

Exploring the Role of Expected Collision Feedback in Crowded Virtual Environments

Haoran Yun*

Jose Luis Ponton†

Alejandro Beacco‡

Carlos Andujar§

Nuria Pelechano¶

Universitat Politècnica de Catalunya, Spain



Figure 1: We explore the effect of collision feedback expectancy in users who move through groups of virtual agents. We tested nine conditions, differing in the combination of types of collision feedback, shown in (a), users *could expect* when colliding with a virtual agent: an audible complaint, a vibration provided by an arm-mounted vibrotactile device, and a physical bump with a real person. We designed three different scenarios: a narrow corridor, shown in (b), where collisions with virtual agents were inevitable, a corridor, shown in (c), where agent movement was easily predictable and collisions could be avoided by taking a detour around the group, and a street crossing, shown in (d), where virtual agents moved in two orthogonal directions. The avatar of the participant is highlighted in red.

ABSTRACT

An increasing number of virtual reality applications require environments that emulate real-world conditions. These environments often involve dynamic virtual humans showing realistic behaviors. Understanding user perception and navigation among these virtual agents is key for designing realistic and effective environments featuring groups of virtual humans. While collision risk significantly influences human locomotion in the real world, this risk is largely absent in virtual settings. This paper studies the impact of the expected collision feedback on user perception and interaction with virtual crowds. We examine the effectiveness of commonly used collision feedback techniques (auditory cues and tactile vibrations) as well as inducing participants to expect that a physical bump with a real person might occur, as if some virtual humans actually correspond to real persons embodied into them and sharing the same physical space. Our results indicate that the expected collision feedback significantly influences both participant behavior—encompassing global navigation and local movements—and subjective perceptions of presence and copresence. Specifically, the introduction of a perceived risk of actual collision was found to significantly impact global navigation strategies and increase the sense of presence. Auditory cues had a similar effect on global navigation and additionally enhanced the sense of copresence. In contrast, vibrotactile feedback was primarily effective in influencing local movements.

Keywords: Virtual Reality, Crowd Simulation, User Studies, Collision Avoidance

Index Terms: Computing methodologies—Computer graphics—Graphics systems and interfaces—Virtual reality, Perception;

*e-mail:haoran.yun@upc.edu

†e-mail:jose.luis.ponton@upc.edu

‡e-mail:alejandro.beacco@upc.edu

§e-mail:carlos.andujar@upc.edu

¶e-mail:nuria.pelechano@upc.edu

Human-centered computing—Human computer interaction (HCI)—HCI design and evaluation methods—User studies;

1 INTRODUCTION

In computer applications such as building evacuation planning, training, and video games, crowd simulation is becoming increasingly important, as it allows hundreds or thousands of agents to navigate through the virtual environment. Furthermore, a number of applications require the user to be immersed in such virtual environments. By using a VR headset, the user can be embodied into a virtual avatar [29], gaining an illusion of body ownership over it, feeling present (sense of presence) in the virtual world and having the experience of being part of the virtual crowd.

Presence can be seen as the addition of place illusion (PI) and plausibility (Psi) [39, 40]. PI is the illusion of being in the place depicted by the VR. Psi is the illusion that the virtual situation and events are really happening. While PI can be achieved in VR through natural sensorimotor contingencies integrated by the hardware on sensory outputs [24], Psi is a far more complex and interesting aspect to study since any unexpected event can produce a break in presence or a failure of expectations in the participant. Similarly, copresence refers to the extent to which a participant has the illusion of being there with other participants [10]. This can also be achieved through sensorimotor contingencies that give the participant the illusion of being in the same space as the other humans, whether real or virtual, as it happens with PI. To achieve Psi, it is also necessary that the virtual characters respond to the user when the user interacts with them. The most difficult requirement for copresence is meeting the user’s expectations, which largely depend on the context (for example, people behave differently in a park vs. in a bus) and appearance (cartoon characters will imply different expectations than realistic ones) [46]. When copresence is achieved, we can expect users to show similar behaviors to those in reality, such as maintaining the same proxemics [4].

Regarding virtual crowds, the breaks in presence or failures of expectations can be due to different factors such as visual or animation artifacts (e.g., foot sliding [27]), detecting virtual character clones [19], strange behaviors, or lack of interaction. If virtual

agents do not respond to the user's actions, the user might quickly notice this and cause a break in presence.

Conversely, we should expect that some actions or events happening in the virtual environment could also affect the user's behavior. In our case, we wonder how a virtual crowd may affect the user's behavior, and more specifically, how collisions with virtual agents may influence the user depending on different types of collision feedback users might expect: an audible complaint (audio), a vibrotactile feedback through an arm-mounted device, or a physical bump with a real person embodied in an avatar. The first two feedback modalities can be easily implemented with commodity VR hardware, so all collisions with a virtual agent could trigger any combination of these feedback modes. In contrast, providing realistic physical bump feedback for all collisions is impractical, as this would imply that all avatars correspond to physical users embodied in them. Instead, we designed the experiment so that users are induced (depending on the condition) to expect a physical bump with a real person, even if such bumps occur only in an initial adaptation scenario. Similar to the Rubber Hand Illusion [8], we physically touch the participants in synchrony with the virtual collision to induce such belief. We assume this adaptation will also induce a collision feedback expectancy from the point of view of the user. Therefore, our study involves mainly exploring the role of this expected collision feedback in crowded virtual environments. Our specific research questions are:

- **RQ1:** Does collision feedback expectancy (users induced to expect a combination of audio, vibrotactile, or physical bump feedback) enhance the **sense of presence and/or copresence** in tasks involving moving across a group of virtual agents?
- **RQ2:** Does collision feedback expectancy influence the **behavioral response** of individuals in the tasks above?

Our results show that expected collision feedback plays an important role in shaping participants' behavior and subjective perceptions towards virtual crowds. In the absence of it, people might tend to ignore other agents. However, believing there could be physical collisions with real people can enhance presence. Also, copresence with the virtual crowds can be enhanced by adding audio feedback. Similarly, global behaviors such as navigation can also be significantly influenced: participants will show a tendency to maintain extended distances from crowd agents and take more time to complete tasks. Contrarily, vibrations can affect local behaviors such as adding more torso motion to actively avoid collisions.

2 RELATED WORK

The existing body of literature on crowd behavior in VR environments encompasses several areas, including crowd perception, realistic crowd simulation, and collision feedback mechanisms. For the purpose of this paper, we divide related work into three broad categories: studies related to crowd perception in VR, proxemics and copresence, and providing collision feedback in virtual crowds.

2.1 Virtual Crowds

An essential factor in crowd simulation is the accurate representation and perception of crowds. Olivier *et al.* [22] establish the efficacy of VR as a tool for the study of crowds. Subsequent studies have delved into various crowd attributes that could influence users' perception in VR. Factors like the visual appearance of the crowd [21] and the realism and diversity of animations [20] have been shown to enhance the perceived authenticity of the virtual crowd, affecting how users navigate around other agents. However, it is worth noting that an exhaustive set of animations is not essential; minimal motion variation is adequate for maintaining the perception of crowd heterogeneity, especially in scenarios involving large crowds [1]. Additionally, the presence of crowds has been shown to influence the decision-making processes of users in a VR environment [35].

Raimbaud *et al.* [31] investigate the impact of agents' gaze on users, highlighting that gaze direction may affect those with social anxiety but not necessarily impact their locomotion. In this line, some works have focused on studying how the emotional state of a virtual crowd may affect the way people navigate and interact [26,44]. Adding basic social interaction can also increase crowd realism and user presence [15]. Similarly, responsive virtual crowd behaviors increase the feeling of presence [14].

Evidence of the necessity for realistic crowd behavior in VR comes from works that simulate real-world scenarios. Arias *et al.* [3] reconstructed the Love Parade disaster in VR to evaluate various disaster prevention strategies. Zhao *et al.* [45] used VR to test different fire evacuation systems for a building under construction, highlighting the cost-effectiveness of VR in such studies.

