Universiteit Utrecht
Department of Information and Computing Sciences

Master Thesis
Game and Media Technology

Social Group Behavior and Path Planning

ICA-3919137

Author:
A.E. Kremyzas

Supervisors:
dr. R.J. Geraerts
N.S. Jaklin, MSc

Second Examiner:
prof. dr. M.J. van Kreveld

February 5, 2015

Abstract

This work discusses the navigation of social pedestrian groups in crowded environments. First, we highlight the necessity for including social group behavior in crowd simulations. We also compare existing group methods against each other. Next, we propose *Social Groups and Navigation* (SGN) method for steering social groups of agents in planar homogeneous environments with polygonal obstacles. Our method provides group-specific details on global and local aspects of pedestrian navigation. SGN is inspired by the social-force model of Moussaïd et al. 2010. In addition, it borrows the vision-based collision-avoidance algorithm of Moussaïd et al. 2011. Both original works have been adjusted to improve the ability of the simulated groups to remain social and coherent, while avoiding obstacles and other groups. SGN is flexible and can be coupled with different navigation meshes and global route-planning algorithms. The method is evaluated through extensive experiments. Results demonstrate the ability of SGN to produce coherent and social-friendly group configurations throughout the simulation. Based on our evaluation metrics, the quality improvement over the works of Moussaïd et al. is significant. When groups of three or four are considered, SGN produces social friendly configurations at a significantly higher rate than the method of of Moussaïd et al. In all tested scenarios, the difference of this rate ranges from 15% to 31% for groups of three. Regarding groups of four, the difference is even greater, ranging from 13% to 53%. When groups of two are considered, the difference is small, ranging from 1% to 4%, but is still statistically significant. Performance results suggest that SGN is capable of simulating in real-time thousands of agents that are organized in small social groups.

This thesis project was conducted as part of a collaboration between Utrecht University and INCONTROL Simulation Solutions. Our proposed method has been integrated both in Pedestrian Dynamics, a crowd flow simulator developed by INCONTROL, and in the crowd simulation framework developed by Utrecht University. This research has been supported by the COMMIT/ project [1].

Contents

1 Introduction

Crowd simulations and computer games can be seen as applications in which virtual characters populate their environments. Motion planning methods are employed in these applications to generate paths that lead all simulated characters to their destinations. Realism is a strong requirement for both application domains and can be roughly estimated as the ability of the application to generate plausible motion behaviors. In practice, this means that the generated motion behavior of each character should exhibit a trade-off between following a short and smooth path and following a path that avoids collisions with other characters and dynamic obstacles.

However, path optimality and smoothness along with collision-free motions are not the only requirements for a simulation model to be considered realistic. When several people navigate in real-life, various phenomena have been observed, such as lane formation, stop-and-go waves, vortices and pushing behavior [27,70,71]. As an effort to enhance naturalness of simulations, these emergent crowd behaviors should also be predicted by the underlying motion planning techniques. More importantly, the simulation of pedestrians should also take into account the fact that people often navigate in small social groups. In these pedestrian groups, the individual paths should be planned in a manner that facilitates social interactions between the group members throughout their navigation. This aspect of pedestrian behavior has been neglected until recently [66] by the proposed simulation models and remains an active research topic (see Section 2.1).

1.1 Project motivation

The research findings in the fields of sociology and psychology make the requirement to model pedestrian groups justifiable in order to generate realistic crowd simulations. According to James [42] and Coleman and James [17], approximately 70% of observed pedestrians walk in groups. This percentage is expected to be even higher in social-friendly environments, for instance, theme parks, such as Disney World [2], shopping malls, museums or fun fairs. Roloff [78] observes the difference in the pedestrian behavior of individuals depending on whether they belong in a small group or not. The application domain of crowd simulations often includes simulating mass event scenarios for safety engineering purposes. Aveni [7] mentions that it is expected for small groups to prevail over individuals at mass events. Furthermore, one could assume that during emergency scenarios, such as building evacuations, panic dominates the psychology of the crowd thus leading to individualistic behavior. However, according to Cocking et al. [16] this is not the case. Aguirre et al. [5] note that, under panic, the social interactions between individuals are further enhanced, thus leading to collective behavior. Mawson [62] has observed that the collective behavior of social groups under stress often increases evacuation times. On a similar note, Drury et al. [22] indicate that it is also expected from strangers to establish social bonds and form groups during emergency situations. Therefore, it is evident that a motion planning method should model social group behavior to generate convincing simulation results. Existing methods for simulating social groups focus on the local behavior of the individuals that form these groups [48, 66, 92]. To the best of our knowledge, no social-group motion-planning method has been proposed that provides group-specific details on both global and local aspects of pedestrian navigation.

1.2 Project goals

The aim of this project is to propose a method for simulating social pedestrian groups throughout their navigation in a planar environment. A predefined set of pedestrians composes each social group. Each member of a social group can be initialized anywhere in the walkable environment. The members of each group must first gather at a meeting area. Next, all group members must embark together and navigate towards their common destination. Each member of a group must traverse a path that is a trade-off between the shortest path to the goal of the group and a path that maintains social interactions with the rest of the group members, while avoiding collisions with obstacles, other groups and other

pedestrians in the crowd. When a group loses its coherence, its members must attempt to regain it. Group coherence is determined based on the separating distance and visibility between the members. Group sociality is determined based on the deviation from a formation that benefits social interactions. Visual stimulus is employed for concluding on whether social interactions between group members exist. The proposed method must provide details on motion-planning for social groups with respect to both global and local aspects of pedestrian navigation. At the same time, we aim at proposing a method that meets the performance requirements of a largely populated environment, where social interactions are expected [7]. Finally, the proposed method must allow for heterogeneity in terms of the body size, preferred speed and visual ability of the simulated pedestrians.

1.3 Crowd simulation

For a comprehensive study of existing motion-planning algorithms and crowd simulation techniques, the reader is referred to the books of LaValle [55] and Thalmann and Musse [86], respectively.

The method proposed in this thesis has been implemented within the Explicit Corridor Map (ECM) Crowd Simulation Framework [4]. This framework consists of five different planning levels; see Figure 1.1. On the top level (Level 5), the framework performs *high-level planning*. This involves handling events, action planning and any other process that is not treated in the lower levels. This is the level in which the start and goal positions of the simulated agents are set. Next (Level 4),there is the *global route-planning* stage, during which the framework is responsible for computing a path for each agent that connects its start with its goal position. During global route-planning, static obstacles define the non-walkable areas of the virtual environment. A path is generated for each agent, based on the walkable areas, the start and the goal of the agent. This path serves as a guide for the agent to its destination. The rest of the simulation levels are repeated sequentially at each simulation step and determine the motion behavior of each agent between two simulation steps. The *route-following* stage determines a preferred velocity for each agent based on its current position and the global route that was computed at Level 4. Next, during the *local-movement* stage (Level 2) the preferred velocity and the local information of each agent are combined to determine its actual velocity. In the simplest case, the local information for each agent involve its neighboring agents and any dynamic obstacles. In this case, the aim of the local-movement stage is to compute for each agent a velocity that deviates as little as possible from the preferred velocity while avoiding collisions with the static and dynamic environment. After a velocity has been computed for an agent, the lowest level (Level 1) handles the animation. Depending on the implemented methods and user input, the framework can return to Level 5 or 4 before repeating the simulation-step cycle. A similar design is expected in alternative crowd simulation frameworks. Hoogendoorn et al. [38], for instance, propose a three level hierarchy to pedestrian navigation: 1) *strategical*, which involves general planning, 2) *tactical*, which involves route selection, and 3) *operational*, which involves basic walking behavior.

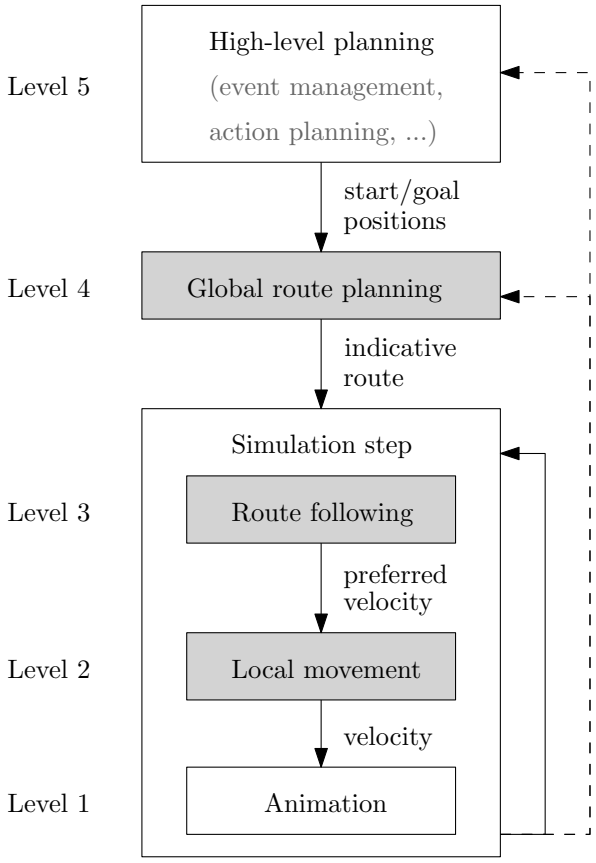

Figure 1.1: ECM Crowd simulation framework [4]

1.4 Report structure

The rest of this work is structured as follows. In Section 2, we present and compare work that is related to our research topic. In Section 3, we discuss in detail the works that are employed by our method. The proposed method for simulating social groups is presented in Section 4. We conduct experiments to evaluate our method and present their results in Section 5. Finally, conclusions are drawn and ideas for future work are given in Section 6.

Our method is designed to allow for crowd heterogeneity in terms of the body size, preferred speed and visual ability of the simulated pedestrians. The method addresses social group behavior on both global and local aspects of pedestrian navigation. At the same time, our evaluation suggests that the proposed method produces convincing crowd flows where social group behavior is exhibited. According to performance results, the proposed method is capable of simulating in real-time thousands of agents that are organized in small social groups.

2 Related work

2.1 Social group behavior

In this section, we discuss existing research on modelling group behavior. We distinguish approaches based on whether they set a constraint on the group size (*large groups*), or whether they consider groups consisting of few members only (*small groups*). We also discuss group behavior approaches based on cellular automata and briefly mention relevant research conducted in the fields of robotics and computer animation. Finally, we draw conclusions in Section 2.1.2.

Large groups. Modelling the motion behavior of agents that navigate in groups has been of research interest in the fields of path planning, animation and crowd simulation. Reynolds [75] presents a method for simulating the group behavior of a flock of birds, a school of fish or a herd of land animals. The motivation is that group behavior should emerge through the aggregation of the individual motion behaviors of the group members. Therefore, a distributed behavioral approach is followed, where each group member navigates based on simple rules such as *separation*, *alignment*, and *cohesion*. Later, Reynolds enhances his model by adding several steering behaviors for autonomous agents [76] and elaborates on how to simulate interesting scenarios where grouped moving entities interact with the environment [77]. A main limitation of the boid models [75–77] is that they describe the motion behavior of large groups of animals. As a result, they neglect the specificity of a small group of humans, where motion planning should maintain social interactions between the group members.

Similar to Reynolds' work [75–77], Musse and Thalmann [67] introduce a rule-based method for modelling the crowd behavior of multiple pedestrian groups. Again, the simulated agents exhibit flocking behavior, but, unlike Reynolds' work, they are also able to switch between groups based on sociological factors. However, no details are provided on how group coherence is addressed on a global motion-planning level. Figures indicate that herding behavior is exhibited among group members, even for small groups and at low crowd densities, which is something that contradicts the empirical observations of Moussaïd et al. [66]. Also, the collision-avoidance scheme of Musse and Thalmann [67] is naive and is expected to generate unnatural motions and emergent crowd behaviors that lack plausibility, since in every potential collision pair only one individual is assigned the responsibility of avoiding the other; this choice is unintuitive.

Kamphuis and Overmars [45] propose a motion planning technique for groups that specifically addresses the requirement to provide a solution path if one exists, while maintaining group coherence and naturalness. This work improves upon the boid models [75–77] in that it guarantees that the members of the group will not be separated. Although the authors deal with the motion behavior of human groups, they focus on large groups where tactical manoeuvres are allowed to maintain group coherence in terms of space occupancy. Therefore, such an approach is more suitable to military applications and does not consider the social and group-size aspects of a small pedestrian group.

On a similar note, Kimmel et al. [49] introduce a velocity-based method for maintaining group coherence. Their work serves as an extension to existing *velocity obstacle* (VO) approaches [9, 10, 24, 84] (see Section 2.2) by enhancing these local collision-avoidance methods so as to consider team behavior. A main limitation is that the agents move at a constant speed throughout the simulation. Also, this method attempts to maintain a single formation for every group based on a proximity graph. This, however, is not the case for real-life pedestrian groups where it has been observed that a group switches between formations depending on its social and spatial context [66].

Unlike the methods mentioned above, efforts have been made to address the group motion planning problem using insight from sociology and psychology research. Park et al. [69] propose a model for social groups that is inspired by *common ground theory* [14, 15]. Their work is not limited to modelling reactive intra-group behavior but also considers the higher level social interactions between group members. A leader is assigned to each group and coordination strategies that follow the common ground theory can be defined at both a

micro and macro simulation level. As a measure of group coherence the authors compute the distance of a follower from the leader projected on the direction of motion of the leader.

As an alternative approach, Qiu and Hu 2010 [74] introduce a method for modelling dynamic groups that is based on utility theory and social comparison theory. Their proposed approach consists of two steps and allows for individuals to dynamically leave a group of pedestrians and join another according to their spatial and social context. During the first step, utility theory is applied for an individual to decide which group to join (*group formation*). During the second step, social comparison theory is applied for an individual to decide which member of the chosen group to follow (*individual selection*). Each agent can perceive its neighbours that lie within an elliptical field of view and a *sociality* parameter determines how likely the agent is to switch groups.

Most of the methods above adopt a leader-follower approach to model group motion behavior [67, 69, 74–77]. However, such an approach does not suffice to generate paths that facilitate the social interactions between the members of small groups. Even when group coherence is guaranteed [45, 49], no social group behavior is exhibited. Finally, although the methods of [69] and Qiu and Hu 2010 [74] deal with the social aspects of group behavior, they fail to provide details for social groups in all involved motion-planning stages. Therefore, our method will attempt to deal with the social aspects of small groups. Whenever the spatial context does not allow for sociality to be exhibited among group members, we will borrow techniques from the methods above to maintain group coherence.

Small social groups. More recently, research has focused on the motion behavior of social groups that consist of few pedestrians. Moussaïd et al. [66] conduct an experimental analysis on the group behavior that is observed within pedestrian crowds. According to their observations, a small group of pedestrians typically consists of up to four members that walk in a horizontal line formation at low crowd densities (*line-abreast* formation). At moderate crowd densities, the formation is adjusted to a V-shape that allows for social interaction between the group members (*V-like* formation). At high crowd densities, the group members form a lane and a leader-follower approach is adopted so that there is enough clearance for the group to proceed towards its destination (*river-like* formation). Based on these observations a social-force model is proposed to model the behavior of small groups. However, their approach ultimately sums the force factors that simulate the social and avoidance interactions of each individual to determine its new velocity vector. As a result, it is expected that under special conditions the sum of the forces can yield a velocity vector that is either not realistic or even leads the simulated characters to get stuck in local minima.

Karamouzas and Overmars [47] present a velocity-based approach for steering agents that are organized in small groups of up to three members. Their method serves as a trade-off between maintaining group coherence and generating optimal individual paths. The proposed method relies on empirical data that were discussed in the work of Moussaïd et al. [66] and explicitly models group formation. The motion of the group members is determined through optimization based on a cost function that leverages collision avoidance and social group coherence.

By combining the work of Karamouzas and Overmars [47] with the vision-based avoidance scheme of Ondřej et al. [68], Wu et al. [92] introduce an alternative method for simulating the behavior of small groups. Such an approach is expected to generate stop-and-go waves in crowd simulations, which is an emergent crowd behavior that the method of Moussaïd et al. fails to generate due to its lack of prediction. On the other hand, the performance of Moussaïd et al. is expected to be significantly better than of Karamouzas et al. and Wu et al., where the velocity of the group and its formation must be jointly optimized. Moreover, unlike Karamouzas et al. or Wu et al., the collective behaviors that the method of Moussaïd et al. generates are emergent, since group formation is not explicitly modelled. Calibrating the parameters of any of the three discussed approaches so as to achieve a robust group behavior under different simulation environments and crowd densities might be laborious.

The empirical observations of Moussaïd et al. [66] on small social groups have yielded important insight on how sociality is exhibited among group members with respect to their

trajectories. Moussaïd et al. [66], Karamouzas and Overmars [47] and Wu et al. [92] propose models to meet these observations. However, none of these methods provides necessary details on the high-level planning, global route-planning and route-following levels of the simulation. Also, no performance evaluation is conducted on the works of Moussaïd et al. [66] and Wu et al. [92]. Performance results of Karamouzas and Overmars [47] indicate that their method is unlikely to be used for real-time simulations of thousands of agents that navigate in small groups. However, being able to efficiently simulate small groups in largely populated areas is expected to be of interest, as the presence of small groups at mass events is dominant [7]. Therefore, our aim is to propose a method that allows for simulating in real-time thousands of agents that are organised in small social groups.

Cellular automata. Another approach that has been used in crowd simulations is based on cellular automata [12, 50, 79]. In this approach, the scene is discretized into cells, where each cell can be occupied by a single agent. A finite number of states are defined for each cell. A set of cells relative to the current cell models its neighbourhood. The transition between cells is based on method-specific rules. More recently, cellular-automata research has also focused on modelling group behavior [80, 82].

Sarmady et al. [80] extend their previous work on modelling crowd behavior using cellular automata [79] so as to consider the navigation of groups. Their model is based on the *principle of least effort* [94] and is adjusted to compensate for leader-follower group behavior. For each follower in a group, a *dependence factor* is introduced to describe the motivation of the follower to track its leader in order to reach its destination. The distance of a neighbouring cell from the leader is used as a metric to compute the preference of a follower to move to that cell in order to maintain group coherence.

Seitz et al. [82] also propose a leader-follower approach that utilizes cellular automata for modelling pedestrian group behavior. The novelty in this approach is that the leader is not a fixed group member. Instead, the leader is defined at each time as the member that has minimum distance from the destination. A speed-adjustment scheme is employed to maintain group coherence. The authors also propose a technique for simulating the motion of larger groups. Their approach treats a large group as a union of small groups where inter-leader attraction potentials keep small groups together.

The cellular-automata approaches are rather simple to implement and offer fast execution times. However, these approaches assume that each cell can be occupied by a single agent at a time. Furthermore, since cellular automata employ a discretization of the scene into a grid, the allowed directions of motion are also discretized. This has a negative effect on the plausibility of the group motions, because the generated paths may not be cost-optimal or smooth. Therefore, we avoid basing our method on cellular automata. However, we will borrow and adjust the concept of the dynamic leader from Seitz et al. [82] to facilitate our high-level planning mechanisms for maintaining group coherence.

Robotics. Modelling group behavior has been a well-studied research topic of the robotics community. Relevant approaches address the problem of coordination between robots without a requirement for communication protocols [6, 39, 60, 61]. Leader-follower models have also been proposed [11, 18, 40, 57]. However, these approaches are expected to be computationally expensive when modelling multiple groups in large crowd simulations, while they do not emphasize on the social context of pedestrian motion in groups. More recently, Morales et al. [63] introduced a utility-based model for a robot that walks side-by-side to a human companion. Their method cannot generalize to the case of small groups consisting of more than two members, while it specifically addresses the requirements of human-robot interaction.

Computer animation. Pedestrian groups have also been studied within the field of computer animation. Peters and Ennis [72] employ navigation fields to steer small pedestrian groups. Takahashi et al. [85] generate group motion for tactical simulations by applying graph Laplacian analysis on a sequence of group-formation instances. Ennis and O'Sullivan [23] conduct experiments on the plausibility of virtual conversing groups based

on the distance of the group members and their relative orientations. Kwon et al. [52] propose a method for editing group motions under formation constraints, allowing to synthesize large group motion clips by stitching smaller clips together.

Recently, example-based approaches [43,56,58] have addressed the problem of simulating group behavior by exploiting video recordings of real-life pedestrian groups. The trajectories of group motions are either extracted manually [58] or semi-automatically [56] from captured data. Lai et al. [53] utilize motion graphs to synthesize the motion of virtual groups. By following example-based approaches, however, the naturalness of the generated motions relies on the size of the available group motion database. Also, due to performance reasons, these methods are often used for offline simulations, while the path planning aspect of the group behavior is neglected.

Computer animation research on social groups mainly focuses on the lowest level of crowd simulation which is outside the scope of this research. However, the findings of Ennis and O'Sullivan [23] can be employed during local-movement phase to decide on whether a group should be avoided as a whole by other groups or agents.

2.1.1 Comparison of group methods

In Tables 2.1 and 2.2, we compare the most relevant of the group methods that were described in this section. Each cell indicates whether the method that corresponds to the cell's column exhibits the property indicated by the cell's row.

In the first three rows, we examine whether the methods under comparison provide group-specific details on the global route-planning, route-following and local-movement levels of the simulation (see Figure 1.1). No method so far deals with groups on these three simulation levels.

The forth row indicates whether a method is designed for a pedestrian group consisting of a few only members (≤ 5). Park et al. [69], for instance, do not explicitly set a constraint on the group size, but their method is mainly focused on small social groups.

In the fifth row, we illustrate whether a method models the social interactions between the members of a pedestrian group. Flocking methods generally do not exhibit this property [67, 75–77, 80], unless additional social rules are incorporated in such methods [69, 74, 82]. Also, the velocity-obstacle approach of Kimmel et al. [49] is not considered to generate social group behavior, because the formation of the group is fixed and does not adjust according to the spatial context. Although the method of Kamphuis and Overmars [45] guarantees group coherence, it does not handle social interactions between the group members.

The sixth row indicates which methods employ a leader for each group. For clarity reasons, Karamouzas and Overmars 2010a [47] employ a leader in their group-model, but do not adopt a leader-follower approach. Instead, the leader is dynamic and is used as a reference for defining a local coordinate system and simplifying math computations. Similarly, Wu et al. [92] borrow this leader-technique from Karamouzas and Overmars 2010a. Seitz et al. [82] present the only method that utilizes a dynamic leader in a leader-follower approach.

The seventh row indicates whether a method allows for agents to form groups dynamically, leave or switch between groups. Only the works of Musse and Thalmann et al. [67] and Qiu and Hu [74] allow for this behavior using high-level decision making techniques.

The eighth row indicates whether a method explicitly uses a formation model to describe the spatial dependencies between the members of a group. The methods of Kimmel et al. [49] and Qiu and Hu [74] do not consider switching between between formation models. Instead, they employ a single formation throughout the navigation of a group.

The ninth and tenth row distinguish methods on whether they utilize parameters to model behavior that reflects the sociality and mental state of the agents, respectively. The boid models [75–77] use the average position and velocity of the group members as social variables, while an example scenario is presented in [77], where the annoyance of simulated birds is used as a psychological variable. Musse and Thalmann et al. [67] define the level of group dominance and the level of group relationship as social variables in their model. They also use the emotional status and the goal interest of an agent to describe its current

mental state. The social-force model of Moussaïd et al. [66] involves among others a group visual factor and a group attraction factor. The weights of these factors in the final formula consist of social parameters that express the degree in which an agent tries to maintain visual contact and spatial proximity with its fellow members, respectively. Karamouzas and Overmars [47] optimize a cost function to determine the new velocity of a group at each simulation step. The cost function involves a term that measures the deformation of the group under a candidate velocity compared to a social-friendly formation. The weight of this term consists of a social parameter describing the degree in which a group tries to remain in a social formation. Wu et al. [92] borrow the optimization technique from Karamouzas and Overmars [47], thus employ the same social variable. Sarmady et al. [80] employ a dependence factor as a social variable to define the degree to which a group member tries to follow its leader instead of following an individual path to the goal of the group. Qiu and Hiu [74] use a sociality parameter to describe how likely it is for an agent to switch between groups.

The eleventh row indicates whether the proposed methods explicitly allow agents to have individual parameters that differ from their fellow members. Musse and Thalmann [67] enable group members to have an individual emotional status. Kimmel et al. [49], Karamouzas and Overmars [47] and Wu et al. [92] allow group members to have a different radius. Moussaïd et al. [66] can simulate groups where the members have different masses and preferred speeds. The method of Sarmady et al. [80] enables each member to have its own dependence factor.

As can be seen in the twelfth row, the work of Sarmady et al. [80] is the only group method that uses density information to guide agents. Finally, as the final row indicates, no method so far takes terrain preferences into consideration for planning the motion of the groups.

For comparison reasons, Table 2.2 also displays the properties of the method we propose in this work. Our method is the only method under study that jointly addresses group-relevant aspects of pedestrian navigation on global route-planning, route-following and local movement simulation levels. Similar to Seitz et al. [82], our method employs a dynamic leader for each group, although no leader-follower approach is adopted in our case. Also, we do not use a formation model. This choice allows for social group behavior to emerge. Social variables are incorporated in our method and define the strength of coherence and social interactions between the group members. Our method is designed to allow for the simulation of heterogeneous groups, with respect to the group sizes as well as the body size, preferred speed and visual ability of the group members. Dynamic group formation, psychological variables, density based planning and terrain preferences are not modelled in our work.

2.1.2 Conclusions

The work of Moussaïd et al. [66] manages to provide a rather simple method for simulating small groups and is stated to adequately meet empirical observations. Furthermore, the method is formulated as a social-force model, thus is expected to allow for real-time simulation of thousands of characters that are organised in small groups. However, the underlying collision-avoidance mechanisms are reactive, therefore lack anticipation. What is more, this work does not provide the details on how global route planning and route following should be addressed from a social-group standpoint. The work of Kamphuis and Overmars [45] discusses these aspects, but it does not handle social interactions between the group members. Furthermore, it is possible for a social group to lose its coherence in densely populated environments or in the presence of bottlenecks. In such cases, a social group should attempt to regain its coherence. However, none of the discussed methods includes basic high-level mechanisms that are necessary for a small social group to successfully traverse its path to its destination while attempting to be coherent. In this thesis, we are going to tackle the aforementioned problems.

	Reynolds 1987, 1999, 2000	Musse and Thalmann 1997	Kamphuis and Overmars 2004	Kimmel et al. 2012	Moussaid et al. 2010	Karamouzas and Overmars 2010a
Group-relevant global route-planning	-	-	+	-	-	-
Group-relevant route following	+	+	+	-	-	-
Group-relevant local movement	+	-	-	+	+	+
Aim at small group size (≤ 5)	-	-	-	-	+	+
Social group behavior	-	-	-	-	+	+
Leader	+	+	-	-	+	+
Dynamic groups	-	+	-	-	-	-
Formation model	-	-	-	+	-	+
Social variables	+	+	-	-	+	+
Psychological variables	+	+	-	-	-	-
Heterogeneous groups	-	+	-	+	+	+
Density based	-	-	-	-	-	-
Terrain preferences	-	-	-	-	-	-

Table 2.1: Comparison between group planning methods.

	Wu et al. 2013	Sarmady et al. 2009	Seitz et al. 2014	Park et al. 2012	Qiu and Hu 2010	Our method
Group-relevant global route-planning	-	-	-	-	-	+
Group-relevant route following	-	+	+	+	+	+
Group-relevant local movement	+	-	+	-	-	+
Aim at small group size (≤ 5)	+	-	+	+	-	+
Social group behavior	+	-	+	+	+	+
Leader	+	+	+	+	+	+
Dynamic groups	-	-	-	-	+	-
Formation model	+	-	-	-	+	-
Social variables	+	+	-	-	+	+
Psychological variables	-	-	-	-	-	-
Heterogeneous groups	+	+	-	-	+	+
Density based	-	+	-	-	-	-
Terrain preferences	-	-	-	-	-	-

Table 2.2: Comparison between group planning methods.

2.2 Local collision avoidance

To achieve realism in crowd simulation, the local behavior of individuals must be modelled. Furthermore, simulating group behavior in pedestrian crowds should effectively combine collision avoidance with group coherence techniques. Local collision avoidance is an important aspect of micro simulation, since individuals are expected to adjust their planned route so as to avoid collisions with dynamic obstacles and other individuals. In the rest of this section, we discuss various approaches for local collision avoidance between simulated characters.

The social-force model. Early research has proposed social-force methods to model avoidance behavior [34,37,76]. Helbing and Molnár [37] introduce an approach for describing pedestrian motion based on social forces. Social forces express the systematic changes over time in the desired velocity of a given individual within a crowd context. The forces are composed of several terms that attempt to capture the spectrum of motion behavior. The method was later extended by its authors to better capture pedestrian behavior in evacuation scenarios [35] and under panic [36].

A main advantage of the social-force model approaches [35–37] is that their implementation is rather intuitive and is expected to perform faster than alternative approaches. Moreover, emergent crowd behavior is observed, since the method implicitly generates lane formation in areas where agents move in opposite directions. On the other hand, future research has reported that the generated motions can be unconvincing when a small number of pedestrians are simulated and the method parameters are not well specified [54]. It is also argued that unrealistic parameters should be used in such cases in order to maintain naturalness. Therefore, before applying this method, experimentation is expected on the manual tuning of such parameters in relation to the virtual scene context. Furthermore, in high densities the sum of the force factors often leads to a direction of motion that rapidly changes over time. As a result, oscillatory motion is exhibited thus reducing the realism of the simulation. Lakoba et al. [54] have also observed, despite the corresponding force factors, that collisions between agents do occur. Although, collisions between pedestrians are expected to happen under certain conditions in real-life, the problem lies in that the method does not resolve those collisions properly, thus allowing agents to pass through each other.

Velocity-based approaches. Velocity-based approaches [9, 10, 24, 48, 84] were introduced as an effort to incorporate anticipation in the generated local behavior, which is an aspect that the social-force approaches failed to consider. Fiorini and Shiller [24] introduce the concept of the *velocity obstacle* (VO) to define the set of velocities that would lead a robot to collide in the future with its static and dynamic environment. This approach considers a single reactive agent (the robot) and the rest of the dynamic obstacles are assumed to be passive. Van den Berg et al. 2008 [10] extend the VO approach so as to achieve collision avoidance between multiple robots that navigate in a planar environment. For this reason they base their method on the concept of reciprocity, i.e. each agent assumes that the rest of the agents follow the same avoidance reasoning. By examining all possible pairs that an agent can form and then by intersecting the allowed new velocity spaces, each agent selects its new velocity as the allowed velocity that is closest to its desired velocity. Their method is further enhanced in van den Berg et al. 2011 [9] to guarantee that the generated motions are oscillation-free while guaranteeing collision avoidance. This work elaborates on the optimization scheme to define the new velocities for all agents. The problem is solved using linear programming and relies on the assumption that for each agent-pair, the avoidance responsibility is evenly shared between the involved agents. A main limitation of the VO approaches is that they do not allow for variance between the individual motion strategies. Reciprocity generates rather homogeneous avoidance behaviors, thus results can be unrealistic. In real-life scenarios, some of the pedestrians are expected to apply a waiting rule similar to Pelechano et al. [71] to give space to the rest so that they can easily progress after some time. Also, VO methods do not allow for pushing behavior since collision avoidance is stated to be guaranteed.

As an attempt to increase the realism of the generated avoidance behavior, Karamouzas and Overmars 2010b [48] present a different velocity-based approach that is based on the empirical observations of Pettré et al. [73]. The authors perform an experimental analysis on the motion capture data of [73], and propose a velocity-based model that can sufficiently reproduce the observed motions. Previous research on human behavior is also taken into account and helps identifying key factors affecting motion behavior. These include the *personal space* [33] each individual tries to maintain from the rest, as well as the energy efficiency that dictates the re-planning decisions of the individuals based on the *principle of least effort* [94]. Only a few most urgent collisions are considered for realism, starting with the earliest one. Similar to the distance thresholds defined in [34], this method employs remaining time-to-collision thresholds to define the permitted velocities and orientations throughout the collision avoidance stages using piecewise functions that match the observations. These permitted velocities and orientations are then optimized based on a criterion that jointly minimizes energy consumption, collision risk, linear and angular acceleration. A limitation of this method is that it employs axis aligned boxes to model the static obstacles, can thus lead to a crude approximation of the non-navigable space.

Vision-based approaches. More recently, vision-based methods have been presented as an attempt to better simulate the perception model of pedestrians and their corresponding motion planning mechanisms [65, 68]. Ondřej et al. [68] propose a collision-avoidance technique that is based on a synthetic-vision perceptual model. The proposed model is inspired by the experimental observations of Cutting et al. [20]. According to Ondřej et al. , collision avoidance is based on the bearing angle between a moving pedestrian and an obstacle, and the remaining time-to-interaction. When the remaining time-to-interaction is large, avoidance is achieved by maintaining the speed and altering the direction of motion. On the other hand, when the remaining time-to-interaction is small, then the agent decelerates until it stops to avoid an imminent collision. The method computes the desired angular velocity by attempting to both avoid the obstacles and deviate as little as possible from the goal direction. Only when the time-to-interaction of the most imminent collision is below a threshold value, then the speed of the agent decreases exponentially.

As an alternative, Moussaïd et al. 2011 [65] introduce a collision-avoidance method that is based on cognitive science approaches. They employ two vision-based behavioral heuristics to simulate the perception of pedestrians and their corresponding decision making processes throughout navigation. At each simulation step, an agent computes the time that corresponds to the most imminent collision for each candidate direction. Then, based on the first heuristic, the agent determines its desired direction of motion by trying to minimize the distance from its destination when a new directional change is expected due to the presence of an obstacle. The new speed of the agent is chosen based on the second heuristic, which dictates that the agent should maintain a distance from the first obstacle on the chosen direction that corresponds to a time-to-collision of greater than or equal to the relaxation time. Should unpredicted collisions occur with other agents or obstacles, physical forces are applied to simulate the interactions between the colliding entities.

Due to the complexity of the perception model, both vision-based approaches are expected to perform slower than social-force models or velocity-based approaches. Also, defining the field of view of an agent using a viewing angle and a fixed maximum viewing distance can be insufficient. In real-life scenarios, the visual stimulus of a pedestrian is heavily dependent on occlusions. However, the view-dependant visual stimulus that vision-based models offer better approximates the perception of humans compared to social-force or velocity-based models. Therefore, we choose to employ the method of Moussaïd et al. [65] as a basis of our collision-avoidance technique.

2.2.1 Conclusions

The vision-based approaches of Ondřej et al. [68] and Moussaïd et al. [65] are promising in that they attempt to model the perceptual and decision making mechanisms of humans regarding collision avoidance. Social-force models superimpose all pair-wise interactions of

an agent to determine the overall influence on its motion. This avoidance scheme is arguable and the vision-based approaches of Ondřej et al. and Moussaïd et al. do not adopt it. The heuristics employed by Moussaïd et al. are rather intuitive and describe pedestrian behavior as a tradeoff between following the shortest path and following a path that minimizes directional changes. Heuristics have been reported to express well human behavior by being an integral part of our decision making processes [31, 32]. The challenge lies in identifying the heuristics that pedestrians employ when planning their motion. In the frame of this work, an additional challenge would be to identify heuristics that describe the social interactions of a small group and effectively combine them with the ones defined by Moussaïd et al. to describe the collision-avoidance behavior of each group member. Another point of consideration is the performance of vision-based techniques. Ondřej et al. report real-time performance for up to 200 simulated pedestrians, after partially implementing their method on the GPU. They also propose guidelines on improving running time. On the other hand, Moussaïd et al. do not evaluate their method in terms of performance. For our method, we will use the collision-avoidance scheme of Moussaïd et al. 2011 [65] which is expected to be easier to combine with the social-force model of Moussaïd et al. 2010 [66] for simulating small social groups than the avoidance scheme of Ondřej et al. [68].

3 Preliminaries

In this section, we present the details of the work that is directly related to our research.. In Section 3.1, we present the details of the vision-based method by Moussaïd et al. 2011 for local collision avoidance. In Section 3.2, we discuss the social-force model of Moussaïd et al. 2010 [66] for simulating social groups. For reference and readability reasons, the current section uses the notation from the original papers [65, 66]. In Section 4, we will provide an updated notation that resolves any conflicts between the methods borrowed.

3.1 Vision-based avoidance algorithm

The work of Moussaïd et al. 2011 [65] is one of the first approaches to employ visual information to determine the avoidance behavior of simulated agents [65, 68]. The motivation behind this choice is that vision is the principal stimulus that affects the local-movement decisions of pedestrians [8, 29, 30, 87].

Let x_i and v_i define the position and velocity of an agent i, respectively. Each agent i is modelled by a disc of radius r_i centered at $\mathbf{x}_i$. The radius r_i of agent i is assumed proportional to its mass m_i and is defined as: $r_i = m_i/320$, where m_i is measured in kg and r_i is measured in meters.

The authors describe the visual stimulus of an agent i using a circular sector to model the vision field. The vision field ranges ϕ° to the left and right of the line of sight vector $\vec{H}_i$. Let d_{Max} define the maximum viewing distance of each agent. Assuming a disc of radius d_{Max} centered at $\mathbf{x}_i$, the vision field is defined as the circular sector of that disc which has a central angle of $2\phi^\circ$ and the line of sight vector $\vec{H}_i$ as the angle bisector of the central angle. Figure 3.1 gives an illustration.

Let v_i^0 define the preferred walking speed of agent i and O_i define the last visible point on the direction of the destination point. Let α_{des} define the angle between the line of sight $\vec{H}_i$ and the desired direction of motion. Let v_{des} denote the desired walking speed on the desired direction of motion. Two heuristics are formulated to determine α_{des} and v_{des}. The authors argue that employing heuristics to determine the local behavior of an agent is based on the observation that people often resort to similar cognitive processes when they have to process a lot of information or decide on their actions in a small amount of time [31, 32].

Let $\alpha \in [-\phi, +\phi]$ be a candidate angle of motion and α_0 be the angle that corresponds to the direction of the destination point O_i. Let $f(\alpha)$ define the distance to the first collision in the direction that corresponds to candidate angle α, assuming agent i moves at its preferred speed v_i^0. For the computation of $f(\alpha)$, the current velocities and radii of other agents are used. For each candidate angle α, if no collision occurs within distance d_{Max}, then $f(\alpha)$ is set to d_{Max}.

Desired angle of motion. The first heuristic determines the desired direction of motion by computing α_{des}. Based on empirical observations [8, 87], an agent i chooses a direction of motion that is a trade-off between avoiding obstacles and minimizing the deviation from the most direct path. The first heuristic is formulated as follows:

> *A pedestrian chooses the direction α_{des} that allows the most direct path to destination point O_i, taking into account the presence of obstacles [65].*

Based on this heuristic, α_{des} is computed from the following equation:

$$\alpha_{des} = \underset{\alpha \in [-\phi, +\phi]}{\operatorname{argmin}} \big(d(\alpha)\big), \tag{3.1}$$

where

$$d(\alpha) = \sqrt{d_{Max}^2 + f(\alpha)^2 - 2 d_{Max} f(\alpha) \cos(\alpha_0 - \alpha)}. \tag{3.2}$$

For each candidate angle α, $d(\alpha)$ represents the Euclidean distance from the destination point O_i, after moving along the corresponding candidate direction for a distance of $f(\alpha)$. Therefore, the first heuristic gives a desired angle of motion α_{des} such that the distance

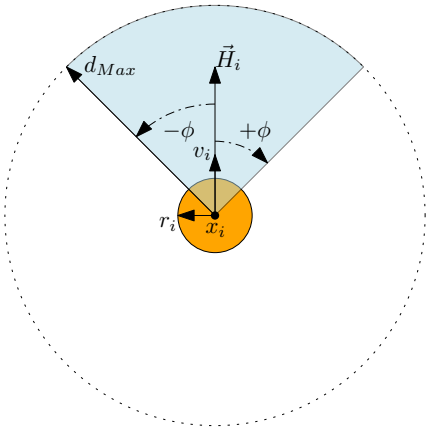

Figure 3.1: The vision field (*light blue*) of an agent i (*orange*).

to the destination is minimum when the next directional change should occur. Figure 3.2 illustrates an example for computing $d(\alpha)$.

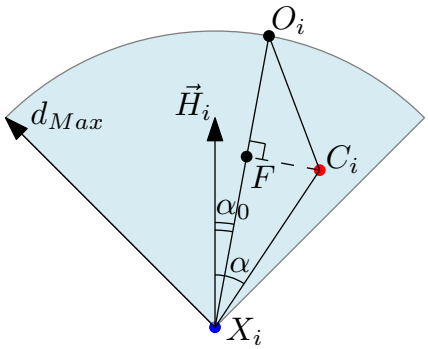

Figure 3.2: Example computation of $d(\alpha)$. Let X_i be the position of point agent i (*blue*) and C_i be the position of a static point agent (*red*). The destination point O_i lies at a bearing angle α_0, while C_i lies at a bearing angle α. Then, $f(\alpha) = |X_iC_i|$ and $d(\alpha) = |C_iO_i|$. Let F be the foot of C_i on X_iO_i. Then, $d(\alpha)$ can be computed as follows:
$|X_iF| = |X_iC_i| \cos(\alpha - \alpha_0) = f(\alpha) \cos(\alpha - \alpha_0)$
$|C_iF| = |X_iC_i| \sin(\alpha - \alpha_0) = f(\alpha) \sin(\alpha - \alpha_0)$
$|O_iF| = |O_iX_i| - |X_iF| = d_{Max} - f(\alpha) \cos(\alpha - \alpha_0)$
$d(\alpha) = |C_iO_i| = \sqrt{|O_iF|^2 + |C_iF|^2} = \sqrt{{d_{Max}}^2 + f(\alpha)^2 - 2d_{Max} f(\alpha) \cos(\alpha - \alpha_0)}$

Desired walking speed. The second heuristic determines the desired walking speed v_{des} and is applied after the desired angle of motion α_{des} has been computed. The relaxation time τ is utilized to determine the desired speed and is estimated at $0.5s$. In practice, a relaxation time $\tau = 0.5s$ indicates that a pedestrian requires $0.5s$ to adjust its current velocity to an arbitrary desired velocity. The second heuristic relates the relaxation time with the desired speed as follows:

> *A pedestrian maintains a distance from the first obstacle in the chosen walking direction that ensures a time to collision of at least τ [65].*

Therefore, an agent chooses its speed so that, in case of an imminent collision, there will always be enough time to avoid that collision. This is achieved by ensuring that at least a time of τ remains for the agent to adjust its current velocity to a desired collision-free velocity. Let d_h be the distance between agent i and the first obstacle on the desired angle of motion α_{des}. Then, according to the second heuristic:

$$v_{des} = \min\left(v_i^0, d_h/\tau\right) \tag{3.3}$$

As a result, at every simulation step, the desired velocity $\vec{v}_{des}$ of an agent i is computed in two steps: 1) the first heuristic is applied to determine the desired angle of motion α_{des} (Equation (3.1)), and 2) the second heuristic is applied to determine the desired walking speed v_{des} (Equation (3.3)).

Physical-contact forces. The authors also introduce physical contact forces in their method to simulate the interactions between an agent and the static environment or between a pair of agents. Those forces are only applied to an agent when there is physical contact (collision).

Let f_{ij} be the force that describes the effect of the physical contact with agent j on agent i. Let k be a global parameter defining the strength of the physical contact forces within the simulation. Also, let $\vec{n}_{ij}$ be the unit vector pointing from j to i. Let d_{ij} define the distance between the centers of mass of agents i and j. Then:

$$f_{ij} = kg(r_i + r_j - d_{ij})\vec{n}_{ij}, \tag{3.4}$$

where

$$g(x) = \begin{cases} x & \text{if } x \geq 0 \\ 0 & \text{if } x < 0 \end{cases} . \tag{3.5}$$

Also, let f_{iW} be the force that describes the effect of the physical contact with wall W on agent i. Let d_{iW} be the distance between agent i and wall W. Let $\vec{n_{iW}}$ define the unit vector that is perpendicular to wall W and points from wall W towards agent i. Then:

$$f_{iW} = kg(r_i - d_{iW})\vec{n}_{iW}. \tag{3.6}$$

Acceleration. At each simulation step, the acceleration $\mathrm{d}\vec{v_i}/\mathrm{d}t$ of agent i is determined by combining its internal motivation to adopt a desired velocity $\vec{v}_{des}$ with the effect of all physical-contact interactions. It is computed using the following equation:

$$\frac{\mathrm{d}\vec{v_i}}{\mathrm{d}t} = \frac{\vec{v}_{des} - \vec{v}_i}{\tau} + \frac{\sum_j \vec{f}_{ij}}{m_i} + \frac{\sum_W \vec{f}_{iW}}{m_i}. \tag{3.7}$$

3.2 Social-force model for simulating social groups

In [66], Moussaïd et al. 2010 study the motion behavior of small pedestrian groups. Initially, they study video recordings of pedestrians and identify groups of people that walk together and exhibit social interactions. Then, they examine the relative trajectories of the identified group members. Conclusions are drawn on the spatial organization of social groups based on the empirical observations. Finally, a social-force model is proposed for simulating the observed walking behavior of social groups.

According to the observations, a small group of pedestrians typically consists of up to four members that walk in a horizontal line formation at low crowd densities (*line-abreast* formation, see Figure 3.3a). At moderate crowd densities, the formation is adjusted to a V-shape that allows for social interaction between the group members (*V-like* formation, see Figure 3.3b). At high crowd densities, the group members form a lane and a leader-follower behavior is adopted so that there is enough clearance for the group to proceed towards its destination (*river-like* formation, see Figure 3.3c).

Additionaly, the authors notice that at moderate crowd densities, the V-like formation that pedestrian social groups tend to adopt does not facilitate crowd flow. A V-like formation acts as a trade-off between maintaining social interactions within a group and keeping a safety distance from the rest of the crowd. The V-like formation is not aerodynamic, because it is concave at its front (with respect to the walking direction of the group). This leads to a reduced crowd flow, which indicates the significance of social interactions when examining crowd dynamics. Next, we discuss the details of the proposed social force model that simulates pedestrian social group behavior.

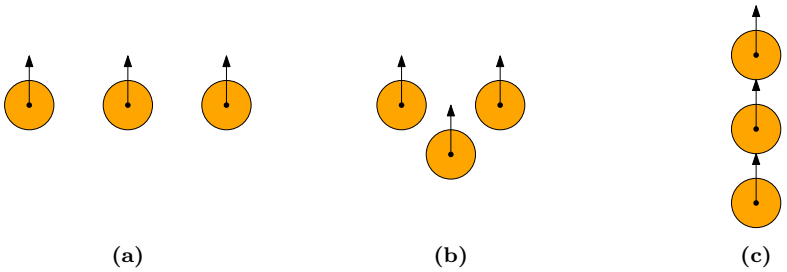

Figure 3.3: Typical formations of a social group of three members according to local crowd density. (a) Line-abreast formation at low crowd densities. (b) V-like formation at moderate crowd densities. (c) River-like formation at high crowd densities.

Acceleration. Let x_i and $\vec{\mathrm{v}}_i$ define the position and velocity of agent i, respectively. At each simulation step, the acceleration $\mathrm{d}\vec{\mathrm{v}}_i/\mathrm{d}t$ of agent i is computed using the following equation:

$$\frac{\mathrm{d}\vec{\mathrm{v}}_i}{\mathrm{d}t} = \overrightarrow{f}_i^{\,0} + \overrightarrow{f}_i^{\,wall} + \sum_j \overrightarrow{f}_{ij} + \overrightarrow{f}_i^{\,group} \ . \tag{3.8}$$

$\overrightarrow{f}_i^{\,0}$ is an attractive force that models the desire of agent i to move towards a certain direction at a preferred speed. $\overrightarrow{f}_i^{\,wall}$ is a repulsive force that models the desire of agent i to keep clear from static obstacles. $\overrightarrow{f}_{ij}$ is a force that describes the effect on agent i from its interaction with agent j. $\overrightarrow{f}_i^{\,group}$ models the effect that the rest of the group members have on agent i. The first three factors of Equation (3.8) are computed based on the work of Moussaïd et al. 2009 [64]. The group force $\overrightarrow{f}_i^{\,group}$ is the sum of three factors:

$$\overrightarrow{f}_i^{\,group} = \overrightarrow{f}_i^{\,vis} + \overrightarrow{f}_i^{\,att} + \overrightarrow{f}_i^{\,rep} \ . \tag{3.9}$$

Visual group force. The term $\overrightarrow{f}_i^{\,vis}$ is a deceleration force applied to the agent by the center of mass of the rest group members c_i. This force represents the desire of agent i to keep c_i within its field of view as an attempt to remain social. Let α_i define the minimum angle of rotation of the gazing vector $\vec{H}_i$ that is required for agent i to have c_i within its field of view. The force factor $\overrightarrow{f}_i^{\,vis}$ is computed using the following equation:

$$\overrightarrow{f}_i^{\,vis} = -\beta_1 \alpha_i \overrightarrow{V}_i, \tag{3.10}$$

where $\overrightarrow{V}_i$ is the velocity vector of agent i and β_1 is a model parameter describing the strength of the social interactions between the group members.

Attractive group force. The term $\overrightarrow{f}_i^{\,att}$ is an attractive force applied to agent i towards the center of mass of the group. This force is applied only when the distance between agent i and the center of mass of the group exceeds a threshold distance. The threshold distance is estimated at $(N-1)/2$ where N is the group size. This force represents the desire of the agent i to stay close to the rest of the group members as an effort to maintain group coherence. The force factor $\overrightarrow{f}_i^{\,att}$ is computed using the following equation:

$$\overrightarrow{f}_i^{\,att} = q_A \beta_2 \vec{U}_i \ , \tag{3.11}$$

where $\vec{U}_i$ is the unit vector pointing from agent i to the center of mass of the group and β_2 is a parameter describing the strength of the modelled attraction effect. If the distance between agent i and the center of mass of the group exceeds the threshold distance, then $q_A = 1$, otherwise $q_A = 0$.

Repulsive group force. The term $\overrightarrow{f}_i^{rep}$ is a repulsive force applied to agent i in order to guarantee that group members do not overlap. Let d_{ik} define the distance between agent i and another group member k. Let d_0 define a threshold value. Agent i tries to maintain a distance of at least d_0 from any group member. The force factor $\overrightarrow{f}_i^{rep}$ is computed using the following equation:

$$\overrightarrow{f}_i^{rep} = \sum_k q_R \beta_3 \vec{W}_{ik} \ , \tag{3.12}$$

where $\vec{W}_{ik}$ is the unit vector pointing from group member k to agent i, β_3 is a parameter describing the strength of the modelled repulsion effect. $q_R = 1$ if agent i overlaps with group member k, otherwise $q_R = 0$.

4 Path planning for social groups

In this section, we propose a method, called *Social Groups and Navigation* (SGN), for steering social groups of agents in a planar homogeneous environment with polygonal obstacles. SGN adopts the hierarchical approach of the ECM Crowd Simulation Framework [4] (see Figure 1.1), thus addressing global and local aspects of navigation in social groups. First, we present all necessary assumptions and definitions of our method (see Section 4.1). In Section 4.2, we formulate high-level mechanisms for coordinating the members of each simulated social group. In Section 4.3, we discuss the global route-planning stage of SGN. This stage describes how each social group computes its path towards its destination. Next, in Section 4.4, we provide group-relevant details for the route-following stage. This stage is performed at each simulation step and is responsible for assigning a preferred velocity to each agent based on the agent's path. Finally, in Section 4.5, the actual velocity and new position of each agent are computed during the local-movement stage.

Our method is extended to account for group-splitting behavior. Details are also provided on how SGN handles single individuals that do not belong to a social group. These additional functionalities of our method are discussed in Section 4.6.

During the local-movement stage, the method combines elements from the work of Moussaïd et al. 2010 [66] and Moussaïd et al. 2011 [65]. In brief, it borrows the vision-based avoidance algorithm and the social-force model of Moussaïd et al. 2011 to describe the interactions of each agent with obstacles and other individuals. At the same time, it employs a social-force model inspired by Moussaïd et al. 2010 [66] to describe the internal social interactions between the members of each group. Adjustments have been made to both algorithms borrowed to effectively combine them and increase the plausibility of the generated motion behaviors. In Section 4.7, we compare our method against the original works of Moussaïd et al. 2010 and Moussaïd et al. 2011.

4.1 Assumptions and definitions

Let n_i be the size of a social group of agents $\mathcal{G}_i$, where $2 \leq n_i \leq n_{max}$, $1 \leq i \leq N_{\mathcal{G}}$. n_{max} defines the maximum allowed group size and $N_{\mathcal{G}}$ is the total number of social groups. Let $\mathcal{A}_{ij}$ denote an agent that is member of group $\mathcal{G}_i$, where $1 \leq j \leq n_i$. Each agent $\mathcal{A}_{ij}$ is modelled by a disc of radius r_{ij} and is centered at $\mathbf{x}_{ij}$. Similar to Moussaïd et al. 2011 [65], we relate the mass m_{ij} of agent $\mathcal{A}_{ij}$ with its radius r_{ij} using the following equation:

$$m_{ij} = 320 r_{ij}, \tag{4.1}$$

where m_{ij} is measured in kg and r_{ij} is measured in meters. Each agent $\mathcal{A}_{ij}$ has a personal space that is modelled by a disc of radius $r_{p,ij}$ centered at $\mathbf{x}_{ij}$, where $r_{p,ij} \geq r_{ij}$. Let N_W define the total number of wall line segments that model the static environment. For simplicity reasons, when we refer to the position of an agent, we refer to the position of the disc center that models the agent. Similarly, when we refer to the distance between two agents, we refer to the distance between their disc centers. Also, the distance between an agent and a wall W_u, where $u \in [1, N_W]$, is the shortest Euclidean distance between the disc center of the agent and the line segment that models the wall.

Let $v_{pref,ij}$ define the preferred speed of agent $\mathcal{A}_{ij}$ and $v_{pref,i}$ define the preferred speed of group $\mathcal{G}_{ij}$. Then, the preferred speed of each group is defined as the minimum preferred speed of all its members, i.e.:

$$v_{pref,i} = \min_{j \in [1,n_i]} (v_{pref,ij}), \forall i \in [1, N_{\mathcal{G}}]. \tag{4.2}$$

We assume, at every simulation step, that the center X_{ij} of the disc that models agent $\mathcal{A}_{ij}$ is a good approximation the location of its head projected on the navigation plane. The vision field of agent $\mathcal{A}_{ij}$ is modelled by a circular sector. The vision field ranges $\phi_{ij}\,^\circ$ to the left and $\phi_{ij}\,^\circ$ to the right of the gazing unit vector $\hat{H}_{ij}$. Let $d_{Max,ij}$ define the maximum viewing distance of agent $\mathcal{A}_{ij}$. We assume that $d_{Max,ij} > r_{ij} \forall i \in [1, N_{\mathcal{G}}] \forall j \in [1, n_i]$. Assuming a disc of radius $d_{Max,ij}$ centered at $\mathbf{x}_{ij}$, the vision field is defined as the circular

sector of that disc which has a central angle of $2\phi_{ij}$ ° and the gazing vector $\hat{H}_{ij}$ as the angle bisector of the central angle. An agent $\mathcal{A}_{ij}$ is considered to be *visible* by another agent $\mathcal{A}_{pq}$, where $p \in [1, N_{\mathcal{G}}], q \in [1, n_p] : (p \neq i) \cup (q \neq j)$, when at least one point in the disc that models agent $\mathcal{A}_{ij}$ is within the vision field of agent $\mathcal{A}_{pq}$. Let Visible$(\mathcal{A}_{ij}, \mathcal{A}_{pq})$ be the predicate that returns: a) *true*, if $\mathcal{A}_{ij}$ is visible by $\mathcal{A}_{pq}$, b) $false$, otherwise.

For each group $\mathcal{G}_i$, the set $S_i = \{\mathbf{s}_{i1}, \ldots, \mathbf{s}_{in_i}\}$ defines the start positions of its members. Let $G_i = (\mathbf{g}_i, r_{g_i})$ define the goal area of group $\mathcal{G}_i$, where $\mathbf{g}_i$ is the position of the center of the goal area and r_{g_i} is the radius of the goal area. Let $\mathbf{a}$, $\mathbf{b}$ define the position of points A and B, respectively. Let Reachable$(\mathbf{a}, \mathbf{b}, c)$ be the predicate that returns: a) *true*, if the line segment with A and B as its endpoints has a shortest Euclidean distance of at least c from every wall segment W_u of the static environment, b) $false$, otherwise. A point X at position $\mathbf{x}$ is inside the goal area G_i if and only if X lies within the disc of radius r_{g_i} centered at $\mathbf{g}_i$ and $\exists \epsilon > 0$: Reachable$(\mathbf{x}, \mathbf{g}_i, \epsilon)$. Finally, let $\rho(\mathbf{x}, r)$ be the density inside a disc of radius r centered at position $\mathbf{x}$ measured in agents per square meter.

