Automatic Parameter Tuning via Reinforcement Learning for Crowd Simulation with Social Distancing

Universiteit Utrecht

Yiran Zhao and Roland Geraerts

Why simulation?

Societal relevance of simulation

- There are many places with big crowds
- Questions
  - In how much time can a train station be evacuated?
  - Where and how can potentially dangerous situations appear?
  - How and where can a city accommodate 0.5M people during an event?
  - How can we populate a virtual world with a believable crowd?

Love Parade 2010: *21 deaths, 510 injuries*

There are many things we could measure

- Questions
  - What should we measure?
  - Crowd density: how many people fit in one squared meter?
    - What is a safe value?

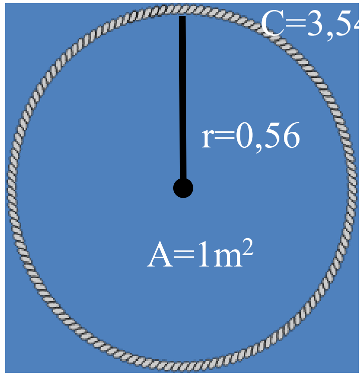

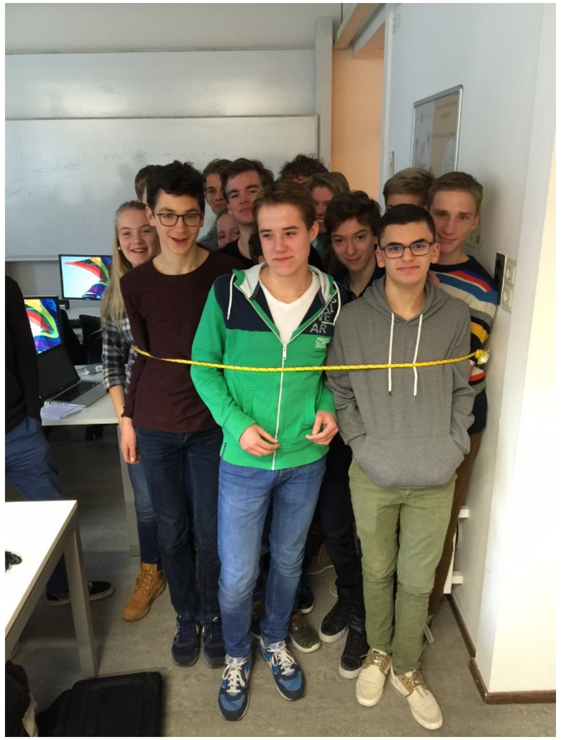

Example projects

Examples of simulation projects

Optimizing pedestrian flows

- Tour de France

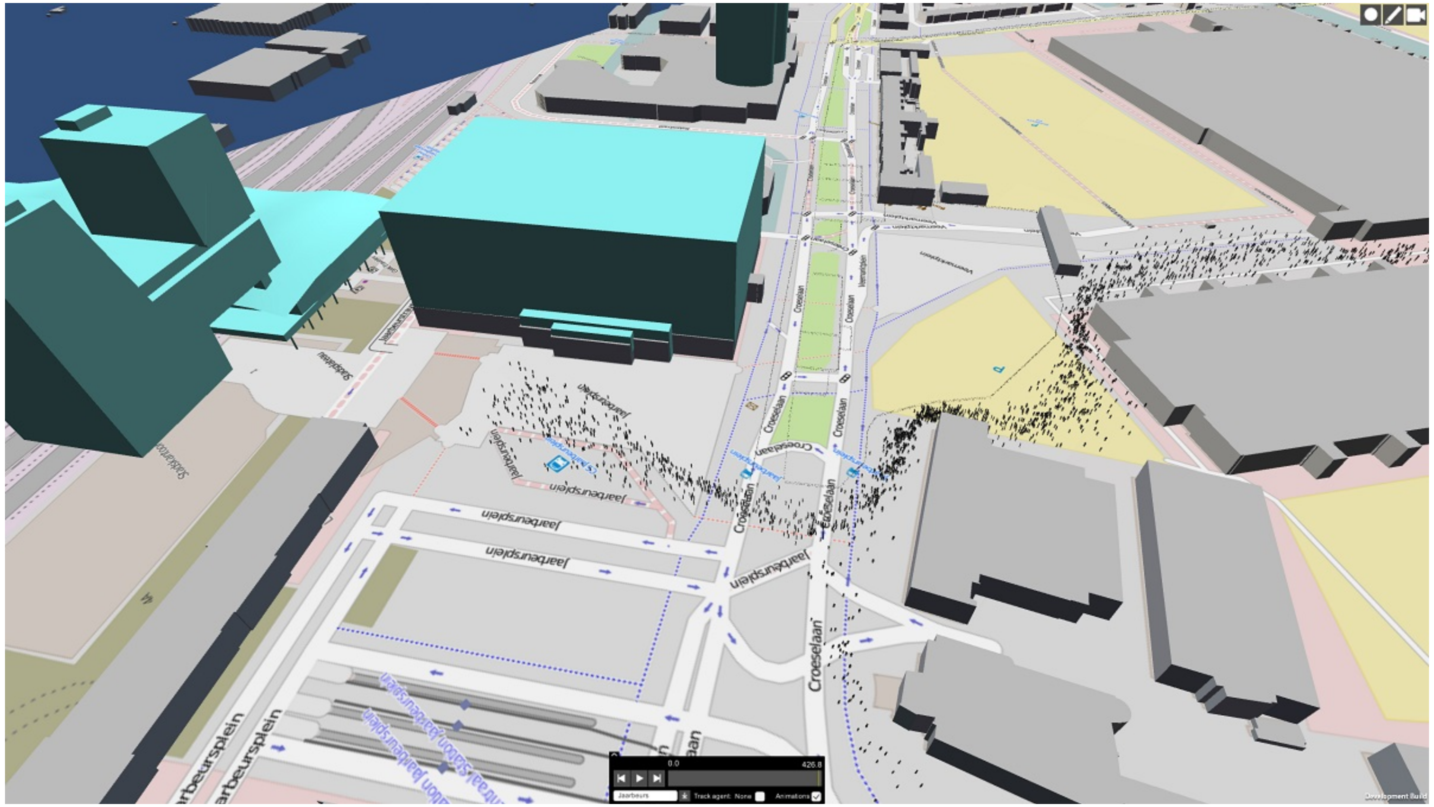

City of Utrecht

Examples of simulation projects

Evacuation of pedestrians (some hold a bicycle)

