

# CrowdAR Table

## An AR system for Real-time Interactive Crowd Simulation

Noud Savenije  
Utrecht University  
Utrecht, the Netherlands  
n.h.savenije@students.uu.nl

Roland Geraerts  
Utrecht University  
Utrecht, the Netherlands  
r.j.geraerts@uu.nl

Wolfgang Hürst  
Utrecht University  
Utrecht, the Netherlands  
huerst@uu.nl

**Abstract**—Spatial augmented reality, where virtual information is projected into a user’s real environment, provides tremendous opportunities for immersive analytics. In this demonstration, we focus on real-time interactive crowd simulation, that is, the illustration of how crowds move under certain circumstances. Our augmented reality system, called CrowdAR, allows users to study a crowd’s motion behavior by projecting the output of our simulation software onto an augmented reality table and objects on this table. Our prototype system is currently being revised and extended to serve as a museum exhibit. Using real-time interaction, it can teach scientific principles about simulations and illustrate how such they, in combination with augmented reality, can be used for crowd behavior analysis.

### I. INTRODUCTION

Interactive crowd simulation software [1] has the potential of being a valuable educational tool for teaching people about the complexities of managing crowds. Examples for which the analysis of such complex data is needed, include crowd behavior in emergency situations (e.g., evacuation of people in a building) or when a certain crowd behavior is desired - such as the social distancing efforts due to the recent COVID-19 outbreak, which requires business owners, municipalities and even countries to ensure that people are never forced to form dense crowds. Yet, to let the public interact with such a comprehensive system and use it for complex data simulation and analysis, we need an interactive interface that is easy to operate, robust, accessible and intuitive to understand. We developed a prototype of an augmented reality installation to enable an immersive data analysis, see Fig 1. Here, simulations are projected in real-time onto a table and the objects placed there. Users can manipulate the simulation (e.g., change parameters such as number of simulated pedestrians and spawn points) as well as directly manipulate the environment (e.g., by placing obstacles onto the table). This demonstration paper describes our current development and extension of this prototype into an installation that will function as a museum exhibit. A number of other augmented reality tables exist for use in a scientific or educational context, for example [2], [3], [4]. There are also many interactive tabletop museum exhibits [5], [6] and even ones using AR [7]. The CrowdAR [8] table differentiates itself by combining robust real-time interaction



Fig. 1. Example of an AR table installation for crowd simulation.

with a complex simulation. This paper will give an update regarding the development of the CrowdAR table. First the requirements for and benefits of this installation in a museum context will be discussed, and second we will go into the implementation of the system and its individual components.

### II. CROWDAR IN A MUSEUM CONTEXT

The CrowdAR table is being developed for the University Museum in Utrecht, which is a museum that focuses on scientific exhibitions. The simulation framework behind the CrowdAR system is developed by our research group. The framework, among other thing, handles the creation of the walkable environment and the global path planning, collision avoidance and animation for the agents in the simulation. Affordances are an important part of designing interactive exhibits. People will know how to interact with an AR sandbox [9], even if they have never seen it before. The environment offers user-centric interaction to the user and the AR aids in the education by projecting extra information in the form of heightmaps or floodplains, for example. The target audience

for this museum varies significantly – ranging for example from older kids to senior citizens. The CrowdAR table should therefore offer an accessible interactive interface suitable for a museum exhibit. It should be clear from the environment which actions can be used to interact with the simulation. Therefore, the museum wants to make use of both physical and virtual objects, known as mixed reality, on this AR table. Being able to project an area people know, like the area around the museum, a running simulation should help people understand what they are looking at and how they can influence this world. Visitors should thus be able to place physical obstacles on the table surface representing buildings or roadblocks in the simulation, mixing the virtual and physical spaces. There are two main goals for this exhibit. The visitors should be educated about the complexities of dynamic crowds. They should gain awareness of the difficulties of quickly altering an entire living space so that everyone can be far enough apart from each other, and also certain design features can act as bottlenecks when trying to evacuate a building or festival area for example. Second, the museum also focuses on engaging the public in the process of science itself, and connecting them to actual research. Thus, the installation should also be used to teach people how simulations are used for scientific research. It should help give the users an intuition for the limits of models and the difficulty of verifying data for example. Due to its educational purpose, the presentation should be interactive and since most visitors will only interact with this installation for a couple of minutes, it should be intuitive and accessible. The installation itself should be able to deal with crowds, that is, large numbers of visitors as we anticipate larger groups, such as school classes, where each individual wants to explore the system and play with it. This, and the varying age ranges, also pose high requirements on the robustness of the implementation, since we cannot expect kids to always follow certain rules or behavior when interacting with it. People will use it in various ways, including things that push its limits and have not been anticipated by the system designers.

