

# Holonomic Collision Avoidance for Virtual Crowds

Rowan Hughes, Jan Ondřej, John Dingliana

Trinity College Dublin

---

## Abstract

*All approaches to simulating human collision avoidance for virtual crowds make simplifications to the underlying behaviour. One of the prevalent simplifications is to ignore its holonomic aspect (i.e. sidestepping, walking backwards). This does not, however, capture the full range of how humans avoid collisions. In real world scenarios we can often observe people sidestepping around each other and obstacles in their environment. In this paper we present a new holonomic collision avoidance algorithm for real-time crowd simulation. Our model is elaborated from experimental data, which allowed us to both observe the conditions under which holonomic interactions occur, as well as the strategies walkers use during such interactions to avoid collision. Our model is general enough to be used with other collision avoidance techniques. We validate our approach by reproducing situations from our experiments and we demonstrate several examples in which our method provides more plausible collision avoidance behaviour.*

Categories and Subject Descriptors (according to ACM CCS): I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence—Multi-Agent Systems

---

## 1. Introduction

One of the key components in creating believable, populated and vibrant virtual worlds is the creation of autonomous human agents that navigate their environment in a plausible manner. With the advent of more powerful graphics and computer hardware, the push for greater levels of realism has kept pace. Many modern video games contain large environments populated by A.I. driven crowds, the agents within which must behave in a realistic, or plausible, manner in order to maintain maximum user immersion within the environment.

With this in mind, in this paper we have developed an approach to add additional collision avoidance strategies to the repertoire of autonomous virtual agents. These strategies allow agents to more closely approximate some of the subtler aspects of true human collision avoidance, in particular we are interested in its holonomic aspect, that is to strafe or sidestep. These types of behaviours can be observed frequently within real crowd scenarios, particularly in dense, complex or constrained scenarios. In such situations pedestrians' time to predict and avoid collisions becomes constrained, this forces them to adopt emergency avoidance strategies.

Our main motivation is to investigate the nature of holonomic collision avoidance, based on observed behaviours in real scenarios. These types of behaviour have been entirely neglected up to now in approaches to collision avoidance, despite being a prevalent, observable feature in real human behaviour. Secondly, we aim to elaborate a model for simulating holonomic behaviour capable of synthesizing realistic trajectories for virtual humans. We believe that through the addition of holonomy we can add an extra level of dynamism and variety that can improve overall plausibility of virtual crowds.

We adopted an experimental approach. We assume that the holonomic aspect of human locomotion is mostly described by the lateral component of velocity and the strategy is mainly used as an emergency collision avoidance strategy when the time to avoid a collision becomes constrained. We therefore propose an experimental protocol to observe the conditions under which humans exhibit holonomic behaviour. We chose to motion capture several participants as they navigated through a set of obstacles during a series of laboratory experiments. These obstacles were modulated to constrain both the participants' time to collision and also the angle of their possible avoidance trajectories.

Our contributions include:




Figure 1: (Left) **Non-holonomic motion**: Direction of motion is supported by body orientation. (Right) **Holonomic motion**: Direction of motion and body orientation are decoupled.

1. An experimental study on the holonomic aspects of human collision avoidance. Using a motion capture set-up we recorded 10 participants as they navigated through a series of obstacles. We obtained a data-set that was used to inform us of the parameters of holonomic collision avoidance.
2. A new, general approach to adding holonomic collision avoidance strategies to the existing set for virtual humans in interactive multi-agent frameworks. For the purposes of this paper we have tested our technique as extensions to the models described by Pettré et al. [POO\*09] and Ondřej et al. [OPOD10] but it could be easily be incorporated into many other crowd simulation models.
3. We demonstrate the benefits of using a proxy object that more closely approximates human physiology, and how in combination with holonomic locomotion, agents are able to navigate highly constrained scenarios.

The rest of the paper is organized as follows: Section 2 will give a brief overview of the related work. Section 3 will give an overview of the experimental protocol we undertook in order to understand the nature of holonomic collision avoidance. Section 4 will present our approach, which extends the current set of strategies for collision avoidance in a velocity-based multi-agent framework. We will then present our main findings and results in Section 5, we compare our model with other examples as well as comparing our simulated trajectories to real data before Discussion and Conclusion.

## 2. Related Work

Collision avoidance has been extensively studied across a wide variety of fields including control theory, robotics, crowd simulation, etc. In the field of crowd simulation they fall, mainly, into four categories: reactive, rule-based, data-driven and geometric.

Helbing's social forces model is [HM95] an example of a reactive approach. The agents are simulated as a collection of velocity controlled particles each undergoing a sum of acceleration forces. The Helbing model has been extended upon by several subsequent works [HFV00, BM-dOB03, LKF05, PAB07].

