



**Health, demographic change and wellbeing
Personalising health and care: Advancing active and healthy ageing
H2020-PHC-19-2014
Research and Innovation Action**



Deliverable 2.2
Human motion models

Deliverable due date: 01.02.2017	Actual submission date: 06.03.2017
Start date of project: February 1, 2015	Duration: 42 months
Lead beneficiary for this deliverable: UNITN	Revision: 1
Authors: Daniele Fontanelli (UNITN), Francesco Farina (UNISI), Andrea Garulli (UNISI) Antonello Giannitrapani (UNISI), Sean Sedwards (INRIA), Thomas Given (INRIA)	
Internal reviewer: Luigi Palopoli (UNITN)	

The research leading to these results has received funding from the European Union's H2020 Research and Innovation Programme - Societal Challenge 1 (DG CONNECT/H) under grant agreement n° 643644.		
Dissemination Level		
PU	Public	X
CO	Confidential, only for members of the consortium (including the Commission Services)	

The contents of this deliverable reflect only the authors' views and the European Union is not liable for any use that may be made of the information contained therein.

Contents

Executive Summary	3
1 Introduction	7
2 Beyond the Social Force Model	9
2.1 The Social Force Model	10
2.2 The Non-Holonomic Nature of Human Locomotion	11
2.3 The Headed Social Force Model	11
2.3.1 Force Inputs	14
2.3.2 Torque Input	14
2.4 Enforcing Group Cohesion in the HSFM	15
3 HSFM Simulations	17
3.1 HSFM Behaviour	17
3.1.1 Nonholonomic behaviour	17
3.1.2 Blending between SFM and Nonholonomic behaviour	20
3.2 Group Cohesion Forces	24
3.3 Model Parameters Tuning	25
3.3.1 Force input	25
3.3.2 Torque input	27
3.3.3 Group cohesion	27
3.4 Comparison with the SFM at the Museum	31
3.5 Comparison with Behavioural Dynamics	34
3.6 Comments	37
4 On-the-fly Group Abstraction	39
4.1 Background	39
4.1.1 Trace inference	41
4.1.2 Group abstraction	41
4.2 Modelling real pedestrians with Group Abstraction	44
4.2.1 Dataset	44
4.2.2 Distance metrics	45
4.2.3 Results	45
5 Relation with the other Work Packages	49

Executive Summary

The deliverable D2.2 “Human motion models”, due at month 24, presents the research activities concerning the mathematical models for human locomotion in crowds and carried out within Task T2.1 “Models for social interaction”. It focuses primarily on Objective 2.1 of WP 2: “identify behavioural patterns that emerge when a group of people moves in shared space and describe them by appropriate mathematical models”. Indeed, a key component of the *FriWalk* developed in the ACANTO project is its ability to move in shared, crowded environments avoiding collisions, respecting safety limits and obeying to unwritten social rules while moving towards the desired destination. To this end, the Activity Planner and, more importantly, the Reactive Planner (both developed in WP5 - “Execution Support of Social Activities”) have to rely on reliable human motion models to carry out predictions on future human motion intentions. Similarly, both the collaborative localisation and the interpretation of social context developed within WP3 - “Perception of Users and Environment” benefit of an effective model of social motions for individuals and groups. Finally, motion patterns have also an important impact on the manoeuvres the *FriWalk* has to generate, hence on the work carried out in WP6 - “Design of Robotic Personal Devices”.

In this deliverable the new model of human locomotion proposed within ACANTO is presented. Our solution, called Headed Social Force Model, relies upon the broadly known Social Force Model and solves some of its descriptive limitations, as shown by simulations in this deliverable. Observations made on actual experiments with human beings and reported in D2.1 have been used to validate the proposed model.

To aid the efficiency of group motion planning, we have developed the notion of *group abstraction* to simplify the models of large numbers of pedestrians. In this deliverable we present the successful results of applying on-the-fly trace inference and group abstraction to observations of real pedestrians in a crowded environment.

Chapter 1

Introduction

In this deliverable, we present the definition of a new mathematical model for human locomotion, which stems from the analysis carried out in D2.1 - “Human motion models (preliminary)”, in which a taxonomy of the different approaches proposed in the literature, inspired to the survey [62], has been presented and deeply detailed. Starting from the historical perspective of building evacuation dynamics in both emergency and normal situations, in which the models are based on macroscopic quantities, such as densities and fluids [23], the analysis in D2.1 focused on the models of relevance for the navigation problems of the *FriWalk*, i.e. where the interactions are not so frequent as in overcrowded evacuation dynamics. In this framework, microscopic models of pedestrians are preferable. In such a case, the proposed approaches can be roughly categorised into four main classes: cellular automata [7], agent-based models [63], graph-based methods [6] and social force models [22]. Cellular automata are represented by a discrete system evolving on a discrete set of cells at discrete time intervals. The value of each cell depends on the modelled behaviour of the agent occupying it, on the neighbouring cell values and on a set of local updating rules (e.g., see [47, 18, 50]). Agent-based approaches model the active and reactive behaviours of the pedestrians according to stochastic models. In this framework, constant velocity models have received large attention since they are easily tractable and allow the direct use of Kalman filters for predictions and belief computations (e.g., see [12, 32]). In graph-based approaches, the environment is subdivided into regions using empirical observations and learning algorithms. The regions are usually mapped as nodes on the graph, while the paths joining them are the arcs. The nodes are usually considered as places in the environment of particular interest, where people stop or make decisions (e.g., see [34, 6]).

In the rather broad set of microscopic solutions, the Social Force Model (SFM) turns out to fit best to the project purposes. Basically, such a model assimilates each individual to a particle subject to social forces which drive her during the motion [22]. A recognised plus of this model is its possibility to describe group of people walking together while maintaining social interactions [10, 57]. An outcome of the D2.1 - “Human motion models (preliminary)”, highlighted also by the empirical evidence of the simulated visit at the museum, was the lack of an appropriate kinematic model for the human locomotion embedded in the SFM framework to generate more realistic trajectories (e.g., alternating forward motions and turnings). As a consequence, an in-depth study has been carried also to properly select additional force terms accounting for group behaviours.

The main outcome of the study summarised in this deliverable was the conception of the Headed Social Force Model (HSFM), which is an enhanced version of the traditional SFM that explicitly accounts for the pedestrians’ heading. To this end, we describe the motion of each individual by means of a dynamic model similar to that adopted in [40] for generating biologically-inspired robot trajectories. The main contribution to the literature in this field is the introduction of suitable model inputs (i.e., forces and torques which drive the dynamics of each pedestrian), with the purpose of maximising the realism of the resulting human trajectories. In doing so, several conflicting objectives have to be taken into account. In low density scenarios, the pedestrians’ motion should be as smooth as possible, consistently with what is observed in practice [45]. Similarly, in these circumstances, lateral motions should be avoided since individuals walk ahead most of the time. On the

contrary, in crowded or cluttered environments, the interaction among pedestrians, as well as between pedestrians and the environment, is stronger and determines most of the pedestrians' behaviour. The adopted solution consists in computing the model inputs as suitable functions of the various force terms in the traditional SFM. Additionally, a new optional force contribution is introduced in order to account for people walking together as a single group. It is shown that considering the pedestrians' heading enhances the fidelity of the model in two ways. Whenever nonholonomic motion patterns naturally arise, the generated trajectories resemble more closely those empirically observed. Typical examples include people walking in open spaces or reaching close targets. More generally, accounting explicitly for the pedestrians' heading helps to increase the regularity of the trajectories, resulting in fewer abrupt changes of direction and a reduced number of collisions. In this deliverable we present the ideas underlying the HSFM and we show that the SFM can be naturally retrieved when the environment becomes very cluttered, due to the presence of obstacles or of an overcrowded area. The proposed blend between the SFM and a more regular trajectory generator addresses an issue highlighted in D2.1 - "Human motion models (preliminary)", that is the necessity to switch between different behaviours according to the users' surroundings. Furthermore, it is worthwhile to note that within the ACANTO project, this is extremely beneficial since the model can be used to predict the motion of the human beings in the surroundings but also as a reference model to *generate* the path for the *FriWalk*, which is intrinsically nonholonomic.

The performance of the HSFM is evaluated via numerical simulation under very different operating conditions, and a sensitivity analysis of the model behaviour with respect to variations in the model parameters is presented. As a byproduct, guidelines on the selection of the parameter values are obtained.

The potentiality of the HSFM has been tested in the same simulations set-up used in D2.1 - "Human motion models (preliminary)" to verify if the proposed model solves some of the descriptive issues of the SFM. Moreover, the HSFM has been tested in a simulation environment mimicking the experimental set-up adopted in D2.1.

We make intensive use of the SFM in our reactive motion planner, developed by WP5, to predict the likely outcome of different actions (directions of motion) of the *FriWalk* users. Our group motion planning problem potentially scales exponentially with the number of participants in an activity and the number of other pedestrians in the vicinity. It is well known, however, that pedestrians often travel in small social groups that maintain their formation in the short to medium term [25]. We observe that it is therefore not necessary or desirable to "micromanage" the motion of *FriWalk* users in a group, and better to model the group as a single agent. This approach significantly eases the computational burden and allows the pedestrians to manage their own social interactions.

In Chapter 4 we present our on-the-fly trace inference and group abstraction algorithms. We first describe our motion planning approach and thus motivate the notion of modelling pedestrian behaviour at the level of groups. We then present our algorithms in detail and visualise the results of applying them to the ETH Zürich BIWI Walking Pedestrians dataset.¹

The deliverable is structured as follows. In Chapter 2 the HSFM is introduced and discussed in details. In Chapter 3 the simulation results of the HSFM in the same scenarios of the D2.1 are presented. In the same chapter, scenarios mimicking the behavioural dynamics experiments presented in D2.1 are also reported. Chapter 4 first recalls the motivation of group abstraction, then describes the algorithms we have developed and finally presents the successful results of applying them to a standard walking pedestrian dataset. Finally, in Chapter 5 the relation of the WP2 results in the general framework of ACANTO is described.

¹www.vision.ee.ethz.ch/datasets/

Chapter 2

Beyond the Social Force Model

The idea of modelling pedestrian motions by using a system of forces describing social interactions dates back to 1979. In [44], magnetic forces acting on a pedestrian and generated by a magnetic pole has been utilised for computer simulations with the purpose of designing building architectures. The Social Force Model (SFM) [22, 20] is one of the most popular human motion models based on social forces. In the SFM, each individual is assimilated to a point-wise particle subject to social forces. Hence, the pedestrians' dynamics are described by means of a system of differential equations. The SFM is especially well suited to reproduce individual motion of pedestrians in high-density scenarios (crowd), as well as the interactions occurring among pedestrians. The potential of the SFM, and in general of models based on social forces, in providing realistic representations of crowd behaviours has been widely acknowledged [21, 56, 17]. Due this, the original formulation of the SFM has been successively refined in the literature. For example, in [26] the authors propose an alternate version considering both relative positions and velocities, which works particularly well for low density cases. Relative velocities between pedestrians are instead considered in [61], while [52] uses pedestrians absolute velocities to govern the user head-on interactions. The relative positions and velocities provide also a way to account for the stop situation, which cannot be modelled by the original model [27, 51]. For example, [27] proposes three different SFM models for agents that are standing still. The models describe the possibility of the agent to avoid incoming humans by coding a step forward/backward behaviour, the ability to recover its desired position as well as changing it according to the environmental situation. The idea of relative velocities is further extended in [58], where the estimate of the "time to collision" is included in the SFM formulation for repulsive forces. Some versions of the SFM take into account the prediction of possible collisions explicitly, as in [53], where the time to collision has been firstly reported in lane-like avoidance by expressing the anticipation time, or in [46], where an additional force term is added to the original SFM as a function of the body and face poses.

However, to the best of the authors' knowledge, the different versions of the SFM have not explicitly modelled the dynamics of the pedestrians' heading so far. In the literature, at any time, a person is supposed to be able to move freely in any direction. On the contrary, empirical evidence shows that, most of the time, pedestrians tend to move forward, i.e. their velocity vector is most often aligned with their heading, due to the biomechanics of humans. This phenomenon has been observed by several studies [3, 2, 15], which come to the conclusion that a *nonholonomic* model may be more appropriate to describe human motion in many cases. For instance, unicycle-like models, widely used in the mobile robotics field, are able to accurately reproduce goal-oriented locomotion of an individual moving in free space [3]. Moreover, the adoption of such models in [2] allow the authors to give a nice interpretation of the mechanism underlying the formation of human trajectories (namely, the minimisation of the time-derivative of the path curvature).

To overcome this limitations, we introduce the Headed Social Force Model (HSFM) in order to enhance the traditional SFM by explicitly accounting for the pedestrians' heading, as discussed in the Introduction and as detailed in Section 2.3.

The findings reported in this chapter are part of two publications [13, 14] written in collaboration within the

ACANTO project between UNISI and UNITN.

2.1 The Social Force Model

In this section we will briefly summarise the SFM, denoting vectors in bold type. Agent i has mass m_i centred at position $\mathbf{x}_i \in \mathbb{R}^2$ in the environment, radius r_i and velocity $\mathbf{v}_i \in \mathbb{R}^2$. The SFM is described by a system of linear differential equations

$$\begin{cases} \dot{\mathbf{x}}_i = \mathbf{v}_i \\ \dot{\mathbf{v}}_i = \frac{\mathbf{v}_i^0 - \mathbf{v}_i}{\tau_i} + \frac{\mathbf{f}_i + \boldsymbol{\xi}_i}{m_i} \end{cases} \quad (2.1)$$

\mathbf{v}_i^0 is the *driving (desired) velocity* of agent i , represented by a product of speed v_i^0 and normalised direction \mathbf{e}_i^0 . Usually, \mathbf{e}_i^0 is given by the line joining the current position and the next via point. Importantly, since v_i^0 is by default set to the user's preferred walking speed, \mathbf{v}_i^0 is time invariant between via points. τ_i is the time taken to react to the difference between desired and actual velocity, while $\boldsymbol{\xi}_i$ is a noise term modelling fluctuations not accounted for by the deterministic part of the model. The noise term can also serve to avoid deadlocks and hypothesise alternative trajectories. Usually, the $\boldsymbol{\xi}_i$ is assumed normally distributed. In the absence of the exogenous inputs \mathbf{f}_i and $\boldsymbol{\xi}_i$, the agent's trajectory simply converges to the driving velocity with time constant τ_i . \mathbf{f}_i is the overall force acting on agent i resulting from other objects in the environment and is given by

$$\mathbf{f}_i = \sum_{j \neq i} [\mathbf{f}_{ij}^{\text{soc}} + \mathbf{f}_{ij}^{\text{att}} + \mathbf{f}_{ij}^{\text{ph}}] + \sum_b [\mathbf{f}_{ib}^{\text{soc}} + \mathbf{f}_{ib}^{\text{ph}}] \quad (2.2)$$

The first term on the right-hand side of (2.2) includes all the forces on agent i resulting from interactions with other agents: $\mathbf{f}_{ij}^{\text{soc}}$ is the repulsive social force that inhibits strangers from getting too close, $\mathbf{f}_{ij}^{\text{att}}$ is the attractive social force that, e.g., brings friends together, $\mathbf{f}_{ij}^{\text{ph}}$ is the physical force that exists when two people come into contact. The second term includes the forces acting on agent i as a result of fixed environmental obstacles (e.g., walls): $\mathbf{f}_{ib}^{\text{soc}}$ is the social force that inhibits agent i from getting too close to the boundaries, $\mathbf{f}_{ib}^{\text{ph}}$ is the physical force that exists when agent i touches the boundary b .

\mathbf{f} is principally a function of the distance between an agent and the other objects in the model. d_{ib} is the minimum distance between the circumference of agent i and fixed object b . d_{ij} is the distance between the centres of mass of agents i and j , i.e., the centres of the discs, while $r_{ij} = r_i + r_j$ is the "touching distance". To aid modelling the different force regimes that exist when agents are not in contact and when they touch (i.e. agents i and j touch if $r_{ij} - d_{ij} \leq 0$) it is customary to choose the function $\Theta(r_{ij}, d_{ij}) = \max(0, r_{ij} - d_{ij})$.

Using these notions, the various repulsive social and physical forces of (2.2) are defined as follows:

$$\mathbf{f}_{ij}^{\text{soc}} = \{A_i \exp[(r_{ij} - d_{ij})/B_i]\} \mathbf{n}_{ij} \Lambda(\lambda_i, \varphi_{ij}) \quad (2.3)$$

$$\mathbf{f}_{ij}^{\text{ph}} = k_1 \Theta(r_{ij} - d_{ij}) \mathbf{n}_{ij} + k_2 \Theta(r_{ij} - d_{ij}) \Delta v_{ji}^t \mathbf{t}_{ij} \quad (2.4)$$

$$\mathbf{f}_{ib}^{\text{soc}} = \{A_i \exp[(r_i - d_{ib})/B_i] + k_1 \Theta(r_i - d_{ib})\} \mathbf{n}_{ib} \quad (2.5)$$

$$\mathbf{f}_{ib}^{\text{ph}} = -k_2 \Theta(r_i - d_{ib}) (\mathbf{v}_i \cdot \mathbf{t}_{ib}) \mathbf{t}_{ib} \quad (2.6)$$

\mathbf{n}_{ij} (\mathbf{n}_{ib}) is a normalised vector pointing from agent j (fixed object b) to agent i , i.e., the direction of the repulsive force. \mathbf{t}_{ij} (\mathbf{t}_{ib}) is a normalised vector tangential to the relative movement of agent i and agent j (fixed obstacle b), i.e., the motion tangential direction. $\Delta v_{ji}^t = (\mathbf{v}_j - \mathbf{v}_i) \cdot \mathbf{t}_{ij}$ is the tangential velocity difference. The social forces (2.3) and (2.5) increase exponentially with reducing distance between objects, with a scale defined by constants A_i and B_i . In particular, A_i is the force acting on agent i at the touching distance; B_i is loosely the distance at which the force takes effect.

