The One-Man-Crowd: Single User Generation of Crowd Motions Using Virtual Reality

Tairan Yin, Ludovic Hoyet, Marc Christie, Marie-Paule Cani, and Julien Pettré

Fig. 1. Snapshot of crowd motions generated using the one-man-crowd approach. A single user successively embodies each displayed virtual agent in the order indicated by the highlighting colour (from blue to yellow). We studied 3 scenarios that replicated existing experiments from left to right: circular unidirectional flow, bottleneck situation, inflow behaviour (entering a lift).

Abstract— Crowd motion data is fundamental for understanding and simulating realistic crowd behaviours. Such data is usually collected through controlled experiments to ensure that both desired individual interactions and collective behaviours can be observed. It is however scarce, due to ethical concerns and logistical difficulties involved in its gathering, and only covers a few typical crowd scenarios. In this work, we propose and evaluate a novel Virtual Reality based approach lifting the limitations of real-world experiments for the acquisition of crowd motion data. Our approach immerses a single user in virtual scenarios where he/she successively acts each crowd member. By recording the past trajectories and body movements of the user, and displaying them on virtual characters, the user progressively builds the overall crowd behaviour by him/herself. We validate the feasibility of our approach by replicating three real experiments, and compare both the resulting emergent phenomena and the individual interactions to existing real datasets. Our results suggest that realistic collective behaviours can naturally emerge from virtual crowd data generated using our approach, even though the variety in behaviours is lower than in real situations. These results provide valuable insights to the building of virtual crowd experiences, and reveal key directions for further improvements.

Index Terms—Crowd motion data, human interaction, virtual reality

✦

1 INTRODUCTION

Crowd datasets, i.e., recordings of trajectories from numerous people moving together in a same location, are paramount in the understanding, modelling and simulation of crowd behaviours. For instance, these datasets are used to study and understand collective behaviours, or to train, calibrate and evaluate simulation models. Such datasets remain however rare in spite of the large interest they yield. This scarcity is due to various reasons, such as costs, logistical, ethical and technical issues. One can imagine the effort it takes to bring many participants in a large enough laboratory equipped with scaled tracking technologies, or, out of labs, the technical difficulties in tracking people in crowded public places, after obtaining the required authorisations. The current lack of valuable crowd datasets is therefore significantly hampering research on understanding and simulating collective behaviours.

Previous work leveraged Virtual Reality (VR) to facilitate the task of collecting crowd data, typically by immersing a single participant in a virtual crowd of autonomous characters [8, 49, 50] to investigate how users behave and react. As it is generally admitted that crowd behaviours emerge from the combination of local interactions, those observations contribute to the understanding of crowd motions. Si-

multaneously, the simplicity in capturing motion data in VR facilitates observations, which avoids time-consuming post-processing steps and eases the control and replication of experimental conditions. The realism of the studied situation is however limited to the capabilities of autonomous characters to properly react to the user's actions. Therefore, using this paradigm to measure and analyse large-scale crowd phenomena presents limitations, as the collective behaviour of the virtual crowd mainly results from the techniques employed to steer such autonomous agents.

In this work, we propose to push further the idea of leveraging VR to collect crowd data, in order to overcome these limitations. To this end, we reconsider the idea of immersing participants among autonomous virtual characters. We propose and evaluate a novel approach where the motion of all virtual agents that populate the Virtual Environment (VE) is directly created by a human participant. Rather than immersing *N* participants simultaneously in the VE to embody *N* virtual agents, which would fall back to the cost- and logistics-related issues, we explore a more radical option by addressing the following question: can we collect a *valuable* crowd dataset using VR by immersing only one single participant successively embodying the *N* agents that, collectively, compose a crowd? We refer to this concept as the *one-man-crowd* paradigm (OMC). From a technical point of view, it is obviously possible to immerse someone multiple times, to embody him/her in a new agent each time, to let him/her see the previously controlled characters moving there, and finally to compose a moving crowd this way. However, what would be the *value* of an OMC-dataset and more crucially, how would the collective behaviour obtained compare to the real ones?

To explore this question, we replicated three real crowd experiments with the OMC paradigm: 1D-flow, bottleneck and inflow experiments

• *Tairan Yin, Ludovic Hoyet, Marc Christie, and Julien Pettré: Inria, Univ Rennes, CNRS, IRISA, France. Corresponding author e-mail: tairan.yin@inria.fr*
• *Marie-Paule Cani: Ecole Polytechnique/CNRS (LIX), IP Paris, France.*

(detailed in following sections), which correspond to situations of great interest in the field of crowd modelling [17, 32, 35]. We however do not aim at *quantitative accuracy* while replicating these experiments, as the use of VR will inevitably introduce behavioural and perceptual biases (see Section 2.2). In this work, we hypothesise that in addition, the OMC paradigm may bring two extra, main biases: (i) Asynchrony biases due to unilateral interactions, since only the immersed participant can react to virtual agents, but not conversely. (ii) Variety biases due to limited behavioural variety, since the paradigm consists in making only one participant repeatedly walk. We used two variants of OMC to study these possible biases: (i) S2N (Single-participant-to-N-agents), where a single participant embodies successively the *N* agents of the virtual crowd experiment (as we described above). (ii) N2N (N-participants-to-N-agents), where *N* different participants successively embody the *N* characters (to study the question of unilateral interactions in isolation). The analysis of our results leads to the following conclusions: the OMC paradigm enables a single user to generate a moving crowd with emerging, realistic collective patterns. The generated data however reflects behavioural redundancy when a single user is repetitively immersed, as well as limited, unilateral interaction. While immersing multiple users (with N2N) can reduce the first bias, further researches need to be conducted to improve OMC, and make it a logistically-simple and time-friendly approach to capture crowd motion data.

In summary, our contributions are the following. (i) We introduce one-man-crowd, a new paradigm to leverage VR to record crowd datasets. (ii) We propose and evaluate two versions of this paradigm to observe variety of behaviours in data. (iii) We evaluate this paradigm on three well-known real-world experiments, and demonstrate that OMC preserves some of the emerging large-scale crowd motion patterns, despite of lack of variety and the unilateral interaction bias.

The remaining of the paper is structured as follows. Section 2 presents the state-of-the-art on the topics of acquiring real crowd datasets and of studying collective behaviours in VR. Section 3 describes the general experimental details. Sections 4, 5 and 6 then present and analyse the results of each specific scenario evaluated using our OMC paradigm. Section 7 provides a general discussion of our approach, and Section 8 finishes with the concluding remarks.

2 RELATED WORK

In this section, we first discuss the existing literature on crowd datasets, as well as how they are usually analysed to study crowd behaviours and to model and simulate crowds. We then discuss how the field leverages VR to further study collective behaviours and collect data. We finally position our work in this context.

2.1 Crowd Datasets

The realistic simulation of crowd behaviours requires the support of real-world data of crowd motion: to study individual behaviours in crowds [40], to understand how collective motion patterns emerge [26], to devise knowledge-driven models [46], to train data-driven models [12], to calibrate models and evaluate simulations [61]. One can distinguish at least 2 types of dataset: first, field datasets capture natural, unbiased, motion (e.g., by using video tracking techniques), and second, laboratory datasets study a specific situation or the effect of some factors on collective behaviours. Our paper focuses on the latter type of dataset because they are better described and can be more accurately replicated. Only a few situations, yet of great interest, were studied in laboratories. We selected three of them for replication according to the one-man-crowd paradigm (unidirectional flow, bottleneck and inflow), because those situations implicitly set an order in the role of each participant that eases OMC replication (extension to other situations is discussed in Section 7). In the following paragraphs, we shortly describe these three situations, explain their value in modelling crowds, detail the analysis methods which apply to each, could that be on the individual or larger scales to study emergent patterns.

Unidirectional (1D-)flows are composed of pedestrians moving in the same direction [15, 62], and are interesting to study because they are frequent in real environments (e.g., queues or traffic in narrow corridors) with applications to safety (e.g., prediction of the evacuation time through a set of security exits). Experimentally, they are often reduced to a single lane of participants following one another [51, 53]. 1D-flow experiments have often been replicated to understand the effect of various independent variables: cultural parameters [13, 25], age [10, 20, 47], lighting conditions [9], environment complexity [14, 29, 55, 58], sound [63, 64], etc. 1D-flow data are analysed in different ways. Individual variables explore relations between speed and headway. Their aggregation leads to relations between flow and density that are captured by fundamental diagrams. Finally, the emergence of stop-and-go waves (similar to the ones emerging in car traffic jams) was also investigated [30, 33].

Bottleneck stands for a flow captured in an environment whose capacity locally decreases (typically a corridor that narrows). As for 1D-flow, the situation is of great interest because it corresponds to daily situations and is used to predict environment flow capacities and jams as well as evacuation times where the environment imposes the presence of bottlenecks [22, 27]. Such situations are analysed first by calculating flows of people before and after the bottleneck, i.e., the number of pedestrians passing through the bottleneck per unit time in relation with the corridor widths [11, 34, 42, 43]. It was proved that only the smaller width determines the output flow [1], and that personal motivation has a large effect on the flow [2]. Finally note that the number of lanes participants can form in the narrow part of the corridor is important to study because it determines the flow. Humans may create lanes with little separation by walking in a staggered pattern, called the *zipper effect* [51, 52].

