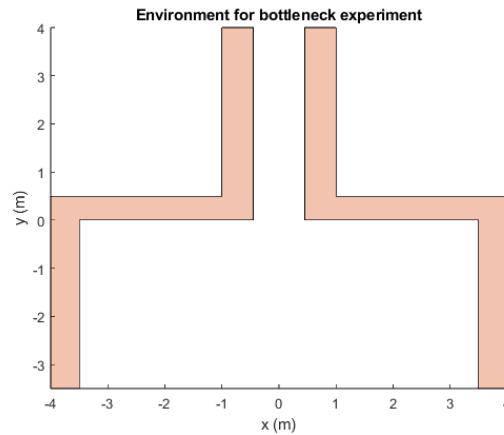


Figure 6-11: Environment for the bottleneck experiment.

A quantitative comparison of the lane formation metrics is depicted in Table 6-2. The OVL was similar but slightly higher in the simulation than in the real data, suggesting that the simulated flows exhibit greater overlap between counterflowing streams and therefore less distinct lane segregation, albeit only slightly. The zero-lag correlation was moderately negative for the simulation and strongly negative for the real data, indicating that while both experiments show a certain degree of alternating lane structure, the effect is more pronounced in reality. Finally, the peak offset between the two directional histograms was very comparable in the simulation and the real data. Taken together, these results indicate that the model captures the qualitative tendency for lanes to emerge in bidirectional flow in a similar manner to the real-life experiment, although the lanes are less pronounced and marginally more overlapping than in the empirical data.

Table 6-2: Lane forming metrics for simulation model and real-life data for bidirectional corridor experiment

Lane Forming Metric	Simulation Model	Real-Life Data
Average overlap coefficient	0.317	0.215
Average zero-lag correlation	-0.332	-0.708
Average Peak offset (m)	0.16	0.20

6-2-2 Bottleneck experiment

In addition to the bidirectional corridor experiment, the model was further validated using a different experimental setup involving a bottleneck scenario, also with a corresponding real-life dataset [10]. The corresponding environment is illustrated in Figure 6-11. In this experiment, agents enter the environment from the negative y -direction and attempt to traverse the narrow corridor toward the positive y -direction. There is a camera mounted on the ceiling which detects and tracks the trajectories of all agents. As in [10], it is not mentioned what area is captured exactly by the camera it has been assumed that the camera tracks agents between $y = -2.7$, since this is the rounded average y -coordinate of the agents entering the experiment and $y = 4$ as this is the boundary of the environment.

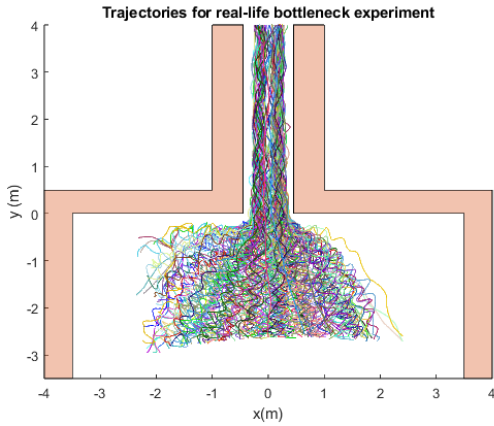


Figure 6-12: Agent trajectories for real-life bottleneck experiment.




Figure 6-13: Agent trajectories for simulated bottleneck experiment.

This experiment has been run three times with the simulation model and the average outputs have been documented. Again, the simulation time reduced for the simulation compared to the real experiment, to decrease the computational burden. Therefore, the real-life data was filtered regarding the agents' spawning times just as with the bidirectional corridor experiment. The resulting trajectories for the real-life experiment and one of the simulations are depicted in Figure 6-12 and Figure 6-13, respectively.

A first observation from the inspection of the simulated agent trajectories is that some trajectories extend beyond the camera's field of view. This occurs because agents occasionally take steps in the negative y -direction when the local density at their spawning location, which is defined as $-2.5 \leq x \leq 2.5$, $y = -2.7$, is very high. In the real experiment, however, such behavior is not observed, since agents enter the environment at a position outside the camera's range. Consequently, any sorting behavior that prevents agents from stepping backward is contained within portions of the trajectories that remain unrecorded. Moreover, even if such behavior were to occur in reality, it would be impossible to be captured by the camera.

Moreover, the previously established KPIs have been used as a validation analysis for the simulation model under the bottleneck conditions. However, since the KPIs regarding the distribution for the lateral position of the trajectories are only valuable for the bidirectional corridor experiment, these KPIs are left out for the bottleneck experiment. This is because of the fact that in the bottleneck experiment, the distribution will always look very similar, as all agents have to pass the narrow corridor, resulting in a large peak at those coordinates. A comparison for the remaining KPIs is presented in Table 6-3.

For most KPIs the outcome is very similar to that of the test case, even under these more congested circumstances. Compared to the real-life data, the simulation model again allows significantly more close encounters, while being very comparable regarding the encounters for a vicinity of 0.5 m and 1.0 m . Furthermore, there are no wall collisions, implying that, also in this environment, the physical boundaries are respected by the agents.

The speed distribution parameters show partial agreement. The α values are close, indicating that the shape of the speed profile is consistent across model and data. However, the β parameter differs substantially, implying that the simulation underestimates the skewness

Table 6-3: KPI values for simulation model and real-life data for bottleneck experiment

KPI	Simulation Model	Real-Life Data
Agent collisions (<0.1 m)	1.4249	0
Agent collisions (<0.3 m)	9.6964	0.6202
Agent collisions (<0.5 m)	18.2232	15.8779
Agent collisions (<1.0 m)	69.1577	61.7155
Wall collisions	0	0
Speed distribution alpha	1.3858	1.4310
Speed distribution beta	4.8795	12.7046
Variance y-direction	1.3129	0.4993
Goal reached percentage	0.3348	0.4471

of the empirical speed distribution, likely due to differences in how deceleration is handled in high-density conditions. The variance in the y-direction is considerably higher in the simulation than in the real data, showing that simulated trajectories are far more laterally variable or “squiggly.” This suggests that the model generates excess oscillatory movement when agents attempt to pass through the bottleneck.

The goal-reached percentage shows that the simulation yields a lower completion rate compared to the empirical data. At first glance, this might suggest that the model underestimates effective throughput. However, it must be taken into account that the simulated experiments contain a larger total number of agents, namely 233 on average, than the real-life experiment, which contains 170 agents. As a result, the lower goal-reached percentage in the simulation does not necessarily indicate poorer congestion resolution, but rather reflects the fact that a larger number of agents compete for passage through the same bottleneck within the fixed simulation time.

This interpretation is supported by the throughput metrics in Table 6-4. The simulation produces both a higher average total throughput and a higher average throughput per second. It can be seen in Figure 6-14, which shows a comparison of the throughput of the real-life experiment and one of the simulated experiments, that during most of the simulation the throughput is very comparable, but the maxima for the simulated are slightly higher. These results indicate that once agents are positioned favorably at the entrance of the bottleneck, the model allows them to pass through slightly more efficiently than in the real-life experiment. The discrepancy therefore arises primarily from differences in total population size and flow demand, rather than from a fundamental misrepresentation of local congestion dynamics.

Table 6-4: Lane forming metrics for simulation model and real-life data for bidirectional corridor experiment

Lane Forming Metric	Simulation Model	Real-Life Data
Average total throughput	85	69
Average throughput per second	2.12	1.73

Taken together, these findings suggest that the model successfully captures the qualitative difficulty of bottleneck navigation although it slightly overestimates the achievable throughput compared to the empirical data.

6-3 Intermediate Conclusion

Overall, the simulation model demonstrates a satisfactory ability to reproduce pedestrian dynamics in the two experiments used for testing and validation. Its primary limitation appears to be an overestimation of close-proximity interactions between agents, particularly in densely crowded conditions. Nevertheless, since the model is able to realistically capture fundamental phenomena such as lane formation and throughput, it is considered sufficiently validated for application in further investigations across different environments.




Figure 6-14: Throughput during bottleneck experiment for simulated experiment and real-life data.

Chapter 7

Sensitivity Analysis

To further examine the model's robustness, a sensitivity analysis is performed around the found optimum as determined in Chapter 4. The analysis is split up into a parameter sensitivity analysis and a local sensitivity analysis.

7-1 Parameter Sensitivity

In this section the parameter sensitivity is explained. To avoid a redundantly long chapter, this section focuses only on the most significant parameters for the sensitivity analysis. A full overview of the parameter sensitivity analysis can be found in Appendix B.

The parameter sensitivity analysis was conducted by gradually perturbing each individual optimization parameter by up to 10 %, while keeping all other parameters fixed. Each perturbed parameter set was then used as input to the surrogate model to predict the values of all KPIs. For every parameter and each KPI, the standard deviation of the predicted KPI values was computed as a measure of the variability induced by that parameter. These standard deviations were subsequently normalized by dividing each value by the maximum standard deviation obtained for the corresponding KPI.

As the simulation model is computationally heavy and the sensitivity analysis requires a large number of simulations, the analysis was carried out using the surrogate model. As demonstrated in Chapter 5, the surrogate model reliably reproduces the behavior of the full simulation model and is therefore considered a valid tool for performing the sensitivity analysis.

For all agent collision related KPIs ($< 0.1\text{ m}$ - $< 1.0\text{ m}$), the sensitivity analysis yielded very similar results. Therefore, only the agent collisions KPI regarding the 0.1 and 1.0 meter vicinity will be analyzed further. Interestingly, Figure 7-1 and Figure 7-2 reveal that agent collision-related KPIs are most strongly influenced by variations in the goal potential weight, whereas changes in the wall potential weight and the obstacle potential weight have comparatively minor effects. This result indicates that the dominant driver of collision dynamics in




Figure 7-1: Sensitivity Analysis for agent collisions (<0.1 m) KPI.




Figure 7-2: Sensitivity Analysis for agent collisions (<1.0 m) KPI.

the model is the strength with which agents are attracted to their goal. When KG is set too high, agents prioritize goal-seeking behavior over avoidance, thus increasing the likelihood of short-range interactions and collisions with other agents. On the other hand, when KG is reduced, agents exhibit more cautious and adaptive trajectories, resulting in fewer collisions but also potentially longer travel times.