$\Lambda : \mathbb{R}^2 \mapsto [0, 1]$ is a function that gives greater weight to the social force (2.3) arising from the agents in front of (notionally, *seen* by) an agent. λ_i is a parameter that regulates the effect of Λ on agent i , while φ_{ij} is the

angle between the directions e_i^0 and $-\mathbf{n}_{ij}$, i.e., the field of view of the agent. The physical force (2.4) between agents comprises a repulsive body compression force (first term) that acts in direction \mathbf{n}_{ij} , plus a frictional force (second term) that acts in direction \mathbf{t}_{ij} to impede the relative tangential movement of two agents in contact. k_1 and k_2 are constants that define the scale of the physical forces. The physical force (2.6) between an agent and a fixed object is solely described by a frictional term.

A more detailed description of the SFM can be found in the original Helbing paper [22] or in the deliverable D2.1.

2.2 The Non-Holonomic Nature of Human Locomotion

The fact that the SFM cannot completely describe the interactions between humans, even in the presence of improvements of the original model, i.e., [21, 26, 52, 58], has been also recognised in [53] where the unnatural persons behaviours, such as oscillations in position when a person approaches a standing or slow obstacle or erratic behaviour when two persons meet in a corridor, have been observed. Some of these issues can be solved using a non-holonomic model.

Using empirical evidence, many authors have recognised the necessity to include a non-holonomic behaviour in the human locomotion, for example in [3] and also similarly adopted in [37]. In particular, [3] noticed that the motion of the torso of walking humans follows the dynamic of a unicycle-like vehicle. More precisely, since the torso posture anticipates the route taken, a more reliable model can be given by the car-like model in which the torso is the first derivative of the steering angle. Such an empirical observation is justified also by the neuroscience approaches modelling the human motion [24, 35, 54]. Indeed, it has been noted an inverse relationship between the path curvature and the walking speed that obeys to the so-called *power law* [35]. Using this perspective, [2] defined an inverse optimal control problem based on a car-like model to generate human-like trajectories in empty spaces. Nonetheless, in both [3, 2] the non-holonomic approach for the human locomotion is considered valid only for a subset of the possible walking behaviours, e.g., it cannot verify the step-aside manoeuvre or the backward manoeuvre, which can be instead fused coherently with a blending approach.

From a different perspective, [38, 49] adopt weighted average filters in their modified SFMs to avoid abrupt and sharp velocity changes in both directions and amplitude. The resulting model very much recalls the findings on nonholonomic kinematic models [37, 3, 2]. Moreover, [50] also noticed that depending on the current human being velocity the mobility changes in a way that very much recall the nonholonomic constraints as defined in [3].

In light of this considerations, the HSFM presented next tends to inject this nonholonomic behaviour in the SFM whenever the surroundings verify the applicability of such a model, i.e., enough free space, while it naturally turns to the standard SFM in overcrowded situations.

2.3 The Headed Social Force Model

Humans walk ahead most of the time, and their motion can be well approximated by nonholonomic models [3]. There are some circumstances, though, in which sideward motions violating nonholonomic constraints, are commonly observed (e.g., avoiding unexpected obstacles, negotiating a narrow passage or navigating in highly crowded places). In these cases, a *holonomic* model is preferable (with a slight abuse of terminology, here we denote by “holonomic model” any model not subject to nonholonomic constraints, thus including unconstrained models). In order to account for such a variability, in the HSFM each individual is modelled by means of a dynamic system like that presented in [40], which is able to reproduce both holonomic and nonholonomic motion patterns by suitably designing the system inputs (i.e., the forces and torques driving the dynamics of the pedestrians’ position and heading). In the HSFM, such inputs are designed as suitable functions of the social

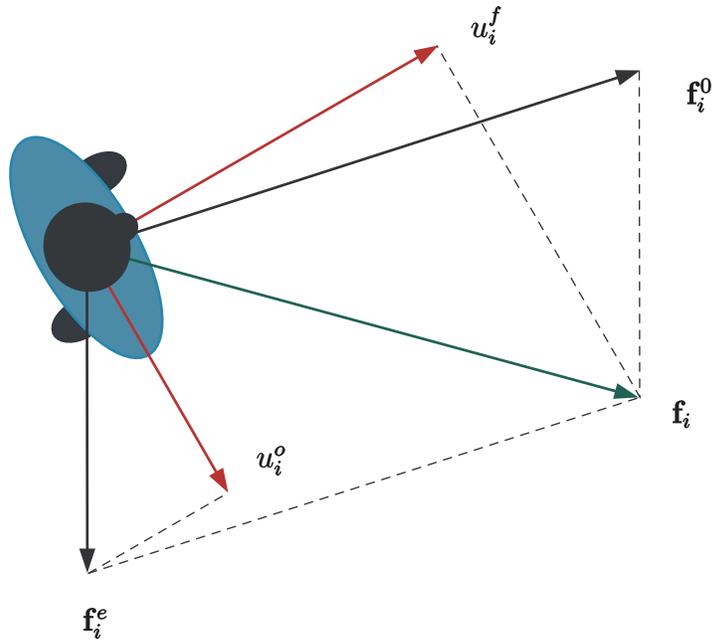


Figure 2.1: **Force decomposition in the Headed Social Force Model.** The force u_i^f , acting along the forward direction, is the projection (along the same direction) of the total force \mathbf{f}_i resulting from the traditional SFM. The force u_i^o , acting along the orthogonal direction, is the projection (along the same direction) of the only \mathbf{f}_i^e force.

forces acting on each individual, computed according to the traditional SFM. Let

$$\mathbf{f}_i = \mathbf{f}_i^0 + \mathbf{f}_i^e$$

denote the total force acting on individual i according to the SFM. The term \mathbf{f}_i^0 is the force attracting the pedestrian towards her target, such as a waypoint, whereas \mathbf{f}_i^e accounts for repulsive and interaction forces among individuals, and between individuals and the environment. In a sense, \mathbf{f}_i^0 models long-term objectives, such as travelling a prescribed path, whereas the force terms in \mathbf{f}_i^e account for short-term corrective actions, such as manoeuvres needed to avoid nearby obstacles or pedestrians. Then, in the HSFM, the motion of pedestrians is generated as follows:

- The forces u_i^f and u_i^o driving the *translational dynamics* are computed by projecting \mathbf{f}_i and \mathbf{f}_i^e along the forward direction of motion (identified by the pedestrian's heading) and the orthogonal direction of motion, respectively (see Fig 2.1).
- The torque driving the *rotational dynamics* is proportional to the projection of the term \mathbf{f}_i^0 along the orthogonal direction of motion.
- An additional force term is added in order to ensure *group cohesion* when simulating people moving together. This is achieved by: i) defining a rectangular region, centred at the group centroid, within which the group members are expected to lie, and ii) exerting a force pushing the pedestrians back into that region whenever they get out of it.

In both the translational and the rotational dynamics, damping terms are included in order to weaken oscillations and obtain smoother trajectories. Only the force \mathbf{f}_i^e is assumed to affect lateral moves, because they are mainly

caused by the interactions with other pedestrians or the environment. On the other hand, body rotations are generated by the lateral component of the force \mathbf{f}_i^0 , which is in charge of driving the pedestrian towards the goal. This choice is derived directly from observations: a person tends to turn faster towards the target the more she/he is attracted by the target itself. The idea of the group cohesion force is inspired by the approach proposed in [42] for modelling small groups of pedestrians (from two to four individuals) walking together and subject to social interaction constraints. In this deliverable, such an approach is adapted to the proposed dynamic model which accounts for pedestrians' heading. In particular, the force term is designed in order to reproduce the formation of larger groups, including many individuals moving together (e.g., like a group of tourists following a guide).

The proposed HSFM enriches the traditional SFM with a more complex human locomotion model which is well suited to represent human trajectories complying with nonholonomic constraints, as typically occurs in large spaces occupied by a limited number of pedestrians. At the same time, the HSFM preserves the power of the SFM in realistically reproducing the flow of a large number of people moving in densely populated environments. A unique feature of the proposed model lies in its ability to adapt to the external conditions, by smoothly switching between holonomic and nonholonomic motion patterns depending on a number of factors, including the pedestrian density, the pedestrians' goal and the clutter of the environment. Notably, this behaviour is achieved without the need of changing online any of the model parameters, but as a natural reaction and adaptation to the external conditions.

We present now the mathematical details of the proposed Headed Social Force Model. Consider a system of n pedestrians moving in a 2D environment. Following the modelling approach of the Social Force Model subsumed in (2.1), the i -th individual, $i = 1, \dots, n$, is assimilated to a particle with mass m_i , whose position and velocity, expressed in a global reference frame, are denoted by $\mathbf{r}_i = [x_i, y_i]^\top$ and $\mathbf{v}_i = [\dot{x}_i, \dot{y}_i]^\top$, respectively. The equations of motion are

$$\begin{aligned}\dot{\mathbf{r}}_i &= \mathbf{v}_i, \\ \dot{\mathbf{v}}_i &= \frac{1}{m_i} \mathbf{u}_i,\end{aligned}$$

where \mathbf{u}_i represents the social force driving the i -th particle. In order to include the pedestrians' heading into the model, it is convenient to attach a body frame to each individual, i.e. a reference frame centred at the pedestrian's position and whose x -axis is aligned with the pedestrian's forward direction of motion. Let $\mathbf{q}_i = [\theta_i, \omega_i]^\top$ be the vector containing the heading θ_i (angle between the x -axis of the body frame and that of the global reference frame) and the angular velocity $\omega_i = \dot{\theta}_i$ of the i -th pedestrian. Denote by $\mathbf{v}_i^B = [v_i^f, v_i^o]^\top$ the velocity vector expressed in the body frame. The components v_i^f and v_i^o of vector \mathbf{v}_i^B correspond to the projection of the velocity vector \mathbf{v}_i along the forward direction and the orthogonal direction, respectively. Clearly, $\mathbf{v}_i = \mathbf{R}(\theta_i) \mathbf{v}_i^B$, where the rotation matrix $\mathbf{R}(\theta_i)$ is defined as

$$\mathbf{R}(\theta_i) = \begin{bmatrix} \cos(\theta_i) & -\sin(\theta_i) \\ \sin(\theta_i) & \cos(\theta_i) \end{bmatrix} \doteq \begin{bmatrix} \mathbf{r}_i^f & \mathbf{r}_i^o \end{bmatrix}.$$

Then, similarly to [40], the human locomotion model can be written as

$$\dot{\mathbf{r}}_i = \mathbf{R}(\theta_i) \mathbf{v}_i^B, \quad (2.7)$$

$$\dot{\mathbf{v}}_i^B = \frac{1}{m_i} \mathbf{u}_i^B, \quad (2.8)$$

$$\dot{\mathbf{q}}_i = \mathbf{A} \mathbf{q}_i + \mathbf{b}_i u_i^\theta, \quad (2.9)$$

where

$$\mathbf{A} = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}, \quad \mathbf{b}_i = \begin{bmatrix} 0 \\ \frac{1}{I_i} \end{bmatrix}, \quad (2.10)$$

and I_i denotes the moment of inertia of pedestrian i . In model (2.7)-(2.9), the inputs are $\mathbf{u}_i^B = [u_i^f, u_i^o]^\top$, whose entries are the forces acting along the forward direction and the orthogonal direction, respectively, as well as the torque u_i^θ about the vertical axis, as depicted in Figure 2.1. Notice that such a model is indeed unconstrained. However, if $v_i^o(0) = 0$ and $u_i^o(t) = 0$, for all t , the dynamic unicycle model is recovered. In general, whenever $v_i^o = 0$, the model features a nonholonomic behaviour, the velocity vector being aligned with the pedestrian's heading.

The basic idea of the HSFM is to compute the model inputs u_i^f , u_i^o and u_i^θ on the basis of the the forces resulting from the traditional SFM. To this purpose, the total force \mathbf{f}_i that acts on the i -th pedestrian according to [20] is decomposed as

$$\mathbf{f}_i = \mathbf{f}_i^0 + \mathbf{f}_i^e. \quad (2.11)$$

The first term

$$\mathbf{f}_i^0 = m_i \frac{\mathbf{v}_i^d - \mathbf{v}_i}{\tau_i} \quad (2.12)$$

accounts for the pedestrian's desire to move with a given velocity vector \mathbf{v}_i^d . In (2.12), the characteristic time $\tau_i > 0$ is a parameter determining the rate of change of the velocity vector. The force

$$\mathbf{f}_i^e = \mathbf{f}_i^p + \mathbf{f}_i^w \quad (2.13)$$

accounts for the pedestrians' interaction. The terms \mathbf{f}_i^p and \mathbf{f}_i^w represent the repulsive forces exerted on individual i by the other pedestrians and by possible obstacles present in the environment (e.g., walls), respectively. The expressions of \mathbf{f}_i^p and \mathbf{f}_i^w are simply given by linear combination of the relations in (2.3)-(2.6) defined for the standard SFM and define the effects of other pedestrians on individual i and the repulsive effects of obstacles or boundaries such as walls on individual i , respectively, i.e.

$$\mathbf{f}_i^p = \sum_{j \neq i} \mathbf{f}_{ij}^{\text{soc}} + \mathbf{f}_{ij}^{\text{ph}} \quad \text{and} \quad \mathbf{f}_i^w = \sum_b \mathbf{f}_{ib}^{\text{soc}} + \mathbf{f}_{ib}^{\text{ph}}.$$

This is an alternative way to include all the effects of the SFM included in (2.2).

2.3.1 Force Inputs

The inputs of the HSFM are computed from \mathbf{f}_i^0 and \mathbf{f}_i^e as follows. The input vector \mathbf{u}_i^B includes the forces acting along the pedestrian's forward direction and the orthogonal direction (see Figure 2.1). Given the total social force \mathbf{f}_i , a natural choice for computing u_i^f is to project \mathbf{f}_i along the forward direction. In order to avoid sideward motions if not strictly needed, the component u_i^o is computed by projecting the interaction force \mathbf{f}_i^e (possibly scaled), along the orthogonal direction. Finally, in order to drive to zero the sideward velocity v_i^o when the sideward force is zero, a damping term proportional to v_i^o is added to u_i^o . Hence, the model inputs u_i^f and u_i^o are computed as

$$u_i^f = (\mathbf{f}_i^0 + \mathbf{f}_i^e)^\top \mathbf{r}_i^f, \quad (2.14)$$

$$u_i^o = k^o (\mathbf{f}_i^e)^\top \mathbf{r}_i^o - k^d v_i^o, \quad (2.15)$$

where $k^o > 0$ and $k^d > 0$.

2.3.2 Torque Input

The input u_i^θ represents the torque about the vertical axis which drives the dynamics of the pedestrian's heading. This term is designed on the basis of the force \mathbf{f}_i^0 defined in (2.12). Denote by f_i^0 and θ_i^0 the magnitude and the

phase in the global reference frame of \mathbf{f}_i^0 . Notice that both quantities are in general time-varying. The input u_i^θ is computed as

$$u_i^\theta = -k^\theta(\theta_i - \theta_i^0) - k^\omega\omega_i. \quad (2.16)$$

The parameters k^θ and k^ω are designed in order to achieve suitable dynamics of the heading. It can be easily verified that, with u_i^θ defined as in (2.16), the orientation error $\tilde{\theta}_i \doteq \theta_i - \theta_i^0$ evolves according to the dynamic model

$$\ddot{\tilde{\theta}}_i + \frac{k^\omega}{I_i}\dot{\tilde{\theta}}_i + \frac{k^\theta}{I_i}\tilde{\theta}_i = -\frac{k^\omega}{I_i}\dot{\theta}_i^0 - \ddot{\theta}_i^0. \quad (2.17)$$

A possible design procedure is to select the values of k^θ and k^ω on the basis of the desired poles λ_1 and λ_2 of the dynamic system (2.17). For the ACANTO model, real poles are considered, so that $\lambda_2 = \alpha\lambda_1 < 0$, for some $\alpha > 1$. In turn, the dominant pole λ_1 is selected as a function of f_i^0

$$\lambda_1 = -\sqrt{\frac{k^\lambda f_i^0}{\alpha}},$$

where $k^\lambda > 0$ is used to tune the dominant time constant of system (2.17). The corresponding expressions of k^θ and k^ω are then

$$k^\theta = I_i k^\lambda f_i^0, \quad k^\omega = I_i(1 + \alpha)\sqrt{\frac{k^\lambda f_i^0}{\alpha}}. \quad (2.18)$$

The choice of time-varying poles allows one to modulate the responsiveness of the system with the intensity of the driving force \mathbf{f}_i^0 . The underlying idea is that the more authoritative the \mathbf{f}_i^0 , the faster the change in the pedestrian's heading. In this way, the heading convergence rate is proportional to f_i^0 .

2.4 Enforcing Group Cohesion in the HSFM

Increasing attention is gaining the description of groups of people socially interacting while moving in an environment, since it has been noticed that in some environments social groups comprise about 70%-80% of the walking population [48]. A social group is a collection of people that move together to a common goal and are engaged in a social relation [39]. The concept of relation is quite blurred, since it refers to both verbal and non-verbal communication, such as gestures and gaze exchange [29], but it also comprises the probably most important interaction of eye contact, which helps estimating the partners' reactions and anticipating their intentions [28]. Eye contact is also responsible of the equilibrium of physical proximity [4]. Social groups have been investigated from a variety of different stand-points. For example, it has been observed that a group of socially interacting and walking pedestrians with more than three persons tends to split into groups of two or three elements with a very high probability [25]. The size of the social group has been also investigated in the past and it has been conjectured that the probability of losing or gaining a group member per unit time is proportional to the group size [8, 41]. The velocity of the group has been also analysed: the more is the density of the crowd, the less would be the velocity of the group, which is evidently true also for individuals. What is of major interest is that the group velocity is inversely proportional to the group dimension [41]. Similar empirical observations have been also collected by [30, 31].