Inflow stands for successive entry of pedestrians into a closed area [17]. Understanding this process helps modelling traffic in public transport systems (e.g., entering buses, trains, elevators) [18, 19, 36]. At the individual level, existing analyses focus on the decision-making process by participants to pick a final position in the confined space. To this end, different calculations have been adopted such as Nearest Neighbour Distance [36], Proxemic Floor Field [19] and more frequently Voronoï Diagram [17, 36]. As shown by these studies, the search for a personal comfort zone plays a dominant role in the decision-making process [16, 56, 59], and locations near the walls of the vehicle are preferred [37].

2.2 VR Crowd Datasets

Recent research leveraged VR to study collective behaviours using a multi-user [41, 54] or, most often, a single-user immersion in a virtually crowded VE [44, 65]. VR offers several benefits, among which: a single participant at a time is needed, conditions can be accurately replicated, data collection is greatly facilitated, post-processing efforts are reduced, etc. VR has for example been used to study group interactions [8], gaze activity in crowds [7], herding effects [50], following behaviours [49], etc. On the one hand VR facilitates the study of collective behaviours, but on the other hand, it simultaneously introduces a number of perceptual and behavioural biases that may compromise the usability of data collected using VR, e.g., affecting depth perception [3, 39, 48, 60], walking speeds [5, 21], balance [28], social distances [4, 8, 23, 38], etc. For this reason, a lot of effort has been put into the evaluation of the effect of those biases on VR behaviours, with a common conclusion that VR induces quantitative differences in most of spatial variables, while participants keep behaving in a qualitatively similar manner.

2.3 Positioning

To our knowledge, our work is an unprecedented attempt to generate a crowd dataset with a single user and no autonomous agents, as well as to observe collective motion patterns in VR. We evaluate the capacities of the OMC paradigm by replicating three real crowd experiments (1D-flow, bottleneck and inflow). We selected those experiments because they are well documented and were replicated multiple times in real-world situations to study various dependent variables, they involve mostly unilateral interactions (a failure to reproduce them would obviously completely invalidate OMC, including cases of bilateral interactions), and we dispose of a solid analysis framework to compare
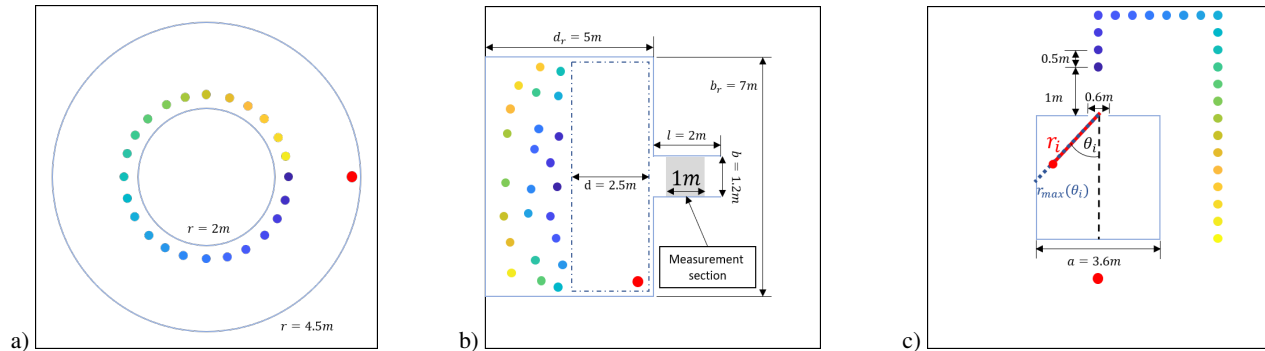
Fig. 2. Configuration of the three virtual experiments. Black squares delimit the $10\,m \times 10\,m$ experimental ground. Blue lines and circles represent walls defining the virtual room of each scene. Red dots represent the initial position and ending position of each trial. The other dots indicate the starting positions of each virtual character, coloured from dark blue (first to move) to bright yellow (last to move). For each trial, participants walked from the red dot to the position of one coloured dot to trigger the experience. a) Characters are uniformly spaced out along the circular path. b) The initial positions of characters are generated by a Poisson distribution, and ordered based on their distance to the entrance of the bottleneck. c) The queue of characters waiting to enter the virtual lift was adapted to our experimental space.

various aspects of data, both at the individual and collective scales. We expect the known perceptual and behavioural biases to influence our data. However, at this stage, we are not interested in quantifying their effect. This question was already covered by previous work on pedestrian behaviours in VR. Our main goal is to explore whether the new biases introduced by the OMC paradigm, namely variety and asynchrony biases, will allow the same collective behaviours to emerge. Will OMC enable the observation of realistic collective patterns in spite of the lack of behaviour variety due to the singularity of the user? In spite of the asynchronous unilateral interactions between agents?

3 OVERVIEW

The goal of this study is to investigate whether VR can be used to acquire novel crowd motion data in an unprecedented manner, that we call the one-man-crowd paradigm, where a single user immersed in a virtual environment either successively acts all the characters of the crowd (S2N) or different characters in different crowds (N2N).

3.1 Apparatus

Participants were immersed in the virtual environment (designed in Unity 2019.4) using a Pimax 5K Head-Mounted Display (specifications: 90 Hz, 200° fov, 2560×1440 resolution), which provided a wide field of view beneficial for such situations involving close proximity with other characters. The HMD was used with 4 SteamVR 2.0 base stations, yielding a tracking area of approximately $10\,m \times 10\,m$. We also used one HTC Vive controller in the Inflow scenario (details below), which was held by the participants and used to trigger the end of each trial. During the experiment, we recorded participants' body motions using an Xsens motion capture system, which was used both to display the participants' motions in real-time on their avatar, as well as to animate the motions of the corresponding character in successive trials. To enable participants to move freely in the environment, they were equipped with a HP Z VR G2 backpack (specifications: NVidia RTX 2080, Intel Core i7-8850H processor, 32GB RAM) running the experiment, on which were physically connected all the devices. This whole setup enabled participants to physically walk in the real space, while interacting with the other characters of the virtual crowd and seeing their own motions displayed on a co-localised avatar.

3.2 Tasks

For the purpose of this study, we designed three virtual scenarios replicating classical real crowd motion experiments: a circular 1D-flow situation, a bottleneck situation, and an inflow situation (Figure 1). Each scenario is described in details below, involved $N =24$ or 25 virtual characters depending on the scenario, and was designed to evaluate whether virtual situations would lead to emergent phenomena and individual interactions similar to those observed in previous real experiments. However, all the tasks followed the same general principle. For a given trial T_i a single participant is immersed in the virtual environment in one specific character c_i, displaying the participant's co-located motions in real-time. Every character $c_{j\in[1..i-1]}$ is displaying global and body motions recorded during previous trials T_j of this specific participant (except for the baseline condition which will be described in Section 3.3). Characters $c_{k\in[i+1..N]}$ are handled differently depending on the scenario. Participants started each trial on a red spot (located at a fixed position during the entire experiment), then walked towards another spot located on the ground among a group of virtual characters (initial position of character c_i). Each scenario had a specific ending condition, which would terminate the trial once activated. Participants would then be guided back to the initial red spot (not displayed during the actual trial) so that they would always start new trials from the same location.

1D-flow. This scenario was designed to replicate a real-world experimental condition used by Lemercier et al. [33], where 24 real participants walked behind each other in a circular space. As in the real-world experiment, our participants were instructed to walk along a circular path (radius: 2.4 m, i.e. trajectory length of 15.1 m) delimited by two circular walls (inner radius: 2 m, outer radius: 4.5 m – see Figure 2-a). The inner wall was 2 m height, to prevent participants from anticipating their future movement based on characters too far down the line. During the experiment, participants were instructed to follow the character in front of them from the beginning of each trial, while avoiding collisions. They were also not allowed to overtake the character in front of them. The following characters were animated using a pre-recorded motion capture animation (see below). They did not react to participants' motions, but became transparent if they overtook or went through the participant's avatar. For the first trial, all the characters (except the participant's avatar) were animated with a pre-recorded motion capture animation so that participants could continuously adjust their speed depending on a previous character as in the real-world experiment, and would not walk alone. This pre-recorded animation was motion captured on a male confederate prior to the experiments. The confederate was instructed to follow a virtual object in the same virtual scene, animated using the real trajectory of one experimental condition of Lemercier et al. [33] so that participants would respond according to accelerations and decelerations observed in a previous real situation. Using the same pre-recorded trajectory and animation for each participant to initially follow also helped to ensure comparable situations in the following analyses. In each trial, participants walked for a duration of 60s, after which they were instructed to return to the initial location indicated by the red spot.

Bottleneck. This scenario was designed to replicate a real-world ex-

perimental condition used by Liddle et al. [35], where a group of 24 participants were asked to walk through a real corridor (bottleneck). As in the real-world experiment, our participants were instructed to walk through a virtual bottleneck (1.2 m wide, 2 m long – see Figure 2-b), while avoiding collisions with virtual characters and attempting to exit the bottleneck quickly. Before starting a trial, participants joined the other virtual characters in a 7 $m \times 2$ m waiting zone, located 2 m from the entrance of bottleneck. Unlike in the previous situation, only the first $c_{j \in [1..i-1]}$ virtual characters were animated during trial T_i. The other characters $c_{k \in [i+1..N]}$ remained idle in the waiting zone, but were presented to participants at the beginning of the trial so that they had a sense of the number of characters involved in the scenario. The trial ended when the participant went beyond a red line on the ground, indicating that they had exited the bottleneck. Also, the order of acting the virtual characters was based on their initial distance to the entrance of the bottleneck (from minimum to maximum), so that participants would first act virtual characters closest to the entrance.