Leader and last member. At each simulation step, every social group has a *leader* and a *last member*. Let $d_g(\mathcal{A}_{ij})$ be the curve-length distance of an agent $\mathcal{A}_{ij}$ from goal position $\mathbf{g}_i$ of the group measured along $\mathcal{A}_{ij}$'s global path (Section 4.3). Let $L_i \in [1, n_i]$ indicate the leader $\mathcal{A}_{iL_i}$ of group $\mathcal{G}_i$. Let $l_i \in [1, n_i]$ indicate the last member $\mathcal{A}_{il_i}$ of group $\mathcal{G}_i$. For each group, the leader is defined as the member whose curve-distance from the goal position of the group measured along the global path is minimum, i.e.:

$$L_i = \underset{j \in [1, n_i]}{\operatorname{argmin}} \left(d_g(A_{ij})\right), \forall i \in [1, N_{\mathcal{G}}]. \tag{4.3}$$

Similarly, for each group, the last member is defined as the member whose curve-distance from the goal position of the group measured along the global path is maximum, i.e.:

$$l_i = \underset{j \in [1, n_i]}{\operatorname{argmax}} \left(d_g(A_{ij})\right), \forall i \in [1, N_{\mathcal{G}}]. \tag{4.4}$$

Group coherence. In order to maintain proximity between the members of a social group throughout their navigation the term *group coherence* is employed. Group coherence expresses the ability of the members of a group to remain close to each other. Depending on the relative positions of the leader and last member, we conclude on whether a social-group $\mathcal{G}_i$ is *coherent* using the following definition.

Definition 4.1. *A social group $\mathcal{G}_i$ is coherent if and only if there exists at least one point in the disc that models the leader $\mathcal{A}_{iL_i}$ that is no further from the last member $\mathcal{A}_{il_i}$ than the maximum viewing distance of the last member d_{Max,il_i}, i.e.:*

$$\text{Coherent}(\mathcal{G}_i) \iff ||\boldsymbol{x}_{iL_i} - \boldsymbol{x}_{il_i}|| \leq d_{Max,il_i} + r_{iL_i}. \tag{4.5}$$

Group sociality. In order to maintain social interactions between the members of a social group throughout their navigation the term *group sociality* is employed. Group sociality is based on the visual stimuli and relative positions of all members of a group and expresses the ability of those members to maintain visual contact and proximity with each other. Based on empirical observations [19, 25, 66, 91], we employ an inter-member *social threshold distance*, d_{social}. We assume that for a pedestrian to develop social interactions with another pedestrian of the same group, it is necessary that the two pedestrians are mutually visible and their separating distance is at most d_{social}. We assume that $d_{social} + r_{ij} \leq d_{Max,ij}$, $\forall i \in [1, N_{\mathcal{G}}] \forall j \in [1, n_i]$. We conclude on whether a group $\mathcal{G}_i$ is in a *social* configuration using the following definitions.

Definition 4.2. *A group $\mathcal{G}_i$ is in a partially social configuration if and only if for every member there exists at least one fellow member, such that these two members are mutually visible and their separating distance is at most the social threshold distance, i.e.:*

$$\text{PartiallySocial}(\mathcal{G}_i) \iff \forall j \in [1, n_i] \exists k \in [1, n_i] : \{(k \neq j) \ \wedge \text{Visible}(\mathcal{A}_{ik}, \mathcal{A}_{ij}) \wedge \text{Visible}(\mathcal{A}_{ij}, \mathcal{A}_{ik}) \wedge \ (||\boldsymbol{x}_{ij} - \boldsymbol{x}_{ik}|| \leq d_{social} + r_{ij} + r_{ik})\} . \quad (4.6)$$

Definition 4.3. *A group $\mathcal{G}_i$ is in a totally social configuration if and only if the group is partially social and every member has all fellow members within its vision field, i.e.:*

$$\text{TotallySocial}(\mathcal{G}_i) \iff \text{PartiallySocial}(\mathcal{G}_i) \wedge \Big(\forall j \in [1, n_i] \forall k \in [1, n_i] : (k \neq j) \ \text{Visible}(\mathcal{A}_{ik}, \mathcal{A}_{ij})\Big) . \quad (4.7)$$

Lemma 4.1 relates the sociality with the coherence of a group. Figures 4.1, 4.2, 4.3, and 4.3 give example configurations of a group and conlude on its coherence and sociality.

Lemma 4.1. *A social group $\mathcal{G}_i$ is in a totally social configuration only if $\mathcal{G}_i$ is coherent.*

Proof. Let $\mathcal{G}_i$ be a social group. We assume that $\mathcal{G}_i$ is in a totally social configuration. From definition 4.3, the group leader $\mathcal{A}_{iL_i}$ is visible by the last member $\mathcal{A}_{il_i}$. Therefore, $||\mathbf{x}_{iL_i} - \mathbf{x}_{il_i}|| \leq d_{Max,il_i} + r_{iL_i}$. Then, from definition 4.1, group $\mathcal{G}_i$ is coherent. □

Both coherence and sociality definitions have been formulated from an individual perspective. In essence, they assess the ability of a group to remain coherent and social by superimposing the ability of each individual member to remain close to its fellow members and develop social interactions with them. We follow this approach, because we consider the individual perception of every member a necessary and sufficient condition for a group to be classified as coherent or social.

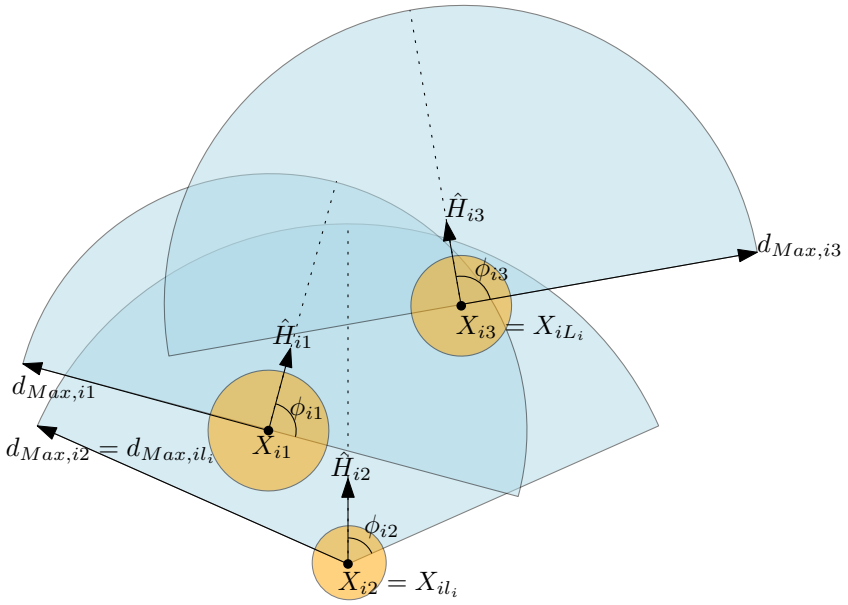

Figure 4.1: A group $\mathcal{G}_i$ of size $n_i = 3$. None of the members of $\mathcal{G}_i$ are mutually visible. Therefore, group $\mathcal{G}_i$ is not in a partially social configuration. As a result, it is not in a totally social configuration. The group is coherent, because $||\mathbf{x}_{iL} - \mathbf{x}_{il}|| \leq d_{Max,il} + r_{iL_i}$.

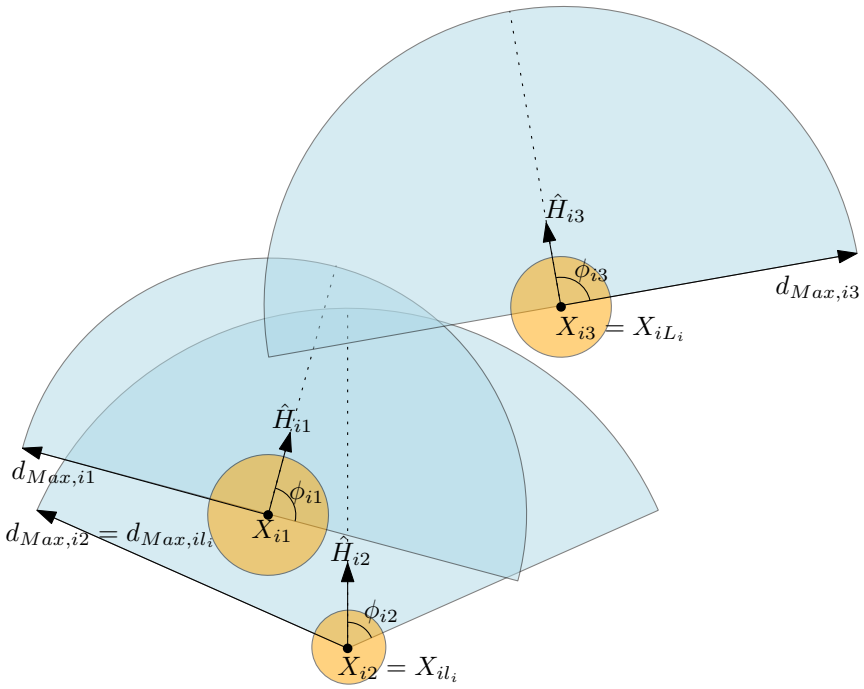

Figure 4.2: Another example, similar to Figure 4.1. The group is not coherent, because $||\mathbf{x}_{iL} - \mathbf{x}_{il}|| > d_{Max,il}$. Therefore, group $\mathcal{G}_i$ is not in totally social configuration. Also, none of the members of $\mathcal{G}_i$ are mutually visible. Therefore, group $\mathcal{G}_i$ is not in a partially social friendly configuration.

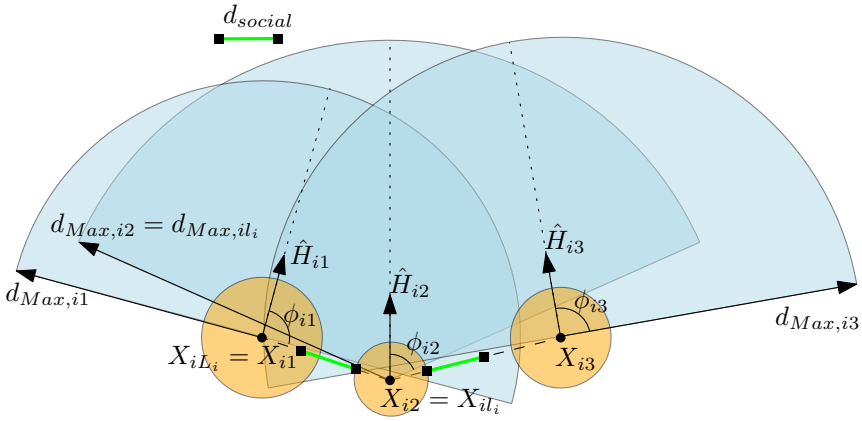

Figure 4.3: Another example, similar to Figures 4.1 and 4.2. The group is coherent, because $||\mathbf{x}_{iL} - \mathbf{x}_{il}|| \leq d_{Max,il} + r_{iL_i}$. All members are visible to each other, but no fellow member exists for $\mathcal{A}_{i3}$ with a separating distance of at most d_{social}. Therefore, the group is not in a partially social configuration. As a result, the group is not in a totally social configuration.

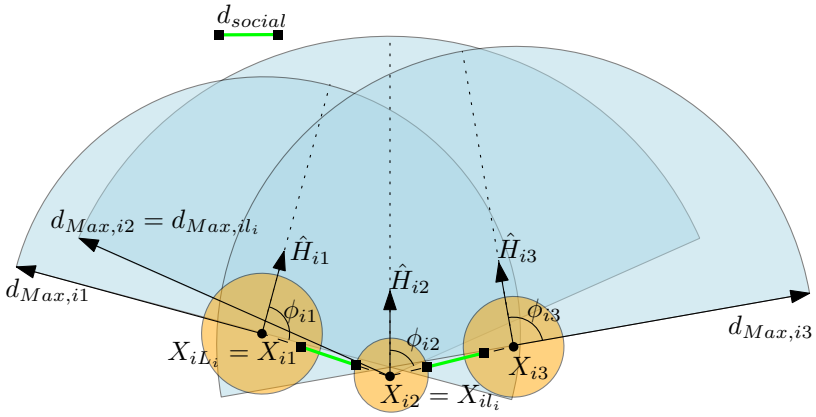

Figure 4.4: Another example, similar to Figures 4.1, 4.2 and 4.3. Here $||\mathbf{x}_{i1} - \mathbf{x}_{i2}|| \leq d_{social} + r_{i1} + r_{i2}$ and $||\mathbf{x}_{i2} - \mathbf{x}_{i3}|| \leq d_{social} + r_{i2} + r_{i3}$. Additionally, all members are visible to each other. Therefore the group is in a partially and totally social configuration. As a result, the group is coherent.

4.2 High-level planning

In the frame of this work, we have avoided using group-coordination mechanisms in the lower levels of the simulation. Approaches like the algorithm of Karamouzas and Overmars [48] require, at each simulation step, that the new velocities of the members of a social group must be jointly optimized. Performance results [48] indicate that such an approach is computationally expensive therefore unlikely to meet our project goals (see Section 1.2). Also, in the algorithm of Karamouzas and Overmars, an agent avoids other groups as single entities to simulate a phenomenon that has been observed in real life by Cheyne and Efran [13]. However, it should be noted that, in the study of Cheyne and Efran [13], only recorded pedestrians that were not walking in groups were taken into consideration [1]. Therefore, their observations have an empirical basis only for single pedestrians. As a result, we have chosen not to model this avoidance behavior in our method. On the other hand, early experiments of our method proved that it is insufficient to solely rely on the group-related lower-level mechanisms of each member to guarantee coherence for each group. It has been observed in those experiments that in densely populated areas and bottlenecks, members of the same group often separate and group coherence is violated. This is something we expect to happen in real-life, although relevant research is required to validate our claim. In any case, we expect groups that have lost their coherence to attempt to regain it.

For these reasons, we have introduced group-coordination mechanisms on the higher-level of the simulation. Navigation of a social group towards its goal and maintenance of group coherence are achieved by alternating between group phases. States are also defined for each

[1] *The behavior of all pedestrians who passed the confederates and who were not walking in groups was recorded by several observers who were positioned where they could obtain a clear view without themselves interfering with the subjects' behavior [13].*

member of a group to implement the group phases. The rest of this section discusses the details of these mechanisms.

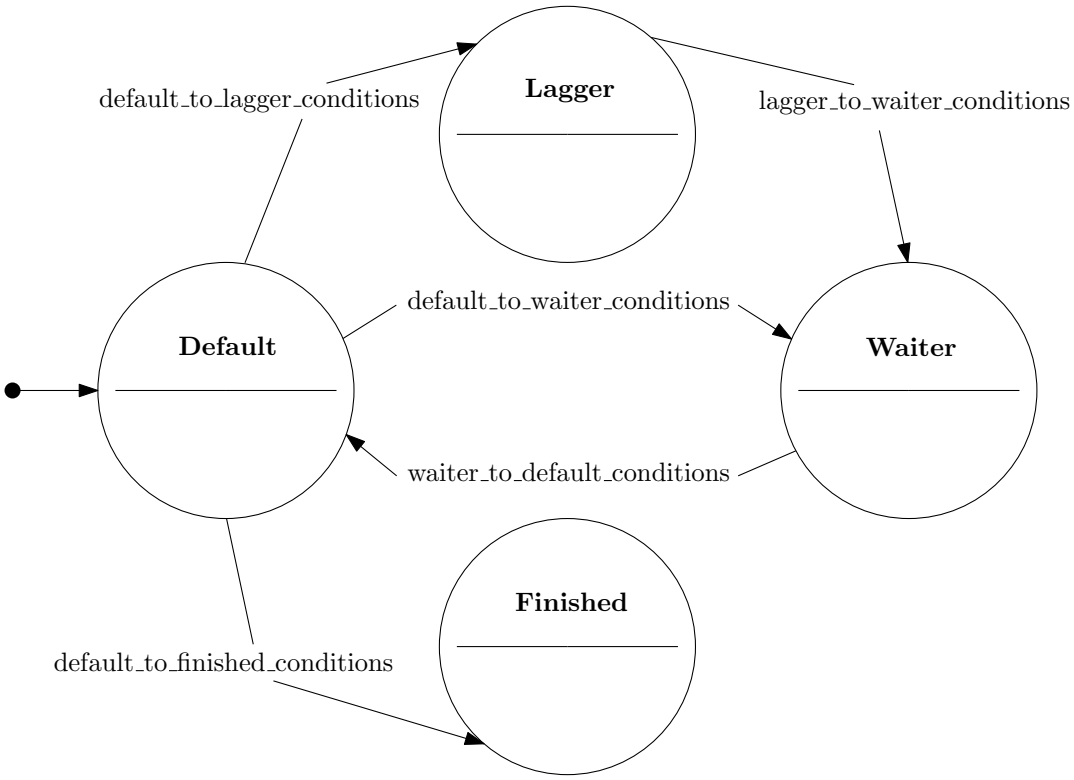

Figure 4.5: State diagram of a social-group member. The transition conditions of member $\mathcal{A}_{ij}$ of group $\mathcal{G}_i$ are:

- default_to_waiter_conditions:
  $\text{Default}(\mathcal{A}_{ij}) \wedge \text{Coordination}(\mathcal{G}_i) \wedge (j = L_i)$
- default_to_lagger_conditions:
  $\text{Default}(\mathcal{A}_{ij}) \wedge \text{Coordination}(\mathcal{G}_i) \wedge (j \neq L_i)$
- default_to_finished_conditions:
  $\text{Default}(\mathcal{A}_{ij}) \wedge$
  $\Big(\Big((||\mathbf{x}_{ij} - \mathbf{g}_i|| \leq r_{ij} + r_g) \wedge \big(\exists \epsilon > 0 : \text{Reachable}(\mathbf{x}_{ij}, \mathbf{g}_i, \epsilon)\big)\Big) \vee$
  $\Big(\exists k \in [1, n_i] : \{\text{Finished}(\mathcal{A}_{ij}) \wedge \big(||\mathbf{x}_{ij} - \mathbf{x}_{ik}|| \leq r_{p,ij} + r_{ik}\big) \wedge \text{Visible}(\mathbf{x}_{ik}, \mathbf{x}_{ij})\}\Big)\Big)$
- lagger_to_waiter_conditions:
  $\text{Lagger}(\mathcal{A}_{ij}) \wedge$
  $\Big(\exists k \in [1, n_i] : \{\text{Waiter}(\mathcal{A}_{ik}) \wedge \big(||\mathbf{x}_{ij} - \mathbf{x}_{ik}|| \leq r_{p,ij} + r_{ik}\big) \wedge \text{Visible}(\mathbf{x}_{ik}, \mathbf{x}_{ij})\}\Big)$
- waiter_to_default_conditions:
  $\text{Waiter}(\mathcal{A}_{ij}) \wedge \text{Coordination}(\mathcal{G}_i) \wedge \big(\nexists k \in [1, n_i] : \text{Lagger}(\mathcal{A}_{ik})\big)$

Member states. Figure 4.5 displays the state diagram of a social-group member. At each simulation step, a member of a group can be in one the following states:

- *Finished*: An agent who has reached the goal. An agent reaches the goal either when the agent intersects the goal area of its group or when the agent has inside its personal space a visible fellow member that is in a finished state. An agent in a finished state remains at its position.
- *Default*: An agent who is not in a finished state and whose group is in *walking* phase. An agent in a default state follows its path to the goal of the group.

- *Waiter*: An agent whose group is in *coordination* phase and is either the leader of the group or has inside its personal space a visible fellow member that is in a waiter state. An agent in a waiter state remains at its position.
- *Lagger*: An agent whose group is in coordination phase and does not have inside its personal space a visible fellow member that is in a waiter state. An agent in a lagger state follows its path to the leader of the group.

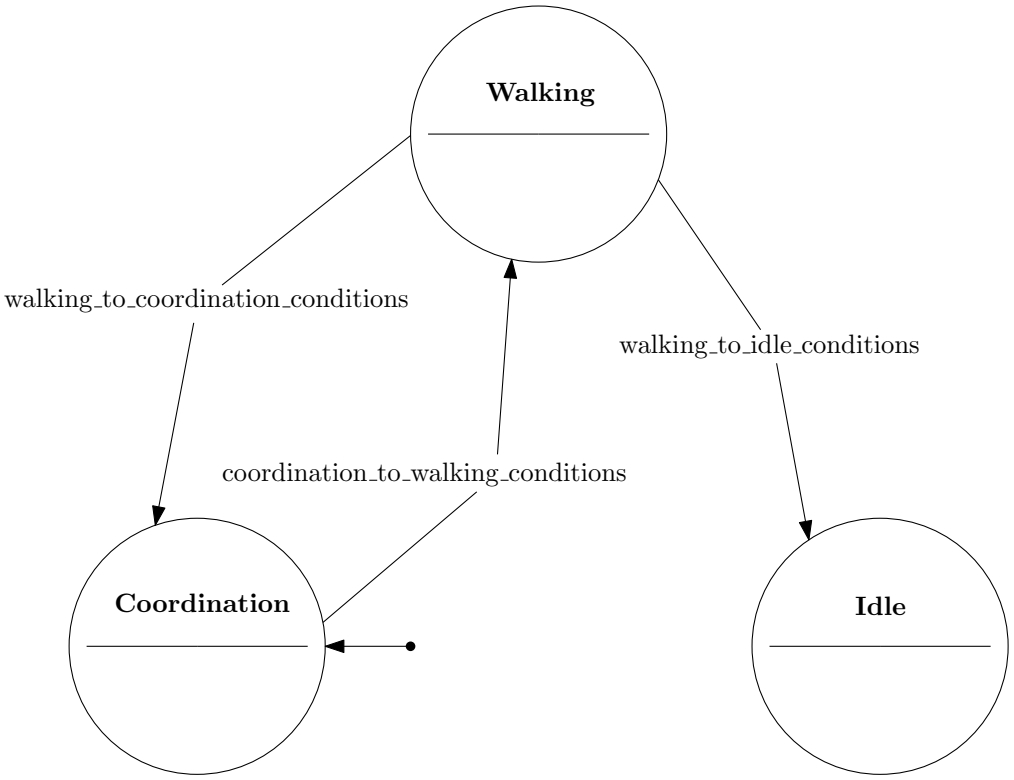

Figure 4.6: State diagram of a social group. The transition conditions of group $\mathcal{G}_i$ are:

- walking_to_idle_conditions:
  $\text{Walking}(\mathcal{G}_i) \wedge \big(\forall j \in [1, n_i]\text{Finished}(\mathcal{A}_{ij})\big)$
- walking_to_coordination_conditions:
  $\text{Walking}(\mathcal{G}_i) \wedge \big(\neg\text{Coherent}(\mathcal{G}_i)\big) \wedge \big(\rho(\mathbf{x}_{iL_i}, 1) < 0.7\big)$
- coordination_to_walking_conditions:
  $\text{Coordination}(\mathcal{G}_i) \wedge \big(\nexists k \in [1, n_i]\text{Lagger}(\mathcal{A}_{ik})\big)$

Group phases. Figure 4.6 displays the state diagram of a social group. At each simulation step, each social group can be in one of the following phases:

- *Coordination phase:* A group enters a coordination phase for two possible reasons: a) the group has just entered the simulation, or b) the group has lost its coherence. When a group $\mathcal{G}_i$ enters the coordination phase, its leader $\mathcal{A}_{iL_i}$ enters a waiter state and remains stationary at its current position $\mathbf{x}_{iL_i}$. The remaining group members $\mathcal{A}_{ik}, \forall k \in [1, n_i] : k \neq L_i$ enter a lagger state, plan individual global paths from their current positions $\mathbf{x}_{ij}$ to the position of the leader $\mathbf{x}_{iL_i}$ and follow these paths. Throughout the coordination phase, all group-relevant aspects of the lower simulation levels (Sections 4.3, 4.4, and 4.5) are not included in the motion-planning computations of the laggers. When a lagger has a waiter in its personal space, the lagger switches its state to waiter. When all members have switched to waiter state (i.e. no more laggers

exist), the leader $\mathcal{A}_{iL_i}$ plans a path from its current position $\mathbf{x}_{iL_i}$ to the goal of the group g_i using the approach detailed in Section 4.3. The remaining group members $\mathcal{A}_{ik}, \forall k \in [1, n_i] : k \neq L_i$ merge their individual paths towards the leader $\mathcal{A}_{ij}$ with the path that the leader has just computed. Finally, all group members switch their state from waiter to default and the group enters a walking phase.

- *Walking phase:* For each social group, a walking phase begins when a coordination phase ends. During the walking phase, at least one member of the group is in a default state and follows its path to the goal of the group (route following, see Section 4.4) while avoiding collisions with other agents and remaining social within its group (local movement, see Section 4.5).

- *Idle phase:* For each social group, an idle phase begins when all members of that group are in a finished state. During the idle phase, all group members remain stationary. In other words, a group in an idle phase is a group whose members have reached the goal.

When a group $\mathcal{G}_i$ first enters the simulation, it also enters a coordination phase. At this time only, the leader $\mathcal{A}_{iL_i}$ cannot be defined using Equation 4.3, because no global path has been computed yet for the members of $\mathcal{G}_i$. At the start of the simulation, we arbitrarily select a member $\mathcal{A}_{iL_{start,i}}$ from each group $\mathcal{G}_i$ to act as the leader of that group, where $1 \leq L_{start,i} \leq n_i$.

At the end of each simulation step, we update the leaders and last members of every group that is in a walking phase, according to Equations 4.3 and 4.4. For each group that is in a walking phase, we check whether the coherence condition is met (see Definition 4.1). We would like every group that has lost its coherence to enter a coordination phase. However, if a leader stopped at a densely populated region due to its group having lost coherence, local minima would be likely to occur to the crowd flow. At the same time, we expect that, in real-life, a member of a group that has lost its coherence waits for its fellow members within a non-congested region for visibility reasons. For these reasons, a group enters a coordination phase if and only if the coherence of the group is lost and the local density around the leader of the group is below 0.7 pedestrians per square meter. However, we have no empirical data to support our claim and future research is required to validate our approach. The density threshold is based on the *Pedestrian Level Of Service* (PLOS), as proposed by Fruin [26], and suggests an acceptable density for pedestrian flow. The local density is measured over a disc that is centered at the position of the leader and has a radius that is equal to one meter. For each group that is in a coordination phase, we check whether no more laggers exist for that group. Any group that is in a coordination phase and has no laggers enters a walking phase.

4.3 Global route planning

In agent-based motion-planning methods, the global route-planning level of the simulation is responsible for computing a path from the starting position of an agent towards its destination. For the path to be feasible, the corresponding algorithm must take into account the presence of static obstacles as well as the size of the moving agent. Depending on the quality of the path, global route planning can be an expensive task that needs to search separately the underlying navigation mesh for each agent.

In the proposed method, all members of each social group are assumed to have a common destination. Therefore, for each group it is efficient to plan a path to the destination for one member (leader) and share its path with its fellow members. For this reason, all members must first gather around the leader, before utilizing the path of the leader towards the destination of the group. The gathering of the members is achieved through the coordination phase as explained in Section 4.2. As soon as all members have reached the leader, the leader

plans a path to the destination of the group. The global path is computed to account for clearance that corresponds to the largest group member. As a result, the computed path is feasible and shared by all members, while guaranteeing that all members will have homotopic trajectories throughout their navigation. Any global route-planning method that accounts for clearance can be employed at this step [28, 44, 46, 89].

4.4 Route following

The *route-following* stage determines a preferred velocity $\vec{v}_{pref,ij}$ for each agent $\mathcal{A}_{ij}$ based on its current position $\mathbf{x}_{ij}$, a preferred speed and its global route that was computed at the global route-planning stage (see Section 4.3). This stage is repeated for all agents at each simulation step. Any agent-based route-following algorithm [41, 46] can be employed to implement this stage for our method. For each agent $\mathcal{A}_{ij}$, if group $\mathcal{G}_i$ is in a coordination phase, we set the individual preffered speed $v_{pref,ij}$ to be the preffered speed that the route-following algorithm utilizes. If group $\mathcal{G}_i$ is in a walking phase, the route-following algorithm uses the preferred speed of the group $v_{pref,i}$ (see Equation 4.2) for the computation of the preferred velocity $\vec{v}_{pref,ij}$ of each member $\mathcal{A}_{ij}$.

According to our route-following approach, it is possible that all group members compete against each other for moving to the same position. Nevertheless, the local-movement simulation stage (see Section 4.5) ensures that each member will minimally adjust its direction of motion so as to avoid collisions. This approach has been proven sufficient to enable group members walk next to each other, despite having preferred velocities that point towards the same point on their global routes.

4.5 Local movement

At each simulation step, the *local-movent* stage determines the new position and velocity of each agent $\mathcal{A}_{ij}$ based on its preferred velocity $\vec{v}_{pref,ij}$ that was computed at the route-following stage (see Section 4.4) and the local conditions of $\mathcal{A}_{ij}$. Local movement is performed in two steps. First, the desired velocity $\vec{v}_{des,ij}$ of agent $\mathcal{A}_{ij}$ is computed during local collision avoidance (Section 4.5.1). Next, the acceleration $\frac{\mathrm{d}\vec{v}_{ij}}{\mathrm{d}t}$ of each agent $\mathcal{A}_{ij}$ is computed based on the proposed social-force model (Section 4.5.2). Finally, the new position and velocity of agent $\mathcal{A}_{ij}$ are computed using Euler integration [3] based on its current position $\mathbf{x}_{ij}$, velocity $\vec{v}_{ij}$ and acceleration $\frac{\mathrm{d}\vec{v}_{ij}}{\mathrm{d}t}$.

4.5.1 Local collision avoidance

The proposed local collision-avoidance algorithm borrows the approach of Moussaïd et al. 2011 [65]. Adjustments have been made to the original algorithm to allow for more efficient avoidance maneuvers. We have observed that the original algorithm fails to implement the avoidance mechanism that the first heuristic dictates (see Section 3.1). According to the heuristic, at each simulation step, the desired angle of motion of an agent should lead to *"the most direct path to destination point taking into account the presence of obstacles [65]"*. Based on Equation 3.1, the algorithm is expected to choose a desired angle of motion such that the distance to the destination is minimum when the next directional change should occur. The authors suggest that for each candidate angle, if no collision occurs within the maximum viewing distance, then the distance to first collision for that angle is set to be equal to the maximum viewing distance (see Section 3.1). This approach might discourage choosing a candidate angle that would lead to an efficient avoidance maneuver towards the destination. The reason is that, on a collision-free candidate direction, it might be more efficient to turn towards the destination before covering the maximum viewing distance on that direction.

Instead, we propose a different approach for the computation of the desired angle of motion. Let X_{ij} be a point at position $\mathbf{x}_{ij}$ of agent $\mathcal{A}_{ij}$ and O_{ij} be the last visible point on the direction of the preferred velocity $\vec{v}_{pref,ij}$. Let $\alpha_{0,ij}$ be the bearing angle of O_{ij}. Let $\Omega_{\alpha_{ij}}$ be the last visible point on the direction that corresponds to candidate angle

$\alpha_{ij} \in [-\phi_{ij}, +\phi_{ij}]$. Let $F_{\alpha_{ij}}$ be the foot of O_{ij} on $X_{ij}\Omega_{\alpha_{ij}}$. If $T_{\alpha_{ij}}$ is a point on the last collision-free position for the agent on the direction that corresponds to α_{ij}, then we compute $f(\alpha_{ij})$ using the following equation:

$$f(\alpha_{ij}) = \min\left(|X_{ij}T_{\alpha_{ij}}|, |X_{ij}F_{\alpha_{ij}}|\right). \tag{4.8}$$

Similar to the original method [65], the remaining distance $d(\alpha_{ij})$ to the destination is given by:

$$d(\alpha_{ij}) = \sqrt{d_{Max,ij}^2 + f(\alpha_{ij})^2 - 2d_{Max,ij}f(\alpha_{ij})\cos\left(\alpha_{0,ij} - \alpha_{ij}\right)}. \tag{4.9}$$

Figures 4.7 and 4.8 give two examples.

In other words, assuming an agent decides to walk along an arbitrary direction, then the most efficient choice is to keep walking on that direction until the destination has a bearing of $\pm 90\,^\circ$. Then, the agent should turn towards its destination and walk straight ahead (Figure 4.7). However, the agent should anticipate collisions in the chosen walking direction and expect to turn towards the destination earlier, if necessary (Figure 4.8). Similar to the method of Moussaïd et al. 2011 [65], the chosen walking direction should correspond to the minimum remaining distance to the destination, when the next directional change is expected. Therefore, the desired angle of motion $\alpha_{des,ij}$ of agent $\mathcal{A}_{ij}$ is given by:

$$\alpha_{des,ij} = \underset{\alpha_{ij} \in [-\phi_{ij}, +\phi_{ij}]}{\operatorname{argmin}} \big(d(\alpha_{ij})\big). \tag{4.10}$$

Similar to Moussaïd et al. 2011 [65], we employ the relaxation time τ to determine the desired speed $||\vec{v}_{des,ij}||$ of each agent $\mathcal{A}_{ij}$:

$$||\vec{v}_{des,ij}|| = \min\left(v_{pref,i}, d_{h_{ij}}/\tau\right), \tag{4.11}$$

where $d_{h_{ij}}$ is the distance between agent $\mathcal{A}_{ij}$ and the first obstacle on the desired angle of motion $\alpha_{des,ij}$.

Therefore, the desired velocity $\vec{v}_{des,ij}$ of each agent $\mathcal{A}_{ij}$ is determined in two steps. First, the desired direction is obtained using Equation 4.10. Next, the desired speed is computed using Equation 4.11.

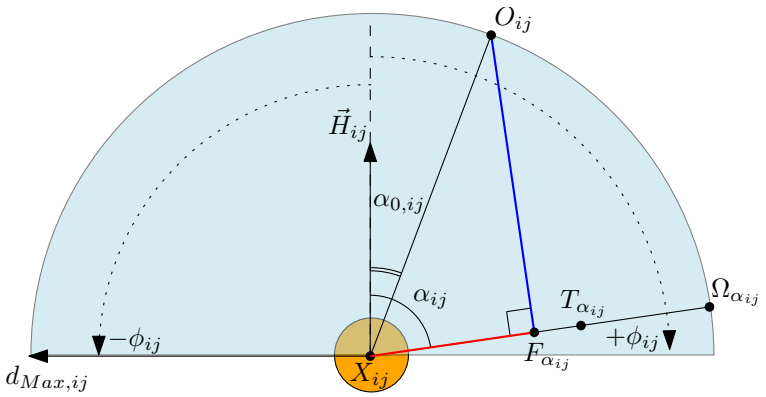

Figure 4.7: Example computation of $f(\alpha_{ij})$. In this example, $|X_{ij}F_{\alpha_{ij}}| < |X_{ij}T_{\alpha_{ij}}|$. Therefore, $f(\alpha_{ij}) = |X_{ij}F_{\alpha_{ij}}|$ (*red*) and $d(\alpha_{ij}) = |F_{\alpha_{ij}}O_{ij}|$ (*blue*).

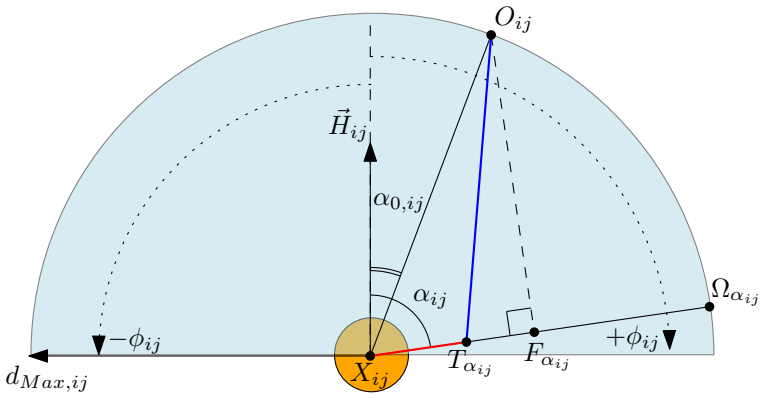

Figure 4.8: Another example computation of $f(\alpha_{ij})$. In this example, $|X_{ij}T_{\alpha_{ij}}| < |X_{ij}F_{\alpha_{ij}}|$. Therefore, $f(\alpha_{ij}) = |X_{ij}T_{\alpha_{ij}}|$ (*red*) and $d(\alpha_{ij}) = |T_{\alpha_{ij}}O_{ij}|$ (*blue*).

4.5.2 Social-force model

At each simulation step, the acceleration $\frac{d\vec{v}_{ij}}{dt}$ of agent $\mathcal{A}_{ij}$ is computed using the following equation:

$$\frac{\mathrm{d}\vec{v}_{ij}}{\mathrm{d}t} = \frac{\vec{v}_{des,ij} - \vec{v}_{ij}}{\tau} + \frac{1}{m_{ij}} \sum_{p=1}^{N_{\mathcal{G}}} \sum_{q=1}^{n_p} \vec{f}_{ij,pq} + \frac{1}{m_{ij}} \sum_{u=1}^{N_W} \vec{f}_{ij,W_u} + \frac{\vec{f}_{ij}^{group}}{m_{ij}} \ , \tag{4.12}$$

where τ is the relaxation time, $\vec{v}_{des,ij}$ is the desired velocity of agent $\mathcal{A}_{ij}$ that is provided by the local collision-avoidance step (see Section 4.5.1), $\vec{v}_{ij}$ is the current velocity of agent $\mathcal{A}_{ij}$, m_{ij} is the mass of agent $\mathcal{A}_{ij}$, and $p \in [1, N_{\mathcal{G}}], q \in [1, n_p] : p = i \Rightarrow q \neq j$. The first three terms of Equation (4.12) are computed based on the work of Moussaïd et al. 2011 [65], while the last term is computed based on the work of Moussaïd et al. 2010 [66]. Below we give the details for the computation of each term.

Physical-contact forces. In Equation 4.12, the force factor $\vec{f}_{ij,pq}$ defines the force applied to agent $\mathcal{A}_{ij}$ due to physical-contact interaction with agent $\mathcal{A}_{pq}$. Let $d_{ij,pq}$ denote the distance between agents $\mathcal{A}_{ij}$ and $\mathcal{A}_{pq}$. Let $\hat{n}_{ij,pq}$ define the unit vector pointing from agent $\mathcal{A}_{pq}$ to agent $\mathcal{A}_{ij}$, i.e. $\hat{n}_{ij,pq} = (\mathbf{x}_{ij} - \mathbf{x}_{pq})/||\mathbf{x}_{ij} - \mathbf{x}_{pq}||$. The force factor $\vec{f}_{ij,pq}$ is computed using the following equation:

$$\vec{f}_{ij,pq} = kg(r_{ij} + r_{pq} - d_{ij,pq})\hat{n}_{ij,pq}, \tag{4.13}$$

where:

$$g(x) = \begin{cases} x & \text{if } x \geq 0 \\ 0 & \text{if } x < 0 \end{cases} \ , \tag{4.14}$$

and k is a global parameter that defines the strength of physical contact forces in the simulation.

In Equation 4.12, the force factor $\vec{f}_{ij,W_u}$ defines the force applied to agent $\mathcal{A}_{ij}$ due to physical interaction with wall line segment W_u. Let d_{ij,W_u} define the distance between agent $\mathcal{A}_{ij}$ and wall W_u. Let $\hat{n}_{ij,W_u}$ define the unit vector that is perpendicular to wall W_u and points from wall W towards agent $\mathcal{A}_{ij}$. The force factor $\vec{f}_{ij,W_u}$ is computed using the following equation:

$$\vec{f}_{ij,W} = kg(r_{ij} - d_{ij,W_u})\hat{n}_{ij,W_u} \ . \tag{4.15}$$

Group forces. In Equation 4.12, the force factor $\vec{f}_{ij}^{group}$ defines the force applied to agent $\mathcal{A}_{ij}$ due to social interaction with the rest of its group members. It is computed using the following equation:

$$\vec{f}_{ij}^{group} = \vec{f}_{ij}^{vis} + \vec{f}_{ij}^{att} \ . \tag{4.16}$$

In Equation 4.16, the term $\vec{f}_{ij}^{vis}$ is a deceleration force applied to the agent as an attempt to remain social within its group. This force represents the desire of agent $\mathcal{A}_{ij}$ to keep the heads of all fellow members within its vision field. Let $\theta_{ij,ik}, k \in [1, n_i] : k \neq j$, be the minimum angle of rotation of the gazing vector $\hat{H}_{ij}$ that is required for agent $\mathcal{A}_{ij}$ to have the center X_{ik} of fellow member $\mathcal{A}_{ik}$ within its vision field. Then, we define θ_{ij} using the following equation:

$$\theta_{ij} = \max_{k \in [1,n_i]: k \neq j} (\theta_{ij,ik}), \tag{4.17}$$

The force factor $\overrightarrow{f}_{ij}^{vis}$ is computed using the following equation:

$$\vec{f}_{ij}^{vis} = -\beta_1 \theta_{ij} \vec{v}_{des,ij}, \tag{4.18}$$

where $\vec{v}_{des,ij}$ is the desired velocity of agent $\mathcal{A}_{ij}$ that the local collision avoidance produces (see Section 4.5.1) and β_1 is a model parameter describing the strength of the social interactions between the group members.

Figure 4.9 illustrates the computation of angle $\theta_{ij,ik}$ for an example scenario. Let $P_{ik,ij}$ be the projection of X_{ik} on the line of the gazing vector $\hat{H}_{ij}$. Also, let $\tau_{AB\Gamma}$ define the triangle with points A, B and Γ as its vertices. $\widehat{BA\Gamma}$ defines the angle of triangle $\tau_{AB\Gamma}$ at vertex A. In this example, the rotation angle $\theta_{ij,ik}$ is computed from the triangle $\tau_{X_{ij}X_{ik}P_{ik,ij}}$ (see Figure 4.9b), as follows:

$$\begin{aligned}
\theta_{ij,ik} &= \widehat{P_{ik,ij}X_{ij}X_{ik}} - \phi_{ij} \Leftrightarrow \\
\theta_{ij,ik} &= \arccos\left(\frac{||\overrightarrow{X_{ij}P_{ik,ij}}||}{||\overrightarrow{X_{ij}X_{ik}}||}\right) - \phi_{ij} \Leftrightarrow \\
\theta_{ij,ik} &= \arccos\left(\frac{||\overrightarrow{X_{ij}X_{ik}} \cdot \hat{H}_{ij}||}{||\overrightarrow{X_{ij}X_{ik}}||}\right) - \phi_{ij} \ .
\end{aligned} \tag{4.19}$$

Figure 4.10 gives a different example, where the bearing angle of X_{ik} is greater than $90\,^\circ$. In this case, the rotation angle $\theta_{ij,ik}$ is computed from the triangle $\tau_{X_{ij}X_{ik}P_{ik,ij}}$ (see Figure 4.10b), as follows:

$$\begin{aligned}
\theta_{ij,ik} &= 180\,^\circ - \widehat{P_{ik,ij}X_{ij}X_{ik}} - \phi_{ij} \Leftrightarrow \\
\theta_{ij,ik} &= 180\,^\circ - \arccos\left(\frac{||\overrightarrow{X_{ij}P_{ik,ij}}||}{||\overrightarrow{X_{ij}X_{ik}}||}\right) - \phi_{ij} \Leftrightarrow \\
\theta_{ij,ik} &= 180\,^\circ - \arccos\left(\frac{||\overrightarrow{X_{ij}X_{ik}} \cdot \hat{H}_{ij}||}{||\overrightarrow{X_{ij}X_{ik}}||}\right) - \phi_{ij}
\end{aligned} \tag{4.20}$$

Equations 4.19 and 4.20 can be combined to provide a general solution for the computation of angle $\theta_{ij,ik}$:

$$\theta_{ij,ik} = \arccos\left(\frac{\overrightarrow{X_{ij}X_{ik}} \cdot \hat{H}_{ij}}{||\overrightarrow{X_{ij}X_{ik}}||}\right) - \phi_{ij} \ . \tag{4.21}$$

Determining whether the rotation by angle $\theta_{ij,ik}$ is clockwise or anticlockwise reduces to checking whether X_{ik} lies to the right or to the left of the agent, respectively. The right and left sides of the agent are defined by the half-planes that the gazing vector $\hat{H}_{ij}$ induces. If $||\overrightarrow{X_{ij}X_{ik}}|| > d_{Max,ij}$ then no rotation is possible for making X_{ik} visible to agent $\mathcal{A}_{ij}$, because the distance between members $\mathcal{A}_{ij}$ and $\mathcal{A}_{ik}$ exceeds the maximum viewing distance of $\mathcal{A}_{ij}$.

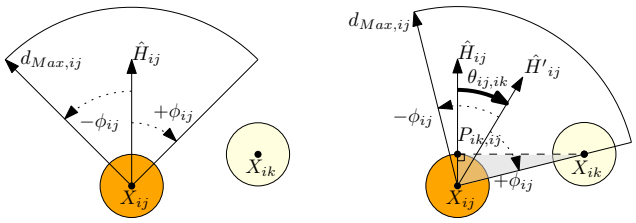

Figure 4.9: Computation of the rotation angle $\theta_{ij,ik}$ that is necessary for the agent $\mathcal{A}_{ij}$ (*orange*) to have X_{ik} within its vision field. (a) The current vision field of agent $\mathcal{A}_{ij}$ and gazing direction $\hat{H}_{ij}$. (b) The desired vision field of agent $\mathcal{A}_{ij}$ based on the rotated gaze vector $\hat{H}'_{ij}$. The desired vision field includes X_{ik} at its border. The angle of rotation $\theta_{ij,ik}$ can be computed from the triangle $\tau_{X_{ij}X_{ik}P_{ik,ij}}$ (*gray*). The bearing angle $(\theta_{ij,ik} + \phi_{ij})$ of X_{ik} is less than 90 °.

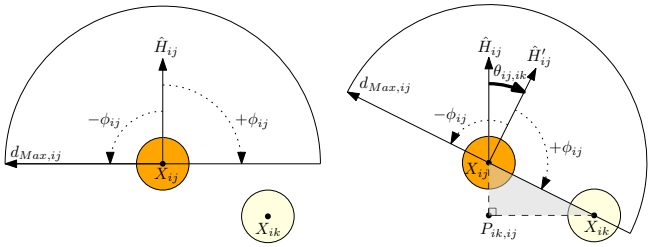

Figure 4.10: Another example similar to Figure 4.9. The bearing angle $(\theta_{ij,ik} + \phi_{ij})$ of X_{ik} is greater than 90 °.

In Equation 4.16, the term $\vec{f}_{ij}^{att}$ is an attractive force applied to agent $\mathcal{A}_{ij}$ by the centroid C_i of group $\mathcal{G}_i$. This force represents the desire of the agent $\mathcal{A}_{ij}$ to stay close to the rest of the group members as an effort to maintain group coherence. Let d_{ij,C_i} define the distance between agent $\mathcal{A}_{ij}$ and centroid C_i. Similar to Moussaïd et al. 2010 [66], let $d_{C_i,threshold} = (n_i - 1)/2$ define a threshold distance. It is reminded that n_i is already defined as the size of group $\mathcal{G}_i$. The agent should be attracted by the centroid of the group when the distance d_{ij,C_i} exceeds the threshold distance. The force factor $\vec{f}_{ij}^{att}$ is computed using the following equation:

$$\vec{f}_{ij}^{att} = q_A(d_{ij,C_i})\beta_2\hat{U}_{ij} \ , \tag{4.22}$$

where $\hat{U}_{ij}$ is the unit vector pointing from agent $\mathcal{A}_{ij}$ to the centroid C_i, β_2 is a parameter describing the strength of the modelled attraction effect, and:

$$q_A(d_{ij,c_i}) = \begin{cases} 1 & \text{if } d_{ij,C_i} \geq d_{C_i,threshold} \\ 0 & \text{if } d_{ij,C_i} < d_{C_i,threshold} \end{cases} \ . \tag{4.23}$$

4.6 Method extensions

In this section, we discuss extensions to our base method. First, we provide details on how group-splitting behavior is modelled in SGN. Next, we describe how single individuals are handled.

Group-splitting behavior. Costa [19] observes that, in real-life, social groups of more than three members often split to single individuals, dyads and triads. Our model can be easily

adjusted to account for this group behavior. The adjustment is applied in the computation of the visual term $\vec{f}_{ij}^{vis}$ of the group force (see Section 4.5.2, Equation 4.18). First, we consider a group that splits into sub-groups of two or three members. For each sub-group member, Equation 4.17 is modified so that it only considers fellow sub-group members. In other words, for each agent we compute the angle of rotation of its gazing vector that is required for making the head of each of its fellow sub-group members visible. The deceleration of the agent is proportional to the maximum of the computed angles. If the group is split to single individuals, then $\vec{f}_{ij}^{vis}$ is omitted from our group force (see Equation 4.16). The remainder of the method details remain identical to the base method.

Navigation of single individuals. Simulating single individuals in SGN is possible with intuitive adjustments. The main concept is to treat single individuals as groups of size one and omit all group-relevant aspects of our method. Therefore, for single individuals, we skip the high-level planning mechanisms (see Section 4.2). Regarding the global route-planning stage, we do not adopt the approach discussed in Section 4.3. Instead, we couple our method with any global route-planning algorithm that accounts for clearance [28, 44, 46, 89] and use it without further adjustments. The same applies to the route-following stage (see Section4.3); any agent-based route-following algorithm [41, 46] can be directly employed. Collision avoidance is performed as described in Section 4.5.1 with no need for modifications. Finally, the group force $\vec{f}_{ij}^{group}$ is omitted in the social-force model (see Section 4.5.2, Equation 4.12).

4.7 Method summary and comparison with works of Moussaïd et al.

The method we have proposed follows a hierarchical approach (see Figure 1.1) to simulate pedestrians that are organized in social groups. SGN addresses motion-planning for social groups both at a global and at a local level. For this reason, we have introduced high-level mechanisms that coordinate the members of each group and aid in maintaining group coherence (see Section 4.2). We have also presented novel techninques for adressing social group behavior during global route-planning and route-following levels (see Sections 4.3 and 4.4). On a local level, SGN employs a social-force model coupled with a vision-based collision-avoidance algorithm (see Section 4.5). The proposed social-force model describes the social interactions between group members and is inspired by the work of Moussaïd et al. 2010 [66] (see Section 3.2). In SGN, the vision based collision-avoidance algorithm enables predictive avoidance behavior and is heavily based on the work of Moussaïd et al. 2011 [65] (see Section 3.1). In the remainder of this section, we first discuss the differences between our social-force model and the model of Moussaïd et al. 2010. Next, our adjustments to the collision-avoidance algorithm of Moussaïd et al. 2011 are summarized.

Social-force model differences. Equation 4.12 defines the social-force model of our method. A similar model has been proposed by Moussaïd et al. [66] and is defined using Equation 3.8. As discussed in Section 3.2, in the method of Moussaïd et al. 2010, the attractive force towards the destination of an agent, the force that describes interactions with walls, and the sum of forces that describe interactions with agents that do not belong to the same group are all based on the work of Moussaïd et al. 2009 [64]. This work leads to reactive avoidance behavior and lacks anticipation. Instead, in SGN, these forces have been replaced by the corresponding forces of Moussaïd et al. 2011 [65], thus introducing predictiveness to the local-behavior of our agents.

What is more, in SGN, for the computation of the attractive term $\vec{f}_{ij}^{att}$ of the group force $\vec{f}_{ij}^{group}$, the centroid C_i of group $\mathcal{G}_i$ is employed (Equation 4.22). Instead, in the method of Moussaïd et al. 2010 [66], the relevant term employs the center of mass of the group (Equation 3.11). We follow this approach, because we assume that the mass variation between the members is irrelevant to the desired effect of this term.

Furthermore, in SGN, for the computation of the visual term $\vec{f}_{ij}^{vis}$ of the group force $\vec{f}_{ij}^{group}$, we compute the required angle of head rotation for making the head of each fellow

member visible (Equation 4.21). Then, we take the maximum of those angles into consideration for computing the amplitude of the deceleration force $\vec{f}_{ij}^{vis}$ (Equations 4.17 and 4.18). The motivation behind this approach is that we would like each member to slow down when realizing that a fellow member is not visible. Instead, in the method of Moussaïd et al. 2010 [66], the relevant term employs the required angle of head rotation for making the center of mass of all fellow members visible (Equation 3.10). Similar to the attractive group term, the effect of the approach of Moussaïd et al. 2010 is determined by the mass variation between the members. What is more, in the approach of Moussaïd et al. 2010, when a single member stays behind and depending on the group size, fellow members do not effectively slow down. This is expected to lead to group-sociality and group-coherence issues in dense scenarios.

Additionally, the visual term in the method of Moussaïd et al. 2010 [66] applies a deceleration force that is based on the current velocity of the agent (Equation 3.10). The resulting force simulates a reactive social group behavior. Instead, the visual term in our method applies a deceleration force that is based on the desired velocity of the agent that has already been computed by the local collision-avoidance simulation stage (Equation 4.18). Therefore, we introduce predictiveness into the social group behavior. Furthermore, this approach is expected to result in an acceleration vector $\frac{\mathrm{d}\vec{v}_{ij}}{\mathrm{d}t}$ whose direction better simulates the desired effect.

Finally, the social-force model of of Moussaïd et al. 2010 [66] includes a repulsive term in its group force for modelling the effect of physical contact between members of the same group (Equation 3.12). Instead, our method does not include such a term. In SGN, the effect of physical contact between two agents is defined using Equation 4.13 and is invariant to the group origin of each agent.

Collision-avoidance adjustments. In the original work of Moussaïd et al. 2011 [65], an agent is expected to move on each candidate direction until a directional change is required due to the presence of an obstacle or another agent. Then, the agent is expected to change its direction of motion and move towards its destination. If no collisions are expected on a candidate direction, then the agent is expected to move in that direction for a distance that is equal to the maximum viewing distance before turning towards its destination. The algorithm of Moussaïd et al. 2011 picks the candidate direction that minimizes the distance to the destination when the next directional change is expected.

In SGN, an adjustment has been made in this avoidance scheme described above. The travel distance on each candidate direction is estimated using Equation 4.8. In essence, this equation dictates that, for each candidate direction, an agent is expected to move on that direction until the agent either has to avoid a collision or the destination has a bearing of $\pm 90\,^{\circ}$. This approach is expected to improve the efficiency of the avoidance maneuvers. The reader is referred to Section 4.5.1, for further discussion on this adjustment. The remainder of the avoidance scheme is identical to Moussaïd et al. 2011 [65].

Finally, in SGN, the notation has been updated to allow for agents to differ from each other in terms of their viewing angle and maximum viewing distance. As a result, crowd heterogeneity can be exhibited regarding the visual ability of the simulated pedestrians.

5 Experiments

In this section we discuss the experiments we have conducted to evaluate our method. First, we provide details on regarding the set-up of the experiments and the implementation of SGN (Section 5.1). Next, we briefly describe the scenarios we have employed for our experiments (Section 5.2). The quality of our method is evaluated in Section 5.3. Finally, in Section 5.4 we provide results regarding running-time performance.

5.1 Set-up and implementation details

The proposed method has been implemented in C++ within the Explicit Corridor Map (ECM) Crowd Simulation Framework [4]. The code was compiled and tested in Visual Studio 2013. All experiments were performed on a PC with an Intel Core i7 860 processor (2.8 GHz), an nVidia GeForce GTX 285 graphics card and 8 GB of RAM. The operating system was Windows 7 Ultimate 64-bit with Service Pack 1. Throughout the experiments a single core was used and we disabled visualization.

For the global route-planning phase of our method, we employed the algorithm described in [28]. For the route-following phase, the Modified Indicative Routes and Navigation (MIRAN) [41] method was used.

We also provide performance results of a parallel implementation of our method. These results have limited comparison value. Nevertheless, they demonstrate the ability of our method to simulate in real-time scenarios that involve thousands of pedestrians that navigate in small social groups.

Steerbench issues. Initially, it was our aim to use Steerbench [83] for the evaluation of the individual behavior that our method generates. Steerbench provides an agent-based benchmark score that accounts for collisions, energy efficiency and travel time. However, after integrating with our framework, some issues were discovered. A limitation of Steerbench is that it can only run benchmarks on static environments that have been modelled by axis-aligned-bounding-boxes. Furthermore, it employs a grid-database for obtaining collision metrics. A limit is also set on the number of agents that can simultaneously occupy each cell. The defined grid resolution combined with the density limit per cell have been proven insufficient for our experiments. The authors of Steerbench suggest that modification of either the grid resolution or the density limit per cell may decrease performance. In practice, any of the two approaches caused Steerbench to crash. For these reasons, we have excluded Steerbench from our evaluation. As a result, an agent-based evaluation of our method is missing from the present report and is left as future-work.