In the seminal work described by Reynolds [Rey99] collisions are solved via a predictive approach. The future trajectories of the walkers are extrapolated and checked for potential imminent collisions. If a potential collision is detected, reactive accelerations are computed for the agents involved to avoid collision.

Data-driven approaches use example behaviours from video or motion capture data to drive the simulation of virtual characters. In the work of Lerner et al. [LFCCO07], a database of human trajectories is learnt from video recordings of real walkers. At each step during a simulation, each agent reacts to its state by searching the database and selecting a trajectory that most closely matches the agent's current state. Similarly, Lee et al. [LCHL07] used a regression-based learning algorithm in order to synthesize realistic group behaviours from crowd videos.

Geometric models are models that adapt the notion of the *velocity obstacle* [FS98]. In these approaches agents and static obstacles, are represented as obstacles in velocity space. More recently, van den Berg et al. [vdBLM08] introduced the *Reciprocal Velocity Obstacle (RVO)*, this technique helps to deal with the issue of unwanted oscillations. Building on the RVO technique, Guy et al. [GCK\*09] introduced a highly-parallel algorithm which uses a discrete optimisation method to greatly improve performance. Similarly, Berg et al. [BGLM11] further improve performance by allowing each agent to efficiently compute a collision-free path by solving a low-dimensional linear problem.

Our technique looks to extend the functionality of existing multi-agent frameworks. There has been a tremendous amount of other work that also seeks to add additional functionality and features, some examples of these include: Kim et al. [KGM13] present a technique to incorporate external physical forces, (i.e.) collisions, pushing, etc.), into a velocity-obstacle based collision avoidance framework. Lemerrier et al. [LJK\*12] present a method for simulating following behaviours in virtual crowds through control of agent velocity based on local pedestrian density. Curtis et al. [?] extends a velocity obstacle-based model through the use of line segments, derived from navigation data structures, rather than points as agent goals.

Holonomic locomotion, specifically for the purposes of multi-agent simulation has not been studied thus far. There have been a number of works in other fields that incorporate holonomic locomotion for simulated agents. Specifically, van Basten et al. [vBSE11] and Lee et al. [CKHL11] present work that includes motion interpolation schemes incorporating holonomic locomotion in order to animate char-

acters. While we adopt a similar animation system, the techniques that these approaches adopt to derive input to the animation systems would prove prohibitive for virtual crowds. Truong et al. [TFP\*10] similarly present a motion interpolation system that would conceivably work well using our model as input.

### 3. Experimental Study

**Objectives:** Collision avoidance is a major component of how humans navigate through their environment, therefore it is important to model it as plausibly and accurately as possible. Our objective is to assess the holonomic/non-holonomic nature of human collision avoidance. Specifically, we are interested in what separates holonomic from non-holonomic collision avoidance events. Previous experiments [ALHB08], demonstrate that the lateral velocities during navigation can be considered negligible enough to be ignored for simulation purposes. In sparse environments, this assumption largely holds true. Once the environment becomes dense, however, the time that an agent has to avoid collisions can be greatly reduced. In these scenarios, it is our hypothesis that this assumption fails. We designed our experimental protocol in order to limit the amount of time that a participant has to avoid a collision.



Figure 2: The experiment set-up in the motion capture studio.

**Protocol:** The proposed experimental protocol is illustrated in Figure 3. At the start of each experiment the participant was given a single instruction, to walk at a comfortable walking speed from the start position to the goal position between two obstacles. The experimental area was  $8m$  long and  $2m$  wide. Two static obstacles were placed in the path of the participant. The constrained area was modulated according to two parameters;  $l$  represents the vertical offset between the obstacles and  $d$  represents the horizontal offset of the obstacles. Overall there were 16 total combinations of  $l$  and  $d$  parameters and each combination was repeated 6 times. 10 subjects took part in these experiments, 8 male and 2 female. Each participant performed 96 trials in total, with the order of the trials being randomized across participants. In order to record the data we used a Vicon MX op-

tical motion-capture system. Trajectories were captured at a 120 Hz sampling rate. Once data has been captured, it must be reconstructed and processed manually in order to ensure complete trajectories throughout the trials.




Figure 3: Visualisation of the experiment protocol. Each red marker denotes the centre of an obstacle. The depicted grid is  $0.5m \times 0.5m$  in size. Some sample obstacle locations are shown.