- Metro stations North/South line

City of Amsterdam

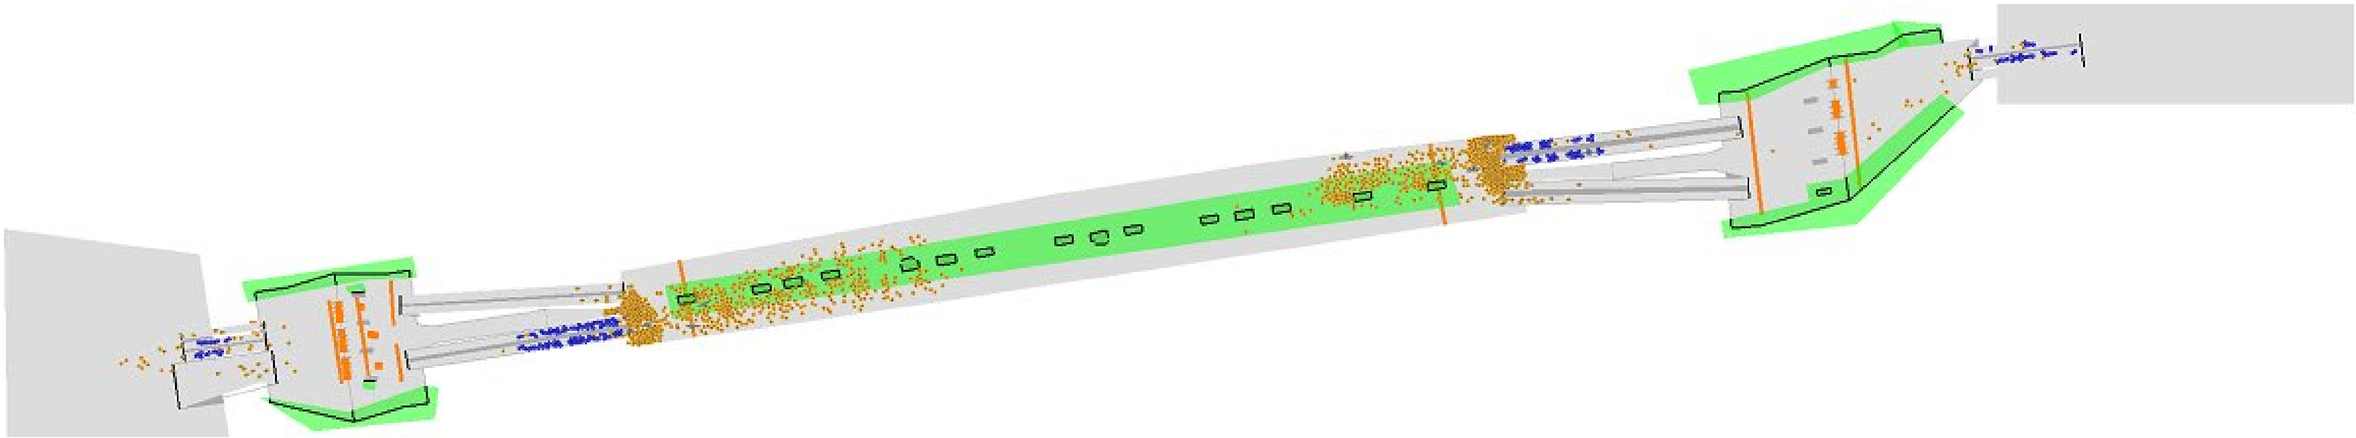

Examples of simulation projects

Carrying out what-if scenarios

- Train station Schiphol

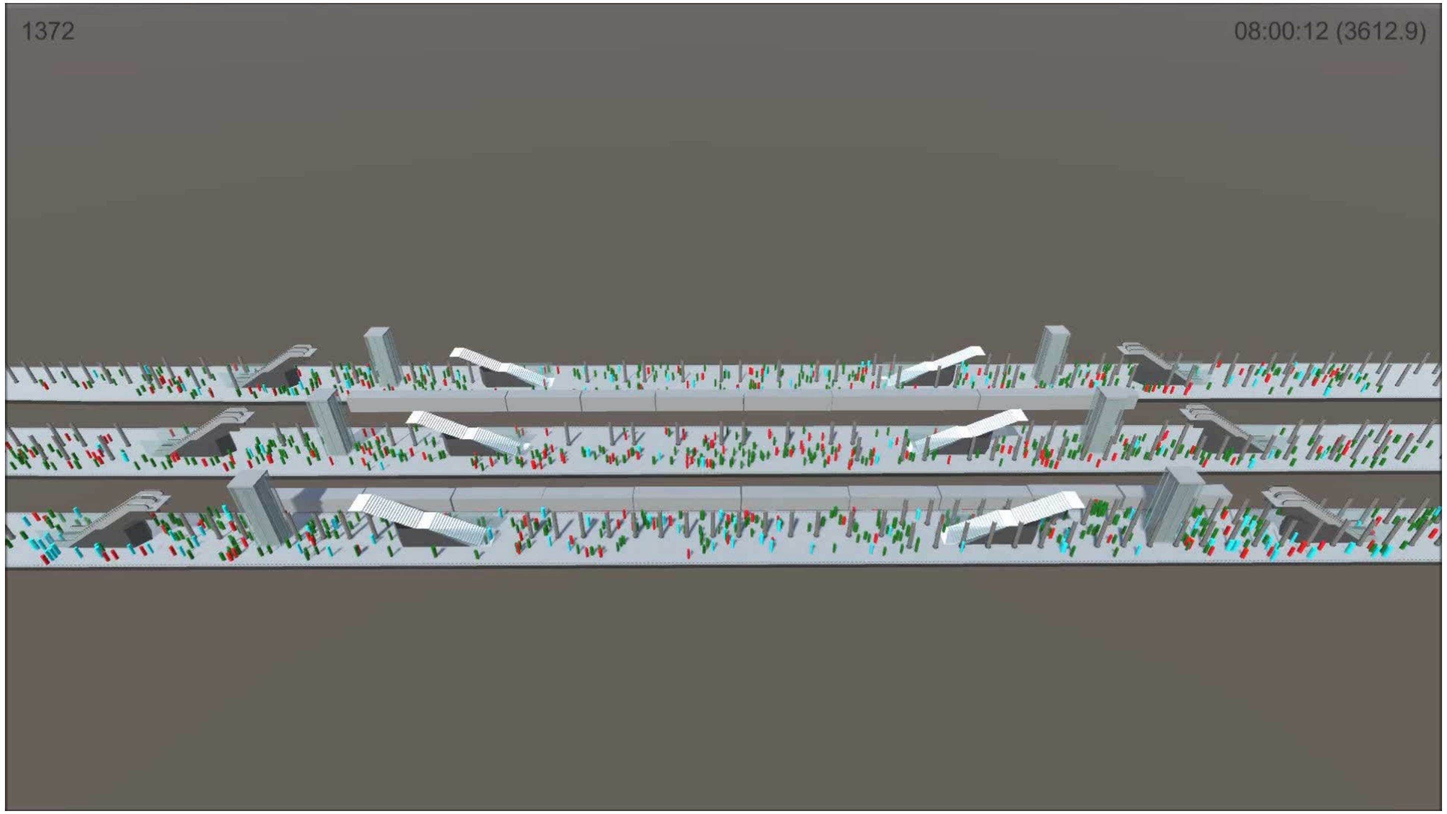

Examples of simulation projects

Studying crowd safety measures

- Vuelta 2022

City of Utrecht

Movares
consultants & engineers

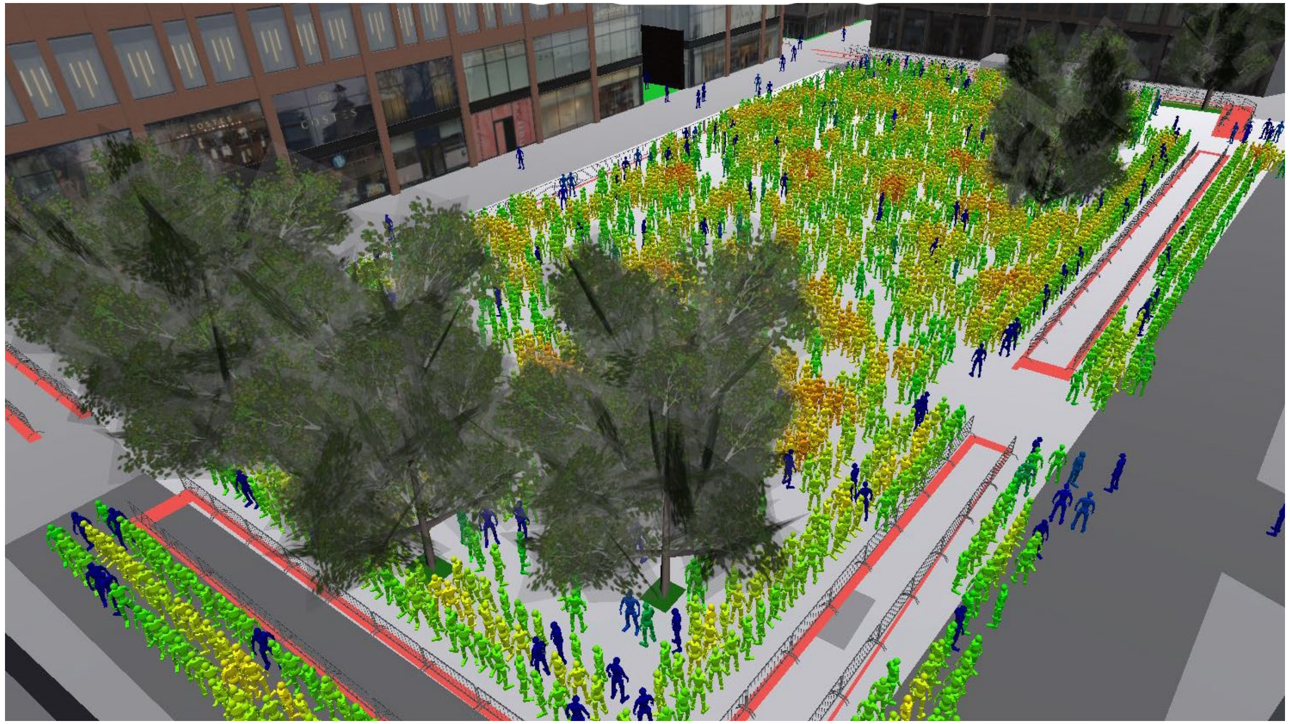

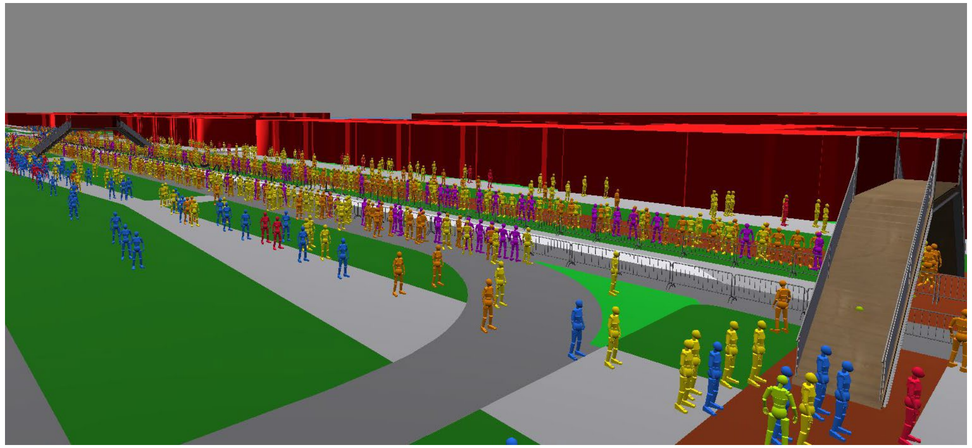

Examples of simulation projects

Studying effects of social distancing

- Covid-19: How much extra voting time (Dutch parliamentary elections)?

Ministerie van Binnenlandse Zaken en Koninkrijksrelaties

NOBLEO technology

How

can you simulate a human crowd?

Where to start?

- Festival?

It started with crowd simulation models for animals

- Reynold’s model
  - Navigate to the center, avoid collisions, blend with local speed / direction

Crowd simulation framework

- Representation environment
- Level 5
  - Plans actions
- Level 4
  - Creates indicative routes
- Level 3
  - Traverses the routes
  - Yields speed/direction pairs
- Level 2
  - Adapts routes
  - E.g., to avoid collisions
- Level 1
  - Moves the agents

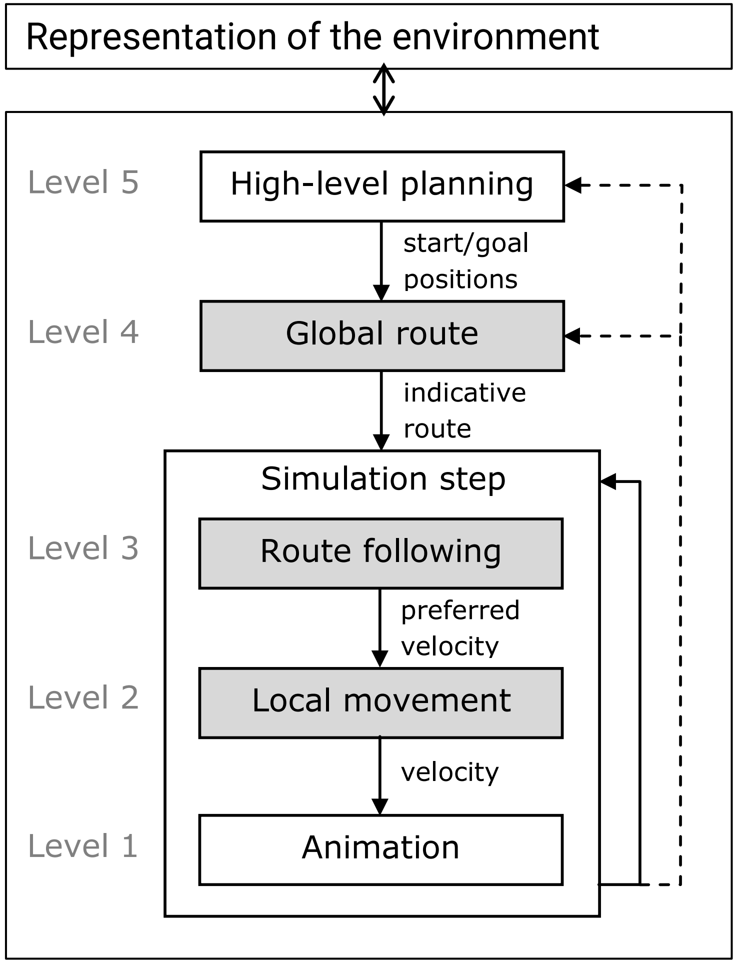

Van Toll, Jaklin, and Geraerts, 2015. Towards Believable Crowds: A Generic Multi-Level Framework for Agent Navigation.

Crowd simulation framework

- Representation environment
- Level 5
  - Plans actions
- Level 4
  - Creates indicative routes
- Level 3
  - Traverses the routes
  - Yields speed/direction pairs
- Level 2
  - Adapts routes
  - E.g., to avoid collisions
- Level 1
  - Moves the agents

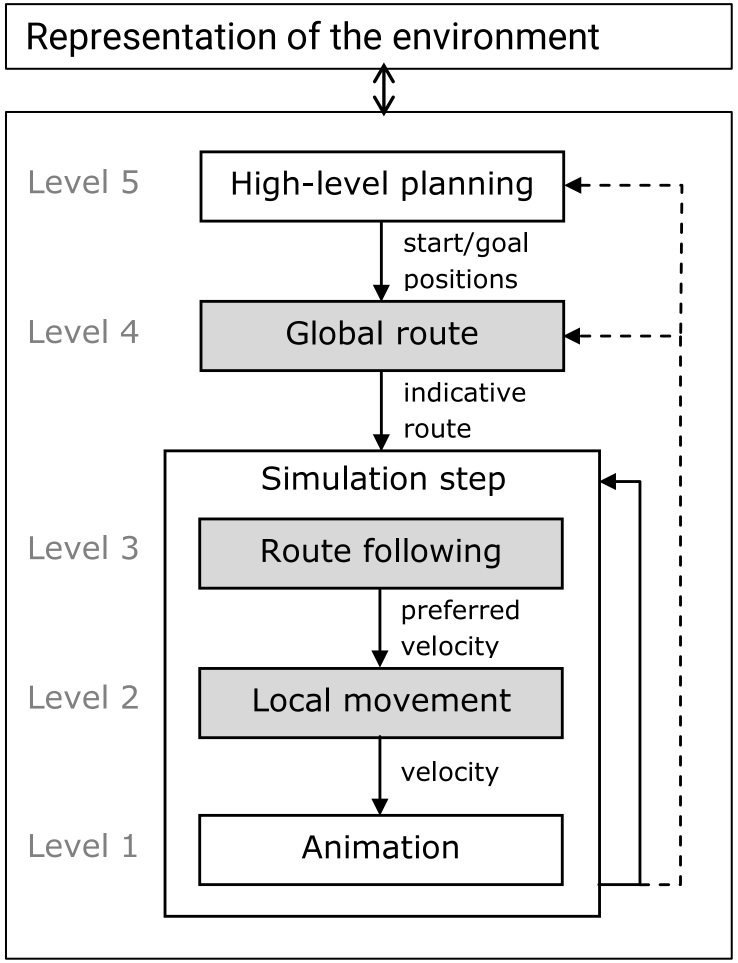

Van Toll, Jaklin, and Geraerts, 2015. Towards Believable Crowds: A Generic Multi-Level Framework for Agent Navigation.

Representing the walkable environment

Representation of the walkable environment

- Inspired by tiles…

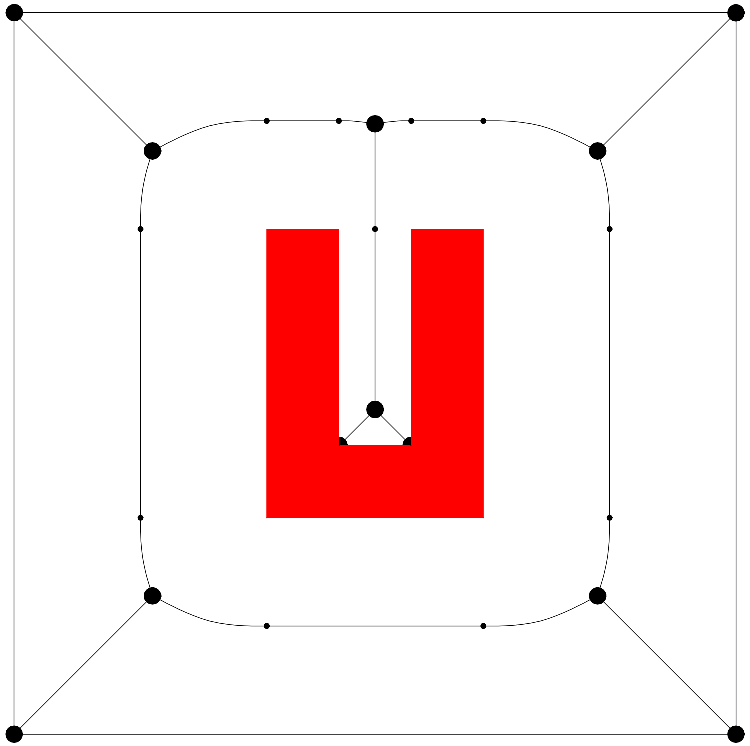

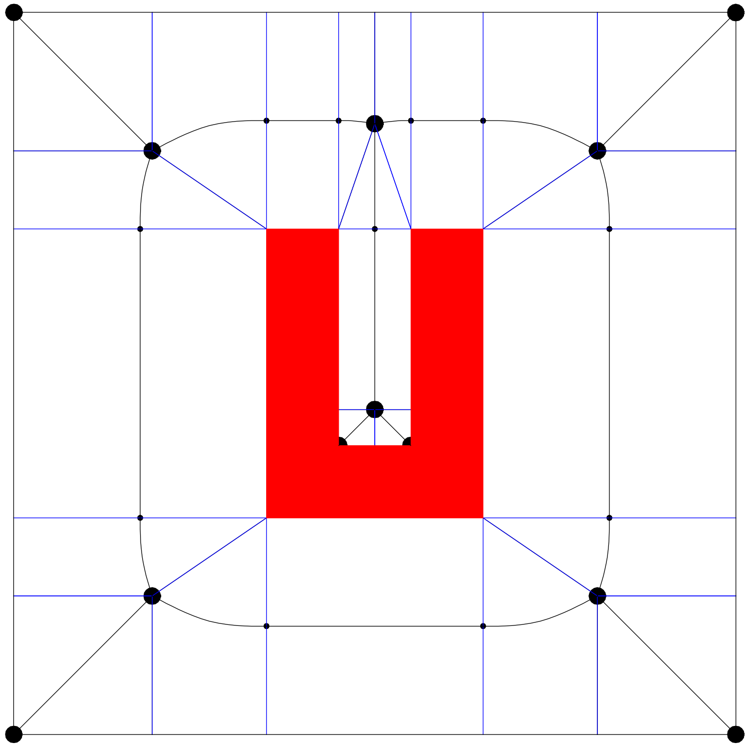

Representation of the walkable environment

What is the best representation for an environment's walkable space?