Inflow. This scenario was designed to replicate a real-world experimental condition used by Ezaki et al. [17], where a group of 25 participants were asked to enter one after another into a lift. As in the real-world experiment, our participants were instructed to stand in a line of virtual characters, then to enter a virtual lift of 3.6 $m \times 3.6$ m (see Figure 2-c) and to choose a place to stay once they entered it. The entry of the lift was 60 cm wide to limit participants' view so that they could not perceive the position of the virtual characters that had already entered the lift. As for the bottleneck scenario, only the first $c_{j \in [1..i-1]}$ virtual characters in front of participants were animated during trial T_i, using participants' recorded animations from previous trials. The other characters $c_{k \in [i+1..N]}$ remained in their original position, but were presented standing in line with an idle motion to participants at the beginning of the trial so that they had a sense of the number of characters that should enter the lift. At the beginning of this scenario, we insisted that participants should take into account that all the 25 virtual characters should be able enter the lift, and therefore that they should choose their position during each trial to enable the remaining characters to enter the lift as well. The trial ended when the participant considered that he had reached his wanted position and was facing the entrance, which was indicated by participants pressing the trigger on the HTC Vive controller that they were holding in this scenario.

3.3 Experimental Design

Upon their arrival, participants read and signed the experiment consent form, during which they were presented the task to perform. They were then equipped with the Xsens motion capture system, the VR backpack, as well as the Pimax 5K HMD. Calibration of the motion capture system was then performed to ensure motion capture quality, as well as to resize the avatar to participants' dimensions. Participants then first performed one trial of each scenario presented in Section 3.2, contributing to the creation of the N2N baseline (N-participants-to-N-agents) described below, after which they performed a full single scenario, referred hereafter as S2N (Single-participant-to-N-agents).

N2N baseline: trajectories and motions of the N characters in the crowd were created successively by N different participants. For each scenario, participant i only took the role of the i-th character c_i, while the $c_{j \in [1..i-1]}$ previous characters were animated based on recorded animations and trajectories of the previous $i - 1$ participants. Participants always performed a trial of the 1D-flow scenario first, then the bottleneck scenario, and finally the inflow scenario.

S2N procedure: a single participant created successively the trajectories and motions of all the N characters in the crowd, by taking the role of one character at a time. Each participant took part in a single scenario, comprised of 24 or 25 trials (depending on the number of characters to act). It is also important to point out that participants were never told during the experiment that they would be interacting with previous versions of themselves (see Section 7 for further discussion).

3.4 General Hypotheses

Our general hypotheses are related to the possible effects of all the perceptual biases (e.g., distance or self-motion perception), behavioural biases (e.g., slower walking speeds), variety biases (e.g., single user in the S2N case) and asynchrony biases (e.g., unilateral interactions) on the data generated by the OMC approach.

- **H1**: At the individual level, quantitative differences are expected between real and OMC virtual data. Aside from this, we expect a high similarity with real data, and to preserve existing relations between analysed variables.

- **H2**: On the collective scale, the OMC paradigm will lead to a lack of variety in data, meaning that S2N data could be considered as a subset of real or N2N data (e.g., S2N data distributions included in real data distributions). Moreover, N2N data will be more similar to real data than S2N data.

- **H3**: We expect that both N2N and S2N conditions will enable us to observe the emergence of crowd motion patterns, as the scenarios that we replicate involve mostly unilateral interactions.

As an important note, we explore those hypotheses in the light of qualitative analyses described in the following sections. As we replicate experiments designed by the crowd modelling community, we also replicate the analysis framework that was performed, and discuss the most striking effects we can observe. This choice is discussed in Section 7.

3.5 Participants

We recruited 25 participants for our experiment (age min=22, max=32, avg=26.8±3.1), which is the minimum number of participants required for the creation of our virtual baseline data. To avoid biases caused by motion characteristics coming from different genders, as well as to embody participants in gender-matched avatars, we recruited only male participants. Participants were recruited through internal emailing lists amongst students and staff, were all naive to the purpose of the experiment, had normal or corrected-to-normal vision, and gave written and informed consent prior to the experiment. The study conformed to the declaration of Helsinki.

For simplicity, as participants took part both in the N2N and S2N procedures, we attribute an ID from #1 to #25 to each participant, with respect to their participation order in our experiment. This means that participant #i took the role of the i-th character in the N2N baseline scenarios. Participants #1 to #8 then performed the 1D-flow S2N scenario, participants #9 to #16 the bottleneck S2N scenario, and participants #17 to #25 the inflow S2N scenario.

4 UNIDIRECTIONAL FLOW

By virtually replicating Lemercier et al.'s experiment [33], we expect to also replicate a number of real-world observations they made: the emergence of stop-and-go-waves, specific relations between individual speed and headway, as well as resulting fundamental diagrams. We also wonder whether the absence of physical constraints in VR will lead to unrealistic situations, such as overlaps between agent positions.

4.1 Analysis & Results

Data processing. For this experiment (only), participants positions $\mathbf{p}_i$ are computed according to a polar coordinate system centred on the circle path they follow: $\mathbf{p}_i^{polar}(t) = [\theta_i(t); r_i(t)]$. We are only interested in the evolution of θ, as $r_i(t)$ is almost constant since participants walk along a circle. We approximate r by its mean value r_{mean}. The position, speed and acceleration can thus be reduced to:

$$\begin{cases} p_i(t) \approx \theta_i(t) r_{mean} \\ v_i(t) \approx \dot{\theta}_i(t) r_{mean} \\ a_i(t) \approx \ddot{\theta}_i(t) r_{mean} \end{cases} \quad (1)$$

Stop-and-go waves. Figure 3 plots all the trajectories recorded for one

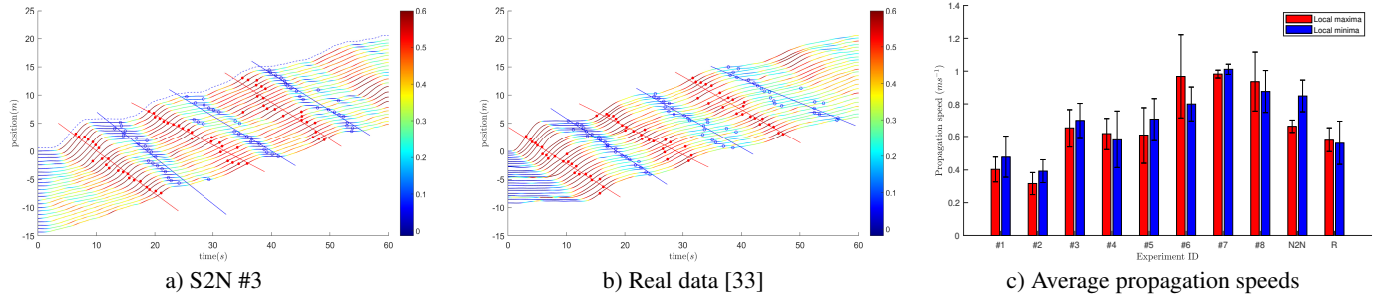a) S2N #3   b) Real data [33]   c) Average propagation speeds

Fig. 3. Illustration and results of the calculation of propagation speeds. a)-b) Local maxima and local minima of pedestrians' speed, and result of linear regression. a) Result of #3 (dashed line: pre-recorded trajectory of the first character). b) Result of Lemercier et al. [33]. c) Average and SDM results of all S2N (#1 to #8), N2N and real-world data (R).

given S2N dataset in parallel with real data from Lemercier et al. [33]. On the S2N plot, the upper blue dashed curve corresponds to the first agent moving in the circle, i.e., the pre-recorded confederate reference trajectory. As we said in Section 3.2, this trajectory reproduces a real case (center plot) and is made of periods where the confederate walks, stops, and walks again, etc. The plot makes evident that the stops of the confederate propagated backward in the 1D-flow, with a certain propagation speed. Jelic et al. [30] proposed a method to compute the speed of counter flow propagation, by computing the local extrema of $v_i(t)$, and by performing a linear regression on the obtained points, as illustrated in Figure 3. The rightmost plot shows the distribution of the obtained speed values for all the trials of each participant, as well as for the example from Lemercier et al.'s real experiment. This figure shows that S2N experiments resulted into both faster and slower waves propagation speeds, depending on participants, and that N2N resulted in in-between values, close to real ones.

Reaction time and following distance. Lemercier et al. [33] found that participants in a 1D-flow were mainly matching the speed of their predecessor, trying to cancel $\Delta v_i(t) = v_i(t) - v_{i-1}(t)$. They also found that the propagation speed of waves was probably determined by participants' reaction time τ when matching speed, and that the wave could be damped or resurge depending on the following distance between individuals $p_i(t) - p_{i-1}(t)$. We followed the procedure proposed in [33] to estimate reaction-times (for each participant, and for each time frame of data). We estimate $\tau^*(t)$ on a moving time window W:

$$\tau_i^*(t) = \arg\max_{\tau}(a_i(t) \star \Delta v_i(t + \tau)) \tag{2}$$

where $a_i(t)$ is the participant's acceleration and $\star$ the cross-correlation. Figure 4 plots the obtained distributions for reaction time as well as following distances for each dataset (S2N and N2N) in comparison to values obtained for Lemercier et al.'s experiments [33]. This figure shows that the range of observed reaction times is very close between OMC-S2N or N2N and real. On average, participants matched the speed of their predecessors with a delay of 0.5 s.