**Method:** Participant trajectory was established from the mean of the two front hip markers,  $P(x, y)$ . Velocity,  $V = dP/dt$ , and acceleration,  $A = dV/dt$ , is noted in every frame. The data was filtered to remove noise and reduce the effect of natural oscillations, (using a Butterworth low-pass second order filter, 1Hz cutoff frequency). Trajectories are decomposed into three periods of time: participants start walking during the *initial phase*, during which they reach their comfort speed. The *interaction phase* starts when the participants is minimally close to the start obstacle, at time  $t = t_s$ , and ends when the participant is minimally close to the finish obstacle, at  $t = t_f$ . Finally, for  $t > t_f$ , the participant heads for their goal during the *recovery phase*. The purpose of our study was to examine locomotion during the *interaction phase*,  $t_s < t < t_f$ .

During the *interaction phase* we expect to observe three distinct locomotion strategies: walk with negligible holonomic motion, turns and sidesteps (or walk with significant holonomic motion). Often when examining human locomotion trajectories, lateral velocities are neglected and locomotion is assumed to be non-holonomic [ALHB08]. We ex-

plicitly decompose the velocity into its tangential and lateral components and take note of the normalised lateral velocity,  $V_l$ , at each frame.

**Experimental Results:** The first stage in development of our model for collision avoidance is to classify what constitutes a holonomic event. We have been able to see from observing the data that the lateral velocity profile for holonomic events is quite different from that of non-holonomic events. Holonomic events tend to exhibit a large, brief spike in lateral velocity that occurs towards the centre of the *interaction phase*. Turns, especially quick turns do show spikes in lateral velocity. These spikes, however, are much less brief and occur toward the end and/or beginning of the *interaction phase*. In order to determine a threshold for what constitutes a holonomic event we examine the maximum lateral velocity during the *interaction phase* for every trial, for every subject:

$$\max(V_l(t)), t \in t_s < t < t_f$$

we then use a simple iterative selection algorithm to find a threshold in the data. The threshold we glean from our data is,  $\tau = 0.47$ .  $\tau$  is the threshold that we consider to delineate a holonomic from a non-holonomic collision avoidance event. We make use of this parameter in building our computational model, described in Section 4.

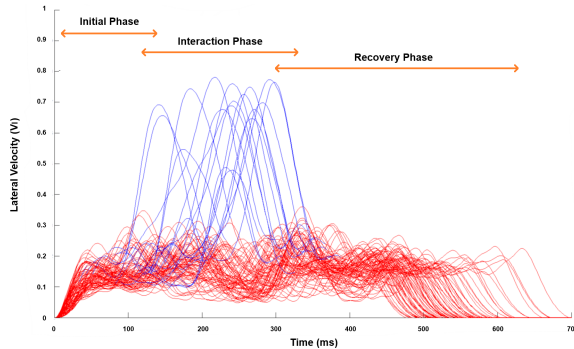


Figure 4: Lateral velocity as recorded for every trial for one participant. Events considered holonomic are marked in blue. Note trial start points are not synchronised in time. Peak of blue trajectories occur, approximately, at the same time.

**Other Observations:** Holonomic events are almost exclusively restricted to instances where the gap between the obstacles is close to or less than the shoulder width of the walker. Figure 4 shows all of the recorded trajectories for one participant. All of the trajectories that contain a holonomic avoidance event, coloured in blue, occur where the distance between the obstacles was 0.5 meters. 0.5 meters represents the closest distance between the obstacles (It is also a fair approximation of average human shoulder width). This suggests that holonomic motion is, primarily, restricted

to the cases when the participant cannot fit through the obstacles using non-holonomic motion.

#### 4. A Holonomic Model

Through our experimental investigation we now have an understanding of the conditions that elicit holonomic avoidance strategies in humans. In this section we present our technique for incorporating holonomy into collision avoidance models. Our solution is general and can be combined with many existing crowd simulation techniques.

**Holonomic Collision Avoidance:** In order to provide our agents the ability to move in a holonomic fashion they must have a method of controlling their orientation with respect to their velocity. We make use of the fact that we have understanding of the following:

- The agents desired velocity,  $V_d$ , is calculated from its current goal position.  $V_d$  is oriented toward this goal. Then we express this in local coordinates as a desired world velocity relative to the agent:

$$V_{d\ w/a} = -V_d.$$

- From our experiments we have observed that the lateral component of velocity,  $V_l$ , during a holonomic event lies between  $\tau < V_l < 1$ .

$$V_l = \text{acos}(\text{Agent}_{\text{velocity}} \cdot \text{Agent}_{\text{direction}})$$

- We can calculate the time to contact,  $ttc$ , to obstacles. This tells us how much time we have, given our current heading and velocity, before colliding with an obstacle. We observe in our experiments holonomic avoidance behaviour is largely an emergency strategy. We can therefore use  $ttc$  to determine whether or not to avoid a collision in a holonomic fashion. The formula for calculating  $ttc$  is:

$$ttc = P_{o/a} + V_{o/a}$$

where  $P_{o/a}$  and  $V_{o/a}$  are the obstacles position and velocity relative to the agent respectively.