- Parallel computation of navigation mesh applied to 243 km^2

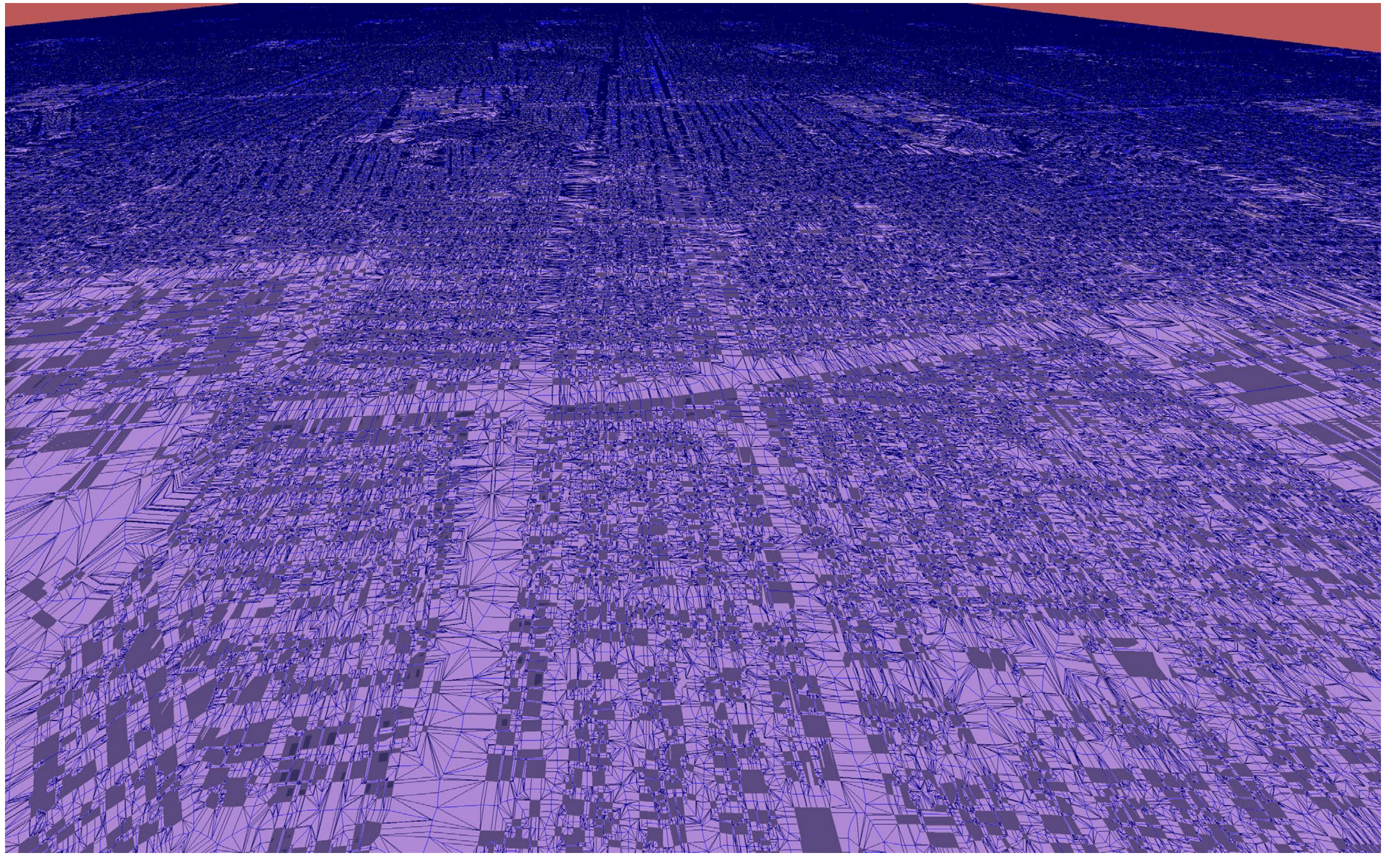

From 3D geometry to a navigation mesh

- Computation of walkable areas in an environment, and navigation map

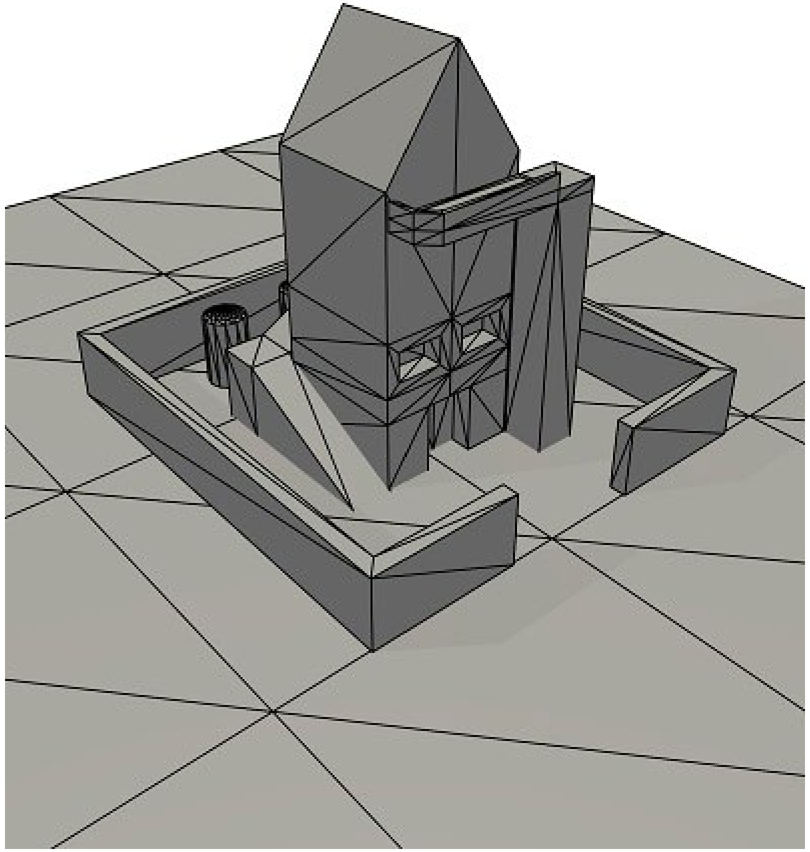

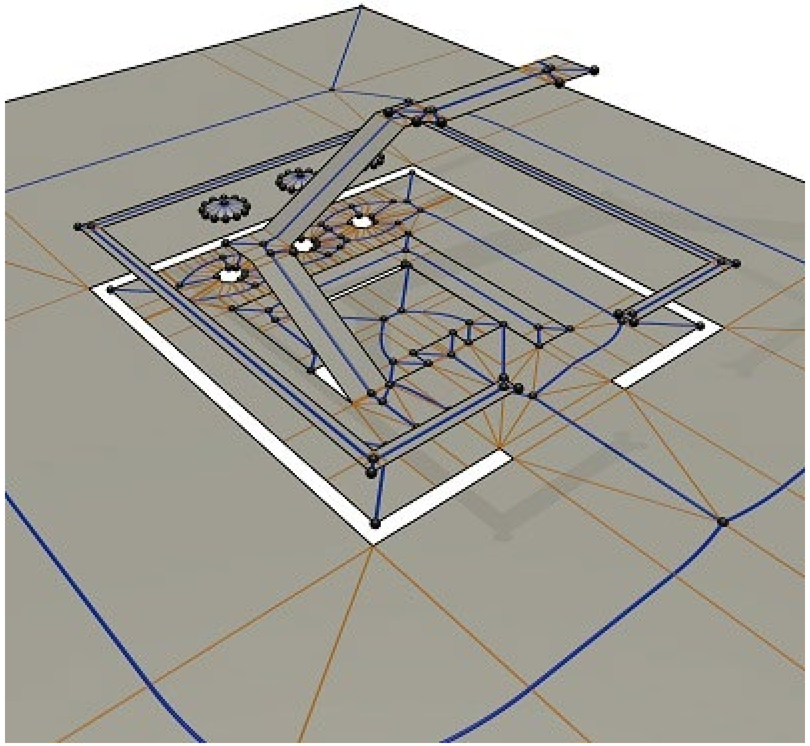

Van Toll et al, 2018: The Medial Axis of a Multi-Layered Environment and its Application as a Navigation Mesh

From other data resources to navigation mesh

- Other data sources
  - Open data from the Dutch Cadastre (BGT, AHN, BAG3D)