In the deliverable D2.1, we have reported that the preferred spatial structure for three interacting persons is the “V” or “U” formation with the vertex behind [10]. Since in low and moderate densities the group structures bends forward in walking directions, and not backwards as a expected from mechanics of flexible structures moving against an opposite flux [41], it is recognised in the literature that the formation is *actively created* to facilitate the interaction process. Instead, as naturally expected, higher densities squeeze the structure. If the density increases more, the structure is lost and the emerging behaviour of the group is to act like a “river” [19], which is an SFM well described behaviour. Similarly to the blend between the nonholonomic motion and the

original SFM, it turns again out that by injecting the group formation idea in the HSFM an accurate description of the social interaction can be also given.

The very first effective approach that modifies the SFM with the purpose of including social interactions during motions has been provided by [41], where the Authors add a new force term, the *group force* term, in order to describe the aggregation forces. The group force term is the sum of three independent forces. The main characteristic of this model is its ability to describe the abreast, the “V” and the “U” group formations in the case of two, three or four group members, respectively. The idea of modifying the SFM by adding an additional force term has been also used [55] for evacuation dynamics, in [17] to describe large group motions or in [60, 59] using potential fields among the group pedestrians.

For the HSFM, we followed a similar path. In order to model a group of people moving together in a group, the force input (2.14)-(2.15) can be modified by adding an additional force term, which forces the pedestrians to lie within a given region.

Let $\mathbf{c} = \frac{1}{n} \sum_{j=1}^n \mathbf{r}_j$ be the centroid of the group, and define $\mathbf{p}_i = \mathbf{c} - \mathbf{r}_i$. The model inputs u_i^f and u_i^o are computed as

$$u_i^f = (\mathbf{f}_i^0 + \mathbf{f}_i^e)^\top \mathbf{r}_i^f + k_1^g h(\mathbf{p}_i, \mathbf{r}_i^f, d^f), \quad (2.19)$$

$$u_i^o = k^o (\mathbf{f}_i^e)^\top \mathbf{r}_i^o - k^d v_i^o + k_2^g h(\mathbf{p}_i, \mathbf{r}_i^o, d^o), \quad (2.20)$$

where $k_1^g > 0$, $k_2^g > 0$, $d^f > 0$, $d^o > 0$ and

$$h(\mathbf{x}, \mathbf{y}, z) = \begin{cases} 1 & \text{if } |\mathbf{x}^\top \mathbf{y}| > z \\ 0 & \text{otherwise.} \end{cases} \quad (2.21)$$

Chapter 3

HSFM Simulations

In this section, the results of a number of numerical simulations are reported, in order to highlight the characteristic features of the proposed model. Three different scenarios are considered. In Scenario I, we simulate two simple case studies, involving a single pedestrian, aimed at showing the high fidelity of the HSFM in reproducing the trajectories of pedestrians moving in free space according to a nonholonomic behaviour. In Scenario II, we consider three different experiments, involving a number of pedestrians ranging from 20 to 200. The purpose is to illustrate the ability of the HSFM to automatically adapt the generated trajectories to the external context, smoothly relaxing the nonholonomic constraints as the pedestrian density increases or unexpected obstacles come into play. In Scenario III, we consider a more articulated case study, by simulating a group of 10 people visiting a museum together. The focus of this study is to show how the group force introduced in the HSFM originates trajectories preserving the cohesion of the group. This section concludes with a discussion on the role played by the parameters of the HSFM. An extensive simulation campaign is performed in order to analyse the effect of parameter variations on the generated trajectories, thus providing useful guidelines for the tuning of the model.

In all the simulations presented hereafter, the reference velocity vector \mathbf{v}_i^d , which is used by the SFM to compute the force \mathbf{f}_i^0 (see (2.12)), is generated as $\mathbf{v}_i^d = v^d \mathbf{e}_i^d$. The desired speed v^d is assumed constant over each simulation run. The unit vector \mathbf{e}_i^d , which identifies the desired direction of motion, is computed from a sequence of way-points encoding the desired pedestrian path, similarly to [22]. Videos of the simulations are available at <http://control.dii.unisi.it/MobileRoboticsPage>.

In the simulations, the radius r_i and the mass m_i of each pedestrian have been randomly generated in the intervals [0.25 m, 0.35 m] and [60 kg, 90 kg], respectively, assuming uniform distributions. The inertia moment I_i in (2.10) is computed as $I_i = \frac{1}{2}m_i r_i^2$, i.e., the pedestrian is assimilated to a cylinder rotating about its main axis. The following parameters entering in the computation of the model inputs (2.14)-(2.18) and (2.19)-(2.21) have been used in all the simulations (unless differently stated): $k^o = 1$, $k_d = 500 \text{ kg}\cdot\text{s}^{-1}$, $\alpha = 3$, $k^\lambda = 0.3 \text{ N}^{-1}\text{s}^{-2}$, $d^f = 2 \text{ m}$, $d^o = 1 \text{ m}$ and $k_1^g = k_2^g = 200 \text{ N}$. The value of the parameters of the SFM used in all the simulations are taken from [20].

3.1 HSFM Behaviour

3.1.1 Nonholonomic behaviour

To evaluate how well the HSFM can reproduce the nonholonomic behaviour empirically observed in [3], we consider two use cases.

In the first example, a single pedestrian walks between two points A and B , multiple times. In this case, the trajectory resulting from the SFM is quite unnatural, as the path boils down to a segment (red line in Figure 3.1). In other words, the SFM models a human being that either walks forward and backward without

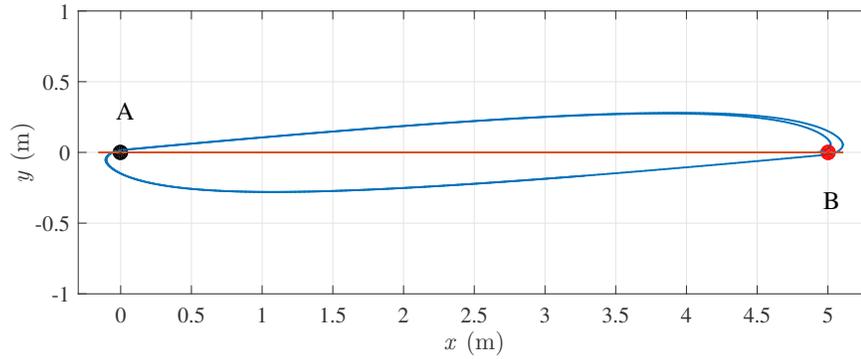


Figure 3.1: **Scenario I, alternate motion between two points.** A single pedestrian has to move back and forth between A and B , starting from A , with a desired speed $v^d = 1.5 \text{ ms}^{-1}$: SFM (red) and HSFM (blue).

turning or turns instantaneously on the spot. This phenomenon is due to the SFM neglecting the information about the pedestrian’s heading, so that forward or backward motions are equivalent. On the contrary, the trajectory generated by the HSFM is more realistic thanks to the existence of a preferred direction of motion (blue line in Figure 3.1). Although the HSFM allows a pedestrian to have her velocity vector not aligned with her heading, the model input tends to drive the orthogonal component of the velocity to zero if no lateral forces are present, thus generating an “almost nonholonomic” behaviour. It can be observed that in the resulting path, the pedestrian approaches the turning point preparing to invert her orientation with a sort of U-turn, as it happens in practice.

In the same setting, consider the case in which a pedestrian has to move from A to B , starting with four different values of the initial heading $\theta(0)$ (see Figure 3.2). When $\theta(0) = \pi$, the goal point B lies behind the pedestrian’s back. In this case, the HSFM makes the pedestrian first take a step back to turn towards the goal, and then move forward to reach the target. Clearly, the SFM trajectory lies on a segment once again, since the heading is neglected.

The previous examples confirm that, in the considered scenario, the HSFM gives rise to a more realistic behaviour, endowing the pedestrian with the ability of moving in a nonholonomic way when she is expected to do so.

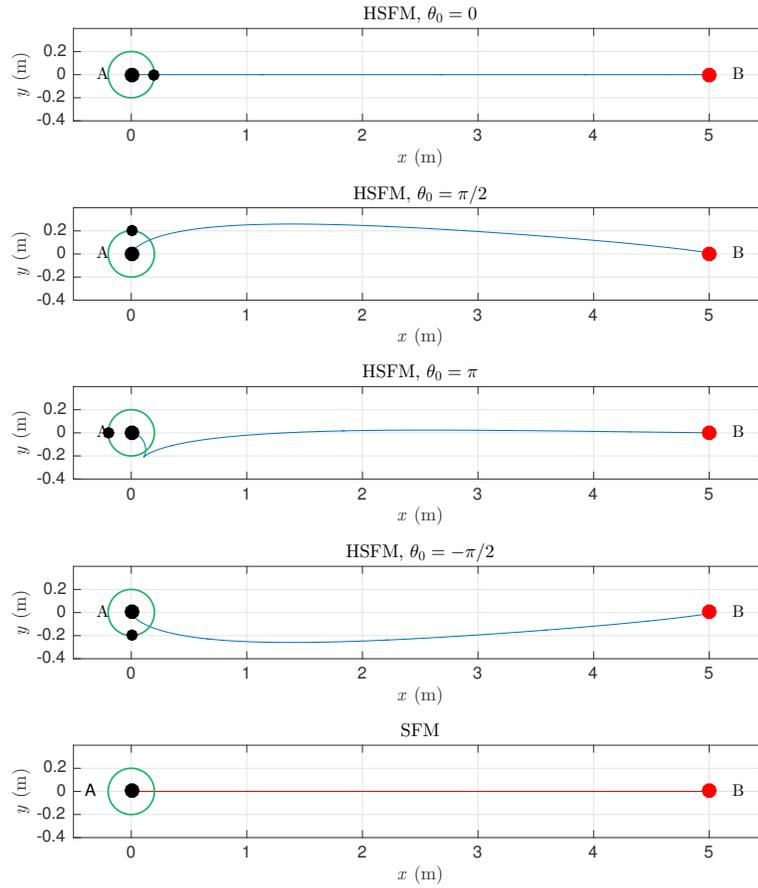


Figure 3.2: **Scenario I, starting with different orientations.** A single pedestrian has to move from A to B , starting with different headings (denoted by the black dot), at a desired speed $v^d = 1.5 \text{ ms}^{-1}$: SFM (red) and HSFM (blue).

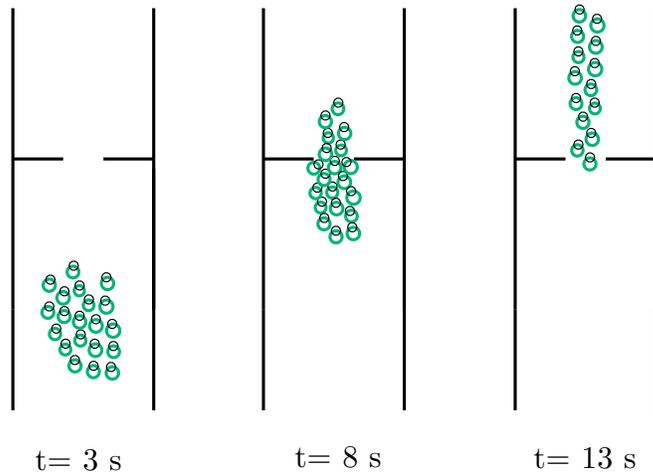


Figure 3.3: **Scenario II, Pedestrians in a corridor.** A group of 20 pedestrians walking in the same direction in a 7.5m-wide corridor at a desired speed $v^d = 1.5 \text{ ms}^{-1}$. Three snapshots of a simulation run of the HSFM, taken at different time instants t .

3.1.2 Blending between SFM and Nonholonomic behaviour

In this scenario, we consider three examples. In the first one, 20 pedestrians walking in a 7.5m-wide corridor have to pass through a 2m-wide door (see Figure 3.3). In the second example, two groups of 10 pedestrians are walking in opposite directions in a 5m-wide corridor (see Figure 3.4). In the third example, we consider an evacuation experiment in which 200 pedestrians must get out of a $15\text{m} \times 15\text{m}$ room through a door of width 1 m, similarly to what presented in [20] (see Figure 3.5). In these experiments no group cohesion forces are included.

For comparison purposes, the following indicators are considered:

- the average *exit frequency* of pedestrians \bar{F} , i.e. the average number of pedestrians that pass through the door per unit time (first and third examples);
- the average *square of the magnitude of the jerk* of the trajectories

$$\bar{J} = \frac{1}{n} \sum_{i=1}^n \frac{1}{T} \int_0^T \|\mathbf{j}_i(t)\|^2 dt, \quad (3.1)$$

where n denotes the total number of pedestrians and \mathbf{j}_i is the jerk vector of the i -th trajectory, (i.e., the third-order derivative of the position).

The first indicator has been selected as a measure of the macroscopic behaviour of the models. The second indicator is used to evaluate both the regularity and the realism of the resulting trajectories. As a matter of fact, it is commonly acknowledged that the motions performed by humans tend to be smooth and to minimise the jerk, as first experimentally verified for hand movements in [16], and extended to the trajectories of walking pedestrians later on in [45]. The obtained results can be summarised as follows.

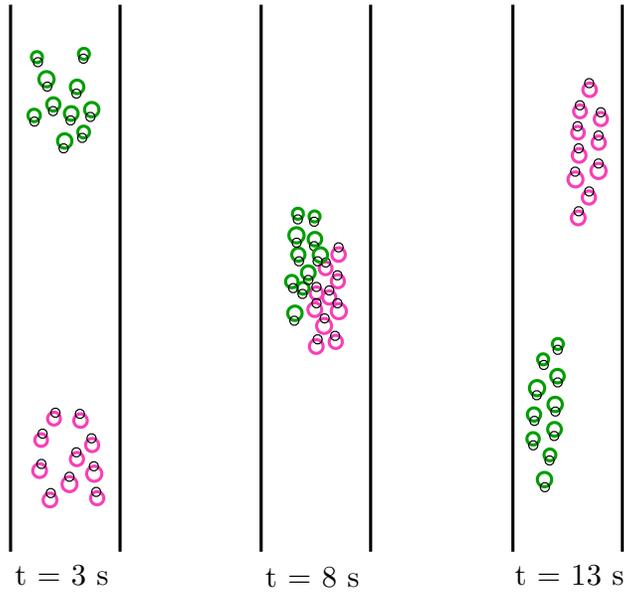


Figure 3.4: **Scenario II, Two groups walking in opposite directions.** Two groups of 10 pedestrians each walking in opposite directions in a 5m-wide corridor at a desired speed $v^d = 1.5 \text{ ms}^{-1}$. Three snapshots of a simulation run of the HSFM, taken at different time instants t .

Pedestrians in a corridor

In order to compare the trajectories generated by the SFM and the HSFM, a Monte Carlo analysis has been performed. Starting from random initial positions and headings of the pedestrians (with zero initial velocity), 100 runs of the SFM and the HSFM have been simulated for 20 s. Concerning the exit frequency, both models give similar results, with average values $\bar{F}_{HSFM} = 2.70 \text{ s}^{-1}$ and $\bar{F}_{SFM} = 2.75 \text{ s}^{-1}$. Overall the two models seem to reproduce the same macroscopic behaviour. However, significant differences can be appreciated by looking at the regularity of the resulting trajectories. The average square of the magnitude of the jerk is very different in the two cases, with average values during the door crossing (time range [6, 10] seconds) of $\bar{J}_{HSFM} = 4.1 \cdot 10^{-4} \text{ m}^2\text{s}^{-6}$ and $\bar{J}_{SFM} = 5.3 \cdot 10^{-3} \text{ m}^2\text{s}^{-6}$. These figures capture the different qualitative behaviours that can be observed by looking at the resulting trajectories. When compared to the HSFM, in the proximity of the door, the SFM tends to generate vibrations, sudden changes of direction and even “bounces” among pedestrians or between pedestrians and walls.

Two groups walking in opposite directions

Also in this case, results are averaged over 100 simulation runs. In this example, the difference of the indicator \bar{J} ($\bar{J}_{HSFM} = 4.3 \cdot 10^{-3} \text{ m}^2\text{s}^{-6}$ for the HSFM vs. $\bar{J}_{SFM} = 2.3 \cdot 10^{-2} \text{ m}^2\text{s}^{-6}$ for the SFM) is mostly due to the very different trajectories over the time range [6, 10] seconds, where the two groups interact to negotiate the traversing of the corridor. In this situation, the pedestrian motion generated by the HSFM is much more regular than that reproduced by the SFM, in which several collisions among pedestrians belonging to different groups are experienced.

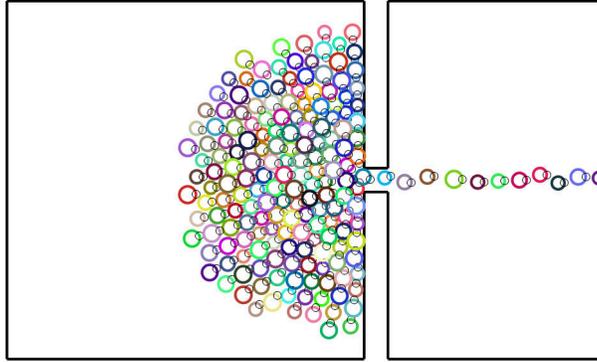


Figure 3.5: **Scenario II, Evacuation of a room.** A group of 200 pedestrians must evacuate from a $15\text{m} \times 15\text{m}$ room through a door of width 1 m (see [20]).

Evacuation of a room

In the considered setting, the door can become a bottleneck and arching and clogging may arise in the proximity of the door. In [20], it was studied how the exit frequency varies with the desired speed v^d of the pedestrians. As expected, at slow speeds, the frequency grows with v^d . However, when v^d exceeds a threshold value (about 1.5 ms^{-1}) the frequency drops due to the increased jam induced by panic (the so called “faster-is-slower” effect [20]). In order to evaluate the ability of the HSFM to reproduce such a phenomenon, the evacuation experiment has been simulated for 60 s, at different desired speeds, ranging from 0.5 ms^{-1} to 6 ms^{-1} . For each simulation run, the average exit frequency resulting from the SFM and the HSFM has been computed. The results are pretty similar (see Fig 3.6), thus confirming the adequateness of the HSFM also in highly crowded environments.