Fundamental Diagram. We computed the fundamental diagram that captures relations between speed and density. As in [30], we considered local density and speed, i.e., where $\rho_i(t) = (p_i(t) - p_{i-1}(t))^{-1}$. We smoothed values by adopting the following binning procedure. We first calculated the density-speed pair $(\rho_i(t), v_i(t))$ for each character i at each time step t. We then sorted all the density-speed pairs from the minimum density to the maximum, and separated them into 960 bins of K elements each ($K = 180$ in our experimentation). We then computed the average density and speed per bin as:

$$\bar{\rho} = \frac{1}{K}\sum_{n=1}^{K}\rho_n \quad \bar{v} = \frac{1}{K}\sum_{n=1}^{K}v_n \tag{3}$$

Figure 5 plots the resulting pairs, for all S2N, N2N as well as the real example from [33].

Global density. In real 1D-flow experiments, participants' positions are restricted by the circle they follow, which gives a global density. In OMC virtual replications, since participants only see characters in front of them and cannot feel the physical constraints imposed by participants behind them, nothing prevents them from forming circles where the positions of the first characters overlap with the positions of the lasts. This explains why we gave an initial position to the participants out of the circle and asked them to get to their position in the queue: doing this allowed them to get an estimate of the global density, and we here check whether this global density was matched all along the trial and whether character positions overlapped. To this end we compute the total length $L(t)$ of the circle formed by characters as the sum of the distance between each pair of predecessor-follower as follows:

$$L(t) = \sum_{i=2}^{N} p_i(t) - p_{i-1}(t) \tag{4}$$

As can be seen in Figure 4-c, the length of the 1-D queue exceeds the perimeter of the circle for a few participants, indicating the emergence of overlapping: the expected global density is not matched.

4.2 Discussion

Figure 3 illustrates the striking similarity in data we could obtain between Lemercier et al.'s real experiment and their OMC replication. The complete set of results are provided as supplementary material. Especially, we were able to observe as expected the propagation of stop-and-go waves for all participants, which validates **H3** (emergence of collective phenomena) for this scenario. The further analysis of wave characteristics however reveals quantitative differences, as expected by **H1**. Real experiments resulted in wave propagation speeds limited to $0.57 \pm 0.09\ m.s^{-1}$, while a larger range could be observed in S2N datasets ($0.69 \pm 0.24\ m.s^{-1}$). Moreover, due to a lack of variety (i.e., the same participant acting all the characters in the crowd), the S2N procedure resulted in participant-specific values. On the contrary, N2N results mixed and averaged these participant-specific behaviours. For instance, the propagation speed of waves in N2N results ($0.66 \pm 0.03\ m.s^{-1}$) is close to the average of all our S2N datasets. As stated in **H2**, N2N lowers the effect of lack of variety, even though N2N wave propagation speeds remain different from real values. However, as real and OMC values were obtained over two different sets of participants, a more strict OMC-real comparison would need additional experiments, which is discussed in Section 7.

According to their 1D-flow model, Lemercier et al. found that wave propagation speeds mainly depend on participants' reaction time to match the speed of their predecessor. This is consistent with our analysis of reaction time. For instance, participants #7 and #8 displayed both the lowest reaction times and the fastest wave propagation. As a result, when different participants mix in the N2N procedure, the resulting wave speed gets averaged as a result of individual differences in reaction times. We can however observe that participants with the lowest reaction times are the ones with the highest following distances. This suggests that they might simply be more careful and attentive, and
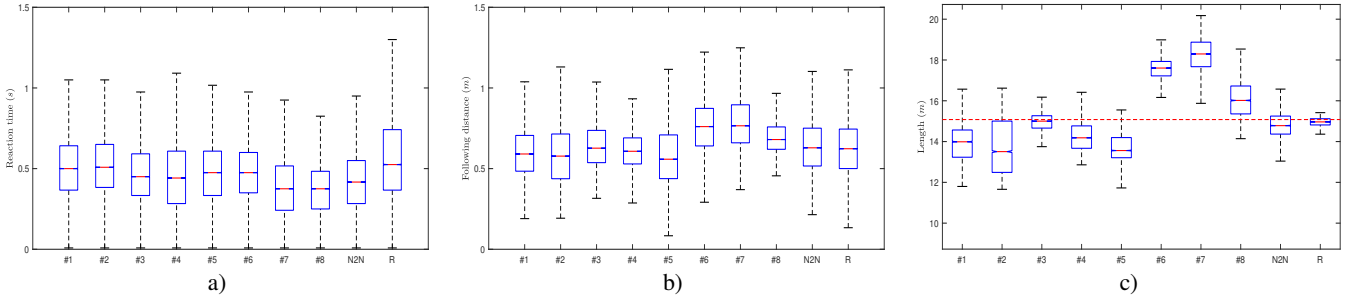
5

Fig. 4. a) Reaction time. b) Following distance. c) Total length of the circle formed by characters (the red dashed line indicates the physical length of the path). #1-#8: S2N results. R: Real-world data [33].
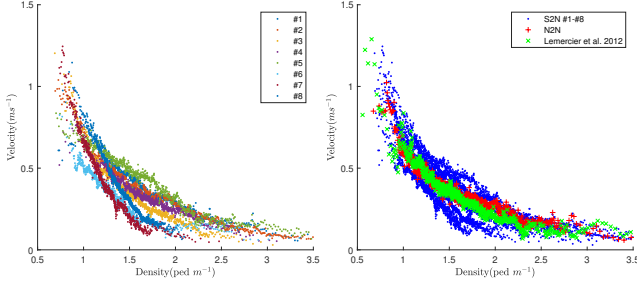

Fig. 5. Fundamental Diagrams showing the velocity-density relationship. Left: Results of the S2N procedure separately plotted. Right: Compared to real-world results, data generated by our method seems to be translated towards higher density level.

that this individuality gets amplified by making them interacting with themselves. However, it is also possible that differences in reaction times in VR could be related to biases in the perception of distances, as objects typically look closer in VR than in real life [39]. Exploring whether perceived proximity to predecessors could make participants more reactive would however need to be explored in future experiments.

The analysis of fundamental diagrams leads us to similar conclusions. From Figure 5 we can see the speed-density pairs of each S2N result locate around the ones of real cases, with a distorted shape caused by the personal preference of each S2N participant and the lack of variety (**H1**). Thus, we could think that the real or N2N fundamental diagrams could result from the combination of the obtained S2N diagrams. Again, the N2N condition partly solves issues related to lack of variety, and the corresponding fundamental diagram has a more similar coverage compared to the real one (**H2**). However, note that fundamental diagrams of both the S2N and N2N procedure are also shifted compared to real data. This result indicates that participants tended to walk faster than in real life with the same local density. While this contradicts previous studies which found slower speeds in VR [21], one possible cause of this phenomenon is the absence of contact rendering combined with HMDs' reduced field of vision: participants may not be able to see the entire body (of themselves and of their predecessor), especially the lower part. The physical space to perform walking motion was also not constrained by the presence of the physical legs of a real predecessor, resulting in larger steps and faster walking speed. Note that in real experiments, these physical constraints result into step locking behaviours as studied by Jelic et al. [31].

Finally, we feared that the absence of physical constraints would lead to overlaps between agents position. For example, the beginning of the queue of agents may overlap its end. During the experiment, we did not render agents in this situation to avoid disturbing participants, as these agents would otherwise block their vision and traverse their avatar. We however compared the length of virtual queues against real ones (which are obviously bounded by the physical length of the circular path), as displayed in Figure 4-c. Our results provide mixed conclusions. Some participants clearly adopted following distances

which, accumulated, violated the limit of the path length (e.g., #6 and #7), whereas most participants finally adopted distances that matched the expected global density. This suggests that our choice of making participants able to observe the global density by giving them a starting position out of the circle, then joining in the space allocated to them in the circle, somehow enabled most participants to be aware of the global density to be matched. This result also suggests that this choice might have had a persistent effect through the experiment, and that physical differences from real situations inherent to this virtual setup might be in some ways attenuated with careful organisation.

As a final point, we discuss the choice to animate the first agent in the circle using motions pre-recorded by a confederate. It was motivated by two reasons. First, in real experiments, stop waves are initiated by the first participant in the circle being blocked by the presence of the last one. This would have raised a problem in the OMC replication since the last agent motion is recorded last. Also, as the confederate replicates the motion of the first participant in Lemercier et al.'s experiment, doing this enabled us to perform a strict comparison between real and virtual data as illustrated in Figure 3, as well as to ensure that all participants were presented with the same initial condition. While requiring to record specific animations prior to an experiment might be limiting in some situations, we believe that approaches replicating characteristics of an observed trial could be also used, e.g., from recorded videos.

5 BOTTLENECK

In this replication of a bottleneck experiment [35] we explore the biases introduced by the unilateral interactions, as participants need to negotiate with people coming from their side to enter the bottleneck. We evaluate the resulting number of lanes in the bottleneck as well as fundamental diagrams, since they are useful outcomes of such an experiment.