2.2 Proxemics and copresence

User's locomotion behavior when moving around other avatars in immersive VR can vary depending on several factors, such as the quality of the animations and the fear of a real collision happening. For example, the work by Ríos *et al.* [34] showed that when users walk around another virtual human while sharing the same virtual space (thus, there is the possibility of a real collision), proxemics could be very similarly to reality, as long as the animations of the avatars were continuous. Previous work observed larger proxemics or clearance not being respected when animation artifacts were present, thus leading to virtual collisions when users knew the other user was located remotely [28]. Trivedi *et al.* [43] found that a high avoidance radius in the simulated crowd agents led to longer paths and lower perceived realism, drastically reducing user experience. Some researchers have investigated the impact of the avatar's appearance on users and observed that photorealism can increase plausibility and presence, but that it does not seem to affect copresence [18]. In fact, their study showed that what increases copresence is the user-avatars interaction, for example, if the avatars gaze at the user.

2.3 Collision Feedback in VR

Understanding collision feedback in virtual environments is critical for simulating realistic crowd behaviors. The work by Olivier *et al.* [23] highlights that pedestrians only adjust their movements if they perceive a future risk of collision. This observation is particularly relevant to our study, as it suggests the necessity for exploring how expected collision feedback might affect collision avoidance behavior in VR crowds.

A common approach for simulating collision feedback is using haptics. Krogmeier *et al.* [12] used a haptic vest equipped with 70 haptic points and showed that haptic feedback enhances presence and embodiment. On the other hand, Berton *et al.* [5] reported that haptic feedback improved collision avoidance but did not significantly affect the sense of presence. Similarly, Koiliias *et al.* [11] demonstrated that users' behavior was affected when haptics was enabled; however, they failed to distinguish between random and accurate haptic feedback. Similarly, Krum *et al.* [13] studied how priming haptic rendering influenced social measures and found that it impacted subjective social experiences but not proxemics. Different types of haptic feedback also led to different levels of emotional response towards virtual crowds. More recently, Venkatesan *et al.* [36] observed that users reported a higher sense of realism and increased nervousness when exposed to haptic stimuli, as opposed to solely tactile stimuli or no feedback at all.

Audio cues have mainly been given as a sound informing the participant about a collision, for example, while navigating using different locomotion techniques [42]. Blom *et al.* introduced the soundfloor, an audio-haptic interface for providing virtual collision feedback with the ground, but their study showed no effect of such feedback on performance [6]. On a different kind of task, insertion

tasks (manipulating a 3D object to get through one or more virtual apertures), Lecuyer *et al.* also studied how adding haptic, visual, and audio feedback affected performance [16]. While none of the added cues improved the task completion time, it seemed that participants moved less when colliding and having additional feedback, as if paying more attention to the possible collisions, and therefore taking more time to complete their task.

Studying social VR, Reinhard and Wolf explored the usability of applications involving virtual avatars and how difficult some tasks can be when they block paths or occlude the participant’s view [32]. Although opposite from our goal, they propose a solution without collision but where the participant is allowed to walk through the other avatars. They then explore multimodal feedback related to that, such as hearing the heartbeat of the avatar you are walking through. In their results, they find that presence significantly increases when multimodal feedback is provided, referring to our real-world experience when bumping into someone, but their study was limited to a two-avatar configuration.

Previous research has primarily explored the impact of single or multimodal collision feedback, such as audio and haptics, on users’ interactions and perceptions within virtual crowds. Yet, there is still a lack of studies investigating the influence on users’ behavior and perception when they expect potential collisions with real individuals in a virtual crowd setting. This scenario presents a difference to the effects evoked by commonly employed collision feedback techniques. Furthermore, our research aims to bridge this gap by examining the combined effects of these feedback modalities on user experience.

3 USER STUDY: OVERALL DESIGN

This section outlines the main decisions regarding the user study: the independent variables we considered, and how we implemented them in the virtual reality application. The scenarios and the activities participants had to undertake within the scenarios are also detailed, along with other relevant aspects concerning the appearance and behavior of the virtual agents.

3.1 Factors: Collision Feedback Expectancy

In the real world, collision risk with other individuals is a critical factor that explains crowd behavior. This risk leads individuals to move cautiously and exhibit distinct behavioral patterns compared to isolated walking. To simulate this aspect in a VR environment, we believe it is necessary to induce in users a certain expectancy of collision feedback, influencing their actions and locomotion choices to align with real-world behavior. Existing methods often use auditory cues or tactile feedback to achieve this outcome. The ultimate goal of these methods is to establish a collision feedback expectancy, encouraging users not to disregard virtual characters but to perceive them as real entities sharing the same physical space (copresent) and thus push them to avoid collisions.

A basic crowd simulation in VR entails rendering diverse human avatars moving along specified paths. Although this alone may facilitate a sense of presence, we believe that it does not prevent users from disregarding the crowd and going through virtual characters. We hypothesize that introducing collision feedback expectancy will increase the sense of presence and copresence, and lead users to move more cautiously.

The following are the specific factors related to collision feedback expectancy examined in this study:

- (a) *AUDIO*: Positional audio (a clearly audible complaint) upon collision with a virtual character.
- (b) *VIBR*: Arm-specific vibration upon collision with a virtual character.

- (c) *COLBLF*: Inducing the feeling or belief that a collision risk with an actual human exists, showing a virtual panel and actually having a gentle physical collision during an adaptation phase.

When the *AUDIO* is enabled, a collision triggers positional audio cues emitted from the location of the involved virtual character. To maintain consistency while minimizing auditory repetition, three gender-specific audio samples were chosen, resulting in a total of six distinct audio cues. Examples of these samples include *Ouch!*, *Careful!*, and *Watch Out!*.

For *VIBR*, participants are equipped with two vibration devices on their upper arm exclusively activated upon collision with virtual avatars. The vibration is selectively initiated in the arm closest to the estimated collision point. Specifically, if a collision with an avatar occurs on the participant’s left side, the left armband will be activated, and correspondingly, collisions on the right arm trigger vibrations in the right armband.

Regarding the *COLBLF* factor, we hypothesize that by nurturing the belief that genuine collisions can occur, users would exhibit collision avoidance behaviors more consistent with real-life crowds. To strengthen this belief, during an adaptation phase, users bump into a real person’s arm when the *COLBLF* setting is enabled.

3.2 Simulation design

3.2.1 Main scenarios

We designed two main scenarios: a single corridor (*Corridor*) and a street crossing (*Crossing*). The corridor had a 4.5 m × 7.5 m walkable area, both physically and virtually, for participants walking in naturally. Similarly, for the Crossing scenario, users walked in the intersection of two corridors, which had a 4.5 m × 4.5 m walkable area. Virtual humans were spawned slightly outside the walkable area to allow users more time to observe the agent’s trajectories.

3.2.2 Crowd trajectory

The crowd simulation was implemented using the RVO2-Unity package, which incorporates the Optimal Reciprocal Collision Avoidance (ORCA) algorithm. The maximum speed was set to 0.75 m/s which is around 62.5 % of normal walking speed 1.2 m/s, according to the U.S. Manual on Uniform Traffic Control Devices. When navigating around other virtual humans, people tend to walk slower [34], especially if haptic feedback is provided then the walking speed drops to 0.4 m/s approximately [5]. Thus, a relatively low speed was chosen in our study to allow participants to have enough time to plan and navigate through the crowd. With these settings, along with the initial position, rotation, and target position, RVO2 generated a 2D global trajectory for each character.

In the Corridor scenario, we confined the spawn and target positions of the characters to ensure that participants had at least two path options across the moving agents. More precisely, the crowd design incorporated two distinct pathways: a 0.3-meter-wide *direct path* placed closer to the user, and a wider 2-meter-wide *detour path* that required a larger detour to avoid the incoming crowd, as shown in Figure 2 top. The direct path served as a more direct route to the goal but increased the likelihood of collision with the crowd agents. Conversely, the detour path offered a safer but slower route. Our hypothesis posits that in the absence of collision-related feedback, participants would be more likely to opt for the *direct path*.

In the Crossing scenario, two orthogonal directions were included resulting in a crossing between two corridors. Agents traversed from one end of the crossing to the other and reversed direction upon reaching the end. To achieve this behavior, four spawn positions were located at four ends of the crossing. Each agent was assigned a specific target. Once reaching the target, another target in the opposite direction was created. This scenario is intended to more closely mimic a real crowd environment in which individuals

move in multiple directions and have different objectives. Virtual agents were programmed to avoid colliding with each other and the participant only in the Crossing scenario.