Comments

Overall, previous results show that at a microscopic level, the HSFM generates smoother, and consequently more realistic, trajectories than the traditional SFM. At the same time, the macroscopic behaviour of the whole system, which is typically well approximated by the SFM, is fully preserved.

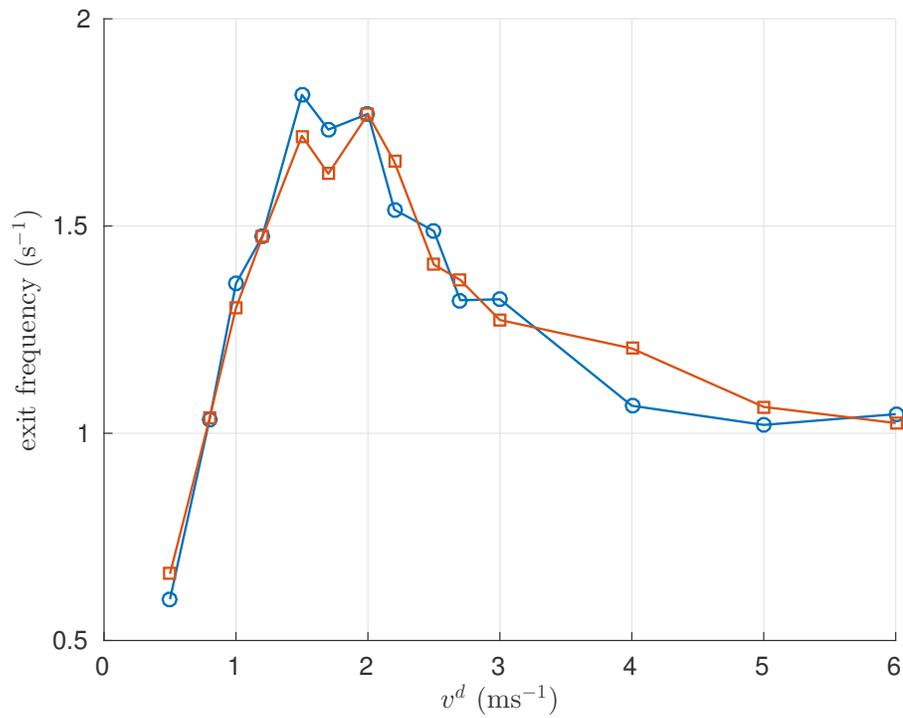


Figure 3.6: **Scenario II, Evacuation of a room.** Average exit frequency: SFM (red) and HSFM (blue) for different values of the desired speed v^d .

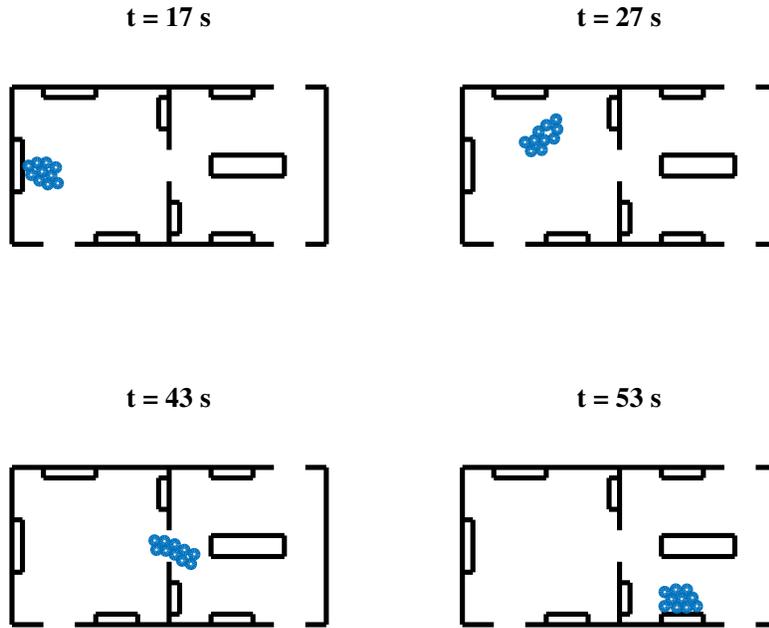


Figure 3.7: **A visit at the Museum with Group Cohesion Forces.** Snapshots of a simulation run of the HSFM with the inclusion of group cohesion forces.

3.2 Group Cohesion Forces

In this scenario we test the ability of the HSFM to reproduce pedestrians moving together. As a case study, we consider the visit of a museum carried out by a group of 10 people. The considered environment is composed of two communicating rooms, each of which contains four artworks on display. Three doors connect the rooms with the rest of the museum (see Figure 3.7). The objective of the group is to visit a selection of the pieces of the exhibition in a given order, while avoiding collisions with obstacles and/or other individuals. Once the visitors reach the selected artwork, they stop in front of it for a predefined amount of time, before moving to the next point of interest.

We compare the results obtained using the HSFM with and without the group forces. In Figure 3.8 and Figure 3.7 four different snapshots of the trajectories from the two cases are shown. The main difference lies in the way the group moves from one exhibition to the other. In the absence of group cohesion forces, the group tends to elongate and the visitors form a line (see Figure 3.8). This unrealistic behaviour is avoided when group forces are included (see Figure 3.7). A measure of the group cohesion is given by the average *distance from the centroid* of the group, defined as

$$\xi(t) = \frac{1}{n} \sum_i^n d_i(t), \quad (3.2)$$

where $d_i(t)$ is the distance at time t of pedestrian i from the centroid of the group. This indicator gives a measure of the dispersion of the pedestrians during their motion. The time evolution of $\xi(t)$ is depicted in Figure 3.9, in both cases. Without group forces, the group radius oscillates between small values (corresponding to the visitors standing still in front of an artwork) and large values (when people switch from one artwork to the

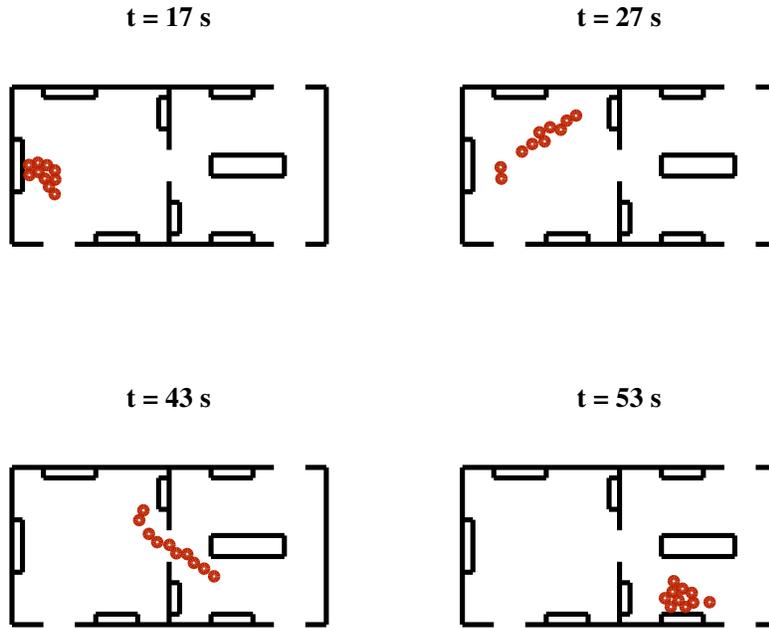


Figure 3.8: **A visit at the Museum.** Snapshots of a simulation run of the HSFM without the inclusion of group cohesion forces.

next one). Conversely, the introduction of the group forces effectively keeps the group together, with a radius smaller than 2 m.

3.3 Model Parameters Tuning

In this section, we study the role of the parameters of the HSFM on the resulting system behaviour. Specifically, we consider separately the parameters which affect the computation of: i) the force input, ii) the torque input and iii) the group cohesion term.

3.3.1 Force input

The force driving the translational dynamics of the pedestrian depends on two parameters, namely k^o and k^d . The first one is a gain that modulates the force acting on the direction orthogonal to the pedestrian's heading. The second one is a damping coefficient on the speed along the same direction. As a case study representative of the HSFM behaviour under most circumstances, the same example, described in the evacuation scenario, involving 20 pedestrian crossing a door in a corridor, is considered (see Figure 3.3). In this analysis, no group cohesion forces are included. Several simulations have been carried out for different combinations of the parameter values. Figure 3.10 depicts a snapshot of the simulations taken when the individuals have almost completely crossed the door. By looking at the different configurations of the pedestrians, the following phenomena can be observed. For a given k^d , the platoon gets wider as k^o increases, since more authoritative lateral repulsive forces among pedestrians are exerted. Parameter k^d has an even greater impact on the width

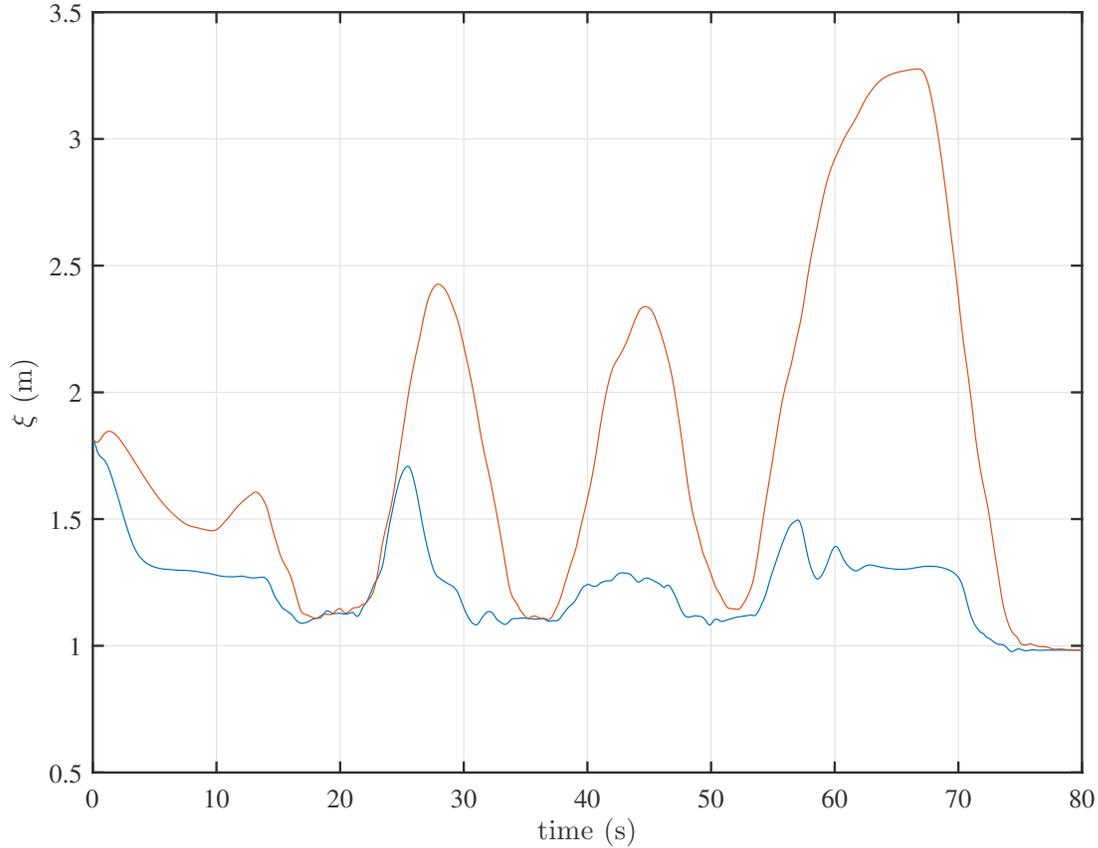


Figure 3.9: **Mean distance from the group centroid over time.** It shows the evolution of ξ with cohesive forces (blue) and without cohesive forces (red).

of the platoon. For a fixed k^o , the larger the value of k^d , the faster the lateral speed is driven towards zero. As a result, with very high values of k^d the pedestrians tend to arrange in a line, which is an unrealistic behaviour. Besides the geometric distribution of the individuals, both parameters have an effect on the smoothness of the generated trajectories. To analyse this feature, 100 simulation runs have been performed, starting from random initial conditions. In Figure 3.11, two indicators are shown as a function of k^d , for different values of k^o . The first one is the average square of the magnitude of the jerk \bar{J} as defined in (3.1), which measures the regularity of the trajectories. The second one is Δ , defined as

$$\Delta = \frac{1}{T} \int_0^T \xi(t) dt,$$

where $\xi(t)$ is given by (3.2). It represents the mean distance of a pedestrian from the centroid, averaged over the whole simulation run. The evolution of \bar{J} suggests that the trajectories become more and more regular as k^o decreases and k^d increases. The tuning of parameter k^d has to take into account also the impact that it has on the geometry of the platoon, which in Figure 3.11 is summarised by the indicator Δ . Too large values of k^d imply a growth of the radius Δ , which, in turns, reflects the tendency of the pedestrians to form a line. Hence, parameter k^d has to be tuned by trading-off these conflicting objectives. Values in a neighbourhood of $k^o = 1$ and $k^d = 500 \text{ kg}\cdot\text{s}^{-1}$ have been observed to ensure regular trajectories and a realistic geometry of the

platoon [50, 1]. Moreover, this choice guarantees very low sensitivity of the indicators \bar{J} and Δ to variations in the model parameters, which suggests robustness of the system behaviour under different scenarios.

3.3.2 Torque input

The torque controlling the heading dynamics is designed via pole placement, so that the closed-loop system has a desired pair of real poles. In this approach, a major role is played by the pole ratio α . The effect of α on the resulting trajectory is clearly visible in Figure 3.12, for the simple case in which a pedestrian goes through four way-points forming a square. Basically, the larger the α , the slower is the dynamics of the pedestrian's heading, which results in larger turning radius. Values of α in the range 3-5 seem appropriate for reproducing a realistic path, the resulting curvature dynamics being neither too aggressive (i.e., $\alpha = 1$) nor too loose (i.e., $\alpha = 10$).

3.3.3 Group cohesion

The parameters defining the force term which aims at keeping together people belonging to the same group, have a clear physical meaning. This makes their tuning much easier than the previous ones. Parameters d^f and d^o are half of the side length of the desired rectangular region along the forward and orthogonal direction, respectively. Parameters k_1^g and k_2^g correspond to the intensity of the cohesion forces acting along the forward and orthogonal direction, respectively. In the simulations presented so far, the following values have been selected: $d^f = 2$ m, $d^o = 1$ m and $k_1^g = k_2^g = 200$ N.

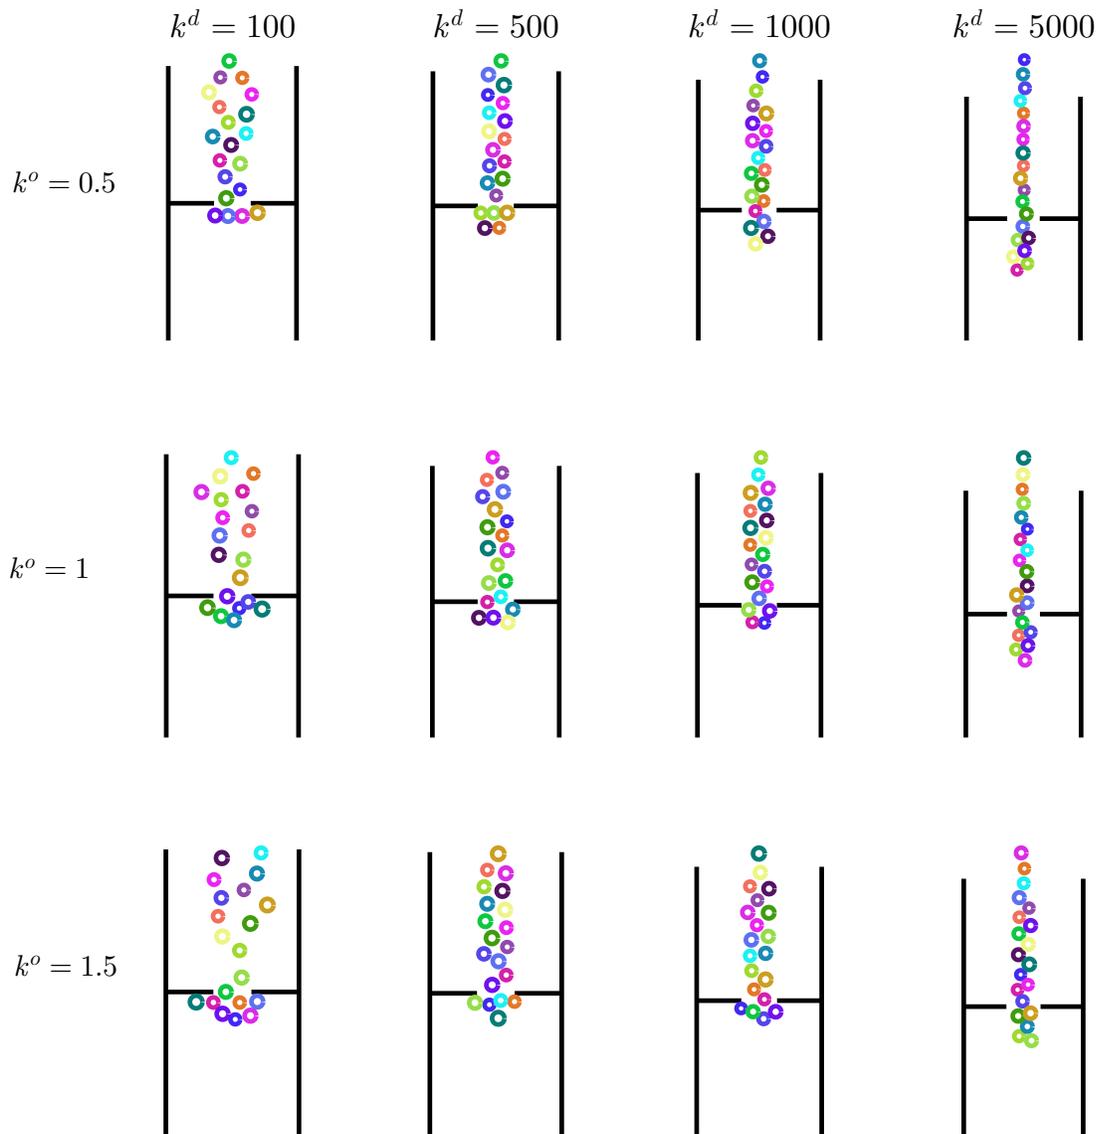


Figure 3.10: **Effect of k^o and k^d on the pedestrian trajectories.** A snapshot of the simulation of 20 pedestrians walking in a corridor, for different values of k^o and k^d [$\text{kg}\cdot\text{s}^{-1}$].