5.2 Scenarios

We test the ability of our method to generate plausible social group behavior against a variety of scenarios:

- **Bidirectional corridor:** The scenario involves six social groups that navigate through a horizontal corridor of length 20 meters. The corridor has a fixed width of 10 meters. Half of the groups are initialized at the left side of the environment and have group goals assigned at the right side of the environment. The other half of the groups are initialized at the right side of the environment and have group goals assigned at the left side of the environment. All members of each group are initialized at random positions within a square region of area 4 square meters. At each side, the initialization square regions are on top of each other with a separating distance of 1 meter. This guarantees that the members of each group are initialized close to each other and that the groups do not overlap during initialization. The goal area of each group is centered at a random position within a square region of area 4 square meters. The radius of the goal area is set to 0.6 meters for all groups. For each group, the initialization and destination square regions are symmetrical with respect to the vertical axis of the environment. The environment is bound by a square region of total area 1600

square meters. Figures 5.1, 5.3 and 5.7 provide simulation instances of this scenario for different group sizes. In this scenario, the challenge for each group is to remain social before and after it encounters groups that move in the opposite direction. All groups are also expected to remain coherent throughout their navigation.

- **Squeezing corridor:** The scenario involves twelve social groups that navigate through a horizontal corridor of length 50 meters. The corridor has a linearly decreasing width. The width of the corridor decreases equally at each of the corridor sides. At the left side of the environment the corridor has a width of 40 meters. At the right side of the environment the corridor has a width of 10 meters. The medial axis of the corridor is parallel to the horizontal axis of the environment and passes through the center of the environment. All groups are initialized at the left side of the environment. All members of each group are initialized at random positions within a square region of area 4 square meters. The initialization square regions are on top of each other with a separating distance of 1 meter. This guarantees that the members of each group are initialized close to each other and that the groups do not overlap during initialization. The goal area of each group is centered at a random position within a square region of area 4 square meters. The radius of the goal area is set to 0.6 meters for all groups. For each group, the initialization and destination square regions are symmetrical with respect to the vertical axis of the planar environment. The environment is bound by a square region of total area 3600 square meters. Figures 5.9, 5.11 and 5.13 provide simulation instances of this scenario for different group sizes. In this scenario, the challenge for each group is to remain coherent and social while the corridor becomes narrower and the crowd density increases. If any of these properties was lost while inside the corridor, the groups are also tested on their ability to regain coherence and sociality after exiting the corridor.

- **Corners:** The scenario involves four social groups that navigate within an square room whose interior is free of static obstacles. Each group is initialized near a different corner of the room and has its goal area assigned at the opposite corner. All members of each group are initialized at random positions within a square region of area 4 square meters. At each corner, the initialization square regions has a separating distance of 1 meter from each of the two walls that form the corner. The goal area of each group is centered at a random position within a square region of area 4 square meters. The radius of the goal area is set to 0.6 meters for all groups. For each group, the initialization and destination square regions are symmetrical with respect to the center of the environment. The environment is bound by a square region of total area 1600 square meters. Figures 5.15, 5.17 and 5.19 provide simulation instances of this scenario for different group sizes. In this scenario, the challenge for each group is to remain social before and after all groups cross each other at a center from four different directions. Groups are also expected to remain coherent throughout their navigation.

- **Building evacuation:** A building of ten rooms, one corridor and two exits is occupied by 490 groups. For each group, all members are initialized at random positions within the same room. All groups must evacuate the building through the nearest exit. One rectangle region is defined above the top exit of the building and another rectangle region is defined below the bottom exit of the building. The goal area of each group is centered at a random position within the region that corresponds to the selected exit. The radius of the goal area is set to 0.6 meters. Each of the destination regions has a width of 40 meters and a height of 20 meters. Each of the destination regions has a separating distance of 10 meters from the nearest environment boundary. The center of each destination region is horizontally aligned with the center of the corresponding exit. The environment is bound by a square region of total area 0.09 square kilometers. The dimensions of the rectangle that minimally bounds the building are 95 meters by 128 meters. Figures 5.21, 5.23 and 5.25 provide simulation instances of this scenario for different group sizes. The exact dimensions of the building elements can be either extracted from any of these figures or they can be provided per request. In this

scenario, the challenge for each group is to remain coherent throughout the evacuation. Each group is also expected to remain social when the local conditions allow for this to happen.

- **Room evacuation:** A crowd of 180 pedestrians evacuate a room through a single exit. The exact environment dimensions and set-up can be found in [51]. Figures 5.28, 5.29, 5.30, 5.31 and 5.32 provide the initial conditions for this scenario for different group sizes. This scenario examines the effect that the group size has on the evacuation time.

After initial experimentation, we have chosen parameter settings that are shared among all our experiments. Table 5.1 lists these settings. As suggested by Weidmann [88], we apply a normal distribution to determine the preferred speed of each pedestrian, with a mean of 1.34 m/s and a standard deviation of 0.26 m/s. We will compare SGN method against the model of Moussaïd et al. [65, 66]. Therefore, the rest of the parameters are assigned intuitive universal values. We follow this approach as an attempt to limit the variance in the benchmarking scores that could be attributed to the heterogeneity of the crowd. Nevertheless, our method has been defined and implemented in a way that such heterogeneity can be exhibited, when the application scenario requires it.

We run each scenario 100 times to obtain results that are statistically significant. At each run, the starting positions of all agents are (possibly) altered due to the different seed values to the employed pseudo-random number generators. However, the starting positions are still constrained to lie within the predefined initialization regions. The preferred speed of each agent remains unaltered throughout the runs, to prevent inducing further variance to our results.

Parameter	Value	Description
r_{ij}	0.24 m	radius of agent $\mathcal{A}_{ij}$
ϕ_{ij}	90 °	vision-field half angle of agent $\mathcal{A}_{ij}$
$d_{Max,ij}$	10 m	maximum viewing distance of agent $\mathcal{A}_{ij}$
$r_{p,ij}$	1 m	personal space radius of agent $\mathcal{A}_{ij}$
d_{social}	1 m	inter-member social distance threshold
mean($v_{pref,ij}$)	1.34 m/s	mean value of agent preferred speed
stdev($v_{pref,ij}$)	0.26 m/s	standard deviation of agent preferred speed
$r_{g,i}$	0.6 m	radius of goal area of group $\mathcal{G}_i$
τ	0.5 s	relaxation time
k	5000	strength of physical contact forces
b_1	1	strength of social interactions between group members
b_2	3	strength of coherence interactions between group members
t_{step}	0.1 s	simulation time-step

Table 5.1: Simulation parameters and their assigned values throughout the experiments.

5.3 Quality evaluation

We evaluate the ability of SGN to produce plausible simulations in two steps. First, we examine the quality of the generated social group behaviors (Section 5.3.1). Next, we perform a macroscopic evaluation of our method (Section 5.3.2).

5.3.1 Social-group-behavior evaluation

We evaluate the ability of our method in terms of quality of the generated social group behavior under the following scenarios: bidirectional corridor, squeezing corridor, corners and building evacuation. Similar to Köster et al. [51], we make visual observations to get a qualitative overview. We agree with Köster et al. that visual observations are insufficient

on their own. For this reason, we define quantitative metrics to get statistical measures on the quality of the generated behavior. The metrics are based on our definitions of group coherence (see Definition 4.1) and group sociality (see Definitions 4.3 and 4.3).

First, we define the *lifetime* of a group as the number of simulation steps that it takes for a group to reach its goal. In our experiments, all groups are initialized at the start of the simulation. Therefore, the lifetime of group $\mathcal{G}_i, \forall i \in [1, N_{\mathcal{G}}]$, is given by the simulation step that the last member $\mathcal{A}_{il_i}$ reaches the goal.

To examine the ability of a method to simulate coherent groups, we define the *coherence* metric. The coherence metric for a group $\mathcal{G}_i$ is the percentage of the simulation steps that $\mathcal{G}_i$ is coherent (see Definition 4.1) measured over its lifetime. To examine the ability of a method to simulate social groups, we propose two *sociality* metrics: the partial sociality metric, and the total sociality metric. The partial sociality metric for a group $\mathcal{G}_i$ is the percentage of the simulation steps that $\mathcal{G}_i$ is in a partially social configuration (see Definition 4.2) measured over its lifetime. Similarly, the total sociality metric for a group $\mathcal{G}_i$ is the percentage of the simulation steps that $\mathcal{G}_i$ is in a totally social configuration (see Definition 4.3) throughout its lifetime.

We employ these metrics throughout the experiments and run the aforementioned scenarios for both our proposed method (SGN) and the method of Moussaïd et al. [66]. As a result, we are also able to compare SGN against an existing method in terms of the generated social group behavior. Throughout the experiments, our method and the method of Moussaïd et al. share the same set of seed values and parameter settings. As a result, both methods run an identical set of scenarios.

For the method of Moussaïd et al. [66] to be able to simulate the scenarios, it is necessary that we extend it so that it captures all necessary crowd-simulation phases (see Section 1.3). For this reason, we integrate the algorithm proposed in [66] in the local-movement level of our framework. We also substitute the collision-avoidance method of Moussaïd et al. 2010 [66] with that of Moussaïd et al. 2011 [65]. This substitution improves the avoidance behavior of the original algorithm by introducing anticipation and has been verified to do so through early experiments. Furthermore, for the implementation of the method of Moussaïd et al. , we enable the high-level coordination mechanism only during initialization. We follow this approach, because we want to have a working social-group model to compare against with the minimum required extensions to original the works of its authors.

We run each scenario 100 times for groups of two, three and four. At each run of a scenario, we measure the coherence, partial sociality and total sociality for every simulated group. For each run of a scenario, we compute the average of each of these measures over the entire group population. For each scenario, we perform Q-Q plots of these averages to test whether the recorded data follow normal distributions over the runs. These plots can be found in Appendix B. We perform this procedure for both our method and the method of Moussaïd et al. For a few scenario runs, it can occur that not all groups reach their goal within reasonable time. The reasonable time is determined based on the average lifetime of the rest of the runs for each scenario-method combination. Therefore, we set a time limit to the simulation of each scenario. If any of the groups has not reached its goal within the time limit, then we consider the run as failed run. Table 5.2 lists all failed runs for each scenario-method combination. We exclude all failed runs from our comparative analysis to maintain statistical significance and examine them separately. The average group coherence, partial sociality and total sociality and their corresponding standard deviation are listed in Tables 5.3, 5.4 and 5.5, respectively. No Q-Q plots have been performed for the average group coherence metric on bidirectional-corridor, squeezing-corridor, and corners scenarios, because group coherence is nearly constant throughout all experiments involving these scenarios (see Table 5.3).

Scenario	Group size	#failed runs	
		SGN	Moussaid
Bidirectional corridor	2	0	0
	3	0	0
	4	0	0
Squeezing corridor	2	0	0
	3	2	2
	4	1	5
Corners	2	0	0
	3	0	0
	4	0	0
Building evacuation	2	7	8
	3	8	16
	4	19	40

Table 5.2: Number of experiments that failed to run successfully. Results are displayed for each scenario-method combination.

Scenario	Group size	Coherence			
		Average (%)		StDev (%)	
		SGN	Moussaid et al.	SGN	Moussaid et al.
Bidirectional corridor	2	100.000	100.000	0.000	0.000
	3	100.000	100.000	0.000	0.000
	4	100.000	100.000	0.000	0.000
Squeezing corridor	2	100.000	100.000	0.000	0.000
	3	99.980	100.000	0.196	0.000
	4	100.000	100.000	0.000	0.001
Corners	2	100.000	100.000	0.000	0.000
	3	100.000	100.000	0.000	0.000
	4	100.000	100.000	0.000	0.000
Building evacuation	2	96.052	92.742	0.426	0.629
	3	90.551	84.113	0.845	0.908
	4	83.347	75.686	0.912	1.048

Table 5.3: Average group coherence (%) and corresponding standard deviation (%). Results are measured over all successful runs for each scenario-method combination.

Scenario	Group size	Partial sociality			
		Average (%)		StDev (%)	
		SGN	Moussaid et al.	SGN	Moussaid et al.
Bidirectional corridor	2	91.720	88.993	2.243	2.639
	3	77.357	51.125	5.128	13.189
	4	64.286	21.847	5.265	6.730
Squeezing corridor	2	92.383	90.353	1.040	1.228
	3	82.502	61.010	3.586	9.547
	4	72.216	20.477	4.598	5.903
Corners	2	91.286	89.595	2.523	2.717
	3	80.499	52.679	4.662	18.953
	4	70.659	19.411	4.689	7.963
Building evacuation	2	54.458	50.470	0.881	0.961
	3	31.189	15.581	0.791	0.654
	4	20.453	6.873	0.677	0.385

Table 5.4: Average group partial sociality (%) and corresponding standard deviation (%). Results are measured over all successful runs for each scenario-method combination.

Scenario	Group size	Total sociality			
		Average (%)		StDev (%)	
		SGN	Moussaid et al.	SGN	Moussaid et al.
Bidirectional corridor	2	91.720	88.993	2.243	2.639
	3	66.738	25.089	6.386	10.660
	4	41.465	0.300	7.251	1.032
Squeezing corridor	2	92.383	90.353	1.040	1.228
	3	75.497	27.213	4.228	10.116
	4	56.169	0.596	5.969	1.221
Corners	2	91.286	89.595	2.523	2.717
	3	71.590	25.366	5.538	15.169
	4	50.359	0.215	6.488	1.285
Building evacuation	2	54.458	50.470	0.881	0.955
	3	26.501	7.166	0.731	0.442
	4	12.403	0.477	0.555	0.088

Table 5.5: Average group total sociality (%) and corresponding standard deviation (%). Results are measured over all successful runs for each scenario-method combination.

Note that, each scenario resolves to three variants, depending on the group size. The variants should not be compared against each other, because the size of the agent population is linear to the group size. As a result, for each scenario variant, the methods under study are tested on their ability to produce coherent and social group behavior at a different crowd density. Nevertheless, by following this approach the size of the group population remains constant between the variants of each scenario. The motivation behind this approach is that we aim at a similar precision in the obtained metrics of all variants of each scenario. Furthermore, the partial sociality and total sociality for groups of two are equivalent by definition. Nevertheless, we have included both metrics in all scenario-variants that involve groups of two as an extra validation step of our experimental analysis.

As suggested by Tables 5.4 and 5.5, for each scenario, the variance of both the partial sociality and total sociality metrics differs between the two methods. As a result, the

assumptions for applying a Student's paired t-test analysis are not met [81]. Therefore, we employ Welch's t-test [90] to compare the methods for each scenario variant. Furthermore, we compute the confidence intervals for these tests to further investigate the difference between the methods under comparison. We follow the same strategy for the coherence metric on the building-evacuation scenario, since unequal variances are detected between the two methods (see Table 5.3). In the following paragraphs, we discuss the results obtained for every variant of each scenario.

Bidirectional-corridor scenario. In this scenario, groups walk along a corridor on two opposing directions. The challenge for each group is to remain social before and after it encounters groups that move in the opposite direction. All groups are also expected to remain coherent throughout their navigation.

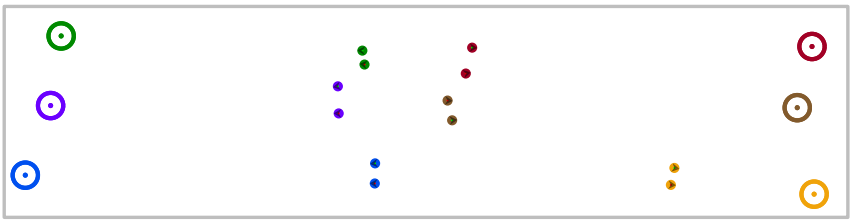

Figure 5.1: Bidirectional-corridor scenario for groups of two: simulation instance using SGN method. For each group, all members share the same color. The goal area of a group is visualized by a ring that shares the color of the group members. A small *green* (*red*) dot at an agent's position indicates that the agent currently acts as a leader (last member) of its group, respectively.

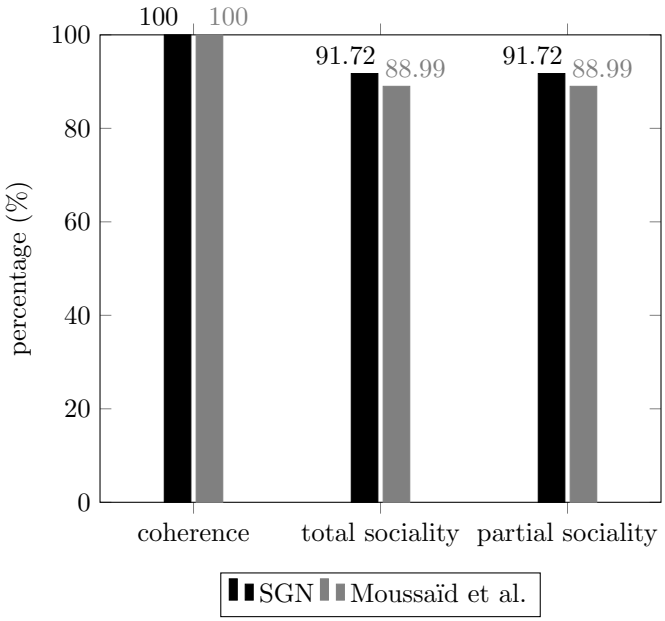

Figure 5.2: Bidirectional-corridor scenario for groups of two: bar plot of group metrics for the two methods under study. Each percentage is the mean percentage of the simulation steps that a group is in a configuration that meets the property of the horizontal axis (see Definitions 4.1, 4.2 and 4.3) measured over the lifetime of the group.

Figure 5.1 provides a simulation instance of this scenario when groups of two are considered and SGN method is used. Figure 5.2 summarizes the obtained results for this scenario variant. For both our method and the method of Moussaïd et al. , the coherence of all groups

is maintained over all runs. In this scenario variant, both methods demonstrate their ability to produce social group behavior, which is briefly sacrificed for collision-avoidance purposes. A Welch's t-test on both partial and total group sociality indicates that the two-tailed P value is less than 0.0001. Therefore, this difference is considered to be highly statistically significant. When employing our method, the test suggests a mean improvement on sociality $\approx 2.73\%$, with 2.04% - 3.41% being a 95% confidence interval of the improvement over the method of Moussaïd et al.

Figure 5.3: Bidirectional-corridor scenario for groups of three: simulation instance using SGN method. For each group, all members share the same color. The goal area of a group is visualized by a ring that shares the color of the group members. A small *green* (*red*) dot at an agent's position indicates that the agent currently acts as a leader (last member) of its group, respectively.

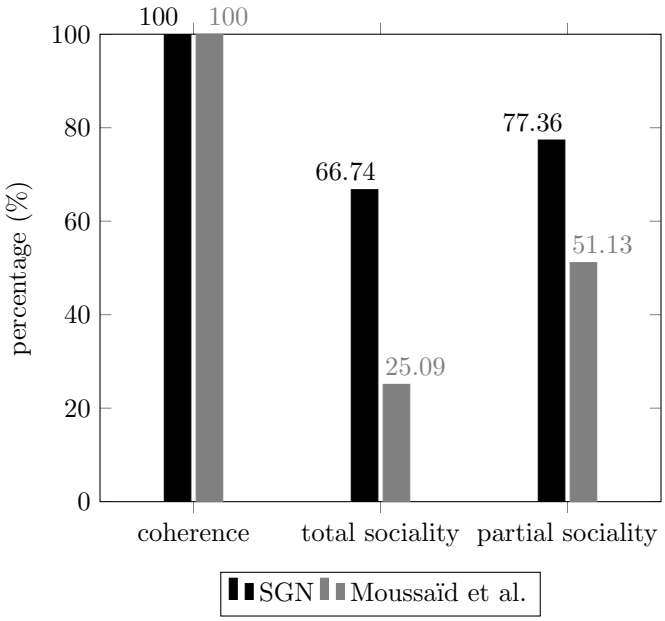

Figure 5.4: Bidirectional-corridor scenario for groups of three: bar plot of group metrics for the two methods under study. Each percentage is the mean percentage of the simulation steps that a group is in a configuration that meets the property of the horizontal axis (see Definitions 4.1, 4.2 and 4.3) measured over the lifetime of the group.

Figure 5.3 provides a simulation instance of this scenario when groups of three are considered and SGN method is used. Figure 5.4 summarizes the obtained results for this scenario variant. For both our method and the method of Moussaïd et al. , the coherence of all groups is maintained over all runs. In this scenario variant, SGN method manages to produce totally social and partially social configurations (66.74% and 77.36%, respectively). The method of Moussaïd et al. scores significantly lower on these metrics (25.09% and 51.13%, respectively).

In detail, regarding partial group sociality, a Welch's t-test indicates that the two-tailed P value is less than 0.0001. Therefore, this difference is considered to be highly statistically

significant. When employing our method, the test suggests a mean improvement on partial sociality $\approx 26.23\%$, with 23.43% - 29.03% being a 95% confidence interval of the improvement over the method of Moussaïd et al.

Regarding total group sociality, a Welch's t-test indicates that the two-tailed P value is less than 0.0001. Therefore, this difference is considered to be highly statistically significant. When employing our method, the test suggests a mean improvement on total sociality $\approx 41.65\%$, with 39.19% - 44.10% being a 95% confidence interval of the improvement over the method of Moussaïd et al.

Before examining the results for the rest of the scenarios, we provide some example images for the groups-of-three variant of the bidirectional scenario, as a visual aid to the reader. Figure 5.5 displays example group configurations that SGN generates. Figure 5.6 displays group configurations for the method of Moussaïd et al. , when the scenario is run using the same seed values. Visual inspection, suggests that, for the method of Moussaïd et al. and the displayed instances, the purple and orange groups do not adopt social configurations. On the other hand, for SGN, only the purple group sacrifices temporarily its sociality for collision-avoidance purposes. These observations do not contradict with the sociality measures we have obtained (see Figure 5.4). Nevertheless, such an evaluation approach is not efficient and will not be further developed.

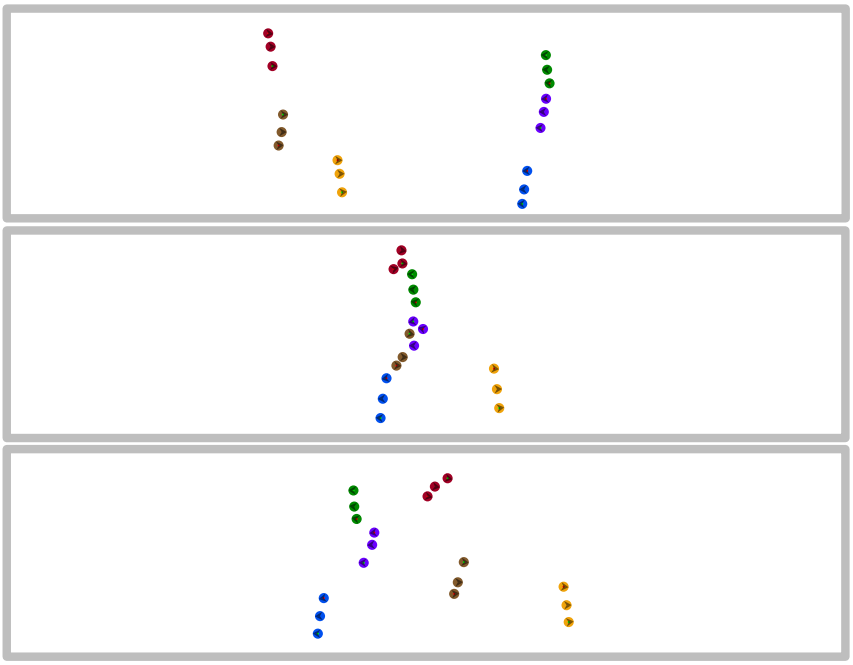

Figure 5.5: Bidirectional-corridor scenario: example configurations that SGN generates for groups of three. The examples are displayed from top to bottom in order of increasing simulation time. For each group, all members share the same color. A small *green* (*red*) dot at an agent's position indicates that the agent currently acts as a leader (last member) of its group, respectively.

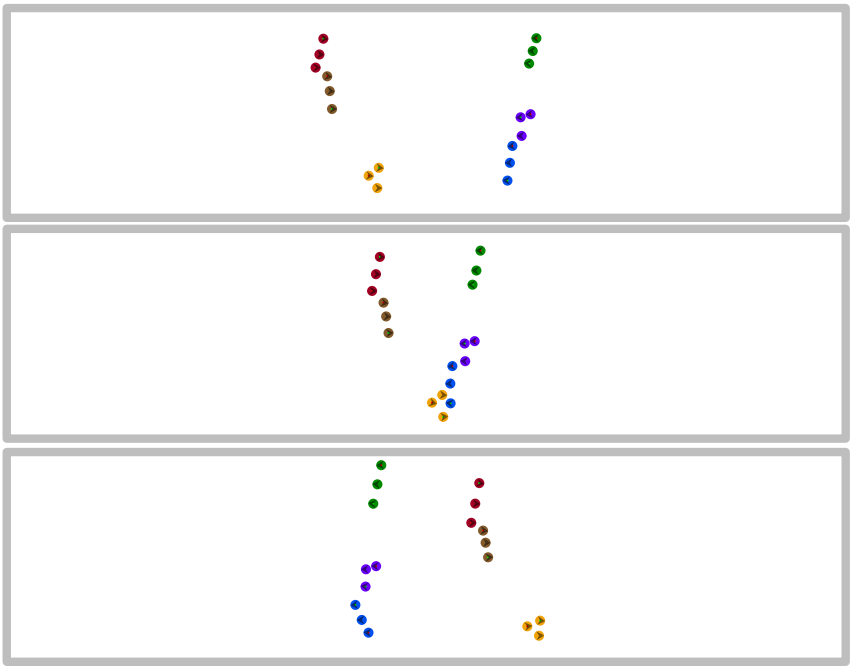

Figure 5.6: Bidirectional-corridor scenario: example configurations that Moussaïd generates for groups of three using the same seed values with Figure 5.5. The examples are displayed from top to bottom in order of increasing simulation time. For each group, all members share the same color. A small *green* (*red*) dot at an agent's position indicates that the agent currently acts as a leader (last member) of its group, respectively.

Figure 5.7: Bidirectional-corridor scenario for groups of four: simulation instance using SGN method. For each group, all members share the same color. The goal area of a group is visualized by a ring that shares the color of the group members. A small *green* (*red*) dot at an agent's position indicates that the agent currently acts as a leader (last member) of its group, respectively.

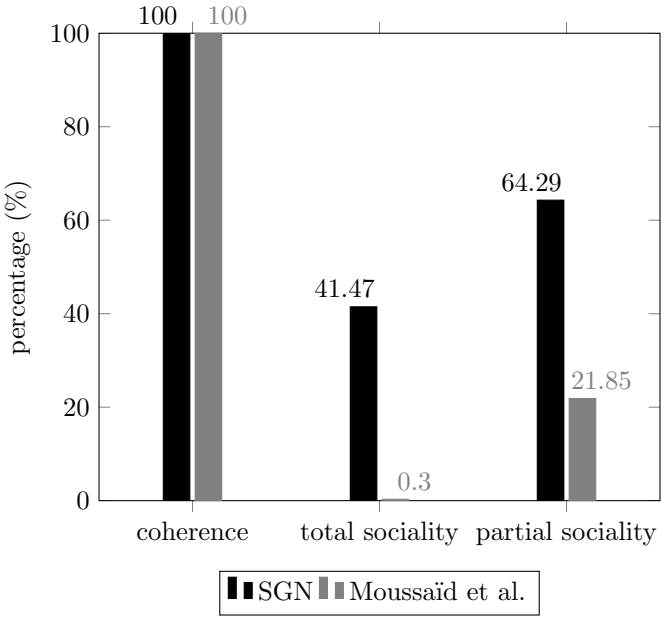

Figure 5.8: Bidirectional-corridor scenario for groups of four: bar plot of group metrics for the two methods under study. Each percentage is the mean percentage of the simulation steps that a group is in a configuration that meets the property of the horizontal axis (see Definitions 4.1, 4.2 and 4.3) measured over the lifetime of the group.

Figure 5.7 provides a simulation instance of this scenario when groups of four are considered and SGN method is used. Figure 5.8 summarizes the obtained results for this scenario variant. For both our method and the method of Moussaïd et al. , the coherence of all groups is maintained over all runs. In this scenario variant, only SGN method manages to produce totally social configurations (41.47%). On the other hand, the method of Moussaïd et al. barely exhibits this ability (0.30%). It manages, however, to produce partially social configurations (21.85%), although its ability is significantly lower when compared against our method (64.29%).

In detail, a Welch's t-test on partial sociality indicates that the two-tailed P value is less than 0.0001. Therefore, this difference is considered to be highly statistically significant. When employing our method, the test suggests a mean improvement on partial sociality $\approx 42.44\%$, with 40.75% - 44.12% being a 95% confidence interval of the improvement over the method of Moussaïd et al.

Regarding total group sociality, a Welch's t-test indicates that the two-tailed P value is less than 0.0001. By conventional criteria, this difference is considered to be highly statistically significant. When employing our method, the test suggests a mean improvement on total sociality $\approx 41.17\%$, with 39.71% - 42.62% being a 95% confidence interval of the improvement over the method of Moussaïd et al.

Squeezing corridor. In this scenario, groups walk along a corridor of decreasing width. The challenge for each group is to remain coherent and social while the corridor becomes narrower and the crowd density increases. If any of these properties was lost while inside the corridor, the groups are also tested on their ability to regain coherence and sociality after exiting the corridor.

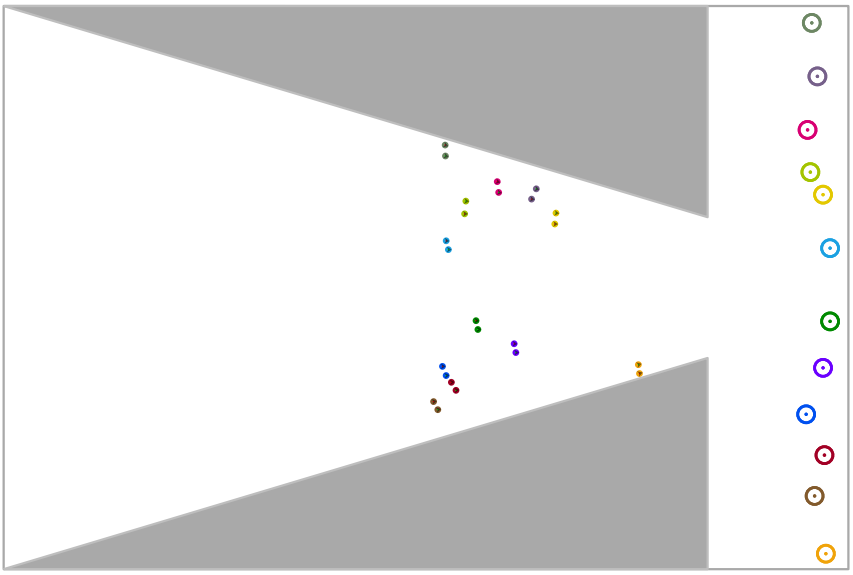

Figure 5.9: Squeezing-corridor scenario for groups of two: simulation instance using SGN method. For each group, all members share the same color. The goal area of a group is visualized by a ring that shares the color of the group members. A small *green* (*red*) dot at an agent's position indicates that the agent currently acts as a leader (last member) of its group, respectively.

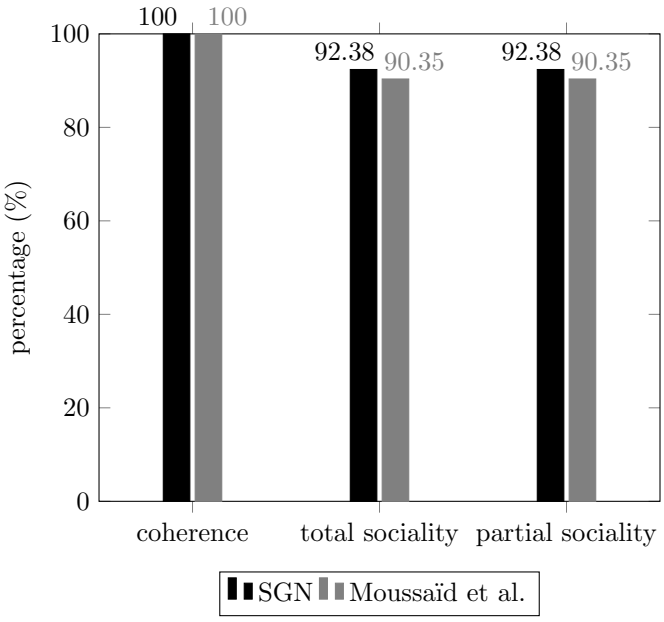

Figure 5.10: Squeezing-corridor scenario for groups of two: bar plot of group metrics for the two methods under study. Each percentage is the mean percentage of the simulation steps that a group is in a configuration that meets the property of the horizontal axis (see Definitions 4.1, 4.2 and 4.3) measured over the lifetime of the group.

Figure 5.9 provides a simulation instance of this scenario when groups of two are considered and SGN method is used. Figure 5.10 summarizes the obtained results for this scenario variant. For both our method and the method of Moussaïd et al. , the coherence of all groups is maintained over all runs. A Welch's t-test on both partial and total group sociality indicates that the two-tailed P value is less than 0.0001. Therefore, this difference is considered

to be highly statistically significant. When employing our method, the test suggests a mean improvement on sociality $\approx 2.03\%$, with 1.71% - 2.35% being a 95% confidence interval of the improvement over the method of Moussaïd et al.

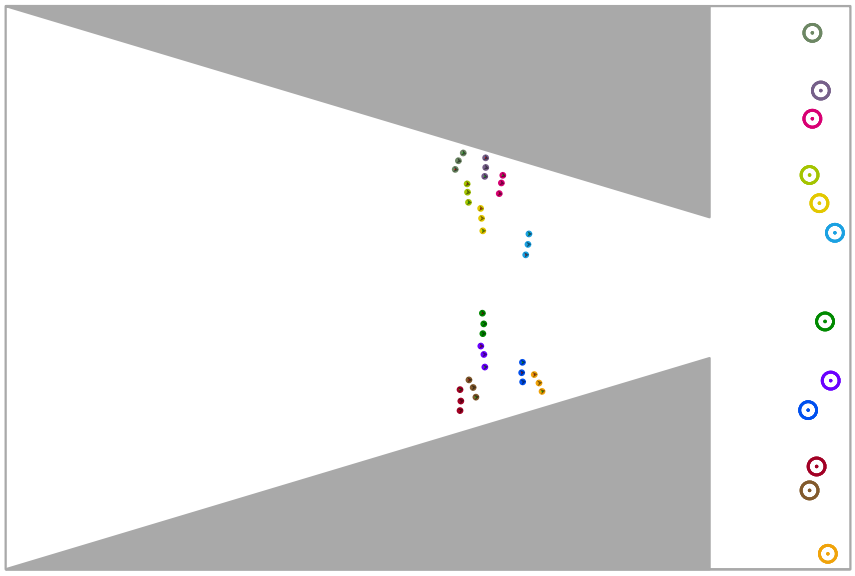

Figure 5.11: Squeezing-corridor scenario for groups of three: simulation instance using SGN method. For each group, all members share the same color. The goal area of a group is visualized by a ring that shares the color of the group members. A small *green* (*red*) dot at an agent's position indicates that the agent currently acts as a leader (last member) of its group, respectively.

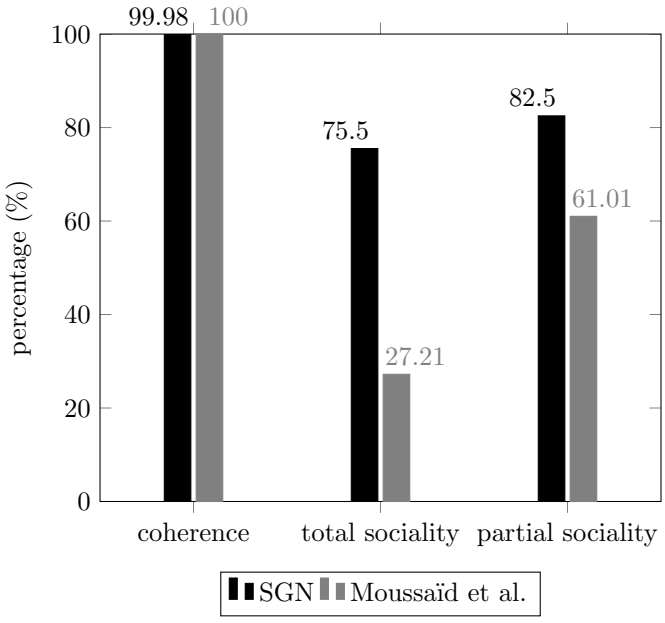

Figure 5.12: Squeezing-corridor scenario for groups of three: bar plot of group metrics for the two methods under study. Each percentage is the mean percentage of the simulation steps that a group is in a configuration that meets the property of the horizontal axis (see Definitions 4.1, 4.2 and 4.3) measured over the lifetime of the group.

Figure 5.11 provides a simulation instance of this scenario when groups of three are

considered and SGN method is used. Figure 5.12 summarizes the obtained results for this scenario variant. For the method of Moussaïd et al. , the coherence of all groups is maintained over all runs (100%). For our method, there is a single run for which coherence is briefly lost for a single group. This results in an average coherence of 99.98%. In this scenario variant, SGN method manages to produce totally social and partially social configurations (75.50% and 82.5%, respectively). The method of Moussaïd et al. scores significantly lower on these metrics (27.21% and 61.01%, respectively).

In detail, a Welch's t-test on coherence indicates that the two-tailed P value is 0.3198. Therefore, this difference is considered to be non statistically significant.

Regarding partial group sociality, a Welch's t-test indicates that the two-tailed P value is less than 0.0001. Therefore, this difference is considered to be highly statistically significant. When employing our method, the test suggests a mean improvement on partial sociality $\approx$ 21.49%, with 19.45% - 23.53% being a 95% confidence interval of the improvement over the method of Moussaïd et al.

Regarding total group sociality, a Welch's t-test indicates that the two-tailed P value is less than 0.0001. Therefore, this difference is considered to be highly statistically significant. When employing our method, the test suggests a mean improvement on total sociality $\approx$ 48.28%, with 46.09% - 50.48% being a 95% confidence interval of the improvement over the method of Moussaïd et al.

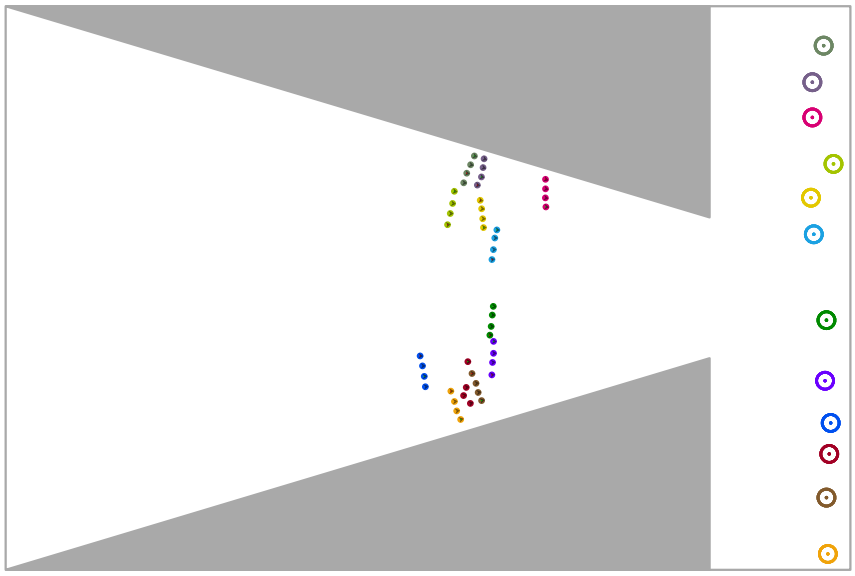

Figure 5.13: Squeezing-corridor scenario for groups of four: simulation instance using SGN method. For each group, all members share the same color. The goal area of a group is visualized by a ring that shares the color of the group members. A small *green* (*red*) dot at an agent's position indicates that the agent currently acts as a leader (last member) of its group, respectively.

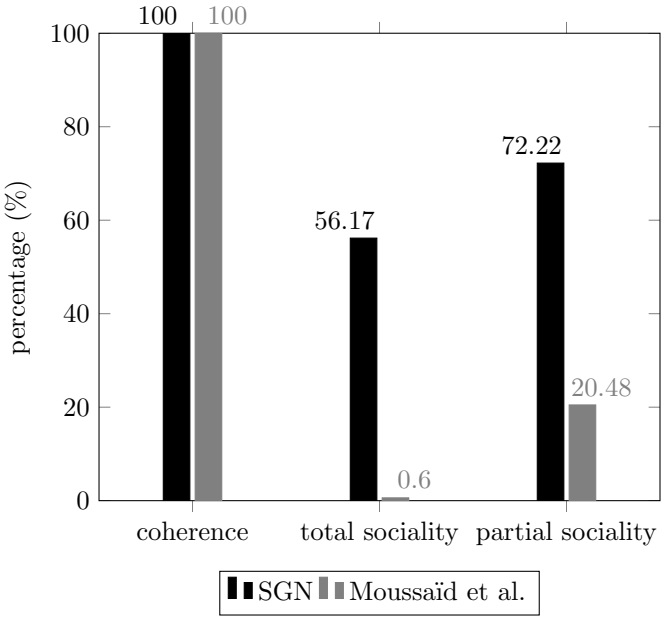

Figure 5.14: Squeezing-corridor scenario for groups of four: bar plot of group metrics for the two methods under study. Each percentage is the mean percentage of the simulation steps that a group is in a configuration that meets the property of the horizontal axis (see Definitions 4.1, 4.2 and 4.3) measured over the lifetime of the group.

Figure 5.13 provides a simulation instance of this scenario when groups of four are considered and SGN method is used. Figure 5.14 summarizes the obtained results for this scenario variant. For both our method and the method of Moussaïd et al. , the coherence of all groups is maintained over all runs. In this scenario variant, only SGN method manages to produce totally social configurations (56.17%). On the other hand, the method of Moussaïd et al. barely exhibits this ability (0.60%). It manages, however, to produce partially social configurations (20.48%), although its ability is significantly lower when compared against our method (72.22%).

In detail, a Welch's t-test on partial sociality indicates that the two-tailed P value is less than 0.0001. Therefore, this difference is considered to be highly statistically significant. When employing our method, the test suggests a mean improvement on partial sociality $\approx 51.74\%$, with 50.24% - 53.24% being a 95% confidence interval of the improvement over the method of Moussaïd et al.

Regarding total group sociality, a Welch's t-test indicates that the two-tailed P value is less than 0.0001. Therefore, this difference is considered to be highly statistically significant. When employing our method, the test suggests a mean improvement on total sociality $\approx 55.57\%$, with 54.36% - 56.79% being a 95% confidence interval of the improvement over the method of Moussaïd et al.

Corners scenario. In this scenario, groups move to the opposite corner of an empty room. The challenge for each group is to remain social before and after all groups cross each other at a center from four different directions. They are also expected to remain coherent throughout their navigation.

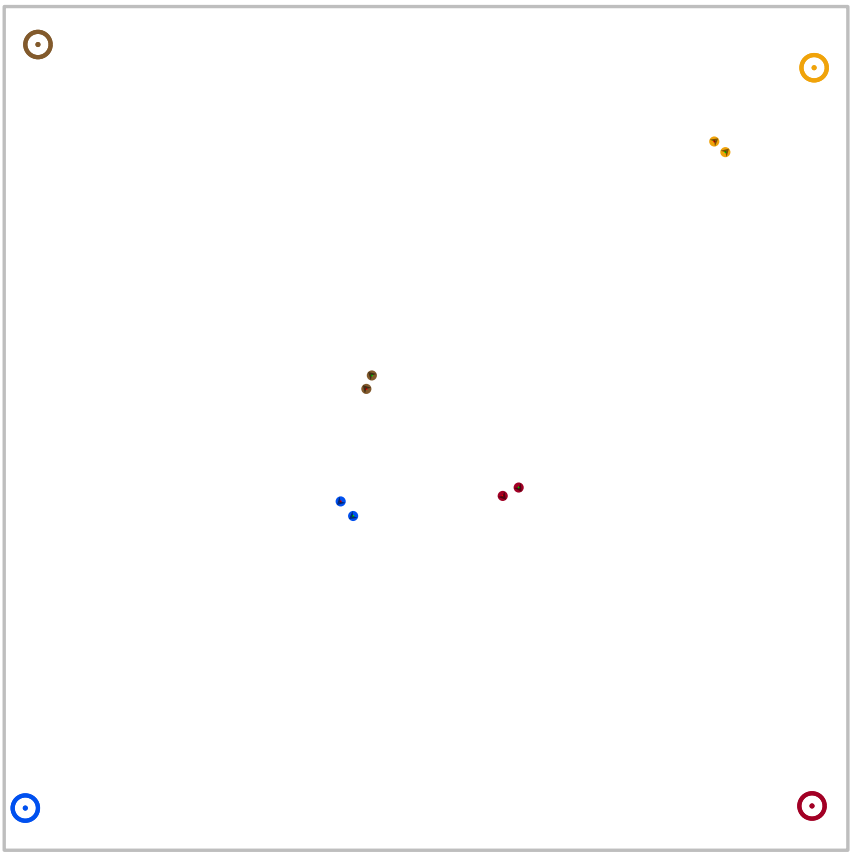

Figure 5.15: Corners-scenario for groups of two: simulation instance using SGN method. For each group, all members share the same color. The goal area of a group is visualized by a ring that shares the color of the group members. A small *green* (*red*) dot at an agent's position indicates that the agent currently acts as a leader (last member) of its group, respectively.

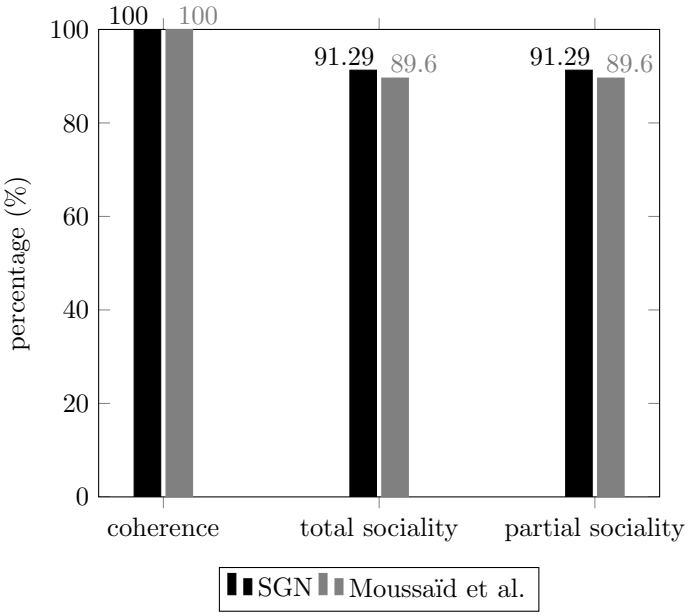

Figure 5.16: Corners scenario for groups of two: bar plot of group metrics for the two methods under study. Each percentage is the mean percentage of the simulation steps that a group is in a configuration that meets the property of the horizontal axis (see Definitions 4.1, 4.2 and 4.3) measured over the lifetime of the group.

Figure 5.15 provides a simulation instance of this scenario when groups of two are considered and SGN method is used. Figure 5.16 summarizes the obtained results for this scenario variant. For both our method and the method of Moussaïd et al. , the coherence of all groups is maintained over all runs. In this scenario variant, both methods demonstrate their ability to produce social group behavior, which is briefly sacrificed for collision-avoidance purposes. A Welch's t-test on both partial and total group sociality indicates that the two-tailed P value is less than 0.0001. Therefore, this difference is considered to be highly statistically significant. When employing our method, the test suggests a mean improvement on sociality $\approx 1.69\%$, with 0.96% - 2.42% being a 95% confidence interval of the improvement over the method of Moussaïd et al.

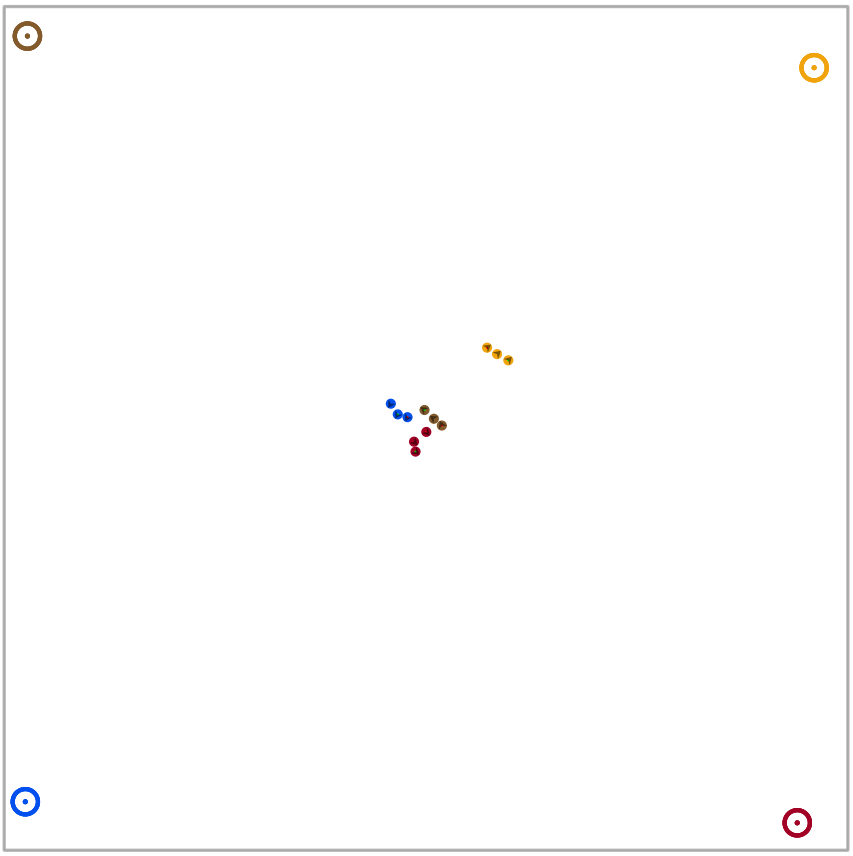

Figure 5.17: Corners scenario for groups of three: simulation instance using SGN method. For each group, all members share the same color. The goal area of a group is visualized by a ring that shares the color of the group members. A small *green* (*red*) dot at an agent's position indicates that the agent currently acts as a leader (last member) of its group, respectively.

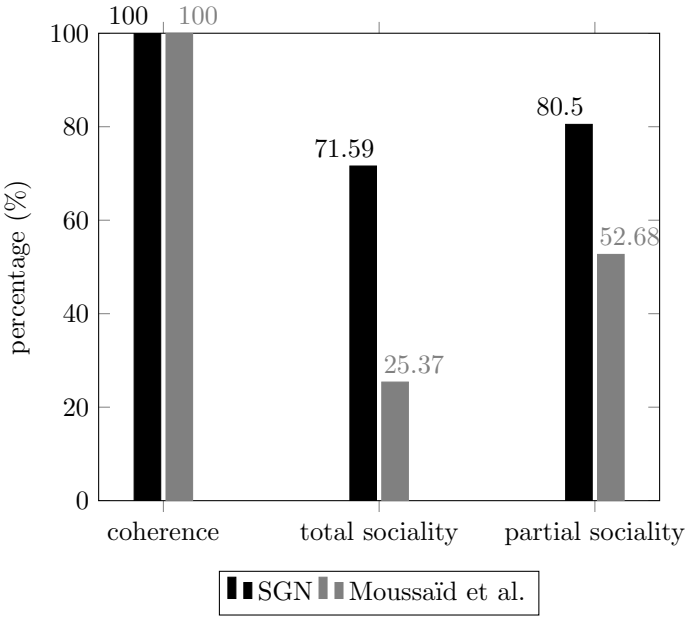

Figure 5.18: Corners scenario for groups of three: bar plot of group metrics for the two methods under study. Each percentage is the mean percentage of the simulation steps that a group is in a configuration that meets the property of the horizontal axis (see Definitions 4.1, 4.2 and 4.3) measured over the lifetime of the group.

Figure 5.17 provides a simulation instance of this scenario when groups of three are considered and SGN method is used. Figure 5.18 summarizes the obtained results for this scenario variant. For both our method and the method of Moussaïd et al. , the coherence of all groups is maintained over all runs. In this scenario variant, SGN method manages to produce totally social and partially social configurations (71.59% and 80.50%, respectively). The method of Moussaïd et al. scores significantly lower on these metrics (25.37% and 52.68%, respectively).

In detail, regarding partial group sociality, a Welch's t-test indicates that the two-tailed P value is less than 0.0001. Therefore, this difference is considered to be highly statistically significant. When employing our method, the test suggests a mean improvement on partial sociality $\approx 27.82\%$, with 23.93% - 31.71% being a 95% confidence interval of the improvement over the method of Moussaïd et al.

Regarding total group sociality, a Welch's t-test indicates that the two-tailed P value is less than 0.0001. Therefore, this difference is considered to be highly statistically significant. When employing our method, the test suggests a mean improvement on total sociality $\approx 46.22\%$, with 43.01% - 49.44% being a 95% confidence interval of the improvement over the method of Moussaïd et al.

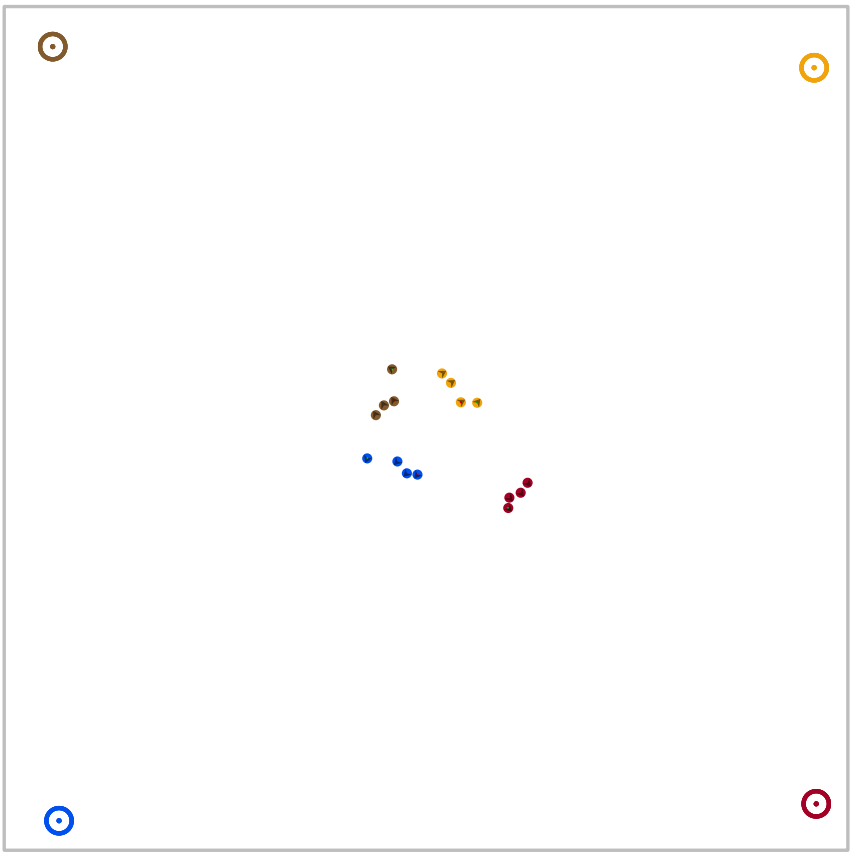

Figure 5.19: Corners scenario for groups of four: simulation instance using SGN method. For each group, all members share the same color. The goal area of a group is visualized by a ring that shares the color of the group members. A small *green* (*red*) dot at an agent's position indicates that the agent currently acts as a leader (last member) of its group, respectively.

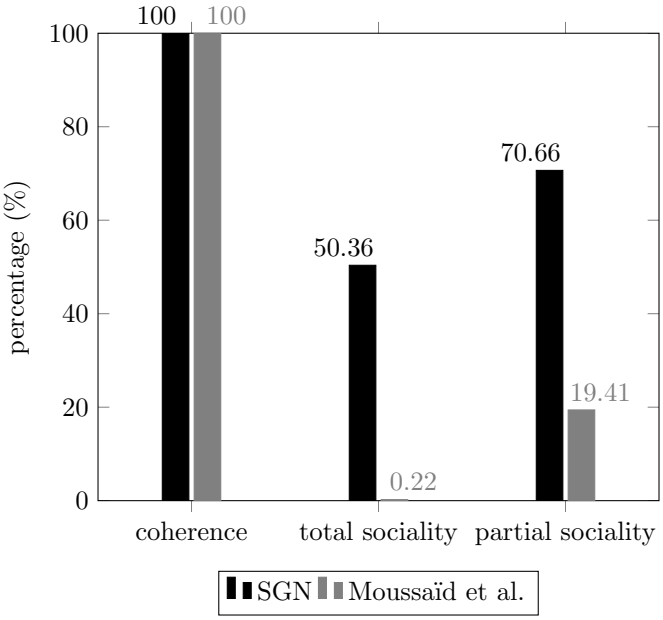

Figure 5.20: Corners scenario for groups of four: bar plot of group metrics for the two methods under study. Each percentage is the mean percentage of the simulation steps that a group is in a configuration that meets the property of the horizontal axis (see Definitions 4.1, 4.2 and 4.3) measured over the lifetime of the group.