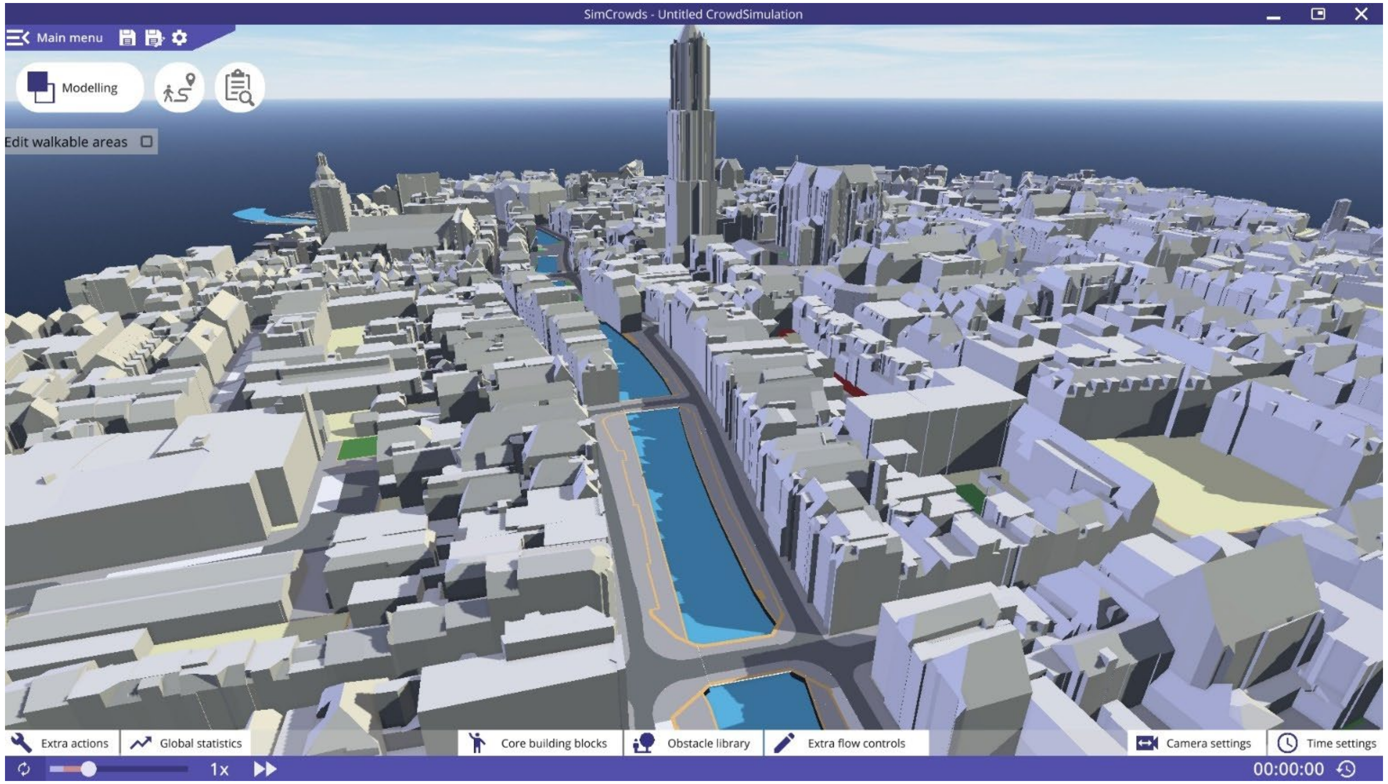

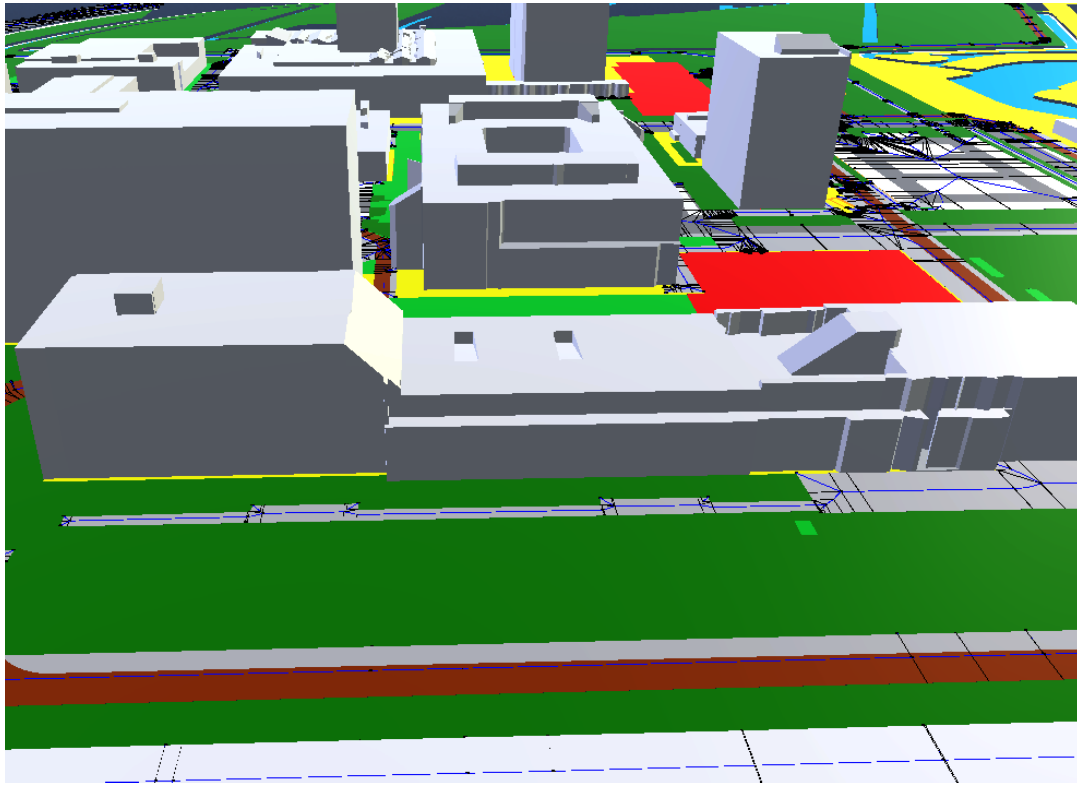

From navigation mesh to simulation of 1 pedestrian

Crowd simulation framework

- Representation environment
- Level 5
  - Plans actions
- Level 4
  - Creates indicative routes
- Level 3
  - Traverses the routes
  - Yields speed/direction pairs
- Level 2
  - Adapts routes
  - E.g., to avoid collisions
- Level 1
  - Moves the agents

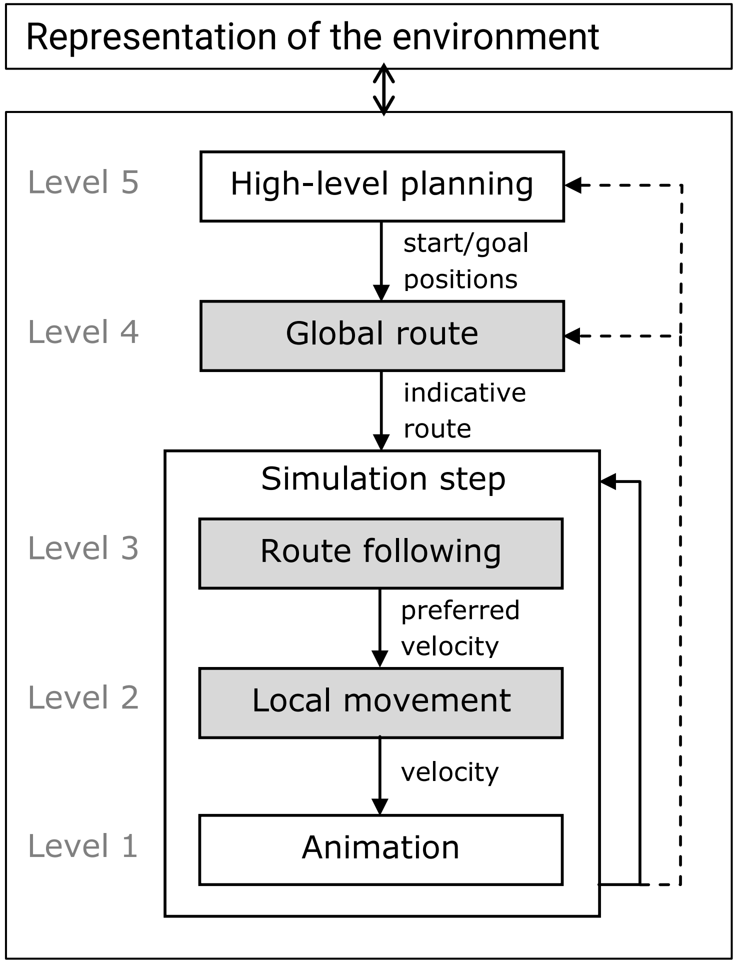

Van Toll, Jaklin, and Geraerts, 2015. Towards Believable Crowds: A Generic Multi-Level Framework for Agent Navigation.

Crowd simulation framework

- Representation environment
- Level 5
  - Plans actions
- Level 4
  - Creates indicative routes
- Level 3
  - Traverses the routes
  - Yields speed/direction pairs
- Level 2
  - Adapts routes
  - E.g., to avoid collisions
- Level 1
  - Moves the agents

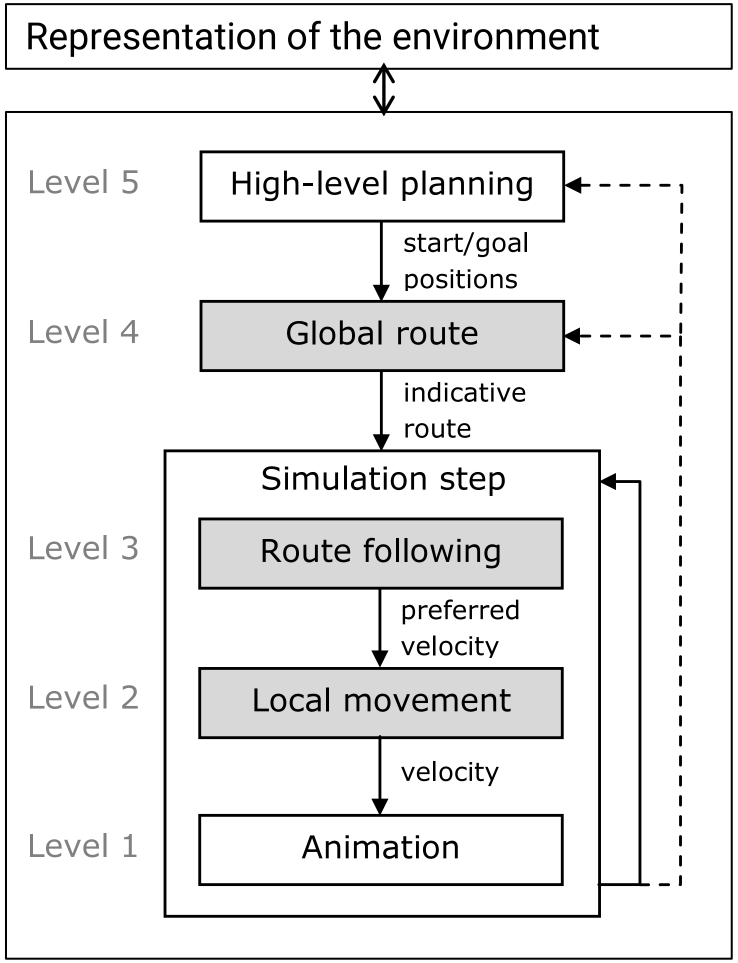

Van Toll, Jaklin, and Geraerts, 2015. Towards Believable Crowds: A Generic Multi-Level Framework for Agent Navigation.

Indicative Routes

- A path planning algorithm should NOT compute a path
  - A one-dimensional path limits the agent's freedom
  - Humans don't do that either
- It should produce
  - An Indicative/Preferred Route
    - Guides agent to goal

Indicative Routes

- Example
  - A shortest path with clearance to obstacles
  - Side preference

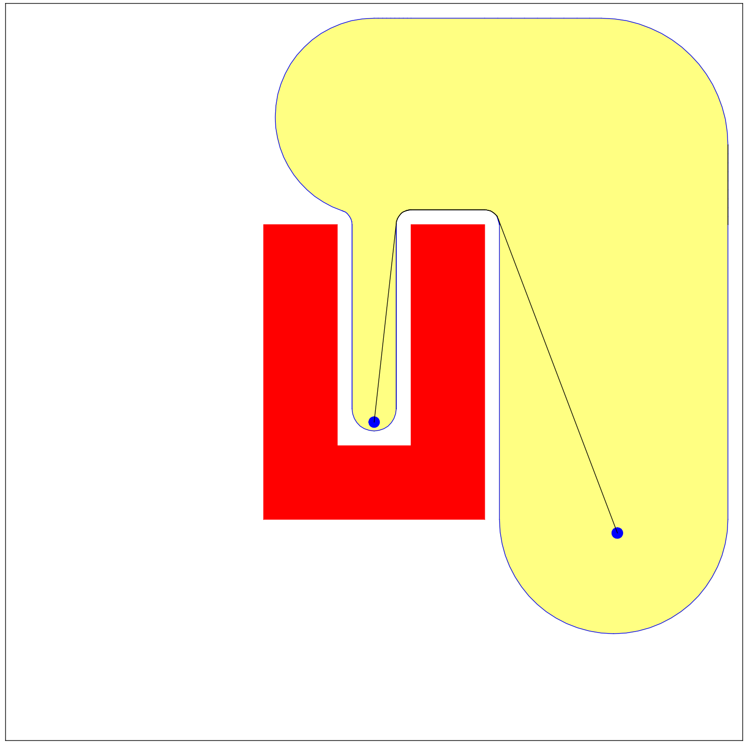

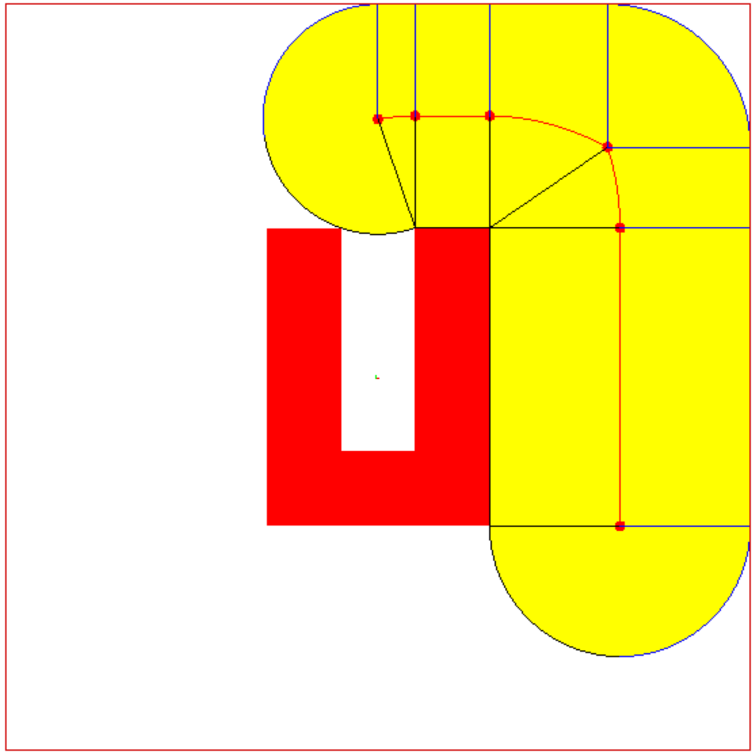

Jaklin et al, 2014: Computing High-Quality Paths in Weighted Regions

Crowd simulation framework

- Representation environment
- Level 5
  - Plans actions
- Level 4
  - Creates indicative routes
- Level 3
  - Traverses the routes
  - Yields speed/direction pairs
- Level 2
  - Adapts routes
  - E.g., to avoid collisions
- Level 1
  - Moves the agents

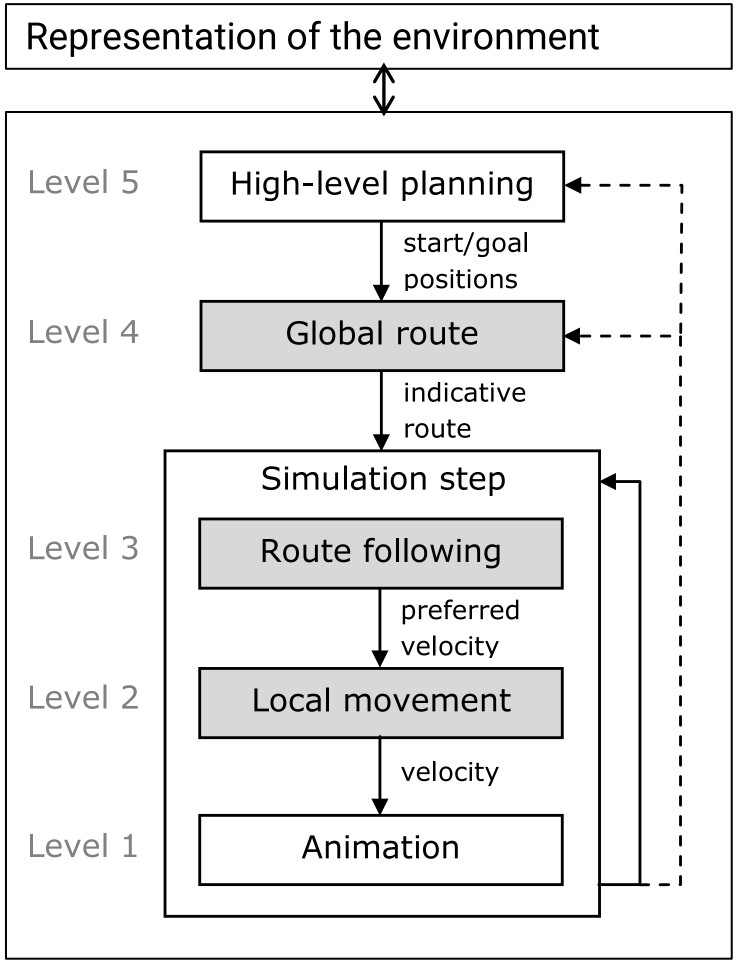

Van Toll, Jaklin, and Geraerts, 2015. Towards Believable Crowds: A Generic Multi-Level Framework for Agent Navigation.