Our technique can conceivably be incorporated, with minimal effort, into many multi-agent frameworks. For the purposes of this paper we adopt a geometric or velocity-obstacle based approach [FS98], with the only distinct difference being that we use a proxy object based more closely on human physiology in velocity obstacle construction.

Once a collision free solution velocity,  $V_{sol}$ , has been calculated we update the position and direction of the agent accordingly. In the non-holonomic case this is very straightforward, with the position of the virtual agent being updated by trajectory and speed. The direction of the agent always being orientated toward the current trajectory. Holonomic locomotion, however, requires a decoupling of velocity and direction. This distinction is a very important feature that most multi-agent frameworks do not take account of.

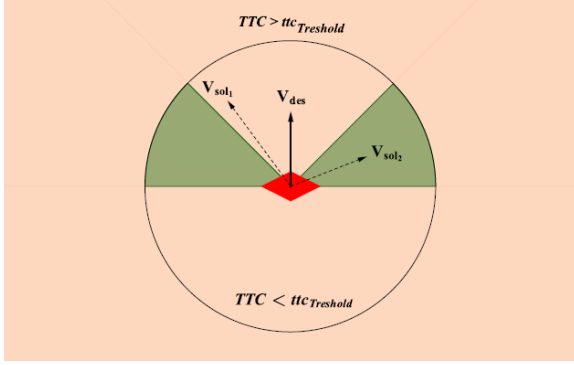


Figure 5: The circle represents the time to collision window within which a solution trajectory will be considered for classification as holonomic. The green and pink shaded areas represent the holonomic and non-holonomic solution spaces respectively.  $V_d$  is the agents desired velocity and  $V_{sol1}$ ,  $V_{sol2}$  represent sample solution trajectories. In these cases  $V_{sol1}$  would be classed as a non-holonomic trajectory with  $V_{sol2}$  being classed as holonomic

Our technique first determines whether or not a particular avoidance event should be considered holonomic. Given  $V_{sol}$ , we calculate  $V_l$ . If  $V_l \geq \tau$  and  $ttc_{min} \leq ttc_{threshold}$  then the agent initiates a holonomic collision avoidance event. If the agent is in a holonomic state then we suspend update of the agent's orientation, preventing the agent from turning.  $\beta_H = 0.6$  represents the time it take for an agent to perform one holonomic step, we discuss the reasoning for this check in the implementation section. Algorithmically:

---

**Algorithm 1** Compute Velocity
 

---

```

 $\vec{V}_{Sol} \leftarrow SolveInteraction()$ 
 $ttc_{min} \leftarrow \min(ttc_i), i \in \{0, \dots, \# Agents\}$ 
if  $Holonomic == true \ \& \ B_H < \beta_H$  then
  return
end if
 $\alpha \leftarrow \arccos(|\vec{V}_{sol}| \cdot \vec{D})$ 
if  $(\tau < \alpha < 1) \ \& \ (ttc_{min} < ttc_{Threshold})$  then
   $Holonomic \leftarrow true$ 
   $B_H \leftarrow 0$ 
end if

```

---

**A Suitable Proxy Object:** For the purposes of real-time crowd simulation, simplifications must be made to the morphology of virtual agents. Agents are commonly represented by simpler proxy geometry. By far the most common representation is a circular proxy object. We have observed in our experiments that humans can and do navigate through gaps that are smaller than their shoulder width. Humans accomplish this via holonomic strategies, exploiting the fact that humans are wider than they are deep. In order to take full advantage of holonomic behaviour and to reproduce our

---

**Algorithm 2** Update Position & Velocity
 

---

```

if  $Holonomic == true$  then
   $\vec{V} \leftarrow \vec{V}_{Sol}$ 
   $B_H += \Delta t$ 
else
   $\vec{V} \leftarrow \vec{V}_{Sol}$ 
   $\vec{D} \leftarrow |\vec{V}|$ 
end if
 $\vec{P} \leftarrow \vec{P} + \vec{V} * \Delta t$ 

```

---

experimental data we need a proxy object that reflects the irregular nature of human physiology.

In the work described by Pettré et al. [POO\*09] a *personal area* is set around the walkers, consisting of a kite shape. The kite shape was chosen as an approximation of the elliptical personal space described in [GLRLR\*05]. We modify this representation through the use of an irregular kite shape. This provides our walker with the ability to navigate spatially through gaps the former representation could not. Our representation attempts to incorporate both human physiology as well as an area of personal space around the agent. The kite's dimension from the centre to each side of the kite is 0.8m, 0.4m to the back and is velocity to the front. The distance from the centre to the front of the kite is calculated as:

$$0.4 + 0.4 \cdot v \cdot u_t$$

where  $v$  is the agents velocity and  $u_t$  is unit of time.