The relatively lower sensitivity to KW and KO suggests that wall and obstacle avoidance mechanisms are sufficiently robust across a broad range of values, ensuring that agents can avoid static boundaries and dynamic obstacles regardless of precise parameter tuning. In contrast, KG directly regulates the balance between goal-directed urgency and social navigation, making it the most critical parameter for calibrating realistic collision avoidance behavior in the model. This finding underscores the importance of careful tuning of KG to achieve a realistic compromise between efficient goal attainment and safe navigation in dense crowd scenarios.

The sensitivity results further show that the relative influence of KW and KO decreases as the collision threshold is relaxed: their relative effects are slightly stronger for the < 0.1 m KPI than for the < 1.0 m KPI. This pattern is consistent with the short-range nature of the repulsive potentials, whose gradients become dominant only when agents approach very close to boundaries or each another. At larger interpersonal distances (e.g., < 1.0 m), global goal attraction, i.e. KG , plays a more significant role in shaping trajectories and inter-agent spacing, thereby weakening the marginal impact of KW and KO . In practice, this implies that calibration of KW and KO should be guided primarily by the tight-proximity KPIs (< 0.1 m, < 0.3 m), while tuning of KG should target medium-to-long range interactions (< 0.5 m, < 1.0 m). The result also suggests that the repulsive fields may be too short-ranged relative to real behavior, because of which a slightly longer decay length or speed-dependent anisotropy for KO could improve realism at intermediate distances without sacrificing near-contact safety. This interpretation is also supported by the fact that the model allows more close encounters between agents than in the real-life experiment, as explained in Chapter 6.

The parameter sensitivity for the variance in the y -direction is shown in Figure 7-3. For this KPI as well, KG emerges as the most influential parameter with respect to sensitivity. This indicates that deviations of agents in the y -direction are primarily governed by the

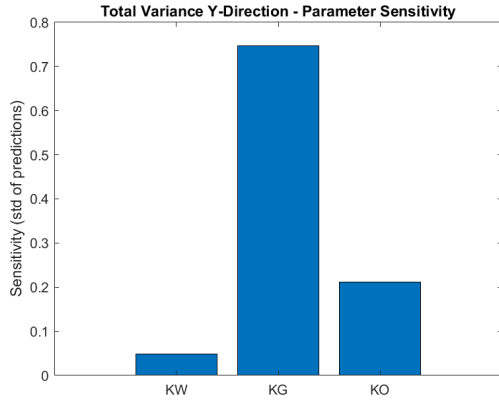


Figure 7-3: Sensitivity Analysis for variance y -direction KPI.

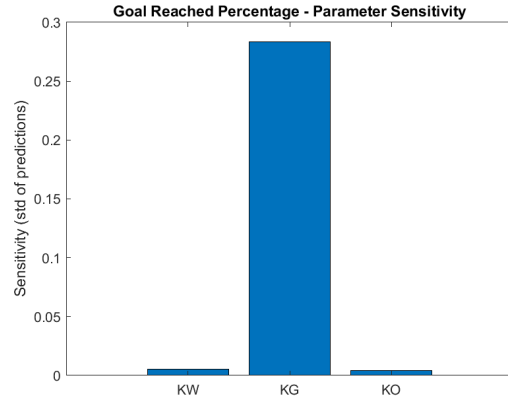


Figure 7-4: Sensitivity Analysis for goal reached percentage KPI.

strength of their attraction towards the goal. This finding is intuitive: increasing KG reduces the variance in the y -direction, as the gradient toward the goal, which is located in the x -direction, becomes steeper, thereby making motion along the x -axis more dominant at each step. Conversely, decreasing KG diminishes the dominance of the x -direction, resulting in greater variance in the y -direction.

Moreover, the results indicate that the influence of KO on the variance in the y -direction is substantially larger than that of KW . This observation is also consistent with expectations: in the bidirectional corridor experiment, agents frequently encounter other agents along their trajectory and must avoid them by sidestepping in the y -direction, whereas wall-induced deflections in the y -direction occur far less frequently.

Finally, for the goal-reached percentage KPI, KG once again appears as the most dominant parameter in terms of sensitivity. This result is intuitive, as variations in KG directly affect the degree to which agents are attracted toward their goal. However, the strong dominance of KG over both KW and KO further implies that the likelihood of agents becoming stuck due to modifications in their repulsive interactions with either other agents or the walls is very low.

7-2 Local Sensitivity

To assess the robustness of the model, a local sensitivity analysis has been performed around the optimum. For this the coefficient of variation (CV) has been used as a quantitative metric. The CV is a normalized measure of statistical dispersion, defined as the ratio of the standard deviation σ to the mean μ :

$$CV = \frac{\sigma}{\mu}. \quad (7-1)$$

Unlike the standard deviation alone, the CV is dimensionless and therefore enables a relative comparison of variability across quantities with different scales or units. In the context of this

sensitivity analysis, the CV is used to quantify the robustness of a surrogate model output with respect to perturbations in the input parameters. Low CV values indicate that the KPI remains relatively stable under parameter variations, whereas higher CV values suggest stronger sensitivity and a potentially less robust response.

Threshold values are often employed to classify the robustness of a KPI based on its CV. Commonly used guidelines include $CV \leq 0.05$ indicating *very robust*, $CV \leq 0.15$ indicating *moderately sensitive*, and $CV \geq 0.30$ indicating *sensitive* behavior [28, 1]. These thresholds are context-dependent, but they provide a practical framework for interpreting sensitivity results.

Figure 7-5 shows the CV for all KPIs around the optimal point. The results demonstrate that most KPIs lie well below the threshold for moderate sensitivity, with only the Wall Collisions KPI exhibiting slightly elevated variability, yet still within a robust to moderate range. In contrast, KPIs such as goal reached percentage, path discrepancy, and total variance in y -direction exhibit extremely low CV values, indicating high robustness. Overall, these results suggest that the simulation model is largely insensitive to small perturbations of the parameters, providing confidence in the stability and reliability of its predictions.




Figure 7-5: Coefficient of variation for each surrogate model

Experiments and Results

In addition to the corridor and bottleneck validation experiments, three exploratory simulations were conducted in a lecture hall environment, as depicted in Figure 8-1. The setup consists of a central corridor with six lecture halls. Four of the lecture halls (A, B, C, and D) are equipped with three separate entrances, while the remaining two lecture halls (E and F) have only a single entrance each. The experiments were designed to assess the performance of the simulation model under different usage scenarios, ranging from normal operation to congested conditions. For all experiments, there are at least two types of agents: agents entering the lecture halls and walking agents. Agents entering the lecture halls enter one of the lecture halls through one of its entrances. Walking agents walk from one end to the other end of the main hall. Both types of agents spawn at either $x = -14.8$ or $x = 52$.

For all experiments a simulation time T of 3600 time steps is used, with time step $dt = \frac{1}{4}$ s. The total simulation time is thus 15 minutes, which is the break time between lectures at TU Delft. Within this time all students are thus entering the lecture halls in a real-life scenario. The spawning probability of the agents is based on a normal distribution, as depicted in Figure 8-2, with mean μ at $\frac{2}{3}T$, assuming that most of the agents enter the environment approximately five minutes before the lectures commence. Moreover, a standard deviation $\sigma = \frac{1}{4}T$ was implemented. Finally, static obstacles will spawn at the beginning of the simulation. These obstacles represent people located at the study areas present in the main hall. The study areas are located at $-6 \leq x \leq 5.5$, $16 \leq x \leq 30.5$ or $42.5 \leq x \leq 50$ with $0.5 \leq y \leq 2$.

8-1 Experiment 1: Normal Conditions

In the first experiment, referred to as *normal conditions*, four of the lecture halls (A–D) are open to the agents, each with three available entrances. Agents either enter one of the open lecture halls or traverse the corridor from one end to the other. The probabilities for each agent type and corresponding spawning location for this experiment are depicted in Table 8-1.

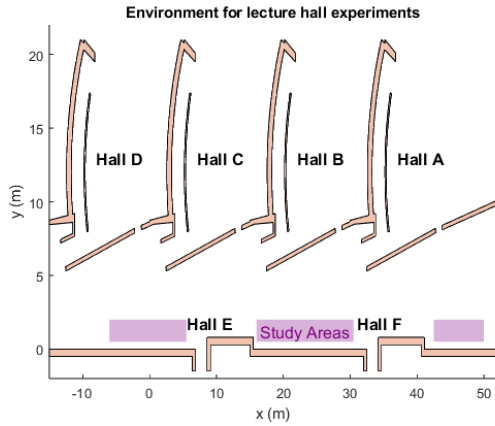


Figure 8-1: Environment for the lecture hall experiments, with indicated hall names and study areas.

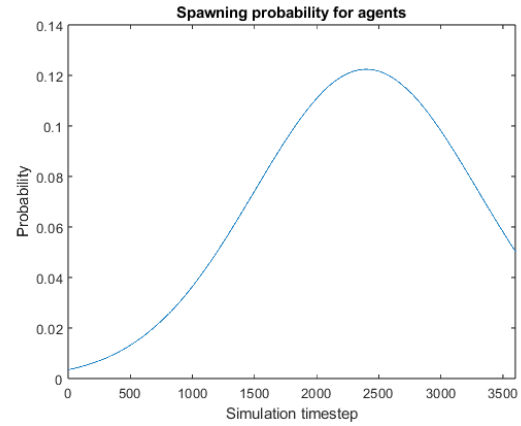


Figure 8-2: Spawning probability density function for agents in lecture hall experiments.

Table 8-1: Absolute spawning probabilities of agents for the busy conditions experiment.