3.2.3 Adaptation Scenario

Besides the two main scenarios described above, we also designed an adaptation scenario to serve during an initial phase conceived to accommodate the participant to the current experimental condition. In this scenario, the participant was placed in a narrow corridor alongside one notice board and two virtual humans standing with an idle animation. The task was to reach a designated target position, experiencing a collision with both virtual characters. The feedback associated with the given condition—a combination of *AUDIO*, *VIBR*, or *COLBLF*—was applied when these collisions occurred.

When the *COLBLF* factor was enabled, collision belief was induced through a visual warning and physical bumps in this adaptation scenario. The visual warning, “There may be REAL people in the crowd” in red, was displayed on the notice board, as shown in Figure 1 (b). The physical bump was facilitated by a confederate standing at the location of one of the two virtual characters in the real world. As participants passed by this position, the experimenter physically bumped into them softly. We ensured that the physical contact was gentle and safe.

In contrast, when the *COLBLF* factor was disabled, the notice board showed “There are NO real people in the crowd.” and the confederate was not standing in the virtual characters’s position.

It is noteworthy that a physical collision between the experimenter and participants occurred only once during the adaptation phase for each condition with *COLBLF* enabled.

3.2.4 Agents’ appearance

We selected a set of 10 character models from Adobe Mixamo [2], comprising 5 male and 5 female characters. The distribution of models across scenarios was as follows: 8 for Corridor and 10 for Crossing. We ensured an even representation of male and female characters in each scenario and no clones in the crowd. In order to achieve a sense of copresence while avoiding the uncanny valley, we opted for a cartoonish visual style for our virtual crowd characters. This stylistic choice was consistently applied to the virtual environment as well. Our decision is supported by prior research indicating that although photorealistic rendering can enhance the sense of presence, non-photorealistic environments are also effective in achieving high levels of user presence [18].

3.2.5 Agents’ animation

We used a data-driven method called Motion Matching to synthesize character motions [30]. This method uses a set of features, including positions and velocities of the feet, the velocity of the hips, and future trajectory directions and positions. Then this information is used as query vectors to search over an animation database for the best matches. The animations included in the dataset were captured using the Xsens Awinda system, with two actors (one female, one male) walking at varying speeds: slow, normal, and fast, as well as walking in different directions like forward, backward, turning, and in-place rotations. In contrast to using single or multiple walking cycles, this data-driven method provides characters with more responsive and realistic walking motions, avoiding potential motion artifacts like foot sliding (see accompanying video).

4 USER STUDY: DESIGN DETAILS

4.1 Apparatus

The experiment was conducted in an isolated laboratory room with 7.5m×4.5m free walking space. The VR application was developed in Unity 2021.3.12f1 and ran on a PC equipped with an Intel Core i7-12700, NVIDIA GeForce RTX 3070 Ti, 32 GB RAM. As a VR headset, we used an HTC Vive Pro, which has a resolution of 1440

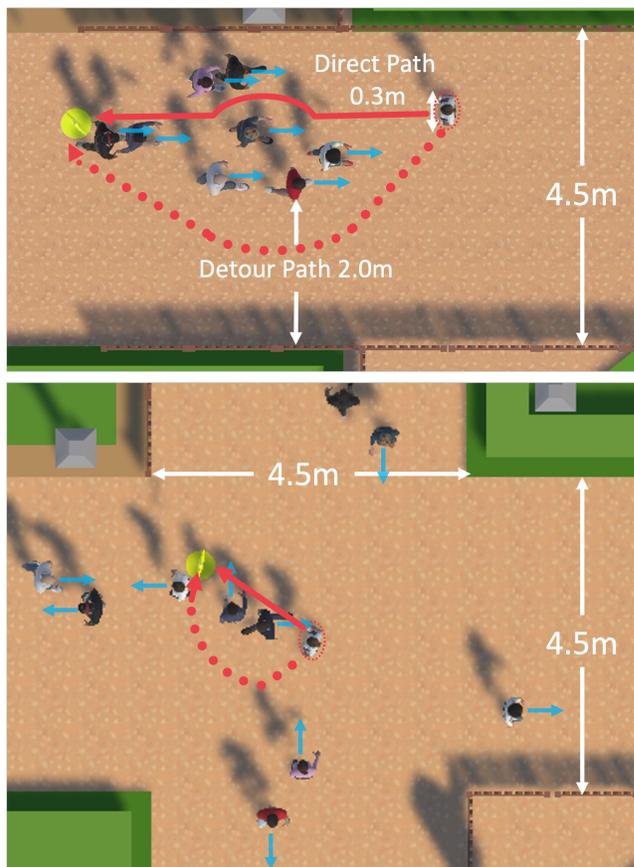


Figure 2: Corridor scenario showing the direct path with a continuous line and the detour path with a dotted line (*top*). Crossing scenario (*bottom*).

× 1600 pixels per eye, 110° field of view, 90 Hz refresh rate, and 3D spatial sound (using the HMD headphones). To provide users with a safe and spacious walking area, a Vive Wireless adapter was used instead of a wired connection. Four SteamVR Base Stations 2.0 were placed at each corner of the room to minimize potential occlusions. A wireless IMU-based motion capture system with 17 sensors, Xsens Awinda, was deployed to give users full-body animation in VR. Myo armbands, gesture control armbands from Thalmic Labs, generated a 1.5-second vibration for each collision, if *VIBR* factor was enabled.

4.2 Study Design

We adopted a within-subjects design, where each of the three factors had two levels, *enabled* and *disabled*, resulting in eight conditions. Participants sequentially completed the Adaptation, Corridor, and Crossing scenarios for each condition. An embodiment-inducing phase, where the participant could get used to the virtual avatar, was presented at the beginning of the experiment. Additionally, a pretest condition was also conducted with all three factors set to *disabled* prior to the main experiment. This pretest [7] allowed us to evaluate the sensitization effects, such as learning or after effects of the main factors. To counterbalance the order of conditions between participants, we used an 8 × 8 Balanced Latin Square. Crowd configurations and target positions remained consistent across different conditions and participants.

4.3 Procedure

A total of 16 participants took part in the experiment (6 female, 10 male, aged 22-39, $\bar{x} = 26.7$, $\sigma = 5.0$). On arrival, participants

Presence	
<i>(Reversed P1 + P2 + Reversed P3 + P4) / 4</i>	
P1	I did not feel present in the virtual space.
P2	I felt present in the crowd.
P3	I still paid attention to the real environment.
P4	Somehow, I felt that the virtual world surrounded me.
Copresence	
<i>(CP1 + CP2 + Reversed CP3 + CP4 + Reversed CP5) / 5</i>	
CP1	I perceived that I was in the presence of other people with me.
CP2	I felt that people were aware of my presence.
CP3	The thought that people were not real crossed my mind often.
CP4	The people appeared to be alive to me.
CP5	I perceived people as being only a computerized image, not as real people.

Table 1: Questionnaire content. The scores are on a 5-Likert scale (1 = strongly disagree, 5 = strongly agree).

were briefed via a document detailing the tasks they would perform, followed by completing a consent form and a demographic questionnaire. After the briefing, they were equipped with the two Myo armbands (one on each upper arm) and the Xsens motion capture suit. After calibrating the Xsens system, participants were equipped with the VR headset.

Next, participants underwent an embodiment phase, immersed in a virtual room with a mirror, allowing them to move and observe their virtual body freely for about three to four minutes. This phase facilitated inducing the sense of embodiment with the virtual body. Afterward, we started with the pretest, followed by the eight conditions of the main experiment, randomized as described above.

At the beginning of each condition, participants started in the middle of a narrow corridor in the Adaptation scenario, embodied in a gender-matched avatar. They were instructed to read the notice board first, then proceed toward a target location marked by a star-shaped object. Post-adaptation, participants transitioned to the Corridor and the Crossing scenarios.