Figure 5.19 provides a simulation instance of this scenario when groups of four are considered and SGN method is used. Figure 5.20 summarizes the obtained results for this scenario variant. For both our method and the method of Moussaïd et al. , the coherence of all groups is maintained over all runs. In this scenario variant, only SGN method manages to produce totally social configurations (50.36%). On the other hand, the method of Moussaïd et al. barely exhibits this ability (0.22%). It manages, however, to produce partially social configurations (19.41%), although its ability is significantly lower when compared against our method (70.66%).

In detail, a Welch's t-test on partial sociality indicates that the two-tailed P value is less than 0.0001. Therefore, this difference is considered to be highly statistically significant. When employing our method, the test suggests a mean improvement on partial sociality $\approx 51.25\%$, with 49.42% - 53.07% being a 95% confidence interval of the improvement over the method of Moussaïd et al.

Regarding total group sociality, a Welch's t-test indicates that the two-tailed P value is less than 0.0001. Therefore, this difference is considered to be highly statistically significant. When employing our method, the test suggests a mean improvement on total sociality $\approx 50.14\%$, with 48.83% - 51.46% being a 95% confidence interval of the improvement over the method of Moussaïd et al.

Building-evacuation scenario. In this scenario, all groups must evacuate a building of several rooms and a wide corridor with two exits. The members of each group are initialized on the same room. For each group, all members must first gather at their leader. Then, they must exit the room as a group and evacuate the building using the nearest exit. The challenge for each group is to remain coherent throughout the evacuation. Each group is also expected to remain social when the local conditions allow for this to happen. However, we do not expect groups to be coherent throughout their lifetime, due to the initial coordination phase.

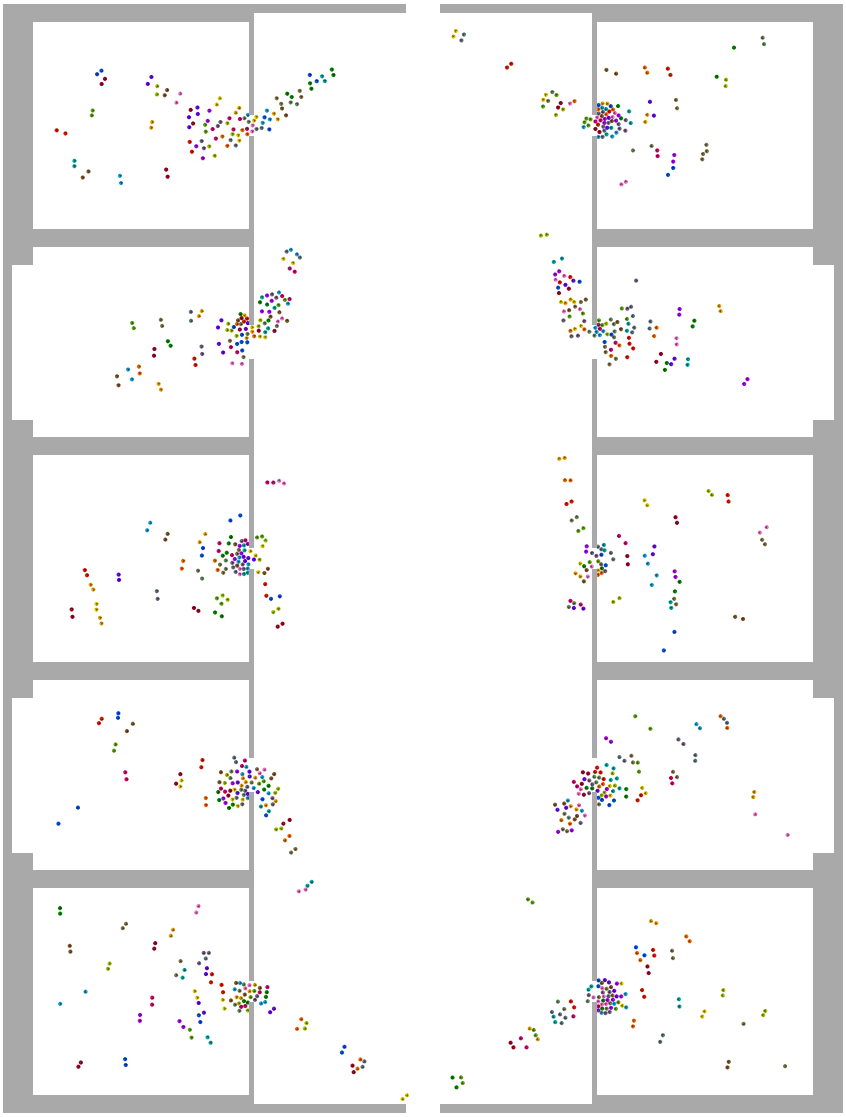

Figure 5.21: Building-evacuation scenario for groups of two: simulation instance using SGN method. For each group, all members share the same color. A 20-color palette has been used for visualizing distinguishable colors. Therefore, the same color has been reused on different groups. A small *green* (*red*) dot at an agent's position indicates that the agent currently acts as a leader (last member) of its group, respectively. For visualization reasons, we have cropped the empty region outside the building, where the goal areas reside.

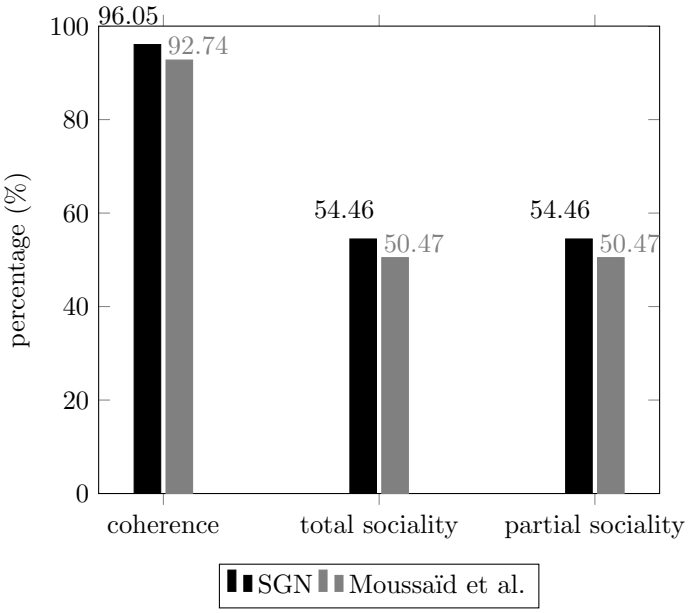

Figure 5.22: Building-evacuation scenario for groups of two: bar plot of group metrics for the two methods under study. Each percentage is the mean percentage of the simulation steps that a group is in a configuration that meets the property of the horizontal axis (see Definitions 4.1, 4.2 and 4.3) measured over the lifetime of the group.

Figure 5.21 provides a simulation instance of this scenario when groups of two are considered and SGN method is used. Figure 5.22 summarizes the obtained results for this scenario variant. For both our method and the method of Moussaïd et al. , the coherence of groups is kept at high levels (96.05% and 92.74% respectively). Visual observation indicates that bottlenecks are formed near the room and corridor exits. Furthermore, due to these bottlenecks, the groups in both methods frequently sacrifice their sociality in order to progress through the crowd.

In detail, regarding group coherence, a Welch's t-test indicates that the two-tailed P value is less than 0.0001. Therefore, this difference is considered to be highly statistically significant. When employing our method, the test suggests a mean improvement on coherence $\approx 3.31\%$, with 3.15% - 3.47% being a 95% confidence interval of the improvement over the method of Moussaïd et al.

A Welch's t-test on both partial and total group sociality indicates that the two-tailed P value is less than 0.0001. Therefore, this difference is considered to be highly statistically significant. When employing our method, the test suggests a mean improvement on sociality $\approx 3.99\%$, with 3.72% - 4.26% being a 95% confidence interval of the improvement over the method of Moussaïd et al.

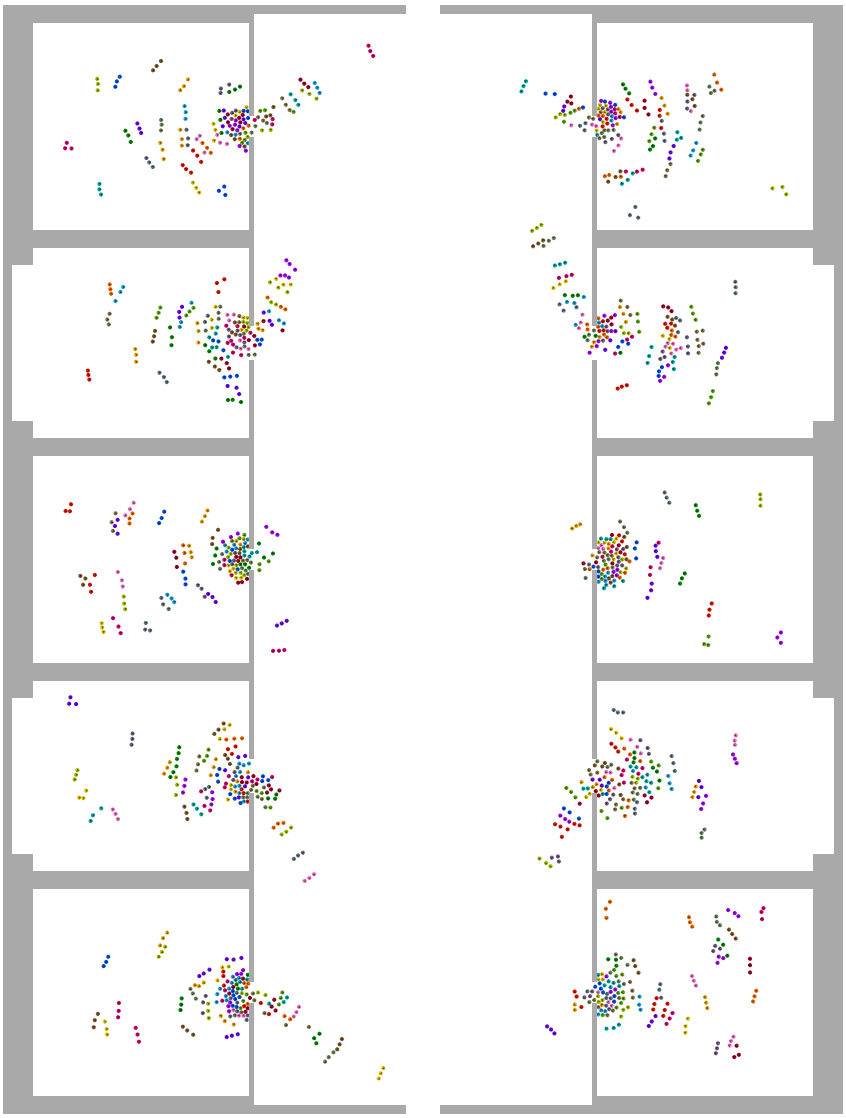

Figure 5.23: Building-evacuation scenario for groups of three: simulation instance using SGN method. For each group, all members share the same color. The same color has been reused on different groups. A small *green* (*red*) dot at an agent's position indicates that the agent currently acts as a leader (last member) of its group, respectively. For visualization reasons, we have cropped the empty region outside the building, where the goal areas reside.

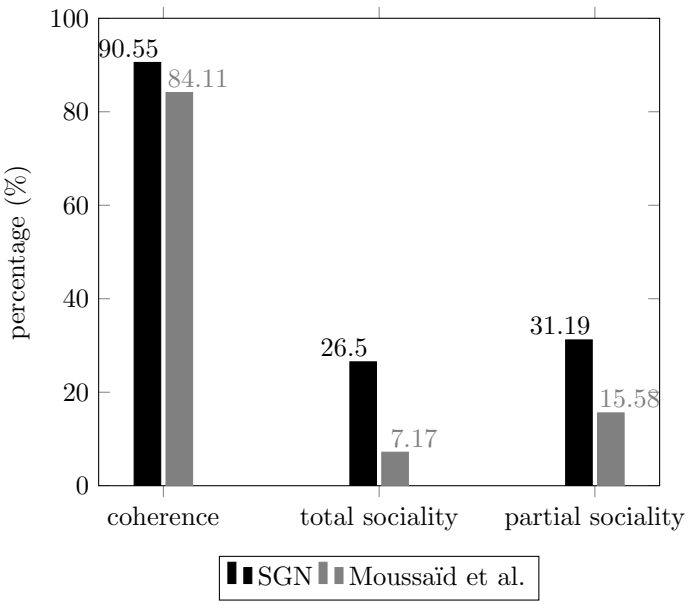

Figure 5.24: Building-evacuation scenario for groups of three: bar plot of group metrics for the two methods under study. Each percentage is the mean percentage of the simulation steps that a group is in a configuration that meets the property of the horizontal axis (see Definitions 4.1, 4.2 and 4.3) measured over the lifetime of the group.

Figure 5.23 provides a simulation instance of this scenario when groups of three are considered and SGN method is used. Figure 5.24 summarizes the obtained results for this scenario variant. Bottlenecks are formed near the room and corridor exits. In this scenario variant, SGN method manages to produce totally social and partially social configurations (26.50% and 31.19%, respectively). The method of Moussaïd et al. scores significantly lower on these metrics (7.17% and 15.58%, respectively).

In detail, a Welch's t-test on group coherence indicates that the two-tailed P value is less than 0.0001. Therefore, this difference is considered to be highly statistically significant. When employing our method, the test suggests a mean improvement on coherence $\approx 6.44\%$, with 6.18% - 6.70% being a 95% confidence interval of the improvement over the method of Moussaïd et al.

Regarding partial group sociality, a Welch's t-test indicates that the two-tailed P value is less than 0.0001. Therefore, this difference is considered to be highly statistically significant. When employing our method, the test suggests a mean improvement on partial sociality $\approx 15.61\%$, with 15.39% - 15.82% being a 95% confidence interval of the improvement over the method of Moussaïd et al.

Regarding total group sociality, a Welch's t-test indicates that the two-tailed P value is less than 0.0001. Therefore, this difference is considered to be highly statistically significant. When employing our method, the test suggests a mean improvement on total sociality $\approx 19.34\%$, with 19.16% - 19.51% being a 95% confidence interval of the improvement over the method of Moussaïd et al.

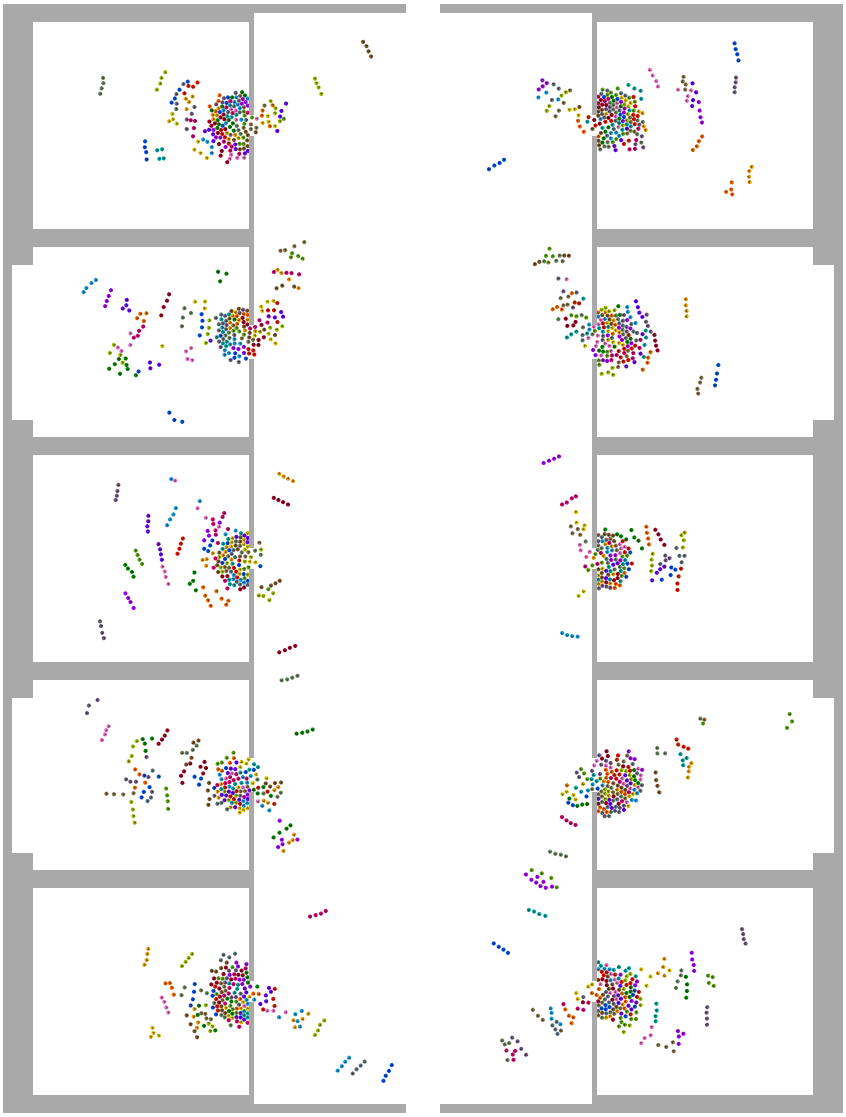

Figure 5.25: Building-evacuation scenario for groups of four: simulation instance using SGN method. For each group, all members share the same color. The same color has been reused on different groups. A small *green* (*red*) dot at an agent's position indicates that the agent currently acts as a leader (last member) of its group, respectively. For visualization reasons, we have cropped the empty region outside the building, where the goal areas reside.

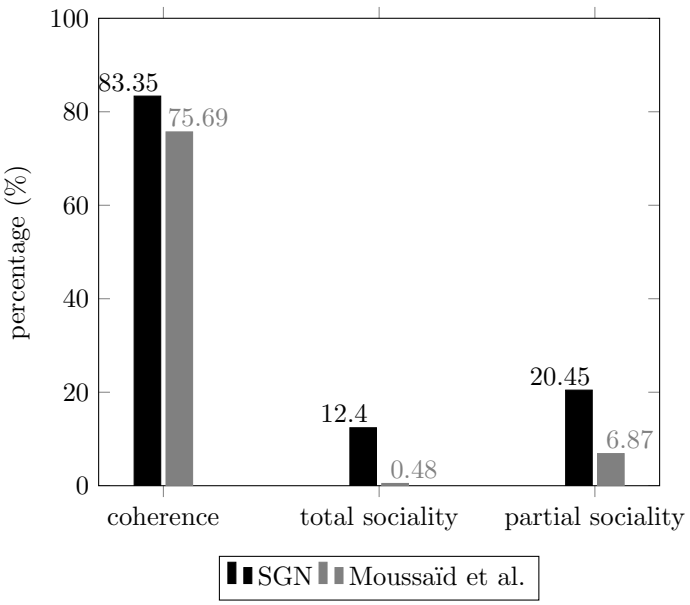

Figure 5.26: Building-evacuation scenario for groups of four: bar plot of group metrics for the two methods under study. Each percentage is the mean percentage of the simulation steps that a group is in a configuration that meets the property of the horizontal axis (see Definitions 4.1, 4.2 and 4.3) measured over the lifetime of the group.

Figure 5.25 provides a simulation instance of this scenario when groups of four are considered and SGN method is used. Figure 5.26 summarizes the obtained results for this scenario variant. In this scenario variant, only SGN method manages to produce totally social configurations (12.40%). On the other hand, the method of Moussaïd et al. barely exhibits this ability (0.48%). It manages, however, to produce partially social configurations (6.87%), although its ability is significantly lower when compared against our method (20.45%).

In detail, a Welch's t-test on coherence indicates that the two-tailed P value is less than 0.0001. Therefore, this difference is considered to be highly statistically significant. When employing our method, the test suggests a mean improvement on coherence $\approx 7.66\%$, with 7.33% - 8.00% being a 95% confidence interval of the improvement over the method of Moussaïd et al.

Regarding partial sociality, a Welch's t-test indicates that the two-tailed P value is less than 0.0001. Therefore, this difference is considered to be highly statistically significant. When employing our method, the test suggests a mean improvement on partial sociality $\approx 13.58\%$, with 13.40% - 13.76% being a 95% confidence interval of the improvement over the method of Moussaïd et al.

Regarding total group sociality, a Welch's t-test indicates that the two-tailed P value is less than 0.0001. Therefore, this difference is considered to be highly statistically significant. When employing our method, the test suggests a mean improvement on total sociality $\approx 11.93\%$, with 11.80% - 12.05% being a 95% confidence interval of the improvement over the method of Moussaïd et al.

Heterogeneous crowds. To demonstrate the ability of our method to simulate heterogeneous crowds, we repeat the building evacuation scenario under different settings. First, we perform a discrete randomization of the size of each group within the interval $[1, 4]$ with a step size of 1 agent. For a group-size value of 1, we treat the group as an individual agent and apply the method details described in Section 4.6. We also perform a discrete randomization of the radius of each agent within the interval [0.20, 0.25] with a step size of 0.01 meters. The rest of our settings remain unchanged and can be found in Table 5.1. Similar to the rest of the experiments, the preferred speed of the agents follows a normal distribu-

tion to allow for further crowd heterogeneity. Under these settings, Figure 5.27 provides a simulation instance obtained using SGN method.

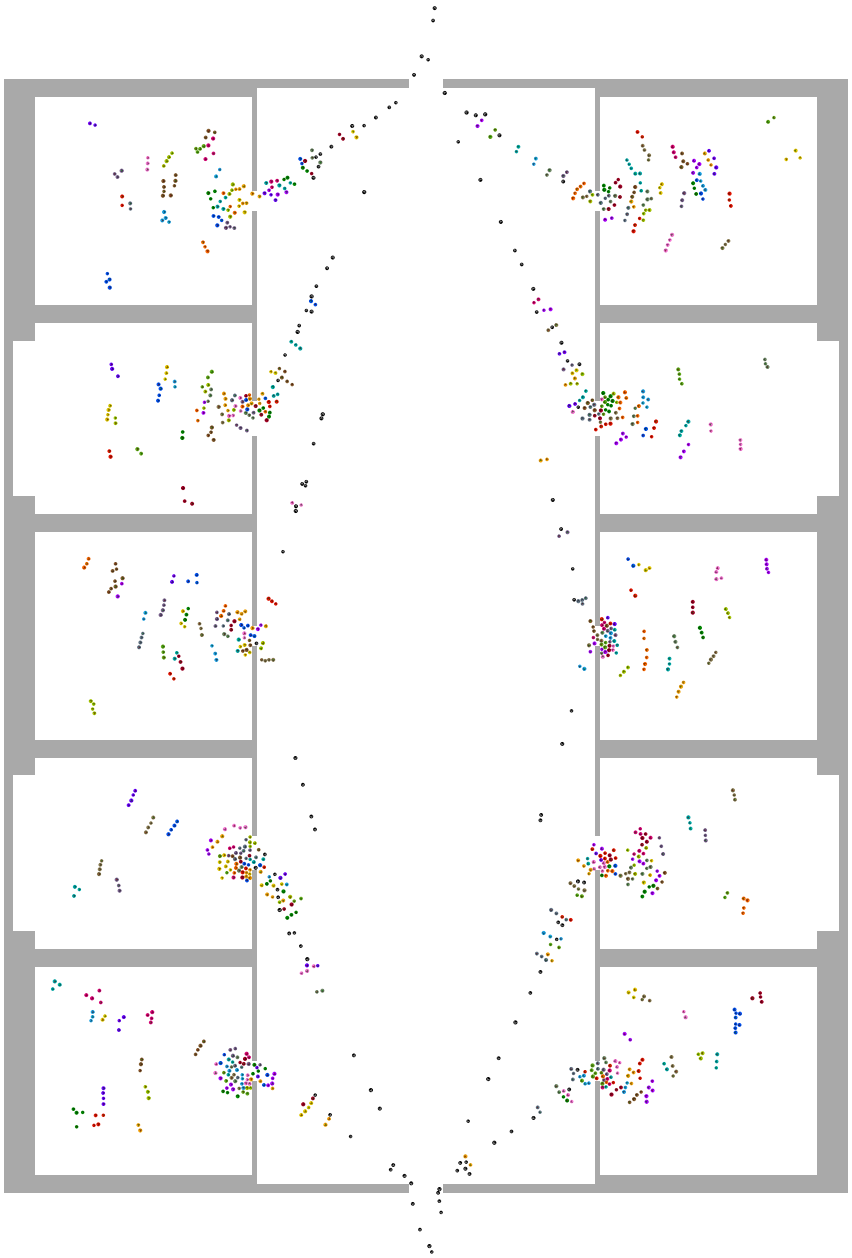

Figure 5.27: Building-evacuation scenario for mixed group sizes and agent radii: simulation instance using SGN method. For each group, all members share the same color. The same color has been reused on different groups. A small *green* (*red*) dot at an agent's position indicates that the agent currently acts as a leader (last member) of its group, respectively. Single individuals are displayed in *black*.

If a crowd involves groups of different sizes, then the average coherence, total-sociality and partial-sociality metrics are all highly sensitive to the distribution of the group sizes. Also, regarding individual agents, these metrics are not applicable. For these reasons, we have chosen not to collect group-evaluation metrics for this senario variant. Nevertheless, based on visual inspection (see Figure 5.27), we observe that our method manages to produce social group behavior in heterogeneous crowds. For this scenario variant, we also observe that individual agents tend to evacuate the building faster than groups. This phenomenon

is possibly attributed to two reasons: 1) individual agents do not have to coordinate at the start of the simulation, and 2) individual agents do not slow down throughout their navigation for coherence and sociality reasons.

Failed runs. We have reproduced all experiment runs that lead to deadlocks. For both SGN method and the method of Moussaïd et al. , it has been detected that the deadlocks occur due to limitations of the employed route-following algorithm [41]. In detail, in all failed runs, at least one agent reaches a configuration that does not allow for an *attraction point* to be computed (see [41] for details). This results in an agent with no preferred velocity. Therefore, both methods fail to steer such agents towards their group goal. As a result, the corresponding groups fail to reach their goals and the simulation reaches a deadlock. This problem is exaggerated in densely populated scenarios. According to our observations, in dense scenarios, it is more likely for an agent to lose visibility of its global route, due to the agent being carried away by the crowd. Therefore, MIRAN algorithm can sometimes pose a limit in the ability to successfully run a scenario, for both SGN method and the method of Moussaïd et al. Nevertheless, MIRAN's use of *reference points* (see [41] for details) allows for an efficient computation of the leader and last member of each group.

5.3.2 Macroscopic evaluation

We run a room-evacuation scenario to test the ability of our method to generate pedestrian dynamics that meet empirical data. We perform the simulation experiment described in Köster et al. [51], which is inspired by the real-life controlled experiment of Liddle et al. [59]. The experiment measures the time that is required for 180 pedestrians to evacuate a room through a narrow corridor. We repeat the experiment for different group sizes. For the exact set-up of this experiment, the reader is referred to the original works [51, 59]. Note that, for the implementation of both SGN method and the method of Moussaïd et al. , we set the radius of all agents to 0.20 m. The rest of the settings remain unaltered and can be found in Table 5.1. Figures 5.28, 5.29, 5.30, 5.31 and 5.32 display the initial conditions of this scenario for all group-size variants.

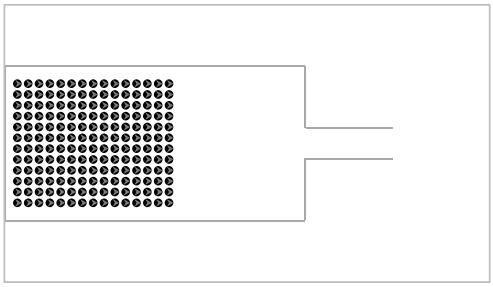

Figure 5.28: Room-evacuation scenario for individuals: initial conditions.

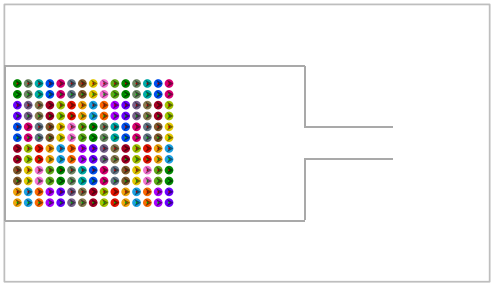

Figure 5.29: Room-evacuation scenario for groups of two: initial conditions. For each group, all members share the same color. The same color has been reused on different groups. A small *green* (*red*) dot at an agent's position indicates that the agent has been initialized as a leader (last member) of its group, respectively.

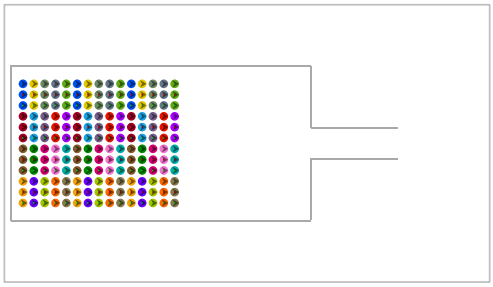

Figure 5.30: Room-evacuation scenario for groups of three: initial conditions. For each group, all members share the same color. The same color has been reused on different groups. A small *green* (*red*) dot at an agent's position indicates that the agent has been initialized as a leader (last member) of its group, respectively.

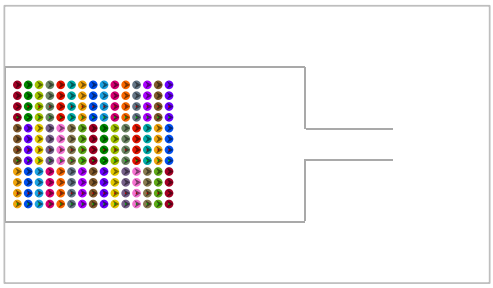

Figure 5.31: Room-evacuation scenario for groups of four: initial conditions. For each group, all members share the same color. The same color has been reused on different groups. A small *green* (*red*) dot at an agent's position indicates that the agent has been initialized as a leader (last member) of its group, respectively.

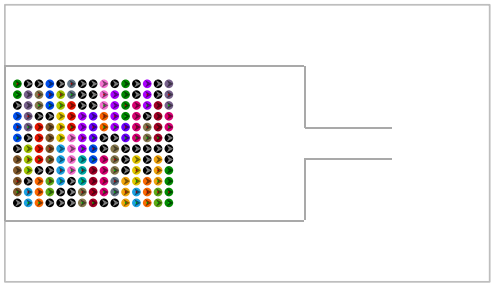

Figure 5.32: Room-evacuation scenario for mixed group sizes: initial conditions. For each group, all members share the same color. The same color has been reused on different groups. A small *green* (*red*) dot at an agent's position indicates that the agent has been initialized as a leader (last member) of its group, respectively. Single individuals are displayed in *black*.

Method	Evacuation time (s)				
	Individuals	Groups of 2	Groups of 3	Groups of 4	Mixed
SGN	73.50	98.10	112.20	115.70	99.20
Moussaïd et al.	82.10	93.50	95.40	84.30	84.10
Köster et al.	73.63	86.93	-	90.59	-

Table 5.6: Evacuation times for the room-evacuation scenario. Results for different methods are displayed. For SGN method and the method of Moussaïd et al. , all agents are modelled using a radius of 0.20 m and we perform a single run per group size. For the method of Köster et al. , we list the mean evacuation times that have been originally reported in [51] (groups of three and mixed group sizes were not included in their work).

Table 5.6 lists the evacuation times for each method. When comparing our method against the method of Köster et al. a similar trend can be observed: the evacuation time increases when the group size increases. This trend cannot be observed in the method of Moussaïd et al. Surprisingly, in the method of Moussaïd et al. , when pedestrians are arranged in groups of four, the evacuation time decreases. In the method of Moussaïd et al. , groups of four exhibit similar behavior with individuals with respect to evacuation times. This is in compliance with the empirical observations of Costa [19], who reports group-splitting behavior when a group is composed by more than three members. This behavior is neither discussed nor explicitly modelled in the method of Moussaïd et al. [66]. On the other hand, both SGN method and the method of Köster et al. indicate a negative effect of social bonds in evacuations. This effect has been previously studied and empirically verified by Song and Hu [93]. Finally, for both our method and the method of Moussaïd et al. , the mixed group-size variant of this scenario produces intermediate evacuation times.

However, we acknowledge a limitation in our experimental approach on this scenario. Due to the strict space and social constraints of this scenario, the initial configuration of all agents had to be defined in a deterministic manner (see Figures 5.28, 5.29, 5.30, and 5.31). In detail, all agents were evenly distributed within the rectangle region specified by the original experiment. The initialization resulted in an grid-like agent arrangement of 12 rows by 15 columns. This effectively allowed us to have all groups initialized in totally social configurations for every group-size variant of this scenario[2]. At the same time, our initialization procedure guarantees that all agents do not initially collide. We have followed this approach, because we consider that our initialization procedure meets space and social properties that are crucial for interpreting the results of this controlled experiment. On the other hand, the lack of variance in the initial conditions does not allow us to perform a statistical analysis similar to Köster et al. [51]. Therefore, we limit ourselves to comparing our single-run metrics with the corresponding means provided by Köster et al. An approach that would allow for inducing variance to our results would be to randomize the preferred speeds of the agents throughout the runs, while respecting the specified normal distribution (see Table 5.1). This is left for future work.

Furthermore, through visual inspection of our room-evacuation simulations, we expect a difference of our method from that of Köster et al. in terms of space utilization. In both approaches, due to high crowd density, a bottleneck occurs within the room, close to the exit. Our method is agent-based and relies on a social-force model. This suggests that competitive behavior emerges between individuals. Our agents do not follow social priority rules, because we have not modelled such behavioral aspects. This results in agents that lie on the back of the crowd to attempt to approach the exit from the sides. On the other hand, the method of Köster et al. is based on a cellular automaton that employs potential functions to steer the individuals. This is expected to lead in a rather coordinated crowd behavior. As a result, in the method of Köster et al. , the space at the sides of the exit is expected to be underutilized.

[2]For the mixed group-size variant (Figure5.32), we employ the same initial positions for the agents. For each column, we require that there exists exactly one group for each group-size under study (i.e. 2, 3 and 4) in a totally social configuration. We fill the remaining three spaces of each column with single individuals.

In a similar real-life scenario, we would expect an intermediate crowd behavior. However, empirical observations are required to validate our claim. Note that, in the original real-life controlled experiment of Liddle et al. [59], density maps are provided. These density maps indicate underutilization of the discussed side areas. Nevertheless, Liddle et al. have used 180 soldiers as their test subjects. As a result, the pedestrian population of the original experiment might be interpreted as a disciplined single large group. Therefore, it is arguable whether the metrics obtained in [59] are relevant to the crowd dynamics of pedestrians that are organized in small social groups.

In addition, we examine the effect that the agent radius has on the evacuation time. We repeat all group-size variants of the room-evacuation scenario for all agent radii between 0.20 m and 0.24 m with a step size 0.01 m. For each group size, we also include a variant of this experiment in which the radius of each agent is assigned discrete random value in the specified interval (see Figure 5.33). Table 5.7 displays the results.

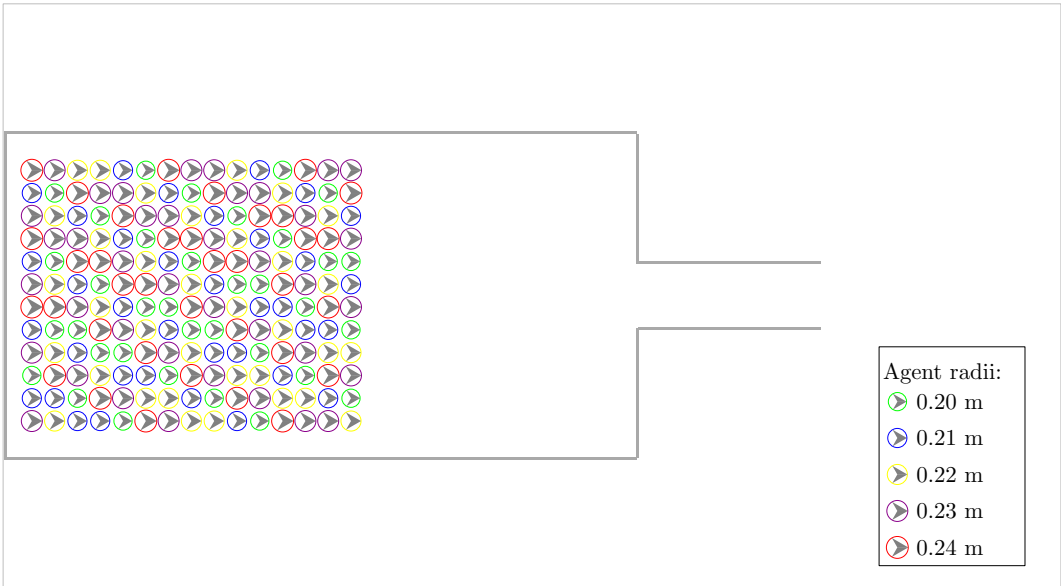

Figure 5.33: Room-evacuation scenario for mixed agent radii: initial positions and radius conditions. The outline color of the disc that models each agent indicates the radius of the agent. The same radii assignment has been adopted in all group-size variants of this scenario.

Agent radius (m)	Evacuation time (s)			
	Individuals	Groups of 2	Groups of 3	Groups of 4
0.20	73.5	98.1	112.2	115.7
0.21	82.9	103.6	120.8	129.6
0.22	86.5	113.8	135.9	135.0
0.23	93.8	293.2	145.0	151.6
0.24	98.9	337.7	170.7	251.8
Mixed	86.3	112.5	131.8	142.8

Table 5.7: Evacuation times for the room-evacuation scenario using SGN method. Results for different agent radii and group sizes are displayed.

We can observe that the radius of the agents has a significant impact on the evacuation time for all group sizes. The greater the radius, the longer it takes for all agents to evacuate the room. For each group-size variant, when agents are modelled using discs of different

radii, our method leads to intermediate evacuation times. Furthermore, regarding groups of two, for a radius of 0.23 m and above, crowd congestion is observed for several seconds. This results in an extreme increase in the evacuation times. Therefore, we conclude that the macroscopic behavior generated by our method is highly sensitive to the size of the radius that models the agents.

It should be investigated if similar differences in evacuation times can be empirically verified. For the real-life controlled experiment of of Liddle et al. [59], it would be interesting to investigate any correlations between the test-subject body sizes and the evacuation times. A repetition of the experiment that also accounts for pedestrian organization in small social groups would provide additional insight. Therefore, we acknowledge a lack of empirical data that would allow for a thorough validation of our method in terms of its generated macroscopic behaviors.

5.4 Performance

To test the performance of our method we execute additional experiments. We investigate the effect that SGN has on the overall performance of ECM crowd simulation framework. As stated in our project goals (see Section 1.2), we are mainly interested in the performance of our method in large-scale scenarios, where small social groups are expected to dominate the crowd. For this reason, we use an extension of the room-evacuation scenario for our performance measurements. Our extended virtual environment includes 11 rooms, similar to the one used in the room-evacuation scenario. Each room is initially occupied with 180 agents. This effectively gives as a crowd of 1980 agents. All agents have to evacuate their room. Four variants of the experiment are executed, depending on the arrangement of the agents as single individuals, groups of two, three or four. We run 100 times each variant of the scenario. Each run is terminated when the first agent reaches its destination. We follow this approach, to prevent artificially high performance scores that would occur if some agents remained stationary at their destination. For each variant, we measure the average simulation step runtime and its corresponding standard deviation.

		Time per step (msec)		Frame-rate (#steps/sec)	
		Average	StDev	Average	StDev
ECM Framework	Individuals	372.11	7.43	2.688	0.048
SGN	Groups of 2	374.96	1.29	2.667	0.009
	Groups of 3	377.37	1.66	2.650	0.012
	Groups of 4	378.21	1.45	2.644	0.010

Table 5.8: Performance results during *serial* execution of the scenario. The scenario involves the evacuation of 1980 agents that are evenly distributed over 11 rooms similar to Figures 5.28, 5.29, 5.30, and 5.31. Average time per step (msec) and frame-rate (#steps/sec) over 100 runs and corresponding standard deviations are displayed per group-size variant of the scenario.

In Table 5.8, we display the obtained results for a *serial* execution of our method. We observe that our method has a minimum impact on the running time of the of the framework. When social groups of two are included in the simulation using SGN, the average increase in the running time of the framework is less than 3 miliseconds per simulation step ($\approx$ 2.85 msec). The increase in the running time is maximized when social groups of four are considered ($\approx$ 6.10 msec). Nevertheless, these increases are only small, if we consider that the single-individuals variant of this scenario requires $\approx$ 372 msec per simulation step.

However, we are also interested in the applicability of our method, when thousands of agents should be simulated in real-time. Therefore, we repeat all performance measures in a parallel implementation of the framework. Agent-wise parallelization is possible in the original framework [4] for each simulation step (Levels 1, 2, and 3, see Figure 1.1). Our

method is fully compatible with this parallelization scheme, therefore can benefit from it. At the same time, we have introduced group-wise parallelization in our implementation of SGN that also encloses higher simulation levels, when necessary (Levels 4 and 5, see Figure 1.1).

		Time per step (msec)		Frame-rate (#steps/sec)	
		Average	StDev	Average	StDev
ECM Framework	Individuals	85.48	1.51	11.702	0.198
SGN	Groups of 2	92.65	1.30	10.795	0.154
	Groups of 3	92.72	1.37	10.788	0.160
	Groups of 4	91.92	1.23	10.881	0.146

Table 5.9: Performance results during *parallel* execution of the scenario. The scenario involves the evacuation of 1980 agents that are evenly distributed over 11 rooms similar to Figures 5.28, 5.29, 5.30, and 5.31. Average time per step (msec) and frame-rate (#steps/sec) over 100 runs and corresponding standard deviations are displayed per group-size variant of the scenario. For both the agent-wise and the group-wise parallel parts of our method all 4 CPU cores were employed, thus allowing a total of 8 threads for simultaneous execution.

Table 5.9 summarizes the obtained results for a *parallel* execution of our method. Again, we observe that our method has a small impact on the running time of the of the framework. When social groups of two are included in the simulation using SGN, the average increase in the running time of the framework is only ≈ 7.17 miliseconds per simulation step. Interestingly, the increase in the running time is maximized when social groups of three are considered (≈ 7.24 msec). Note that, in a parallel implementation, subdividing the crowd of 1980 agents into social triads has almost the same performance impact with subdividing it to social couples. More surprisingly, when social groups of four model the crowd, the running time reaches its minimum (≈ 91.92 msec) among all tested variants that include social group behavior. We expect this phenomenon to be triggered by performance gains on the global route-planning phase. As discussed in our method details (see Section 4), in SGN, for each social group, only the leader is assigned the computationally expensive task of finding a path to the group goal. The rest of the group members simply plan their path to the leader which is usually close-by. When all members have reached the leader, the leader forwards its path to its fellow members. This leads to performance gains over a method that simulates a crowd as single-individuals. The bigger the size of the social-group, the higher these performance gains are. On the other hand, a member of a large social-group has to consider a proportionally large number of fellow members for computing its social deceleration force (see Section 4.5.2). As a result, for a constant number of agents, an increased group size leads to more computations at the local-movement level and less computations at the global route-planning level of our framework.

Furthermore, for the parallel implementation, results indicate that our method can perform on average more than 10 simulation steps per second (see Table 5.9). We remind that the simulation step is set to 0.1s throughout all experiments and has been proven sufficient to generate social group behavior (see Section 5.3.1). Therefore, we conclude that our method can be used for real-time simulations that involve thousands of agents organized in small social-groups.

6 Discussion and future work

In this section, we conclude on our work on social-group navigation. First, we summarize the contributions of our work (Section 6.1). Next, we identify its limitations (Section 6.2). Finally, ideas for future work are presented in Section 6.3.

6.1 Contributions

We have proposed a method for simulating the navigation of social groups in crowded environments. To the best of our knowledge, SGN is the first simulation method to provide group-specific details on all relevant aspects of pedestrian navigation. Our method is also flexible. SGN is designed in a hierarchical manner and allows the developer or researcher to couple it with different navigation meshes or global route-planning algorithms.

At the same time, SGN is an agent-based method that is designed to allow for high crowd heterogeneity. Group members can differ in terms of their body size, preferred speed, and vision-field. This effectively allows our method to be considered for applications where variance in the behavior of individuals should be exhibited. Therefore, application scenarios can include crowd simulations that are based on a collection of empirical data, as well as gaming applications in which virtual characters should exhibit social human-like motion behavior.

In addition, our method avoids utilizing computationally expensive group-coordination mechanisms at the local-movement level. Alternative methods that do not follow this approach [47, 92] are expected to perform worse than SGN in terms of running time. What is more, SGN employs a vision-based collision-avoidance algorithm. This algorithm is an adjusted version of the algorithm of Moussaïd et al. 2011 [65] that improves the efficiency of the generated collision-avoidance maneuvers. It also maintains the ability of the original algorithm to produce stop-and-go waves. This makes our approach highly competitive, since it aims at high-quality crowd behaviors in the microscopic scale without introducing the additional overhead of group-velocity optimization.

Furthermore, our quality evaluation indicates significant improvements over the works of Moussaïd et al. [65, 66]. In all tested scenarios, SGN proved to generate social-friendly group configurations at a higher rate than the method of Moussaïd et al. The differences in quality were small but significant in scenarios that involved groups of two. In scenarios were groups of three or four were considered, the differences in quality were even bigger.

What is more, SGN does not rely on a formation model for steering social-groups. Simulated groups switch between line-abreast, V-like, and river-like formations. However, these navigation choices are not explicitly modelled. Instead, each member of a group adjusts its motion in an effort to remain social within its group. This implies that the generated social-group behaviors are emergent. At the same time, this modelling choice allows for a group to temporarily split, for instance due to overtaking another group. This type of behavior is not exhibited in methods that employ group formation models [47, 92]. Nevertheless, it has been suggested that such avoidance maneuvers are an accepted behavior for social pedestrian groups [51].

Also, SGN does not set an upper bound on the number of members that form a social group. We have evaluated group behavior for groups sizes two to four, because these are the group sizes that have been observed to prevail in pedestrian areas [19, 66]. At the same time, we provide details on how SGN can be easily adjusted, when we desire to split large social groups into smaller sub-groups.

Our performance evaluation indicates that our method can be used for real-time simulations that involve thousands of agents organized in small social-groups. As a result, SGN can be used for mass-event simulations or gaming applications that should exhibit social group behavior.

Finally, another contribution of our work are the quality-evaluation metrics with respect to social group behavior. We have proposed an intuitive way to define and measure the ability of a pedestrian group to remain coherent and social. Coherence and sociality are examined from the perspective of each individual member of a group. Therefore, for a group

to be classified as coherent or social, all members must perceive it as such. What is more, our definitions of group coherence and sociality are on par with empirical data [19, 66, 91], regarding the spatial arrangements of social pedestrian groups.

6.2 Limitations

Despite the strengths of our method, some limitations can be observed. As discussed in the previous section, SGN does not utilize a group-formation model. As a result, there is no control over the observed formation changes. In a simulation scenario, it might be desired that a specific group never splits, but is allowed to switch to a river-like formation instead. A mother with her child can be considered as an example. Currently, our method cannot guarantee this sort of behavior.

What is more, in SGN, social groups do not always avoid other groups as a whole. The motivation behind our choice is that, in densely-populated areas, we expect groups to invade the social space that another group induces. This is also expected to be the case when a single group remains stationary inside a corridor and no alternative routes exist for moving pedestrians. Therefore, we have not explicitly modelled this behavior.

Another limitation of our method is attributed to its vision-based avoidance algorithm. Similar to Moussaïd et al. [65], SGN uses a circular sector to model the vision-field of an agent. In SGN, each agent is also allowed to have a different maximum viewing distance. Nevertheless, this distance is fixed. In real-life however, the vision-field is sensitive to occlusions. In an open and rather unoccupied environment, a pedestrian might be able to see hundreds of meters ahead. Also a pedestrian might or might not be able to see past another pedestrian, depending on their height difference. These visual-stimulus aspects are considered too expensive and are not modelled by our method. However, we are are not aware of any avoidance method that does so.

Furthermore, our method might fail to deal with static agents under certain conditions. For instance, when several groups remain static within an area, another group might try to pass through this occupied area. In detail, the employed global route-planning method [28] does not account for crowd density and produces paths that might cross areas were static agents reside. This is a reasonable choice, since density can change several times by the time the group approaches the congested area. However, our collision-avoidance algorithm relies on the vision-field model we discussed above. As a result, the maximum viewing distance might be insufficient to allow for a selection of a direction of motion that avoids congested areas.

In addition, our method models agents using a disc shape. We have followed this approach for several reasons. First, we consider a disc to be a good approximation of the space that a pedestrian occupies, when projected on the navigation plane. At the same time, it allows for collision checks that are less expensive than those required when employing a more complex shape. This property of a disc-shaped agent allows our method to perform fast enough to be considered for simulating thousands of pedestrians that are organised in social groups. Also, the disc shape allows our method to be coupled with existing global route-planning algorithms that account for clearance [28, 44, 46, 89] without further adjustments. Nevertheless, it has been suggested that an elliptical shape can be a better approximation of the space occupied by agents whose shape is long and narrow [21].

Finally, our collision-avoidance algorithm does not utilize the personal space of each agent. No personal space is used in the original avoidance algorithm of Moussaïd et al. [65]. We have experimented with including personal space in our collision-avoidance algorithm by accounting for it during the computation of $f(\alpha_{ij})$ (see Equation 4.8). However, in dense regions, this led agents to stop moving, because their personal spaces were violated. Therefore, we decided not to follow a collision-avoidance approach that uses personal spaces. As a result, when large crowds evacuate areas through narrow passages, density increases drastically at bottlenecks. This, in turn, leads to crowd congestion. This phenomenon has been empirically verified [88]. However, we expect our lack of personal-space utilization during local-movement phase to cause congestions more frequently than what is observed in real-life.

6.3 Future work

An area for future work is to perform research on heuristics that members of social groups employ throughout their navigation. Also, it would be interesting to investigate how these heuristics interact with the heuristics that pedestrians employ at an individual level [65]. An effective combination of avoidance heuristics with group-relevant heuristics is expected to further increase the ability of agents to remain social, while avoiding obstacles and other agents in the simulation.

Furthermore, SGN could be extended to account for the demographics of the simulated pedestrians groups. Fridman et al. [25] provide empirical data regarding the size of groups, their dominance over individuals, and the gender of their members for several countries. They also discuss avoidance-side preference, walking speed, and personal spaces observed in these countries. Willis et al. [91] provide information on the effect of gender, group size and age on the inter-personal distances and walking speeds. It would be interesting to create agent profiles that model these differences and use them in our method.

As discussed in the limitations of our method, our vision-based collision-avoidance algorithm should be extended to account for the personal space of the agents. This would allow agents to adopt social priority rules avoid and mitigate congestion. As a result, we expect coordination behavior to emerge. The challenge in this approach is to efficiently determine the conditions under which an agent should take its personal space into account.

Regarding control over formation changes, SGN can be extended to selectively adopt a leader-follower approach. This would allow for simulation of social groups that never split, but switch to a river-like formation instead. Again, the challenge is to efficiently determine the conditions under which our method switches from regular route following to a leader-follower model and vice-versa.

Furthermore, the incorporation of psychological variables in SGN will enhance the ability of our method to simulate safety-engineering scenarios. Panic as well as panic propagation are expected to have a significant impact on the crowd dynamics in such simulations. Aguirre et al. [5] suggest that, under panic, the social interactions between individuals are further enhanced thus leading to collective behavior. Also, Mawson [62] has observed that the collective behavior of social groups under stress often increases evacuation times. However, more research should be conducted on how panic affects existing social groups in the microscopic scale.

In addition, SGN can be extended to improve on treatment of static agents and conversing groups. A density-based planning approach seems crucial for achieving this effect. The main challenge in this approach are the expensive computations required for deciding whether it is possible or not to avoid invading the social space of conversing groups.

Finally, our method could possibly be extended to account for terrain preferences. MIRAN algorithm [41], that was employed by SGN for the route-following level of the simulation, already accounts for terrain preference for each individual. When social-groups are considered in heterogeneous environments, it would be interesting to investigate how global route-planning and route-following levels are affected.

References

[1] COMMIT/ project. http://www.commit-nl.nl/, 2014. [Online; accessed 2014].

[2] Disney world theme parks. http://disneyworld.disney.go.com/, 2014. [Online; accessed 2014].

[3] Euler method. http://en.wikipedia.org/wiki/Euler_method, 2014. [Online; accessed 2014].

[4] Explicit corridor map crowd simulation framework. http://www.staff.science.uu.nl/~gerae101/motion_planning/ecm/, 2014. [Online; accessed 2014].

[5] Benigno E. Aguirre, Dennis Wenger, and Gabriela Vigo. A test of the emergent norm theory of collective behavior. *Sociological Forum*, 13(2):301–320, 1998.

[6] Ronald C. Arkin. Cooperation without communication: Multiagent schema-based robot navigation. *Journal of Robotic Systems*, 9(3):351–364, 1992.

[7] Adrian F. Aveni. The not-so-lonely crowd: Friendship groups in collective behavior. *Sociometry*, 40(1):96–99, 1977.

[8] Michael Batty. Predicting where we walk. *Nature*, 388:19–20, 1997.

[9] Jur van den Berg, Stephen J. Guy, Ming C. Lin, and Dinesh Manocha. Reciprocal n-body collision avoidance. In *International Symposium of Robotics Research*, 2009.

[10] Jur van den Berg, Ming Lin, and Dinesh Manocha. Reciprocal velocity obstacles for real-time multi-agent navigation. In *Robotics and Automation, 2008. ICRA 2008. IEEE International Conference on*, pages 1928–1935, 2008.

[11] David C. Brogan and Jessica K. Hodgins. Group behaviors for systems with significant dynamics. *Autonomous Robots*, 4(1):137–153, 1997.

[12] Carsten Burstedde, Kai Klauck, Andreas Schadschneider, and Johannes Zittartz. Simulation of pedestrian dynamics using a two-dimensional cellular automaton. *Physica A: Statistical Mechanics and its Applications*, 295(34):507–525, 2001.

[13] James A. Cheyne and Michael G. Efran. The effect of spatial and interpersonal variables on the invasion of group controlled territories. *Sociometry*, 35(3):477–489, 1972.

[14] Herbert H. Clark. *Using language*. Cambridge University Press, 1996.

[15] Herbert H. Clark and Susan E. Brennan. Grounding in communication. *Perspectives on socially shared cognition*, 13:127–149, 1991.

[16] Chris Cocking, John Drury, and Steve Reicher. The psychology of crowd behaviour in emergency evacuations: Results from two interview studies and implications for the fire and rescue services. *The Irish Journal of Psychology*, 30(1-2):59–73, 2009.

[17] James S. Coleman and John James. The equilibrium size distribution of freely-forming groups. *Sociometry*, 24(1):36–45, 1961.

[18] Luca Consolini, Fabio Morbidi, Domenico Prattichizzo, and Mario Tosques. Leader-follower formation control of nonholonomic mobile robots with input constraints. *Automatica*, 44(5):1343–1349, 2008.

[19] Marco Costa. Interpersonal distances in group walking. *Journal of Nonverbal Behavior*, 34(1):15–26, 2010.

[20] James E. Cutting, Peter M. Vishton, and Paul A. Braren. How we avoid collisions with stationary and moving objects. *Psychological review*, 102(4):627, 1995.

[21] Khac Duc Do. Flocking for multiple elliptical agents with limited communication ranges. *Robotics, IEEE Transactions on*, 27(5):931–942, 2011.

[22] John Drury, Chris Cocking, and Steve Reicher. Everyone for themselves? a comparative study of crowd solidarity among emergency survivors. *British Journal of Social Psychology*, 48(3):487–506, 2009.

[23] Cathy Ennis and Carol O'Sullivan. Perceptually plausible formations for virtual conversers. *Computer Animation and Virtual Worlds*, 23(3-4):321–329, 2012.

[24] Paolo Fiorini and Zvi Shiller. Motion planning in dynamic environments using velocity obstacles. *The International Journal of Robotics Research*, 17(7):760–772, 1998.

[25] Natalie Fridman, Gal A. Kaminka, and Avishay Zilka. The impact of culture on crowd dynamics: An empirical approach. In *Proceedings of the 2013 International Conference on Autonomous Agents and Multi-agent Systems*, AAMAS '13, pages 143–150, 2013.

[26] John J. Fruin. *Pedestrian planning and design.* New York: Metropolitan Association of Urban Designers and Environmental Planners, 1971.

[27] John J. Fruin. The causes and prevention of crowd disasters. In *Smith RA, Dickie JF (eds) Engineering for Crowd Safety.* Amsterdam, Elsevier, 1993.

[28] Roland Geraerts. Planning short paths with clearance using explicit corridors. In *Robotics and Automation (ICRA), 2010 IEEE International Conference on*, pages 1997–2004, 2010.

[29] James J. Gibson. Visually controlled locomotion and visual orientation in animals*. *British Journal of Psychology*, 49(3):182–194, 1958.

[30] James J. Gibson. *The ecological approach to visual perception.* Houghton Mifflin, 1979.

[31] Gerd Gigerenzer. Why heuristics work. *Perspectives on Psychological Science*, 3(1):20–29, 2008.