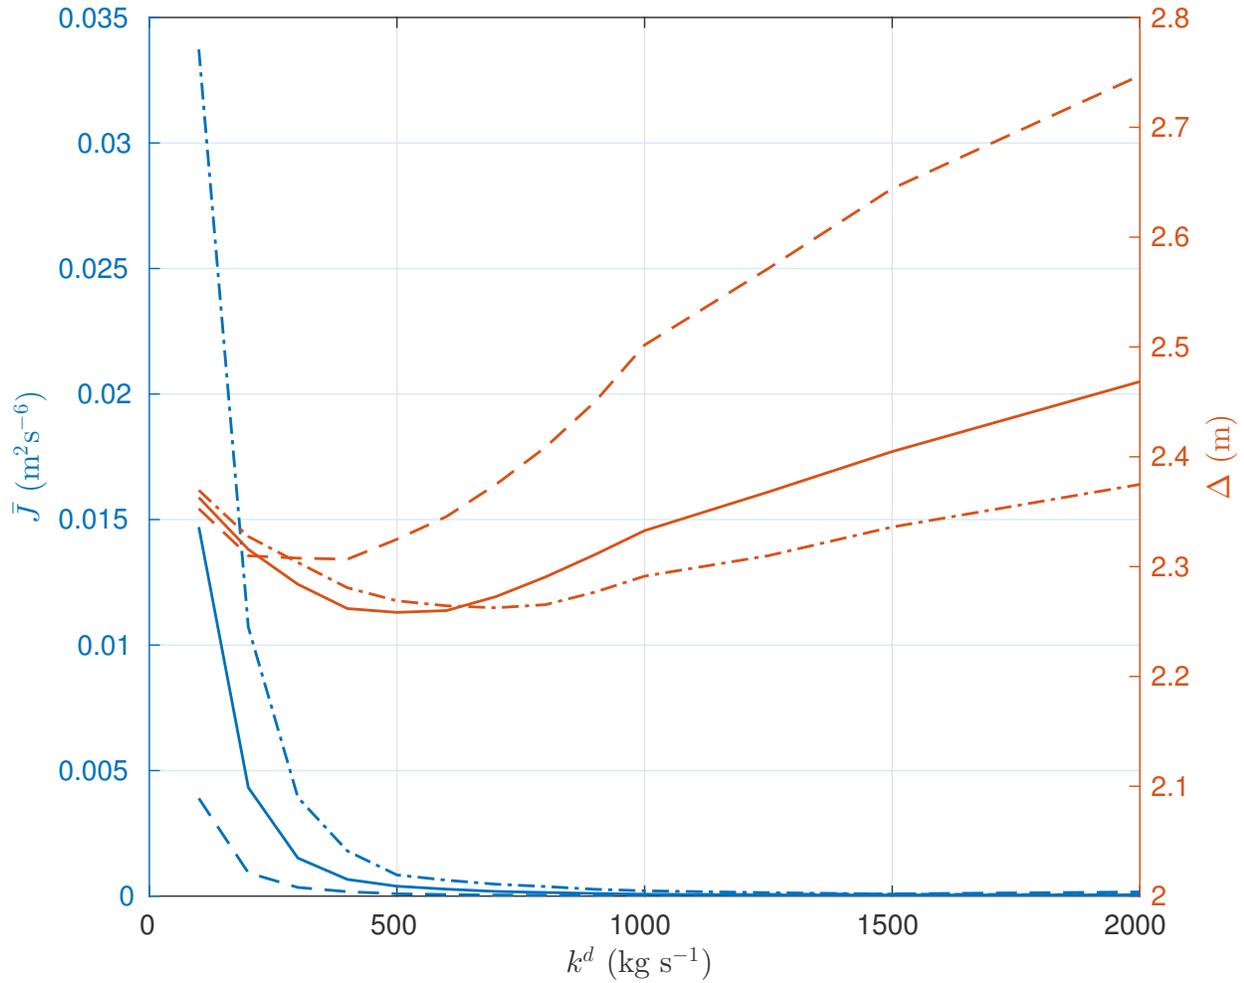


Figure 3.11: **Effect of k^o and k^d on trajectory regularity and distribution of the pedestrians.** Average square of the magnitude of the jerk \bar{J} and average distance Δ of a pedestrian from the group centroid for $k^o = 0.5$ (dashed), $k^o = 1$ (solid) and $k^o = 1.5$ (dash-dotted).

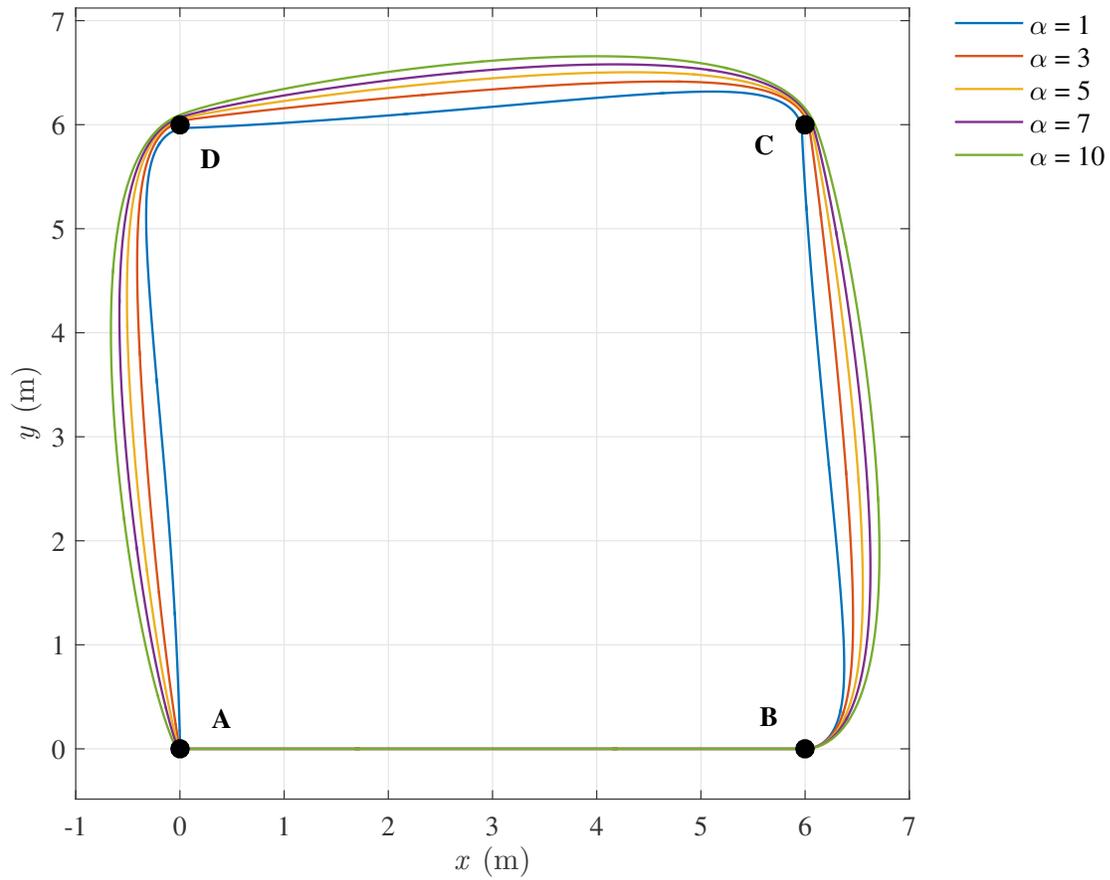


Figure 3.12: **Effect of α on the pedestrian trajectories.** The path followed by a pedestrian passing through the sequence of way-points A-B-C-D, for different values of the parameter α .

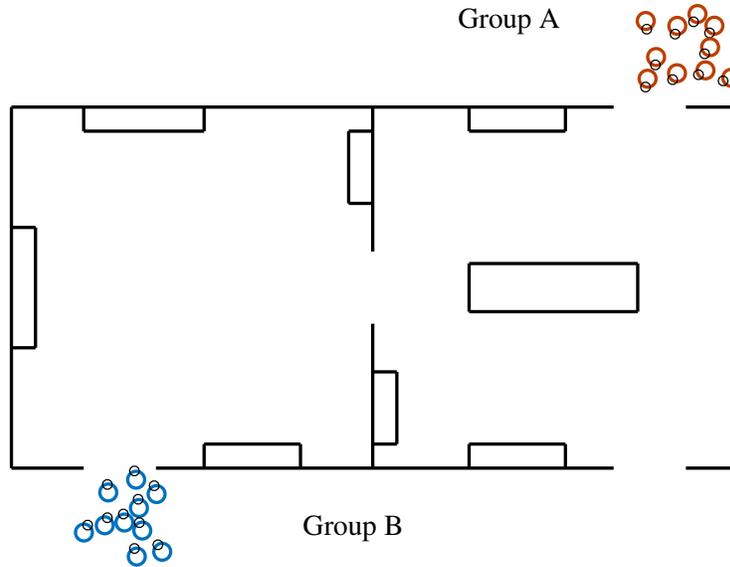


Figure 3.13: Scenario II of D2.1: Initial configuration of the two groups.

3.4 Comparison with the SFM at the Museum

We now discuss the benefit of the proposed HSFM in comparison with the SFM results presented in the deliverable D2.1, here succinctly reported for reference. The considered environment is composed of two communicating rooms, each of which contains four artworks on display. Three doors connect the rooms with the rest of the museum (see Figure 3.7). The objective of the pedestrians is to visit a selection of the pieces of the exhibition in a given order, while avoiding collisions with obstacles and/or other individuals. Once the visitors reach the selected artwork, they stop in front of it for a predefined amount of time, before moving to the next point of interest. This experiment, very similar to Scenario I of deliverable D2.1, has already been presented in Sec. 3.2, where it is shown that the introduction of group cohesion forces is able to effectively keep the group together.

Additional simulations involving more than one group, have been performed. The same experiment described in Scenario II of deliverable D2.1 has been replicated. In this case, two groups (named A and B), with 10 visitors each, walk in the museum, entering from two different doors, as shown in Figure 3.13. About 35 seconds later, the two groups simultaneously try to traverse (in opposite directions) the passage joining the two rooms (see Figure 3.13). In the deliverable D2.1 it was shown that the trajectories generated according to the traditional SFM result in a jam which causes a dispersion in the two groups (see Figure 3.14). On the contrary, the new kinematic model embedded in the HSFM, together with the group cohesion forces (with parameters $k_1^g = k_2^g = 200$ N in (2.19) and (2.20)), preserve the group cohesion after traversing the passage between the two rooms (see Figure 3.15), resulting in amore natural group behaviour.

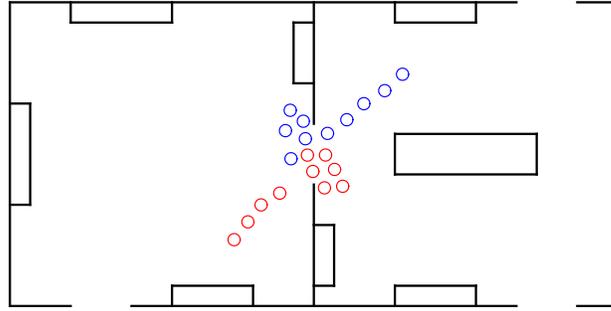
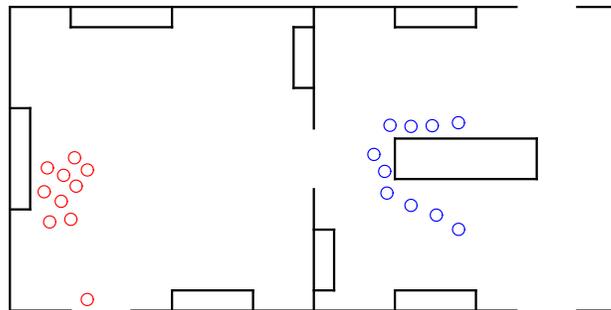
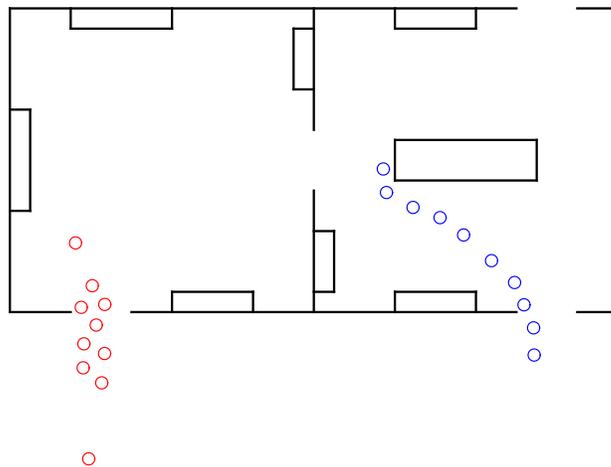
(a) $t = 35$ s(b) $t = 60$ s(c) $t = 70$ s

Figure 3.14: Scenario II and the SFM reported in D2.1: Three snapshots of the simulation taken at different time instants.

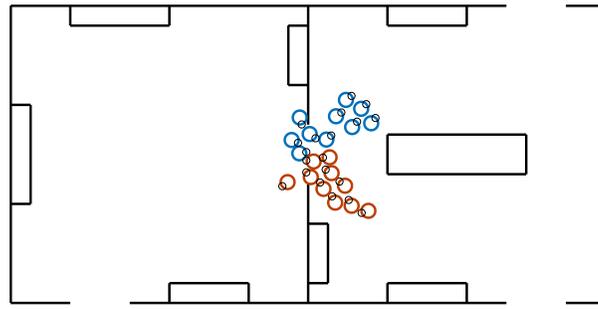
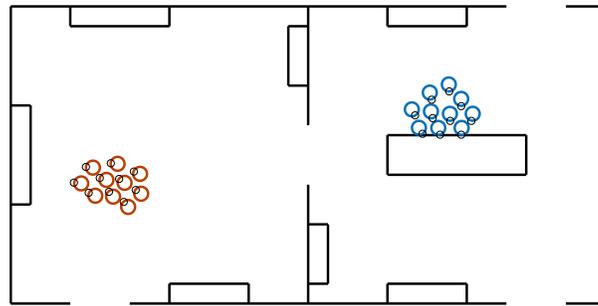
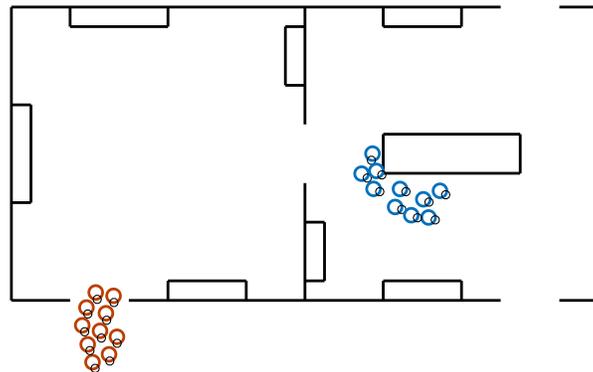
(a) $t = 35$ s(b) $t = 60$ s(c) $t = 70$ s

Figure 3.15: Scenario II IN D2.1 with the HSFM: Three snapshots of the simulation taken at different time instants.

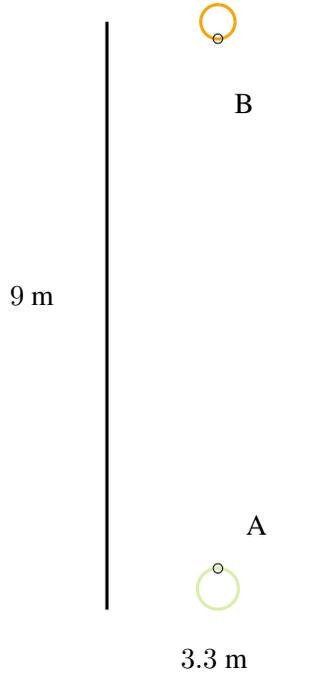


Figure 3.16: Experimental area and initial condition for Study 2 of D2.1, Sec. 4.2.

3.5 Comparison with Behavioural Dynamics

In this section, we evaluate the ability of the developed HSFM to reproduce the results of two experimental studies reported in previous deliverable D2.1 - “Human motion models (preliminary)”. In particular, the purpose of the second study (Study 2 in D2.1, Sec. 4.2) was to investigate the reaction of a pedestrian to movements of another one. Specifically, that experiment was carried out in order to observe pedestrian A’s behaviour when requested to walk towards the target agent (pedestrian B) who changed his position during A’s motion.

The experiment was performed in an area of $9.0 \times 3.3 \text{ m}^2$. Pedestrians A and B were positioned in front of each other, at a distance of 9 m, in the middle of the shortest edge of the arena (see Figure 3.16). At time $t = 0$, pedestrian A started to move towards B. After a time gap ranging from 1 to 5 seconds, pedestrian B started to move from its position to one of the two sides of the experimental area, until he reached the limit of the arena (see Figure 3.17-(c)). Both the time gaps and the side to which pedestrian B had to move were randomly chosen.