Agent Type	Spawn Location	Destination	Absolute Probability
Walking	$x = -14.8$	-	$0.40 \times 0.60 = 0.24$
Walking	$x = 52$	-	$0.40 \times 0.40 = 0.16$
Main Entrance	$x = -14.8$	Hall A	$0.45 \times 0.60 \times 0.37 = 0.0999$
Main Entrance	$x = -14.8$	Hall B	$0.45 \times 0.60 \times 0.27 = 0.0729$
Main Entrance	$x = -14.8$	Hall C	$0.45 \times 0.60 \times 0.18 = 0.0486$
Main Entrance	$x = -14.8$	Hall D	$0.45 \times 0.60 \times 0.18 = 0.0486$
Main Entrance	$x = 52$	Hall A	$0.45 \times 0.40 \times 0.37 = 0.0666$
Main Entrance	$x = 52$	Hall B	$0.45 \times 0.40 \times 0.27 = 0.0486$
Main Entrance	$x = 52$	Hall C	$0.45 \times 0.40 \times 0.18 = 0.0324$
Main Entrance	$x = 52$	Hall D	$0.45 \times 0.40 \times 0.18 = 0.0324$
Back Entrance	$x = -14.8$	Upper Back Entrance	$0.15 \times 0.60 \times 0.50 = 0.045$
Back Entrance	$x = -14.8$	Lower Back Entrance	$0.15 \times 0.60 \times 0.50 = 0.045$
Back Entrance	$x = -14.8$	Hall A	$0.15 \times 0.60 \times 0.37 = 0.0333$
Back Entrance	$x = -14.8$	Hall B	$0.15 \times 0.60 \times 0.27 = 0.0243$
Back Entrance	$x = -14.8$	Hall C	$0.15 \times 0.60 \times 0.18 = 0.0162$
Back Entrance	$x = -14.8$	Hall D	$0.15 \times 0.60 \times 0.18 = 0.0162$
Back Entrance	$x = 52$	Upper Back Entrance	$0.15 \times 0.40 \times 0.50 = 0.03$
Back Entrance	$x = 52$	Lower Back Entrance	$0.15 \times 0.40 \times 0.50 = 0.03$
Back Entrance	$x = 52$	Hall A	$0.15 \times 0.40 \times 0.37 = 0.0222$
Back Entrance	$x = 52$	Hall B	$0.15 \times 0.40 \times 0.27 = 0.0162$
Back Entrance	$x = 52$	Hall C	$0.15 \times 0.40 \times 0.18 = 0.0108$
Back Entrance	$x = 52$	Hall D	$0.15 \times 0.40 \times 0.18 = 0.0108$

The normal conditions experiment serves as a baseline, representing a standard level of activity in the corridor. Agent trajectories for one of the simulations for this experiment are depicted in Figure 8-3. Moreover, a quantitative representation of related KPIs is presented in Table 8-2.

The results demonstrate that the model produces smooth and efficient pedestrian flows when

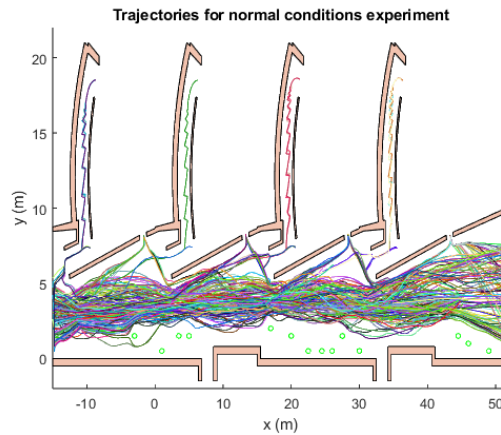


Figure 8-3: Agent trajectories for one of the normal conditions experiments.

lecture halls are readily accessible and no artificial constraints are introduced. Collisions at very short distances are rare and only slightly higher for the < 0.3 m threshold. At intermediate and larger distance thresholds (< 0.5 m and < 1.0 m), the values remain modest, consistent with a setting characterized by moderate local interactions but no significant congestion. Notably, no wall collisions occur, showing that agents respect the boundary conditions of this environment.

Table 8-2: Average KPI values for normal conditions experiment

KPI	Average Value
Agent collisions (< 0.1 m)	0.0926
Agent collisions (< 0.3 m)	0.6546
Agent collisions (< 0.5 m)	1.2586
Agent collisions (< 1.0 m)	5.4990
Wall collisions	0
Speed distribution alpha	2.0190
Speed distribution beta	3.0480
Goal reached percentage	0.9542

The speed distribution parameters indicate a stable velocity profile across agents, with a slight bias towards walking faster, which aligns with expectations for a free-flow environment where movement is not substantially delayed by bottlenecks or closures. Furthermore, it can be seen that, although the vast majority of the agents does reach their goal successfully, the goal reached percentage is lower than 100%. However, this does not necessarily imply that agents get stuck in the environment. This is supported by the depicted spawning times of agents that have not reached their goal for one of the delayed opening experiments depicted in Table 8-3. It can be seen that all agents that have not reached their goal yet have spawned in the latter part of the simulation. Furthermore, it only regards agents that are either of the walking type and thus have to walk through the entire environment, or are spawning on the opposite side of their goal destination. Thirdly, especially for the agents that spawned earliest out of all agents in Table 8-3, the alpha values, as explained in Chapter 3, are quite low. This means that their step size is relatively small, inducing a longer travel time before

successfully reaching their destination. Therefore, they simply have not had enough time to reach their goal yet. This thus represents agents that are late for the lecture or, in the case of the ‘walking’ agents, have not reached the other side of the corridor yet.

Table 8-3: Characteristics of agents that have not reached their goal for one of the normal conditions experiments.

Agent ID	Agent Type	Spawn Location	Alpha Value	Spawning Timestep
226	Hall D	x = 52	8.9423	3074
233	Walking	x = 52	18.6159	3136
238	Hall D	x = 52	25.4062	3203
247	Walking	x = 52	39.6489	3297
248	Walking	x = 52	32.3886	3309
251	Walking	x = 52	44.5988	3348
252	Walking	x = -14.8	18.3900	3354
255	Walking	x = 52	33.2503	3379
259	Hall A	x = -14.8	51.9390	3463
260	Walking	x = 52	35.1939	3478
261	Walking	x = 52	36.3508	3509
262	Walking	x = 52	46.2687	3536
263	Walking	x = -14.8	32.0927	3543

Overall, these findings suggest that the model is able to reproduce realistic and efficient pedestrian dynamics in environments where capacity is not artificially restricted. The relatively low incidence of close-proximity events, combined with a stable speed distribution and high completion rate, highlight the robustness of the model in simulating baseline pedestrian traffic patterns without emergent congestion effects.



8-2 Experiment 2: Delayed Opening

The second experiment, *delayed opening*, uses the same initial conditions as the normal conditions experiment, but introduces a dynamic change in accessibility. Specifically, lecture hall A keeps its entrances closed for half of the simulation time, i.e. for 1800 timesteps. Agents assigned to this hall must wait in front of its main entrance until it opens, after which they may enter. In this scenario, agents only enter hall A via its main entrance. This scenario is designed to mimic real-life conditions where a classroom becomes available only after a delay, leading to local crowding and temporary congestion. The average main KPIs for this experiment are depicted in Table 8-4. Furthermore, the agent trajectories for one of the simulations are depicted in Figure 8-4.

The results show that, although they increase compared to the normal conditions experiment, the frequency of close interactions remains limited. Collisions at distances below 0.1 m are rare, and even when extending the threshold to 0.3 m and 0.5 m, the number of events remains low. This indicates that agents maintain realistic interpersonal spacing despite temporary queuing in front of Hall A. However, due to the more congested conditions, the inter-agent encounters are increased compare to the normal conditions. At the larger distance threshold

Table 8-4: Average KPI values for delayed opening experiment with relative change compared to baseline (normal conditions)

KPI	Average Value	Relative Change (%)
Agent collisions (<0.1 m)	0.1680	+81.5 %
Agent collisions (<0.3 m)	1.0963	+67.5 %
Agent collisions (<0.5 m)	1.9949	+58.5 %
Agent collisions (<1.0 m)	7.1066	+29.2 %
Wall collisions	0	0.0 %
Speed distribution alpha	2.1090	+4.5 %
Speed distribution beta	3.0135	-1.1 %
Goal reached percentage	0.9590	+0.5 %

**Figure 8-4:** Agent trajectories for one of the delayed opening experiments.**Figure 8-5:** Waiting times for agents entering hall A in delayed opening experiment.

of 1.0 m, more frequent interactions are observed, which is consistent with moderate local crowding around the blocked entrance.

No wall collisions occur, confirming that agents respect the physical boundaries of the environment even under waiting conditions. The fitted parameters of the speed distribution show a slight deviation from the normal conditions experiment, indicating that the distribution is slightly more skewed toward lower speeds as can be seen from the increase in the parameter α and the decrease in the parameter β . Such a shift is intuitive, as the average walking speed of the agents is expected to be lower in this scenario due to the congestion caused by agents waiting in front of hall A.

Importantly, the goal reached percentage is very similar to that of the normal conditions experiment, showing that the temporary closure of Hall A does not prevent the agents from successfully reaching their destinations once the entrance becomes available. Furthermore, a plot of waiting times for the agents entering hall A against the simulation timestep is depicted in Figure 8-5. It can be seen that agents spawning earlier in the simulation undergo long waiting times and, as expected, these waiting times are significantly lower for agents spawning after hall A has opened.




Figure 8-6: Agent trajectories for one of the busy conditions experiments.

Overall, these results suggest that the model produces realistic dynamics under delayed accessibility conditions: queuing behavior leads to mild local congestion but without excessive collisions or boundary violations, and nearly all agents eventually reach their goals. This indicates that the potential-field framework can reproduce adaptive crowd behavior in scenarios where access is temporarily restricted.

8-3 Experiment 3: Busy Conditions

The third experiment, *busy conditions*, represents a high-demand scenario. Here, all six lecture halls (A–F) are open, and the expected number of agents is increased compared to the first two experiments from 250 to 450, of which 200 are leaving lecture hall A to move out of the environment via the main corridor. In this setup, agents are either entering lecture halls or traversing the corridor, but overall density levels are higher, producing more frequent interactions and potential congestion effects. Furthermore, an additional mechanism inducing local crowding effects arises from the interaction between opposing flows: agents attempting to enter Hall A are likely required to wait until agents exiting Hall A have completed their departure. The spawning probabilities for this experiment are depicted in Table 8-5.

The trajectories of one of the busy conditions simulations are depicted in Figure 8-6. The results of the busy conditions experiment, shown in Table 8-6, highlight the effects of increased congestion caused by the additional outflow of agents from lecture hall A. The most notable impact is observed in the collision metrics. Very short-range collisions (< 0.1 m) increase by more than a factor of three (+281%), while medium-range collisions (< 0.3 m and < 0.5 m) approximately double compared to the baseline experiment. Even at the larger threshold of < 1.0 m, collisions rise by 68.1%, confirming that dense conditions substantially elevate local interaction frequencies. Wall collisions remain absent, indicating that, although the additional density affects agent-to-agent interactions, boundary adherence is kept in this experiment.