[32] Gerd Gigerenzer, Peter M. Todd, and ABC Research Group. *Simple Heuristics That Make Us Smart.* Oxford University Press, 1999.

[33] Erving Goffman. *Relations in public: microstudies of the public order.* New York : Basic Books, 1971.

[34] Laure Heïgeas, Annie Luciani, Joelle Thollot, and Nicolas Castagné. A physically-based particle model of emergent crowd behaviors. In *Graphicon*, pages 5–10, 2003.

[35] Dirk Helbing, Illés Farkas, Péter Molnár, and Tamás Vicsek. Simulation of pedestrian crowds in normal and evacuation situations. 2000.

[36] Dirk Helbing, Illés Farkas, and Tamás Vicsek. Simulating dynamical features of escape panic. *Nature*, 407:487–490, 2000.

[37] Dirk Helbing and Péter Molnár. Social force model for pedestrian dynamics. *Physical Review E*, 51:4282–4286, 1995.

[38] Serge P. Hoogendoorn, Piet HL. Bovy, and Winnie Daamen. Microscopic pedestrian wayfinding and dynamics modelling. *Pedestrian and evacuation dynamics*, pages 123–154, 2002.

[39] M. Ani Hsieh, Anthony Cowley, James F. Keller, Luiz Chaimowicz, Ben Grocholsky, Vijay Kumar, Camillo J. Taylor, Yoichiro Endo, Ronald C. Arkin, Boyoon Jung, Denis F. Wolf, Gaurav S. Sukhatme, and Douglas C. MacKenzie. Adaptive teams of autonomous aerial and ground robots for situational awareness. *Journal of Field Robotics*, 24(11-12):991–1014, 2007.

[40] Jiangyang Huang, S.M. Farritor, A. Qadi, and S. Goddard. Localization and follow-the-leader control of a heterogeneous group of mobile robots. *Mechatronics, IEEE/ASME Transactions on*, 11(2):205–215, April 2006.

[41] Norman Jaklin, Atlas Cook IV, and Roland Geraerts. Real-time path planning in heterogeneous environments. *Computer Animation and Virtual Worlds*, 24(3-4):285–295, 2013.

[42] John James. The distribution of free-forming small group size. *American Sociological Review*, 1953.

[43] Eunjung Ju, Myung Geol Choi, Minji Park, Jehee Lee, Kang Hoon Lee, and Shigeo Takahashi. Morphable crowds. *ACM Transactions on Graphics*, 29(6):140:1–140:10, 2010.

[44] Marcelo Kallmann. Shortest paths with arbitrary clearance from navigation meshes. In *Proceedings of the 2010 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, SCA '10, pages 159–168, Aire-la-Ville, Switzerland, 2010. Eurographics Association.

[45] Arno Kamphuis and Mark H. Overmars. Finding paths for coherent groups using clearance. In *Proceedings of the 2004 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, SCA '04, pages 19–28, Aire-la-Ville, Switzerland, Switzerland, 2004. Eurographics Association.

[46] Ioannis Karamouzas, Roland Geraerts, and Mark H. Overmars. Indicative routes for path planning and crowd simulation. In *Proceedings of the 4th International Conference on Foundations of Digital Games*, pages 113–120. ACM, 2009.

[47] Ioannis Karamouzas and Mark H. Overmars. Simulating human collision avoidance using a velocity-based approach. In *VRIPHYS 10: 7th Workshop on Virtual Reality Interactions and Physical Simulations*, pages 125–134. Eurographics Association, 2010.

[48] Ioannis Karamouzas and Mark H. Overmars. Simulating the local behaviour of small pedestrian groups. In *VRST '10: Proceedings of the 17th ACM Symposium on Virtual Reality Software and Technology*, pages 183–190. ACM, 2010.

[49] Andrew Kimmel, Andrew Dobson, and Kostas Bekris. Maintaining team coherence under the velocity obstacle framework. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems - Volume 1*, AAMAS '12, pages 247–256, Richland, SC, 2012. International Foundation for Autonomous Agents and Multiagent Systems.

[50] Hubert Klüpfel. *A cellular automaton model for crowd movement and egress simulation.* PhD thesis, Universitätsbibliothek Duisburg, 2003.

[51] Gerta Köster, Franz Treml, Michael Seitz, and Wolfram Klein. Validation of crowd models including social groups. In Ulrich Weidmann, Uwe Kirsch, and Michael Schreckenberg, editors, *Pedestrian and Evacuation Dynamics 2012*, pages 1051–1063. Springer International Publishing, 2014.

[52] Taesoo Kwon, Kang Hoon Lee, Jehee Lee, and Shigeo Takahashi. Group motion editing. *ACM Transactions on Graphics*, 27(3):80:1–80:8, 2008.

[53] Yu-Chi Lai, Stephen Chenney, and Shaohua Fan. Group motion graphs. In *Proceedings of the 2005 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, SCA '05, pages 281–290, New York, NY, USA, 2005. ACM.

[54] Taras I. Lakoba, David. J. Kaup, and Neal M. Finkelstein. Modifications of the helbing-molnár-farkas-vicsek social force model for pedestrian evolution. *Simulation*, 81(5):339–352, May 2005.

[55] Steven M. LaValle. *Planning Algorithms*. Cambridge University Press, Cambridge, U.K., 2006. Available at http://planning.cs.uiuc.edu/.

[56] Kang Hoon Lee, Myung Geol Choi, Qyoun Hong, and Jehee Lee. Group behavior from video: A data-driven approach to crowd simulation. In *Proceedings of the 2007 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, SCA '07, pages 109–118, Aire-la-Ville, Switzerland, Switzerland, 2007. Eurographics Association.

[57] Naomi E. Leonard and Edward Fiorelli. Virtual leaders, artificial potentials and coordinated control of groups. In *Decision and Control, 2001. Proceedings of the 40th IEEE Conference on*, volume 3, pages 2968–2973, 2001.

[58] Alon Lerner, Yiorgos Chrysanthou, and Dani Lischinski. Crowds by example. *Computer Graphics Forum*, 26(3):655–664, 2007.

[59] Jack Liddle, Armin Seyfried, Bernhard Steffen, Wolfram Klingsch, Tobias Rupprecht, Andreas Winkens, and Maik Boltes. Microscopic insights into pedestrian motion through a bottleneck, resolving spatial and temporal variations. *arXiv:1105.1532v1*, 2011.

[60] Maja J. Mataric. Minimizing complexity in controlling a mobile robot population. In *Robotics and Automation, 1992. Proceedings., 1992 IEEE International Conference on*, volume 1, pages 830–835, 1992.

[61] Maja J. Mataric. Designing emergent behaviors: From local interactions to collective intelligence. In *International Conference on From Animals to Animats 2 : Simulation of Adaptive Behavior: Simulation of Adaptive Behavior*, pages 432–441, Cambridge, MA, USA, 1993. MIT Press.

[62] Anthony R. Mawson. Understanding mass panic and other collective responses to threat and disaster. *Psychiatry: Interpersonal Biological Processes*, 68(2):95–113, 2005.

[63] Luis Yoichi Morales Saiki, Satoru Satake, Rajibul Huq, Dylan Glas, Takayuki Kanda, and Norihiro Hagita. How do people walk side-by-side?: Using a computational model of human behavior for a social robot. In *Proceedings of the Seventh Annual ACM/IEEE International Conference on Human-Robot Interaction*, HRI '12, pages 301–308, New York, NY, USA, 2012. ACM.

[64] Mehdi Moussaïd, Dirk Helbing, Simon Garnier, Anders Johansson, Maud Combe, and Guy Theraulaz. Experimental study of the behavioural mechanisms underlying self-organization in human crowds. *Proceedings of the Royal Society B: Biological Sciences*, 2009.

[65] Mehdi Moussaïd, Dirk Helbing, and Guy Theraulaz. How simple rules determine pedestrian behavior and crowd disasters. *Proceedings of the National Academy of Sciences*, 108(17):6884–6888, 2011.

[66] Mehdi Moussaïd, Niriaska 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.

[67] Soraia R. Musse and Daniel Thalmann. A model of human crowd behavior : Group inter-relationship and collision detection analysis. In Daniel Thalmann and Michiel Panne, editors, *Computer Animation and Simulation 97*, Eurographics, pages 39–51. Springer Vienna, 1997.

[68] Jan Ondřej, Julien Pettré, Anne-Hélène Olivier, and Stéphane Donikian. A synthetic-vision based steering approach for crowd simulation. In *ACM SIGGRAPH 2010 Papers*, SIGGRAPH '10, pages 1–9, New York, NY, USA, 2010. ACM.

[69] Seung In Park, Francis Quek, and Yong Cao. Modeling social groups in crowds using common ground theory. In *Proceedings of the Winter Simulation Conference*, WSC '12, pages 113:1–113:12. Winter Simulation Conference, 2012.

[70] Sachin Patil, Jur van den Berg, Sean Curtis, Ming C. Lin, and Dinesh Manocha. Directing crowd simulations using navigation fields. *Visualization and Computer Graphics, IEEE Transactions on*, 17(2):244–254, 2011.

[71] Nuria Pelechano, Jan M. Allbeck, and Norm I. Badler. Controlling individual agents in high-density crowd simulation. In *Proceedings of the 2007 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, SCA '07, pages 99–108, Aire-la-Ville, Switzerland, Switzerland, 2007. Eurographics Association.

[72] Christopher Peters and Cathy Ennis. Modeling groups of plausible virtual pedestrians. *IEEE Computer Graphics and Applications*, 29(4):54–63, 2009.

[73] Julien Pettré, Jan Ondřej, Anne-Hélène Olivier, Armel Cretual, and Stéphane Donikian. Experiment-based modeling, simulation and validation of interactions between virtual walkers. In *Proceedings of the 2009 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, SCA '09, pages 189–198, New York, NY, USA, 2009. ACM.

[74] Fasheng Qiu and Xiaolin Hu. Modeling group structures in pedestrian crowd simulation. *Simulation Modelling Practice and Theory*, 18(2):190–205, 2010.

[75] Craig W. Reynolds. Flocks, herds and schools: A distributed behavioral model. *SIGGRAPH Computer Graphics*, 21(4):25–34, 1987.

[76] Craig W. Reynolds. Steering behaviors for autonomous characters. In *Game Developers Conference 1999*, pages 763–782, 1999.

[77] Craig W. Reynolds. Interaction with groups of autonomous characters. In *Game Developers Conference 2000*, pages 449–460, 2000.

[78] Michael E. Roloff. *Interpersonal communication: The social exchange approach*, volume 6 of *Commtext Series*. SAGE Publications, London, 1981.

[79] Siamak Sarmady, Fazilah Haron, and Abdullah Zawawi H. Talib. Multi-agent simulation of circular pedestrian movements using cellular automata. In *Modeling Simulation, 2008. AICMS 08. Second Asia International Conference on*, pages 654–659, 2008.

[80] Siamak Sarmady, Fazilah Haron, and Abdullah Zawawi H. Talib. Modeling groups of pedestrians in least effort crowd movements using cellular automata. In *Modelling Simulation, 2009. AMS '09. Third Asia International Conference on*, pages 520–525, 2009.

[81] Shlomo S. Sawilowsky. Fermat, Schubert, Einstein, and Behrens-Fisher: The probable difference between two means when $\sigma_1^2 \neq \sigma_2^2$. *Journal of Modern Applied Statistical Methods*, 1(2), 2002.

[82] Michael Seitz, Gerta Köster, and Alexander Pfaffinger. Pedestrian group behavior in a cellular automaton. In Ulrich Weidmann, Uwe Kirsch, and Michael Schreckenberg, editors, *Pedestrian and Evacuation Dynamics 2012*, pages 807–814. Springer International Publishing, 2014.

[83] Shawn Singh, Mubbasir Kapadia, Petros Faloutsos, and Glenn Reinman. SteerBench: a benchmark suite for evaluating steering behaviors. *Computer Animation and Virtual Worlds*, 20(5-6):533–548, 2009.

[84] Jamie Snape, Stephen J. Guy, and Jur van den Berg. Independent navigation of multiple robots and virtual agents. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: Volume 1 - Volume 1*, AAMAS '10, pages 1645–1646, Richland, SC, 2010. International Foundation for Autonomous Agents and Multiagent Systems.

[85] Shigeo Takahashi, Kenichi Yoshida, Taesoo Kwon, Kang Hoon Lee, Jehee Lee, and Sung Yong Shin. Spectral-based group formation control. *Computer Graphics Forum*, 28(2):639–648, 2009.

[86] Daniel Thalmann and Soraia R. Musse. *Crowd Simulation*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2007.

[87] Alasdair Turner and Alan Penn. Encoding natural movement as an agent-based system: an investigation into human pedestrian behaviour in the built environment. *Environment and Planning B: Planning and Design*, 29(4):473–490, 2002.

[88] Ulrich Weidmann. Transporttechnik der fussgänger. 1992.

[89] Ron Wein, Jur van den Berg, and Dan Halperin. The visibility–voronoi complex and its applications. In *Proceedings of the 21st Annual Symposium on Computational Geometry*, SCG '05, pages 63–72, New York, NY, USA, 2005. ACM.

[90] B. L. Welch. The significance of the difference between two means when the population variances are unequal. *Biometrika*, 29(3/4):350–362, 1938.

[91] Alexandra Willis, Nathalia Gjersoe, Catriona Havard, Jon Kerridge, and Robert Kukla. Human movement behaviour in urban spaces: implications for the design and modelling of effective pedestrian environments. *Environment and Planning B: Planning and Design*, 31(6):805828, 2004.

[92] Qianqian Wu, Qingge Ji, Jinghong Du, and Xiaolian Li. Simulating the local behavior of small pedestrian groups using synthetic-vision based steering approach. In *Proceedings of the 12th ACM SIGGRAPH International Conference on Virtual-Reality Continuum and Its Applications in Industry*, VRCAI '13, pages 41–50, New York, NY, USA, 2013. ACM.

[93] Song Xu and H.B.-L. Duh. A simulation of bonding effects and their impacts on pedestrian dynamics. *Intelligent Transportation Systems, IEEE Transactions on*, 11(1):153–161, 2010.

[94] George Kingsley Zipf. *Human behavior and the principle of least effort*. Addison-Wesley, 1949.

A Variables and parameters

n_{max}	the maximum allowed group size
$N_{\mathcal{G}}$	the total number of groups
N_W	the total number of wall segments
$\mathcal{G}_i$	the i-th social group in the simulation
n_i	the size of group $\mathcal{G}_i$
$\mathcal{A}_{ij}$	the j-th member of group $\mathcal{G}_i$
$\mathbf{s}_{ij}$	the start position of agent $\mathcal{A}_{ij}$
S_i	the set of start positions of all members of group $\mathcal{G}_i$
$\mathbf{g}_i$	the goal position of group $\mathcal{G}_i$
$r_{g,i}$	the goal radius of group $\mathcal{G}_i$
G_i	the goal area $(\mathbf{g}_i, r_{g,i})$ of group $\mathcal{G}_i$
$\rho(\mathbf{x}, r)$	the density inside a disc of radius r centered at position $\mathbf{x}$
$\mathbf{x}_{ij}$	the position of agent $\mathcal{A}_{ij}$
r_{ij}	the radius of agent $\mathcal{A}_{ij}$
$r_{p,ij}$	the personal-space radius of agent $\mathcal{A}_{ij}$
m_{ij}	the mass of agent $\mathcal{A}_{ij}$
$\vec{v}_{ij}$	the velocity of agent $\mathcal{A}_{ij}$
$v_{pref,ij}$	the preferred speed of agent $\mathcal{A}_{ij}$
$\vec{v}_{pref,ij}$	the preferred velocity of agent $\mathcal{A}_{ij}$
$v_{pref,i}$	the preferred speed of group $\mathcal{G}_i$
$\vec{v}_{des,ij}$	the desired velocity of agent $\mathcal{A}_{ij}$
$\hat{H}_{ij}$	the gazing unit vector of agent $\mathcal{A}_{ij}$
ϕ_{ij}	the half-angle of the vision field of agent $\mathcal{A}_{ij}$
$d_{Max,ij}$	the maximum viewing distance of the vision field of agent $\mathcal{A}_{ij}$
$d_{g(\mathcal{A}_{ij})}$	the curve-length distance of agent $\mathcal{A}_{ij}$ from group goal position $\mathbf{g}_i$ measured along the global path of $\mathcal{A}_{ij}$
L_i	the index of leader of group $\mathcal{G}_i$
$L_{start,i}$	the index of leader of group $\mathcal{G}_i$ during initial coordination phase
$\mathcal{A}_{iL_{start,i}}$	the leader of group $\mathcal{G}_i$ during initial coordination phase
l_i	the index of the last member of group $\mathcal{G}_i$
$\mathcal{A}_{il_i}$	the last member of group $\mathcal{G}_i$
X_{ij}	a point at the position $\mathbf{x}_i$ of agent $\mathcal{A}_{ij}$
α_{ij}	a candidate angle of motion for agent $\mathcal{A}_{ij}$
O_{ij}	the last visible point on the direction of preferred velocity $\vec{v}_{pref,ij}$ for agent $\mathcal{A}_{ij}$
$\alpha_{0,ij}$	the bearing angle of point O_{ij} with respect to agent $\mathcal{A}_{ij}$
$\Omega_{\alpha_{ij}}$	the last visible point on the direction that corresponds to candidate angle of motion α_{ij}
$F_{\alpha_{ij}}$	the foot of O_{ij} on $X_{ij}\Omega_{\alpha_{ij}}$
$T_{\alpha_{ij}}$	a point on the last collision-free position for the agent on the direction that corresponds to candidate angle of motion α_{ij}
$f(\alpha_{ij})$	the expected travel distance for agent $\mathcal{A}_{ij}$ on candidate angle of motion α_{ij}
$d(\alpha_{ij})$	the expected remaining distance for agent $\mathcal{A}_{ij}$ from point O_{ij} when the next directional change is expected if candidate angle of motion α_{ij} is chosen
τ	the relaxation time of all agents
k	a parameter defining the strength of physical contact forces in the simulation
$\vec{f}_{ij,pq}$	the physical contact force applied to agent $\mathcal{A}_{ij}$ due to collision with agent $\mathcal{A}_{pq}$
$\hat{n}_{ij,pq}$	a unit vector pointing from agent $\mathcal{A}_{pq}$ to agent $\mathcal{A}_{ij}$

W_u	the u-th wall line segment in the simulation
$\vec{f}_{ij,W_u}$	the physical contact force applied to agent $\mathcal{A}_{ij}$ due to collision with wall segment W_u
$\hat{n}_{ij,W_u}$	a unit vector pointing from wall line segment W_u to agent $\mathcal{A}_{ij}$
$\vec{f}_{ij}^{group}$	the group force applied to agent $\mathcal{A}_{ij}$
$\vec{f}_{ij}^{vis}$	the visual group force applied to agent $\mathcal{A}_{ij}$ for maintaining social interactions with fellow group members
$\vec{f}_{ij}^{att}$	the attractive group force applied to agent $\mathcal{A}_{ij}$ for maintaining group coherence
$P_{ij,ik}$	the projection of the center X_{ik} of agent $\mathcal{A}_{ik}$ on the line of the gazing vector $\hat{H}_{ij}$ of fellow member $\mathcal{A}_{ij}$
$\theta_{ij,ik}$	the required angle of rotation of the gazing vector $\hat{H}_{ij}$ for X_{ik} to be visible by agent $\mathcal{A}_{ij}$
β_1	a model parameter describing the strength of the social interactions between the members of each group
C_i	the centroid of all members of group $\mathcal{G}_i$
d_{ij,C_i}	the distance between agent $\mathcal{A}_{ij}$ and centroid C_i
$d_{C_i,threshold}$	the threshold distance between agent $\mathcal{A}_{ij}$ and centroid C_i
$\hat{U}_{ij}$	a unit vector pointing from agent $\mathcal{A}_{ij}$ to centroid C_i
β_2	a model parameter describing the strength of the coherence interactions between the members of each group

B Experiment results

B.1 Bidirectional corridor

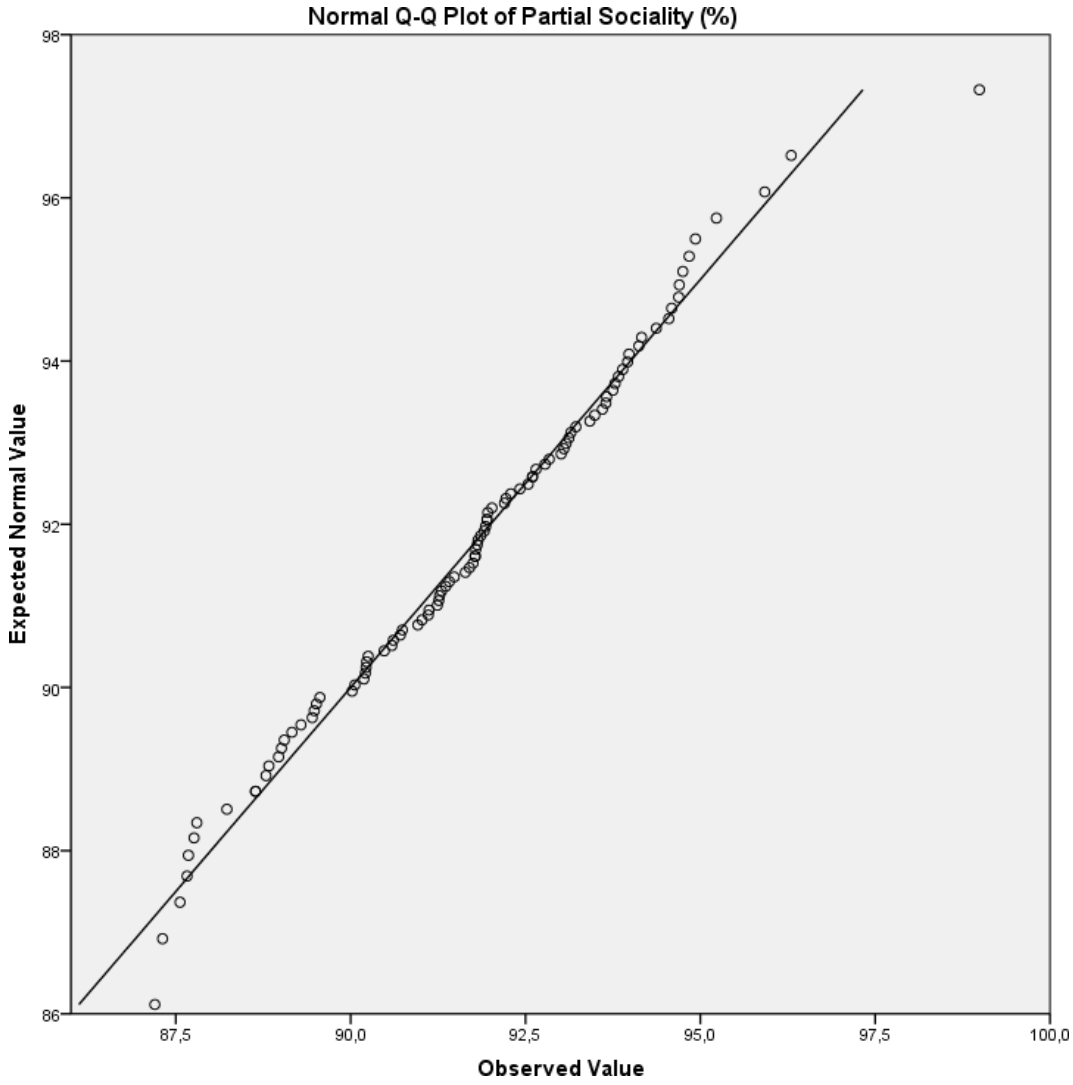

Figure B.1: Q-Q plot of average partial group sociality (%) for the bidirectional corridor scenario with groups of two using SGN method.

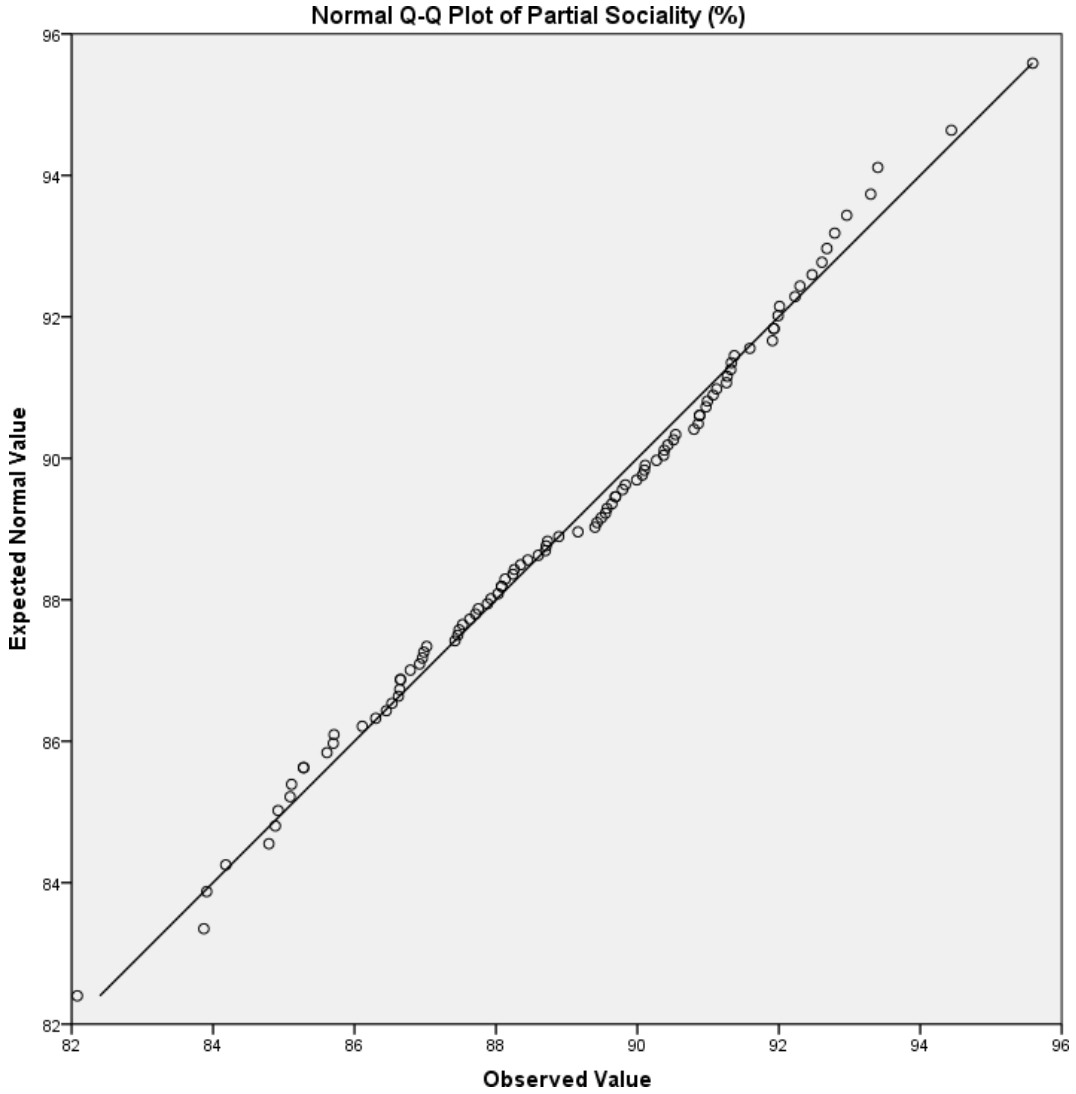

Figure B.2: Q-Q plot of average partial group sociality (%) for the bidirectional corridor scenario with groups of two using SGN method.

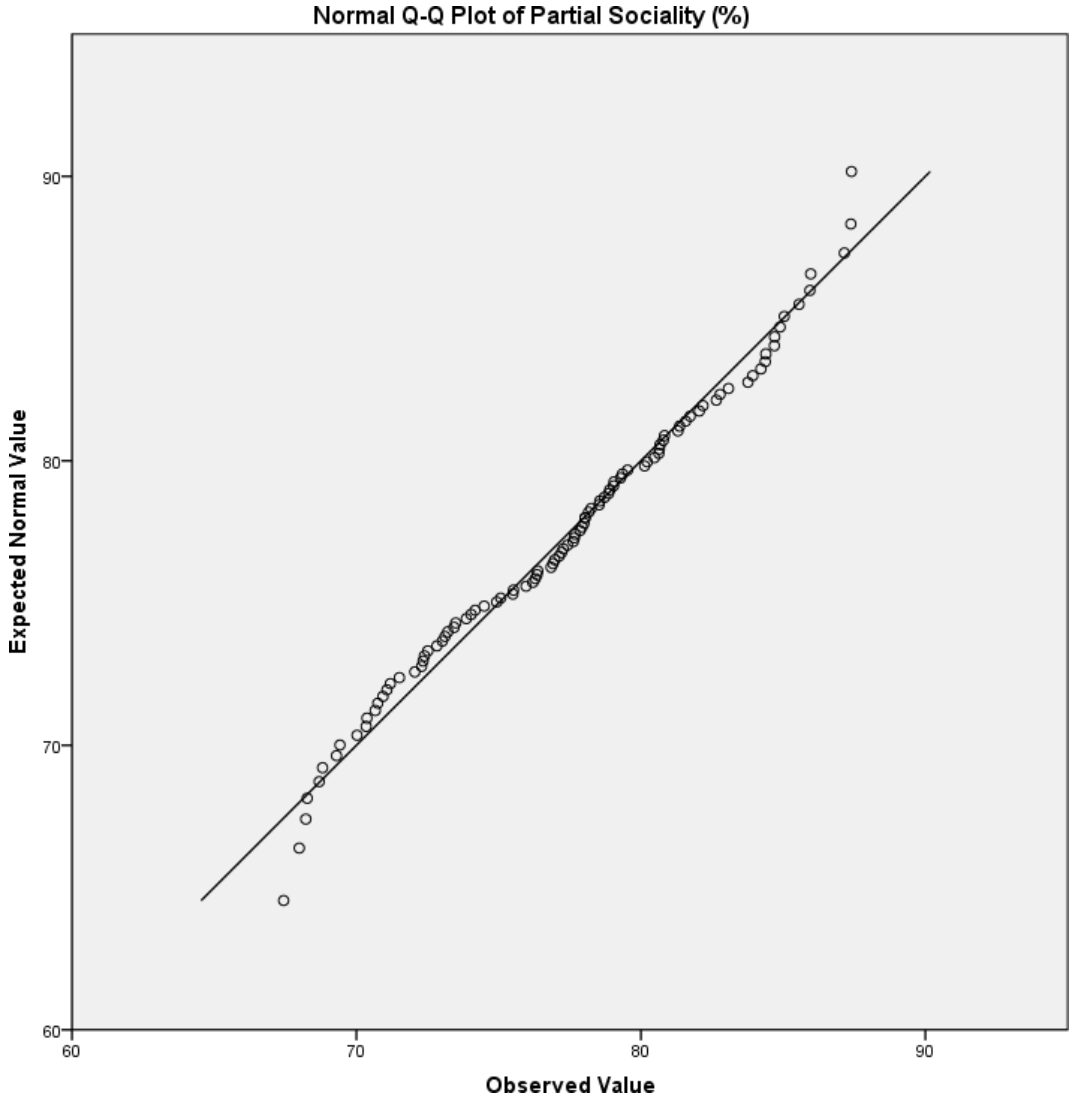

Figure B.3: Q-Q plot of average partial group sociality (%) for the bidirectional corridor scenario with groups of three using SGN method.

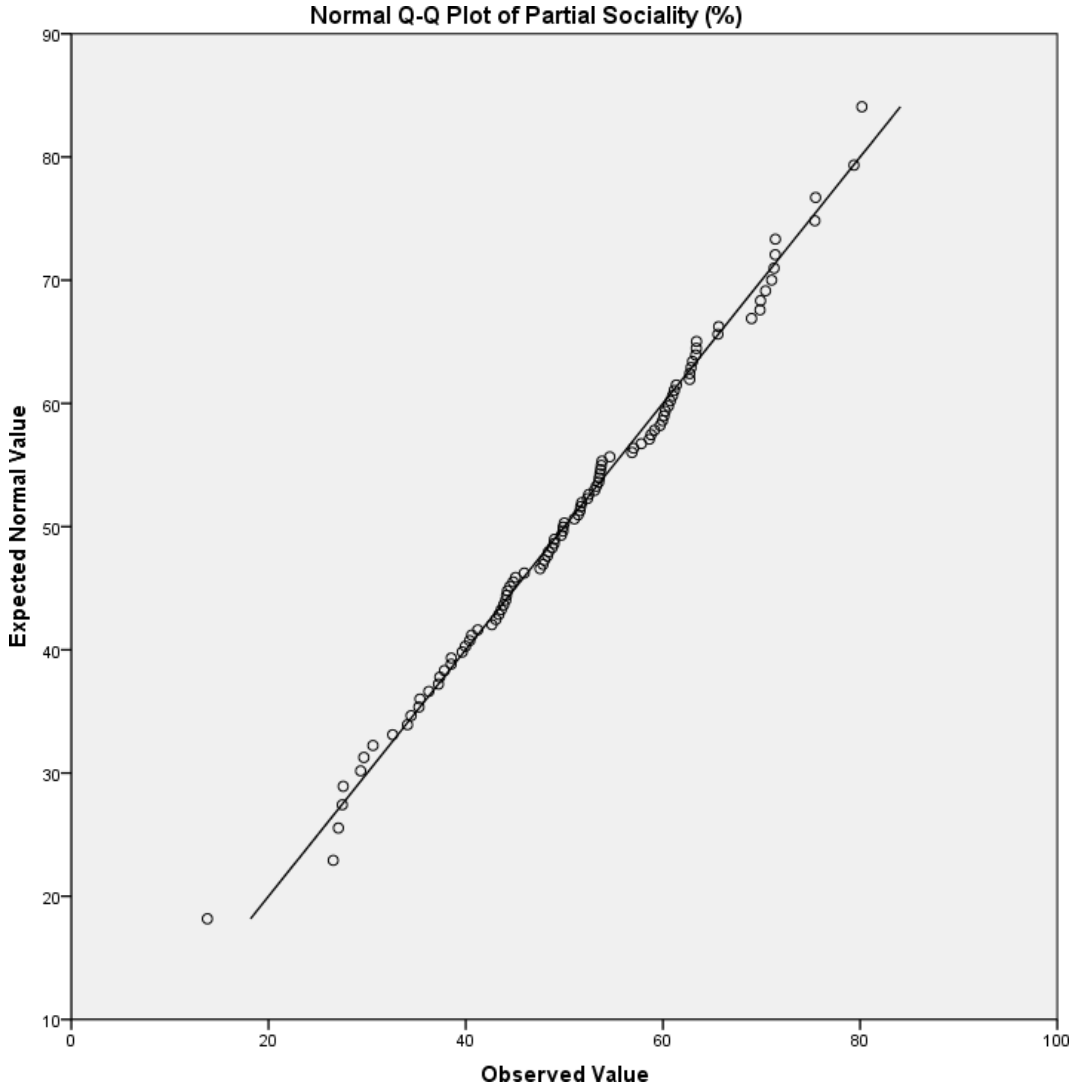

Figure B.4: Q-Q plot of average partial group sociality (%) for the bidirectional corridor scenario with groups of three using the method of Moussaïd et al.

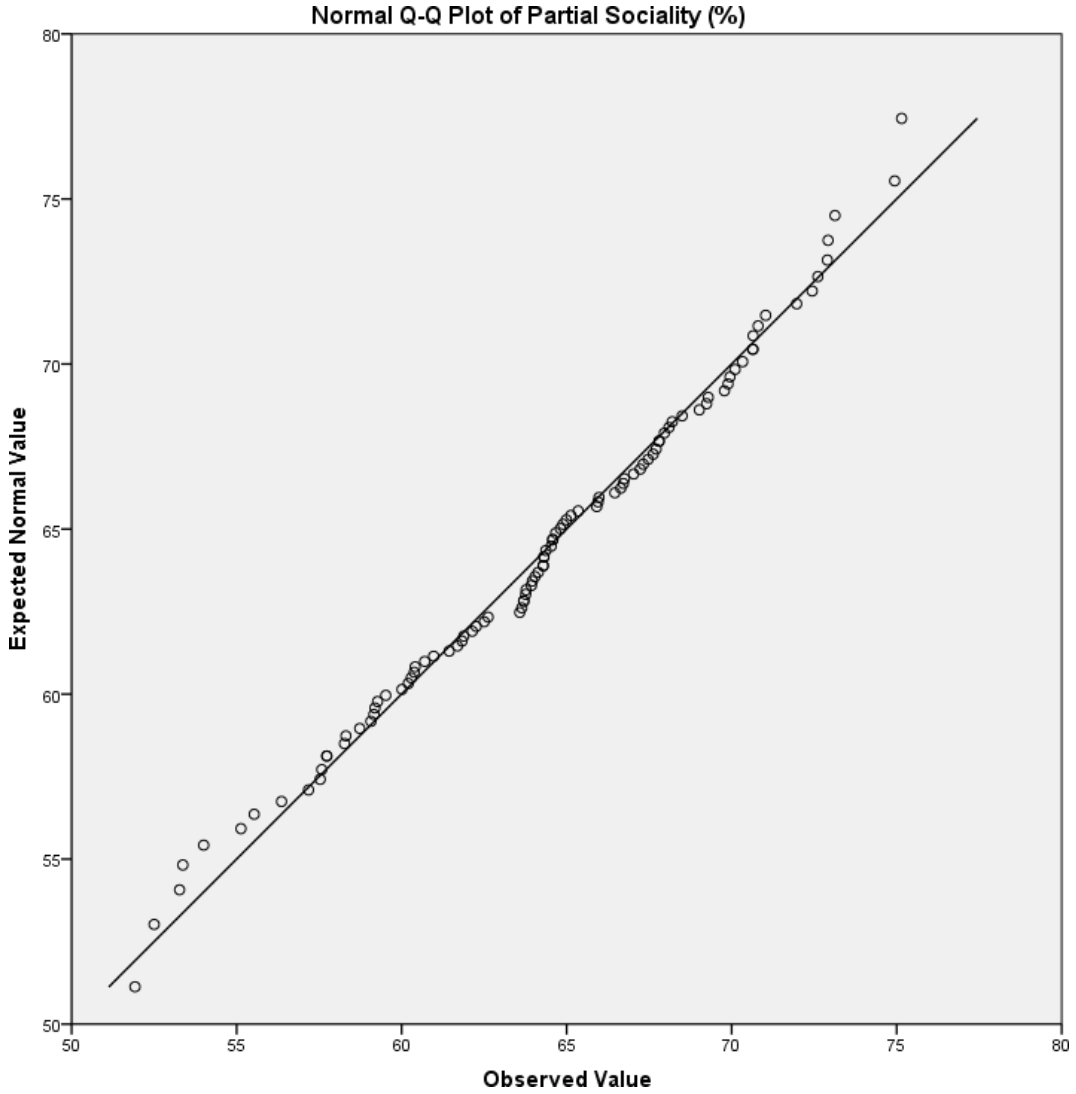

Figure B.5: Q-Q plot of average partial group sociality (%) for the bidirectional corridor scenario with groups of four using SGN method.

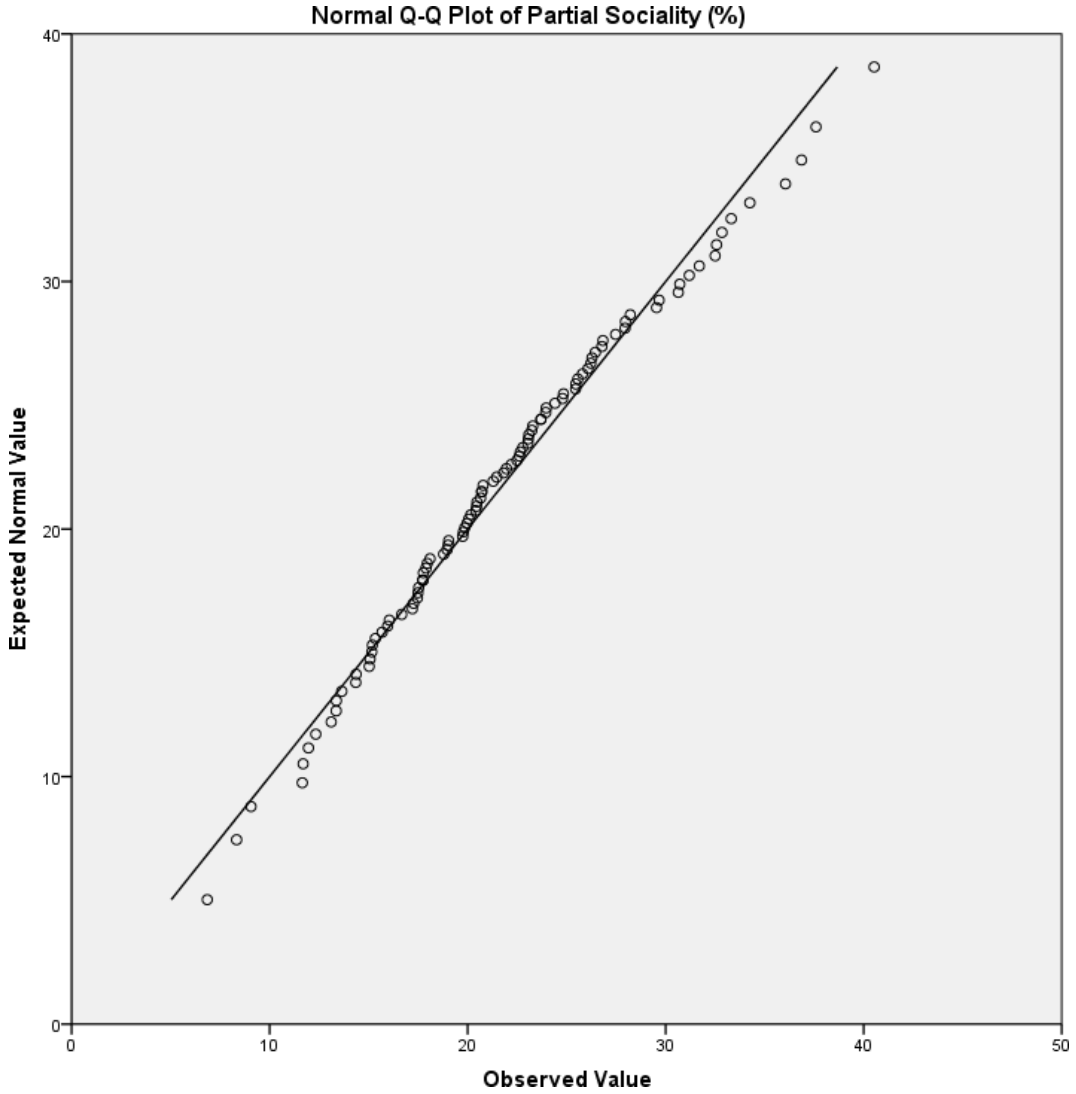

Figure B.6: Q-Q plot of average partial group sociality (%) for the bidirectional corridor scenario with groups of four using the method of Moussaïd et al.

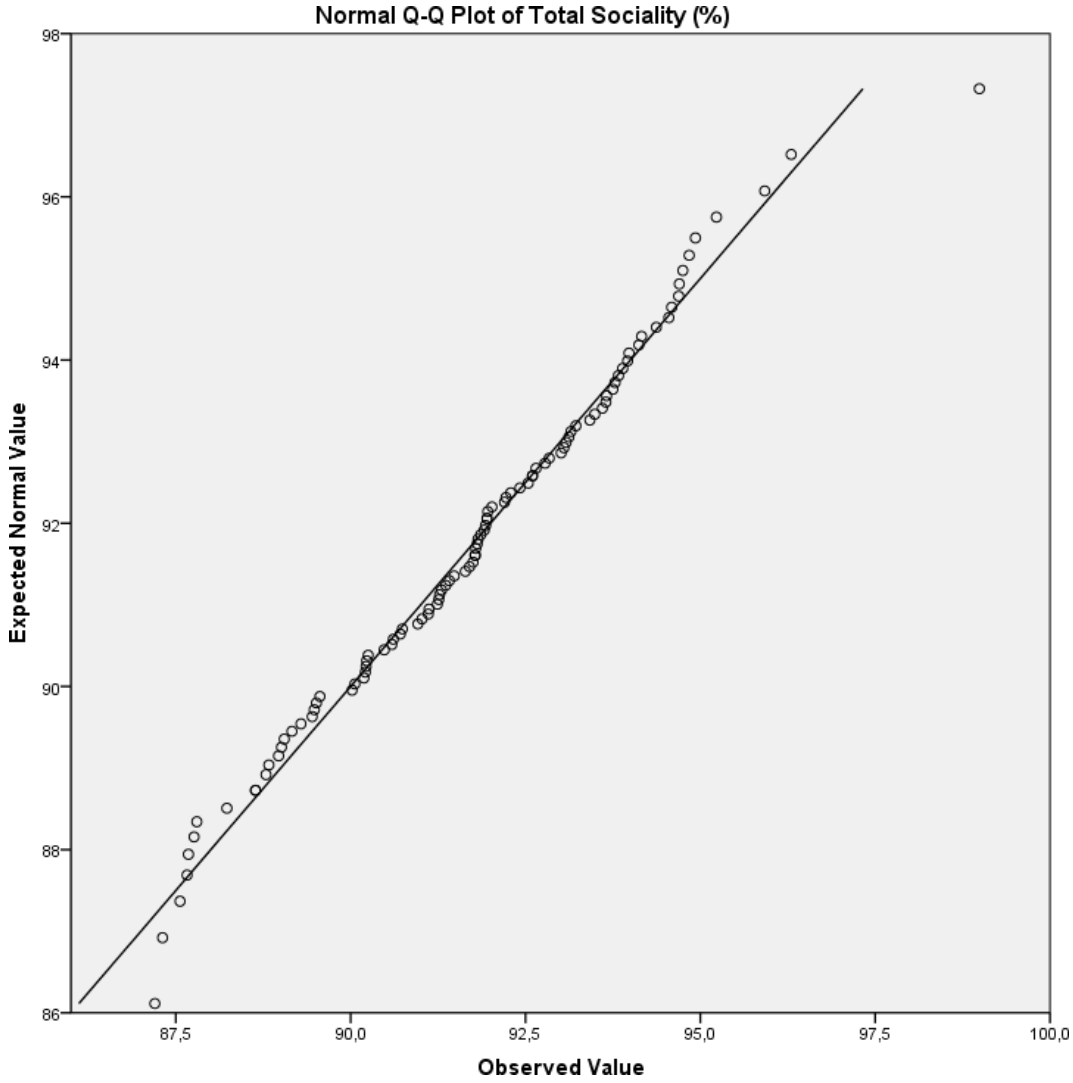

Figure B.7: Q-Q plot of average total group sociality (%) for the bidirectional corridor scenario with groups of two using SGN method.

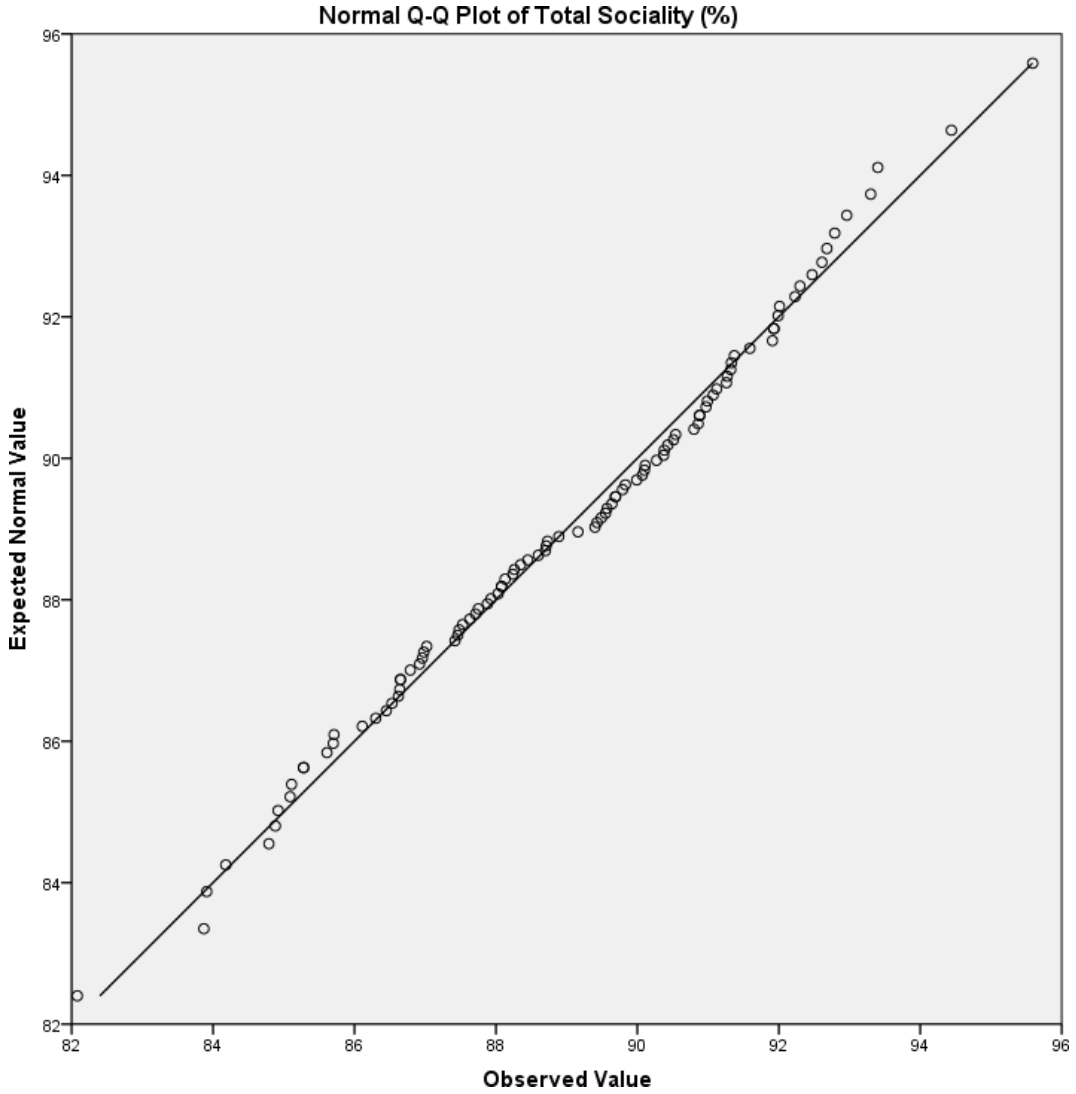

Figure B.8: Q-Q plot of average totall group sociality (%) for the bidirectional corridor scenario with groups of two using SGN method.

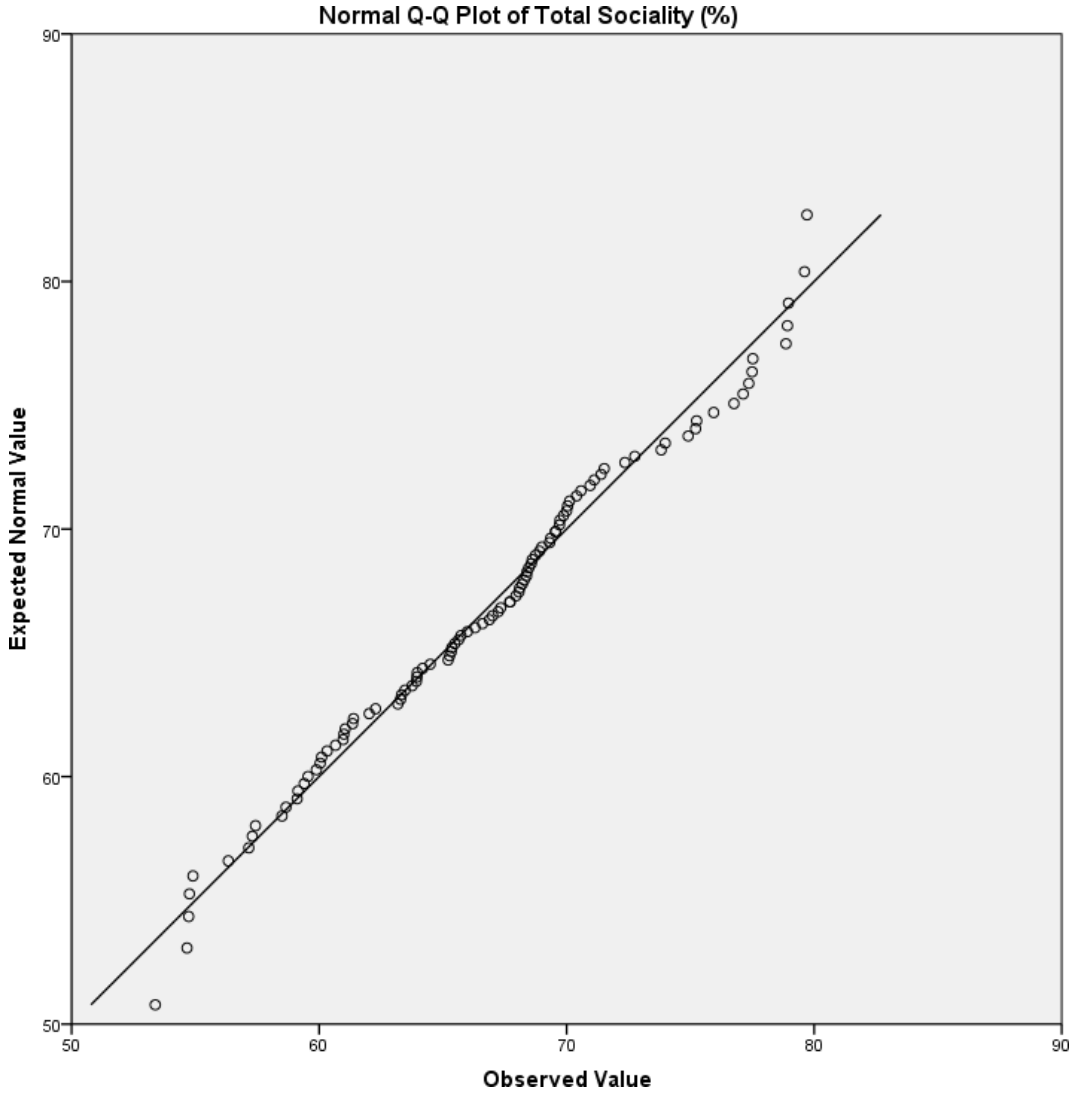

Figure B.9: Q-Q plot of average total group sociality (%) for the bidirectional corridor scenario with groups of three using SGN method.

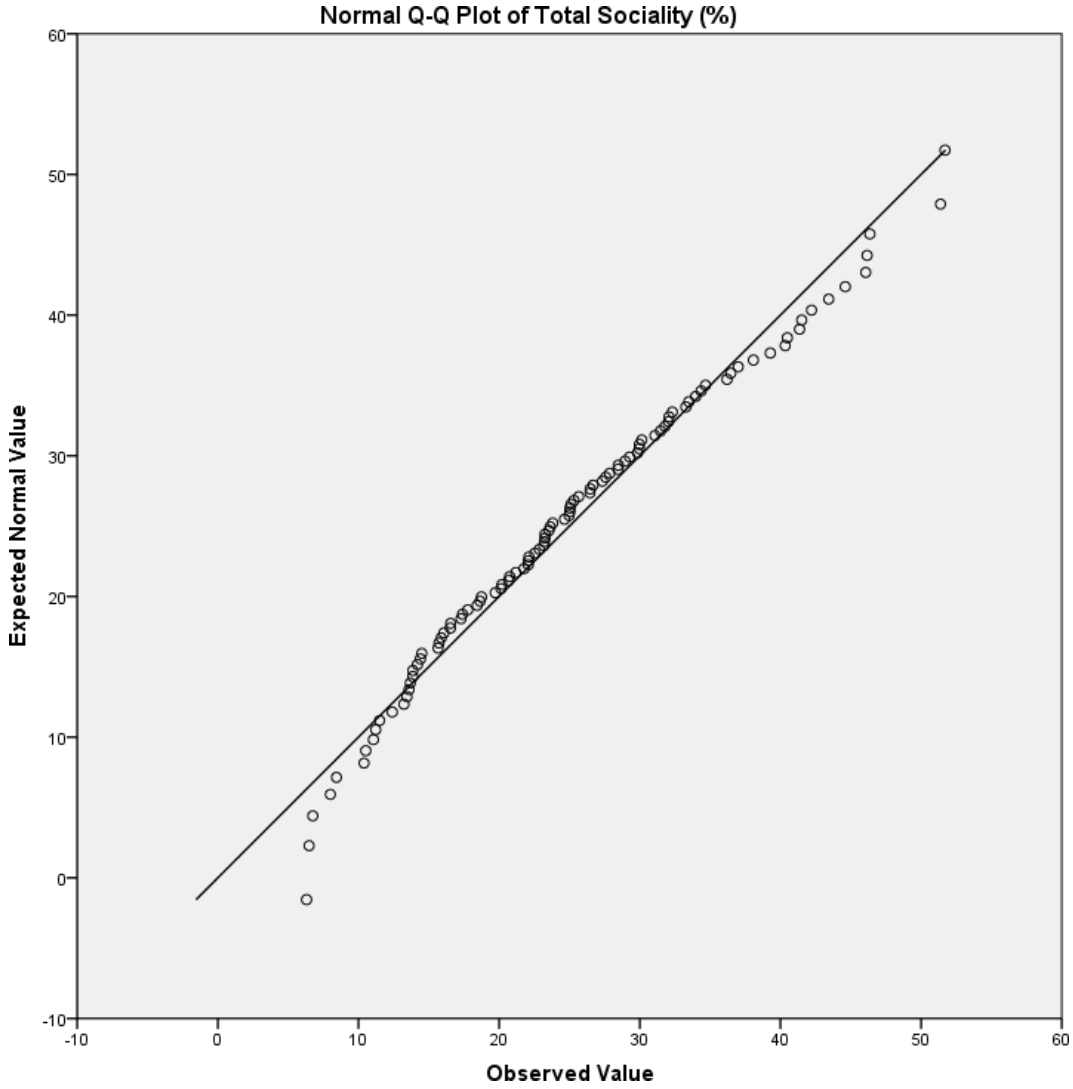

Figure B.10: Q-Q plot of average total group sociality (%) for the bidirectional corridor scenario with groups of three using the method of Moussaïd et al.