In both scenarios, participants were instructed to walk toward a clearly visible target. Once reached, a new target and a new virtual crowd were generated. No instructions were given about how participants should behave in terms of trying to avoid collisions.

For the corridor scenario, each new target was created at the other end of the corridor, so that in each trial, participants had to walk through a new group of agents walking against them. The simulated group performed no collision avoidance, thus leading to inevitable collisions if the participant remained stationary. The goal of this simulated behavior was to study participants' route choices between the two visible paths left by the crowd formation (the *Direct Path* or the *Detour Path*). Regarding the crossing scenario, participants were positioned at the intersection between the corridors and were also asked to reach the designated target positions. In this scenario, simulated characters would perform collision avoidance against each other and participants.

At the end of each condition, they were required to complete a questionnaire within the VR environment (see Table 1), and then they could take a short break. After finishing all the conditions, they were asked to finish a short survey of open-ended questions (see Table 2).

4.4 Measures

4.4.1 Objective metrics

To objectively evaluate participants' behavior, we recorded diverse data from the simulation, including the trajectories of participants and agents on the ground plane within the virtual environment, the completion time for each task, and the number of collisions with

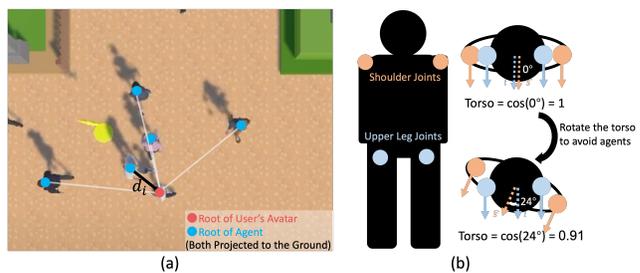


Figure 3: Visual illustration of *Clearance* and *Torso* measures. (a) The distances between the user avatar's root position and every agent's root position are calculated at timestamp i (projected onto the ground); d_i is the shortest distance at this timestamp. (b) *Torso* refers to the cosine of the angle, α , between the torso forward vector (in light salmon), \hat{s} , and the hips forward vector (in light blue), \hat{h} . The torso and hips forward vectors are computed from the forward vectors of shoulders and upper legs respectively. The *Torso* value changes when a user rotates the torso to avoid nearby agents.

virtual agents. Concurrently, we captured body motion using the Xsens motion capture system for further postural analysis (e.g., torso rotation). The collected data was further processed to obtain more intuitive metrics as follows:

- *Collisions* = $\sum_{i=1}^N c_i$, where N is the number of trials, and c_i refers to the number of detected contacts between the participant's avatar and the agents, for each trial.
 - *Time* = $\frac{1}{N} \sum_{i=1}^N t_i$, where N is the number of trials, and t_i refers to completion time for each trial.
 - *Clearance* = $\frac{1}{0.1T} \sum_{i=1}^{0.1T} d_i$, where T is the number of timestamps (with a rate of 50 per second) captured during the simulation. Let $\{d_0, d_1, \dots, d_T\}$ be the array of clearances per timestamp, sorted in ascending order. Here, d_i represents the distance between the root position of user's avatar projected onto the ground and the root of the nearest agent projected onto the ground at each timestamp, as depicted in Figure 3 (a). As the focus is on moments of minimum clearance and considering that the user may often be distant from the crowd, the average distance is calculated using only the lowest 10% of the samples.
 - *Torso* = $\hat{h} \cdot \hat{s}$, which provides the cosine of the angle, α , between the torso forward vector, \hat{s} , and the hips forward vector, \hat{h} . The torso forward vector \hat{s} is computed as the average of both shoulder joints forward vectors. Similarly, \hat{h} is computed from both upper leg joints forward vectors (this information is extracted from the motion capture suit), see Figure 3 (b). Similarly to *Clearance*, we order this value for all timestamps and consider only the lowest 10% of the samples. The cosine provides an intuitive estimate of the user intention to reduce the effective width, eW , which would be the result of projecting the vector between the shoulder joints onto a direction perpendicular to the displacement, and it can be calculated as $eW = \cos(\alpha) \cdot sW$, with sW being the shoulder width.
 - *Path* = $\begin{cases} 0 & \text{if the participant chose a direct path} \\ 1 & \text{if the participant chose a detour path} \end{cases}$
- This binary value is defined for each trial of the Corridor scenario. A manual labeling process was carried out post-experiment to classify participants' path choice—direct or detour—based on the recorded trajectories.

Open-Ended Question	n	\bar{x}	σ
Q1 What can you tell us about the overall experience?	886	0.09	0.36
Q2 How did you feel when people reacted to collisions with audio?	431	0.09	0.30
Q3 How did you feel when you felt vibrations in your arms?	471	-0.09	0.32
Q4 Believing you may collide with a real person, did it make you behave more like in real life?	433	0.21	0.36

Table 2: Post-experiment open-ended questions. We report the results of the sentiment analysis (i.e., the positivity of the answers) performed with the *sentimentr* R package [33]; n is the number of words analyzed, \bar{x} is the average sentiment, and σ is the standard deviation.

Due to the distinct characteristics of the Corridor and Crossing scenarios, specifically in terms of crowd simulation, we opted to analyze certain metrics separately for each scenario. Specifically, we examined $Collisions_{corridor}$, $Collisions_{crossing}$, $Time_{corridor}$, $Time_{crossing}$, $Clearance_{corridor}$ and $Clearance_{crossing}$. *Torso* was analyzed once per condition regardless of the scenario, and *Path* was measured only for the Corridor.

Furthermore, our analysis focuses on the interactions between users and nearby agents across all tasks. Notably, users often remain idle, awaiting subsequent tasks, or maintain a large distance from agents. Our initial strategy involved setting thresholds to omit irrelevant clearance values—particularly when users are significantly distant from agents. However, this approach proved insufficient, as certain participants consistently maintained clearance values above our predetermined threshold, thereby yielding no comparative data. Consequently, we adopted a more universally applicable method suitable for diverse tasks. This method involves organizing the data and selectively extracting the most pertinent percentage. Specifically, for the *Clearance* task, we extract the minimal distances to agents, and for tasks involving torso rotation, we focus on the maximal angle deviations. Upon analyzing the data distribution, we determined that the chosen 10% subset effectively balances comprehensiveness and relevance for our experimental needs.

Additionally, it is important to highlight that we conducted the statistical analysis both with and without the implemented cutoffs, encompassing all *Clearance* distances and *Torso* values. The results remained consistent in both scenarios.

4.4.2 Subjective responses

We also collected subjective ratings with a questionnaire at the end of each condition. The questionnaire included nine items, four addressing *Presence* and five addressing *Copresence*, as detailed in Table 1. Questions were adapted from the I-Group Presence Questionnaire (IPQ) [37, 38] and the copresence questionnaire by Bailenson *et al.* [4].

To further analyze participants’ experiences, we posed several open-ended questions concerning the overall experiment and each factor, as shown in Table 2. We encouraged participants to write down their opinions without limits in a text file. Responses were categorized into positive and negative sentiments using *sentimentr* R package [33].

5 RESULTS

In this section, we present an overview of the results obtained from the statistical analysis. Figure 4 shows the results of the metrics detailed in Section 4.4. For an in-depth discussion and interpretation of these findings, please refer to Section 6.

Shapiro-Wilk tests indicated significant deviations from normality in some instances. As a result, all analyses are carried out using non-parametric tests.

Additionally, we compare the differences between the pretest and all-factors-disabled conditions for all dependent variables. We employ two-sided Wilcoxon tests to compare between groups, and in the case of the dependent variable *Path*, we use a binomial generalized linear mixed-effects model. In all cases, we accept the null hypothesis. Hence, it seems that participants were not affected by sensitization effects such as learning, after effects of main factors (e.g., *VIBR*), or any other effect due to the repeated tasks.