### III. IMPLEMENTATION

A prototype of the CrowdAR was developed last year[8]. This prototype gave confidence in the feasibility of the project. The past year was spent optimizing the solutions put forth, making the system both more robust and optimized. Such optimization is required to provide more complex and sophisticated simulations and interactions. In the following section, we describe the individual parts of the installation.

#### A. Physical Setup

The setup consists of *three* distinct parts: A projector, a depth camera and a table. The *table* has to be white in order to be able to project the largest range of colors and it also has to be a convex quadrilateral in this installation. Our table is approximately 2m by 1m. The *projector* has a full HD resolution, meaning that each pixel projected on the table has a minimal size of  $1mm^2$ . The *depth camera* is the Azure Kinect [10] and the framework also supports the discontinued *Kinect*



Fig. 2. Result of calibrating the projector.

v2 as well as the *Intel Realsense*. Since this is only using one projector and one camera, there will always be shadows we have to deal with. These problems are exacerbated if the projector and sensor are not perpendicular to the table surface. When reaching the edges of the table the shadows inevitably become larger. The use of lower blocks at the edges makes this problem less bad.

#### B. Calibration

Calibration of the system is required to enable communication from the sensor to the simulation and the projector. This calibration is done in two distinct steps. First, we calibrate the sensors in the camera (color, infra-red and depth) using contour detection. For all three sensors, the corners are extracted from the largest contour by finding the longest diagonals from the convex hull of this contour. In order to detect the contours, contrast is needed between the table surface and the background. For the color calibration, this is done by projecting a bright white image on the white table surface. For the other cameras, the table surface needs to be at a different height from the background.

The calibration of the projector space is somewhat more involved. All corners need to be found individually in three steps. First three lines are drawn from predetermined starting points, then the angle of these lines are adjusted until they intersect the corner. The intersection of these lines results in the corner point. Both the camera space and the projector space are mapped using OpenCV's `warpPerspective` function.

#### C. Real-time Object Detection

Users of the installation can place physical objects on the table to act as obstacles for the pedestrians in the simulation. In order for these objects to be recognized by the system, they need to be at least  $3cm$  high, wide and deep. Since the simulation is done on the flat surface of the table, it is sufficient to detect the contours of each object to correctly detect the situation at surface level. To filter out noise and the moving arms of the visitors, it seems best not to update as often as possible. Instead, every object is only counted when it is stationary for a short while. This is done on a pixel




Fig. 3. Objects detected shown in our simulation framework.




Fig. 4. The contours of the detected blocks are being projected back on the table. The contours are only accurate on ground level due to the lack of projection mapping, which takes the blocks' height into account

level by testing if an object was detected at that place for at least  $n$  frames in a row. More frequent updates improve responsiveness, but run a greater risk of showing arms of users or objects that have just stopped moving when this is not intentional.

If in a future version we want to implement projection mapping, true shape recognition needs to be implemented. This can be used for providing extra information and detail, by projecting building names and facades for example. Figure 4 shows the limits of the system without projection mapping.