The speed distribution parameters exhibit only minor deviations, with increases of 4.5% in α and 3.8% in β , suggesting that the overall velocity profile of the population is robust to crowding effects. Counterintuitively, the goal reached percentage increases slightly to 97.1%, compared to 95.4% under baseline conditions. This outcome can be attributed to the larger

Table 8-5: Absolute spawning probabilities of agents by type, x-location, and destination.

Agent Type	x-location	Destination	Absolute Probability
Walking	-14.8	-	$0.40 \times 0.60 = 0.24$
Walking	52	-	$0.40 \times 0.40 = 0.16$
Main Entrance	-14.8	Hall A	$0.45 \times 0.60 \times 0.325 = 0.08775$
Main Entrance	-14.8	Hall B	$0.45 \times 0.60 \times 0.225 = 0.06075$
Main Entrance	-14.8	Hall C	$0.45 \times 0.60 \times 0.15 = 0.0405$
Main Entrance	-14.8	Hall D	$0.45 \times 0.60 \times 0.15 = 0.0405$
Main Entrance	-14.8	Hall E	$0.45 \times 0.60 \times 0.075 = 0.02025$
Main Entrance	-14.8	Hall F	$0.45 \times 0.60 \times 0.075 = 0.02025$
Main Entrance	52	Hall A	$0.45 \times 0.40 \times 0.325 = 0.0585$
Main Entrance	52	Hall B	$0.45 \times 0.40 \times 0.225 = 0.0405$
Main Entrance	52	Hall C	$0.45 \times 0.40 \times 0.15 = 0.027$
Main Entrance	52	Hall D	$0.45 \times 0.40 \times 0.15 = 0.027$
Main Entrance	52	Hall E	$0.45 \times 0.40 \times 0.075 = 0.0135$
Main Entrance	52	Hall F	$0.45 \times 0.40 \times 0.075 = 0.0135$
Back Entrance	-14.8	Upper Back Entrance	$0.15 \times 0.60 \times 0.50 = 0.045$
Back Entrance	-14.8	Lower Back Entrance	$0.15 \times 0.60 \times 0.50 = 0.045$
Back Entrance	-14.8	Hall A	$0.15 \times 0.60 \times 0.325 = 0.02925$
Back Entrance	-14.8	Hall B	$0.15 \times 0.60 \times 0.225 = 0.02025$
Back Entrance	-14.8	Hall C	$0.15 \times 0.60 \times 0.15 = 0.0135$
Back Entrance	-14.8	Hall D	$0.15 \times 0.60 \times 0.15 = 0.0135$
Back Entrance	52	Upper Back Entrance	$0.15 \times 0.40 \times 0.50 = 0.03$
Back Entrance	52	Lower Back Entrance	$0.15 \times 0.40 \times 0.50 = 0.03$
Back Entrance	52	Hall A	$0.15 \times 0.40 \times 0.325 = 0.0195$
Back Entrance	52	Hall B	$0.15 \times 0.40 \times 0.225 = 0.0135$
Back Entrance	52	Hall C	$0.15 \times 0.40 \times 0.15 = 0.009$
Back Entrance	52	Hall D	$0.15 \times 0.40 \times 0.15 = 0.009$

number of agents entering from lecture hall A, who spawn into the environment towards the beginning of the simulation and thus are highly unlikely to not reach their goal within the simulation time. As a result, the completion rate increases despite the higher density and collision frequency.

Overall, the busy conditions experiment demonstrates that the model appropriately captures the nonlinear effects of congestion: higher collision frequencies emerge as density rises, while global efficiency indicators such as speed distribution and completion rates remain relatively stable or even improve due to changes in population composition.

8-4 Computation Time

Although the simulation model is capable of producing realistic results, as demonstrated in the previous sections of this chapter, a major drawback of this approach lies in its computational burden. Specifically, at each time step, the potential function must be recalculated for every agent, since agent positions change dynamically over time. While the static component of

Table 8-6: Average KPI values for busy conditions experiment with relative change compared to baseline (normal conditions)

KPI	Average Value	Relative Change (%)
Agent collisions (<0.1 m)	0.3528	+281.0 %
Agent collisions (<0.3 m)	1.8053	+175.7 %
Agent collisions (<0.5 m)	2.7380	+117.6 %
Agent collisions (<1.0 m)	9.2448	+68.1 %
Wall collisions	0	0.0 %
Speed distribution alpha	2.1092	+4.5 %
Speed distribution beta	3.1648	+3.8 %
Goal reached percentage	0.9708	+1.7 %

the potential function remains constant and therefore does not require recomputation, the dynamic component still results in substantial computational demands, leading to very long computation times, as shown in Table 8-7.

Table 8-7: Average computation times for all experiments

Experiment	Average Computation Time (minutes)
Normal conditions	61
Delayed opening	87
Busy conditions	132

Conclusion and Discussion

The research is concluded by answering the research questions as formulated in Chapter 1. Additionally, key limitations of the research are discussed and suggestions for future research are presented.

9-1 Research Findings

This study regards augmenting and optimizing an existing pedestrian movement model based on a potential field. Multiple experiments were performed to answer the research question: “To what extent can pedestrian movement in complex indoor environments be effectively modeled using a potential field model and calibrated using real-life trajectory data?” To answer this question, four subquestions were formulated.

- 1. What is the state-of-the-art for modeling pedestrian dynamics within a multi-agent system?*

The literature review showed that there are generally three types of pedestrian simulation models: macroscopic models, microscopic models and hybrid models. Macroscopic models are best used for applications in which computational costs are a constraint and behavioral realism on an individual level is not required. Microscopic models are fit for applications in which high fidelity is necessary at the cost of a large computational burden. Hybrid models attempt to exploit the strengths of both of these model types. However, hybrid models are generally designed for very specific applications or the advantages of the macro- and microscopic models are only partially implemented. There thus seems to be a research gap in the development of general-purpose hybrid pedestrian simulation models.

2. What makes a potential field model suitable for simulating pedestrian dynamics in complex environments?

Potential field models are suitable for simulating pedestrian dynamics because they provide a principled mathematical framework in which pedestrian behavior is represented as movement along the gradient of a potential function. This formulation enables the simultaneous modeling of attractive forces, such as the tendency of agents to move toward a goal, and repulsive forces, such as avoidance of other pedestrians and obstacles. Through gradient descent optimization, pedestrian trajectories naturally emerge from the minimization of the underlying potential landscape. In addition, as discussed in the review of previous work, potential field models can form a foundation for hybrid approaches that incorporate complementary modeling techniques.

3. How can a potential field model for pedestrian simulation be augmented and optimized using real-life trajectory data?

The potential field model was augmented and optimized through the integration of real-life trajectory data. By training a surrogate model and employing experimental datasets from bidirectional corridor and bottleneck scenarios, parameter calibration was systematically performed. This data-driven approach enhanced the realism of the simulations and also enabled a sensitivity analysis, highlighting the influence of parameters KG , KW , and KO on key performance indicators. The results demonstrate that coupling potential field models with empirical data constitutes a promising avenue for bridging the gap between theoretical models and observed pedestrian behavior.

4. What are the strengths and limitations of using a potential field model for pedestrian simulation?

Finally, the strengths and limitations of the potential field approach were critically assessed. On the one hand, the model succeeds in reproducing collective crowd phenomena and provides interpretable dynamics that align with empirical findings. It also offers flexibility and extensibility, allowing for adaptation to a variety of scenarios. On the other hand, a major limitation lies in its computational burden: the need to recompute potential functions for every agent at every time step renders large-scale simulations resource-intensive. Additionally, while the model captures macroscopic patterns effectively, certain aspects of microscopic pedestrian behavior may be oversimplified.

To what extent can pedestrian movement in complex indoor environments be effectively modeled using a potential field model and calibrated using real-life trajectory data?

This work has demonstrated that potential field models, when combined with real-world data and systematic optimization, constitute a valid and powerful framework for simulating pedestrian dynamics in complex environments. Nevertheless, their computational cost and inherent simplifications highlight the importance of ongoing efforts to balance realism, scalability, and efficiency in pedestrian simulation research.

9-2 Limitations and Recommendations

The most significant limitation of this research is the computational cost of the simulation model. Due to its slow performance, a surrogate model was introduced to facilitate the optimization process. Although the surrogate model provided accurate approximations of the simulation, its construction required the generation of a dataset by executing the simulation model over a wide range of parameter values. Because this procedure was computationally expensive, the parameter space could not be explored exhaustively. Consequently, the optimization was restricted to three parameters, namely the potential function weights, while all other parameters were assumed to be fixed at their default values. As a result, the optimization focused only on a subset of the full parameter space, which may limit the generality of the conclusions. For example, the obstacle potential function appears to underestimate repulsive effects, leading to an overestimation of close inter-agent encounters. This could for instance be improved by also optimizing over a number of different obstacle potential function types. Future work should therefore extend the optimization framework to include additional parameters once more computationally efficient implementations of the simulation model are available.

Another limitation of the current simulation model is that agents are unable to change their assigned goal destination once it has been assigned. This restricts the model's ability to capture adaptive route choice behavior, which is a central characteristic of pedestrians navigating complex and dynamic environments. In reality, individuals frequently adjust their trajectories in response to local crowding, congestion, or perceived alternative routes.

To overcome this limitation, future work could explore the integration of the presented potential field model within a hybrid modeling framework. In such a setup, the potential field formulation would serve as the microscopic layer, governing local pedestrian interactions through attractive and repulsive potential functions. At the same time, a macroscopic route choice algorithm could be implemented to dynamically determine agents' paths. For example, the environment could be represented as a network of nodes and edges, where the pedestrian density along each edge provides information about congestion levels. These congestion estimates could then influence route selection at the macroscopic level, enabling agents to adapt their decisions throughout the simulation.

This hybrid approach would combine the strengths of both modeling paradigms: the microscopic realism of potential fields in capturing inter-agent interactions and the macroscopic adaptability of network-based route choice. Such a framework would not only improve the realism of pedestrian simulations but also expand their applicability to scenarios involving dynamic environments.

Appendix A

Overview of Uncertainty Analysis

This is an overview of all plots for the uncertainty analysis of the surrogate model.




Figure A-1: Uncertainty Analysis for agent collisions (<0.1 m) KPI.




Figure A-2: Uncertainty Analysis for agent collisions (<0.3 m) KPI.




Figure A-3: Uncertainty Analysis for agent collisions (<0.5 m) KPI.