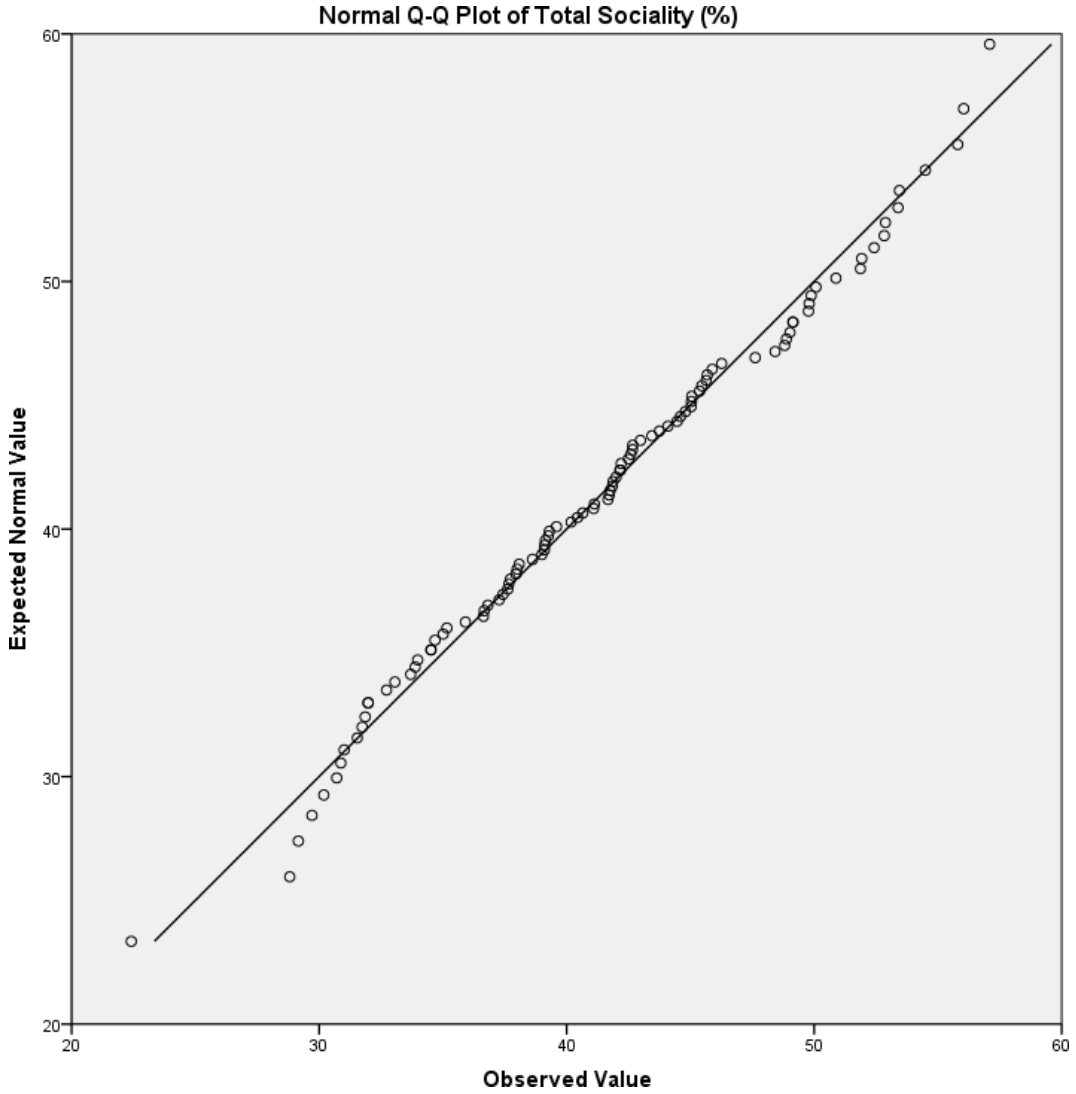

Figure B.11: Q-Q plot of average total group sociality (%) for the bidirectional corridor scenario with groups of four using SGN method.

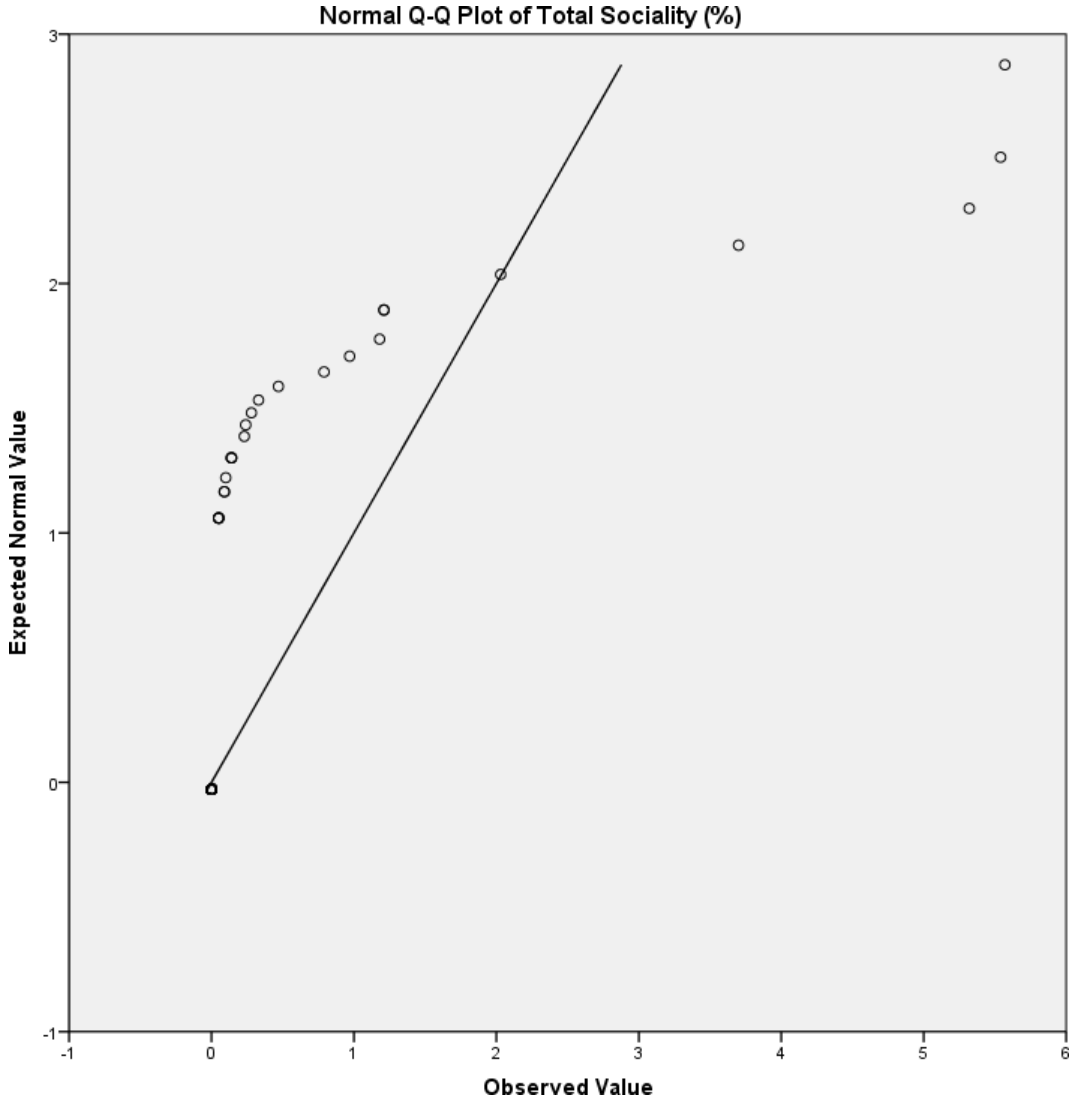

Figure B.12: Q-Q plot of average total group sociality (%) for the bidirectional corridor scenario with groups of four using the method of Moussaïd et al.

B.2 Squeezing corridor

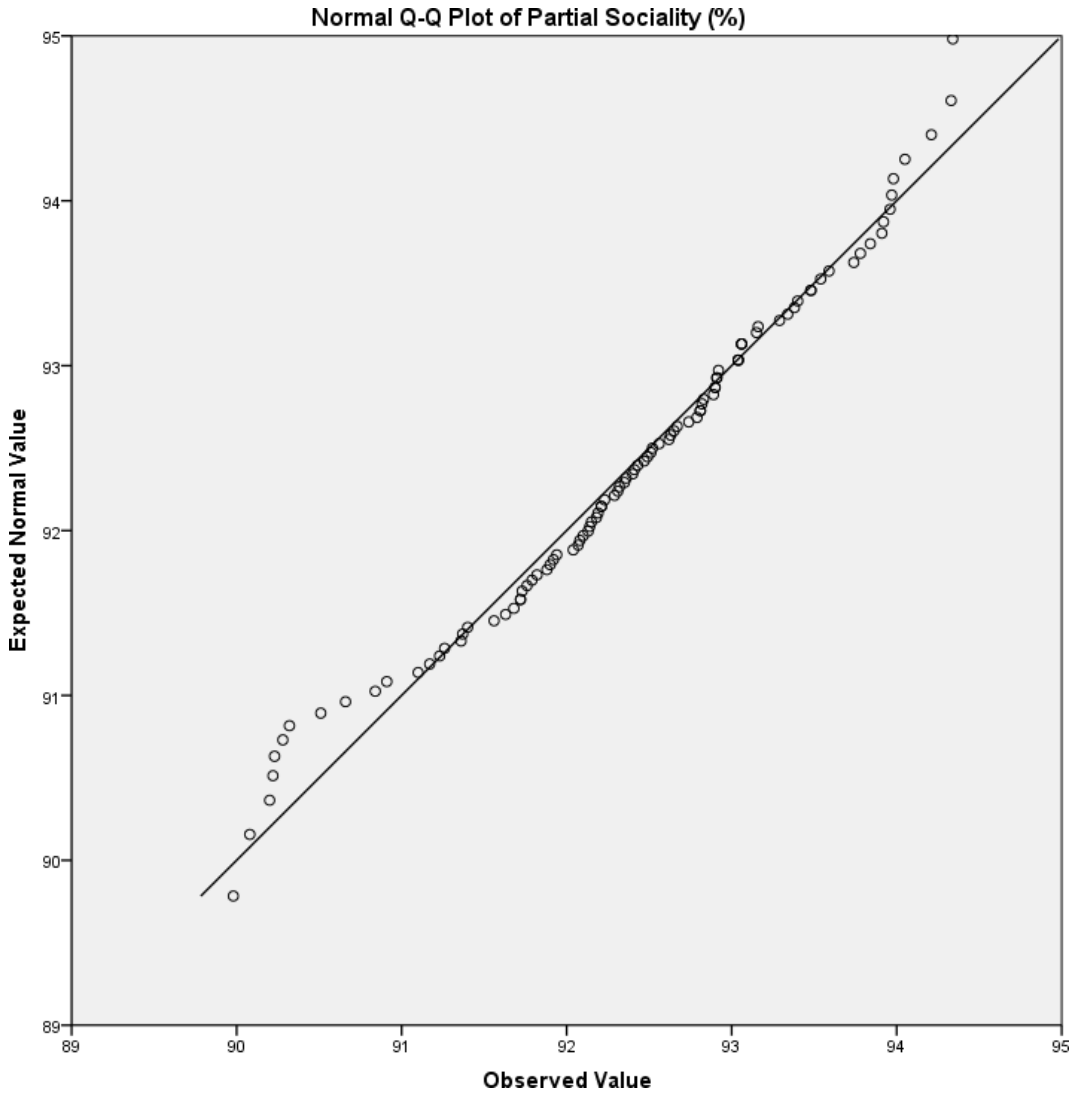

Figure B.13: Q-Q plot of average partial group sociality (%) for the squeezing corridor scenario with groups of two using SGN method.

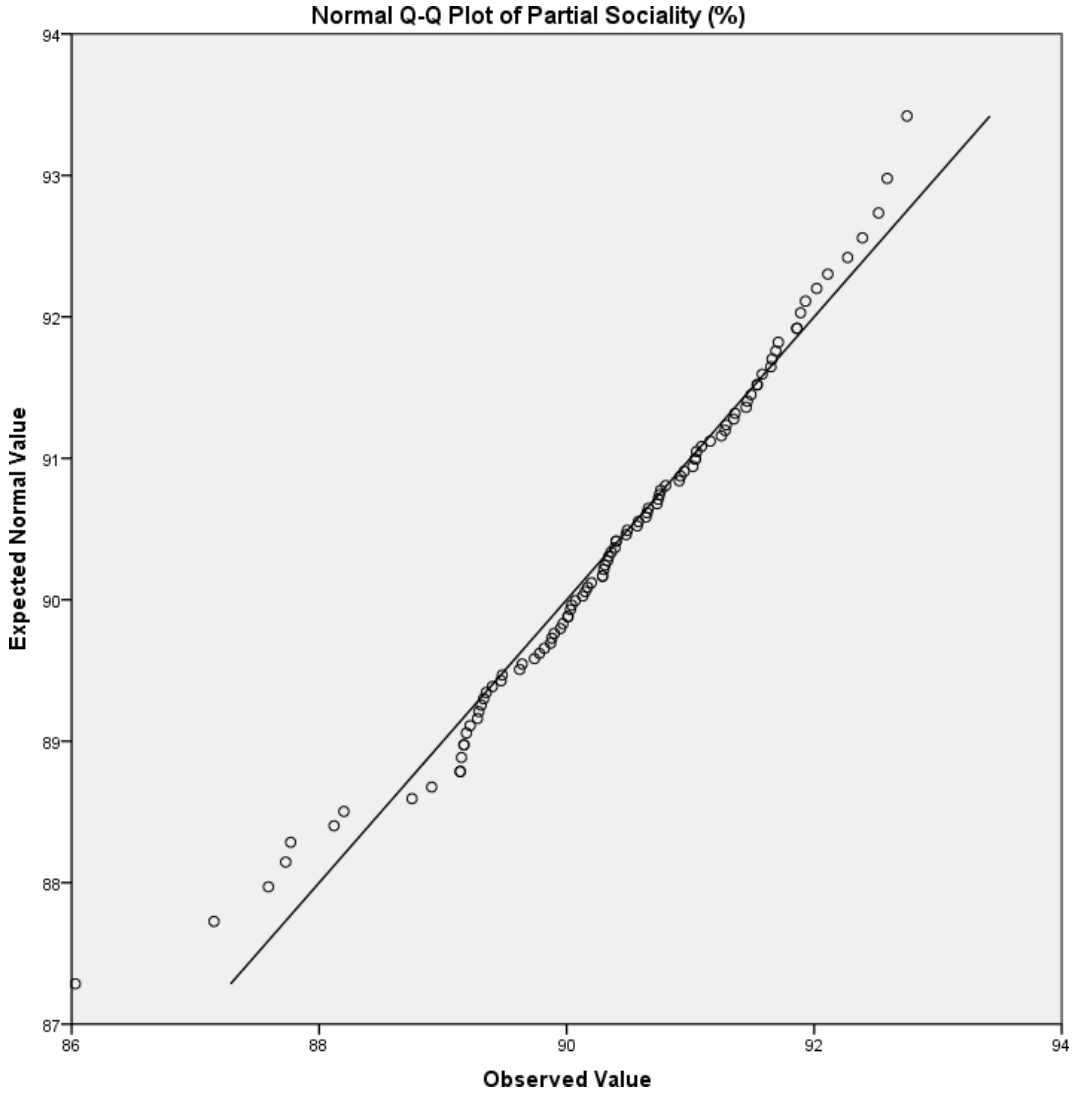

Figure B.14: Q-Q plot of average partial group sociality (%) for the squeezing corridor scenario with groups of two using SGN method.

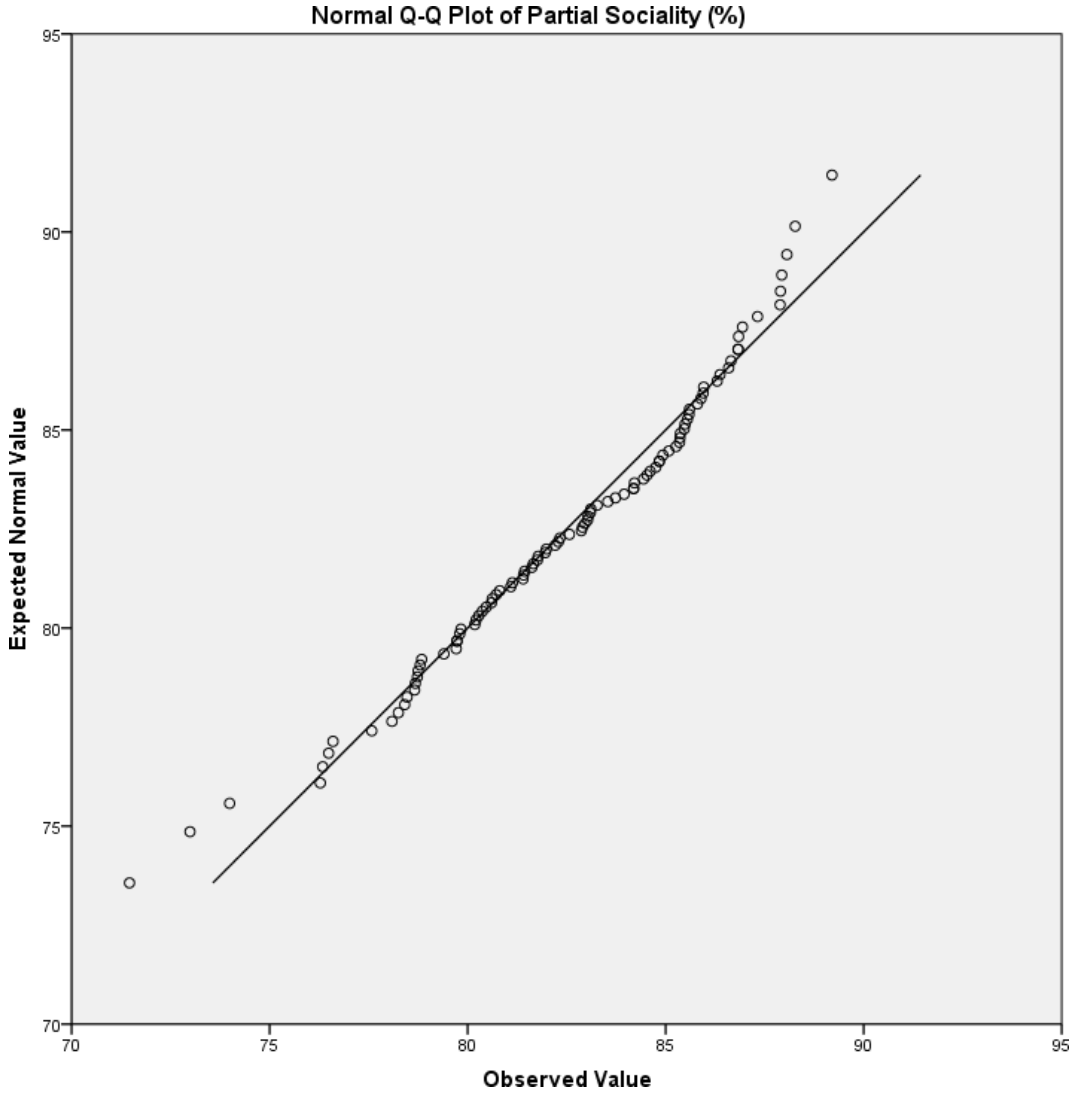

Figure B.15: Q-Q plot of average partial group sociality (%) for the squeezing corridor scenario with groups of three using SGN method.

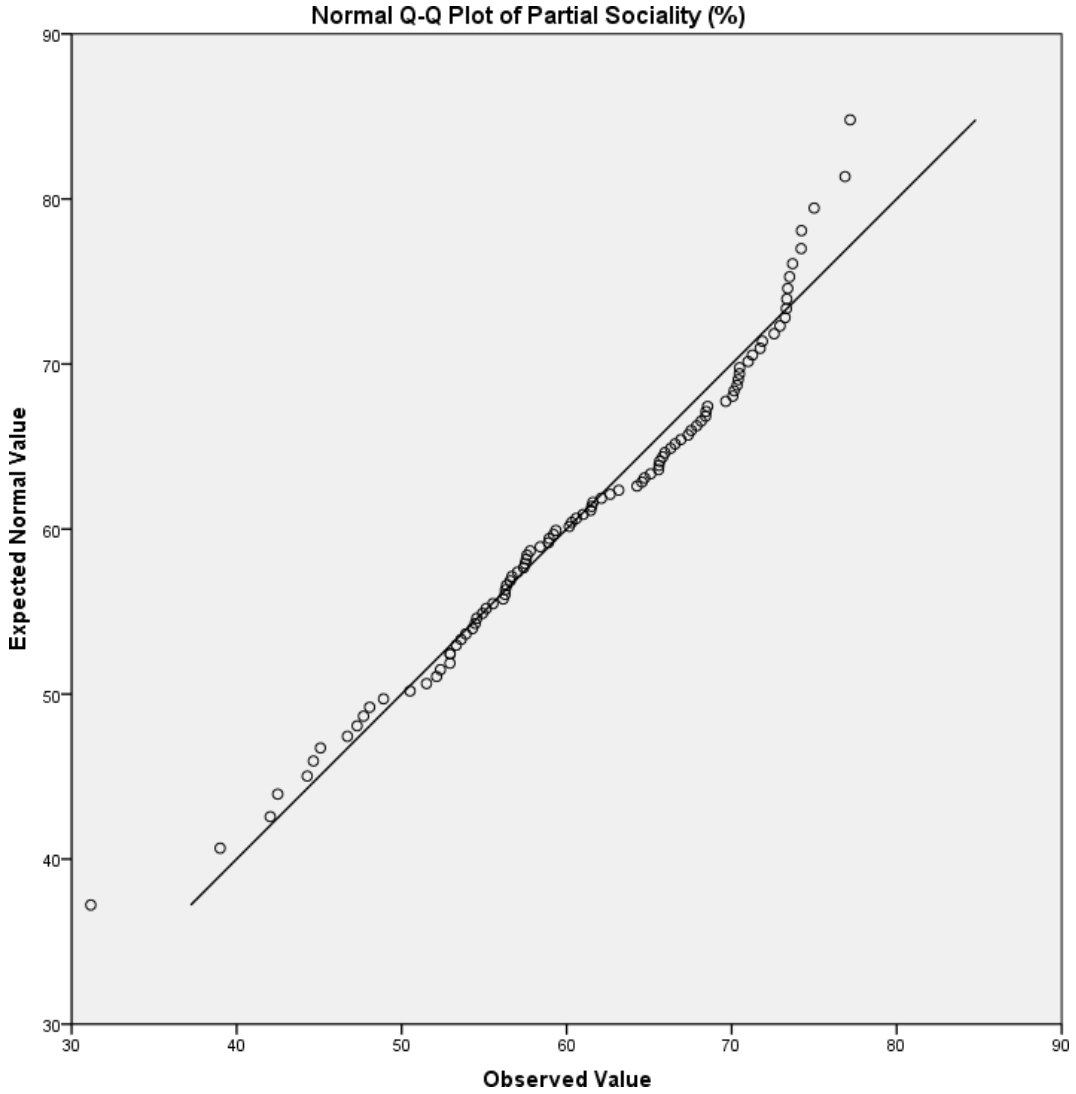

Figure B.16: Q-Q plot of average partial group sociality (%) for the squeezing corridor scenario with groups of three using the method of Moussaïd et al.

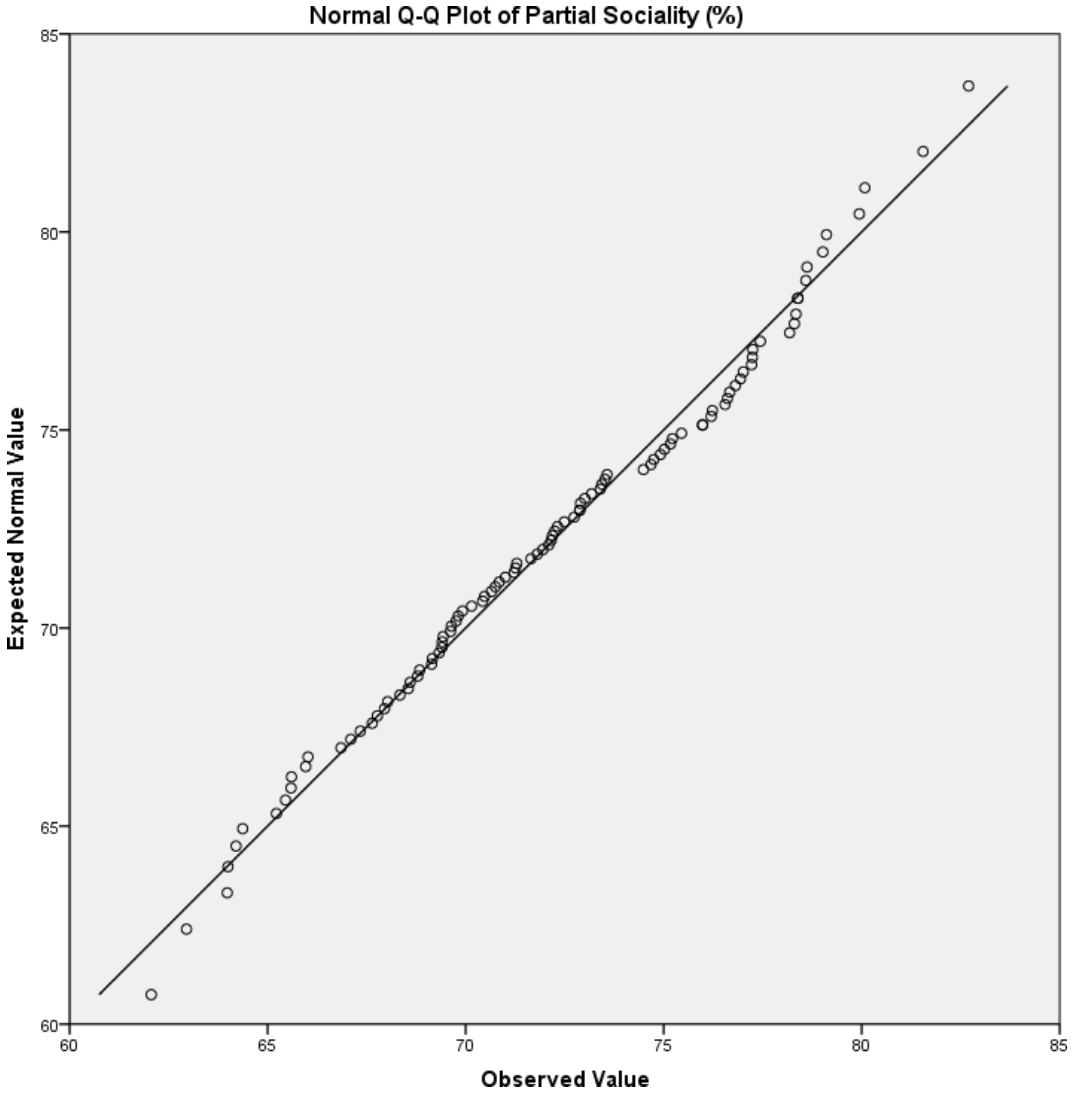

Figure B.17: Q-Q plot of average partial group sociality (%) for the squeezing corridor scenario with groups of four using SGN method.

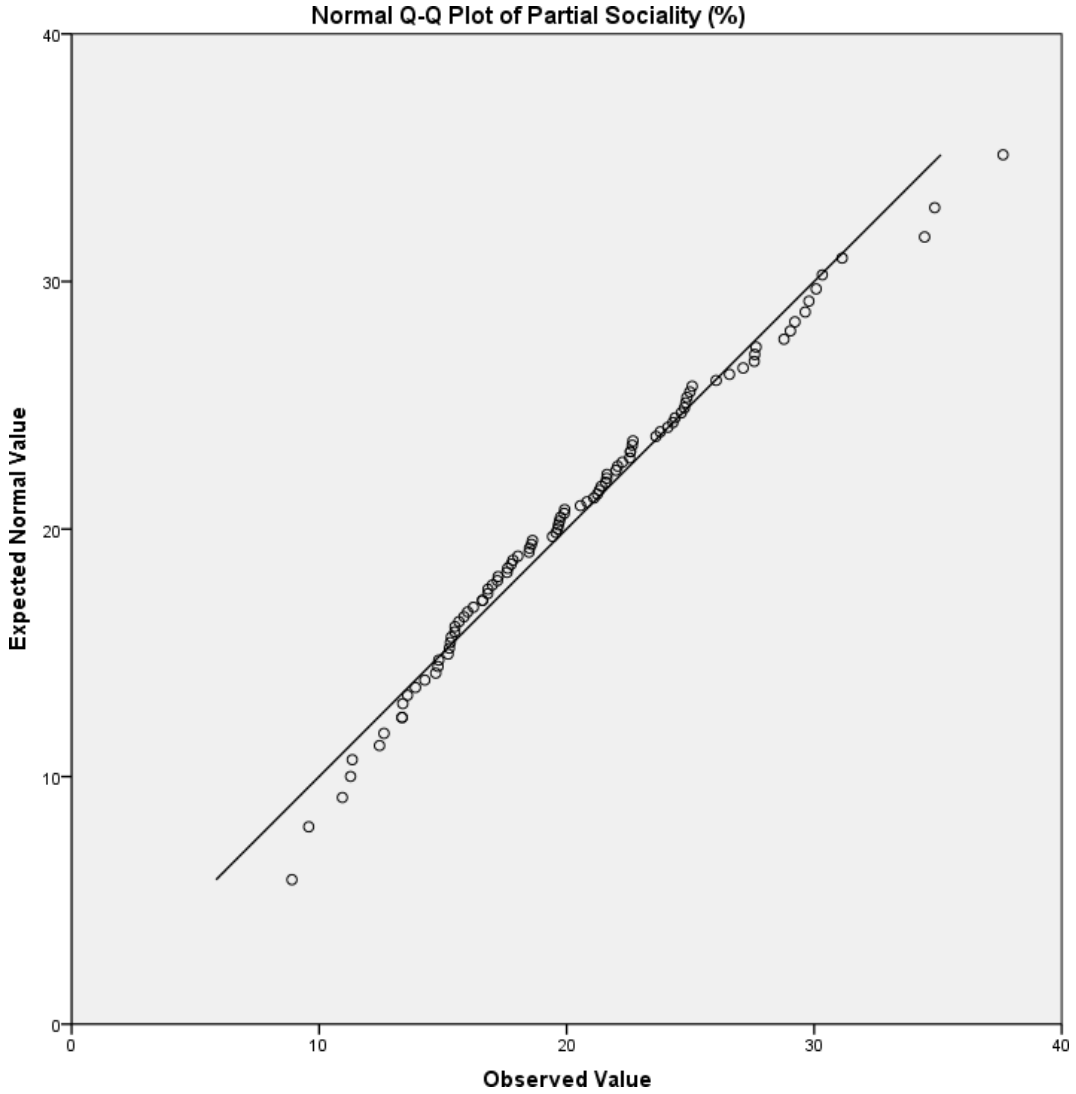

Figure B.18: Q-Q plot of average partial group sociality (%) for the squeezing corridor scenario with groups of four using the method of Moussaïd et al.

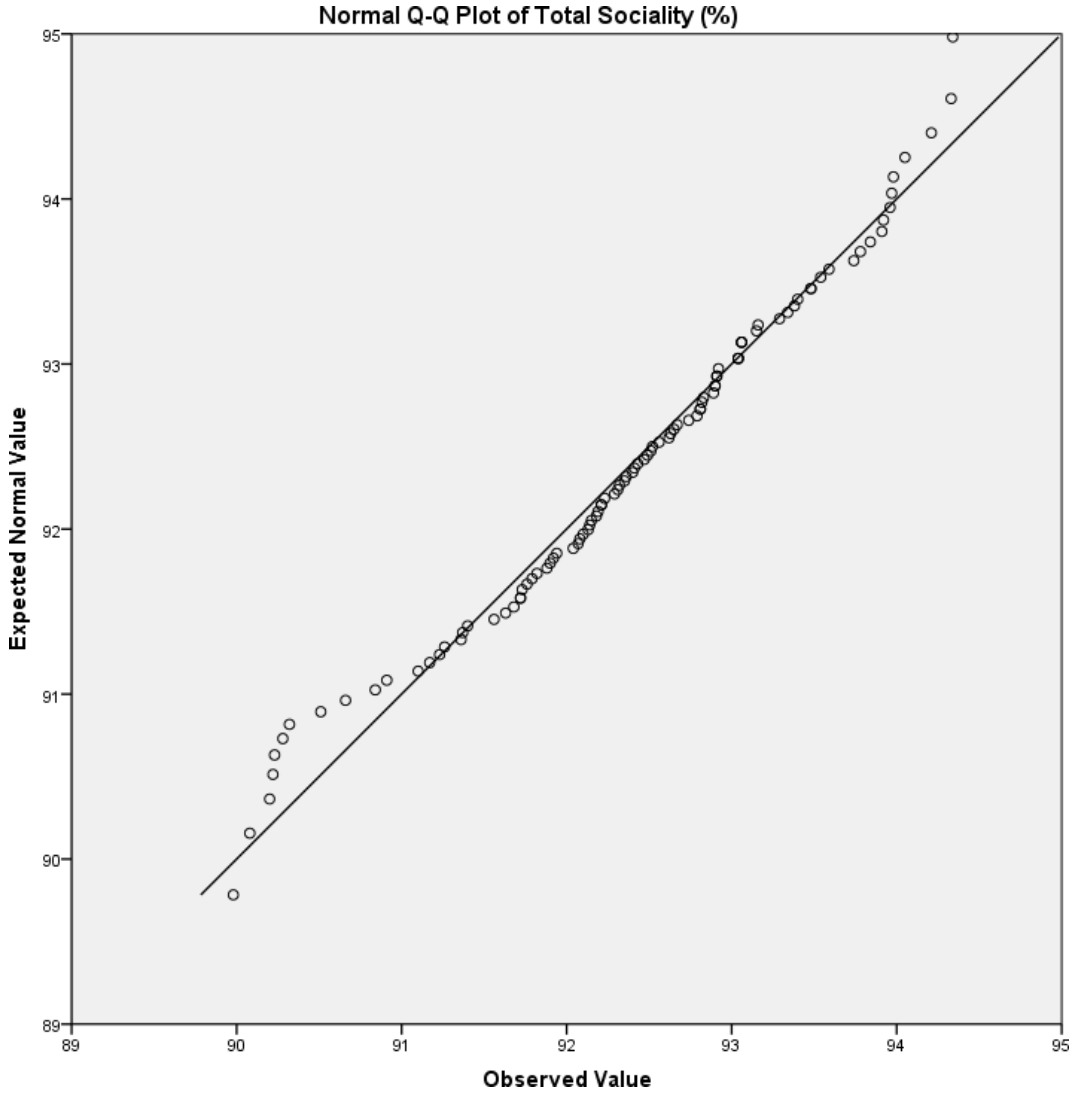

Figure B.19: Q-Q plot of average total group sociality (%) for the squeezing corridor scenario with groups of two using SGN method.

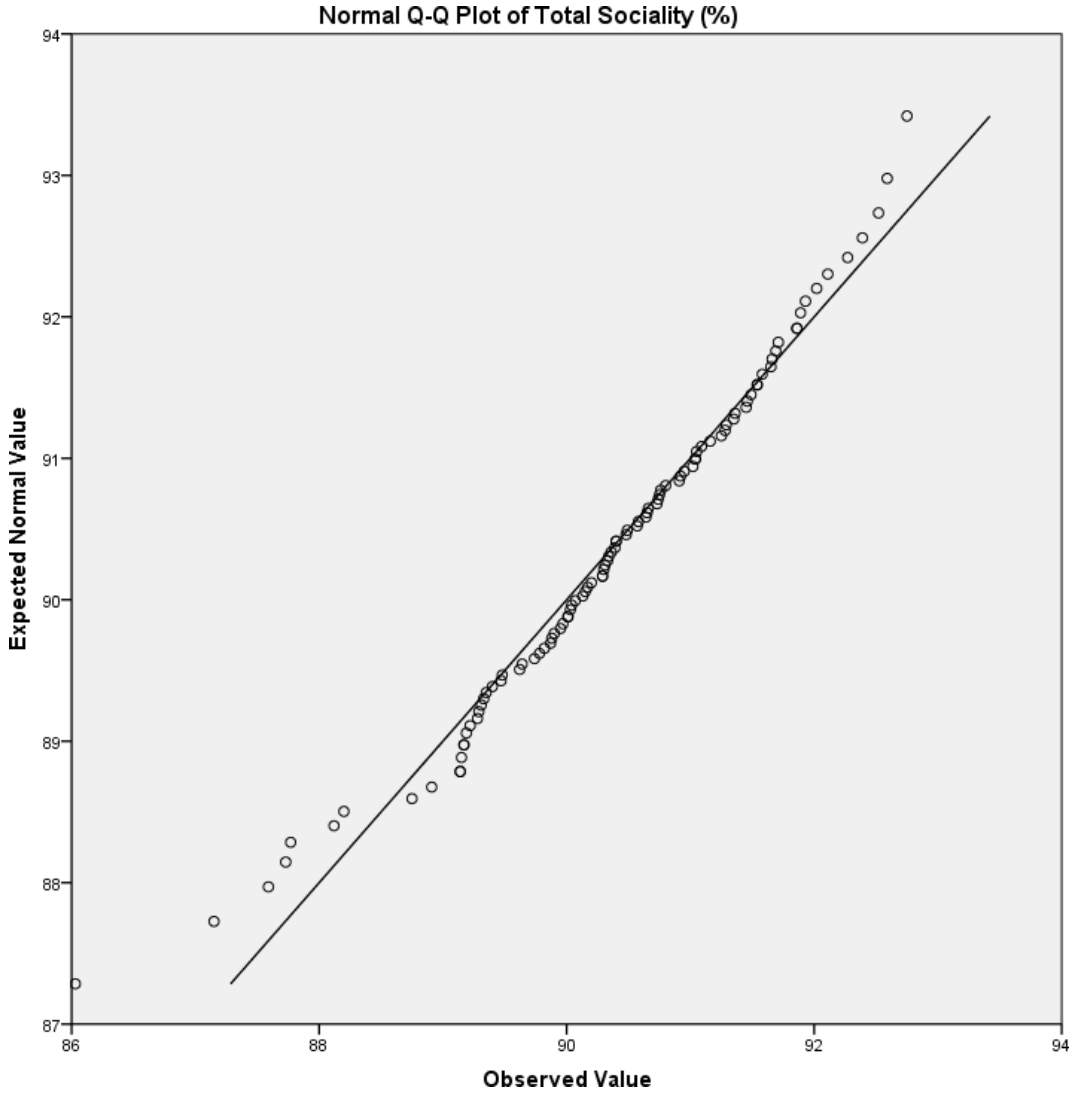

Figure B.20: Q-Q plot of average totall group sociality (%) for the squeezing corridor scenario with groups of two using SGN method.

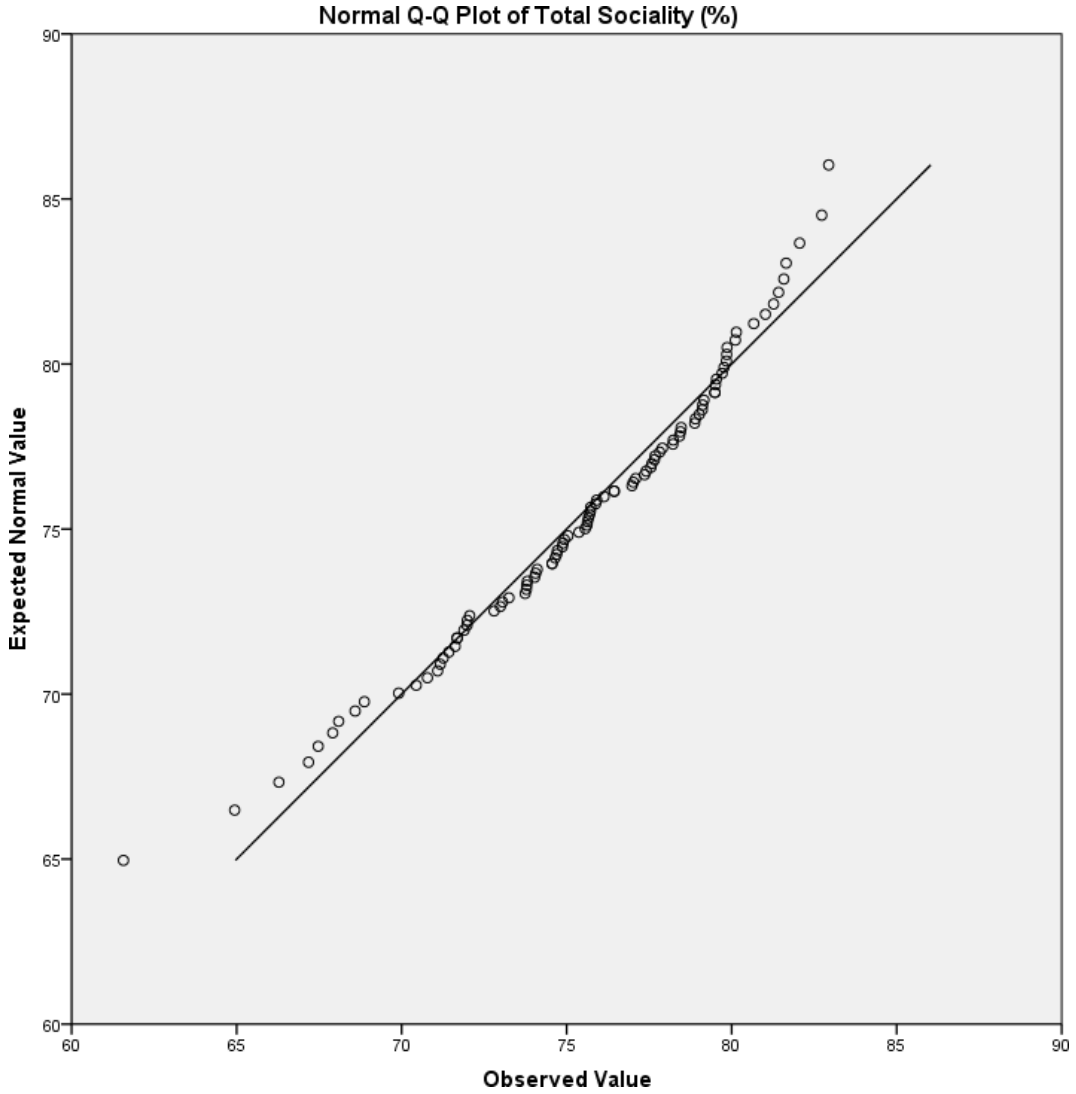

Figure B.21: Q-Q plot of average total group sociality (%) for the squeezing corridor scenario with groups of three using SGN method.

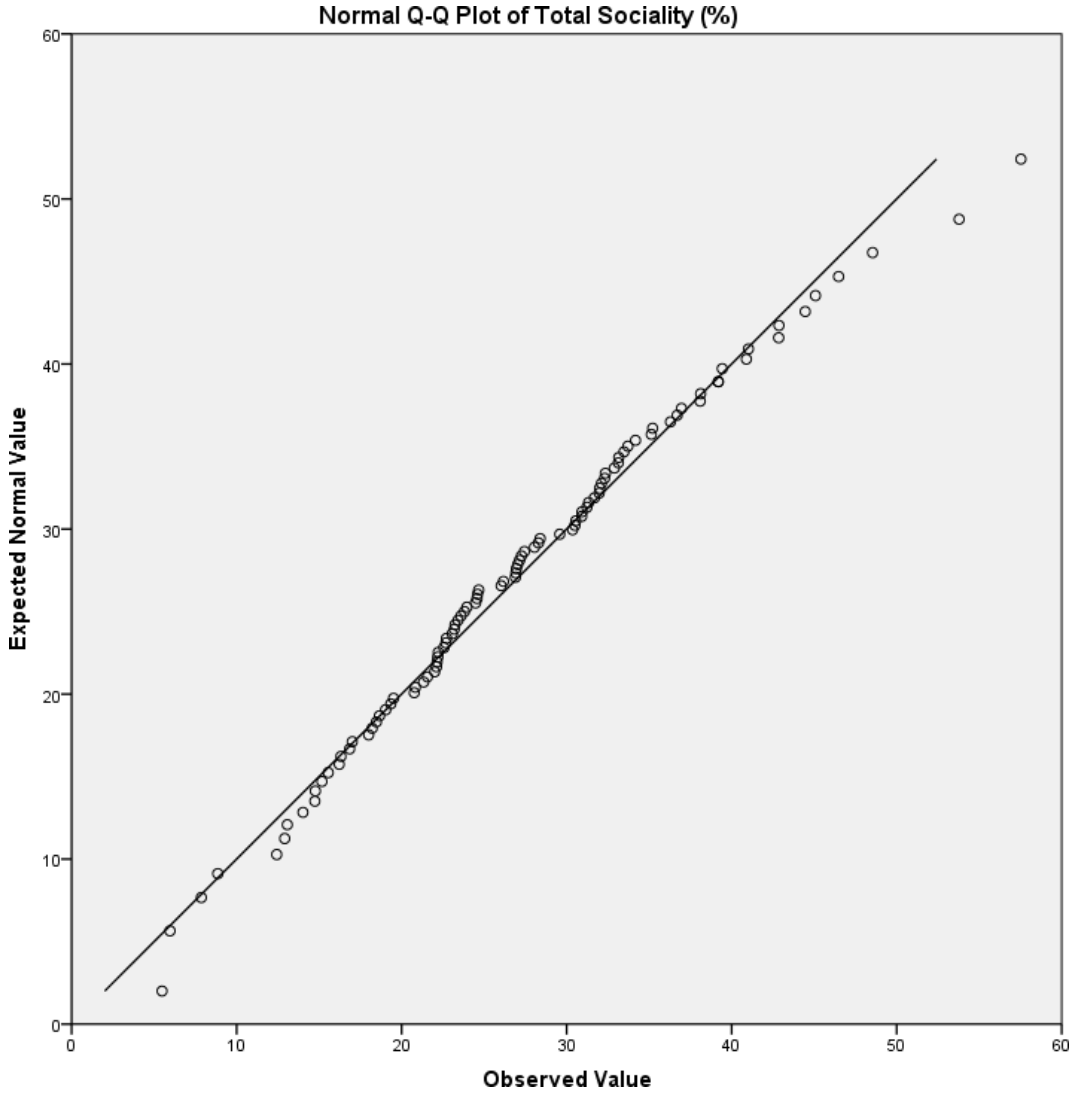

Figure B.22: Q-Q plot of average total group sociality (%) for the squeezing corridor scenario with groups of three using the method of Moussaïd et al.

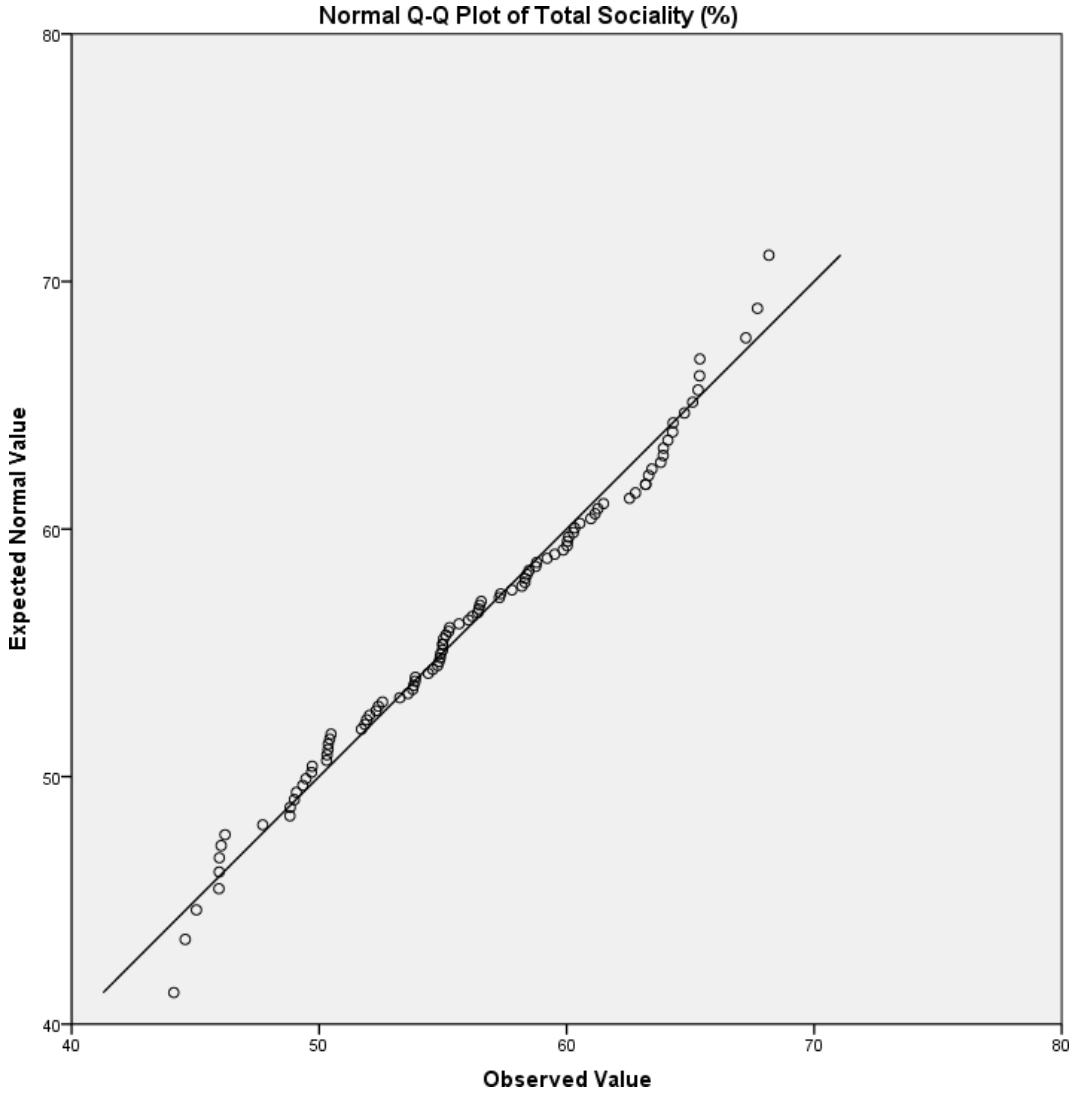

Figure B.23: Q-Q plot of average total group sociality (%) for the squeezing corridor scenario with groups of four using SGN method.

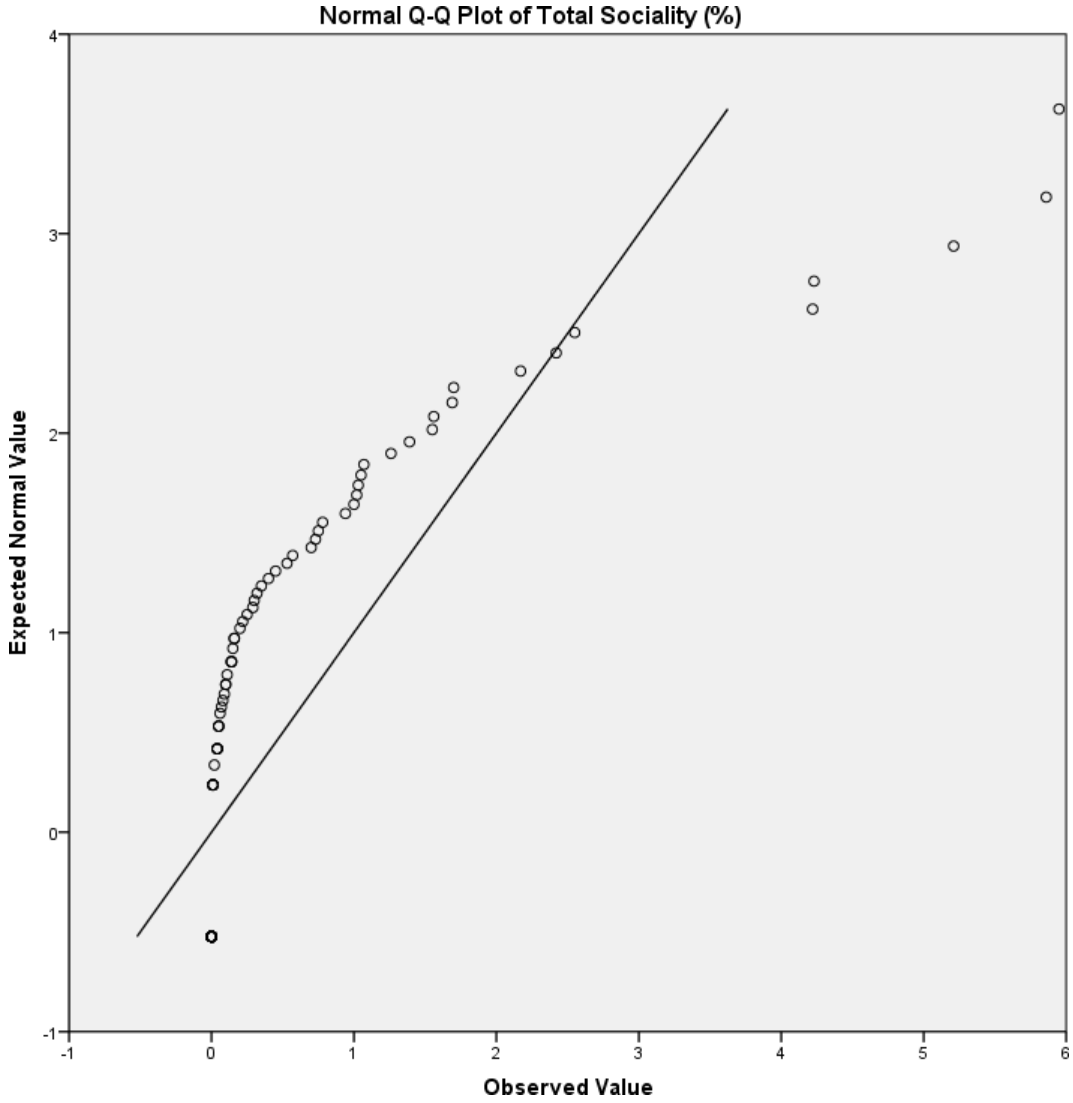

Figure B.24: Q-Q plot of average total group sociality (%) for the squeezing corridor scenario with groups of four using the method of Moussaïd et al.