5.1 Objective metrics

5.1.1 Task metrics: Collisions, Time, Clearance

To examine the within-subjects factors *AUDIO*, *VIBR*, *COLBLF*, we conducted a repeated measures ART analysis of variance (ANOVA). We report effect sizes using partial eta-squared (η_p^2). Following Cohen’s guidelines [9], these can be interpreted as follows: large (>0.14), medium (>0.06) and small (>0.01). The results indicated that users completed the Corridor scenario significantly slower (higher $Time_{corridor}$) when *AUDIO* was enabled ($F_{(1,105)} = 4.57$, $p = 0.035$, $\eta_p^2 = 0.04$). However, no significant effect was found for *VIBR* and *COLBLF* on these metrics. Additionally, the interaction between *VIBR* and *COLBLF* was significant ($F_{(1,105)} = 5.58$, $p = 0.020$, $\eta_p^2 = 0.05$). When both factors were enabled, users completed the Corridor scenario more slowly. Similarly, considering $Clearance_{crossing}$, we found that users significantly increased clearance when *AUDIO* was enabled ($F_{(1,105)} = 6.27$, $p = 0.014$, $\eta_p^2 = 0.06$; without 10% cutoff: $F_{(1,105)} = 4.46$, $p = 0.036$, $\eta_p^2 = 0.04$). We did not find significant differences for the other dependent variables.

5.1.2 Users’ motion: Torso rotation

We conducted the same tests for the motion-captured data variable *Torso*. In this case, we only found a significant main effect for *VIBR* ($F_{(1,105)} = 6.28$, $p = 0.014$, $\eta_p^2 = 0.06$; without 10% cutoff: $F_{(1,105)} = 5.90$, $p = 0.016$, $\eta_p^2 = 0.05$). Torso rotations were significantly higher when *VIBR* was enabled.

5.1.3 Path choice

Regarding the binary *Path* dependent variable, we conducted a binomial generalized linear mixed-effects model for within-subjects factors. For a complete understanding of the model and its interpretation, we report the results in Figure 5 and Table 3. The percentage of choices per condition can be seen in Figure 6. *AUDIO* and *COLBLF* influenced participants to choose the detour path in the corridor scenario. The interaction between *AUDIO* and *COLBLF*, i.e., when both factors were enabled, negatively affected the variable *Path*, and thus, influenced users to choose the direct path. Contrarily, the interaction between all factors positively affected *Path*.

5.2 Subjective responses

5.2.1 Presence and Copresence

We also report the results when considering the questionnaire answers. As before, we conducted a repeated measures ART ANOVA, using *Presence* and *Copresence* as dependent variables. For *Presence*, we only found a significant main effect for *COLBLF* ($F_{(1,105)} = 6.63$, $p = 0.011$, $\eta_p^2 = 0.06$). Regarding *Copresence*, results reported a significant main effect for *AUDIO* ($F_{(1,105)} = 9.51$, $p = 0.002$, $\eta_p^2 = 0.08$). Therefore, *COLBLF* increased the sense of presence, and *AUDIO* that of copresence.

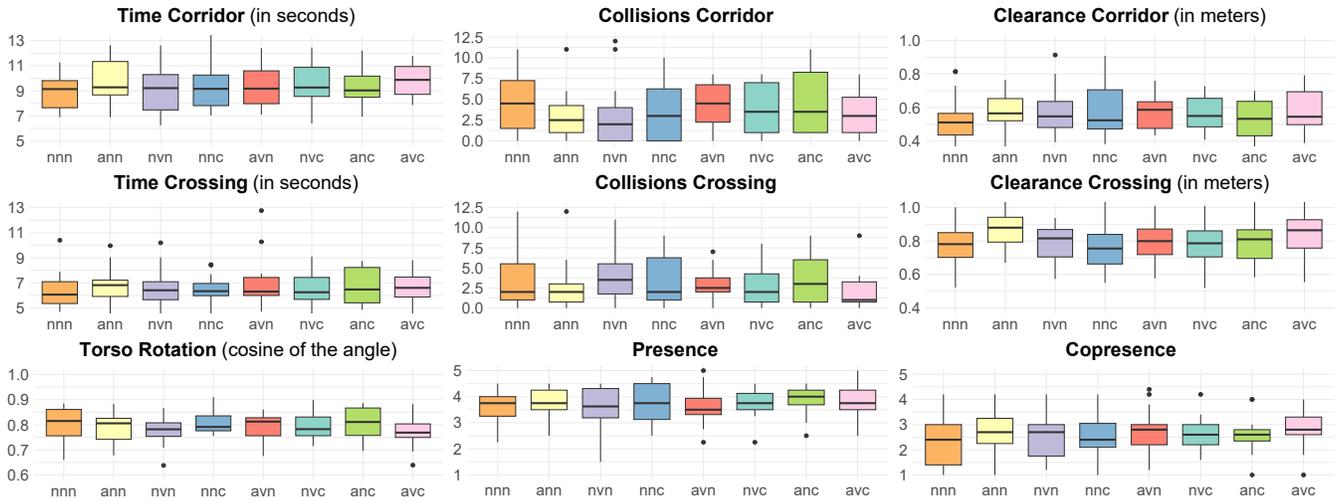


Figure 4: Box plots for all objective and subjective measurements, except *Path*. Each box encloses the middle 50% of the data for each condition. The thick horizontal lines denote the medians. *a* stands for enabled *AUDIO*, *v* for *VIBR*, *c* for *COLBLF*. *n* means the corresponding factor is disabled.

5.2.2 Open-ended questions

We now examine open-ended questions posed to participants upon the conclusion of the study. To assess the sentiment behind these responses, we employ sentiment analysis techniques [17, 25] as employed in recent virtual reality studies [41]. We used the *sentimentr* R package [33] to analyze the positivity of responses for each individual question. This package considers valence shifters, enabling us to assess positivity at the sentence level and subsequently average these values to obtain an overall sentiment score for each answer. We report the number of words analyzed n , the standard deviation σ , and the average sentiment \bar{x} in Table 2. Our analysis reveals an overall positive sentiment for questions **Q1** ($\bar{x} = 0.09$), **Q2** ($\bar{x} = 0.09$) and **Q4** ($\bar{x} = 0.21$). Conversely, negative sentiment was observed for question **Q3** ($\bar{x} = -0.09$). These results indicate that participants generally had a positive experience and found the auditory feedback to be positive. Notably, participants largely agreed with the statements in **Q4**, yielding a high average sentiment score. For **Q4**, which targeted a more specific question, we further conducted manual classification of the responses to discern whether participants agreed or disagreed with the statement. Consistent with the high sentiment score, 81% of participants answered affirmatively to **Q4**.

6 DISCUSSION

In this section, we revisit the original research questions of the study and analyze the results to summarize the main findings.

6.1 People response to virtual crowds

Expected collision feedback plays an important role in shaping participants' behavior towards virtual crowds. As detailed in Section 5.1, in the absence of expected collision feedback, it seems that people tend to walk closer to the virtual crowds and even pass through them. It becomes particularly evident in situations where participants are presented with both *Direct* and *Detour* pathways. When any form of collision feedback—*AUDIO* (odds ratio = 2.60), *VIBR* (odds ratio = 2.19), or *COLBLF* (odds ratio = 3.32)—is introduced, participants demonstrate a tendency to opt for the detour yet safer route (see Figure 6). In particular, we found *AUDIO* and *COLBLF* to significantly affect the path decision. Similarly, in the crossing scenario, an increased sense of caution is observed, with individuals maintaining larger *Clearance* from virtual characters. Notably, with *VIBR* enabled, participants rotate their torso to avoid

collision more often, or react to the vibrotactile feedback, aligning their actions with findings from prior studies [5, 11]. However, it is important to note that in the case of *Clearance* and *Torso*, we observed medium to small effect sizes. These results suggest the need for further studies to fully understand their implications.

We found no significant effect (for none of the three factors) on the *Collisions* metric, even though individual or combined factors showed noticeable impacts on other measures. This contrasts with a previous study by Berton *et al.* [5]. A potential reason for this, besides our smaller sample size, could be the dynamics of our experiment: while the prior research used a static crowd with more agents, giving participants enough time to build a strategy, our setup featured moving agents and a relatively smaller number of them, leading to fewer collisions overall.

Beyond behavioral changes, the expectation of collision feedback also appears to shape participants' subjective perceptions. The presented *Presence* and *Copresence* results underscore this, suggesting participants could have an increased sense of presence when believing there may be physical collisions with real people and also appear to be more copresent with the virtual crowds when they heard virtual characters make sounds during collisions. As before, we observed small to medium effect sizes.