Following routes

- Basic algorithm
  - An attraction point on the indicative route guides the pedestrian to its goal
  - Obstacles repulse pedestrians when they are too close

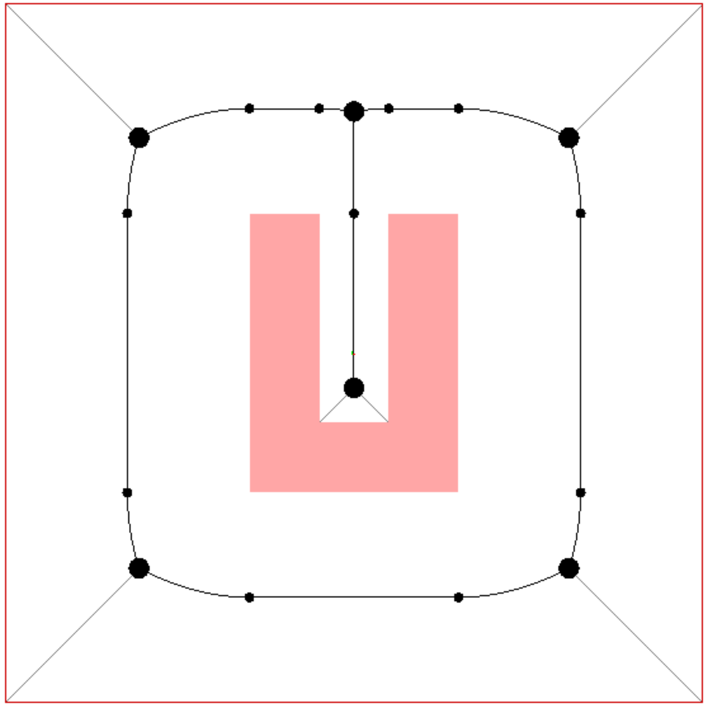

From simulation of 1 pedestrian to a crowd

Crowd simulation framework

- Representation environment
- Level 5
  - Plans actions
- Level 4
  - Creates indicative routes
- Level 3
  - Traverses the routes
  - Yields speed/direction pairs
- Level 2
  - Adapts routes
  - E.g., to avoid collisions
- Level 1
  - Moves the agents

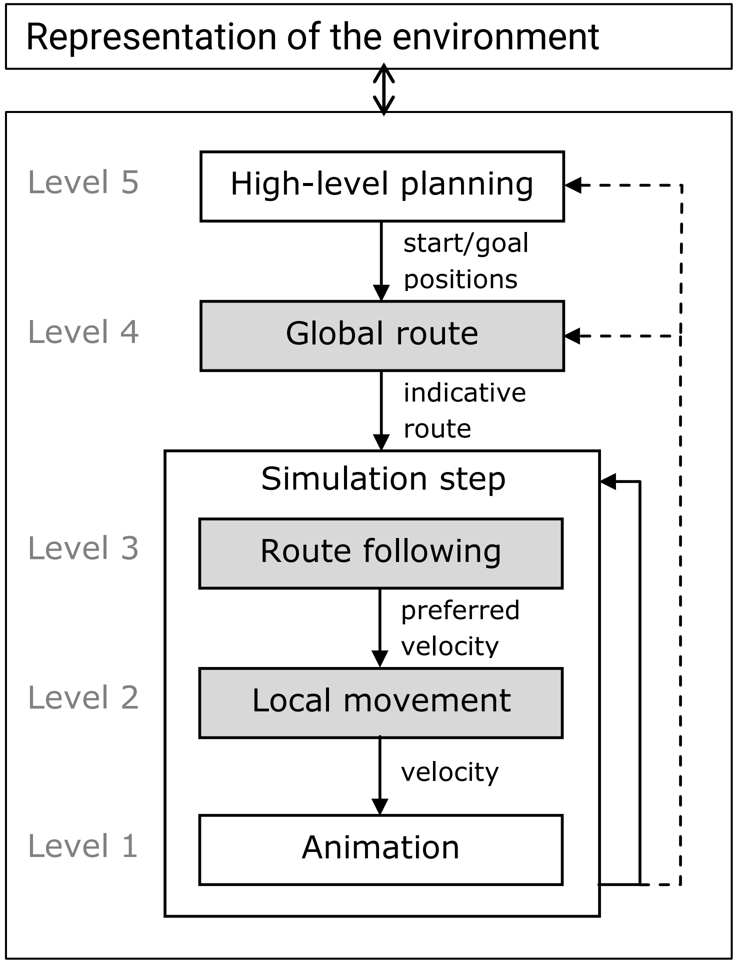

Van Toll, Jaklin, and Geraerts, 2015. Towards Believable Crowds: A Generic Multi-Level Framework for Agent Navigation.

Adapting the routes: Collision avoidance

- Our model is derived from experiments in the MOCAP lab

PhD students: Wouter van Toll and Norman Jaklin

Adapting the routes: Collision avoidance

- Our model slightly adjusts the people's movements
  - Improvement: Optimal Reciprocal Collision Avoidance (ORCA)

Karamouzas et al, 2009: A Predictive Collision Avoidance Model for Pedestrian Simulation; van den Berg et al 2010: ORCA.

Adapting the routes: Social groups

- The group members stay close and visible to each other

Kremyzas et al, 2016: Towards Social Behavior in Virtual-Agent Navigation

Adapting the routes: Social distancing

- Keeping social distances between pedestrians (based on profile)
  - Small children tend to obey smaller distances than adults
  - 1.5 meter is perceived differently by humans
- Members in a social groups do not socially distance to each other

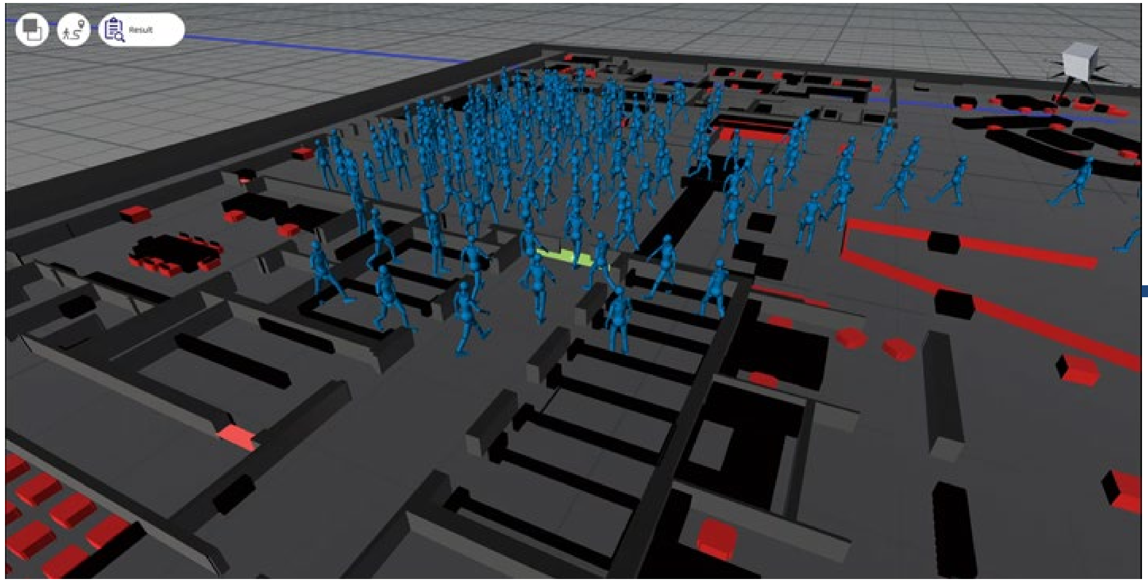

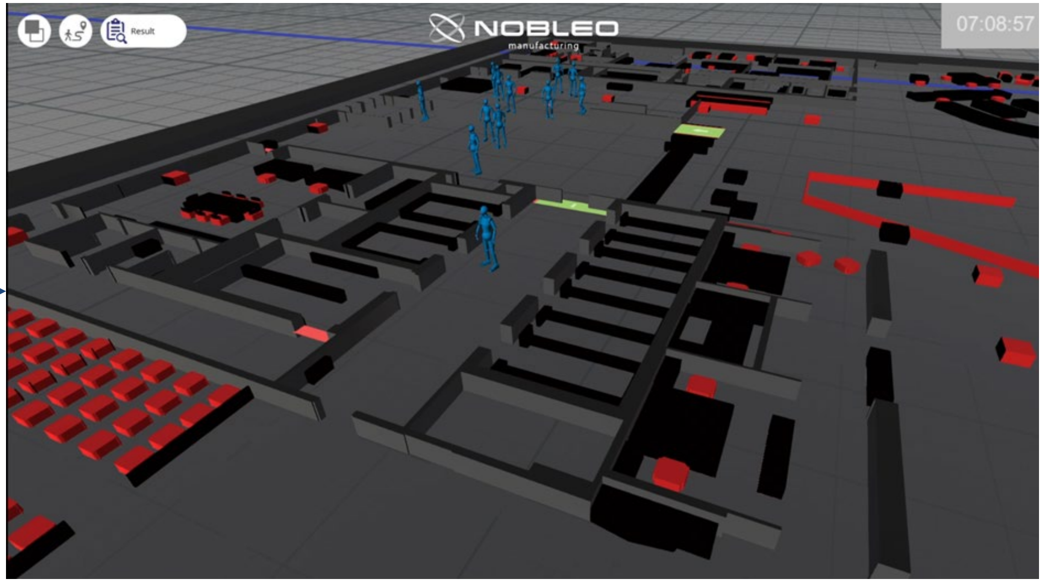

Crowd simulation framework

- Representation environment
- Level 5
  - Plans actions
- Level 4
  - Creates indicative routes
- Level 3
  - Traverses the routes
  - Yields speed/direction pairs
- Level 2
  - Adapts routes
  - E.g., to avoid collisions
- Level 1
  - Moves the agents

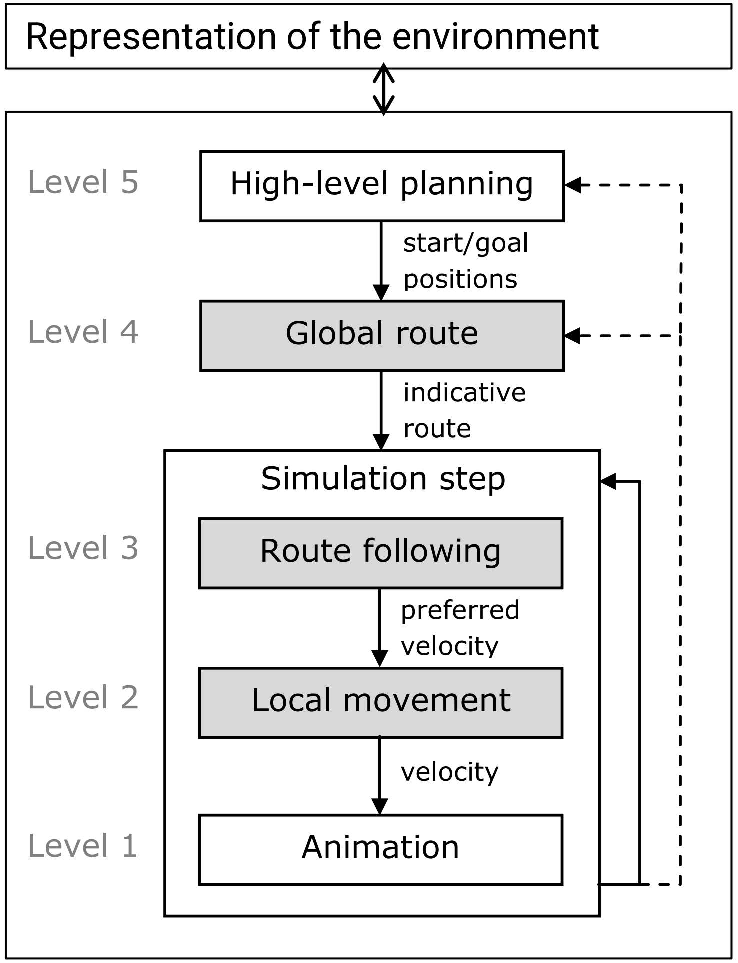

Van Toll, Jaklin, and Geraerts, 2015. Towards Believable Crowds: A Generic Multi-Level Framework for Agent Navigation.

Our paper's goal

- How can we tune the parameters' values dynamically for more realistic behaviors?
  - Density-based planning
  - Optimal Reciprocal Collision Avoidance (ORCA)
  - Path planning (shortest path versus side preference)
  - Social distance

Approach

- Deep reinforcement learning with Proximal Policy Optimization
  - State: position, goal, velocity, neighbors, etc.