Figure A-4: Uncertainty Analysis for agent collisions (<1.0 m) KPI.




Figure A-5: Uncertainty Analysis for wall collisions KPI.

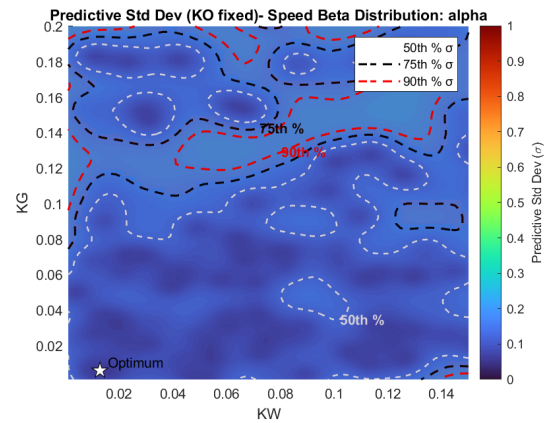


Figure A-6: Uncertainty Analysis for speed distribution alpha KPI.




Figure A-7: Uncertainty Analysis for speed distribution beta KPI.




Figure A-8: Uncertainty Analysis for paths y-axis distribution alpha KPI.




Figure A-9: Uncertainty Analysis for paths y-axis distribution beta KPI.




Figure A-10: Uncertainty Analysis for variance y-direction KPI.

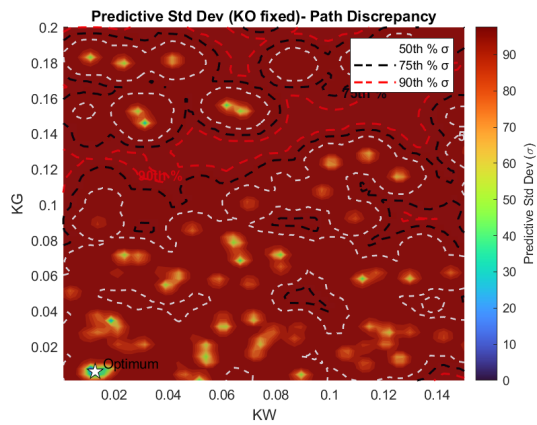


Figure A-11: Uncertainty Analysis for discrepancies KPI.




Figure A-12: Uncertainty Analysis for goal reached percentage KPI.

Appendix B

Overview of Parameter Sensitivity Analysis

This is an overview of all plots for the sensitivity analysis of the surrogate model.




Figure B-1: Sensitivity Analysis for agent collisions (<0.1 m) KPI.




Figure B-2: Sensitivity Analysis for agent collisions (<0.3 m) KPI.




Figure B-3: Sensitivity Analysis for agent collisions ($<0.5\text{ m}$) KPI.




Figure B-4: Sensitivity Analysis for agent collisions ($<1.0\text{ m}$) KPI.




Figure B-5: Sensitivity Analysis for wall collisions KPI.




Figure B-6: Sensitivity Analysis for speed distribution alpha KPI.




Figure B-7: Sensitivity Analysis for speed distribution beta KPI.




Figure B-8: Sensitivity Analysis for paths y-axis distribution alpha KPI.




Figure B-9: Sensitivity Analysis for paths y-axis distribution beta KPI.




Figure B-10: Sensitivity Analysis for variance y-direction KPI.




Figure B-11: Sensitivity Analysis for discrepancies KPI.




Figure B-12: Sensitivity Analysis for goal reached percentage KPI.

Appendix C

Scientific Paper

The scientific paper commences at the next page.

Pedestrian Movement Simulation through Augmented Potential Field Model

Jesse Leijdekker

Delft Center for Systems and Control, Delft University of Technology

Abstract—This research investigates the development and application of an augmented potential field model for simulating pedestrian dynamics in complex indoor environments. A comprehensive literature review identified a research gap in general-purpose hybrid pedestrian movement models. As a foundation for such an approach, a microscopic pedestrian dynamics model was developed, utilizing potential fields and gradient descent optimization for trajectory selection, thereby balancing goal-seeking behavior with obstacle and inter-agent avoidance.

To calibrate and validate the model, real-world trajectory data from bidirectional corridor and bottleneck experiments were utilized. A surrogate model was constructed to accelerate the optimization process, given the high computational cost of the original simulation model. The surrogate model enabled systematic parameter calibration and sensitivity analysis, focusing on three key parameters: the goal potential function weight (K_G), the wall potential function weight (K_W), and the obstacle potential function weight (K_O). The results demonstrated that the optimized model is capable of reproducing key crowd phenomena observed in the empirical datasets.

However, the research also identified significant limitations, primarily the computational burden of the simulation, which constrained parameter optimization to a subset of parameters and potentially led to an underestimation of obstacle repulsion, resulting in overly frequent close inter-agent encounters. Despite these challenges, the work establishes potential field models as a valid and powerful framework for pedestrian simulation, offering extensibility towards future hybrid models, while highlighting the need for more efficient implementations and broader parameter coverage in subsequent research.

Index Terms—Pedestrian Dynamics, Potential Field Model, Multi-Agent Systems, Gradient Descent Optimization, Surrogate Modeling, Gaussian Process Regression, Crowd Simulation.

I. INTRODUCTION

Understanding pedestrian dynamics is crucial for the effective planning, design, and operation of public spaces. As urban environments become more complex and densely populated, the need for accurate models and simulations of pedestrian movement has intensified across various domains, including transportation engineering, crowd safety, emergency evacuation planning, and smart infrastructure

development. Accurate modeling helps optimize circulation spaces, enhance safety and comfort, and inform data-driven decisions in infrastructure projects.

This research specifically focuses on pedestrian dynamics within a university environment, particularly the Mechanical Engineering faculty building at Delft University of Technology. With an anticipated increase in student population from approximately 25,000 to 40,000 [1], understanding pedestrian movement through corridors, lecture halls, and communal spaces is vital for designing efficient circulation routes and preventing congestion during peak transition periods. These insights are essential for designing future campus infrastructure and improving the overall user experience.

However, modeling pedestrian behavior is inherently complex due to its context-dependent, adaptive, and often non-linear nature. Pedestrians continuously adapt to both static environmental features and dynamic interactions with others, leading to emergent collective behaviors such as lane formation, bottlenecks, or stop-and-go waves, which are difficult to reproduce using simple rules or flow-based models.

A literature review identified three main categories of pedestrian dynamics models: macroscopic, microscopic, and hybrid models [2]. Macroscopic models are computationally efficient but sacrifice individual behavioral precision, as pedestrians are treated as continuous flows [3]–[6] and not as individuals. Microscopic models, on the other hand, offer higher fidelity by simulating individual agents separately but incur increased computational demands [3], [7]–[12]. Hybrid models aim to combine the advantages of both approaches, seeking a balance between simulation speed and behavioral realism [3]. However, most existing hybrid models are tailored to specific applications or only partially integrate the merits of their base approaches [13]–[15], revealing a research gap in the development of general-purpose hybrid pedestrian simulation models.

To address this gap, this study introduces a microscopic pedestrian dynamics model based on agent-based potential fields. In this framework, each pedestrian is modeled as an individual agent navigating a scalar potential landscape that encodes goals, obstacles, and interactions with other pedestrians. The agents determine their paths using a gradient descent optimization approach, allowing them to adaptively

Master of Science Thesis, September 15, 2025. This work was conducted at the Faculty of Mechanical Engineering (ME).

choose directions leading towards their goals while avoiding collisions. This approach leverages principles from classical optimization and control theory, offering a flexible and computationally tractable method for simulating realistic pedestrian behavior in complex indoor environments. The potential field approach is also discussed as a foundational step towards developing future hybrid models.

The primary research question guiding this research is: "To what extent can pedestrian movement in complex indoor environments be effectively modeled using a potential field model and calibrated using real-life trajectory data?". To answer this, four sub-questions were formulated:

- 1) What is the state-of-the-art for modeling pedestrian dynamics within a multi-agent system?
- 2) What makes a potential field model suitable for simulating pedestrian dynamics in complex environments?
- 3) How can a potential field model for pedestrian simulation be augmented and optimized using real-life trajectory data?
- 4) What are the strengths and limitations of using a potential field model for pedestrian simulation?

II. PROPOSED MODEL AND METHODOLOGY

The proposed model builds upon an earlier microscopic potential field-based approach [16]. In this model, individual agents, represented as single points in space, find their goals by optimizing a potential function using gradient descent optimization. The agent-specific goals are defined as the global minimum of this potential function. At each time step, agents take a step in the direction of the computed negative gradient. Furthermore, to prevent agents from getting trapped, intermediate goals are also defined along their trajectories.

The total potential function ($V_{\text{tot},i}$) for an agent i is a summation of three separate potential functions:

$$V_{\text{tot},i}(x_{i,t}, x_{j,t}, t) = \sum_{j \neq i}^{N(t)} VO_j(x_{j,t}) + VW(x_{i,t}) + VG_i(x_{i,t}, t) \quad (1)$$

where $x_{i,t} = [x_{i,t} \ y_{i,t}]^T \in \mathbb{R}^2$ is the coordinate vector for agent i at time t , and $N(t)$ is the number of agents present at time t . The obstacle potential function (VO) is designed to prevent collisions with both agents and dynamic or static obstacles. It defines a potential function that exhibits a maximum at the location of the obstacle, creating repulsion at that location:

$$VO_{i,j,t}(x_{i,t}, x_{o,j,t}) = \frac{KO_j}{\|x_{o,j,t} - x_{i,t}\|^2 + 1} \quad (2)$$

Here, $VO_{i,j}$ is the obstacle potential function computed for agent i , considering obstacle j , KO_j is the gain for obstacle j , $x_{o,j,t} = [x_{o,j,t} \ y_{o,j,t}]^T$ denotes the location, i.e.

coordinates, of obstacle j at time t . Note that the obstacle is treated as a single point in space.