5.1 Analysis & Results

Lane formation. To count the number of lanes, we reproduced the method described in [52] to estimate the probability to find a pedestrian at a certain lateral position p_x inside a measurement zone (a corridor section along the y axis) with $-0.6\ m \leq x \leq 0.6\ m$ and $0.5\ m \leq y \leq 1.5\ m$ (see Fig. 2). Example trajectories are illustrated in Figure 6. The distribution of lateral positions (binned) is presented in Figure 7. Real data corresponds to the results of the experiment performed in [35]. The number of formed lanes can be easily observed by counting the number of peaks in the distribution. For the real data, we can see that 2 lanes emerge. More difficult to appreciate from static plots only, the little distance separating the lanes results from the zipper effect (staggered walking pattern, see the companion video for illustration). Interestingly, OMC led to two extreme kinds of lane formation: a parade neatly lined up and a perfect formation of 2 lanes going through the bottleneck (as it can be seen in Figure 6-a and -b respectively). Meanwhile, most of our results present a mixed pattern between these two extreme results, resembling the N2N result presented in Figure 6-c which exhibits a normal distribution of lateral positions. Figure 6-a displays what can happen in the case of a single lane, where participants

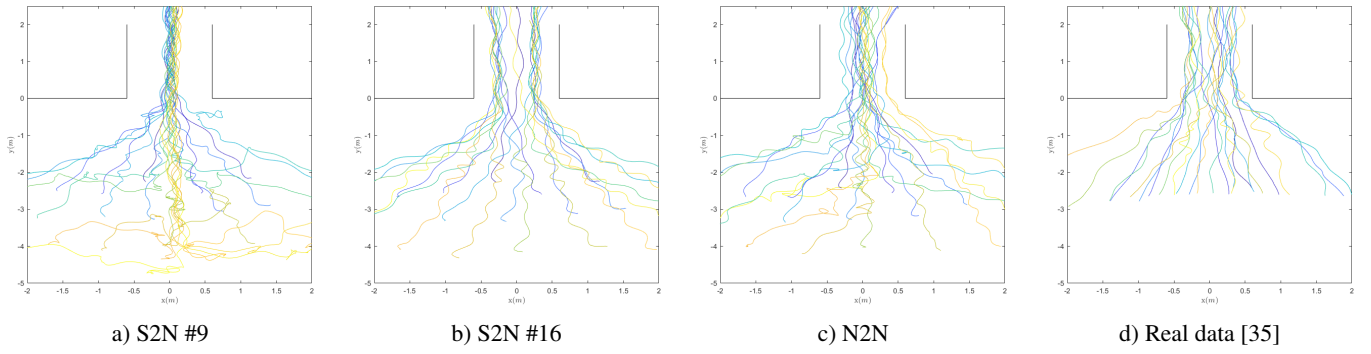a) S2N #9   b) S2N #16   c) N2N   d) Real data [35]

Fig. 6. Comparisons of lane formations in different conditions. Trajectories are coloured by the order in which participants acted as virtual pedestrians (see Figure 2). a) A typical result where virtual pedestrians walk neatly on a single lane (participant #9). b) A typical result showing clearly the formation of 2 major lanes in the bottleneck (participant #16). c) Trajectories resulting from the N2N procedure. d) Trajectories of 24 pedestrians entering the bottleneck in the real-world experiment [35].

lined up by first walking directly towards the end of the parade, instead of gathering together around the entrance (as observed in most other trials, e.g., the N2N one). Additional S2N examples are also included in the supplementary material.

Bottleneck fundamental diagram. Fundamental diagrams capture relations between density and flow, but flow in a bottleneck linearly depends on its width [52]. To make the diagram independent to the bottleneck width, we study the *specific flow* J_s (as defined in [52]) which represents the flow per unit of space, computed as the product of density ρ and the flow speed v:

$$J_s = \rho v \qquad (5)$$

The estimation of ρ and v is performed on a section of the corridor ($y \subset [0.5\,m, 1.5\,m]$). ρ is computed according to the method proposed by [57] (that guaranties continuity with respect to time):

$$\rho(t) = \frac{\int_A P(\mathbf{x},t)d\mathbf{x}}{|A|} = \sum_i \frac{\int_A P_i(\mathbf{x},t)d\mathbf{x}}{|A|} \qquad (6)$$

where A represents the measurement area, $P(\mathbf{x},t)$ is the probability density of having a pedestrian at position $\mathbf{x}$ at time t, $P_i(\mathbf{x},t)$ is the corresponding probability density for pedestrian i at time t, and is supposed to be uniform on the Voronoï cell of the pedestrian. With this definition, $N(t) = \rho(t)|A|$ denotes the number of pedestrians in the considered section of the corridor and is continuous even when an agent enters in or exits from the section. We then compute the flow speed by a weighted averaging method as follows:

$$f(t) = \sum_i \frac{\int_A P_i(\mathbf{x},t)d\mathbf{x}}{N(t)} v_i(t) \qquad (7)$$

where the speed v_i of pedestrian i is numerically derived from its position. We present the fundamental diagram resulting from specific flow against local density in Figure 8. This figure shows that the range of speeds participants walked at are similar between S2N and N2N. The specific flow values are also clearly lower for OMC than real data, confirming the effect of a 1-lane traffic where a 2-lane one was expected, resulting into a loss of flow. We also consider that walking speed, known to be slower in VR, may have accentuated this loss of flow.

5.2 Discussion

In this second experiment, we attempt to validate the idea of generating a crowd walking through a bottleneck with the OMC paradigm. We selected a bottleneck width where, in real experiments, a 2-lane traffic starts emerging (the full real experiment studied corridor width as a dependent variable). Our attempt was not always successful: few participants formed 2 lanes, most only 1 lane, and we could observe
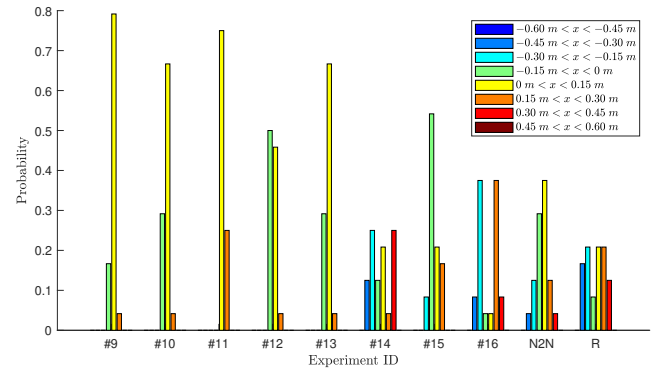

Fig. 7. Probability to find a pedestrian at position x inside the exit measurement zone $-0.6\,m \leq x \leq 0.6\,m$, $0.5\,m \leq y \leq 1.5\,m$.
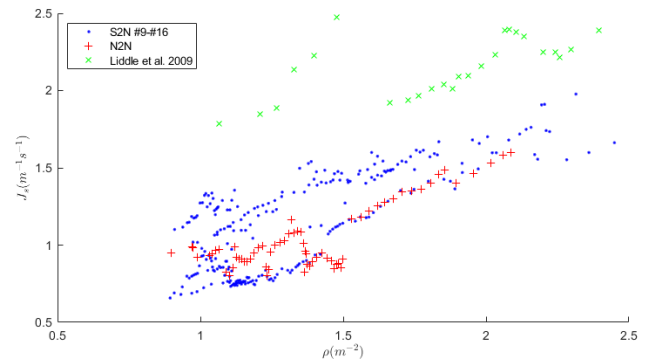

Fig. 8. Estimated specific flow with respect to instantaneous local density.

specific queuing up behaviours. We believe that the non-emergence of a 2-lane traffic can be explained by perceptual and variety biases: i) some participants may have perceived the corridor to be smaller and considered that only a 1-lane traffic would fit, and ii) the S2N procedure resulted in the same decision made by the participant trial after trial, and sometimes resulted in specific trajectories like the one showed in Figure 6-a. This validates hypothesis **H1**, showing quantitative differences between real and OMC virtual data at the individual level. However, we believe that this effect could be further studied based on previous work on passability judgement tasks in VR [24]. Indeed, the question of the number of lanes fitting a given corridor width could be turned into one of passability into a gap. It might therefore be relevant to replicate the full Liddle et al. [35] experiment with corridor width as a dependent variable, so as to observe transition width (from
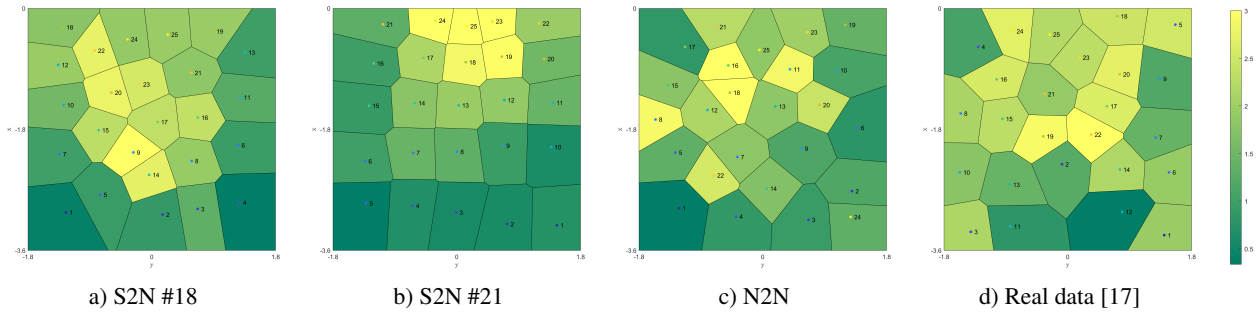
7

a) S2N #18   b) S2N #21   c) N2N   d) Real data [17]

Fig. 9. Final position for different experiences. The entrance of the lift is located on top-middle. Each cell is colored by the local density of each position, truncated at $3\ ped.m^{-2}$. a)-b) Two typical S2N results. c) N2N result. d) One real-world result [17].
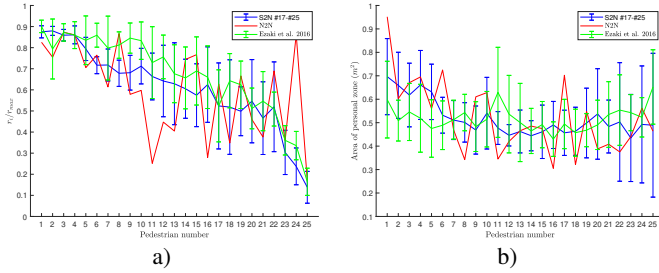

Fig. 10. Evolution of a) how far pedestrians walked along the direction of their final position and b) personal space area, with respect to the order of entering, for S2N, N2N and real-world results, respectively averaged over 8, 1 and 9 experiences.