1. Run the simulation with the parameters sampled from the policy.
2. Store the state, action and reward of each agent in each simulation step in the buffer, under the current parameter policy.
3. At the end of the training episode, calculate the advantage of the current policy, then optimize it and clear the buffer.
4. Repeat from 1.

Algorithm 1 Automatic parameter tuning

$t \leftarrow L_{run}$
while simulation not end **do**
    if $t \geq L_{episode}$ **then**
        for all q in Q **do**
            Calculate advantages
            Optimize policy π_q
            Clear replay buffer B_q
        end for
        Restart simulation
        $t \leftarrow 0$
    end if
    if $t \mod L_{run} = 0$ **then**
        for all q in Q **do**
            Run policy π_q
            Update replay buffer B_q
        end for
    end if
    Perform simulation
    $t \leftarrow t + 1$
end while

Results

- Counter-flow scenario

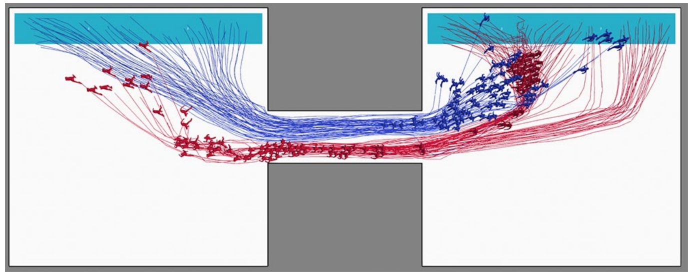

side preference 0.5 + ORCA
Average travelling time 58.85 s

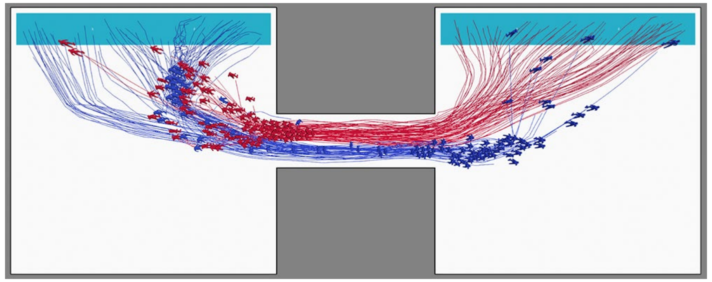

side preference 0.35 + ORCA
Average travelling time 55.99 s

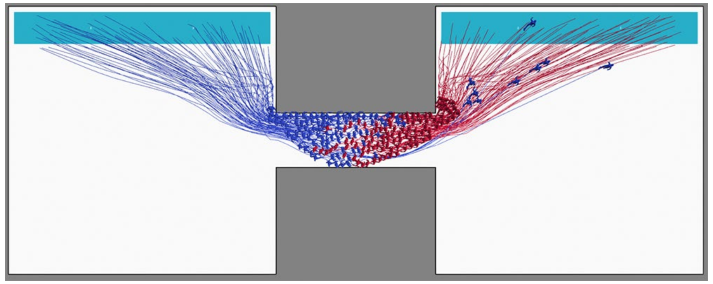

Shortest path + ORCA
Average travelling time 83.01 s

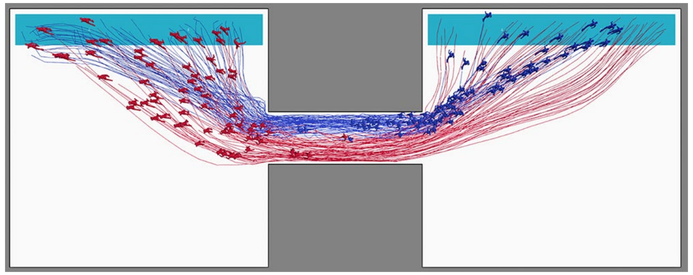

train side preference for 5000 episodes
+ ORCA
Average travelling time 45.05 s

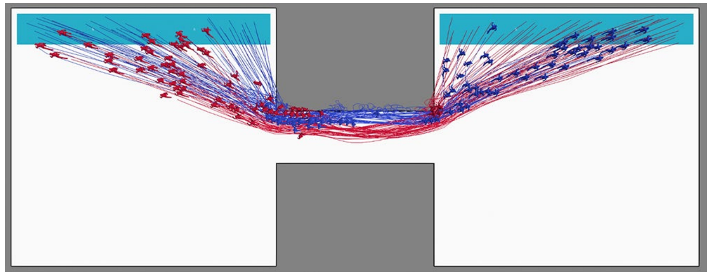

Shortest path + Implicit Crowds
Average travelling time 52.01 s

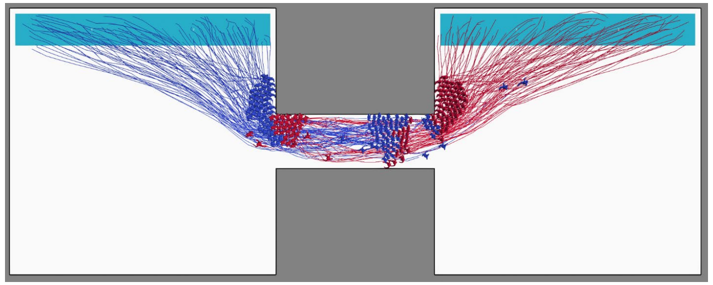

Shortest path
+ SGN [Jaklin et al. 2015]
Average travelling time 297.89 s

Results

- Social distancing scenario

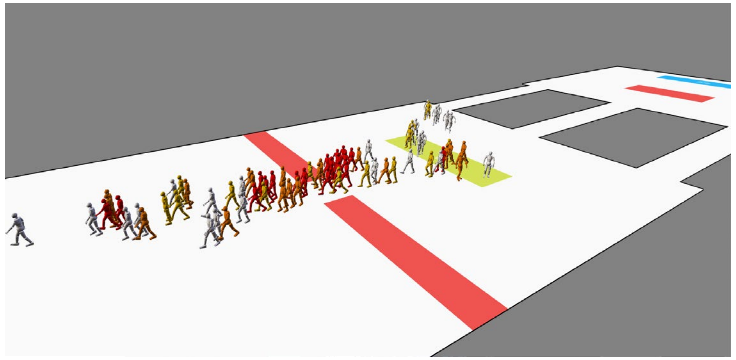

Density weight=0, social preference=0.01
Travelling time=84.95 s, collisions=143676

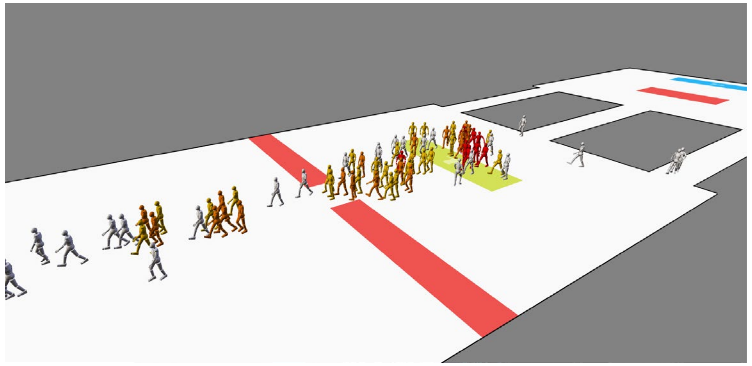

Density weight=30, social preference=0.3
Travelling time=87.25 s, collisions= 74596

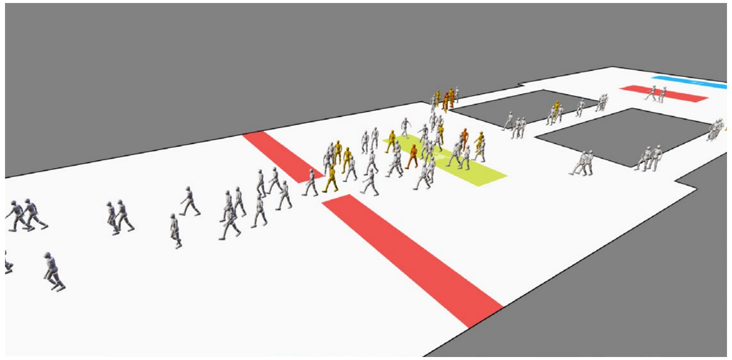

Density weight=100, social preference=0.8
Travelling time=118.25 s, collisions= 57904

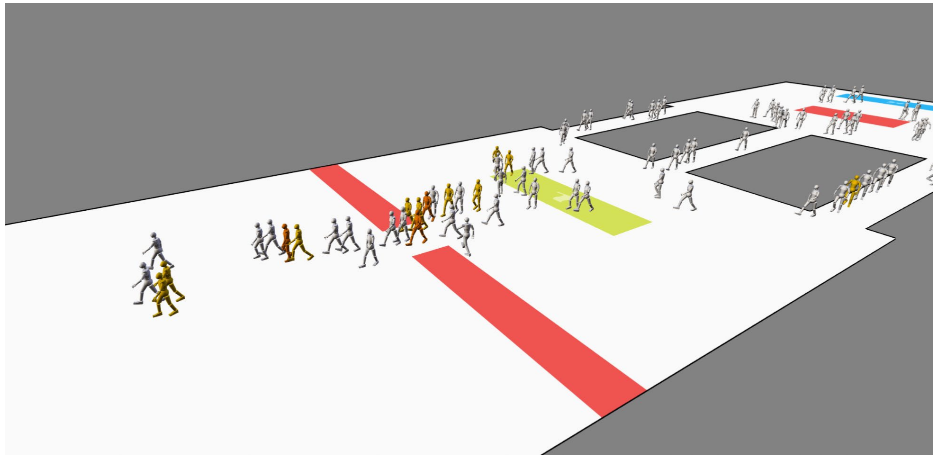

Training parameters: density weight and social distance preference.

Training episodes: 5000
social distance awareness: 1.0

Travelling time: 102.01
Collisions: 32368

Results

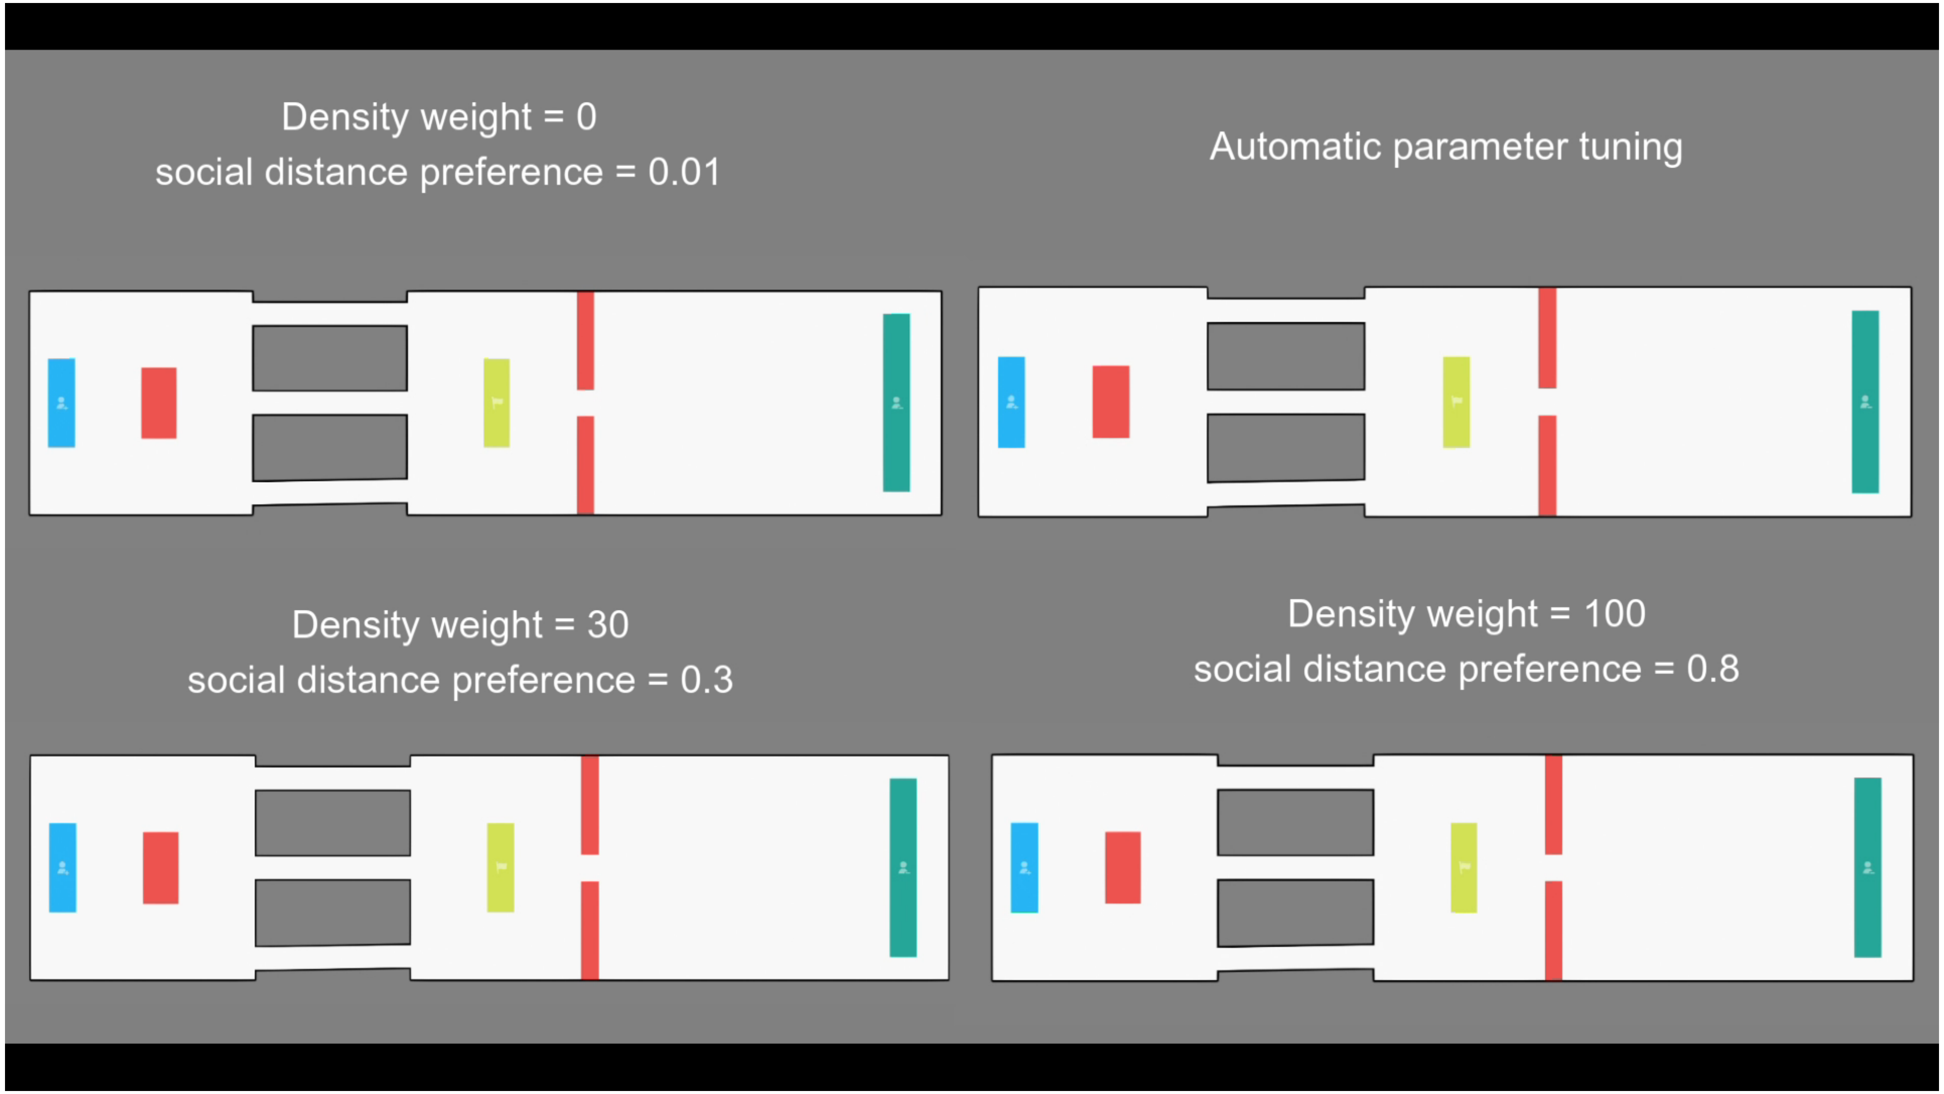
Density weight = 0
social distance preference = 0.01
Automatic parameter tuning
Density weight = 30
social distance preference = 0.3
Density weight = 100
social distance preference = 0.8