The wall potential function (VW) is similar to the obstacle potential but employs a more elongated peak shape for each wall of the environment, guiding agents away from walls while still maintaining realistic movement patterns:

$$VW_k(x_{i,t}) = \left(\frac{1}{1 + e^{-c_{1,k} f_{1,k}(x_{i,t})}} + \frac{1}{1 + e^{-c_{2,k} f_{2,k}(x_{i,t})}} - 1 \right) \cdot \left(\frac{1}{1 + e^{-c_{3,k} f_{3,k}(x_{i,t})}} + \frac{1}{1 + e^{-c_{4,k} f_{4,k}(x_{i,t})}} - 1 \right) \quad (3)$$

Here, VW_k is the wall potential for wall k , $c_{i,k}$ is the gain controlling the steepness of the wall for the appurtenant side, $f_{1,k}(x, y)$ and $f_{2,k}(x, y)$ are the functions defining the location of the borders of the wall and $f_{3,k}(x, y)$ and $f_{4,k}(x, y)$ are the functions defining the location of the borders of the wall, generally perpendicular to $f_{1,k}(x, y)$ and $f_{2,k}(x, y)$.

Finally, the goal potential function (VG) attracts agents towards their designated goal location:

$$VG_i(x_{i,t}, x_{g,i,t}) = \frac{KG}{\|x_{i,t} - x_{g,i,t}\|} \quad (4)$$

Here, VG_i the goal potential for agent i , KG is the goal gain, $x_{g,i,t}$ denotes the goal location coordinates for agent i and at time t .

The movement of agent i at time t is thus determined by the following gradient descent algorithm:

$$x_{i,t+1} = x_{i,t} - \alpha_i \nabla V_{\text{tot},i}(x_{i,t}, x_{j,t}, t) \quad (5)$$

Here, α_i is an agent-specific parameter to control the step size.

A. Model Improvements

Model Improvements were implemented to enhance realism and computational efficiency. A significant performance enhancement was achieved through the parameterization of symbolic expressions. Initially, symbolic variables for potential fields were redefined and differentiated at every time step, causing inefficiency. This was resolved by computing static symbolic expressions (e.g., wall potentials) once outside the main simulation loop and converting them to fast numerical functions. Time-varying components (e.g., agent-specific obstacle potentials) were parameterized with agent-dependent coefficients, allowing reuse of a generic symbolic template. This restructuring reduced simulation

runtime by an order of magnitude. Furthermore, the wall potential function was simplified for straight walls from the complex form in Equation 3 to:

$$VW_k(\mathbf{x}_{i,t}) = \left(\frac{1}{1 + e^{-c_{1,k}f_{1,k}(\mathbf{x}_{i,t})}} + \frac{1}{1 + e^{-c_{2,k}f_{2,k}(\mathbf{x}_{i,t})}} - 1 \right) \cdot \left(\frac{1}{1 + c_{3,k}f_{3,k}(\mathbf{x}_{i,t})} \right) \quad (6)$$

where $f_3(x, y)$ is typically a linear function with very few terms, which significantly reduces the complexity of derivative operations and improves computational speed.

To address pedestrian behavior realism, two key adaptations were made:

- A speed-density relationship was introduced, based on the empirically observed phenomenon that free-flowing walking speeds of pedestrians decrease significantly as the density of people walking around them increases [17]. The model adopted an exponentially decreasing function:

$$v_i = \alpha_i e^{-0.015c_i} \quad (7)$$

Here, c_i is a measure of density, namely the number of agents within a three-meter radius of agent i . $\alpha_i \sim N(\mu_\alpha, \sigma_\alpha)$ is the agent-based desired free-flow speed scaling factor drawn from a Gaussian distribution with mean $\mu_\alpha = 40$ and standard deviation $\sigma_\alpha = 10$.

- The obstacle potential function was augmented to account for the decrease of interpersonal distances in high-density environments [18], thereby improving the realism of the model under crowded conditions. The augmented obstacle potential function is defined as:

$$VO_{i,j,t}(\mathbf{x}_{i,t}, \mathbf{x}_{o,j,t}) = \frac{KO}{0.05(1+0.1(5+c_j))\|\mathbf{x}_{o,j,t}-\mathbf{x}_{i,t}\|^2+0.25} \quad (8)$$

This formulation reduces the spatial reach of the repulsive influence from obstacles as local crowd density (c_j) rises, allowing agents to navigate closer to one another in congested areas, reflecting more natural behavior in congested environments.

B. Model Optimization

The model optimization process aimed to improve the realism of the simulation. Due to the high computational cost of running the full simulation, the optimization was limited to three key parameters: the goal potential function weight (KG), the wall potential function weight (KW), and the obstacle potential function weight (KO). Other parameters, such as time step size, speed distribution, and types of potential functions, were assumed to be sufficiently calibrated to manage the high-dimensional search space. The search space for these parameters was determined

as: $0.001 \leq KW \leq 0.15$, $0.001 \leq KG \leq 0.2$, and $0.0001 \leq KO \leq 0.015$.

A surrogate modeling strategy was chosen over brute force optimization or simple curve fitting due to computational burden of the full simulation model and the complex, non-linear relationships involved. A surrogate model provides a data-driven approximation that can predict outputs with significantly less computational cost than the actual model. The surrogate model was trained on data generated by running the simulation model in a bidirectional corridor experiment setup. This experiment setup will be explained in more detail in Section III. The input parameters used to run these simulations were determined by Latin Hypercube Sampling (LHS) of the input parameter space, ensuring a uniform and efficient sampling distribution. Gaussian Process Regression (GPR) was selected as the regression model for the surrogate due to its ability to provide predictions and quantify the confidence (uncertainty) of each prediction, which is valuable for optimization and understanding model reliability. The GPR model used a Matern-3/2 kernel, which is better suited for capturing localized variations in the noisy and nonlinear nature of pedestrian dynamics compared to over-smoothing kernels [19].

The surrogate model was then optimized using the Particle Swarm Optimization (PSO) algorithm. PSO is a population-based, stochastic method known for balancing exploration and exploitation in high-dimensional, non-convex objective functions where gradient information might be unavailable or unreliable [20]. The optimization minimized an objective function that quantified the difference between the surrogate model's predicted Key Performance Indicator (KPI) values and real-life KPI values, weighted by a covariance matrix to account for correlations between parameters. The KPIs were defined as follows:

- **Agent collisions <0.1 m, <0.3 m, <0.5 m, <1.0 m:** These four KPIs quantify the frequency of inter-agent proximity events at different distance thresholds. They capture not only hard collisions but also near misses and general spacing behavior, allowing the model to reflect realistic human tendencies to maintain personal space and avoid crowding.
- **Wall collisions:** This KPI measures the frequency with which agents collide with walls or static obstacles. It penalizes unrealistic behaviors such as agents walking into impenetrable boundaries, thereby ensuring that the simulated trajectories remain physically plausible.
- **Speed distribution parameters (α , β):** The speed distribution of agents is summarized by fitting a beta distribution to observed speed samples and extracting its α and β parameters. These parameters provide a compact statistical representation of pedestrian speed patterns and allow straightforward comparison between simulated and empirical data.

- **Lateral distribution parameters (α , β):** Since the corridor experiment involves straight-line motion along the x-axis, only the y-coordinates are relevant for lateral positioning. By fitting a beta distribution to the y-positions, the α and β parameters serve as robust indicators of how agents distribute themselves laterally within the corridor.
- **Variance in the y-direction:** This KPI quantifies the variability of agent trajectories along the lateral axis. Higher variance indicates irregular or unstable movement patterns, while lower variance corresponds to smoother, more controlled pedestrian motion as observed in real experiments.
- **Step-by-step positional discrepancy:** At each timestep, the model's predicted next position for each agent is compared to the actual next position observed in the empirical data. This discrepancy provides a fine-grained measure of how accurately the model replicates pedestrian decision-making at the individual level.
- **Goal reached percentage:** This KPI measures the proportion of agents that successfully reach their assigned goals. It serves as a sanity check to detect failure cases, such as agents becoming stuck in local minima of the potential field or blocked by unrealistic gradient configurations.

All sampling rate-dependent and agent count-dependent KPIs were normalized to ensure meaningful comparisons between simulated and empirical data, reflecting per-agent and per-second rates. The computational work was performed using MATLAB R2021b on a machine with a Ryzen 3700x CPU and 16GB of RAM.

III. EXPERIMENTAL SETUPS

Three experimental setups were used to test and validate the augmented potential field model: a bidirectional corridor, a bottleneck, and a lecture hall environment. Each setup targeted different aspects of pedestrian dynamics.

A. Bidirectional Corridor

The bidirectional corridor experiment is an earlier studied setup for pedestrian dynamics [21]. In this setup, two groups of agents walk in opposite directions through a straight corridor, resulting in head-on interactions and potential self-organized lane formation. The environment for this experiment is depicted in Figure 1.

Empirical data from controlled laboratory studies [21] were used for comparison. The corridor had a width of 4 meters, and agents were instructed to walk from one side to the other in a bidirectional manner. This setup is particularly suited to analyze interpersonal distances, collision avoidance, and the emergence of lane structures.




Fig. 1. Environment of the bidirectional corridor experiment.

B. Bottleneck

The bottleneck scenario represents a common source of congestion, where a crowd must pass through a narrow opening [22]. The setup consisted of a corridor section leading to a constriction of 0.9 m width. This configuration created high-density interactions and queuing behavior upstream of the bottleneck.

The bottleneck experiment tests the model's ability to reproduce throughput, congestion buildup, and close-range interactions under pressure. Real-world trajectory data from laboratory bottleneck experiments served as reference [22].




Fig. 2. Environment of the bottleneck experiment.

C. Lecture Hall Environment

To assess the model's applicability in a realistic and complex indoor setting relevant for a university setting, a lecture hall environment was designed. The setup includes a central corridor with six adjacent lecture halls with a

probability of there being static obstacles representing study areas and furniture as depicted in Figure 3.




Fig. 3. Environment for the lecture hall experiments, with indicated hall names and study areas.

Three scenarios were simulated: (i) normal conditions with four lecture halls open, (ii) delayed opening of Hall A, and (iii) busy conditions with all lecture halls active and a larger agent population together with agents leaving Hall A at the beginning of the simulation. Each scenario simulated a 15-minute break period between lectures, being the period in which students and teachers enter the lecture halls. The agents' spawning location, -time and end goal were determined according to a probability distribution.

The lecture hall setup highlights how the model performs in multi-room environments with dynamic demand, congestion, and multiple entry- and exit points.

IV. RESULTS

A. Optimization and Surrogate Validation

The surrogate-based optimization yielded final parameters of $K_W = 0.0128$, $K_G = 0.0060$, and $K_O = 0.0014$, reflecting a balance between goal pursuit, wall avoidance, and inter-agent repulsion. The Gaussian Process surrogate accurately reproduced the behavior of the full potential field model, achieving Coefficient of Determination (R^2) values above 0.9 for most KPIs as depicted in Table I. As for the comparison between the surrogate model and the simulation model for the bidirectional corridor experiment using the optimal parameters, minor deviations occurred in the prediction of lateral trajectory variance and medium-range interactions, as can be seen in Table II.