**Implementation:** In order to limit undesirable oscillations resulting from our technique we incorporated some additional constraints.

Human locomotion is driven by footsteps, our technique differentiates between two modes of locomotion. We therefore need to ensure that once a holonomic trajectory is undertaken that the virtual agent travels along that trajectory for at least a period of one complete footstep. In order to animate our characters we make use of a series of cyclic motion captured clips that form the input for a motion graph. We define the footstep period as the duration of one half-cycle, (i.e one footstep), of animation given the input parameters.

A resulting issue of implementing our footstep constraint was that simulated agents would sometimes select undesirable holonomic trajectories. As a new solution trajectory is computed at every step of simulation, occasionally a series of solutions will flicker between holonomic and non-holonomic. In order to solve this problem we filter forward by a number of frames in the simulation in order to ensure that unwanted holonomic trajectories are filtered out, resulting in a smoother overall trajectory.

## 5. Results

In this section we present the performance of our technique in various scenarios. We also analyse the technique and compare it with other techniques.

**Experimental Scenarios:** We demonstrate a series of virtual scenarios that correspond to examples from our experimental protocol described in Section 3. The purpose of these simulations was to ensure that we could plausibly recreate observed behaviour. We give the virtual agents three waypoints in each scenario. In that section we describe the *interaction phase* during which the participant interacted with the obstacles. The first two waypoints that we give to the virtual agents correspond to the start point and end points of the interaction phase given the obstacle configuration, the third waypoint is the agents goal. Figure 6 shows the results two example non-holonomic and holonomic simulations respectively.

**Dense Crowd Scenarios:** In these scenarios from two hundred to one thousand agents attempt to reach their opposite position on a circle. This scenario demonstrates that our approach is scalable to large virtual crowds, which can be seen in the accompanying media. Our approach has minimal performance impact and consistently leads to the agents solving the scenario in fewer simulation steps.

**Crossing Scenario:** In this scenario two large groups of agents intersect one another at a crossing, as illustrated in Figure 7. This scenario provides a more realistic dense crowd scene than the previous scenario, agents time to avoid collisions is greatly minimised but the scene is not so dense that the agents become so restricted that they can not reach their goals realistically. Using our approach we greatly reduce visible artifacts in agent behaviour, with agent navigating the scene in a more plausible manner.

**Corridor Scenario:** In this scenario we placed several groups of agents in a narrow corridor. This particular example is difficult due to relative agent density and the size of the obstacles in the scene. Using our approach the agents were able to navigate the corridor in a plausible fashion. Figure 8 demonstrates the holonomic behaviour of our agents and some of the issues encountered by other models.

## 6. Discussion

**Experiment** The experiment described in Section 3 was the second set of data that we worked with. The first set of data, though the protocol was near identical, did not produce any observed holonomic events. While the lateral velocity profile was what we might expect, the extrema of  $V_l$  during interaction with the obstacles was quite low. Upon examining the data more closely we realised that because the obstacles that were used during the first experiments were significantly shorter than the participants shoulder height, they did not have to reduce their shoulder diameter with respect to the obstacle. This is because the widest part of the participant, the

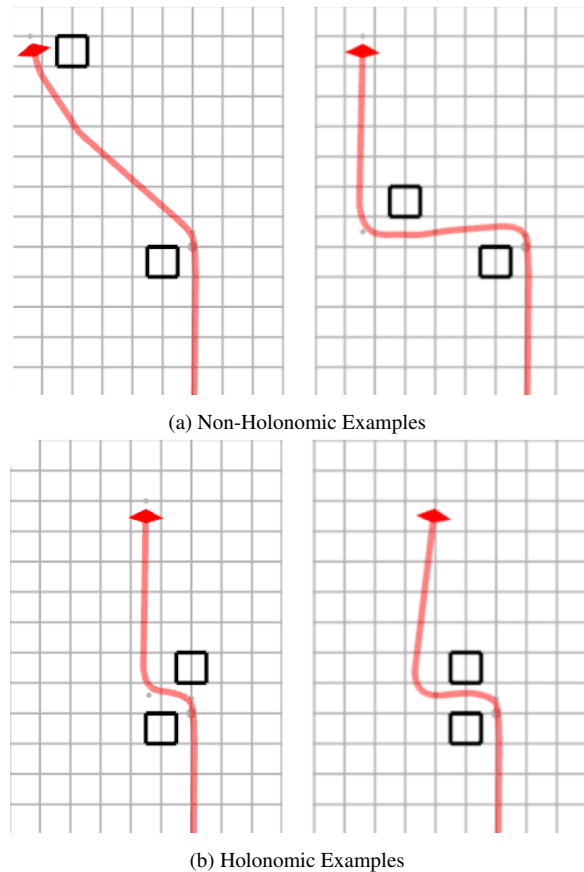


Figure 6: **Experiment Re-creation:** (a) In these examples, the virtual characters navigate through the obstacles in a non-holonomic fashion. In both of these cases the obstacles are places far enough apart that the agent can successfully turn while avoiding collisions. (b) In these cases the obstacles do not allow agent to navigate through in a non-holonomic fashion. The agent successfully sidestep through the obstacles before travelling to the goal.