1 lane to 2 lanes for example) in a VR context. As expected from **H2**, N2N prevents this individual-specific decision on the number of lanes to form to be repeated trial after trial. Our interpretation is that, all together, participants varied the decision to form a 1-lane or a 2-lane traffic, which resulted in a uniform distribution of lateral positions. **H3** is therefore only partly validated, as the expected pattern could be observed in some trials, but this was not always the case.

Finally, it is quite clear that the fundamental diagram analysis enables observing the large quantitative differences in the measured flows: the absence of emergence of a second lane in most cases completely offsets flow values in similar densities. This highlights the need for awareness of quantitative differences of experiments performed in VR. Nevertheless, we also observed a consistency between the S2N and N2N results, as well as that the relation between specific flow and density is preserved compared to real world results (flow decreases when density increases, with similar slopes).

6 INFLOW

Different from the two previous experiments where we study moving crowds, we here focus on entering into a virtual confined space by replicating [17] and investigate how participants decide which position to occupy. While previous experiments explored motion interactions with predecessors, here we explore more how people interact with their environment and the whole crowd: people pick the position they would prefer, also being conscious that a number of people have to fit. We here present the corresponding analysis in terms of participants' final positions and personal space.

6.1 Analysis & Results

Final positions. Final positions picked by participants are illustrated in Figure 9, along the corresponding Voronoï diagram, for two S2N examples, the N2N and a real data example from [17]. The figures show both the order in which the positions were picked, and the personal space participants preserved around them.

Following the analysis performed in [17], we can further analyse the position picking strategy by representing positions according to the polar coordinate system (r_i, θ_i) centred at the entrance. For each direction θ_i, we calculate the distance $r_{\max}(\theta_i)$ between the pole and the wall in the corresponding direction. As the maximum distance is different in each direction, we compare the normalised distance $r_{n_i} = r_i/r_{\max}(\theta_i)$, to figure out how far, between the entrance and the wall, participants went. Figure 10 shows the normalised distance to the walls with respect to the order of entrance. In comparison to OMC, the real results show a higher preference towards positions close to the wall. The more noisy N2N plot (but with no repetition) exhibits more varied strategies, and little influence of the walls.

Personal space. To further explore individual behaviours, we computed the final personal space for each agent, which is here defined as the area of each agent Voronoï cell (see Figure 9). The result is presented in Figure 10-b, with personal space presented with respect to the order of entering.

6.2 Discussion

In this experiment we explored the strategy of picking a position in a virtual room when being filled by a crowd. By analysing the position picking order, it is clear that emergent patterns appear, but differ between real and virtual ones. Pedestrians in the real-world experiment chose their position in a sequence of U-shape formations with outer, middle and inner layers (Figure 9-d). In contrast, pedestrian in our OMC experiment would rather fill the room from the back to the front in a succession of flat rows, more or less organised depending on participants (e.g., Figure 9-a and -b, as well as additional examples in supplementary material). Looking at all OMC results, the corner area near the entrance seems to be rarely used, which differs from real-life scenarios. This difference can be explained by hypothesis **H1**. We have two possible reasons in mind. First, real and virtual walls play different roles: real walls can for example be used to lean on, which is obviously not the case of virtual ones. This can explain the preference to stand close to them first in real experiments (U-shape formations). Also, in the OMC procedure, participants do not have to exit the room at the end of the trial: the VE disappears, and they move back to the initial position. This might decrease the importance of standing close to the entrance (that a U-shape formation will give). We also observed that OMC participants stood further from walls than in reality, suggesting that the perception of the distance to objects in VR might have had quite an important effect on final positions, at least those close to walls.

More generally, the evolution of the virtual pedestrians' personal space (Figure 10-b) indicates that during each S2N experience, participants usually left themselves a larger personal space than in real-life. Furthermore, once they found out that the first virtual agents had occupied a too large space for everyone to fit comfortably, they would overreact by enormously sacrificing their personal space. This is responsible for the fast dropping of personal area in the S2N results, and the relatively large personal space in the N2N result, which is not observed in real experiments.

The lack of variety also led to specific emergent patterns, validating **H2**. This can be seen in Figure 9-b, where the participant clearly applied a strict strategy to fill the room, ending up in a very regular squared pattern. The N2N procedure successfully breaks this regularity (Figure 9-c), and we can see some irregular strategies, such as participant #24 picking the corner position in the back of the room whereas other last participants more preferably picked positions next to the door. Nevertheless, emergent patterns resulted from the OMC procedure, validating **H3**.

7 GENERAL DISCUSSION

Methodology. In this work, our methodological approach consisted in replicating several existing experiments in VR with our OMC paradigm. We relied on existing analysis frameworks typically used in the crowd modelling community to evaluate our results according to common practice from this community. While this involved to rely on qualitative evaluations, we believe this approach was more relevant to provide a first evaluation of the OMC paradigm. Apart from some limitations discussed below, our conclusion is that we *successfully* replicated three experiments that involve, mostly, *unilateral interactions*. Successfully, because we were able to replicate the observation of some existing, already known and observed, emergent phenomena. This is a major difference with previous studies in VR on collective behaviours that focused on individual analyses of local variables (e.g., [8, 49, 50]). It is also clear that the biases we expected influenced to some extent the observed emergent phenomena. However they mostly affected the characteristics of those phenomena without preventing their natural emergence. We were also able to observe which characteristics were affected and how, and could provide interpretations based on some formulated hypotheses.

While our analyses relied on *qualitative analyses*, we also believe that *quantitative analyses* will be important to explore in the future. For instance, a quantitative analysis would ideally require to perform again the real experiments with the same group of participants experimenting OMC. Given that our goal was to evaluate the potentialities of the OMC paradigm, the effort required to conduct such a real experiment would have prevented us from exploring the variety of scenarios we could cover here. At this stage, it would also be relevant to quantitatively evaluate the effect of several specific biases on our results. One interesting direction would be for instance to replicate studies on the effect of known dependent variables (e.g., age, density, corridor width). E.g., does the effect of age on 1D-flow can be reproduced using OMC? While characteristics of emergent phenomena are biased, would the effect of dependent variables still be observed and exploitable?

S2N vs. N2N. Our results suggest that the variety bias seemed to have influenced the most the characteristics of emergent collective phenomena. More specifically, the S2N procedure led in some cases to a replication across all agents of some specific individual behaviour, that could be observed as a sort of pattern exaggeration. Such extreme observations, which were reported in the discussion of each experiment, included the formation of single or extremely distinct lanes (Figure 6-a and -b), or entering the virtual room in a overly-organised manner (Figure 9-b). Since N2N solves the effect of this bias efficiently, we could claim that OMC-N2N is the only viable solution here. However, S2N is more convenient in practice because N2N still requires to involve many participants to create new crowd motion datasets. Therefore, to further explore potentialities of OMC-S2N, we still have two main questions opened to be addressed in the future. First, would the average of data resulting from multiple S2N datasets be equivalent to a N2N dataset? In other words, could we deduce the characteristics of a phenomenon emerging from interactions between various people given multiple ones emerging from interactions with oneself? The second is about possibilities for a M2N procedure, where M participants would control N agents, while keeping M as small as possible compared to N in order to keep the advantage of recruiting few participants to generate large datasets. Future work would then investigate what would be the required M value to cancel the effect of the variety bias.

Character animation details. OMC agents replicate participants' behaviours in a limited way. E.g., we captured their full-body motions using a motion capture suit, but neglected their gaze activity, which is known to play a great role in human interactions, while difficult to study in practice. Other modalities, like sound, were not recorded. Therefore, we would like to explore the possible benefit of capturing more detailed behaviours and incorporating them in the VE, as well as to explore novel ways of making participants aware of the physical presence of characters (e.g., evaluating the benefits of rendering contacts using haptics [6]). Another option would be to edit recorded behaviours, for example to change agents' motion style, which may provide another solution to lower the effect of variety bias. This would also raise novel questions, e.g., would that prevent participants to understand that they are interacting with themselves? Would a single participant behave in a more varied way given that he/she visualizes more varied motions?

Generalising interactions. Did participants realise they were interacting with themselves? The gross answer is that half of them reported yes, while this proportion depended on scenarios (all participants for 1D-flow, 2 for bottleneck, and 3 for inflow). For them, it was clear that interactions would be limited to unilateral ones, and that visible agents would never respond to their actions. While this is a limitation of OMC that will be difficult to overcome, we believe a number of solutions can be explored in the future to generalise the type of interactions that can be performed with OMC. The first one is to decompose interactions into multiple sections of time where participants control alternatively agents, similar to what Osimo et al. [45] did for building conversations with one-self (which was one source of inspiration for OMC). The second solution would be to perform the S2N procedure multiple times. For example, in a dual interaction between agents A and B, let the participant controlling A, then B to unilaterally interact with A, and then A again to interact unilaterally with B, etc. This raises new important questions, such as whether multiple unilateral interactions can converge (or not) toward the result of a bidirectional interaction. The third solution is to use a collision avoidance techniques to override previously recorded trajectories when a character is on a short-term collision course with the user. This consists in making a local and interactive adaptation of the trajectory, only when necessary, to avoid unrealistic situations. This solution has however a clear drawback: as mentioned in the beginning of this paper, modifying trajectories in such a procedural manner might reduce the realism of the generated dataset. With three possible choices, we believe that those solutions will enable us to generalise the use of OMC for other types of scenarios, e.g., studying collision avoidance strategies, large-crowd interactions, etc.