B. Bidirectional Corridor Experiment

The bidirectional corridor experiment served as the primary test case for evaluating model realism against empirical

TABLE I
 R^2 VALUES FOR ALL SURROGATE MODELS

KPI	R^2
Agent collisions (<0.1 m)	0.9298
Agent collisions (<0.3 m)	0.9917
Agent collisions (<0.5 m)	0.9883
Agent collisions (<1.0 m)	0.9867
Wall collisions	0.8129
Speed distribution alpha	0.9184
Speed distribution beta	0.9284
Paths y-axis distribution alpha	0.9035
Paths y-axis distribution beta	0.7891
Variance y-direction	0.9840
Discrepancies	1.00
Goal reached percentage	0.9957

TABLE II
KPI VALUES FOR SIMULATION MODEL AND SURROGATE MODEL AT THE OPTIMUM

KPI	Simulation Model	Surrogate Model
Agent collisions (<0.1 m)	0.0823	0.0850
Agent collisions (<0.3 m)	1.7215	1.7335
Agent collisions (<0.5 m)	4.3509	4.7346
Agent collisions (<1.0 m)	15.6158	19.5395
Wall collisions	0	7.5813e-04
Speed distribution alpha	0.8199	0.7955
Speed distribution beta	20.3820	20.2791
Paths y-axis distribution alpha	1.1252	1.2179
Paths y-axis distribution beta	1.1736	1.2724
Variance y-direction	1.3246	1.5661
Discrepancies	2.2124e+05	2.1420e+05
Goal reached percentage	0.9576	0.9086

trajectory data. The model successfully reproduced medium-range collision avoidance and general locomotion patterns. However, it overestimated the frequency of very close interpersonal encounters (< 0.1 m and < 0.3 m) and generated higher lateral variance compared to real observations. The complete comparison of the KPIs outputted by the simulation model and the real-life data is depicted in Table III.

TABLE III
KPI VALUES FOR SIMULATION MODEL AND REAL-LIFE DATA FOR BIDIRECTIONAL CORRIDOR EXPERIMENT

KPI	Simulation Model	Real-Life Data
Agent collisions (<0.1 m)	0.0823	0
Agent collisions (<0.3 m)	1.7215	0.3273
Agent collisions (<0.5 m)	4.3509	4.138
Agent collisions (<1.0 m)	15.6158	33.7264
Wall collisions	0	0
Speed distribution alpha	0.8199	0.7048
Speed distribution beta	20.3820	20.2518
Paths y-axis distribution alpha	1.1950	1.3777
Paths y-axis distribution beta	1.2696	1.7309
Variance y-direction	1.3246	0.4993
Goal reached percentage	0.9576	0.9432

C. Lane Formation

Lane formation, a hallmark of bidirectional flow, emerged in both the empirical and simulated datasets of the bidirectional corridor experiment. The simulation captured the general tendency toward self-organized lanes, though the simulated lanes were less distinct and slightly more overlapping than in the empirical data as is shown in Table IV.

TABLE IV
LANE FORMING METRICS FOR SIMULATION MODEL AND REAL-LIFE DATA FOR BIDIRECTIONAL CORRIDOR EXPERIMENT

Lane Forming Metric	Simulation Model	Real-Life Data
Average overlap coefficient	0.317	0.215
Average zero-lag correlation	-0.332	-0.708
Average Peak offset (m)	0.16	0.20

D. Bottleneck Experiment

In the bottleneck scenario [22], the model reproduced the qualitative difficulty of passage and congestion buildup. Compared to real-world data, the simulation produced higher average throughput than the real-life experiments (85 vs 69), but also significantly more short-range encounters, just as in the bidirectional corridor experiment, indicating that agents passed more efficiently yet with reduced interpersonal spacing. The comparison of KPIs for the bottleneck experiment is shown in Table V.

TABLE V
KPI VALUES FOR SIMULATION MODEL AND REAL-LIFE DATA FOR BOTTLENECK EXPERIMENT

KPI	Simulation Model	Real-Life Data
Agent collisions (<0.1 m)	1.4249	0
Agent collisions (<0.3 m)	9.6964	0.6202
Agent collisions (<0.5 m)	18.2232	15.8779
Agent collisions (<1.0 m)	69.1577	61.7155
Wall collisions	0	0
Speed distribution alpha	1.3858	1.4310
Speed distribution beta	4.8795	12.7046
Variance y-direction	2.9129	0.4993
Goal reached percentage	0.3348	0.4471

E. Application in Lecture Hall Environment

To assess the applicability of the model in a more complex environment, three exploratory experiments were conducted in a lecture hall setting consisting of a central corridor and six lecture halls. Agents either traversed the corridor or entered one of the lecture halls, with spawning probabilities drawn from a normal distribution peaking five minutes before the lecture start.

1) *Normal Conditions*: In the baseline scenario, four lecture halls (A–D) were open. The model produced smooth pedestrian flows with few close encounters and no wall collisions. Speed distributions were stable and nearly all agents reached their goals, indicating realistic behavior under free-flow conditions.

2) *Delayed Opening*: In this experiment, lecture hall A remained closed for the first half of the simulation. This created temporary queuing and moderate local congestion near the blocked entrance, reflected in increased collision frequencies across all thresholds (see Table VI). However, wall collisions were absent and completion rates remained high, confirming that agents adapt realistically to delayed accessibility.

3) *Busy Conditions*: Here, all six lecture halls were open and the number of agents was increased from 250 to 450, including 200 agents exiting hall A into the corridor. As expected, collisions rose sharply (up to +281% at <0.1 m) due to crowding effects, but wall collisions remained zero and speed distributions were only slightly affected. Interestingly, the goal-reached percentage increased, since many agents exited hall A early in the simulation and thus completed their trajectories within the available time.

4) *Computation Time*: Across all lecture hall experiments, computation time was significant due to the need to update dynamic potentials at each time step. Average runtimes ranged from 61 minutes (normal conditions) to 132 minutes (busy conditions), highlighting a major drawback of the approach.

V. CONCLUSION AND DISCUSSION

This work has shown that potential field models, when systematically calibrated with empirical data and optimized using surrogate modeling, provide a viable framework for simulating pedestrian dynamics in complex environments. The model successfully reproduced key collective phenomena such as lane formation in bidirectional flows and congestion effects in bottleneck scenarios, and demonstrated realistic behavior in a lecture hall environment under varying usage conditions. These findings highlight both the potential and the current limitations of the approach.

A major strength of the model lies in its interpretability: the balance of attractive and repulsive forces offers a transparent representation of pedestrian decision-making, and systematic tuning of key parameters (K_G , K_W , K_O) proved effective in aligning the simulation with real-world data. Furthermore, the experiments in the lecture hall environment illustrated how the framework can be applied to study practical scenarios such as delayed accessibility and crowding during peak demand.

At the same time, several limitations were identified. The computational burden of recalculating dynamic potentials at each time step remains a significant drawback, limiting

TABLE VI
AVERAGE KPI VALUES FOR THE THREE LECTURE HALL EXPERIMENTS. RELATIVE CHANGE (%) FOR DELAYED OPENING AND BUSY CONDITIONS IS SHOWN COMPARED TO BASELINE (NORMAL CONDITIONS).

KPI	Normal Conditions Value	Delayed Opening		Busy Conditions	
		Value	Rel. Change	Value	Rel. Change
Agent collisions (<0.1 m)	0.0926	0.1680	+81.5%	0.3528	+281.0%
Agent collisions (<0.3 m)	0.6546	1.0963	+67.5%	1.8053	+175.7%
Agent collisions (<0.5 m)	1.2586	1.9949	+58.5%	2.7380	+117.6%
Agent collisions (<1.0 m)	5.4990	7.1066	+29.2%	9.2448	+68.1%
Wall collisions	0	0	0.0%	0	0.0%
Speed distribution α	2.0190	2.1090	+4.5%	2.1092	+4.5%
Speed distribution β	3.0480	3.0135	-1.1%	3.1648	+3.8%
Goal reached percentage	0.9542	0.9590	+0.5%	0.9708	+1.7%

the parameter space that can be explored and restricting optimization to a small subset of variables. Moreover, while the model captures short- and medium-range interactions, it tends to overestimate close interpersonal distances under crowded conditions. Finally, the current formulation assumes fixed destinations, limiting the ability to reproduce adaptive route choice behavior that is characteristic of real pedestrians.

A. Recommendations

As stated, a key limitation of this work is the high computational cost of the simulation model, which required the use of a surrogate model to enable optimization. While effective, this approach restricted optimization to three potential function weights, leaving other parameters fixed and potentially limiting the generality of the results. For example, the obstacle potential function tended to underestimate repulsion, leading to an overestimation of close encounters. Future research should extend the optimization framework to a broader parameter set, ideally supported by more computationally efficient implementations.

Furthermore, incorporating adaptive destination and route choice mechanisms would be another improvement to the model. A promising direction is to embed the potential field formulation within a hybrid framework, where microscopic interactions are governed by potential fields and macroscopic route choice adapts dynamically to congestion levels.

In summary, potential field models offer a promising foundation for hybrid pedestrian modeling approaches. With further refinement to improve computational efficiency and behavioral flexibility, they can provide valuable insights into pedestrian movement in complex real-world environments, such as university buildings undergoing expansion.