6.2 Auditory and vibrotactile collision feedback

In Section 5.1, the results highlighted the distinct ways in which the *AUDIO* and *VIBR* factors shaped participants' avoidance behaviors.

The *AUDIO* factor significantly influenced *global* behaviors, which refer to high-level strategies participants use when walking towards the target. This was evidenced by participants' small to medium tendency to maintain larger distances from crowd agents and an increased task completion time.

Conversely, the *VIBR* factor appears to have affected *local* behaviors, i.e., immediate and specific reactions to nearby agents. For instance, the *Torso* metric suggests that participants rotated their torso to narrowly avoid nearby virtual characters, indicating that vibrotactile feedback may have played a role in their immediate, moment-to-moment decisions. Such findings align with previous work, which found that haptic feedback only caused people to rotate their torso more often but did not affect their *global* behaviors [5].

When participants provided feedback on their experiences, auditory cues received a generally positive response. We believe that

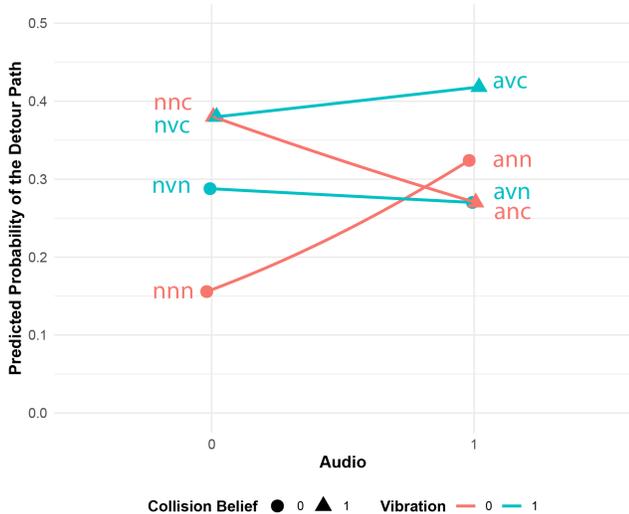


Figure 5: Predicted probability of the detour path from binomial generalized linear mixed-effects models of *AUDIO*, *VIBR* and *COLBLF* on *Path* (0 stands for factor disabled, 1 stands for factor enabled).

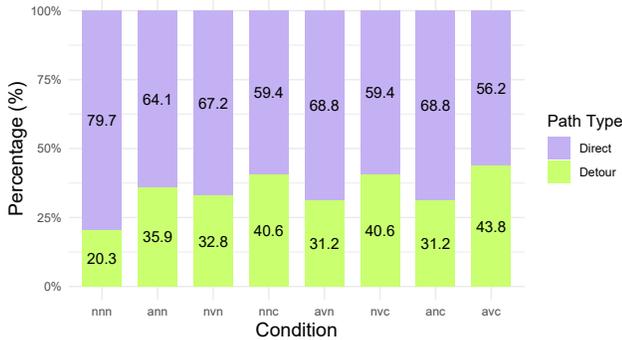


Figure 6: Percentage of time participants walked, in the Corridor scenario, using a detour path (green) and direct path (purple) for each condition.

voices like “Watch out!” and “Ouch” raised empathy towards virtual characters, prompting participants to move through virtual crowds more carefully instead of ignoring their presence. Conversely, vibrotactile feedback received more negative opinions. Some participants mentioned that the vibrations felt intrusive or startling, detracting from the immersive experience. Another point to ponder is the distinction between the haptic sensations produced by the armbands and the genuine tactile experience one might encounter during real-world collisions.

6.3 Perceived risk of collision

In the Corridor scenario, the *COLBLF* factor notably influenced participants towards choosing the detour path when presented with a group of agents offering both detour and direct path options. Additionally, this factor resulted in a significantly higher *Presence* score. This suggests that the presence feeling may have been intensified when participants believed real people existed within the virtual crowd. This may be explained because participants increased their focus on the task of avoiding virtual agents. Participants could only avoid the potential real human(s) by avoiding the virtual avatars. This sentiment was further corroborated in the positive post-experiment responses to question Q4. However, as noted in Figure 5, when both *COLBLF* and *AUDIO* are enabled, the detour path

Fixed Effect	Estimate	Std. Error	Odds R.	<i>p</i>
<i>(Intercept)</i>	-1.69	0.44	0.18	< 0.001
<i>AUDIO</i>	0.95	0.44	2.60	0.031
<i>VIBR</i>	0.78	0.45	2.19	0.08
<i>COLBLF</i>	1.20	0.44	3.32	0.006
<i>A:V</i>	-1.04	0.61	0.35	0.09
<i>A:C</i>	-1.46	0.60	0.23	0.015
<i>V:C</i>	-0.78	0.60	0.46	0.19
<i>A:V:C</i>	1.71	0.84	5.51	0.041

Table 3: Binomial generalized linear mixed-effects models of *AUDIO*, *VIBR* and *COLBLF* on *Path*. Estimate is the coefficient for the predictor in the logistic model. Std. Error is the standard error of the coefficient estimate. Odds R. represents the odds of the user choosing the detour path. Values are reported for enabled fixed effects, e.g., Odds R. in *AUDIO* refers to the probability of a user choosing the detour path when *AUDIO* is on.

choice probability decreases when compared to having only *COLBLF* enabled. Our best explanation for this is that when participants started listening to the virtual avatars, they may have realized no real person was walking in the crowd, thus breaking the collision belief. This does not happen when *VIBR* is enabled, significantly increasing the time to finish the corridor scenario. We believe having vibrotactile feedback further reinforces collision belief, as such vibration is somehow close to a gentle collision or graze with a real human. Nonetheless, this perspective is opposed to the negative feedback from participants concerning the plausibility of the vibrations. Thus, more research is needed to investigate this issue. In addition, while these interactions are statistically significant, it is important to note that the associated effect sizes are relatively small.

Interestingly, while the *COLBLF* factor appears to have influenced the global behavior of the participants and how immersed they felt, it did not seem to change their local behavior. This suggests that while participants might change their overall approach due to *COLBLF*, their natural reactions when moving through crowds remained consistent, regardless of whether they believed they might bump into real people in the virtual space.

7 CONCLUSIONS AND FUTURE WORK

In this paper, we have explored the effect of collision feedback expectancy on VR users moving through virtual crowds. For several reasons, we have considered explicitly the type of feedback the user expects vs. the type of feedback the user gets. First, this allowed us to compare against a novel factor (*COLBLF*) that is mostly based on what the users believe may happen (a collision with a real person embodied in one of the avatars) instead of the true feedback they will actually receive. Second, because actual collision feedback is always approximate: even if the full geometry of the participant avatar and the other avatars is taken into consideration to detect collisions (which is not typically the case for performance reasons), it is very unlikely that the self-avatar faithfully reproduces the shape and motion of the VR user, even if wearing a MoCap suit. This means that subtle collisions will not be detected by the system, and thus no collision feedback will be produced, no matter the condition (*AUDIO*, *VIBR*, *COLBLF*). In other words, although *COLBLF* provides no actual feedback during the experiment (unlike *AUDIO* and *VIBR*), in practice such a difference might go unnoticed unless the user makes absolutely no effort to avoid large collisions.