Generalising scenarios. As previously mentioned, our study prioritised the replication of crowd scenarios mainly involving unilateral interactions. However, daily situations would include multi-lateral interactions as well, which makes scenarios extending beyond unilateral interactions our next target. While we anticipate new artefacts resulting from those scenarios, we have conducted a preliminary experiment exploring them to gain insights about their real nature and importance. We refer the reader to the supplementary material for a detailed description of these preliminary attempts. We summarise them and provide our conclusions hereafter. We designed 2 different scenarios which involved multi-directional interactions, using the same OMC-S2N setup. The first scenario consists in a *Crossing Circle*, where participants stand on the perimeter of a circle and are asked to reach the diagonally opposite position. We performed this scenario under two different density levels (12 or 24 characters around the circle). The second scenario consists in a 90° *Crossing Flow*, where two groups of 12 characters each cross at a 90° angle. We compared both the OMC results to real-world recordings of similar situations. Our main observations, which remain to be validated by a formal experiment with quantitative analysis, are the following. In multi-lateral interaction OMC scenarios, we observed that the participant displayed strategies similar to real-world ones. For example, users chose to cross the circle either by going straight through the circle centre, or by going around the centre, avoiding this area of higher friction between people. In other words, they used strategy

consisting in playing on speed or curvature of their trajectory to avoid collisions. In the *Crossing Flow* scenario, we observed that the participant passed often through the scene's edge to avoid interactions with others characters. This was not observed in real situations. Noticeably however, in both of these two scenarios, we observed that despite the complexity of these scenarios, the participant rarely collided with the virtual characters (no more than once per scenario). This is a very important observation, given that the two scenarios have different difficulties in terms of collision avoidance. For instance, in the *Crossing Circle*, characters arrive from a large number of directions, whereas in *Crossing Flow* the density might occasionally increase and decrease. These two features can both make difficult collision avoidance. These observations suggest that the OMC paradigm is promising to study situations of multi-lateral interactions, where participants seem to be able to appropriately react to complex cases. Meanwhile, although the effect need to be verified by future experiments, it's interesting to evaluate the use of local adjustments of previously recorded trajectories to make characters reactive to the user. With this, we may correct the tendency of users to unrealistically avoid zones of to high frictions that we observed in the crossing flow scenario.

8 CONCLUSION

In this paper, we explored the use of VR to overcome cost and logistics concerns raised by real crowd experiments, which are paramount in the crowd simulation community, and proposed the OMC paradigm. To evaluate this approach, we replicated three typical real-world experiments (1D-flow, bottleneck and inflow scenarios). For each experiment, we evaluated two versions of the paradigm, one where a single participant generates the whole dataset, and one where multiple participants generate it asynchronously. We compared the generated data with real data and our analysis shows that both versions of the OMC paradigm lead to crowd motion data with realistic emergent patterns, despite some quantitative differences. One novel issue we faced is the role of variety and unilaterality biases that were specific to our approach. We reached the same general conclusion than previous work that used VR to generate behavioural datasets: qualitative similarity in behaviours are observed, while quantitative differences also appear in the phenomena we considered. Furthermore, unlike real experiments which generally only capture global positions, our experiments resulted in novel VR crowd datasets, describing full-body movements as well as global trajectories in crowds, which we will share with the community in the hope to foster novel ideas and applications. While the OMC paradigm presents some limitations, it opens new opportunities for research on collective behaviours. Previous work demonstrated the benefit of VR to explore collective behaviours at the individual scale. We here open ways to overcome limitations imposed by autonomous agents capabilities used to populate VR scenes, and also enable observing collective large-scale phenomena.

REFERENCES

[1] J. Adrian, M. Boltes, S. Holl, A. Sieben, and A. Seyfried. Crowding and queuing in entrance scenarios: influence of corridor width in front of bottlenecks. *arXiv preprint arXiv:1810.07424*, 2018.

[2] J. Adrian, A. Seyfried, and A. Sieben. Crowds in front of bottlenecks at entrances from the perspective of physics and social psychology. *Journal of the Royal Society Interface*, 17(165):20190871, 2020.

[3] C. Armbrüster, M. Wolter, T. Kuhlen, W. Spijkers, and B. Fimm. Depth Perception in Virtual Reality: Distance Estimations in Peri- and Extrapersonal Space. *CyberPsychology & Behavior*, 11(1):9–15, Feb. 2008. doi: 10.1089/cpb.2007.9935

[4] J. N. Bailenson, J. Blascovich, A. C. Beall, and J. M. Loomis. Equilibrium theory revisited: Mutual gaze and personal space in virtual environments. *Presence: Teleoperators & Virtual Environments*, 10(6):583–598, 2001.

[5] T. Banton, J. Stefanucci, F. Durgin, A. Fass, and D. Proffitt. The perception of walking speed in a virtual environment. *Presence*, 14(4):394–406, 2005.

[6] F. Berton, F. Grzeskowiak, A. Bonneau, A. Jovane, M. Aggravi, L. Hoyet, A.-H. Olivier, C. Pacchierotti, and J. Pettre. Crowd navigation in vr: exploring haptic rendering of collisions. *IEEE Transactions on Visualization and Computer Graphics*, pp. 1–1, 2020. doi: 10.1109/TVCG.2020. 3041341

[7] F. Berton, L. Hoyet, A.-H. Olivier, J. Bruneau, O. Le Meur, and J. Pettre. Eye-gaze activity in crowds: Impact of virtual reality and density. In *2020 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pp. 322–331, 2020. doi: 10.1109/VR46266.2020.00052

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

[9] S. Cao, P. Wang, M. Yao, and W. Song. Dynamic analysis of pedestrian movement in single-file experiment under limited visibility. *Communications in Nonlinear Science and Numerical Simulation*, 69:329–342, 2019.

[10] S. Cao, J. Zhang, D. Salden, J. Ma, C. Shi, and R. Zhang. Pedestrian dynamics in single-file movement of crowd with different age compositions. *Phys. Rev. E*, 94:012312, Jul 2016. doi: 10.1103/PhysRevE.94.012312

[11] R. Carstens and S. Ring. Pedestrian capacities of shelter entrances. *Traffic Engineering, Inst Traffic Engr*, 1970.

[12] P. Charalambous and Y. Chrysanthou. The pag crowd: A graph based approach for efficient data-driven crowd simulation. In *Computer Graphics Forum*, vol. 33, pp. 95–108. Wiley Online Library, 2014.

[13] U. Chattaraj, A. Seyfried, and P. Chakroborty. Comparison of pedestrian fundamental diagram across cultures. *Advances in Complex Systems*, 12:393–405, 2009.

[14] J. Chen, S. Lo, and J. Ma. Pedestrian ascent and descent fundamental diagram on stairway. *Journal of Statistical Mechanics: Theory and Experiment*, 2017(8):083403, 2017.

[15] W. Daamen and S. Hoogendoorn. Qualitative results from pedestrian laboratory experiments. In *Pedestrian and evacuation dynamics 2003, Greenwich*, pp. 121–132. CMS Press, 2003.

[16] G. W. Evans and R. B. Howard. Personal space. *Psychological bulletin*, 80(4):334, 1973.

[17] T. Ezaki, K. Ohtsuka, M. Chraibi, M. Boltes, D. Yanagisawa, A. Seyfried, A. Schadschneider, and K. Nishinari. Inflow process of pedestrians to a confined space. *Collective Dynamics*, 1:1–18, 2016. doi: 10.17815/CD. 2016.4

[18] T. Ezaki, K. Ohtsuka, D. Yanagisawa, and K. Nishinari. Inflow process: A counterpart of evacuation. In *Traffic and Granular Flow'13*, pp. 227–231. Springer, 2015.

[19] T. Ezaki, D. Yanagisawa, K. Ohtsuka, and K. Nishinari. Simulation of space acquisition process of pedestrians using proxemic floor field model. *Physica A: Statistical Mechanics and its Applications*, 391(1-2):291–299, 2012.

[20] Z. M. Fang, L. X. Jiang, X. L. Li, W. Qi, and L. Z. Chen. Experimental study on the movement characteristics of 5–6 years old chinese children when egressing from a pre-school building. *Safety science*, 113:264–275, 2019.

[21] P. W. Fink, P. S. Foo, and W. H. Warren. Obstacle avoidance during walking in real and virtual environments. *ACM Transactions on Applied Perception (TAP)*, 4(1):2–es, 2007.

[22] J. J. Fruin. Pedestrian planning and design. Technical report, 1971.

[23] M. Gérin-Lajoie, C. L. Richards, J. Fung, and B. J. McFadyen. Characteristics of personal space during obstacle circumvention in physical and virtual environments. *Gait & posture*, 27(2):239–247, 2008.

[24] M. Geuss, J. Stefanucci, S. Creem-Regehr, and W. B. Thompson. Can i pass? using affordances to measure perceived size in virtual environments. In *Proceedings of the 7th Symposium on Applied Perception in Graphics and Visualization*, pp. 61–64, 2010.

[25] S. Gulhare, A. Verma, and P. Chakroborty. Comparison of pedestrian data of single file movement collected from controlled pedestrian experiment and from field in mass religious gathering. *Collective Dynamics*, 3:1–14, 2018.

[26] D. Helbing, L. Buzna, A. Johansson, and T. Werner. Self-organized pedestrian crowd dynamics: Experiments, simulations, and design solutions. *Transportation science*, 39(1):1–24, 2005.