REFERENCES

- [1] S. Bongers, "TU Delft wants to grow to 40 thousand students," 2022. [Online]. Available: <https://delta.tudelft.nl/en/article/tu-delft-wants-grow-40-thousand-students>
- [2] J. Leijdekker, "Pedestrian movement simulation," 2024.
- [3] D. C. Duives, W. Daamen, and S. P. Hoogendoorn, "State-of-the-art crowd motion simulation models," *Transportation Research Part C: Emerging Technologies*, vol. 37, pp. 193–209, 2013.
- [4] R. Hughes, "The flow of large crowds of pedestrians," *Mathematics and Computers in Simulation*, vol. 53, no. 4, pp. 367–370, 2000.
- [5] R. L. Hughes, "A continuum theory for the flow of pedestrians," *Transportation Research Part B: Methodological*, vol. 36, no. 6, pp. 507–535, 2002.
- [6] A. Treuille, S. Cooper, and Z. Popović, "Continuum crowds," *ACM Trans. Graph.*, vol. 25, no. 3, p. 1160–1168, 2006.
- [7] G. N. Gilbert and K. G. Troitzsch, *Simulation for the social scientist*. Open University Press, 2005.
- [8] S. Jamshidi, M. Ensafi, and D. Pati, "Wayfinding in interior environments: An integrative review," *Frontiers in Psychology*, vol. 11, 2020.
- [9] C. Burstedde, K. Klauck, A. Schadschneider, and J. Zittartz, "Simulation of pedestrian dynamics using a two-dimensional cellular automaton," *Physica A: Statistical Mechanics and its Applications*, vol. 295, no. 3, pp. 507–525, 2001.
- [10] A. Kirchner and A. Schadschneider, "Simulation of evacuation processes using a bionics-inspired cellular automaton model for pedestrian dynamics," *Physica A: Statistical Mechanics and its Applications*, vol. 312, no. 1, pp. 260–276, 2002.
- [11] D. Helbing, I. Farkas, and T. Vicsek, "Simulating dynamic features of escape panic," *Nature*, vol. 407, pp. 487–490, 2000.
- [12] P. Fiorini and Z. Shiller, "Motion planning in dynamic environments using velocity obstacles," *The International Journal of Robotics Research*, vol. 17, no. 7, pp. 760–772, 1998.
- [13] M. Xiong, W. Cai, S. Zhou, M. Y. H. Low, F. Tian, D. Chen, D. Ong, and B. Hamilton, "A case study of multi-resolution modeling for crowd simulation," in *Proceedings of the 2009 Spring Simulation Multiconference*, 2009.
- [14] M. Xiong, M. Lees, W. Cai, S. Zhou, and M. Y. H. Low, "Hybrid modelling of crowd simulation," *Procedia Computer Science*, vol. 1, no. 1, pp. 57–65, 2010, iCCS 2010.
- [15] B. Banerjee, A. Abukmail, and L. Kraemer, "Advancing the layered approach to agent-based crowd simulation," in *2008 22nd Workshop on Principles of Advanced and Distributed Simulation*, 2008, pp. 185–192.
- [16] J. Leijdekker, "Pedestrian movement simulation through potential field model," 2024.
- [17] M. Nazir, K. Razi, Q. Hossain, and S. Adhikary, "Pedestrian flow characteristics at walkways in rajshahi metropolitan city of bangladesh," in *2nd International Conference on Civil Engineering for Sustainable Development*, 2014.
- [18] A. Gorrini, S. Bandini, and M. Sarvi, "Group dynamics in pedestrian crowds estimating proxemic behavior," *Transportation Research Record Journal of the Transportation Research Board*, vol. 2421, pp. 51–56, 2014.
- [19] M. L. Stein, *Interpolation of Spatial Data: Some Theory for Kriging*. Springer, 1999.
- [20] J. Kennedy and R. Eberhart, "Particle swarm optimization," in *Proceedings of ICNN'95 - International Conference on Neural Networks*, ser. ICNN-95, vol. 4. IEEE, 1995, p. 1942–1948. [Online]. Available: <http://dx.doi.org/10.1109/ICNN.1995.488968>

-
- [21] Forschungszentrum Jülich, University Of Wuppertal, University Of Siegen, S. Holl, M. Boltes, W. Mehner, and A. Seyfried, "Corridor, bidirectional flow," 2013. [Online]. Available: <http://ped.fz-juelich.de/da/2013bidirectional>
- [22] Forschungszentrum Jülich, University of Wuppertal, Cologne University, Boltes, Maik, Seyfried, Armin, Schadschneider, Andreas, and Winkens, Andreas, "Bottleneck, caserne," 2006. [Online]. Available: <http://ped.fz-juelich.de/da/2006bottleneck>

Bibliography

- [1] A. C. Atkinson and A. N. Donev. *Optimum Experimental Designs*. Oxford University Press, Oxford, UK, 1992.
- [2] B. Banerjee, A. Abukmail, and L. Kraemer. Advancing the layered approach to agent-based crowd simulation. In *2008 22nd Workshop on Principles of Advanced and Distributed Simulation*, pages 185–192, 2008.
- [3] S. Bongers. Tu delft wants to grow to 40 thousand students, 2022.
- [4] M. D. Buhmann. *Radial Basis Functions: Theory and Implementations*. Cambridge University Press, 2003.
- [5] C. Burstedde, K. Klauck, A. Schadschneider, and J. Zittartz. Simulation of pedestrian dynamics using a two-dimensional cellular automaton. *Physica A: Statistical Mechanics and its Applications*, 295(3):507–525, 2001.
- [6] D. Chicco, M. J. Warrens, and G. Jurman. The coefficient of determination r-squared is more informative than smape, mae, mape, mse and rmse in regression analysis evaluation. *PeerJ Computer Science*, 7:e623, 2021.
- [7] D. C. Duives, W. Daamen, and S. P. Hoogendoorn. State-of-the-art crowd motion simulation models. *Transportation Research Part C: Emerging Technologies*, 37:193–209, 2013.
- [8] P. Fiorini and Z. Shiller. Motion planning in dynamic environments using velocity obstacles. *The International Journal of Robotics Research*, 17(7):760–772, 1998.
- [9] G. Flötteröd and G. Lämmel. Bidirectional pedestrian fundamental diagram. *Transportation Research Part B: Methodological*, 71:194–212, 2015.
- [10] Forschungszentrum Jülich, University of Wuppertal, Cologne University, Boltes, Maik, Seyfried, Armin, Schadschneider, Andreas, and Winkens, Andreas. Bottleneck, caserne, 2006.

- [11] Forschungszentrum Jülich, University Of Wuppertal, University Of Siegen, S. Holl, M. Boltes, W. Mehner, and A. Seyfried. Corridor, bidirectional flow, 2013.
- [12] P. I. Frazier. A tutorial on bayesian optimization, 2018.
- [13] G. N. Gilbert and K. G. Troitzsch. *Simulation for the social scientist*. Open University Press, 2005.
- [14] I. Goodfellow, Y. Bengio, and A. Courville. *Deep Learning*. MIT Press, 2016.
- [15] A. Gorrini, S. Bandini, and M. Sarvi. Group dynamics in pedestrian crowds estimating proxemic behavior. *Transportation Research Record Journal of the Transportation Research Board*, 2421:51–56, 2014.
- [16] D. Helbing, I. Farkas, and T. Vicsek. Simulating dynamic features of escape panic. *Nature*, 407:487–490, 2000.
- [17] D. Helbing, P. Molnár, I. J. Farkas, and K. Bolay. Self-organized pedestrian crowd dynamics: Experiments, simulations, and design solutions. *Transportation science*, 39(1):1–24, 2001.
- [18] S. Hoogendoorn, P. Bovy, and W. Daamen. Microscopic pedestrian wayfinding and dynamics modelling. In M. Schreckenberg and S. Sharma, editors, *Pedestrian and evacuation dynamics*, pages 123–154. Springer, 2001.
- [19] R. Hughes. The flow of large crowds of pedestrians. *Mathematics and Computers in Simulation*, 53(4):367–370, 2000.
- [20] R. Hughes. The flow of large crowds of pedestrians. *Mathematics and Computers in Simulation*, 53(4):367–370, 2000.
- [21] R. L. Hughes. A continuum theory for the flow of pedestrians. *Transportation Research Part B: Methodological*, 36(6):507–535, 2002.
- [22] S. Jamshidi, M. Ensafi, and D. Pati. Wayfinding in interior environments: An integrative review. *Frontiers in Psychology*, 11, 2020.
- [23] J. Kennedy and R. Eberhart. Particle swarm optimization. In *Proceedings of ICNN’95 - International Conference on Neural Networks*, volume 4 of *ICNN-95*, page 1942–1948. IEEE, 1995.
- [24] A. Kirchner and A. Schadschneider. Simulation of evacuation processes using a bionics-inspired cellular automaton model for pedestrian dynamics. *Physica A: Statistical Mechanics and its Applications*, 312(1):260–276, 2002.
- [25] J. Leijdekker. Pedestrian movement simulation, 2024.
- [26] J. Leijdekker. Pedestrian movement simulation through potential field model, 2024.
- [27] M. D. McKay, R. J. Beckman, and W. J. Conover. A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics*, 21(2):239–245, 1979.

-
- [28] D. C. Montgomery. *Design and Analysis of Experiments*. Wiley, Hoboken, NJ, 9th edition, 2017.
 - [29] M. Nazir, K. Razi, Q. Hossain, and S. Adhikary. Pedestrian flow characteristics at walkways in rajshahi metropolitan city of bangladesh. In *2nd International Conference on Civil Engineering for Sustainable Development*, 2014.
 - [30] C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning*. MIT Press, 2006.
 - [31] W. S. Correlation and causation. *Journal of agricultural research*, 20(7):557, 1921.
 - [32] S. Sengupta, S. Basak, and R. Peters. Particle swarm optimization: A survey of historical and recent developments with hybridization perspectives. *Machine Learning and Knowledge Extraction*, 1(1):157–191, Oct. 2018.
 - [33] D. Stauffer. Cellular automata: Applications. In J. M. L. M. Palma, J. Dongarra, and V. Hernández, editors, *Vector and Parallel Processing — VECPAR 2000*, pages 199–206, Berlin, Heidelberg, 2001. Springer Berlin Heidelberg.
 - [34] M. L. Stein. *Interpolation of Spatial Data: Some Theory for Kriging*. Springer, 1999.
 - [35] A. Treuille, S. Cooper, and Z. Popović. Continuum crowds. *ACM Trans. Graph.*, 25(3):1160–1168, 2006.
 - [36] U. Weidmann. *Transporttechnik der Fußgänger*. ETH Zürich, 1993.
 - [37] M. Xiong, W. Cai, S. Zhou, M. Y. H. Low, F. Tian, D. Chen, D. Ong, and B. Hamilton. A case study of multi-resolution modeling for crowd simulation. In *Proceedings of the 2009 Spring Simulation Multiconference*, 2009.
 - [38] M. Xiong, M. Lees, W. Cai, S. Zhou, and M. Y. H. Low. Hybrid modelling of crowd simulation. *Procedia Computer Science*, 1(1):57–65, 2010. ICCS 2010.