shoulders, was not going to collide with the obstacle, so the amount the participant needed to turn in order to comfortably avoid the obstacle was reduced. This further reaffirms our intuition that humans only exhibit holonomic collision avoidance when necessary to avoid a collision.

Our experiments were limited to examining human interactions with static obstacles. While we were able to derive a set of parameters as to what constitutes a holonomic event under these scenarios, it remains an open question as to whether these parameters hold true in scenario with moving obstacles. As future work we plan to carry out a series of experiments examining holonomy in dynamic scenarios. In addition holonomic motion does not simply consist of side-stepping, moving backwards is also an example of holo-

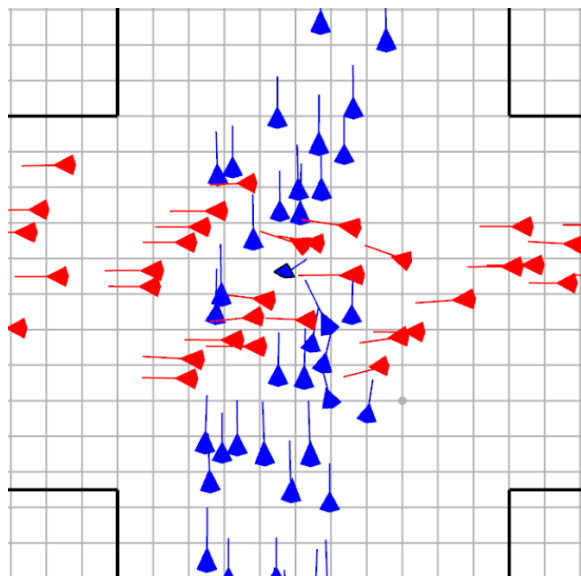


Figure 7: **Crossroads**: In this scenario two large groups of agents intersect one another at a crossroads. Using our method, agents avoid one another elegantly even in dense scenarios.

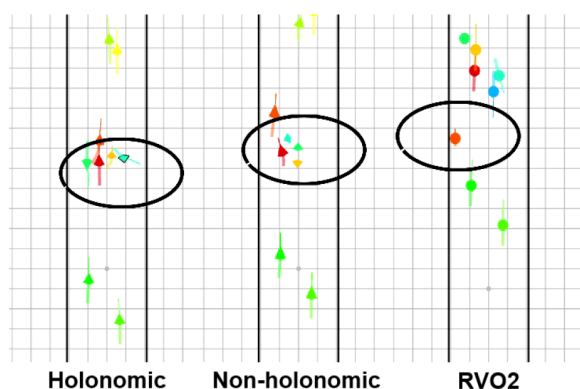


Figure 8: **Corridor**: (Left) Our technique allows the agents to avoid bottlenecks and navigate plausibly to their goal. (Middle) Without holonomic strategies the agents become stuck and navigate in an implausible manner. (Right) Using this model some agents can not find a solution in a timely manner and become fixed in place until other agents have left the scene. (see accompanying media)

nomonic motion, it would be interesting to understand the conditions under which this type of behaviour can be observed.

**Animation** One of the goals of our research was to create more realistic and varied agent collision avoidance. With this in mind it is critical that the animation system be able to handle the addition of holonomic behaviours in a seamless manner. As described in Section 2 there has been some work

in creating 3-dimensional motion graphs that incorporate holonomic behaviour. Due to some dependencies with the animation system we used, we blend between two distinct 2-dimensional motion graphs. One that blends straight line walking with turning behaviour and one that blends straight walking with side-stepping behaviour. While our animation system does a reasonable job of animating transitions there are some noticeable artifacts such as foot skating, particularly at low speeds. We are in the process of building a true 3-dimensional motion graph, following the method described by Truong et. al [TFP\*10] that should largely eliminate these issues. Going forward we plan to investigate perceptually the effect of our technique, it is therefore essential to have robust animation.

**Use with other collision avoidance models** Our solution consists of two main components: decomposition of velocity from direction and an asymmetrical agent representation. To incorporate the first component of our algorithm with other velocity-based methods (other methods as well) is straightforward, as it uses a final result of collision avoidance algorithm (i.e. the solution velocity). However, to fully benefit from holonomic behaviour an asymmetrical personal area is necessary. Without it agents can still side-step but it will not extend the solution space of the model (i.e. go through narrow areas).