As reported in D2.1, that study did not show a significant difference in participants’ performance in relation to the side to which the target pedestrian moved. On the contrary, it highlighted a linear dependence between the participants’ walked time (defined as the time at which the participant’s trajectory begins deviating from a straight line) and the time gap after which the target pedestrian begins to move. It was observed a delay in the onset of participants’ reaction that ranged from 2.2 s (for a time gap of 5 s) to 3 s (for a time gap of 1 s). For example, this means that if the target starts to move at time $t = 1 \text{ s}$, then the pedestrian begins changing its direction at time $t = 4 \text{ s}$ (i.e., the walked time is 4 s).

In order to evaluate the capability of the HSFM to reproduce the behaviour observed in this experimental study, a simulation campaign consisting of 1000 runs has been performed. The radius and the mass of each pedestrian

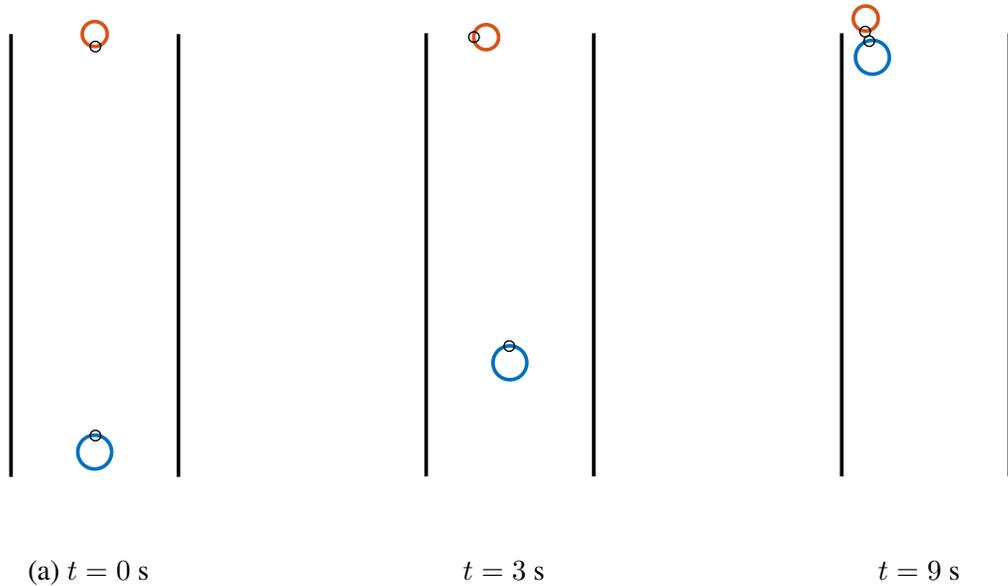


Figure 3.17: Three snapshots of a simulation run.

have been randomly generated in the intervals $[0.25, 0.35]$ m and $[60, 90]$ kg, respectively, assuming uniform distributions. Moreover the desired speed of the pedestrians have been generated randomly with uniform distribution in the interval $[1.1, 1.3]$ m/s. All the other parameters of the model have been chosen as reported in the beginning of this Chapter 3. Obviously, being the model completely symmetric, there is no difference in the performance in relation to the side to which pedestrian B moves. Moreover, as shown in Figure 3.18, the delay in the onset of the reaction of A to the movement of B grows approximately linearly with the time gap, in agreement with what was observed in the experimental study. Different simulation runs have been repeated by using different values for a parameter of the HSFM, namely, the characteristic time τ (see (2.12)). It is a measure of how fast the pedestrian velocity changes as a result of the social forces acting on it (the lower the τ , the more reactive the pedestrian). Values of τ ranging from 0.5 s to 3 s were tested. As shown in Figure 3.19, the average walked time increases with τ and it grows linearly with the time gap for each value of τ . In particular, when $\tau = 3$ s, the delay in the onset of the reaction of A to the movement of B is now about 3 s for a time gap of 1 s, and grows linearly with the time gap, decreasing of about 0.2 s for each step of the time gap, which is exactly what was obtained from experimental evidence in D2.1. These results, although by no means exhaustive, confirm the accuracy of the trajectories generated by the developed HSFM.

In the other experimental study presented in D2.1 (see Study 1 in D2.1.1, Sec. 4.2), a participant had to walk towards a fixed target while two actors were coming towards him from the opposite direction. The objective was to observe participants' behaviour at the point of interaction. Specifically, whether the participant passed either on one side (left or right) or between the two actors was taken into account. Results showed that in a large experimental area, several times the participants decided to pass through the pair of actors which corresponds to a sort of optimal trajectory (minimum length trajectory). In a smaller arena, the participants clearly preferred to avoid the pair of actors by deviating on the right from the optimal trajectory, in agreement with the cultural average avoidance side. By replicating this experiment with the proposed HSFM, it turns out that, no matter the size of the experimental area, whenever the starting point of the participant is between the position of the actors, the model makes the participant pass through them. This shows an expected limitation of current version of the proposed model, which does not embed cultural and social conventions yet.

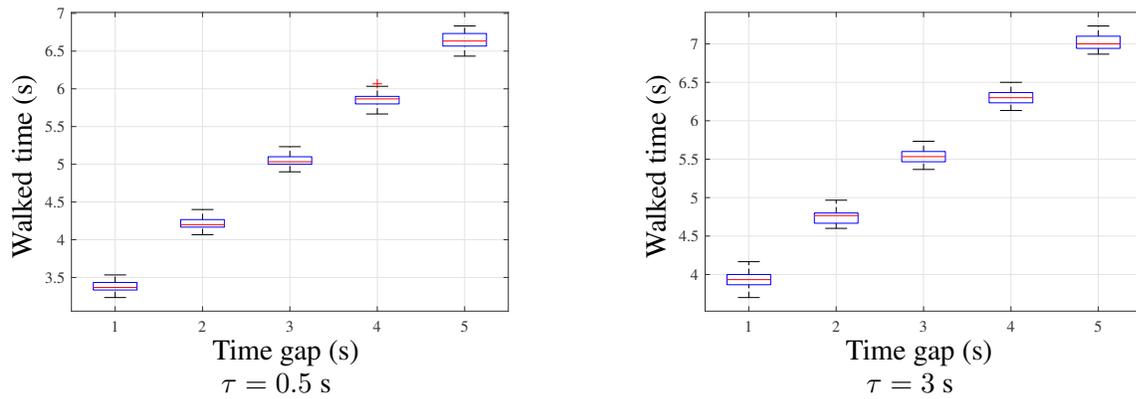


Figure 3.18: Walked time of pedestrian A as a function of the time gap between the start of pedestrian A and the movement of pedestrian B, with two different characteristic times τ . The red line represents the median, the upper and lower edges of the blue box the 25th and the 75th percentile respectively and the black whiskers extend to the most extreme data points.

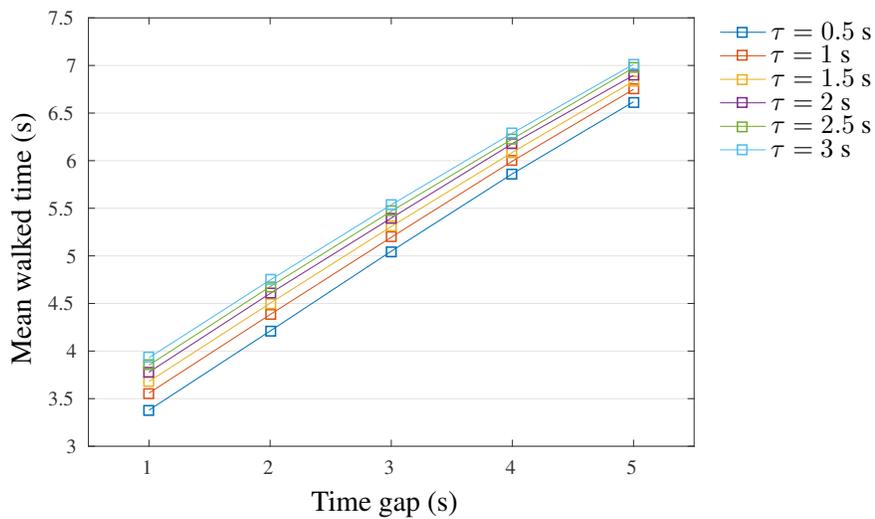


Figure 3.19: Mean walked time of pedestrian A as a function of the time gap between the start of pedestrian A and the movement of pedestrian B for different values of τ .

3.6 Comments

In this deliverable section the Headed Social Force Model and its characteristics have been presented and discussed in details. It has been shown that the HSFM enhances the traditional Social Force Model with the inclusion of the pedestrians' heading. A more complex model of the human dynamics is adopted, whose inputs are computed as suitable functions of the force terms resulting from the traditional Social Force Model. An optional force term has been introduced in order to model pedestrians moving together as a group. Numerical simulations show that considering the heading of the individuals improves the realism of the resulting trajectories, in both low pedestrian density scenarios and crowded environments.

The Headed Social Force Model developed in WP2 will be used within the ACANTO project to simulate pedestrian dynamics in the envisaged use cases. The inclusion of pedestrians' heading, as well as the addition of a force term for modelling the dynamics of people moving in a group, increase the flexibility and realism of the traditional SFM model. Ongoing work is focused on setting up a systematic procedure for tuning the model parameters, depending on the typology of pedestrians at hand (e.g., young people, fully autonomous seniors, or even older adults pushing the *FriWalk*). It is expected that a proper selection of the parameter values will allow the model to reproduce a variety of behaviours ranging from people taking a stroll to individuals walking in a hurry.

Chapter 4

On-the-fly Group Abstraction

Optimising the motion of groups of agents scales exponentially with the number of agents, so in D5.1 we proposed the use of a *group abstraction* social force model, to avoid the combinatorial complexity. Having now developed and tested the idea, in this chapter we report results of applying group abstraction to a dataset obtained by observing pedestrians in a crowded hotel lobby (ETH Zürich BIWI Walking Pedestrians dataset¹).

4.1 Background

The ACANTO Reactive Planner draws on the architecture and technology of the ‘short term planner’ developed in the DALi project [11]. In DALi the short term planner is user-centric and essentially selfish. Visual sensors attached to the user’s walker locate the user, fixed objects and pedestrians in the local environment. Then, using the “predictor-corrector” architecture shown in Fig. 4.1, the short term planner suggests an instantaneous direction that approximately maximises the probability of satisfying the users objectives, but *without* considering the objectives of other pedestrians, who might also be using the DALi platform.

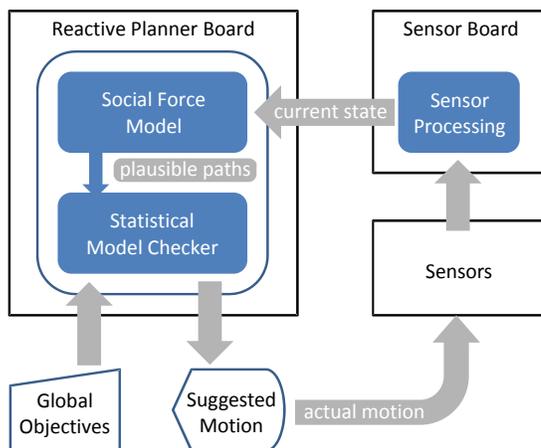


Figure 4.1: Common short-term / Reactive Planner architecture.

In the cases of both DALi and ACANTO, the position and (possibly zero) velocity of objects in the environment are identified by the sensors and used to parametrise the human motion model, i.e., the SFM, that uses stochasticity to account for unpredictable behaviour. In DALi, given this current snapshot of the dynamic environment, the short term planner stochastically simulates multiple future trajectories of the model up to some time horizon, given different hypothesised initial directions of the user. Each trajectory is validated against the user’s

objectives and thus the planner is able to estimate the probability of success for each hypothesised direction. The direction suggested to the user is that which both maximises the probability of success and minimises the deviation from a direct path to the next waypoint.

ACANTO is focused on social activities with groups, so the ACANTO Reactive Planner explicitly considers group motion. In contrast to the DALi short term planner, the ACANTO Reactive Planner must therefore optimise the motion of groups of users with respect to their personal objectives and with respect to group (social) cohesion. Sensor information obtained by each *FriWalk* user is curated and shared with the other users of the ACANTO platform, effectively giving each *FriWalk* a much wider view and much more data to handle. A naive extension of the DALi short term planner to optimise the trajectories of all *FriWalks* with respect to each other incurs an exponential increase in complexity because we must hypothesise the possible actions of one *FriWalk* with respect to the possible actions of all other *FriWalks*.

To address this problem we construct a group abstraction of a crowded environment, taking advantage of the well known phenomenon that pedestrians tend to move in groups and synchronise their motion, even if not known to each other. Such groups are not necessarily related to the notional group of participants in an activity, but are created spontaneously when pedestrians negotiate a crowded environment [25]. These de facto groups are strongly predictive of human motion in the short term, with pedestrians essentially maintaining their relative positions within the group. Hence we may model the behaviour of such groups as a single entity and make the reasonable assumption that individuals will walk in semi-rigid formation. Note that we do not enforce such rigid formations, but merely use this assumption to predict short term future behaviour, for the purpose of estimating the probability of satisfying objectives. The pedestrians are completely free to move around within the group or to leave the group entirely.

For medium to long term prediction we need to know the intentions of pedestrians. That is, we need to know where they're trying to go and with whom they're travelling. This information is well defined for participants of an ACANTO activity, but not for unknown pedestrians. We therefore observe the past behaviour of pedestrians to infer their future behaviour, using the notion of a pedestrian's *trace*.

The frame rate of our visual sensor technology makes it conceivably easy to infer the traces of pedestrians at a visual level, however there will inevitably be natural occlusions that prevent long term visual continuity. Combining the output of sensors on multiple *FriWalks* incurs the additional challenge of identifying pedestrians who leave the view of one sensor and appear in another. Pedestrians may also appear and disappear as a result of sensors being obscured, because of communication unreliability or because users just leave. Our trace inference algorithm, described in Section 4.1.1, therefore makes use of a human motion model to reliably link the *behaviour* of pedestrians over much longer times than that between successive video frames. The presented trace inference algorithm (Algorithm 1) makes minimal a priori assumptions about its input, but in practice will take advantage of whatever other information is available (e.g., high reliability visual continuity).

Having inferred a set of traces from multiple observations of agents in the environment, it is necessary to infer their de facto groups. Instantaneous physical proximity is not a sufficient indicator, since two close pedestrians could actually be trying to get away from each other. Our group inference algorithm (Algorithm 2) thus uses a notion of proximity that includes both position and velocity: if pedestrians are physically close, walking at the same speed in the same direction, it is reasonable to assume (by definition) that they are walking together. To find an optimal partition of traces into groups, our algorithm uses *k*-means clustering [36]. The full technical aspects of are outlined in Section 4.1.2.

Identifying de facto groups allows us to plan motion at a more efficient level of abstraction. When hypothesising the alternative directions for a number of users of the platform, we believe that it is a reasonable compromise to only hypothesise the overall motion of the groups to which they belong. We feel it is not necessary to consider all the possible combinations of suggestions to those within the same group given that, by virtue of how we define a group, their motion is strongly correlated. Note that suggestions are nevertheless tailored to the actual position of an individual within the group, in order to maintain its "social" structure.

4.1.1 Trace inference

Our trace inference algorithm (Algorithm 1) constructs sets of active and inactive traces, where a trace is a sequence of timestamped observations of the position and velocity of pedestrians detected by the sensors. Active traces are those for which the algorithm has reliably inferred continuity and/or there is currently a pedestrian in the field of view of the sensors (a trace may consist of a single observation). Inactive traces are those for which the algorithm could find no valid continuation, so there is no current view of the corresponding pedestrian. Since inactive traces do not contain a current point, further trace inference applies only to active traces. In practice a trace may become inactive due to the obfuscation of a sensor or failure of communications. Our algorithm assumes that the human motion model makes valid predictions up to a maximum time interval of Δt_{max} , hence traces whose endpoints are older than this become inactive.

The initial set of active traces comprise the initial set of observations, while the initial set of inactive traces is empty. Each step of the algorithm appends new observations to active traces or starts new active traces with single observations that cannot be assigned to existing active traces. Active traces to which no new point can be assigned become inactive once their endpoint is older than Δt_{max} with respect to the current time. Using the human motion model, each iterative step of the algorithm generates a set of projections to the current time from the the endpoints of the active traces, then tries to match the projected points to the new observations by finding a minimum distance assignment according to distance metric D_{trace} (see Section 4.2.2).

To solve the *assignment problem* of optimally matching new observations to projected points, our implementation makes use of an $\mathcal{O}(n^3)$ implementation of the Hungarian method [33]. We first construct a square matrix of distances between projected points and the new observations, according to distance metric D_{trace} . The numbers of projected points and active traces may be different, so the smaller of the two sets is padded with dummy entries whose distance to all members of the other set is, by convention, made equal to the maximum observed distance between non-dummy entries. The Hungarian method is guaranteed to find an assignment that minimises the overall distance between the projected points and new observations, but not all individual assignments are close enough to be accepted and some assignments include non-existent (dummy) entries. Observations whose assignment has a distance up to threshold θ_{trace} are appended to the ends of the corresponding active traces. Observations whose assignment has a distance greater than θ_{trace} become the initial points of new active traces. Assignments involving dummy entries are discarded.

4.1.2 Group abstraction

To infer groups we use k -means clustering [36] over the set of active traces. The k -means algorithm partitions a set of $n \geq k$ data points into $\leq k$ clusters, according to a specified metric over the points. Although the k -means problem is computationally hard, there are good heuristics that make it expedient for our on-the-fly inference application (e.g., the k -means++ algorithm [5] of the Apache Commons Math library¹).

Given a specified value of k , the k -means algorithm first defines a set of tentative cluster means (centroids). This may be done randomly or heuristically. It then executes a series of alternating assignment and update steps that (re-)allocate points to clusters, until further steps produce no modifications. Assignment steps assign data points to clusters with the nearest mean, according to a problem-specific distance metric. Update steps re-calculate the means of the clusters. The k -means algorithm is guaranteed to terminate, however the results are generally local optima that are dependent on the initialisation. Heuristics therefore focus on finding good initialisations.