[27] D. Helbing, A. Johansson, J. Mathiesen, M. H. Jensen, and A. Hansen. Analytical approach to continuous and intermittent bottleneck flows. *Physical review letters*, 97(16):168001, 2006.

[28] J. H. Hollman, R. H. Brey, T. J. Bang, and K. R. Kaufman. Does walking in a virtual environment induce unstable gait?: An examination of vertical ground reaction forces. *Gait & posture*, 26(2):289–294, 2007.

[29] S. Huang, T. Zhang, S. Lo, S. Lu, and C. Li. Experimental study of individual and single-file pedestrian movement in narrow seat aisle. *Physica A: Statistical Mechanics and its Applications*, 509:1023–1033, 2018.

[30] A. Jelić, C. Appert-Rolland, S. Lemercier, and J. Pettré. Properties of pedestrians walking in line – fundamental diagrams. *Phys. Rev. E*, 85:036111, 2012. doi: 10.1103/PhysRevE.85.036111

[31] A. Jelić, C. Appert-Rolland, S. Lemercier, and J. Pettré. Properties of pedestrians walking in line. ii. stepping behavior. *Phys. Rev. E*, 86:046111, 2012.

[32] S. Lemercier, A. Jelic, R. Kulpa, J. Hua, J. Fehrenbach, P. Degond, C. Appert-Rolland, S. Donikian, and J. Pettré. Realistic following behaviors for crowd simulation. In *Computer Graphics Forum*, vol. 31, pp. 489–498. Wiley Online Library, 2012.

[33] S. Lemercier, A. Jelic, R. Kulpa, J. Hua, J. Fehrenbach, P. Degond, C. Appert-Rolland, S. Donikian, and J. Pettré. Realistic following behaviors for crowd simulation. *COMPUTER GRAPHICS FORUM*, 31:489–498, 2012. doi: 10.1111/j.1467-8659.2012.03028.x

[34] W. Liao, A. Seyfried, J. Zhang, M. Boltes, X. Zheng, and Y. Zhao. Experimental study on pedestrian flow through wide bottleneck. *Transportation Research Procedia*, 2:26–33, 2014. The Conference on Pedestrian and Evacuation Dynamics 2014 (PED 2014), 22-24 October 2014, Delft, The Netherlands. doi: 10.1016/j.trpro.2014.09.005

[35] J. Liddle, A. Seyfried, W. Klingsch, T. Rupprecht, A. Schadschneider, and A. Winkens. An experimental study of pedestrian congestions: influence of bottleneck width and length. *arXiv preprint arXiv:0911.4350*, 2009.

[36] X. Liu, W. Song, L. Fu, and Z. Fang. Experimental study of pedestrian inflow in a room with a separate entrance and exit. *Physica A: Statistical Mechanics and its Applications*, 442:224–238, 2016.

[37] X. Liu, W. Song, L. Fu, W. Lv, and Z. Fang. Typical features of pedestrian spatial distribution in the inflow process. *Physics Letters A*, 380(17):1526–1534, 2016.

[38] J. Llobera, B. Spanlang, G. Ruffini, and M. Slater. Proxemics with multiple dynamic characters in an immersive virtual environment. *ACM Transactions on Applied Perception (TAP)*, 8(1):1–12, 2010.

[39] J. M. Loomis, J. M. Knapp, et al. Visual perception of egocentric distance in real and virtual environments. *Virtual and adaptive environments*, 11:21–46, 2003.

[40] M. Moussaïd, D. Helbing, S. Garnier, A. Johansson, M. Combe, and G. Theraulaz. Experimental study of the behavioural mechanisms underlying self-organization in human crowds. *Proceedings of the Royal Society B: Biological Sciences*, 276(1668):2755–2762, 2009.

[41] M. Moussaïd, M. Kapadia, T. Thrash, R. W. Sumner, M. Gross, D. Helbing, and C. Hölscher. Crowd behaviour during high-stress evacuations in an immersive virtual environment. *Journal of The Royal Society Interface*, 13(122):20160414, 2016.

[42] F. P. Navin and R. J. Wheeler. Pedestrian flow characteristics. *Traffic Engineering, Inst Traffic Engr*, 39, 1969.

[43] S. Older. Movement of pedestrians on footways in shopping streets. *Traffic engineering & control*, 10(4), 1968.

[44] A.-H. Olivier, J. Bruneau, G. Cirio, and J. Pettré. A virtual reality platform to study crowd behaviors. *Transportation Research Procedia*, 2:114–122, 2014. The Conference on Pedestrian and Evacuation Dynamics 2014 (PED 2014), 22-24 October 2014, Delft, The Netherlands. doi: 10.1016/j.trpro.2014.09.015

[45] S. A. Osimo, R. Pizarro, B. Spanlang, and M. Slater. Conversations between self and self as sigmund freud—a virtual body ownership paradigm for self counselling. *Scientific reports*, 5(1):1–14, 2015.

[46] J. Pettré, J. Ondřej, A.-H. Olivier, A. Cretual, and S. Donikian. Experiment-based modeling, simulation and validation of interactions between virtual walkers. In *Proceedings of the 2009 ACM SIGGRAPH/Eurographics symposium on computer animation*, pp. 189–198, 2009.

[47] X. Ren, J. Zhang, W. Song, and S. Cao. The fundamental diagrams of elderly pedestrian flow in straight corridors under different densities. *Journal of Statistical Mechanics: Theory and Experiment*, 2019(2):023403, 2019.

[48] R. S. Renner, B. M. Velichkovsky, and J. R. Helmert. The perception of egocentric distances in virtual environments-a review. *ACM Computing Surveys (CSUR)*, 46(2):1–40, 2013.

[49] K. W. Rio, G. C. Dachner, and W. H. Warren. Local interactions underlying collective motion in human crowds. *Proceedings of the Royal Society B: Biological Sciences*, 285(1878):20180611, 2018.

[50] A. Ríos and N. Pelechano. Follower behavior under stress in immersive vr. *Virtual Reality*, 24(4):683–694, 2020.

[51] A. Seyfried, M. Boltes, J. Kahler, W. Klingsch, A. Portz, T. Rupprecht, A. Schadschneider, B. Steffen, and A. Winkens. Enhanced empirical data for the fundamental diagram and the flow through bottlenecks. In W. Klingsch, A. Schadschneider, and M. Schreckenberg, eds., *Pedestrian and evacuation dynamics 2008*, pp. 145–156. Springer-Verlag Berlin, 2010. doi: 10.1007/978-3-642-04504-2_11

[52] A. Seyfried, O. Passon, B. Steffen, M. Boltes, T. Rupprecht, and W. Klingsch. New insights into pedestrian flow through bottlenecks. *Transportation Science*, 43(3):395–406, 2009.

[53] A. Seyfried, B. Steffen, W. Klingsch, and M. Boltes. The fundamental diagram of pedestrian movement revisited. *J. Stat. Mech.*, p. P10002, 2005.

[54] S. Sharma, S. Otunba, and J. Han. Crowd simulation in emergency aircraft evacuation using virtual reality. In *2011 16th International Conference on Computer Games (CGAMES)*, pp. 12–17, 2011. doi: 10.1109/CGAMES.2011.6000319

[55] X. Shi, Z. Ye, N. Shiwakoti, and O. Grembek. A state-of-the-art review on empirical data collection for external governed pedestrians complex movement. *Journal of Advanced Transportation*, 2018, 2018.

[56] R. Sommer. Studies in personal space. *Sociometry*, 22(3):247–260, 1959.

[57] B. Steffen and A. Seyfried. Methods for measuring pedestrian density, flow, speed and direction with minimal scatter. *Physica A: Statistical mechanics and its applications*, 389(9):1902–1910, 2010.

[58] J. Sun, S. Lu, S. Lo, J. Ma, and Q. Xie. Moving characteristics of single file passengers considering the effect of ship trim and heeling. *Physica A: Statistical Mechanics and its Applications*, 490:476–487, 2018.

[59] J. Was, B. Gudowski, and P. J. Matuszyk. Social distances model of pedestrian dynamics. In *International Conference on Cellular Automata*, pp. 492–501. Springer, 2006.

[60] P. Willemsen, M. B. Colton, S. H. Creem-Regehr, and W. B. Thompson. The effects of head-mounted display mechanics on distance judgments in virtual environments. In *Proceedings of the 1st Symposium on Applied Perception in Graphics and Visualization*, pp. 35–38, 2004.

[61] D. Wolinski, S. J. Guy, A.-H. Olivier, M. Lin, D. Manocha, and J. Pettré. Parameter estimation and comparative evaluation of crowd simulations. In *Computer Graphics Forum*, vol. 33, pp. 303–312. Wiley Online Library, 2014.

[62] K. Yamori. Going with the flow: Micro–macro dynamics in the macrobehavioral patterns of pedestrian crowds. *Psychological review*, 105(3):530, 1998.

[63] D. Yanagisawa, A. Tomoeda, and K. Nishinari. Improvement of pedestrian flow by slow rhythm. *Phys. Rev. E*, 85:016111, 2012.

[64] G. Zeng, A. Schadschneider, J. Zhang, S. Wei, W. Song, and R. Ba. Experimental study on the effect of background music on pedestrian movement at high density. *Physics Letters A*, 383(10):1011–1018, 2019.

[65] H. Zhao, T. Thrash, M. Kapadia, K. Wolff, C. Hölscher, D. Helbing, and V. R. Schinazi. Assessing crowd management strategies for the 2010 love parade disaster using computer simulations and virtual reality. *Journal of the Royal Society Interface*, 17(167):20200116, 2020.