**Extended solution space** The use of asymmetrical area with side-stepping allowed us to extend the solution space of our model. The agent is able to go through narrow areas, (i.e. the space between the two obstacles shown in Figure 6), which would not be possible previously. However, to keep the algorithm general and efficient we do not try to re-orientate the personal area to fit narrow spaces. In the case where an agent is facing a narrow space our approach would not provide a realistic solution and the agent would stop. We see this as an important area for future work, allowing virtual agents to find "smarter" solution trajectories.

## 7. Conclusion

In laboratory experiments we observed that holonomic behaviour in humans is prominent in constrained environments where the time to avoid collisions is small. Based on these observations we present a simple, yet powerful approach to simulating holonomic collision avoidance strategies in large multi-agent frameworks. Our method allows virtual agents to navigate, holonomically, in highly constrained environments in a plausible manner.

In future work we intend to perform a series of experiments to understand the nature of holonomic navigation in dynamic and multi-agent scenarios. As previously mentioned, side-stepping is simply one aspect of holonomic locomotion. We intend to extend the range of holonomic strategies to provide a broader range of plausible agent behaviour. We further intend to extend our approach to allow

agents to select holonomic solution trajectories regardless of agent orientation, thus reducing the frequency of undesirable agent trajectories.

## 8. Acknowledgements

This research has been funded by the European Commission, Seventh Framework Programme, through project VERVE (ICT, grant: 288914).