B.3 Corners

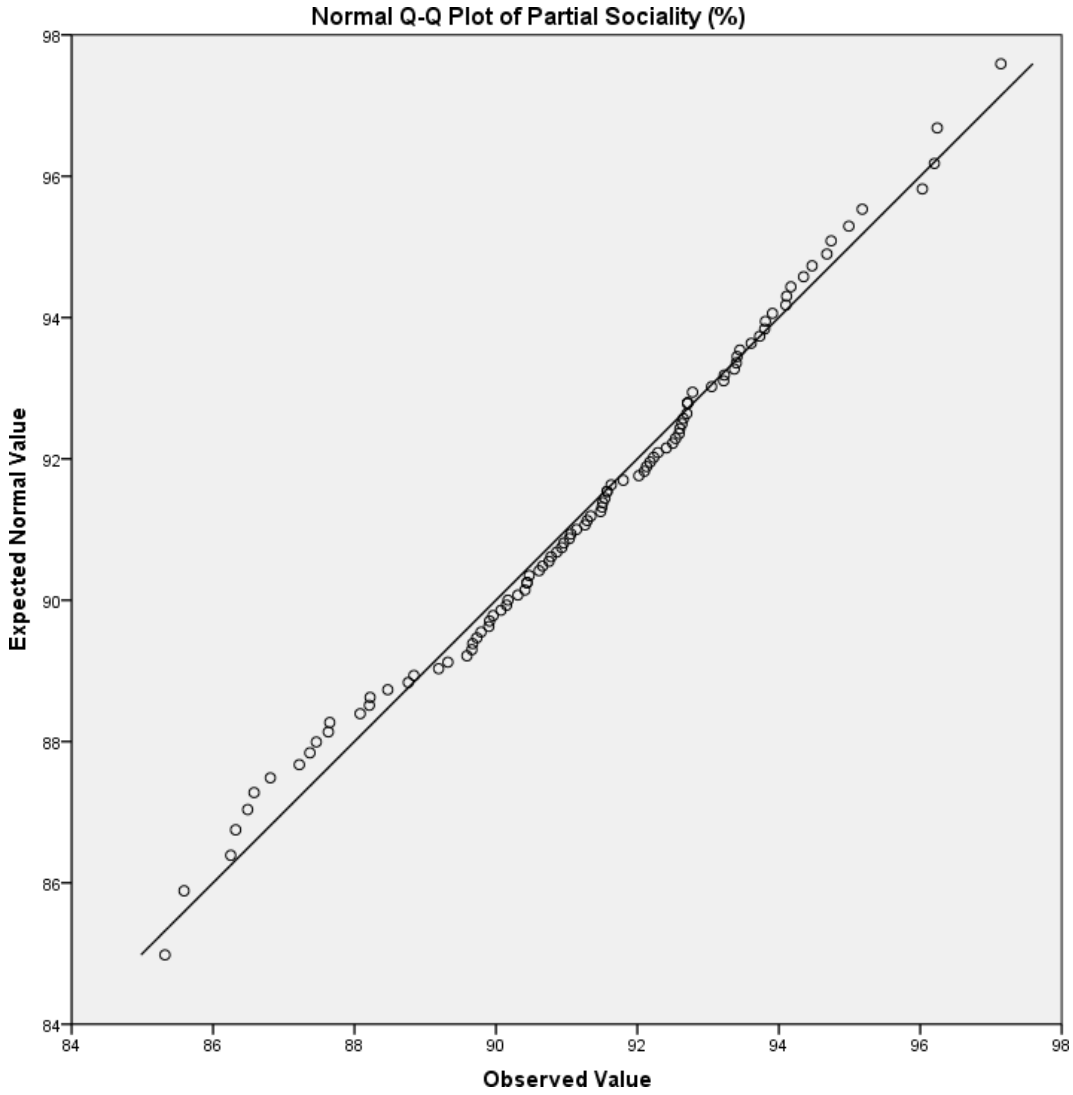

Figure B.25: Q-Q plot of average partial group sociality (%) for the corners scenario with groups of two using SGN method.

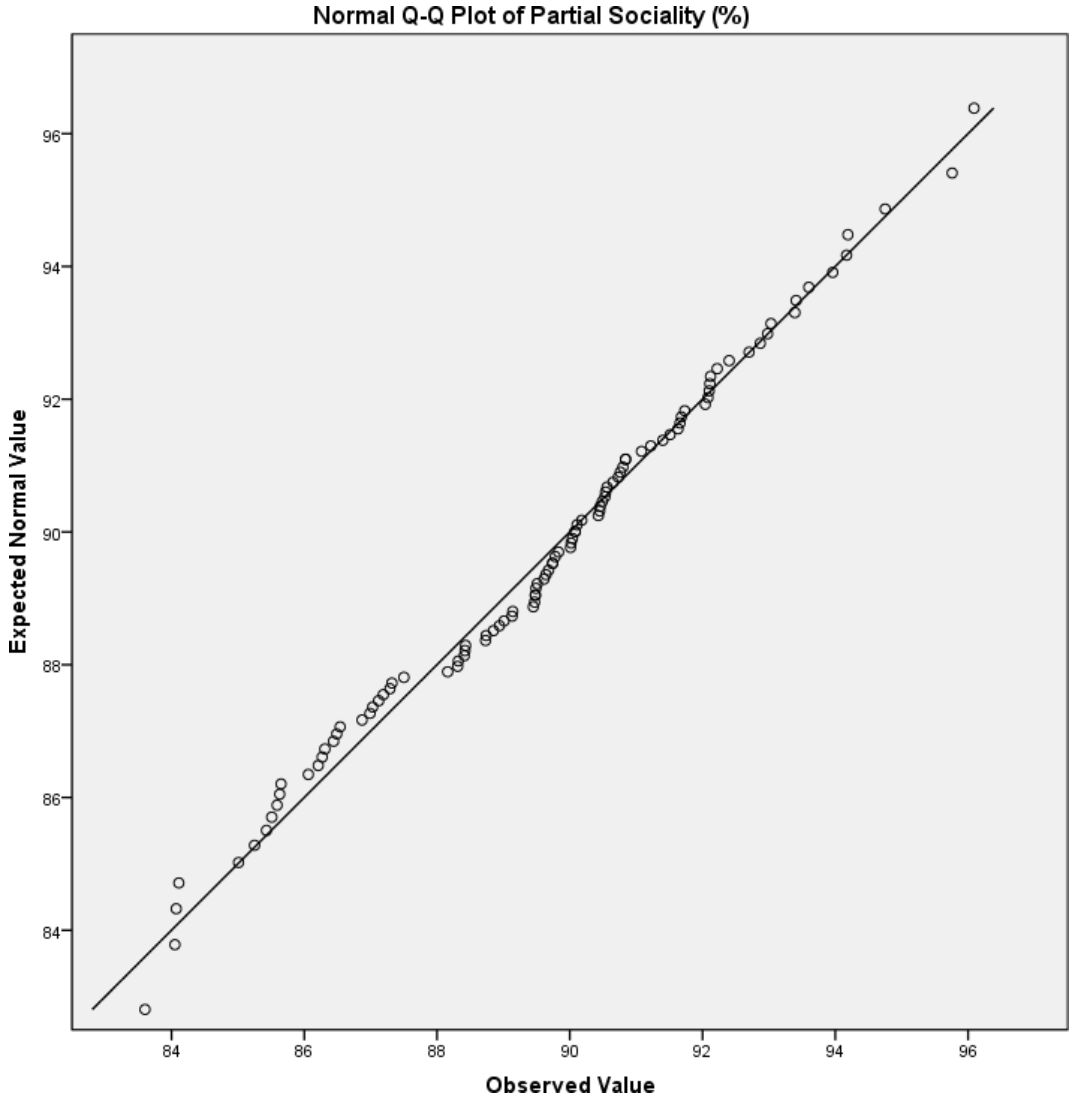

Figure B.26: Q-Q plot of average partial group sociality (%) for the corners scenario with groups of two using SGN method.

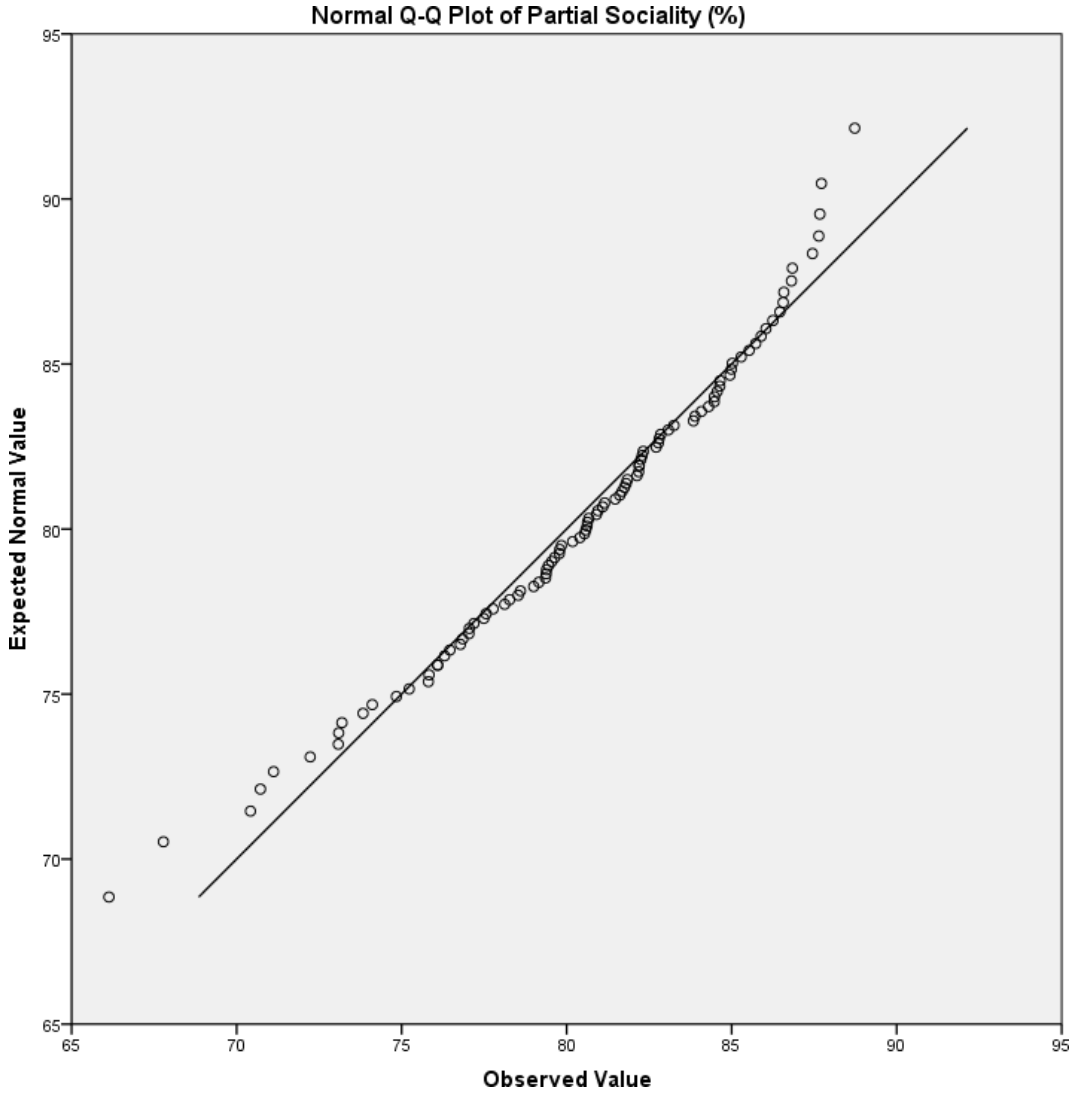

Figure B.27: Q-Q plot of average partial group sociality (%) for the corners scenario with groups of three using SGN method.

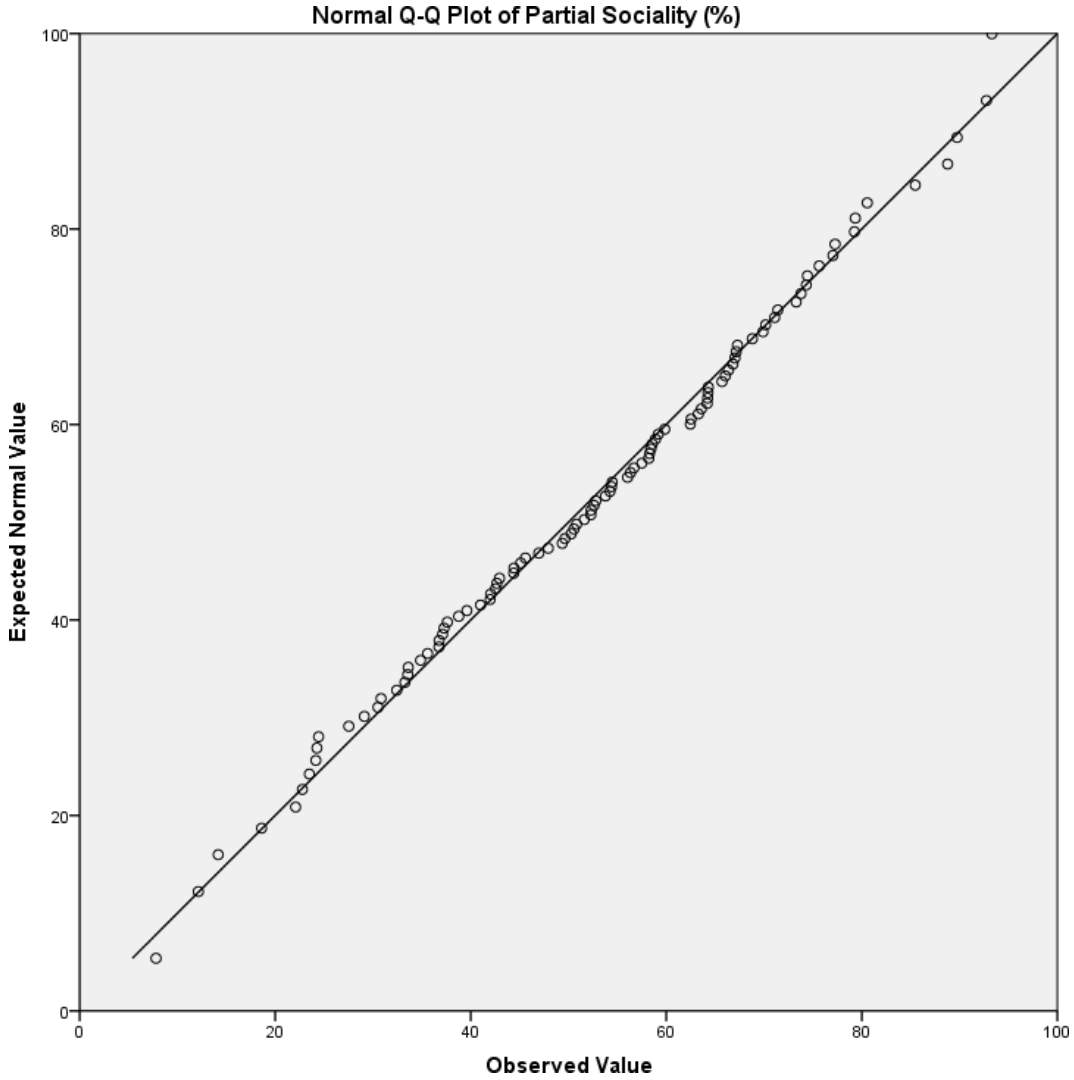

Figure B.28: Q-Q plot of average partial group sociality (%) for the corners scenario with groups of three using the method of Moussaïd et al.

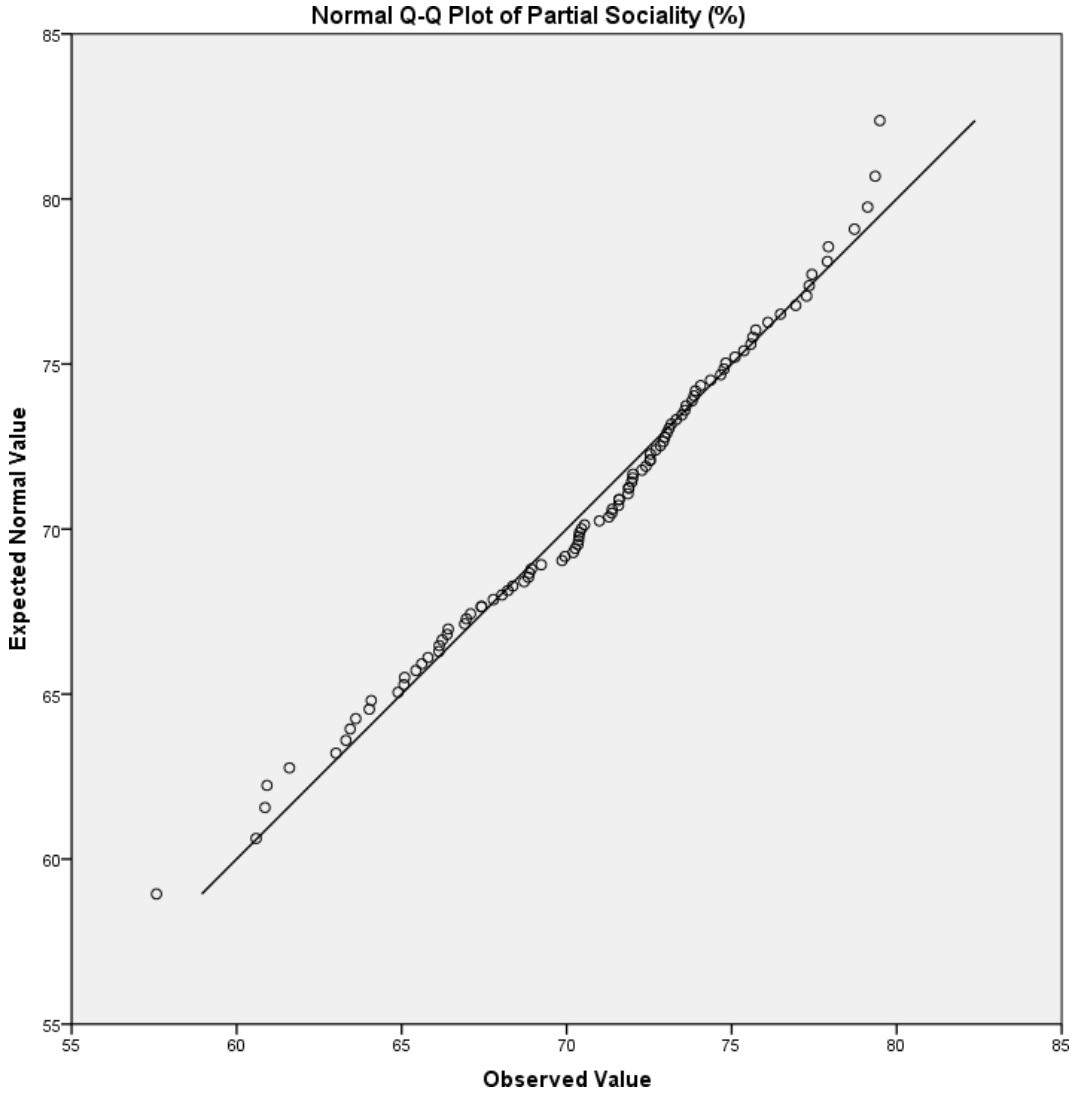

Figure B.29: Q-Q plot of average partial group sociality (%) for the corners scenario with groups of four using SGN method.

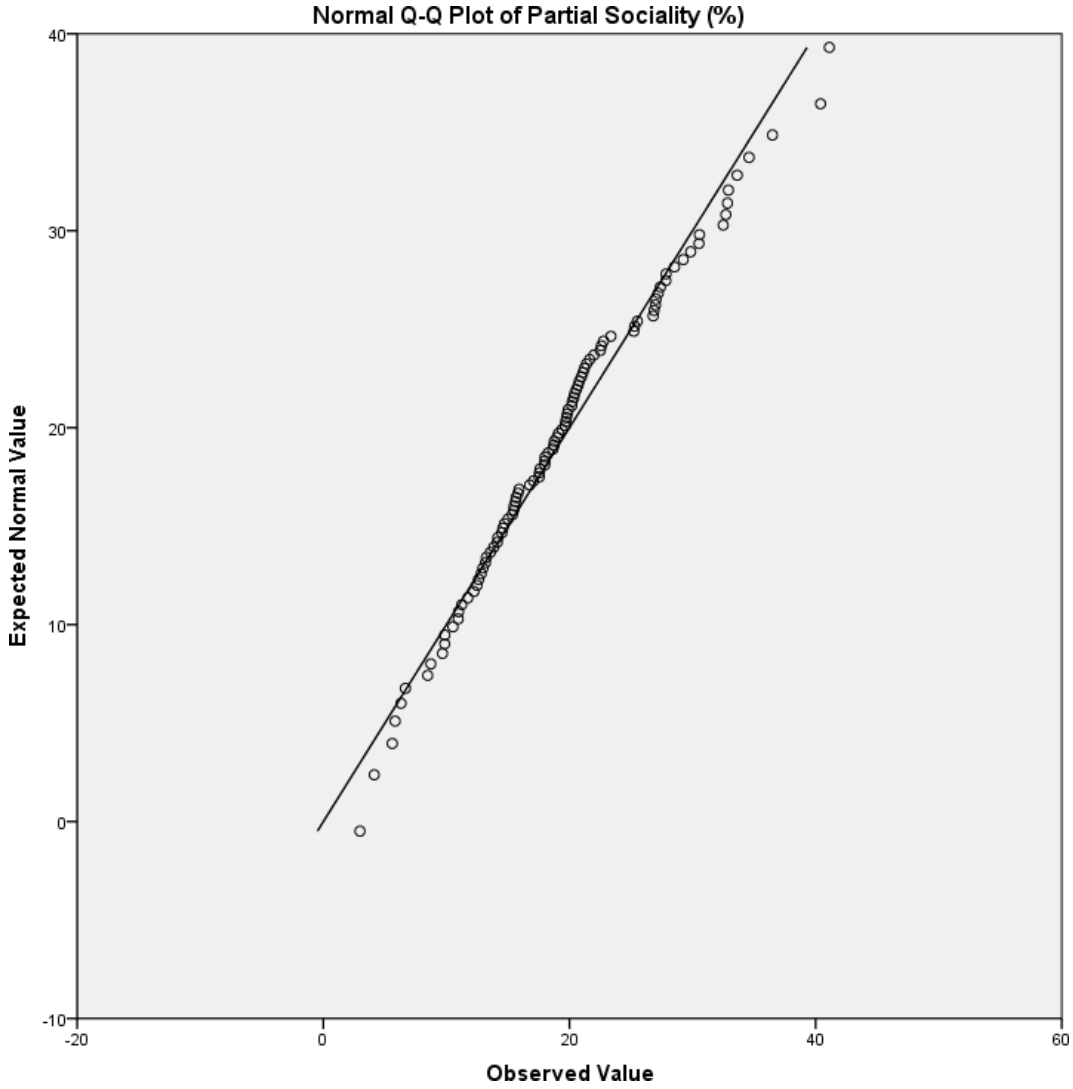

Figure B.30: Q-Q plot of average partial group sociality (%) for the corners scenario with groups of four using the method of Moussaïd et al.

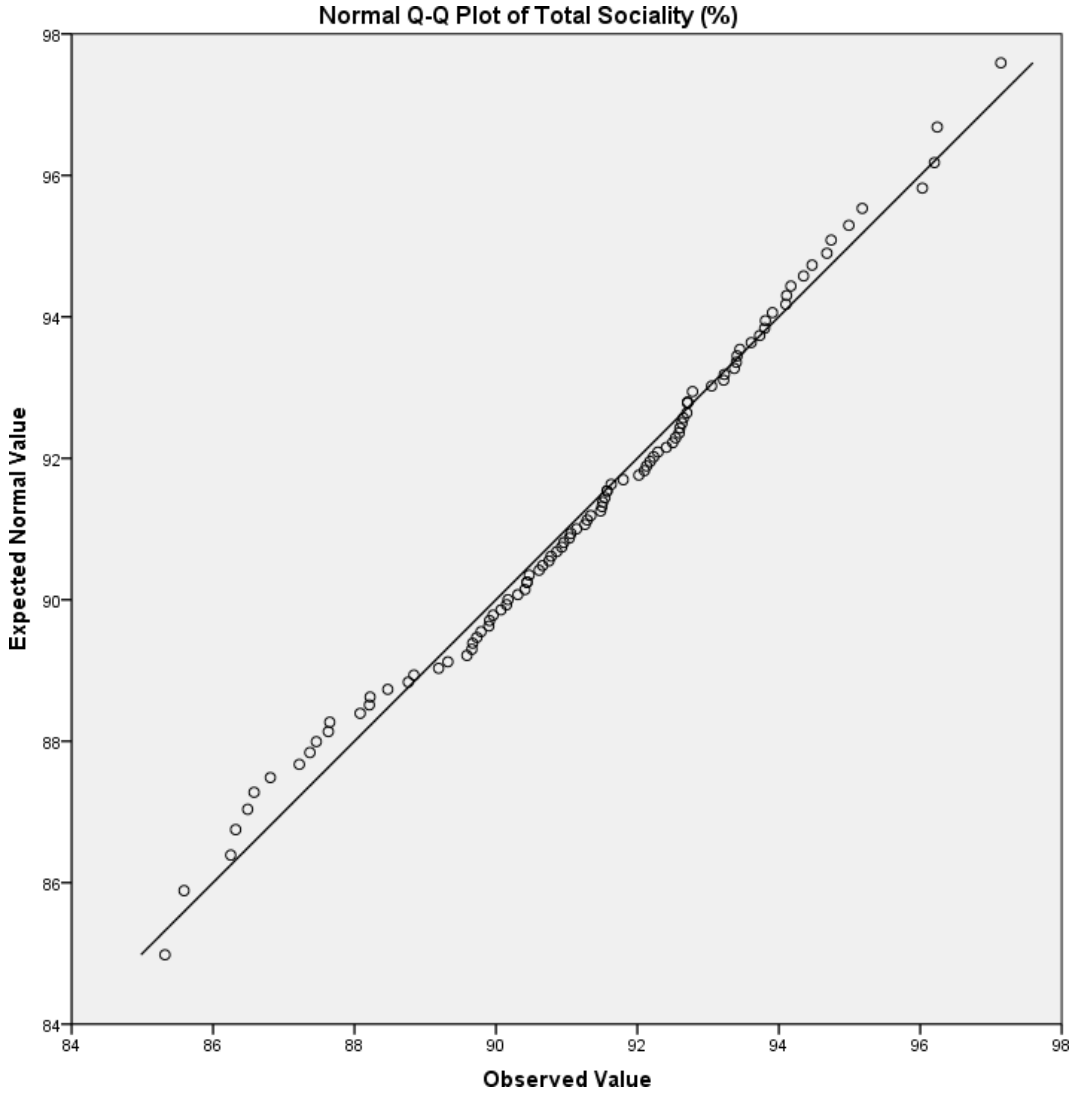

Figure B.31: Q-Q plot of average total group sociality (%) for the corners scenario with groups of two using SGN method.

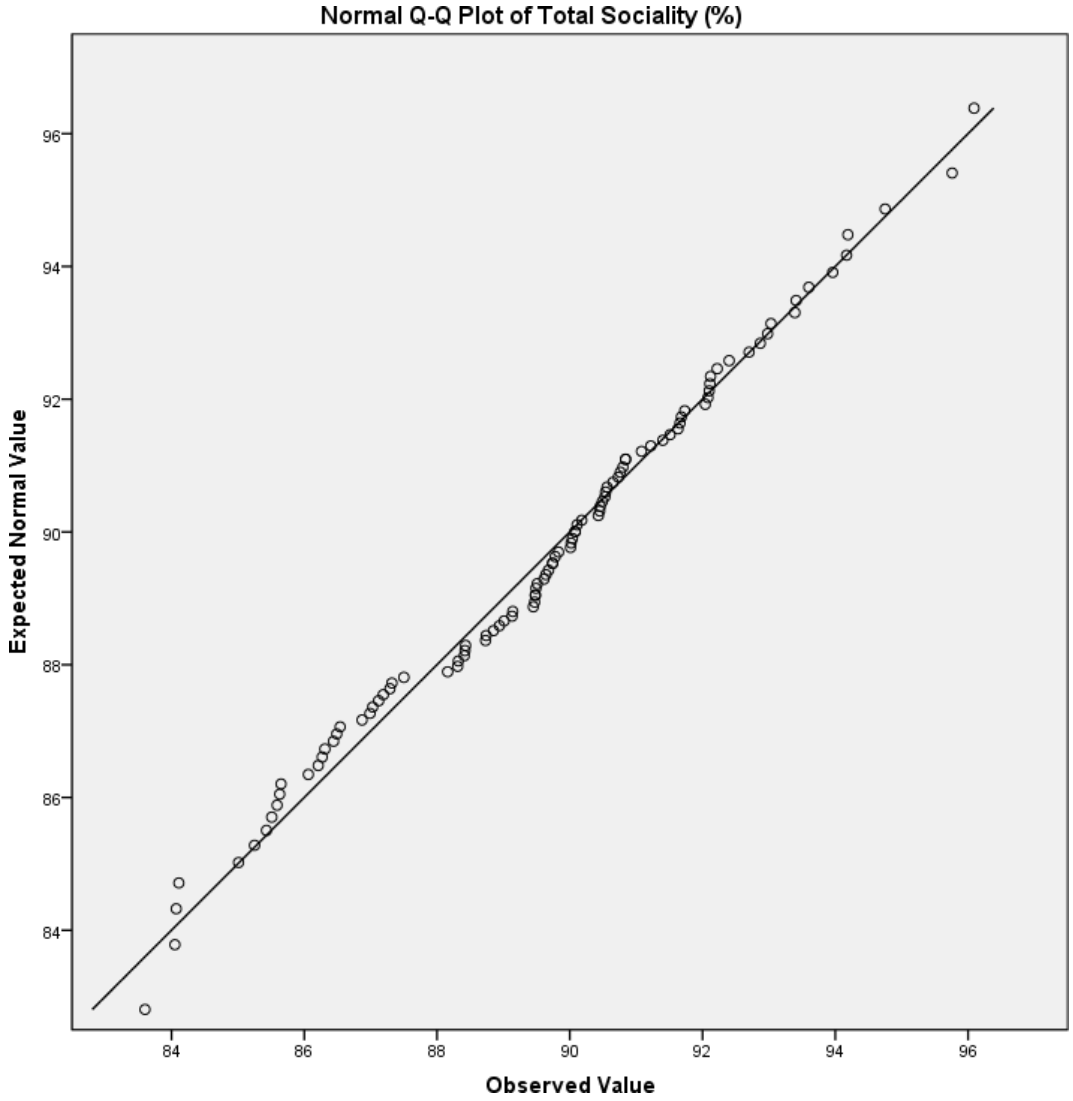

Figure B.32: Q-Q plot of average totall group sociality (%) for the corners scenario with groups of two using SGN method.

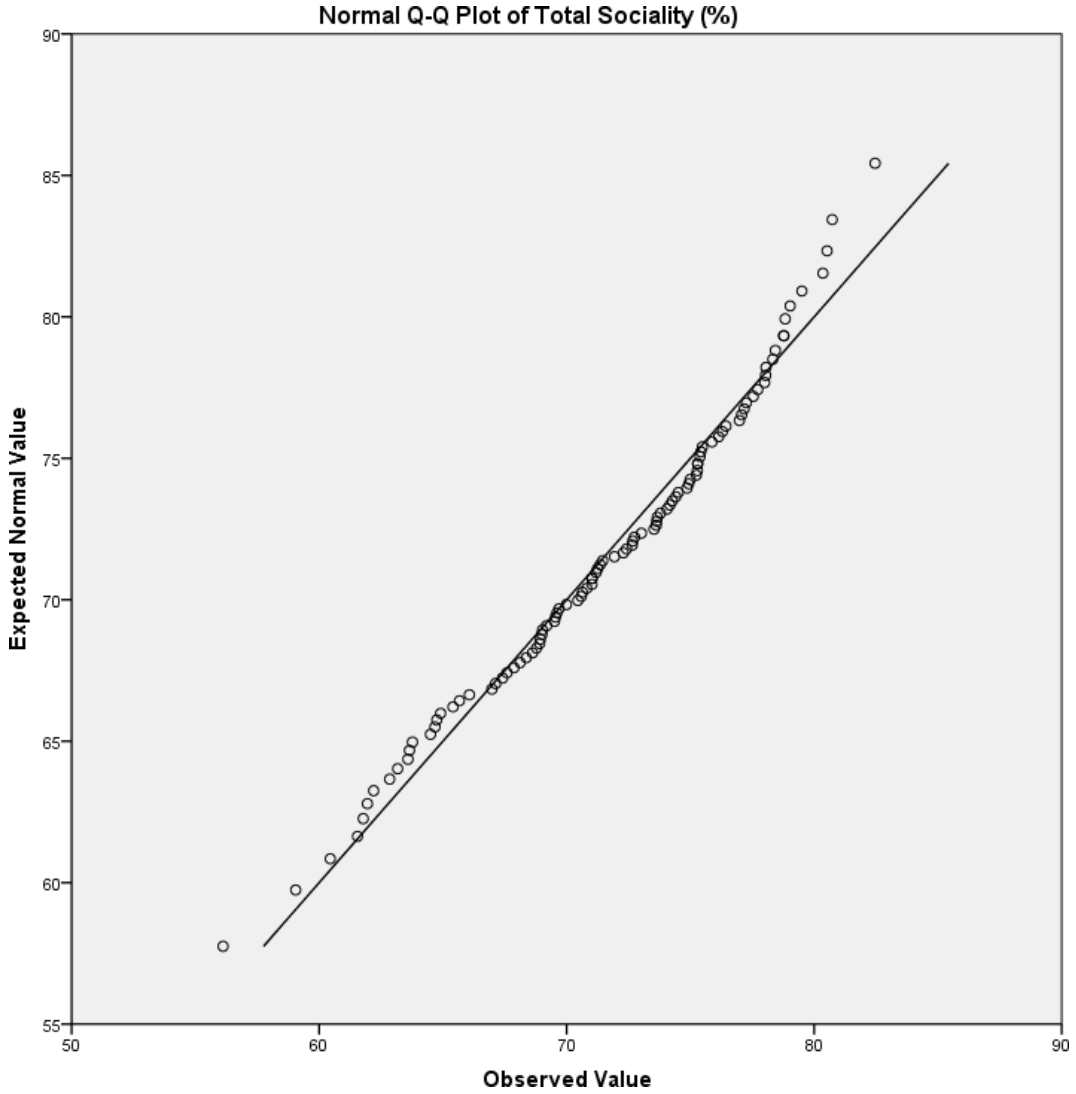

Figure B.33: Q-Q plot of average total group sociality (%) for the corners scenario with groups of three using SGN method.

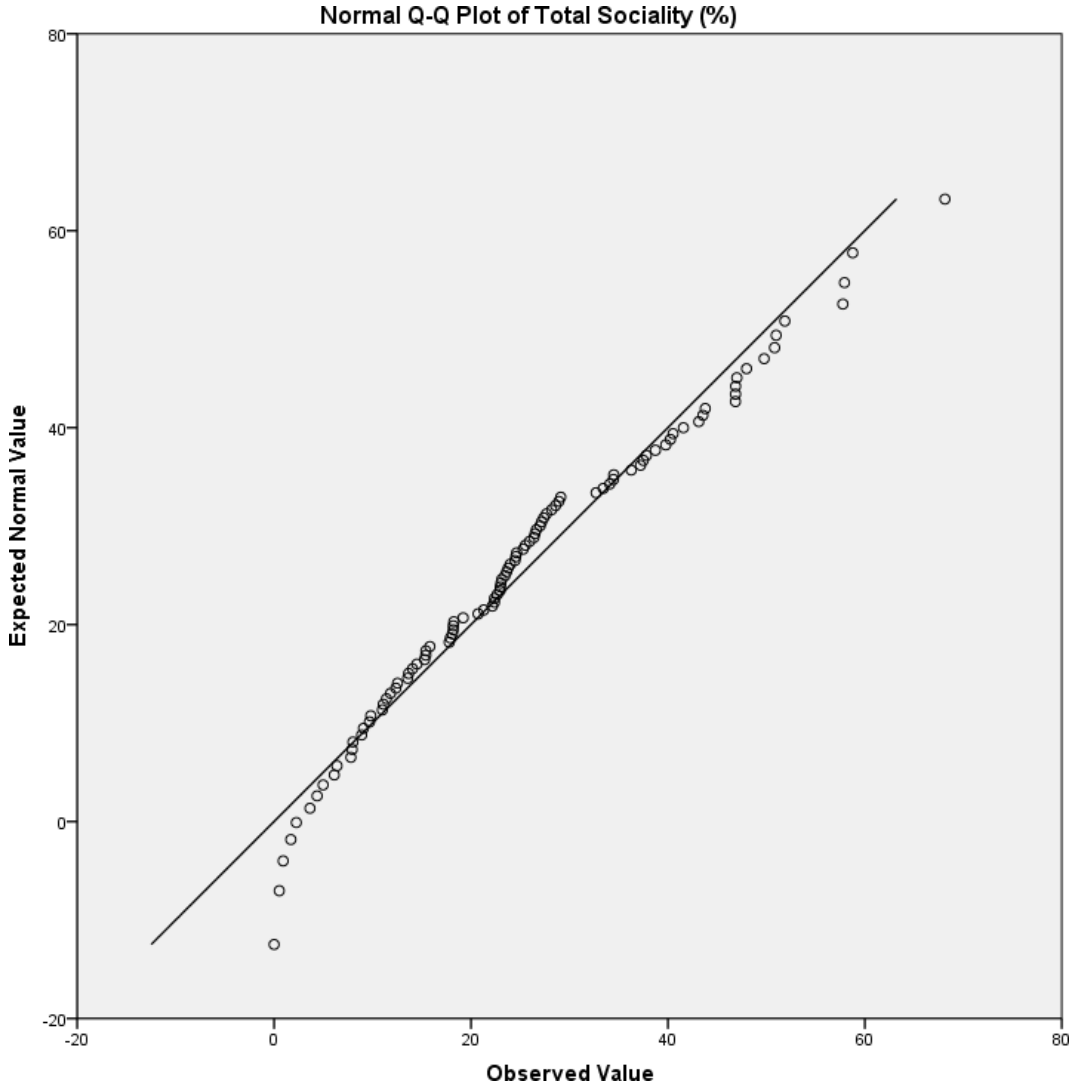

Figure B.34: Q-Q plot of average total group sociality (%) for the corners scenario with groups of three using the method of Moussaïd et al.

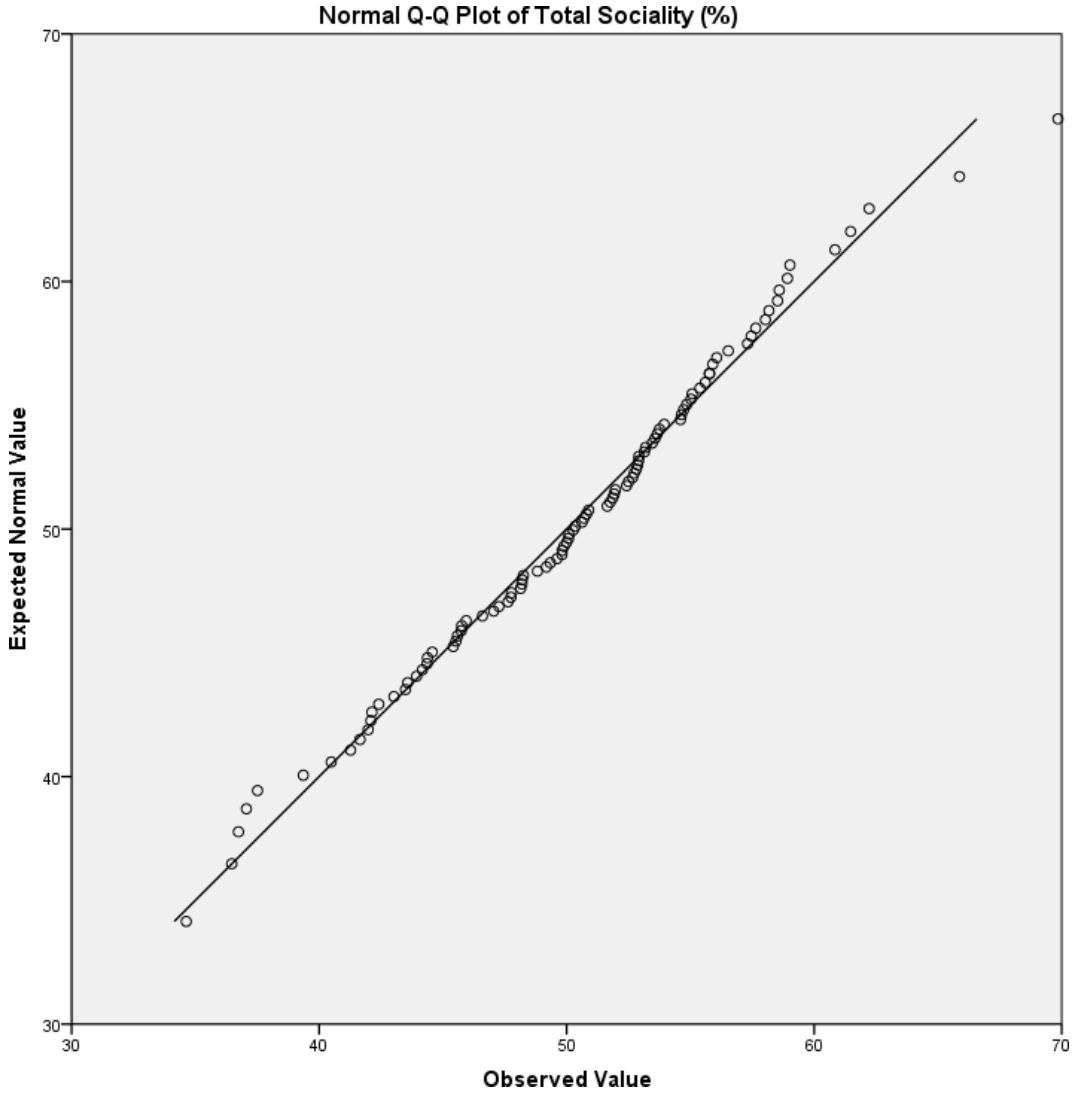

Figure B.35: Q-Q plot of average total group sociality (%) for the corners scenario with groups of four using SGN method.

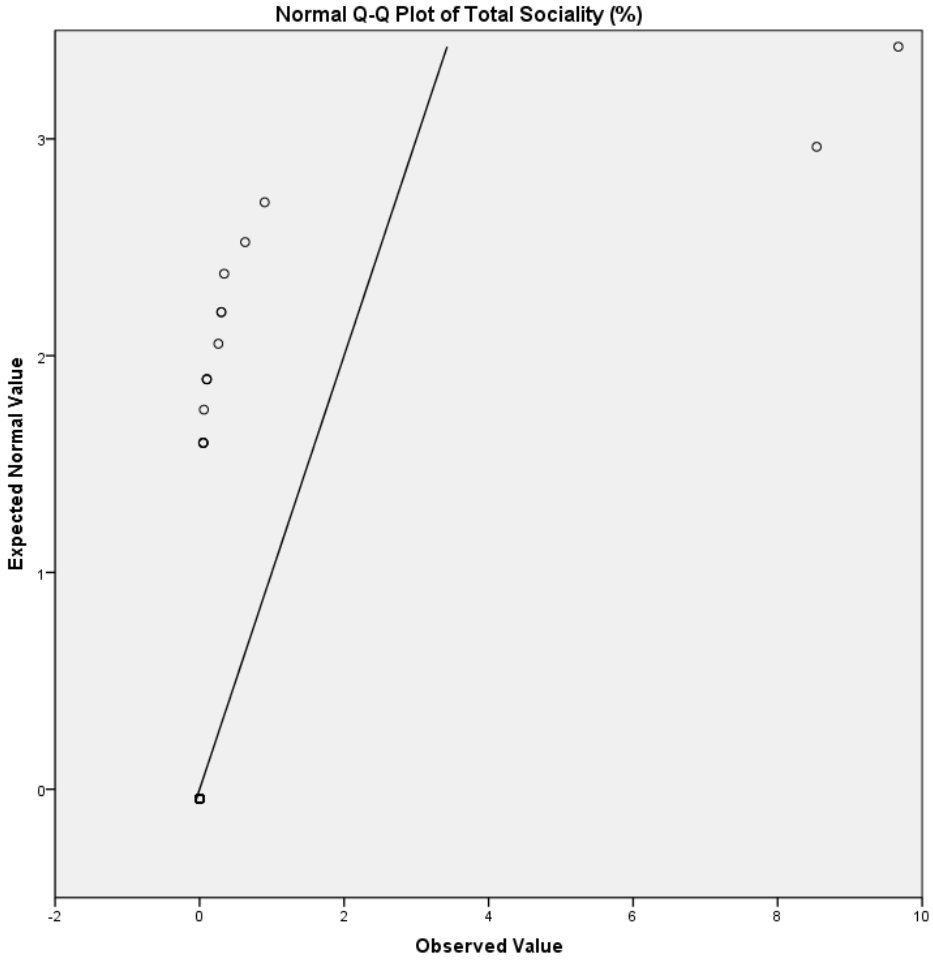

Figure B.36: Q-Q plot of average total group sociality (%) for the corners scenario with groups of four using the method of Moussaïd et al.

B.4 Building evacuation

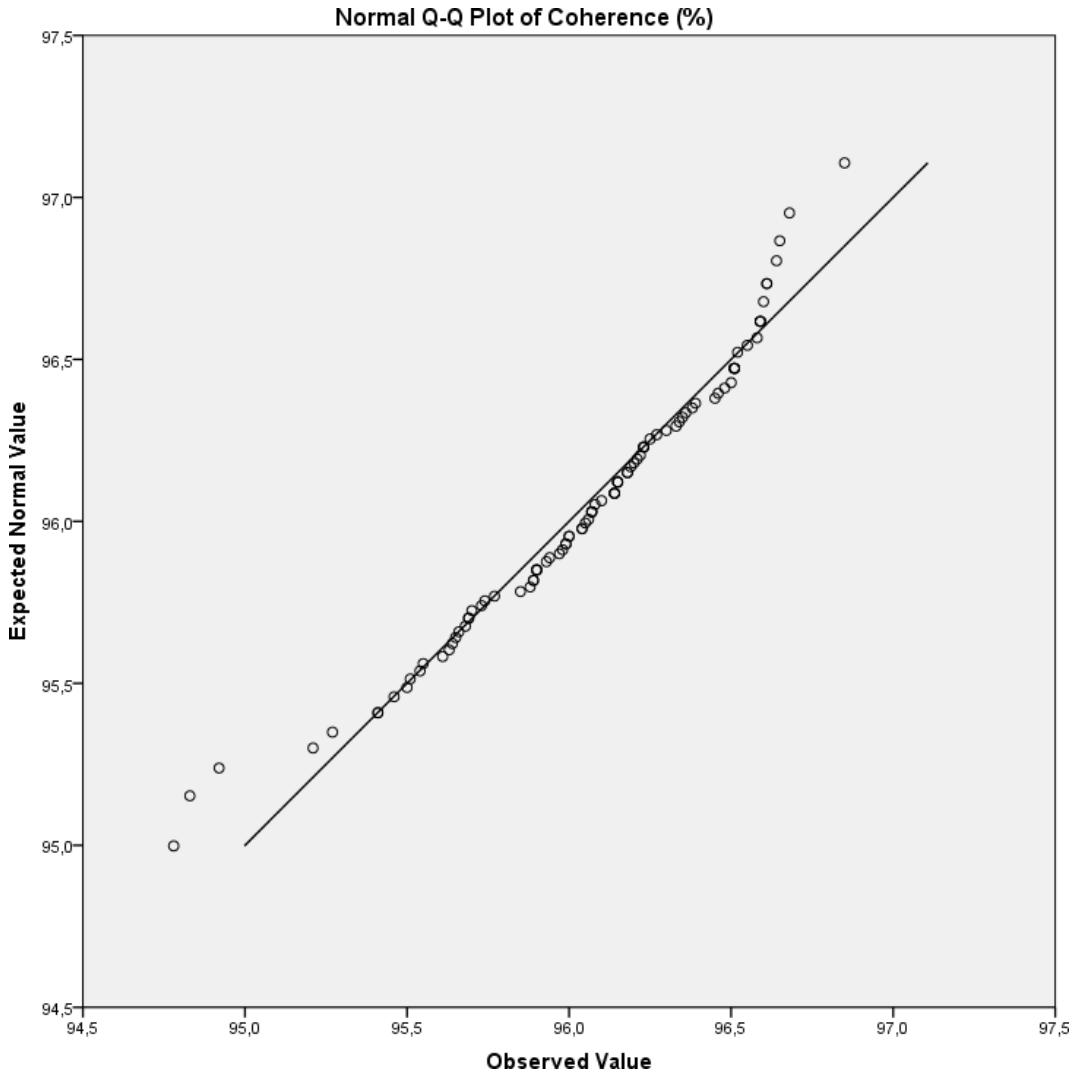

Figure B.37: Q-Q plot of average group coherence (%) for the building-evacuation scenario with groups of two using SGN method.

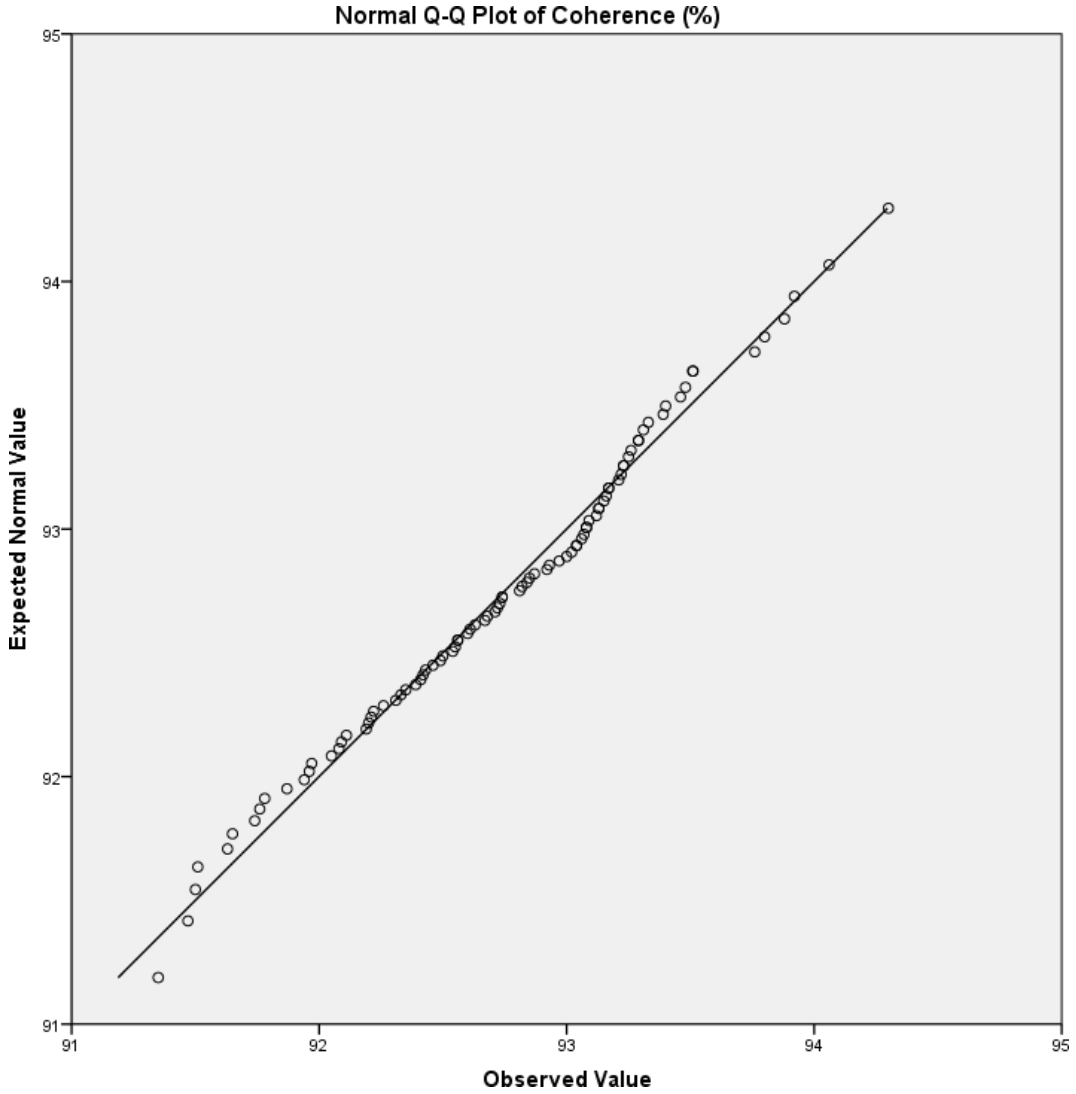

Figure B.38: Q-Q plot of average group coherence (%) for the building-evacuation scenario with groups of two using the method of Moussaïd et al.

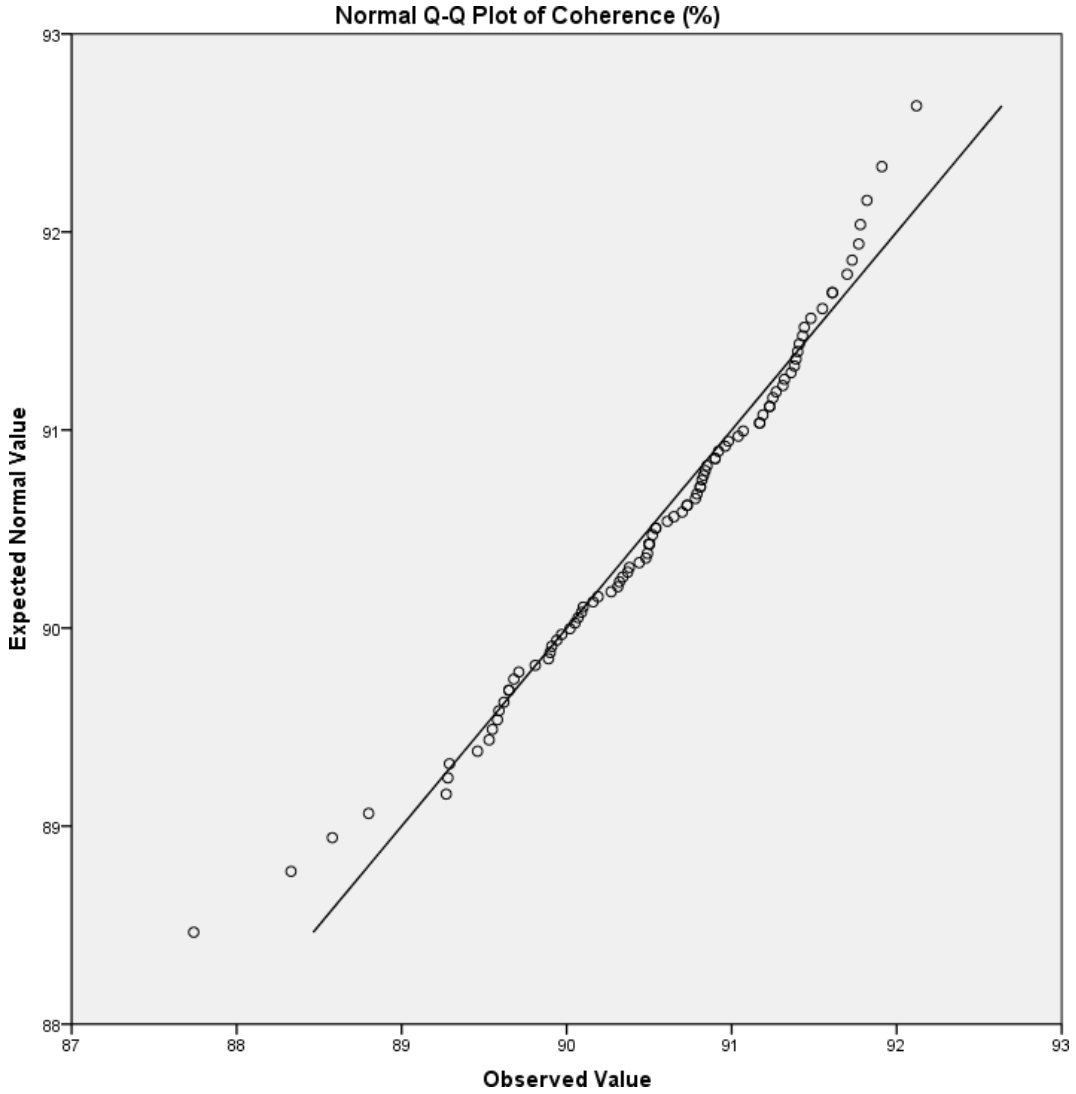

Figure B.39: Q-Q plot of average group coherence (%) for the building-evacuation scenario with groups of three using SGN method.