The main finding of our study is that the expected collision feedback significantly influences both the behavior of the participants and their subjective perceptions of presence and copresence. Our best explanation for the results distinguishes between metrics referring to global behavior (time, clearance, path) vs. local behavior

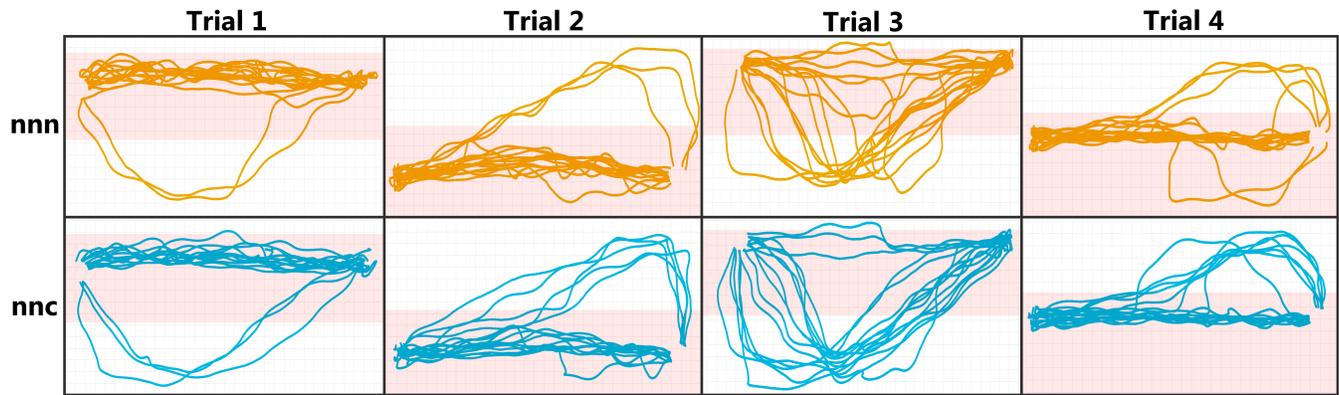


Figure 7: The trajectories taken by participants in the Corridor task across four trials are depicted, with the first row representing the condition where all factors are disabled and the second row highlighting the condition with only *COLBLF* enabled. Areas marked in red indicate high agent density, thus, the area where the *Direct Path* appears, whereas the white area indicates where the *Detour Path* appears. Notably, when *COLBLF* is activated, participants generally opt for detour paths to navigate through the crowd. Specifically, in the third trial, a visible empty space can be seen in the middle of the corridor in the second row, suggesting that participants actively avoided this area. In the fourth trial, in the condition without *COLBLF*, some participants chose to navigate toward the bottom of the corridor where the agent density is highest, a behavior that is absent when *COLBLF* is enabled.

(torso rotation). The introduction of a perceived risk of actual collision seems to significantly affect global behavior and increase the sense of presence. Auditory feedback, in the form of clearly audible complaints, appears to have a similar effect on global behavior and additionally increases the sense of copresence. In contrast, vibrotactile feedback was primarily effective in influencing local behavior, connected to immediate decisions.

We hope that these findings will help researchers and developers both in designing suitable VR environments involving virtual crowds, as well as anticipating the effect of collision feedback strategies on user behavior and experience.

Limitations: Since our sample size ($N=16$) was relatively small, we might have missed some significant effects. In addition, it should be noted that some statistically significant findings exhibit a small to moderate effect size. Notice also that the participants of our study were young adults, so some of our findings may not apply to the general population. As with many other user studies in VR, our results are closely dependent on the chosen scenarios, tasks, crowd behavior, as well as the cultural context of the participants. For example, significantly increasing crowd density and/or agent speed might considerably increase the physical effort to avoid collisions, and this might cause some users to decide to disregard collisions, especially if only audio or vibrotactile feedback is expected. Other aspects, such as the cognitive load of the task, might also influence the results. In our experiment, the task required a minimal cognitive load, since users just had to walk towards a clearly visible target. We guess that more complex scenarios involving, for example, search and/or wayfinding tasks could potentially lessen the effect of collision feedback expectancy.

An interesting avenue for future work is actually conducting further experiments to study which of our results apply to scenarios and tasks significantly departing from ours. For example, we are interested in knowing for how long the different collision feedback expectancy conditions have an effect on the measured variables, as we suspect that tasks with longer completion times might require periodic reinforcements of the expected collision feedback (as in the adaptation scenario), for example in the case of *COLBLF*. We also plan to explore more general strategies for inducing collision belief; although in some settings *COLBLF* is easy and even natural to adopt (for example, in collaborative VR applications with a mixture of human and NPC avatars), since the adaptation phase involves another person, it does not apply for example to VR applications

used at home.

ACKNOWLEDGMENTS

This work has received funding from the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 860768 (CLIFE project), and from MCIN/AEI/10.13039/501100011033/FEDER, UE (PID2021-122136OB-C21), “A way to make Europe”. Jose Luis Ponton was also funded by the Spanish Ministry of Universities (FPU21/01927).

REFERENCES

- [1] R. Adili, B. Niay, K. Zibrek, A.-H. Olivier, J. Pettré, and L. Hoyet. Perception of Motion Variations in Large-Scale Virtual Human Crowds. *Proc. MIG*, pp. 1–7. ACM, New York, Nov. 2021. doi: 10.1145/3487983.3488288 2
- [2] Adobe. Mixamo. <https://www.mixamo.com/>, 2023. Accessed: 2023-07-01. 4
- [3] S. Arias, S. La Mendola, J. Wahlqvist, O. Rios, D. Nilsson, and E. Ronchi. Virtual Reality Evacuation Experiments on Way-Finding Systems for the Future Circular Collider. *Fire Technology*, 55(6):2319–2340, Nov. 2019. doi: 10.1007/s10694-019-00868-y 2
- [4] J. N. Bailenson, J. Blascovich, A. C. Beall, and J. M. Loomis. Interpersonal distance in immersive virtual environments. *Personality and Social Psychology Bulletin*, 29(7):819–833, July 2003. doi: 10.1177/0146167203029007002 1, 6
- [5] F. Berton, F. Grzeskowiak, A. Bonneau, A. Jovane, M. Aggravi, L. Hoyet, A.-H. Olivier, C. Pacchierotti, and J. Pettré. Crowd Navigation in VR: Exploring Haptic Rendering of Collisions. *IEEE Transactions on Visualization and Computer Graphics*, 28(7):2589–2601, July 2022. doi: 10.1109/TVCG.2020.3041341 2, 3, 7
- [6] K. J. Blom, M. Haringer, and S. Beckhaus. Floor-based Audio-Haptic Virtual Collision Responses. In *Proc. JVR12*. The Eurographics Association, Madrid, Spain, 2012. doi: 10.2312/EGVE/JVR12/057-064 2
- [7] P. L. Bonate. *Analysis of pretest-posttest designs*. CRC Press, 2000. doi: 10.1201/9781420035926 4
- [8] M. Botvinick and J. Cohen. Rubber hands ‘feel’ touch that eyes see. *Nature*, 391(6669):756, Feb. 1998. doi: 10.1038/35784 2
- [9] J. Cohen. *Statistical Power Analysis for the Behavioral Sciences*. Routledge, New York, 2 ed., July 1988. doi: 10.4324/9780203771587 6