## References

- [ALHB08] ARECHAVALETA G., LAUMOND J.-P., HICHEUR H., BERTHOZ A.: On the nonholonomic nature of human locomotion. *Auton. Robots* 25, 1-2 (Aug. 2008), 25–35. URL: <http://dx.doi.org/10.1007/s10514-007-9075-2>, doi:10.1007/s10514-007-9075-2. 3
- [BGLM11] BERG J., GUY S., LIN M., MANOCHA D.: Reciprocal n-body collision avoidance. In *Robotics Research*, Pradalier C., Siegwart R., Hirzinger G., (Eds.), vol. 70 of *Springer Tracts in Advanced Robotics*. Springer Berlin Heidelberg, 2011, pp. 3–19. URL: [http://dx.doi.org/10.1007/978-3-642-19457-3\\_1](http://dx.doi.org/10.1007/978-3-642-19457-3_1), doi:10.1007/978-3-642-19457-3\_1. 2
- [BMdOB03] BRAUN A., MUSSE S., DE OLIVEIRA L., BODMANN B.: Modeling individual behaviors in crowd simulation. In *Computer Animation and Social Agents, 2003. 16th International Conference on* (2003), pp. 143–148. doi:10.1109/CASA.2003.1199317. 2
- [CKHL11] CHOI M. G., KIM M., HYUN K., LEE J.: Deformable motion: Squeezing into cluttered environments. *Comput. Graph. Forum* 30, 2 (2011), 445–453. 2
- [CSM12] CURTIS S., SNAPE J., MANOCHA D.: Way portals: Efficient multi-agent navigation with line-segment goals. In *Proceedings of the ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games* (New York, NY, USA, 2012), I3D '12, ACM, pp. 15–22. URL: <http://doi.acm.org/10.1145/2159616.2159619>, doi:10.1145/2159616.2159619. 2
- [FS98] FIORINI P., SHILLERT Z.: Motion planning in dynamic environments using velocity obstacles. *International Journal of Robotics Research* 17 (1998), 760–772. 2, 4
- [GCK\*09] GUY S. J., CHHUGANI J., KIM C., SATISH N., LIN M., MANOCHA D., DUBEY P.: Clearpath: highly parallel collision avoidance for multi-agent simulation. In *Proceedings of the 2009 ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (New York, NY, USA, 2009), SCA '09, ACM, pp. 177–187. URL: <http://doi.acm.org/10.1145/1599470.1599494>. 2
- [GLRLR\*05] GÄLLRIN-LAJOIE M., RONSKY J. L., LOITZ-RAMAGE B., ROBU I., RICHARDS C. L., MCFADYEN B. J.: The negotiation of stationary and moving obstructions during walking: anticipatory locomotor adaptations and preservation of personal space. *Motor Control* 9 (2005), 242–269. URL: <http://www.sciencedirect.com/science/article/pii/S0966636206005728>, doi:http://dx.doi.org/10.1016/j.gaitpost.2006.11.209. 5
- [HFV00] HELBING D., FARKAS I., VICSEK T.: Simulating dynamical features of escape panic. *Nature* (2000), 2000. 2
- [HM95] HELBING D., MOLNÁR P.: Social force model for pedestrian dynamics. *Phys. Rev. E* 51 (May 1995), 4282–4286. URL: <http://link.aps.org/doi/10.1103/PhysRevE.51.4282>, doi:10.1103/PhysRevE.51.4282. 2
- [KGM13] KIM S., GUY S. J., MANOCHA D.: Velocity-based modeling of physical interactions in multi-agent simulations. In *Proceedings of the 12th ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (New York, NY, USA, 2013), SCA '13, ACM, pp. 125–133. 2
- [LCHL07] LEE K. H., CHOI M. G., HONG Q., LEE J.: Group behavior from video: a data-driven approach to crowd simulation. In *Proceedings of the 2007 ACM SIGGRAPH/Eurographics symposium on Computer animation* (Aire-la-Ville, Switzerland, Switzerland, 2007), SCA '07, Eurographics Association, pp. 109–118. URL: <http://dl.acm.org/citation.cfm?id=1272690.1272706>. 2
- [LFCCO07] LERNER A., FITUSI E., CHRYSANTHOU Y., COHEN-OR D.: Behavior-graph for crowd simulation. In *ACM SIGGRAPH 2007 Posters* (New York, NY, USA, 2007), SIGGRAPH '07, ACM. URL: <http://doi.acm.org/10.1145/1280720.1280815>, doi:10.1145/1280720.1280815. 2
- [LJK\*12] LEMERCIER S., JELIC A., KULPA R., HUA J., FEHRENBACH J., DEGOND P., APPERT-ROLLAND C., DONIKIAN S., PETTRÉ J.: Realistic following behaviors for crowd simulation. *Comput. Graph. Forum* 31, 2 (2012), 489–498. 2
- [LKF05] LAKOBA T. I., KAUP D. J., FINKELSTEIN N. M.: Modifications of the helbing-moln&#225;r-farkas-vicsek social force model for pedestrian evolution. *Simulation* 81, 5 (May 2005), 339–352. URL: <http://dx.doi.org/10.1177/0037549705052772>, doi:10.1177/0037549705052772. 2
- [OPOD10] ONDŘEJ J., PETTRÉ J., OLIVIER A.-H., DONIKIAN S.: A synthetic-vision based steering approach for crowd simulation. In *ACM SIGGRAPH 2010 papers* (New York, NY, USA, 2010), SIGGRAPH '10, ACM, pp. 123:1–123:9. URL: <http://doi.acm.org/10.1145/1833349.1778860>, doi:10.1145/1833349.1778860. 2
- [PAB07] PELECHANO N., ALLBECK J. M., BADLER N. I.: Controlling individual agents in high-density crowd simulation. In *Proceedings of the 2007 ACM SIGGRAPH/Eurographics symposium on Computer animation* (Aire-la-Ville, Switzerland, Switzerland, 2007), SCA '07, Eurographics Association, pp. 99–108. URL: <http://dl.acm.org/citation.cfm?id=1272690.1272705>. 2
- [POO\*09] PETTRÉ J., ONDŘEJ J., OLIVIER A.-H., CRETUAL A., DONIKIAN S.: Experiment-based modeling, simulation and validation of interactions between virtual walkers. In *Proceedings of the 2009 ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (2009), SCA '09, ACM, ACM, pp. 189–198. URL: <http://doi.acm.org/10.1145/1599470.1599495>, doi:10.1145/1599470.1599495. 2, 5
- [Rey99] REYNOLDS C.: Steering behaviors for autonomous characters, 1999. 2
- [TFP\*10] TRUONG T.-V.-A., FLAVIGNE D., PETTRE J., MOMBAUR K., LAUMOND J.-P.: Reactive synthesizing of human locomotion combining nonholonomic and holonomic behaviors. In *Biomedical Robotics and Biomechatronics (BioRob), 2010 3rd IEEE RAS and EMBS International Conference on* (Sept 2010), pp. 632–637. doi:10.1109/BIOROB.2010.5626817. 3, 7
- [vBSE11] VAN BASTEN B. J. H., STÜVEL S. A., EGGES A.: A hybrid interpolation scheme for footprint-driven walking synthesis. In *Proceedings of Graphics Interface 2011* (School of



Computer Science, University of Waterloo, Waterloo, Ontario, Canada, 2011), GI '11, Canadian Human-Computer Communications Society, pp. 9–16. URL: <http://dl.acm.org/citation.cfm?id=1992917.1992920>. 2

[vdBLM08] VAN DEN BERG J., LIN M. C., MANOCHA D.: Reciprocal velocity obstacles for real-time multi-agent navigation. In *IEEE INTERNATIONAL CONFERENCE ON ROBOTICS AND AUTOMATION* (2008), IEEE, pp. 1928–1935. 2