To find an optimal group abstraction we first define a group distance metric, D_{group} , and define a group cohesion threshold distance, $\theta_{group} > 0$, that specifies the maximum permissible distance from the centroid of a group. Note that the group distance metric is more concerned with similar motion than physical proximity, so being close to the centroid implies primarily that members of the group are moving in the same way. We do not know

¹commons.apache.org/proper/commons-math/

Algorithm 1: Trace inference

Let \mathcal{H} be the human motion model

Let Act be the current set of active traces (initialised with the first set of observations)

Let $Inact$ be the current set of inactive traces (initially empty)

Let Δt_{max} the maximum trace projection time

Let D_{trace} be the trace matching distance metric

Let θ_{trace} be the trace matching threshold distance

while there are new observations **do**

 Let Obs be the set of new observations at current time t

 Let $Old \subseteq Act$ be the set of traces whose end points are older than $t - \Delta t_{max}$

$Act \leftarrow Act \setminus Old$

$Inact \leftarrow Inact \cup Old$

 Let $Proj$ be the set of points generated by projecting the ends of all traces in Act to time t using \mathcal{H}

if $|Obs| < |Proj|$ **then**

 └ Pad Obs with dummy entries so $|Obs| = |Proj|$

else if $|Proj| < |Obs|$ **then**

 └ Pad $Proj$ with dummy entries so $|Obs| = |Proj|$

 Construct a square matrix $Dist$ of distances between Obs and $Proj$ using D_{trace}

 (set the distance to or from any dummy entry to be the maximum non-dummy distance)

 Apply the Hungarian Method to $Dist$ to find a minimum distance set of assignments $Assign$

 Remove from $Assign$ all assignments involving dummy entries

for $((trace, projection), observation) \in Assign$ **do**

if $D_{trace}(projection, observation) \leq \theta_{trace}$ **then**

 └ Append $observation$ to $trace$

else

 └ Add $observation$ to Act

in advance the optimal number of clusters (i.e., the optimal value of k), so we iterate from $k = 1$ to $k = |Act|$, where Act is the set of traces to cluster, stopping when the set of clusters are sufficiently cohesive. For the purposes of efficient motion planning we would like k as small as possible, so the algorithm aims to find the fewest number of sufficiently cohesive clusters. To judge the cohesiveness of a cluster, the algorithm calculates the distance between each member and the cluster’s centroid using D_{group} . If any member of any cluster is too far from its corresponding centroid, the current set of clusters is abandoned and a new set is generated using $k \leftarrow k + 1$. The algorithm is guaranteed to terminate because when $k = |Act|$, all clusters contain a single element whose distance from their corresponding centroid is guaranteed to be $< \theta_{group}$.

Algorithm 2: Group inference

Let Act be the set of active traces

Let D_{group} be the group distance metric

Let θ_{group} be the group cohesion threshold distance

Set $done \leftarrow false$

Set $k \leftarrow 1$

while $\neg done \wedge k \leq |Act|$ **do**

 Perform k -means clustering on Act

 Let $Clust$ be the resulting set of $\leq k$ clusters of traces

$done \leftarrow true$

for $cluster \in Clust$ **do**

 Let $centroid$ be the centroid of $cluster$

for $trace \in cluster$ **do**

if $D_{group}(trace, centroid) > \theta$ **then**

$done \leftarrow false$

return $Clust$

Although the k -means algorithm partitions data points into Voronoi cells that are disjoint in the multi-dimensional space of the distance metric, groups may physically overlap. This arises, for example, when two groups walking in opposite directions pass through each other. This phenomenon does not occur at the level of individuals and is therefore not considered in the original social force model, however it is nevertheless possible to model it with forces in the SFM framework. To do this we reduce the repulsive social force between the groups (\mathbf{f}^{soc} in (2.2)) and use the physical component (\mathbf{f}^{ph} in (2.2)) to model the “friction” between them. To accurately model the momentum of different sized groups, the mass term (m in (2.1)) will be the sum of the masses of the individuals. The latency parameter (τ in (2.1)) is also likely to be greater for groups, however some of the latency is simply accounted for by the increased mass.

Figure 4.2 illustrates an hypothetical scenario of groups diverging and coalescing over time, noting that these phenomena are also evident in the automatically-generated visualisations of the output of our on-the-fly algorithms applied to real observations. The groups labelled 1, 2 and 3 are assumed to comprise a single pedestrian. Group 4 contains two pedestrians who initially just happen to be walking close to one another in the same direction. Pedestrians 1 and 3 know each other, so they move closer and eventually form group 5. Pedestrian 2 knows one of the members of group 4, so they also move closer to one another and eventually form group 6. The other member of group 4 is just passing through and eventually leaves the view of the sensors (smaller dashed circle). At some time before then, however, he gets very close to pedestrians 2 and 3 (in the region denoted by the larger dashed circle), but no new group is detected because they are all travelling in different directions. The members of groups 5 and 6 are actually part of the same activity, so the system guides them closer, thus eventually forming group 7.

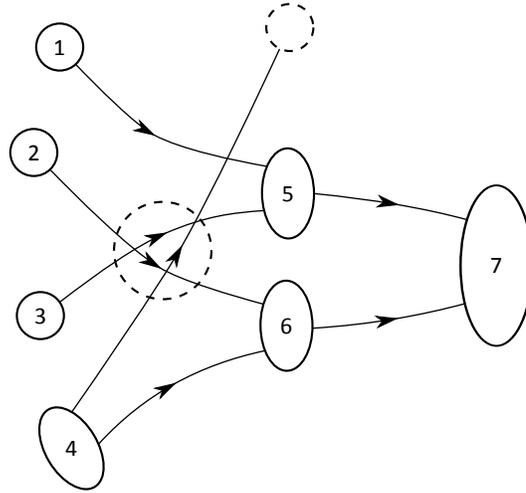


Figure 4.2: Groups diverging and coalescing.

Graph-theoretic group behavioural templates

The SFM uses stochasticity to model unpredictable behaviour, but previous work [9] has identified “behavioural templates” that can deterministically account for certain types of non-smooth interactions between pedestrians. Replacing stochastic uncertainty with deterministic templates can potentially improve the predictive power and efficiency of the human motion model. In the present context we are particularly interested in predicting when and how groups may pass by, split or pass through each other. Such behaviour does not, of course, apply to individual pedestrians, but is common with groups of pedestrians. For example, Figures 4.4 and 4.6 illustrate apparently similar scenarios where one moving group chooses to split to avoid a collision with a stationary group (Figure 4.4), while a similar moving group chooses to stay together to avoid the same stationary group (Figure 4.6). We propose to use graph-theoretic metrics of clustering, cohesion and centrality (see, e.g., [43, Chap. 7]) to identify the conditions under which the different group behaviours are likely and to identify if and how groups will fragment. Such graph-theoretic metrics are efficient to calculate, with minimal additional computational cost above the inherent cost of the SFM and k -means clustering.

4.2 Modelling real pedestrians with Group Abstraction

In this section we give results of applying our on-the-fly algorithms to the ETH Zürich BIWI walking pedestrians dataset¹. We thus demonstrate how our approach infers traces and groups from observations of real pedestrians and justify its use as a means to simplify motion planning of groups or individuals in complex crowded environments.

4.2.1 Dataset

The chosen dataset comprises hand-annotated motion-capture observations of pedestrians in two environments: a hotel lobby and a corridor within the ETH premises. The annotations link the observations into traces and groups, which our algorithms do automatically. We make use of the original annotations only to compare with the automated annotations, noting here that our on-the-fly algorithms successfully identify all the traces and groups identified by hand. In both environments the observations are made using a fixed camera directly overhead, with an original video frame rate of 25fps (0.04s between frames). Here we focus on the hotel data, which contains more interesting and varied interactions between pedestrians. In what follows we thus use the

term ‘dataset’ to refer exclusively to the hotel data.

The dataset contains observations sampled at 2.5fps (0.4s), with observations divided into 27 contiguous intervals separated by more than 0.4s. We suppose that the times not covered in the dataset are excluded because they contain no moving pedestrians, however we believe that some of the omitted frames nevertheless contained stationary pedestrians. Moving pedestrians must avoid stationary pedestrians, so for the purpose of motion planning they cannot be ignored. We illustrate this in the results described below.

4.2.2 Distance metrics

The presented results make use of two different distance metrics that are based on both position and velocity. In the case of the trace inference algorithm (Algorithm 1), to decide whether a projected point is “close” to an actual observation the metric D_{trace} is based on Euclidean distance in the 4-dimensional space of (x, y) position and *discounted* (x, y) velocity. Precisely, we divide the velocity dimension by 2, thus making position more significant than velocity when inferring traces. The value of 2 was chosen empirically and found to work well, but is not critical. Projected points have an implicit velocity that should be similar to that of the observation, but giving too great an emphasis to velocity-matching risks erroneously cross-linking the trajectories of pedestrians walking in formation.

In the case of the group inference algorithm (Algorithm 2), the metric D_{group} is based on Euclidean distance in the 4-dimensional space of *discounted* (x, y) position and (x, y) velocity. Hence, in contrast to trace inference, we make velocity more significant than position to infer groups using k -means clustering. To justify our group abstraction motion planning, we require that pedestrians in a group are moving in a similar way, hence velocity is clearly important. We also require that pedestrians in a group are proximal, but this is less important. A discount factor of 2 was, once again, found empirically to be good and not critical. In fact, we achieved very similar results with no discounting and usable results when we discounted velocity instead of position.

We note here that our chosen distance metrics are “memoryless”, meaning that they consider only the instantaneous position and velocity of pedestrians. We make this choice simply for the purpose of exposition—the visualisations thus show how our algorithms interpret the instantaneous behaviour of the pedestrians. In practise, however, we believe that it will be beneficial to use “history-dependent” metrics that include the previous position and velocity of pedestrians. Thus, for example, two pedestrians would be considered closer with respect to grouping if they were grouped immediately before. This approach smooths out (avoids) the case of groups appearing to momentarily split and motivates the need to link observations into traces.

4.2.3 Results

Figures 4.3, 4.4, 4.5 and 4.6 visualise traces and group abstractions for four intervals from the dataset, produced automatically by our on-the-fly algorithms. The x and y axes give the spatial coordinates (in metres) of the groups, with respect to the origin defined in the dataset. The traces inferred by our algorithm are denoted by black lines, marked at their start by blue discs and at their end by red discs. Each black circle denotes an inferred group at a particular time point. A (trivial) group may consist of a single pedestrian (the smallest circle). Most non-trivial groups in these figures consist of pairs of pedestrians (medium-sized circles), with group containing three pedestrians (largest circle) evident in only Figures 4.3 and 4.4. No groups greater than three pedestrians were identified by the original hand-annotations of the dataset, nor by our automatic annotation. The relative proportions of group sizes in both annotations are consistent with long-standing group size distribution results [25]. Note that the figures abstract away from time, such that the starts and ends of different traces are not necessarily synchronous, while crossing traces do not necessarily imply a collision.

Figure 4.3 The figure visualises an interval of 47.6 seconds starting at time 34.8 seconds with respect to the first frame of the dataset. The traces and groups identified by our algorithms include all those identified in the

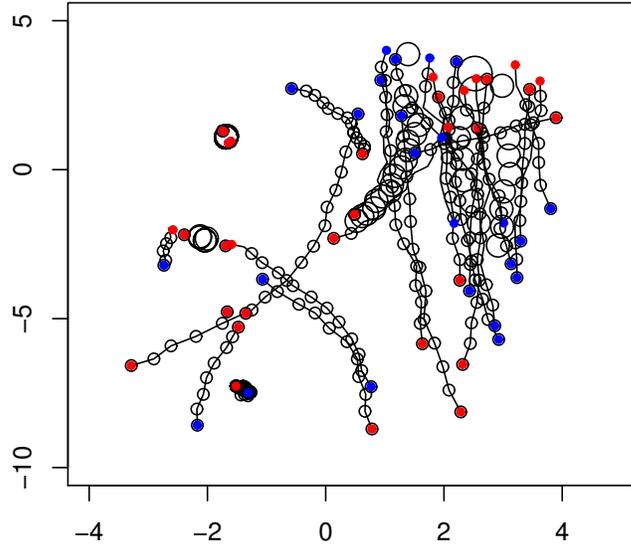


Figure 4.3: Group abstraction of interval 3 of hotel dataset (34.8 to 82.4s).

hand annotation of the original dataset, as well as other groupings that are useful for motion planning. In particular, our algorithms identify stationary groups near coordinates $(-1.7, 1.1)$, $(-2.1, -2.3)$ and $(-1.4, -7.4)$, where the “group” near $(-1.4, -7.4)$ is a trivial group comprising a single pedestrian. We illustrate the effect of stationary groups on moving groups in Figures 4.4 to 4.6, which are described below.

Figure 4.4 The figure visualises an interval of 8.4 seconds starting at time 110.8 seconds with respect to the first frame of the dataset. We first note a stationary group around coordinate $(1.5, -9)$. Not apparent from the figure, they first appear at 113.6 seconds and (approximately) maintain their positions for 5.6 seconds until the end of the interval. Since both their positions and motion are very close, our algorithm correctly identifies these pedestrians as a group. In the original dataset, however, they are not identified as such. We observe that the existence of this stationary group has a significant effect on the motion of the group starting near coordinate $(2.3, 3.3)$. For some reason, the moving group heads directly *towards* the stationary group and then splits near coordinate $(1.7, -4.5)$ to avoid a collision. Importantly, however, up to the point at which the group splits, the group abstraction of the two moving pedestrians provides a good prediction of their behaviour. Following the split, the group abstraction *continues* to provide a good prediction of the moving pedestrians’ behaviour because it detects that they are no longer moving together. The other groups in this figure, within the box created by $x \in [2, 4]$ and $y \in [-8, -3]$, are not identified by the hand-annotation of the original dataset, but are nevertheless valid and useful for the purposes of motion planning.

Figure 4.5 The figure visualises an interval of 23.2 seconds, 58 seconds after the interval illustrated in Figure 4.4. We see that there is once again a stationary group near coordinate $(2.8, -6.9)$, which is not identified in the hand annotations. Given that this group also exists in the three intervening intervals between those of Figures 4.4 and 4.5, we presume that it is the same group. We do not make use of this fact here, but such information can be useful in predicting the intentions of pedestrians. The instantaneous intentions of pedestrians, required by the *driving* or *desired* velocity term \mathbf{v}^0 in (2.1), may not reflect the known medium term objectives of known pedestrians and are even more difficult to infer if the pedestrians are unknown to the system. Knowing that an observed group is more than just instantaneous and ad hoc can therefore help to improve the inference. The motion of the group starting near coordinate $(2.9, -9.5)$ is interesting in comparison to the long group trajectories in the other figures. In this figure the pedestrians are walking slower (closer observations made at

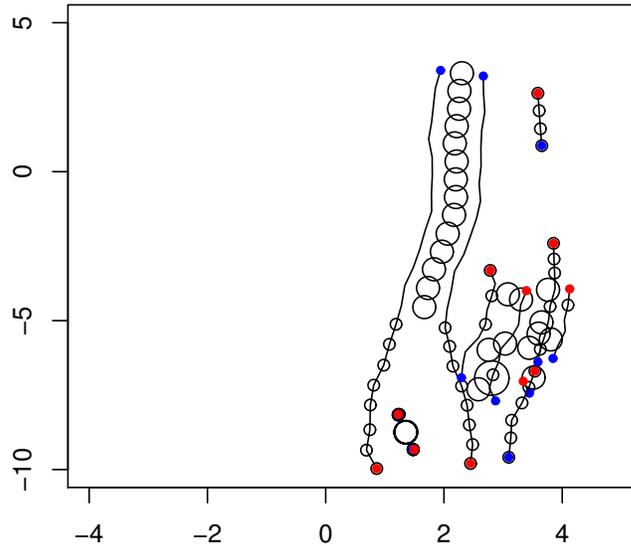


Figure 4.4: Group abstraction of interval 6 of hotel dataset (110.8 to 119.2s).

the same frequency as the other figures), closer together and while not following a smooth path, they appear to very closely maintain their separation distance. We might infer from this that they have a strong social bond. From the perspective of motion planning they are effectively moving as a single agent, thus fully justifying the group abstraction.

Figure 4.6 The figure visualises an interval of 8.8 seconds, 26.8 seconds after the interval illustrated in Figure 4.5. The stationary group seen in Figures 4.4 and 4.5 seems also to be present in this interval, however the group first appears at time 232.8 seconds, 5.6 after the beginning of the interval, and is not present in the single interval between those illustrated in Figures 4.5 and 4.6. From only the recorded observations in the dataset, we cannot disambiguate the possibilities that the group moved away and then returned to the same spot, that the camera was temporarily obscured or that there was a data processing error. Temporary occlusions and measurement errors are to be expected, especially when inferring a global view from ground-level cameras mounted on individual *FriWalks* in the case of ACANTO, and do not invalidate the approach unless the occlusions and errors persist indefinitely.

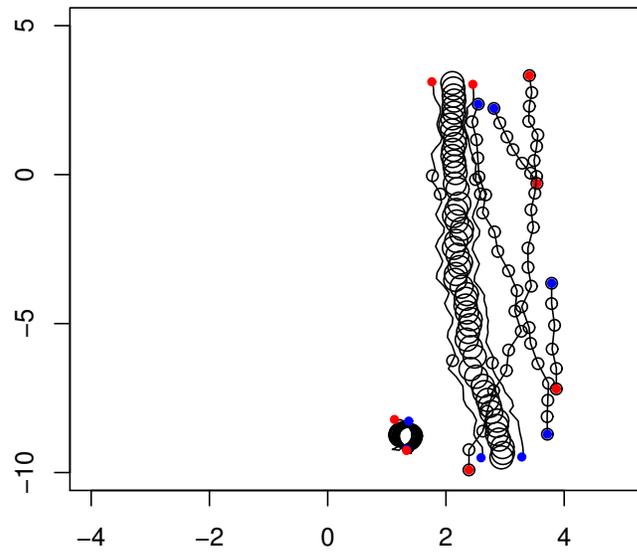


Figure 4.5: Group abstraction of part 10 of hotel dataset (177.2 to 200.4s).