- [10] N. Durlach and M. Slater. Presence in Shared Virtual Environments and Virtual Togetherness. *Presence: Teleoperators and Virtual Environments*, 9(2):214–217, 04 2000. doi: 10.1162/1054746005667361
- [11] A. Koiliias, C. Mousas, and C.-N. Anagnostopoulos. I feel a moving crowd surrounds me: Exploring tactile feedback during immersive walking in a virtual crowd. *Computer Animation and Virtual Worlds*, 31(4-5):e1963, Sept. 2020. doi: 10.1002/cav.1963_2_7
- [12] C. Krogmeier, C. Mousas, and D. Whittinghill. Human, Virtual Human, Bump! A Preliminary Study on Haptic Feedback. In *Proc. IEEE VR*, pp. 1032–1033. IEEE, New York, Mar. 2019. doi: 10.1109/VR.2019.8798139_2
- [13] D. M. Krum, S.-H. Kang, and T. Phan. Influences on the Elicitation of Interpersonal Space with Virtual Humans. In *Proc. IEEE VR*, pp. 223–9. IEEE, New York, Mar. 2018. doi: 10.1109/VR.2018.8446235_2
- [14] M. Kyriakou and Y. Chrysanthou. How responsiveness, group membership and gender affect the feeling of presence in immersive virtual environments populated with virtual crowds. In *Proc. MIG*, pp. 1–9. ACM, New York, Nov. 2018. doi: 10.1145/3274247.3274509_2
- [15] M. Kyriakou, X. Pan, and Y. Chrysanthou. Interaction with virtual crowd in Immersive and semi-Immersive Virtual Reality systems. *Computer Animation and Virtual Worlds*, 28(5):e1729, Aug. 2017. doi: 10.1002/cav.1729_2
- [16] A. Lécuyer, C. Mégard, J.-M. Burkhardt, T. Lim, S. Coquillard, P. Coiffet, and L. Graux. The effect of haptic, visual and auditory feedback on an insertion task on a 2-screen workbench. In *Proc. IPT*. ACM, New York, 2002. 3
- [17] B. Liu. *Sentiment Analysis and Opinion Mining*. Synthesis Lectures on Human Language Technologies. Springer International Publishing, Cham, Switzerland, 2012. doi: 10.1007/978-3-031-02145-9_7
- [18] R. McDonnell, M. Breidt, and H. H. Bülthoff. Render me Real? Investigating the Effect of Render Style on the Perception of Animated Virtual Humans. *ACM Transaction on Graphics*, 31(4):1–11, jul 2012. doi: 10.1145/2185520.2185587_2_4
- [19] R. McDonnell, M. Larkin, S. Dobbyn, S. Collins, and C. O’Sullivan. Clone Attack! Perception of Crowd Variety. In *ACM SIGGRAPH 2008 Papers*, SIGGRAPH ’08. ACM, New York, 2008. doi: 10.1145/1399504.1360625_1
- [20] E. Molina, A. Ríos, and N. Pelechano. The Impact of Animations in the Perception of a Simulated Crowd. In *Proc. CGI*, pp. 25–38. Springer International Publishing, Cham, Switzerland, Sept. 2021. doi: 10.1007/978-3-030-89029-2_2_2
- [21] C. Mousas, A. Koiliias, B. Rekabdar, D. Kao, and D. Anastaslou. Toward Understanding the Effects of Virtual Character Appearance on Avoidance Movement Behavior. In *Proc. IEEE VR*, pp. 40–49. IEEE, New York, Mar. 2021. doi: 10.1109/VR50410.2021.00024_2
- [22] A.-H. Olivier, J. Bruneau, G. Cirio, and J. Pettré. A Virtual Reality Platform to Study Crowd Behaviors. *Transportation Research Procedia*, 2:114–122, Jan. 2014. doi: 10.1016/j.trpro.2014.09.015_2
- [23] A.-H. Olivier, A. Marin, A. Crétual, and J. Pettré. Minimal predicted distance: A common metric for collision avoidance during pairwise interactions between walkers. *Gait & Posture*, 36(3):399–404, July 2012. doi: 10.1016/j.gaitpost.2012.03.021_2
- [24] J. K. O’Regan and A. Noë. A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences*, 24(5):939–973, Oct. 2001. doi: 10.1017/S0140525X01000115_1
- [25] B. Pang and L. Lee. Opinion mining and sentiment analysis. *Foundations and Trends® in information retrieval*, 2(1–2):1–135, 2008. doi: 10.1561/1500000011_7
- [26] Y. Patotskaya, L. Hoyet, A.-H. Olivier, J. Pettré, and K. Zibrek. Avoiding virtual humans in a constrained environment: Exploration of novel behavioural measures. *Computers & Graphics*, 110:162–172, Feb. 2023. doi: 10.1016/j.cag.2023.01.001_2
- [27] N. Pelechano, B. Spanlang, and A. Beacco. Avatar locomotion in crowd simulation. *International Journal of Virtual Reality*, 10(1):13–19, Jan. 2011. doi: 10.20870/IJVR.2011.10.1.2796_1
- [28] I. Podkosova and H. Kaufmann. Mutual collision avoidance during walking in real and collaborative virtual environments. In *Proc. 13D*, pp. 1–9. ACM, New York, 2018. 2
- [29] J. L. Ponton, V. Ceballos, L. Acosta, A. Ríos, E. Monclús, and N. Pelechano. Fitted avatars: automatic skeleton adjustment for self-avatars in virtual reality. *Virtual Reality*, 27:1–20, July 2023. doi: 10.1007/s10055-023-00821-z_1
- [30] J. L. Ponton, H. Yun, C. Andujar, and N. Pelechano. Combining Motion Matching and Orientation Prediction to Animate Avatars for Consumer-Grade VR Devices. *Computer Graphics Forum*, 41:107–118, Mar. 2022. doi: 10.1111/cgf.14628_4
- [31] P. Raimbaud, A. Jovane, K. Zibrek, C. Pacchierotti, M. Christie, L. Hoyet, J. Pettré, and A.-H. Olivier. The Stare-in-the-Crowd Effect When Navigating a Crowd in Virtual Reality. In *Proc. SAP*, pp. 1–10. ACM, New York, Aug. 2023. doi: 10.1145/3605495.3605796_2
- [32] J. Reinhardt and K. Wolf. Go-through: Disabling collision to access obstructed paths and open occluded views in social vr. In *Proc. AHS*, pp. 1–10. ACM, New York, 2020. 3
- [33] T. W. Rinker. *sentimentr: Calculate Text Polarity Sentiment*. Buffalo, New York, 2021. version 2.9.1. 6, 7
- [34] A. Ríos, M. Palomar, and N. Pelechano. Users’ locomotor behavior in collaborative virtual reality. In *Proc. MIG*, pp. 1–9. ACM, New York, 2018. doi: 10.1145/3274247.3274513_2_3
- [35] A. Ríos and N. Pelechano. Follower behavior under stress in immersive VR. *Virtual Reality*, 24(4):683–694, Dec. 2020. doi: 10.1007/s10055-020-00428-8_2
- [36] V. RK, D. Banakou, M. Slater, and M. Muniyandi. Haptic feedback in a virtual crowd scenario improves emotional response. *Frontiers in Virtual Reality*, 4(1242587), Nov. 2023. 2
- [37] T. Schubert, F. Friedmann, and H. Regenbrecht. The experience of presence: Factor analytic insights. *Presence*, 10:266–281, 06 2001. doi: 10.1162/105474601300343603_6
- [38] T. W. Schubert. The sense of presence in virtual environments: A three-component scale measuring spatial presence, involvement, and realness. *Z. für Medienpsychologie*, 15(2):69–71, Apr. 2003. 6
- [39] M. Slater. Place illusion and plausibility can lead to realistic behaviour in immersive virtual environments. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 364(1535):3549–3557, Dec 2009. doi: 10.1098/rstb.2009.0138_1
- [40] M. Slater, D. Banakou, A. Beacco, J. Gallego, F. Macia-Varela, and R. Oliva. A separate reality: An update on place illusion and plausibility in virtual reality. *Frontiers in Virtual Reality*, 3:914392, June 2022. doi: 10.3389/frvir.2022.914392_1
- [41] M. Slater, C. Cabriera, G. Senel, D. Banakou, A. Beacco, R. Oliva, and J. Gallego. The sentiment of a virtual rock concert. *Virtual Reality*, 27(2):651–675, June 2023. doi: 10.1007/s10055-022-00685-9_7
- [42] E. A. Suma, S. Babu, and L. F. Hodges. Comparison of travel techniques in a complex, multi-level 3d environment. In *Proc. 3DUI*. IEEE, New York, 2007. doi: 10.1109/3DUI.2007.340788_2
- [43] H. Trivedi and C. Mousas. Human-virtual crowd interaction: Towards understanding the effects of crowd avoidance proximity in an immersive virtual environment. *Computer Animation and Virtual Worlds*, 34(3-4):e2169, May 2023. doi: 10.1002/cav.2169_2
- [44] M. Volonte, Y.-C. Hsu, K.-Y. Liu, J. P. Mazer, S.-K. Wong, and S. V. Babu. Effects of Interacting with a Crowd of Emotional Virtual Humans on Users’ Affective and Non-Verbal Behaviors. In *Proc. IEEE VR*, pp. 293–302, Mar. 2020. doi: 10.1109/VR46266.2020.00049_2
- [45] 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, June 2020. doi: 10.1098/rsif.2020.0116_2
- [46] K. Zibrek, E. Kokkinara, and R. McDonnell. The Effect of Realistic Appearance of Virtual Characters in Immersive Environments - Does the Character’s Personality Play a Role? *IEEE Transactions on Visualization and Computer Graphics*, 24(4):1681–1690, Jan. 2018. doi: 10.1109/TVCG.2018.2794638_1