#### D. Multi-Touch interaction

For the touch interaction, there are still two possibilities. The multi-touch interaction can be based on the method presented by R. Xiao et al. [11]. It is able to detect an arbitrary number of hands and should be able to recognize most deliberate touches performed. There are still problems regarding optimization and accuracy. The other possibility is by using the skeleton detection software from the Kinect sensor. The advantages of using the Kinect are the efficiency of their code and the lack of false positives in arm detection. Initial testing showed that both work fairly robust, but further tuning and improvements are needed to also process the uncommon and unintended inputs mentioned above.

#### IV. CONCLUSION

We presented a system for accessible interaction with a crowd simulation framework by using an augmented-reality setup suited for use in a museum. The goals for this installation

are to involve the museum visitors with the scientific method and to teach them about the complexities of crowd simulation and management.

Our proof of concept [8] showed that such an immersive augmented-reality setting is well suited for this. In this demonstration paper and the accompanying video, we describe our development of an optimized and robust version of the original prototype. Detailed user testing with this revised version is currently put on hold due to the COVID-19 imposed restrictions, but will be taken up again in the future.

#### V. ACKNOWLEDGMENT

We would like to thank the group of Game and Media Technology students from Utrecht University who implemented the prototype for this project: Mark de Jong, Patrick Groot Koerkamp, Anestis Latos, Niels Mulder, Tiziano Natali, Anca-Mihaela Negulescu, Dimitris Valeris, Floris de Vries, Sietse van der Werf and Clifton Zhuo. In addition, we thank Nick Brouwer for providing his endless support. Finally, the project was supported by Jet Blokhuis and Aniek Bax from Universiteitsmuseum Utrecht.

#### REFERENCES

- [1] ucrowds b.v. <https://www.ucrowds.com>. Accessed: 2020-10-01.
- [2] Laura Tateosian, Helena Mitasova, Brendan Harmon, Brent Fogleman, Katherine Weaver, and Russel Harmon. Tangeoms: Tangible geospatial modeling system. *IEEE transactions on visualization and computer graphics*, 16(6):1605–1612, 2010.
- [3] Bertrand Schneider, Patrick Jermann, Guillaume Zufferey, and Pierre Dillenbourg. Benefits of a tangible interface for collaborative learning and interaction. *IEEE Transactions on Learning Technologies*, 4(3):222–232, 2010.
- [4] Charles R Amburn, Nathan L Vey, Michael W Boyce, and Jerry R Mize. The augmented reality sandtable (ares). Technical report, Army research lab Aberdeen probing ground MD human research and engineering ..., 2015.
- [5] Joyce Ma, Lisa Sindorf, Isaac Liao, and Jennifer Frazier. Using a tangible versus a multi-touch graphical user interface to support data exploration at a museum exhibit. In *Proceedings of the Ninth International Conference on Tangible, Embedded, and Embodied Interaction*, pages 33–40, 2015.
- [6] Nuno Correia, Tarquínio Mota, Rui Nóbrega, Luís Silva, and Andreia Almeida. A multi-touch tabletop for robust multimedia interaction in museums. In *ACM International Conference on Interactive Tabletops and Surfaces*, pages 117–120, 2010.
- [7] Tom Geller. Interactive tabletop exhibits in museums and galleries. *IEEE Computer Graphics and Applications*, 26(5):6–11, 2006.
- [8] Wolfgang Hürst, Roland Geraerts, and Yiran Zhao. Crowdar table-an ar table for interactive crowd simulation. In *2019 IEEE International Conference on Artificial Intelligence and Virtual Reality (AIVR)*, pages 302–3023. IEEE, 2019.
- [9] Augmented reality sandbox. <https://arsandbox.ucdavis.edu/>. Accessed: 2020-10-01.
- [10] Azure kinect. <https://azure.microsoft.com/en-us/services/kinect-dk/>. Accessed: 2020-10-01.
- [11] Robert Xiao, Scott Hudson, and Chris Harrison. Direct: Making touch tracking on ordinary surfaces practical with hybrid depth-infrared sensing. In *Proceedings of the 2016 ACM International Conference on Interactive Surfaces and Spaces*, pages 85–94, 2016.