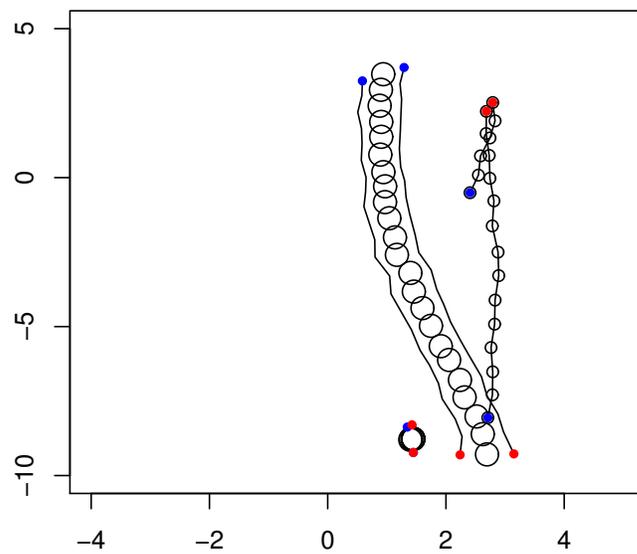


Figure 4.6: Group abstraction of part 12 of hotel dataset (227.2 to 236s).

Chapter 5

Relation with the other Work Packages

In what follows, we propose a collection of potential links other work packages could benefit from.

The results gained in this deliverable are beneficial, to some extent, to the work carried out in WP3 - “Perception of Users and Environment”, in particular to what concerns the collaborative platform localisation and for the interpretation of the social context. Indeed, for collaborative localisation (Task 3.1.3), the knowledge coming from the way in which human beings move as a social group in the environment helps in identifying which sensor can be used to detect other persons in the group and, hence, increase the robustness and accuracy of the collaborative localisation algorithm. In the other direction, the high level knowledge of social groups obtained in WP3 can be used to inform the human motion model and group abstraction algorithm about pedestrians’ intentions.

Moreover, the model of motion of social groups, from a geometric view-point, helps in the interpretation of the social context (Task 3.3). From the platform view (Task 3.3.1), the motion and the trajectory about the behaviour of each individual is fundamental for the learning of the high-level behavioural models as well as model the dynamics and the social relationships between the people in the scene. In particular, the knowledge gained in this deliverable helps in monitor the activities of people (individuals and groups) and the roles they play in the scene (leaders, followers, etc.), and analyse the group dynamics by estimating the group cohesion. Furthermore, since Task 3.3.2 deals with the interpretation of the site-wide social context, the knowledge of the motion models is the basic pillar to identify the behaviour of crowds in terms of forming, movement, and disaggregation. There is thus also scope for cross-fertilisation between WP3 and the graph-theoretic metrics we use to infer group structure.

WP5 - “Execution Support of Social Activities” is the principal beneficiary of the results of this deliverable and, in particular, the Reactive Planner of Task 5.2. The Reactive Planner makes extensive use of the human motion model, while our group abstraction approach is specifically intended to increase its efficiency. Besides the possibility of planning deviations for individuals according to the crowd density in the surroundings (hence, embedding the blend between the SFM and the nonholonomic behaviour), when the Reactive Planner works for group of users and respect their roles inside the group. Notice that the choice of modelling the group cohesion by means of a cohesive force reduces the complexity of the Reactive Planner, that has to plan just for the group centre of mass to steer the whole formation on the planned deviation. The found models are also important to identify the presence of other pedestrians or other group of pedestrians and then predict their motions and reactions to the re-planned paths.

The notion of the motion models, for individuals and groups, is also of paramount importance for the Activity Planner (Task 5.1), which is responsible to produce a concatenation of control actions to carry out correctly an activity. For instance, the knowledge gained on the nonholonomic behaviour is applied directly to the planner to generate paths that satisfy such a constraint by construction. Moreover, the size and density of the group and the type of activity may affect the optimal paths generated by the planner in case of group activity.

Finally, the models here described are tightly connected to the monitor of the activities (Task 5.3). In fact, the

role of the monitor is to evaluate the compliance of the plan with the requirements at the individual level and, when required, at the social level. The knowledge of the motion models allows the detection of misbehaving or uncooperative group members, and hence generate a failure for the individual and/or group activity.

Since the individual or group motion patterns have an impact on the manoeuvres the *FriWalk* has to generate, the findings of this deliverable will define the motion patterns considered as in put in WP6 - “Design of Robotic Personal Devices”. This connection is mainly enforced by the role of Task 2.3, i.e. the bridge between the models and the control actions.

Bibliography

- [1] Gianluca Antonini, Michel Bierlaire, and Mats Weber. Discrete choice models of pedestrian walking behavior. *Transportation Research Part B: Methodological*, 40(8):667 – 687, 2006.
- [2] G. Arechavaleta, J.-P. Laumond, H. Hicheur, and A. Berthoz. An optimality principle governing human walking. *IEEE Transactions on Robotics*, 24(1):5–14, 2008.
- [3] Gustavo Arechavaleta, Jean-Paul Laumond, Halim Hicheur, and Alain Berthoz. On the nonholonomic nature of human locomotion. *Autonomous Robots*, 25(1-2):25–35, 2008.
- [4] Michael Argyle and Janet Dean. Eye-contact, distance and affiliation. *Sociometry*, pages 289–304, 1965.
- [5] David Arthur and Sergei Vassilvitskii. k-means++: The advantages of careful seeding. In *Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 1027–1035. Society for Industrial and Applied Mathematics, 2007.
- [6] Maren Bennewitz, Wolfram Burgard, Grzegorz Cielniak, and Sebastian Thrun. Learning motion patterns of people for compliant robot motion. *The International Journal of Robotics Research*, 24(1):31–48, 2005.
- [7] 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(3-4):507 – 525, 2001.
- [8] James S. Coleman and John James. The Equilibrium Size Distribution of Freely-Forming Groups. *Sociometry*, 24(1):pp. 36–45, 1961.
- [9] Alessio Colombo, Daniele Fontanelli, Dhaval Gandhi, Antonella De Angeli, Luigi Palopoli, Sean Sedwards, and Axel Legay. Behavioural Templates Improve Robot Motion Planning with Social Force Model in Human Environments. In *Proc. IEEE Int. Conf. on Emerging Technologies & Factory Automation (ETFA)*, pages 1–6, Cagliari, Italy, 10-13 Sep. 2013. IEEE.
- [10] Marco Costa. Interpersonal distances in group walking. *Journal of Nonverbal Behavior*, 34(1):15–26, 2010.
- [11] DALi project web site. <http://www.ict-dali.eu/dali>, November 2011.
- [12] N.E. Du Toit and J.W. Burdick. Robot motion planning in dynamic, uncertain environments. *IEEE Transactions on Robotics*, 28(1):101–115, Feb 2012.
- [13] Francesco Farina, Daniele Fontanelli, Andrea Garulli, Antonio Giannitrapani, and Domenico Prattichizzo. When Helbing Meets Laumond: The Headed Social Force Model. In *Proc. IEEE Int. Conf. on Decision and Control (CDC)*, pages 3548–3553, Las Vegas, Nevada, US, Dec. 2016. IEEE.

- [14] Francesco Farina, Daniele Fontanelli, Andrea Garulli, Antonio Giannitrapani, and Domenico Prattichizzo. Walking Ahead: The Headed Social Force Model. *PLOS ONE*, 12(1):1–23, 01 2017.
- [15] G. Ferrer and A. Sanfeliu. Proactive kinodynamic planning using the extended social force model and human motion prediction in urban environments. In *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1730–1735, Sept 2014.
- [16] Tamar Flash and Neville Hogan. The coordination of arm movements: an experimentally confirmed mathematical model. *The journal of Neuroscience*, 5(7):1688–1703, 1985.
- [17] Ning Guo, Jian-Xun Ding, Xiang Ling, Qin Shi, and Imamura Takashi. The walking behavior of pedestrian crowd under impact of static and movable targets. *The European Physical Journal B*, 86(7), 2013.
- [18] Wei Guo, Xiaolu Wang, and Xiaoping Zheng. Lane formation in pedestrian counterflows driven by a potential field considering following and avoidance behaviours. *Physica A: Statistical Mechanics and its Applications*, 432(0):87 – 101, 2015.
- [19] Dirk Helbing, Lubos Buzna, Anders Johansson, and Torsten Werner. Self-Organized Pedestrian Crowd Dynamics: Experiments, Simulations, and Design Solutions. *Transportation Science*, 39(1):1–24, 2005.
- [20] Dirk Helbing, Illés Farkas, and Tamás Vicsek. Simulating dynamical features of escape panic. *Nature*, 407:487–490, September 2000.
- [21] Dirk Helbing and Anders Johansson. Pedestrian, crowd and evacuation dynamics. In Robert A. Meyers, editor, *Encyclopedia of Complexity and Systems Science*, pages 6476–6495. Springer New York, 2009.
- [22] Dirk Helbing and Peter Molnár. Social force model for pedestrian dynamics. *Phys. Rev. E*, 51:4282–4286, 1995.
- [23] LF Henderson. The statistics of crowd fluids. *Nature*, 229:381–383, 1971.
- [24] H. Hicheur, S. Vieilledent, M.J.E. Richardson, T. Flash, and A. Berthoz. Velocity and curvature in human locomotion along complex curved paths: a comparison with hand movements. *Experimental Brain Research*, 162(2):145–154, 2005.
- [25] John James. The Distribution of Free-Forming Small Group Size. *American Sociological Review*, 18(5):569–570, 1953.
- [26] Anders Johansson, Dirk Helbing, and Pradyumn K Shukla. Specification of the social force pedestrian model by evolutionary adjustment to video tracking data. *Advances in Complex Systems*, 10(supp02):271–288, 2007.
- [27] Fredrik Johansson, Anders Peterson, and Andreas Tapani. Waiting pedestrians in the social force model. *Physica A: Statistical Mechanics and its Applications*, 419:95 – 107, 2015.
- [28] Chris L Kleinke. Gaze and eye contact: a research review. *Psychological bulletin*, 100(1):78, 1986.
- [29] Mark Knapp, Judith Hall, and Terrence Horgan. *Nonverbal communication in human interaction*. Cengage Learning, 2013.
- [30] Tobias Kretz, Anna Grünebohm, Maike Kaufman, Florian Mazur, and Michael Schreckenberg. Experimental study of pedestrian counterflow in a corridor. *Journal of Statistical Mechanics: Theory and Experiment*, 2006(10):P10001, 2006.

- [31] Tobias Kretz, Anna Grünebohm, and Michael Schreckenberg. Experimental study of pedestrian flow through a bottleneck. *Journal of Statistical Mechanics: Theory and Experiment*, 2006(10):P10014, 2006.
- [32] H. Kretschmar, M. Kuderer, and W. Burgard. Learning to predict trajectories of cooperatively navigating agents. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4015–4020, May 2014.
- [33] H. W. Kuhn. The Hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, 2(1-2):83–97, 1955.
- [34] Benjamin Kuipers. The spatial semantic hierarchy. *Artificial intelligence*, 119(1):191–233, 2000.
- [35] Francesco Lacquaniti, Carlo Terzuolo, and Paolo Viviani. The law relating the kinematic and figural aspects of drawing movements. *Acta Psychologica*, 54(1 - 3):115 – 130, 1983.
- [36] James MacQueen. Some methods for classification and analysis of multivariate observations. In L. M. Le Cam and J. Neyman, editors, *Proceedings of the fifth Berkeley symposium on mathematical statistics and probability*, volume 1, pages 281–297, 1967.
- [37] Francisco Madrigal, Jean-Bernard Hayet, and Mariano Rivera. Motion priors for multiple target visual tracking. *Machine Vision and Applications*, 26(2-3):141–160, 2015.
- [38] Riccardo Mazzon and Andrea Cavallaro. Multi-camera tracking using a multi-goal social force model. *Neurocomputing*, 100(0):41 – 50, 2013. Special issue: Behaviours in video.
- [39] Clark McPhail and Ronald T Wohlstein. Using film to analyze pedestrian behavior. *Sociological Methods & Research*, 10(3):347–375, 1982.
- [40] Katja Mombaur, Anh Truong, and Jean-Paul Laumond. From human to humanoid locomotion – an inverse optimal control approach. *Autonomous robots*, 28(3):369–383, 2010.
- [41] Mehdi Moussaïd, Niriasca Perozo, Simon Garnier, Dirk Helbing, and Guy Theraulaz. The Walking Behaviour of Pedestrian Social Groups and Its Impact on Crowd Dynamics. *PLoS ONE*, 5(4):e10047, April 2010.
- [42] Mehdi Moussaïd, Niriasca Perozo, Simon Garnier, Dirk Helbing, and Guy Theraulaz. The walking behaviour of pedestrian social groups and its impact on crowd dynamics. *PloS one*, 5(4):e10047, 2010.
- [43] Mark Newman. *Networks: An Introduction*. Oxford University Press, 2010.
- [44] S. Okazaki. A study of simulation model for pedestrian movement in architectural space. Part 1: Pedestrian movement by the application of magnetic model. *Transactions of Architectural Institute of Japan*, 283:111–119, 1979.
- [45] Quang-Cuong Pham, Halim Hicheur, Gustavo Arechavaleta, Jean-Paul Laumond, and Alain Berthoz. The formation of trajectories during goal-oriented locomotion in humans. ii. a maximum smoothness model. *European Journal of Neuroscience*, 26(8):2391–2403, 2007.
- [46] Photchara Ratsamee, Yasushi Mae, Kenichi Ohara, Tomohito Takubo, and Tatsuo Arai. Human-Robot Collision Avoidance using a Modified Social Force Model with Body Pose and Face Orientation. *International Journal of Humanoid Robotics*, 10(01):1350008–1–24, 2013.
- [47] Andreas Schadschneider. Cellular automaton approach to pedestrian dynamics-theory. In *Pedestrian and Evacuation Dynamics*, 2002.

- [48] Michael Schultz, Lars Rössger, Hartmut Fricke, and Bernhard Schlag. Group Dynamic Behavior and Psychometric Profiles as Substantial Driver for Pedestrian Dynamics. In Ulrich Weidmann, Uwe Kirsch, and Michael Schreckenberg, editors, *Pedestrian and Evacuation Dynamics 2012*, pages 1097–1111. Springer International Publishing, 2014.
- [49] Paul Scovanner and Marshall F Tappen. Learning pedestrian dynamics from the real world. In *ICCV*, volume 9, pages 381–388, 2009.
- [50] Michael J. Seitz, Felix Dietrich, and Gerta Köster. The effect of stepping on pedestrian trajectories. *Physica A: Statistical Mechanics and its Applications*, 421(0):594 – 604, 2015.
- [51] Armin Seyfried, Bernhard Steffen, and Thomas Lippert. Basics of modelling the pedestrian flow. *Physica A: Statistical Mechanics and its Applications*, 368(1):232 – 238, 2006.
- [52] PradyumnKumar Shukla. On Modeling and Evolutionary Optimization of Nonlinearly Coupled Pedestrian Interactions. In *Applications of Evolutionary Computation*, volume 6024 of *Lecture Notes in Computer Science*, pages 21–30. Springer Berlin Heidelberg, 2010.
- [53] Bernhard Steffen. A Modification of the Social Force Model by Foresight. In Wolfram W. F. Klingsch, Christian Rogsch, Andreas Schadschneider, and Michael Schreckenberg, editors, *Pedestrian and Evacuation Dynamics 2008*, pages 677–682. Springer Berlin Heidelberg, 2010.
- [54] Emanuel Todorov and Michael I. Jordan. Smoothness Maximization Along a Predefined Path Accurately Predicts the Speed Profiles of Complex Arm Movements. *Journal of Neurophysiology*, 80(2):696–714, 1998.
- [55] Song Xu and H.B.-L. Duh. A Simulation of Bonding Effects and Their Impacts on Pedestrian Dynamics. *IEEE Transactions on Intelligent Transportation Systems*, 11(1):153–161, March 2010.
- [56] W. J. Yu, R. Chen, L. Y. Dong, and S. Q. Dai. Centrifugal force model for pedestrian dynamics. *Phys. Rev. E*, 72:026112, Aug 2005.
- [57] Zeynep Yücel, Francesco Zanlungo, Tetsushi Ikeda, Takahiro Miyashita, and Norihiro Hagita. Deciphering the Crowd: Modeling and Identification of Pedestrian Group Motion. *Sensors*, 13(1):875–897, 2013.
- [58] F. Zanlungo, T. Ikeda, and T. Kanda. Social Force Model with explicit collision prediction. *EPL (Europhysics Letters)*, 93(6):68005, 2011.
- [59] Francesco Zanlungo, Dra žen Brščić, and Takayuki Kanda. Spatial-size scaling of pedestrian groups under growing density conditions. *Phys. Rev. E*, 91:062810, Jun 2015.
- [60] Francesco Zanlungo, Tetsushi Ikeda, and Takayuki Kanda. Potential for the dynamics of pedestrians in a socially interacting group. *Phys. Rev. E*, 89:012811, Jan 2014.
- [61] Jing Zhao, Yi Xu, Xiaokang Yang, and Qing Yan. Crowd instability analysis using velocity-field based Social Force Model. In *IEEE Visual Communications and Image Processing (VCIP)*, pages 1–4, Nov 2011.
- [62] Xiaoping Zheng, Tingkuan Zhong, and Mengting Liu. Modeling crowd evacuation of a building based on seven methodological approaches. *Building and Environment*, 44(3):437–445, 2009.
- [63] Qiuming Zhu. Hidden Markov model for dynamic obstacle avoidance of mobile robot navigation. *IEEE Transactions on Robotics and Automation*, 7(3):390–397, Jun 1991.