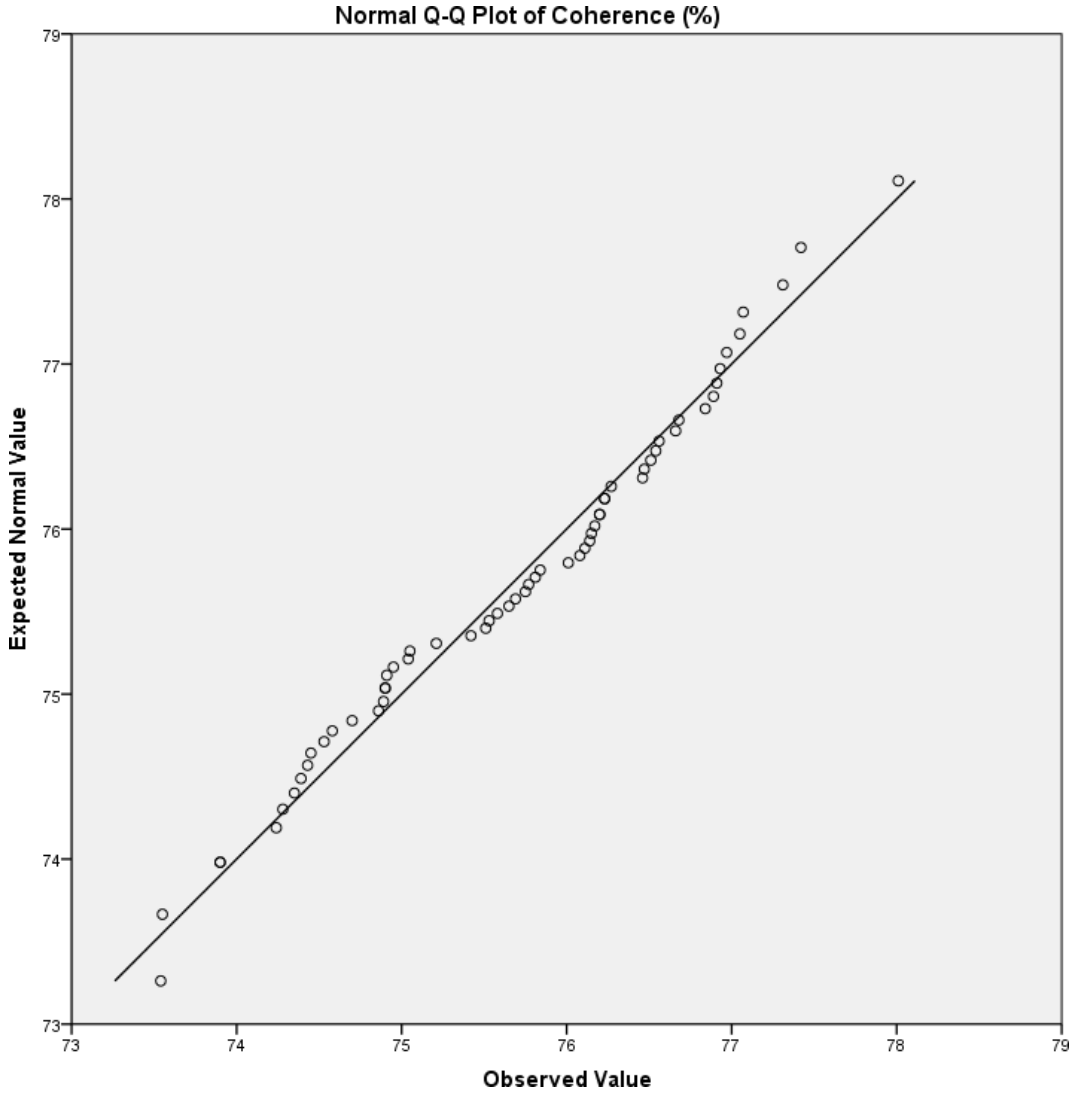

Figure B.40: Q-Q plot of average group coherence (%) for the building-evacuation scenario with groups of three using the method of Moussaïd et al.

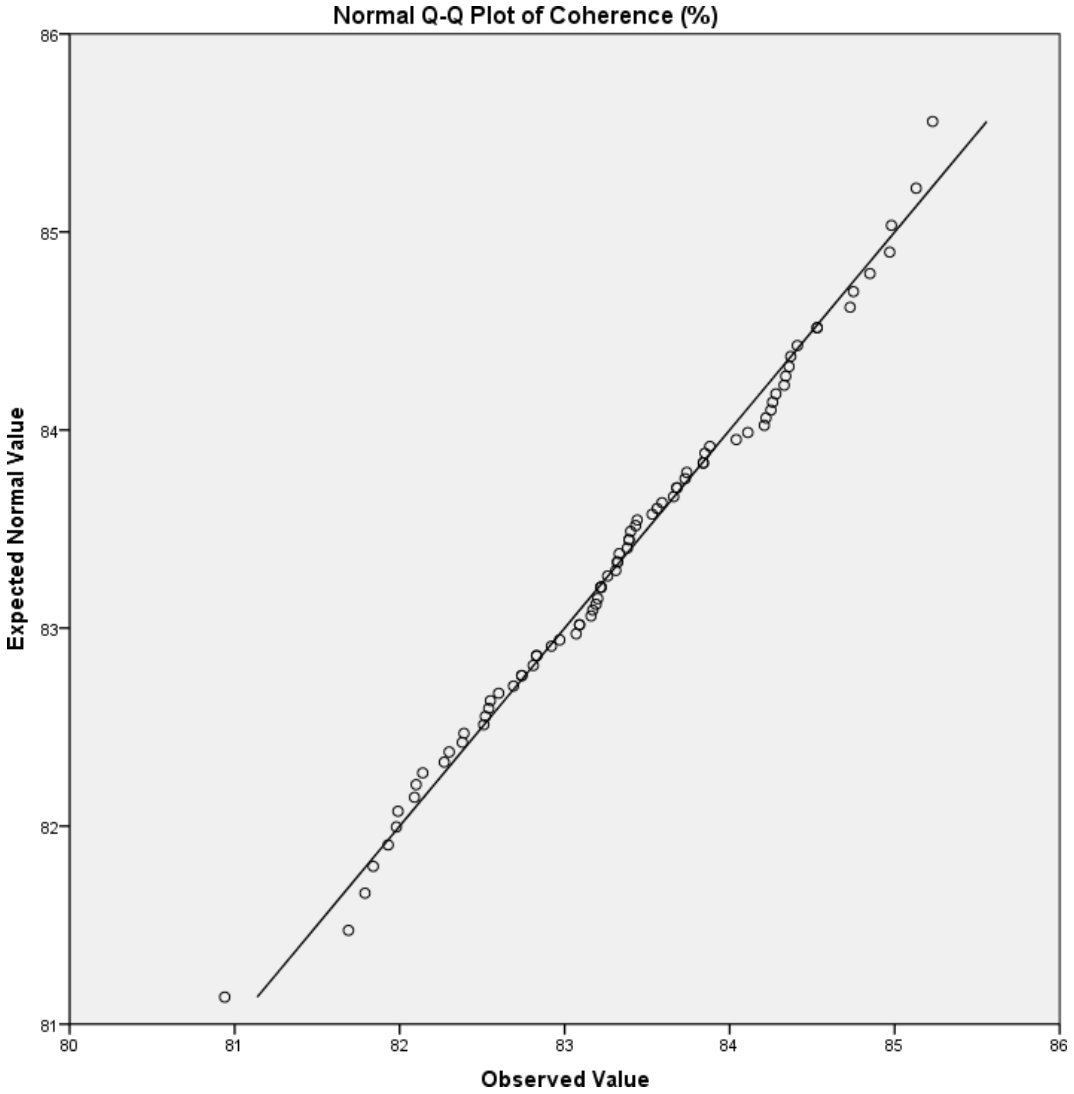

Figure B.41: Q-Q plot of average group coherence (%) for the building-evacuation scenario with groups of four using SGN method.

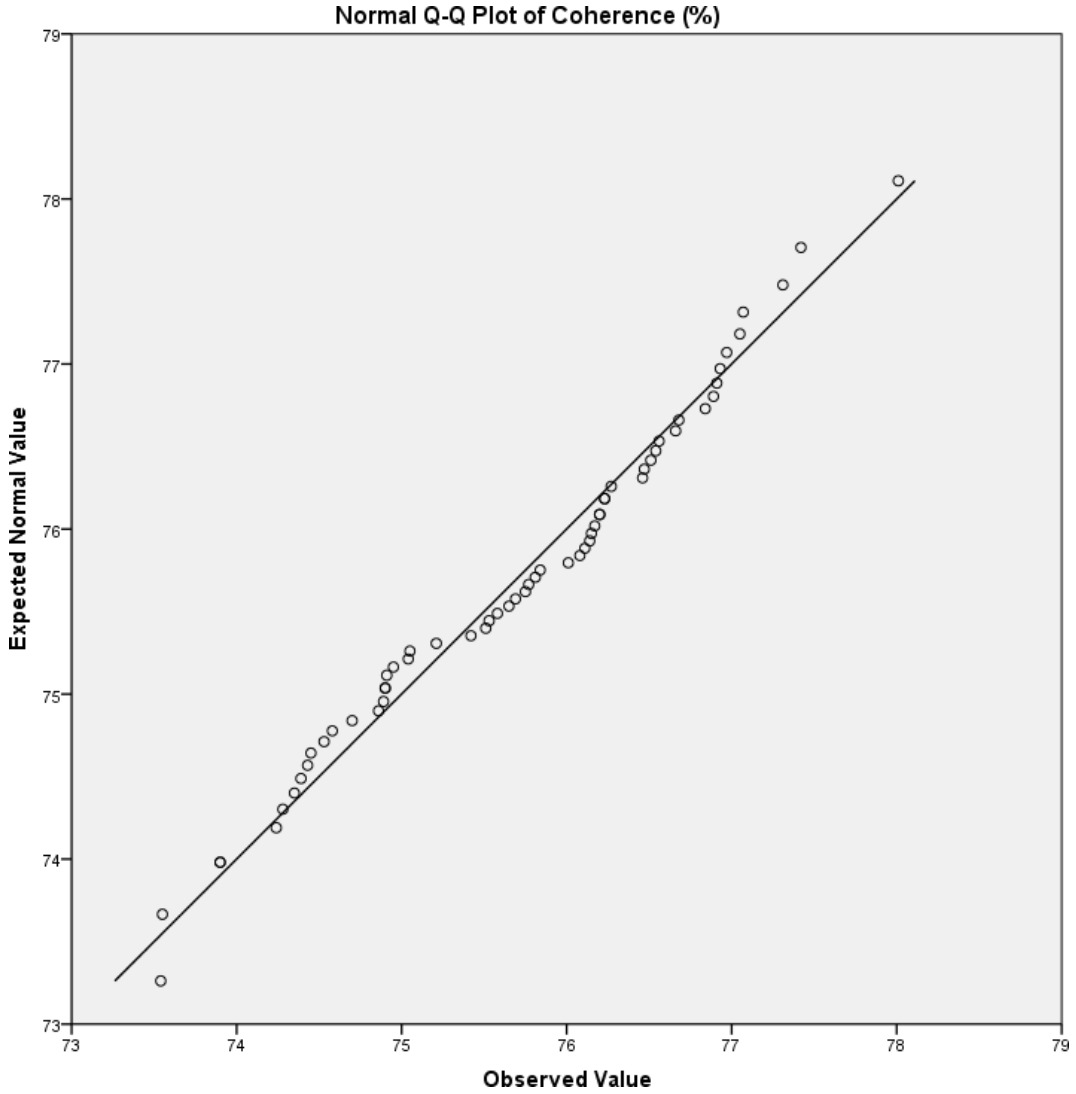

Figure B.42: Q-Q plot of average group coherence (%) for the building-evacuation scenario with groups of four using the method of Moussaïd et al.

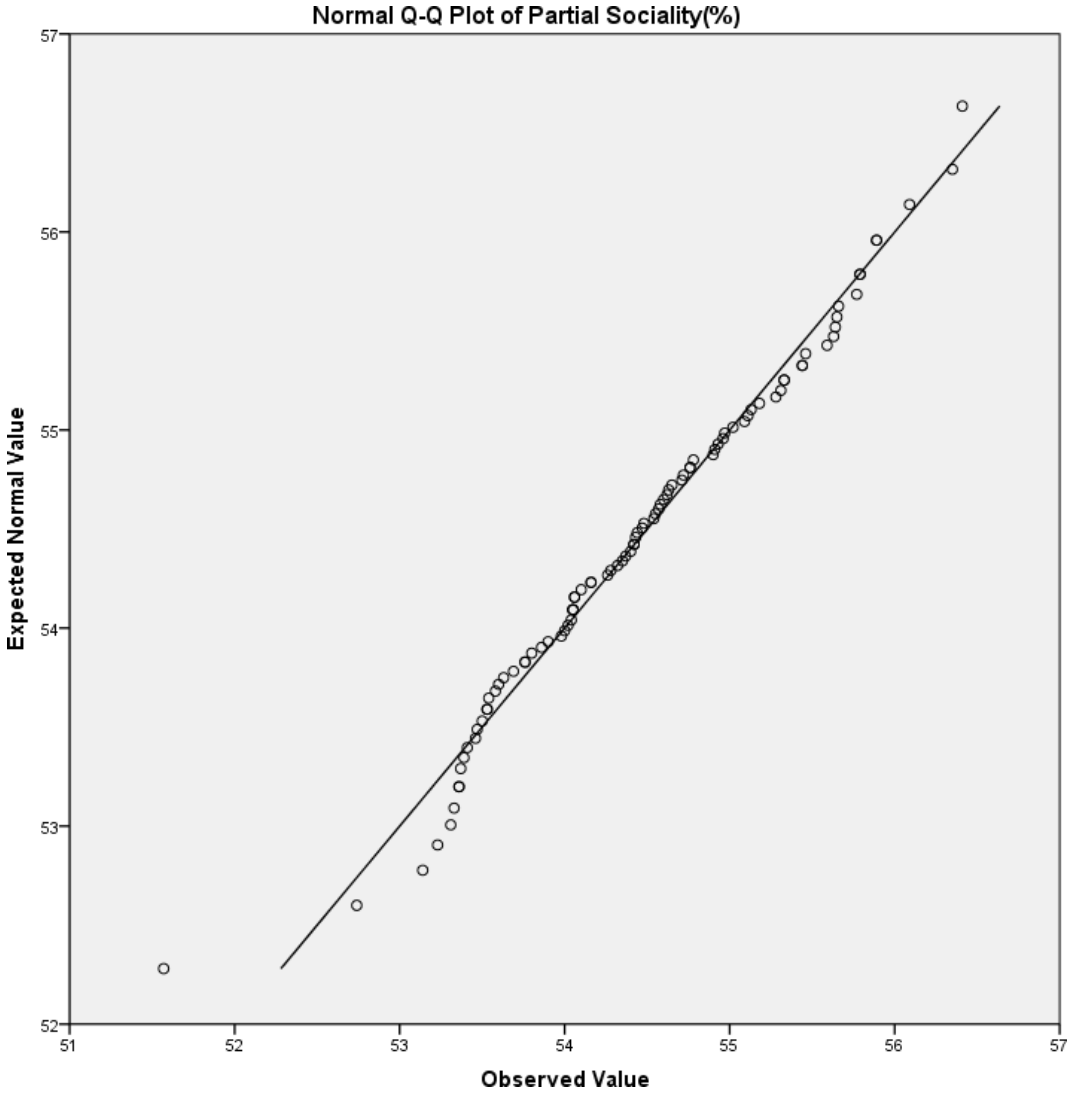

Figure B.43: Q-Q plot of average partial group sociality (%) for the building-evacuation scenario with groups of two using SGN method.

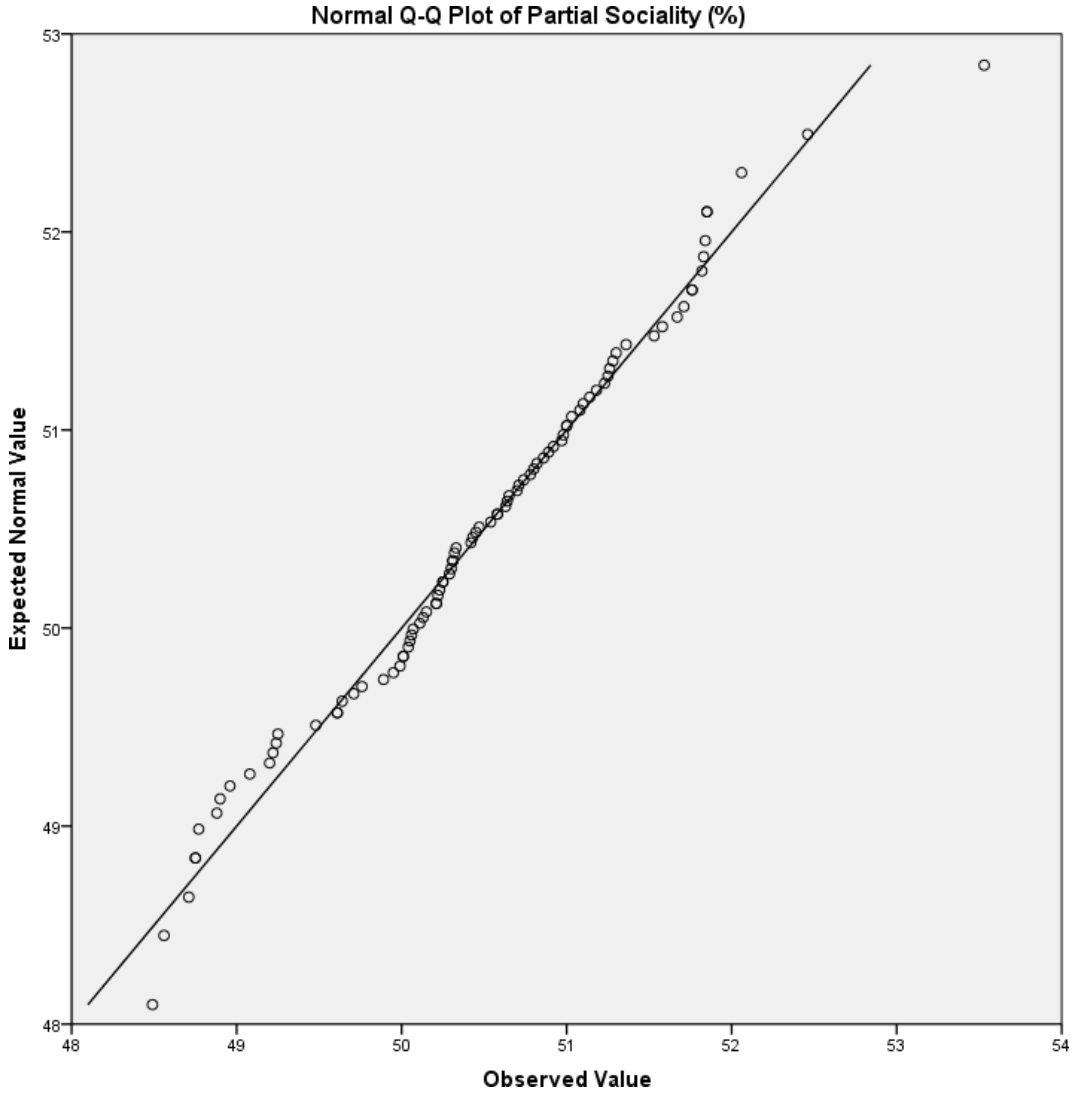

Figure B.44: Q-Q plot of average partial group sociality (%) for the building-evacuation scenario with groups of two using SGN method.

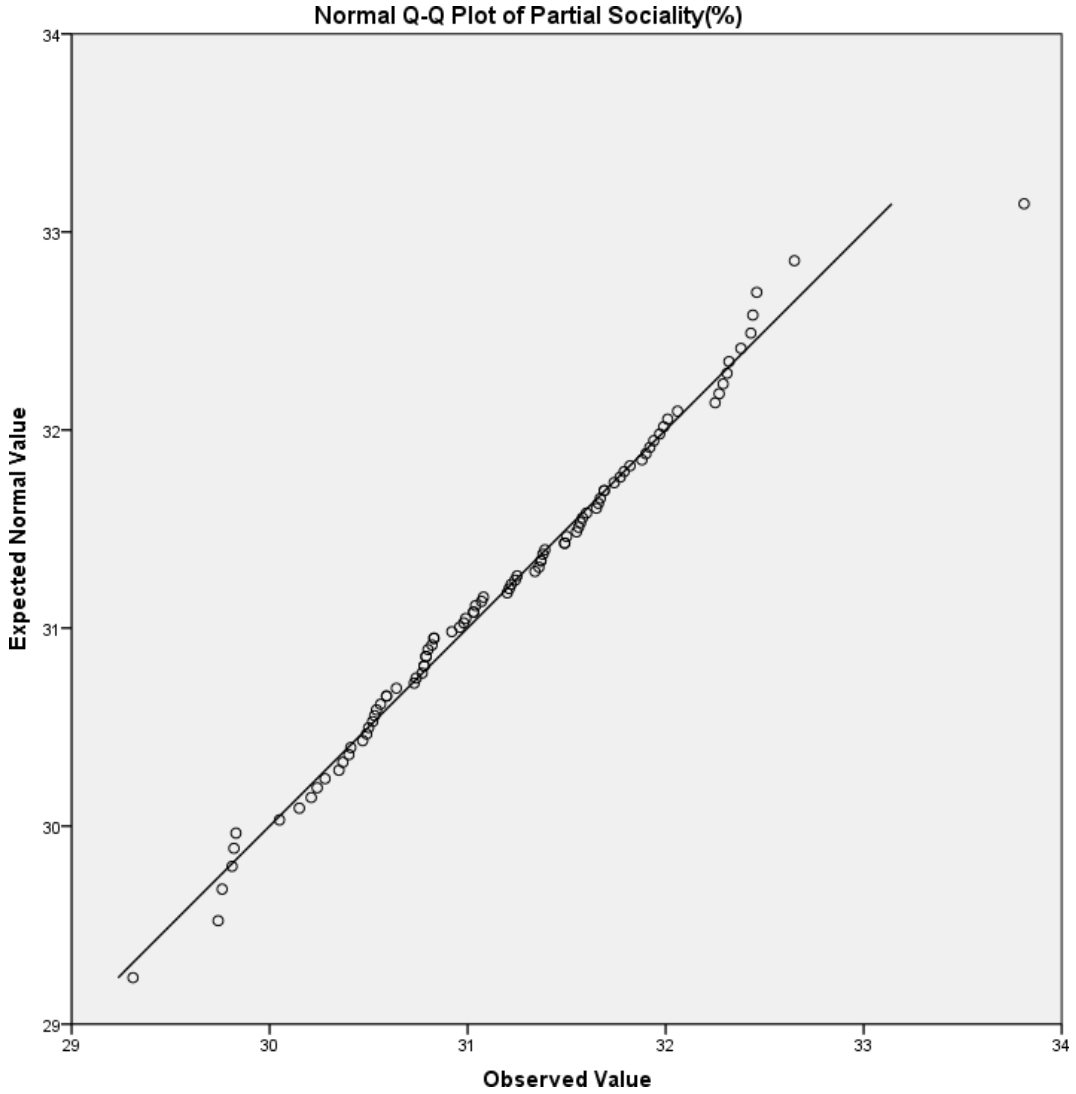

Figure B.45: Q-Q plot of average partial group sociality (%) for the building-evacuation scenario with groups of three using SGN method.

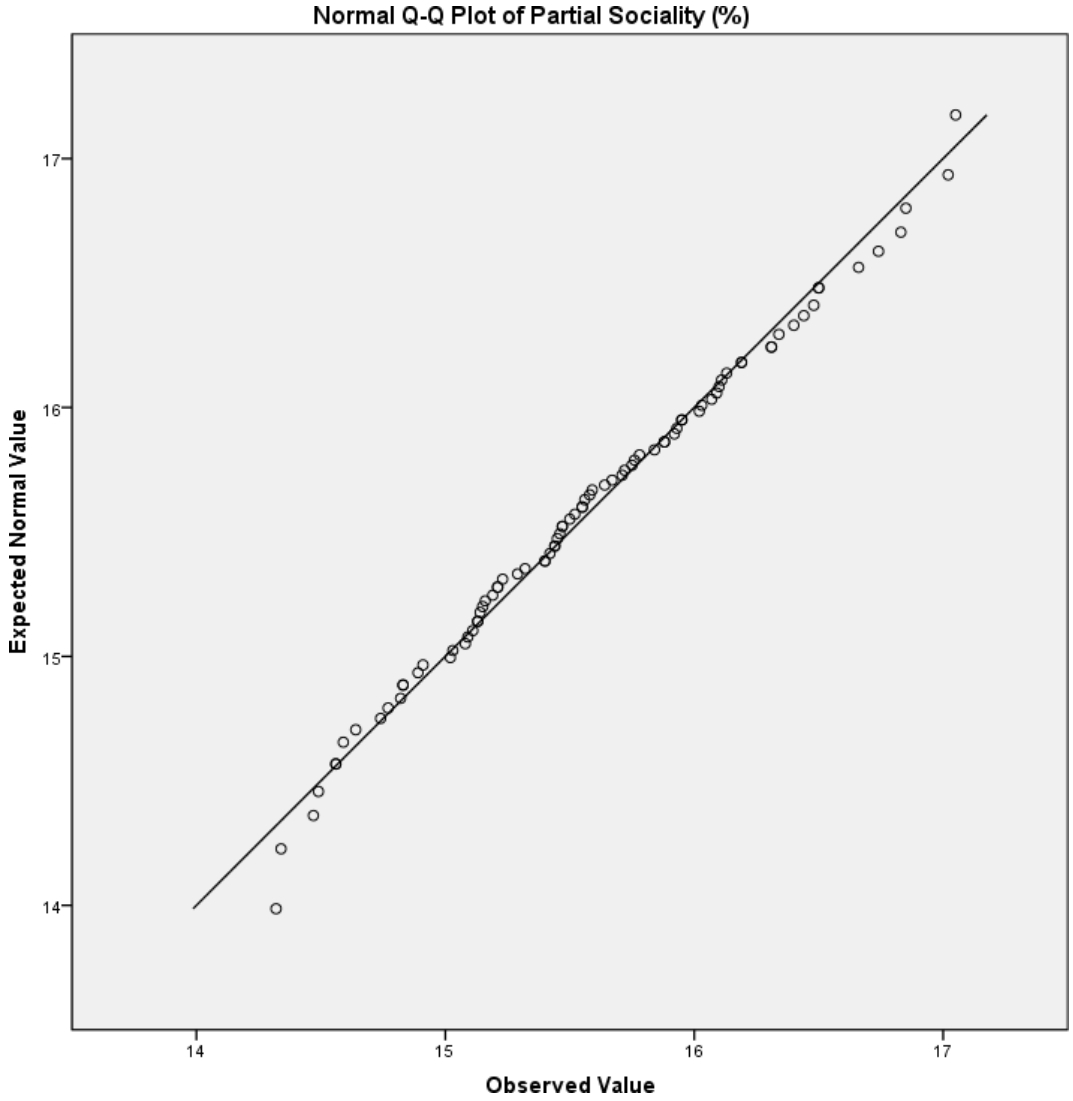

Figure B.46: Q-Q plot of average partial group sociality (%) for the building-evacuation scenario with groups of three using the method of Moussaïd et al.

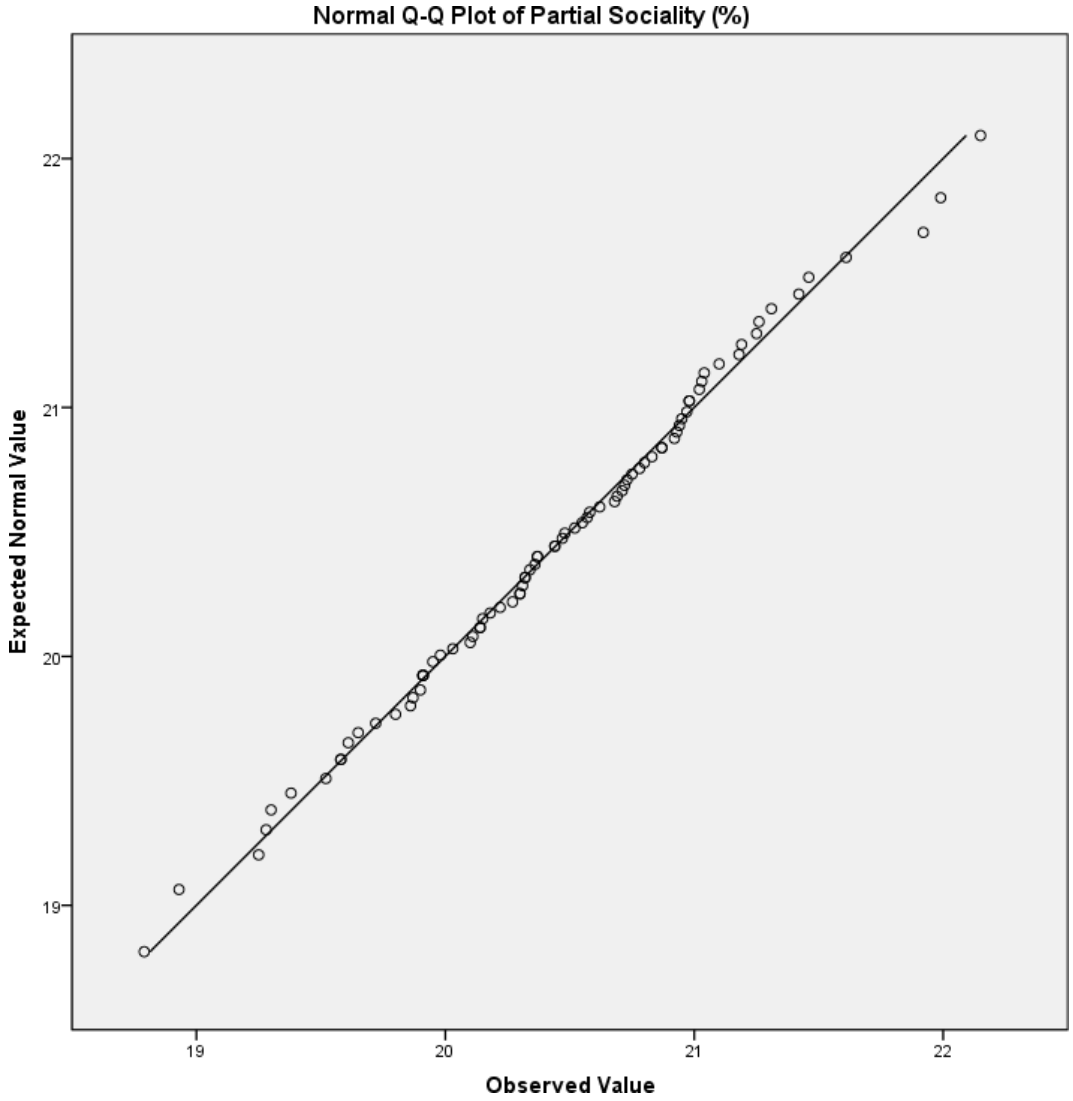

Figure B.47: Q-Q plot of average partial group sociality (%) for the building-evacuation scenario with groups of four using SGN method.

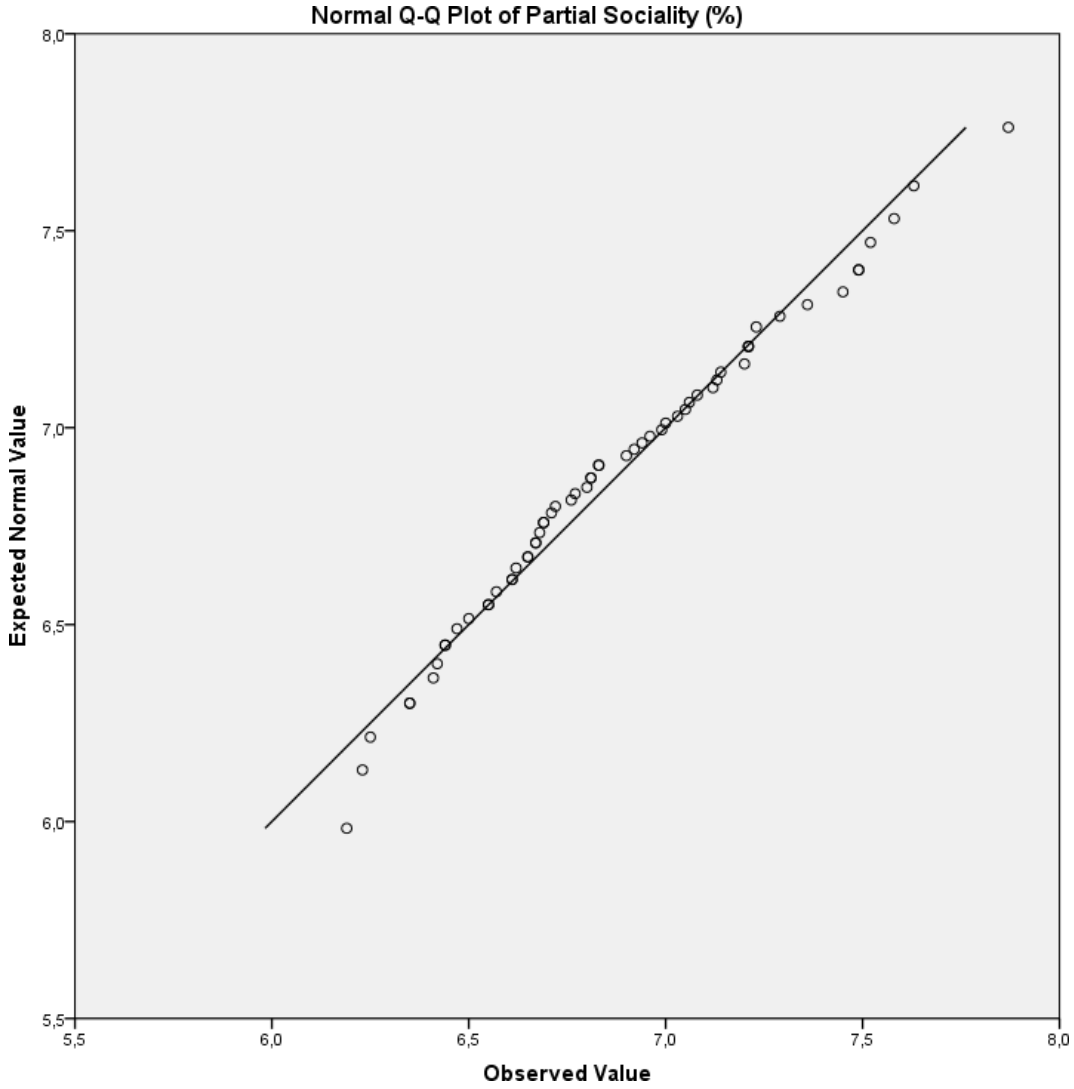

Figure B.48: Q-Q plot of average partial group sociality (%) for the building-evacuation scenario with groups of four using the method of Moussaïd et al.

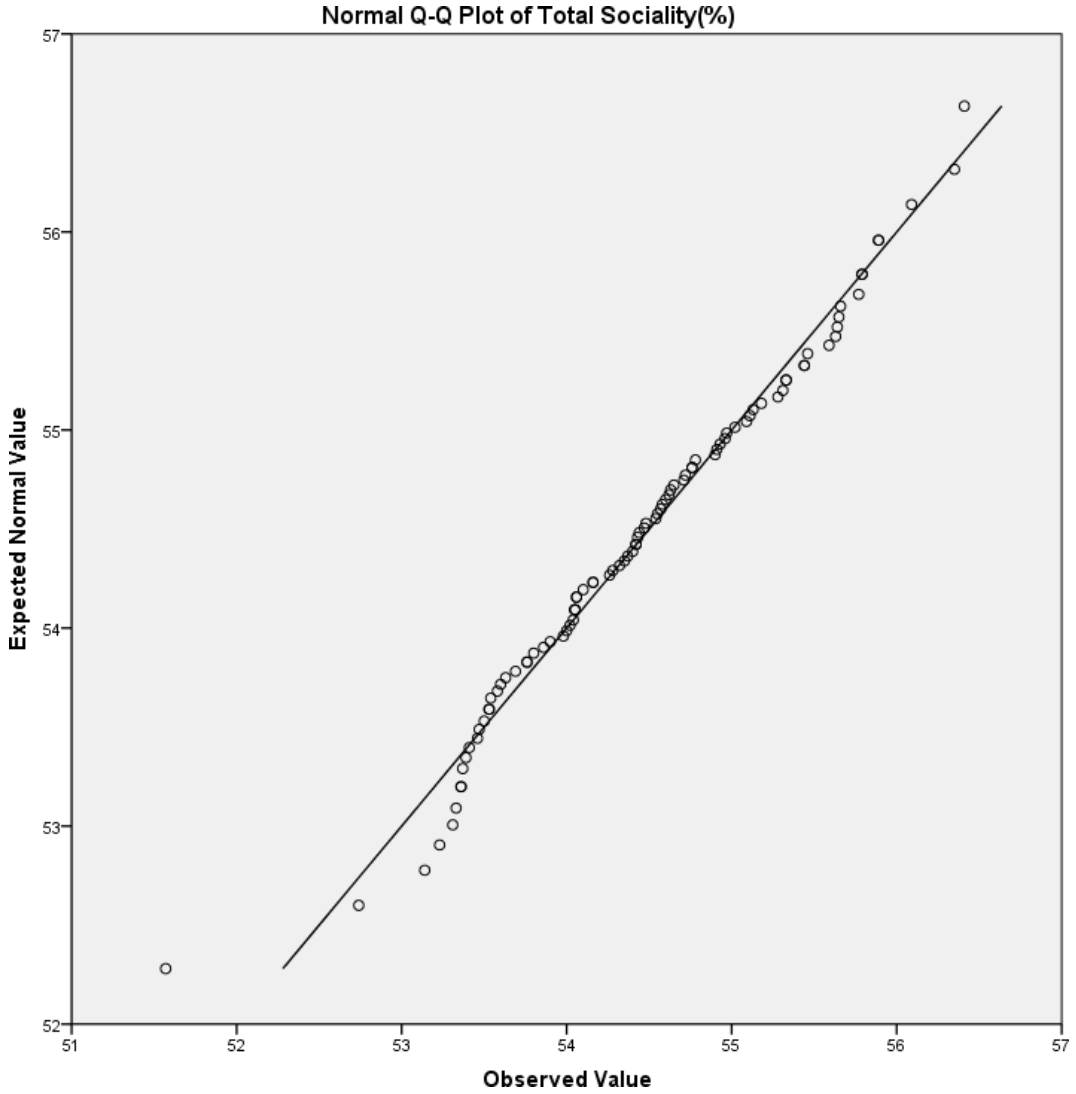

Figure B.49: Q-Q plot of average total group sociality (%) for the building-evacuation scenario with groups of two using SGN method.

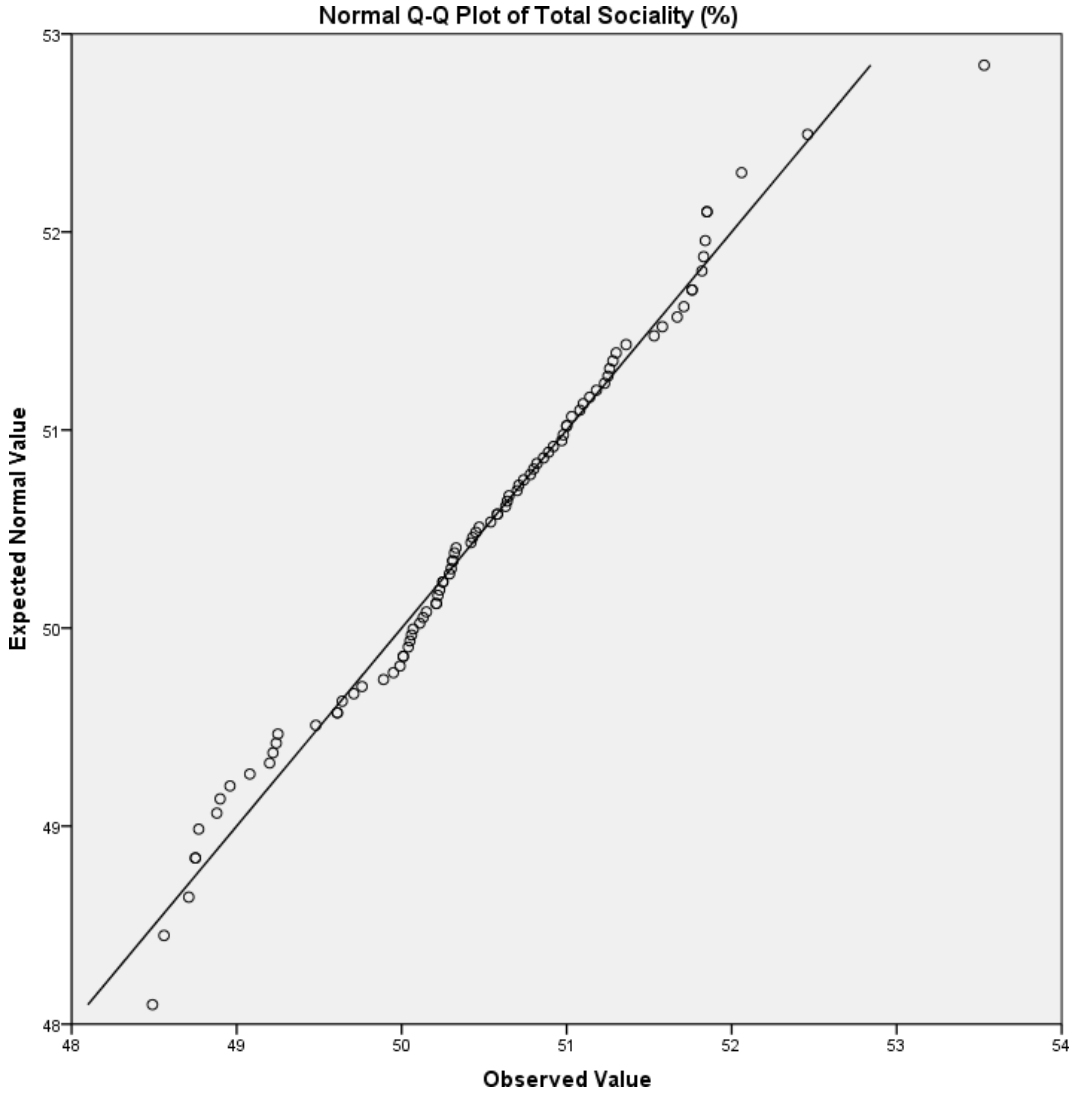

Figure B.50: Q-Q plot of average totall group sociality (%) for the building-evacuation scenario with groups of two using SGN method.

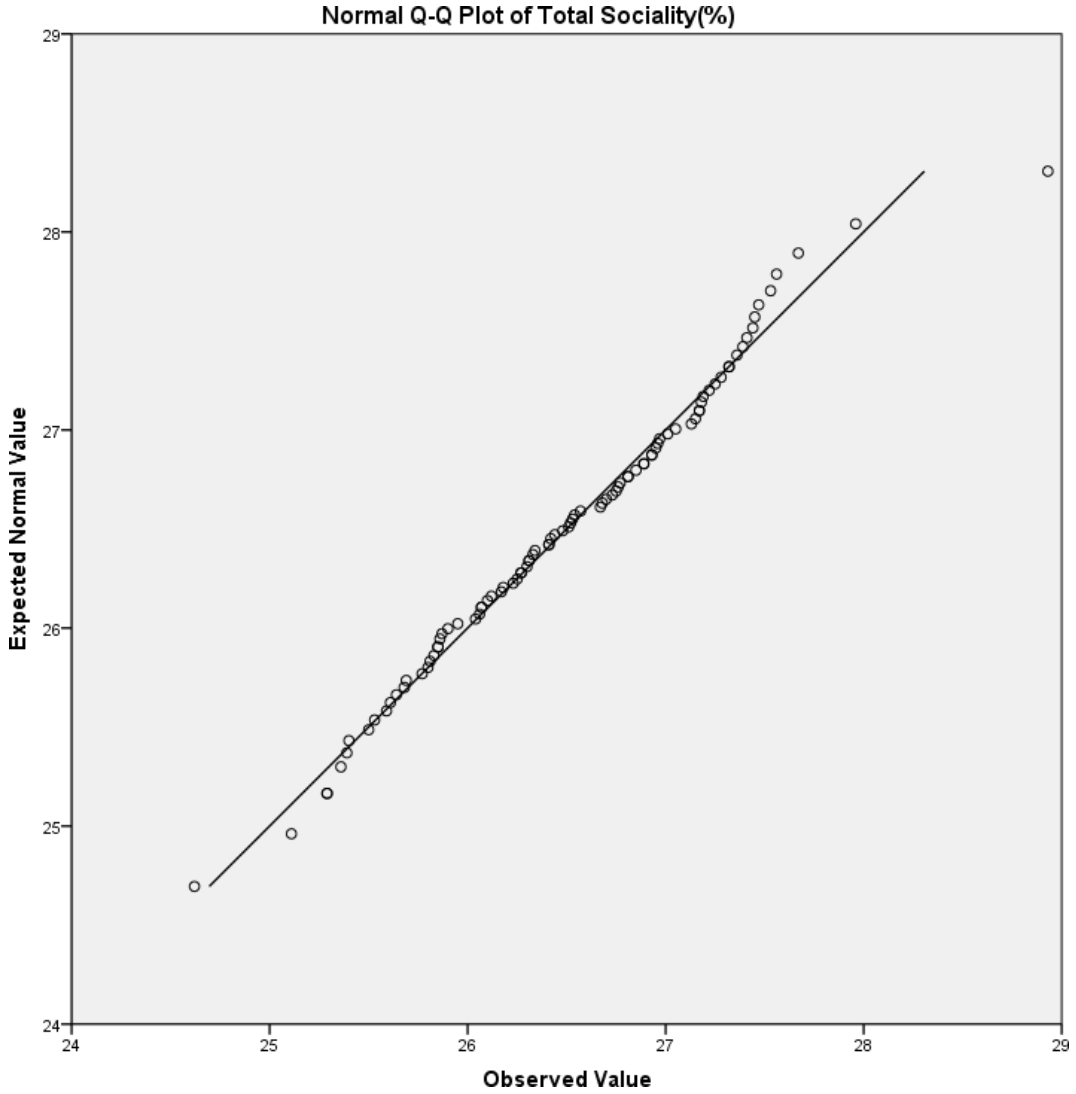

Figure B.51: Q-Q plot of average total group sociality (%) for the building-evacuation scenario with groups of three using SGN method.

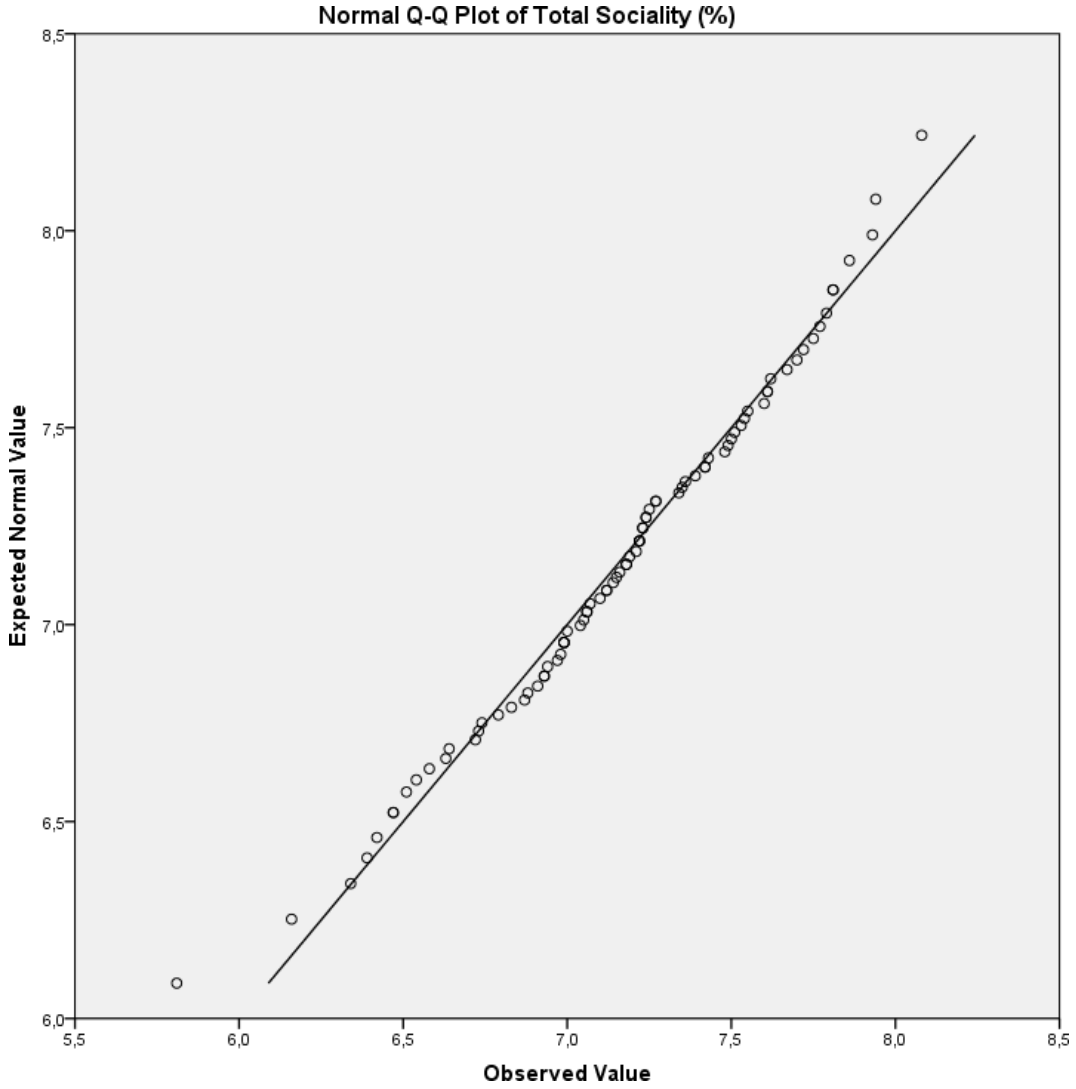

Figure B.52: Q-Q plot of average total group sociality (%) for the building-evacuation scenario with groups of three using the method of Moussaïd et al.

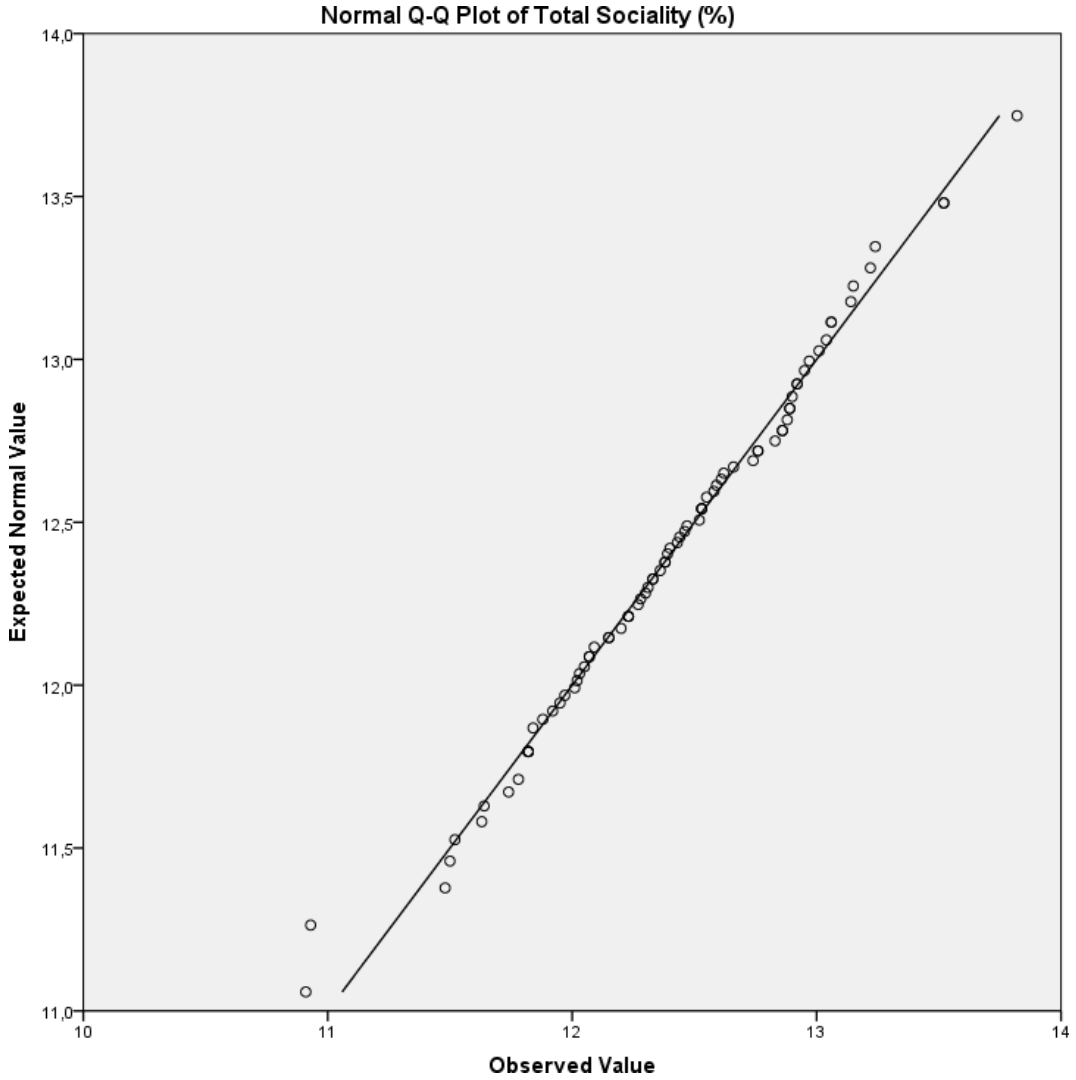

Figure B.53: Q-Q plot of average total group sociality (%) for the building-evacuation scenario with groups of four using SGN method.

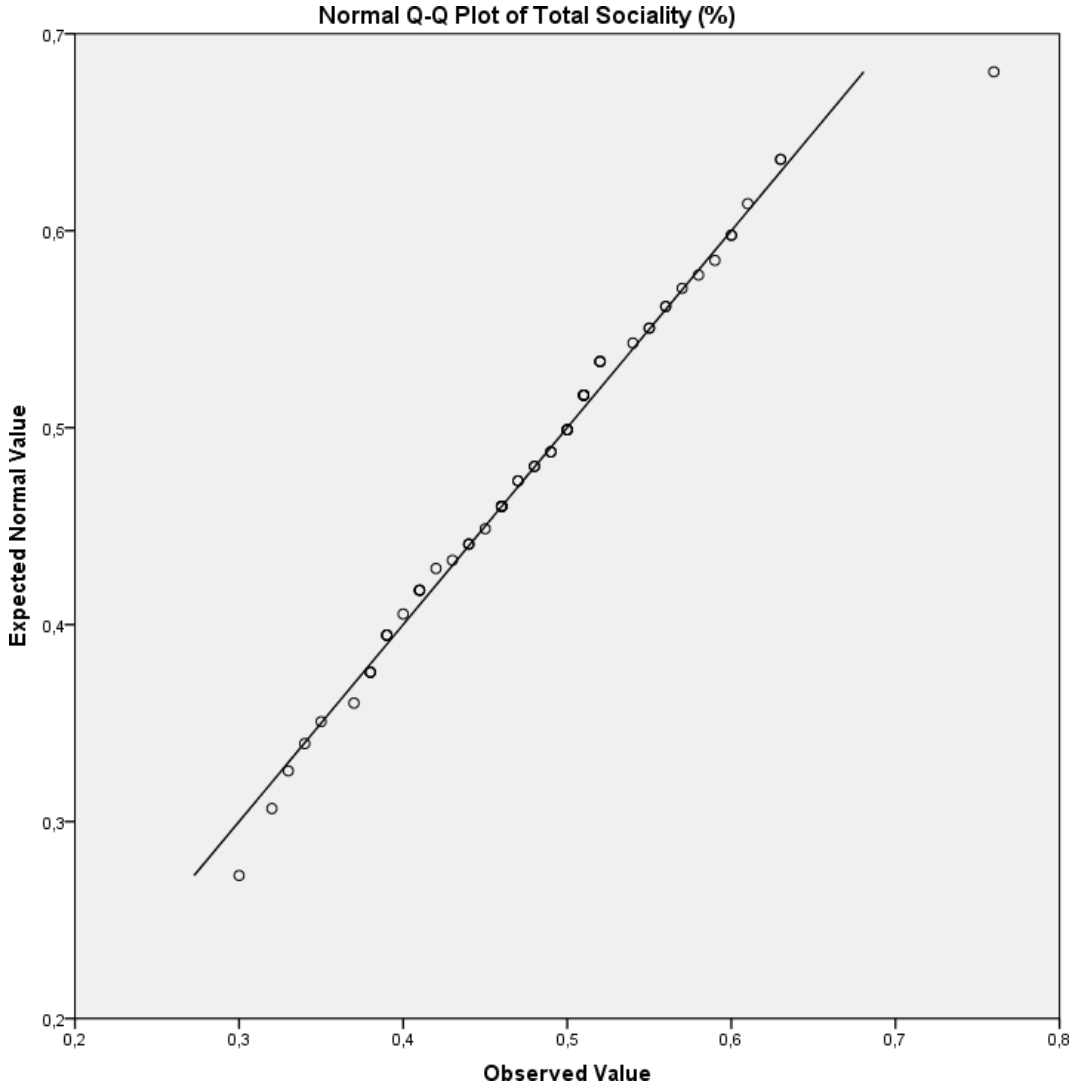

Figure B.54: Q-Q plot of average total group sociality (%) for the building-evacuation scenario with groups of four using the method of Moussaïd et al.