Current research and developments

Digital twins and the Metaverse

- Synthetic environment of a city with 1 mln pedestrians and traffic
  - Training for evacuations in emergency situations (London)

CAE

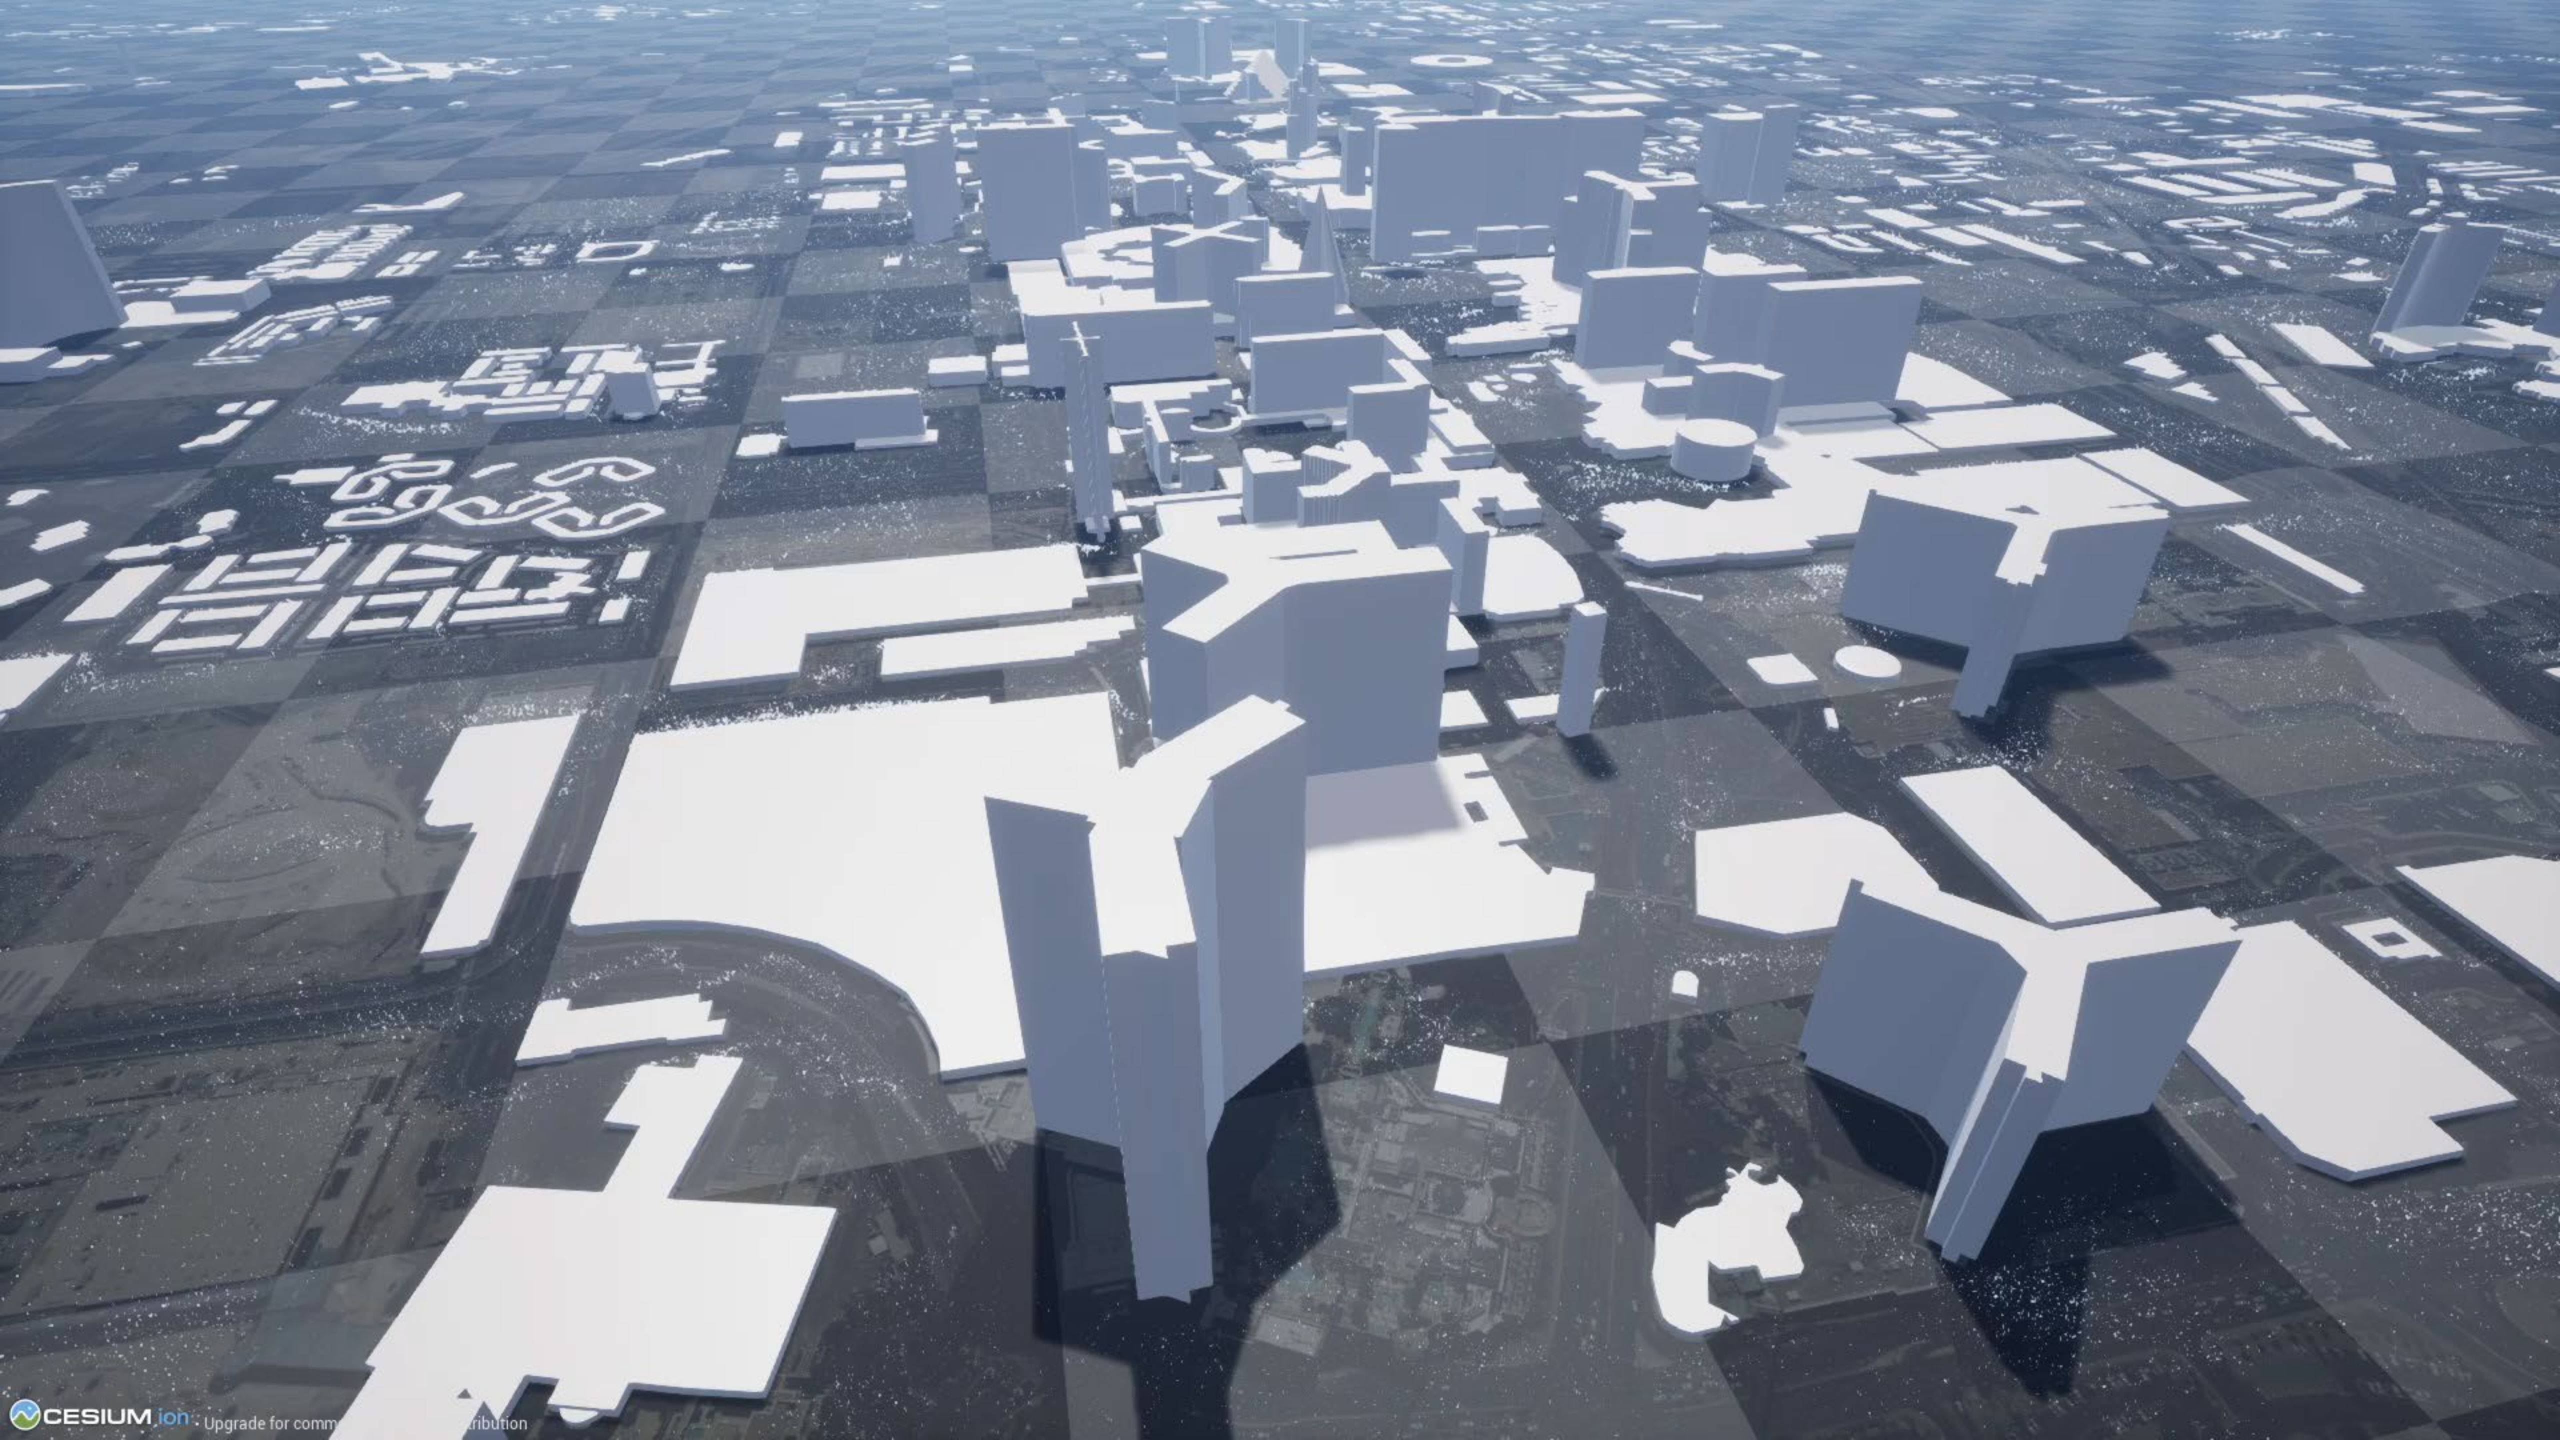
CESIUM ion

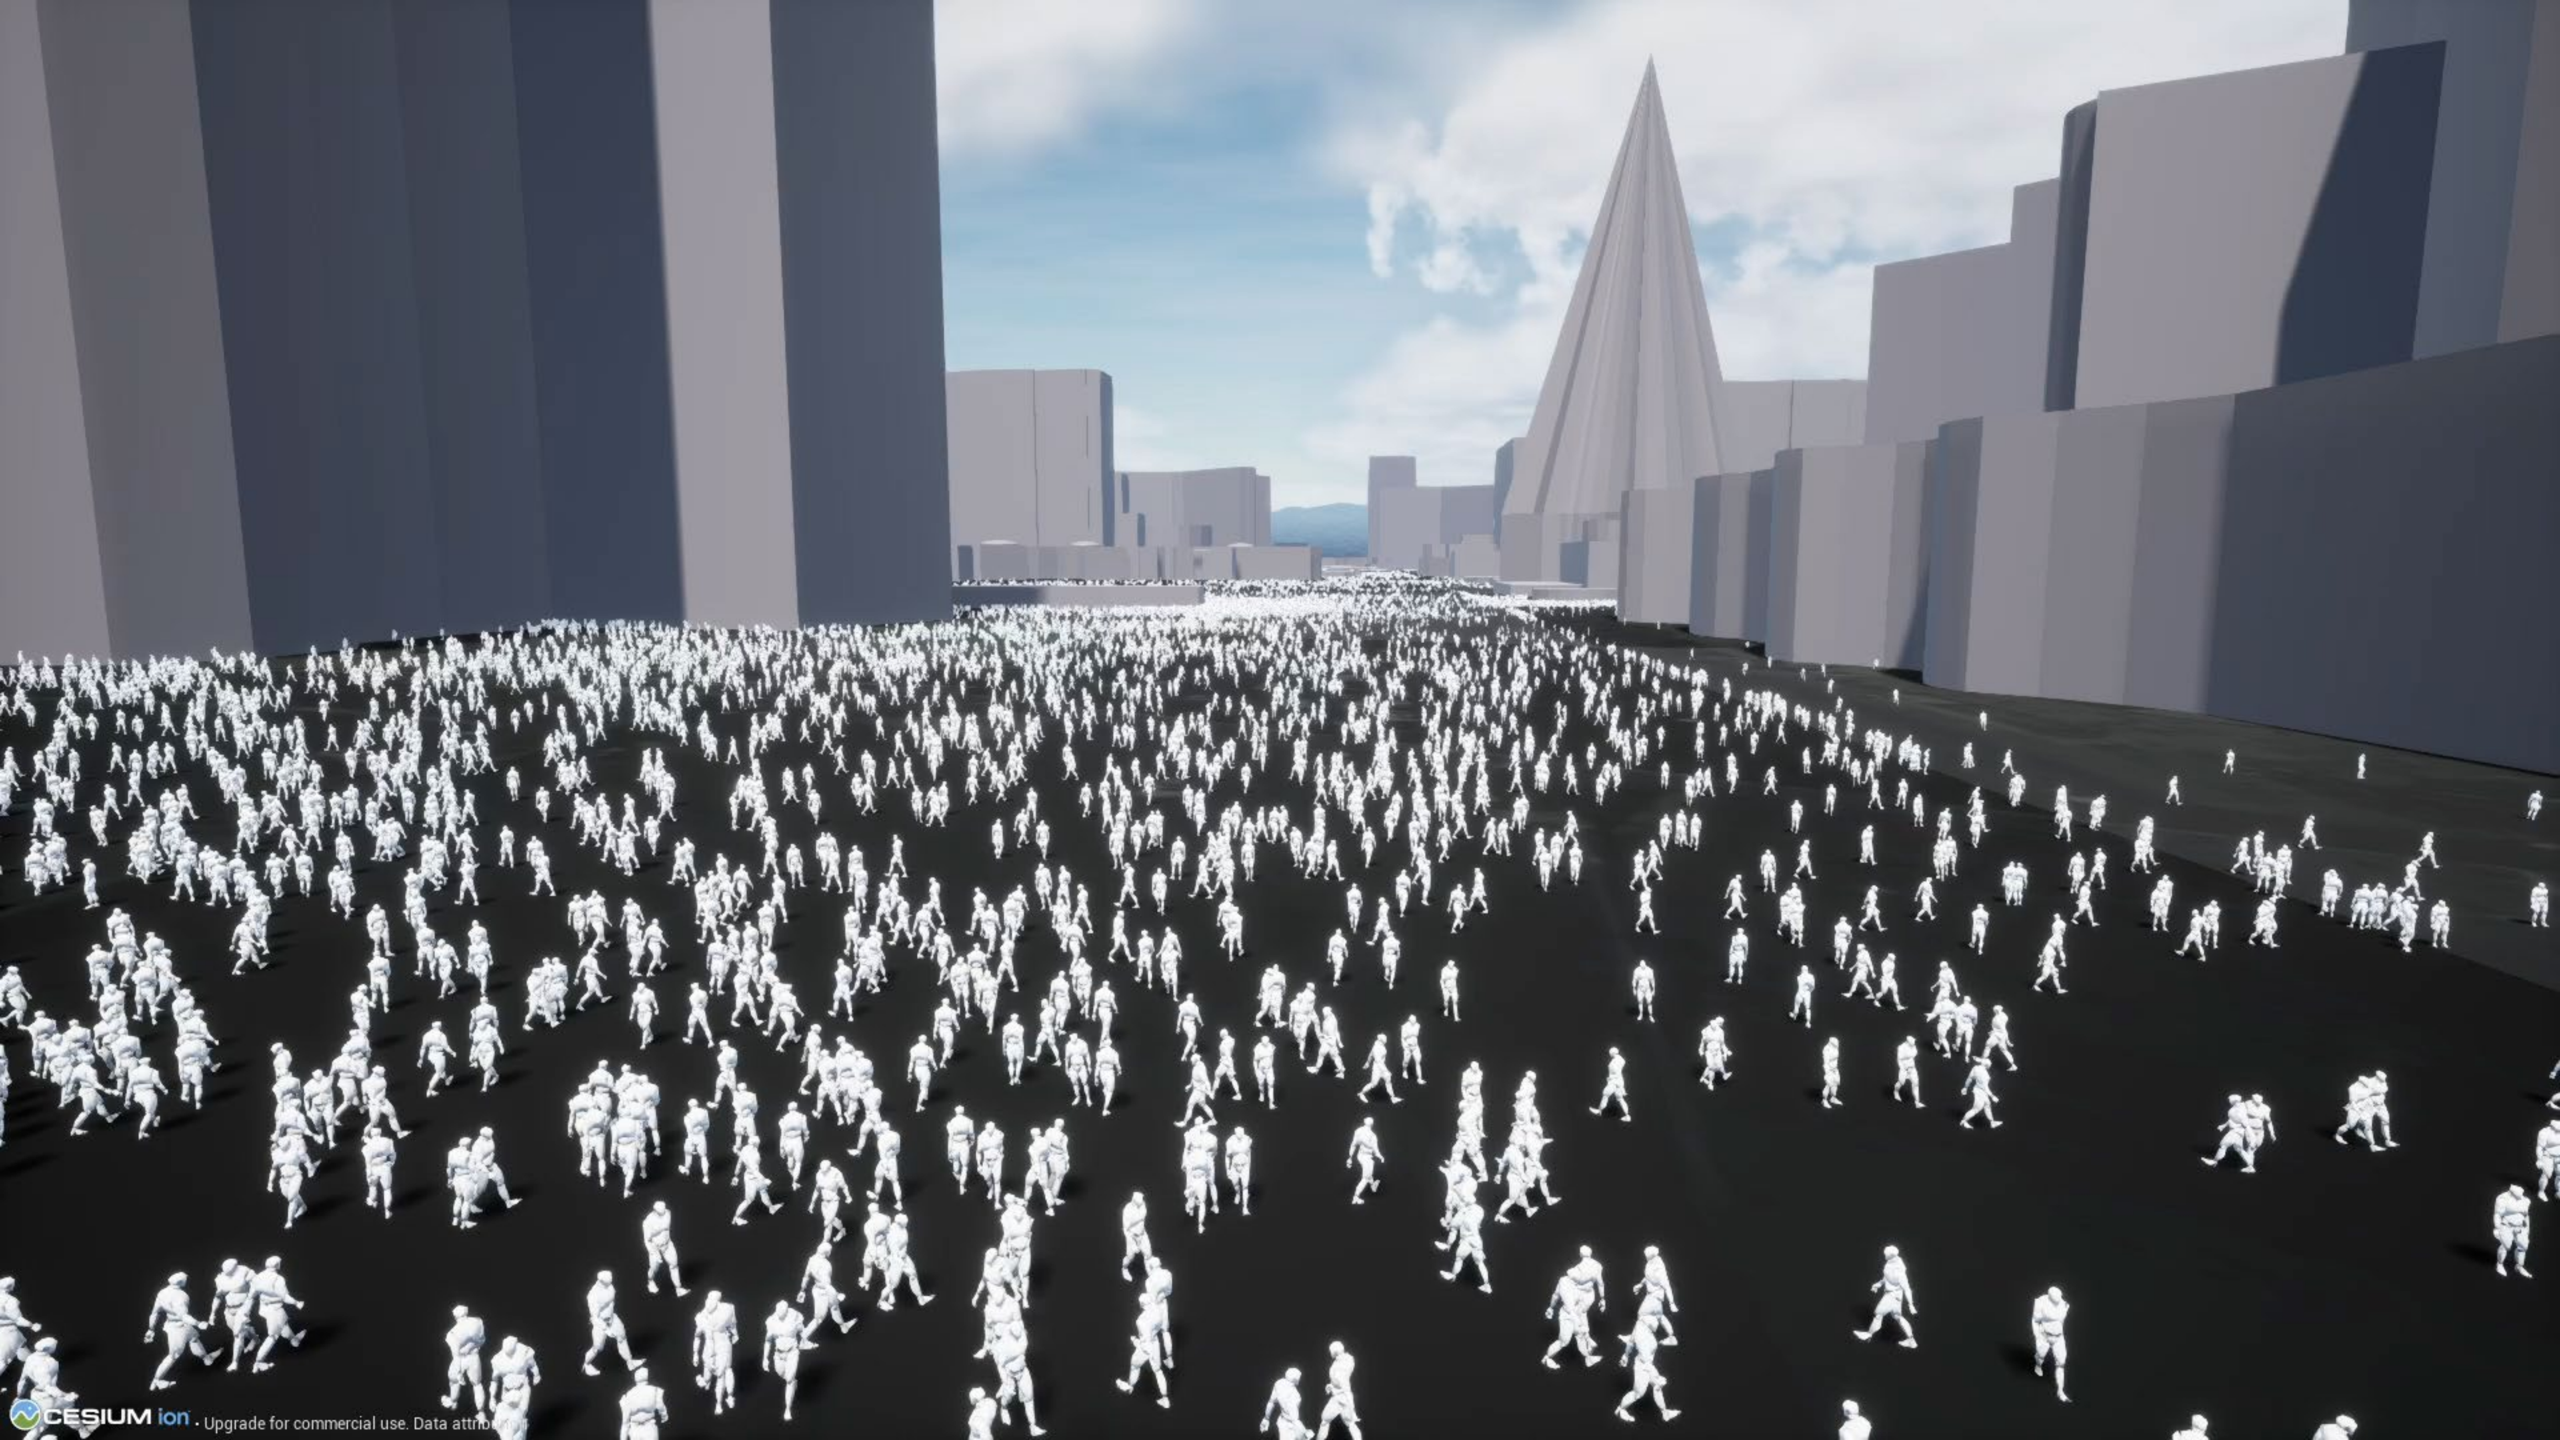
CESIUM ion · Upgrade for commercial use. Data attrib

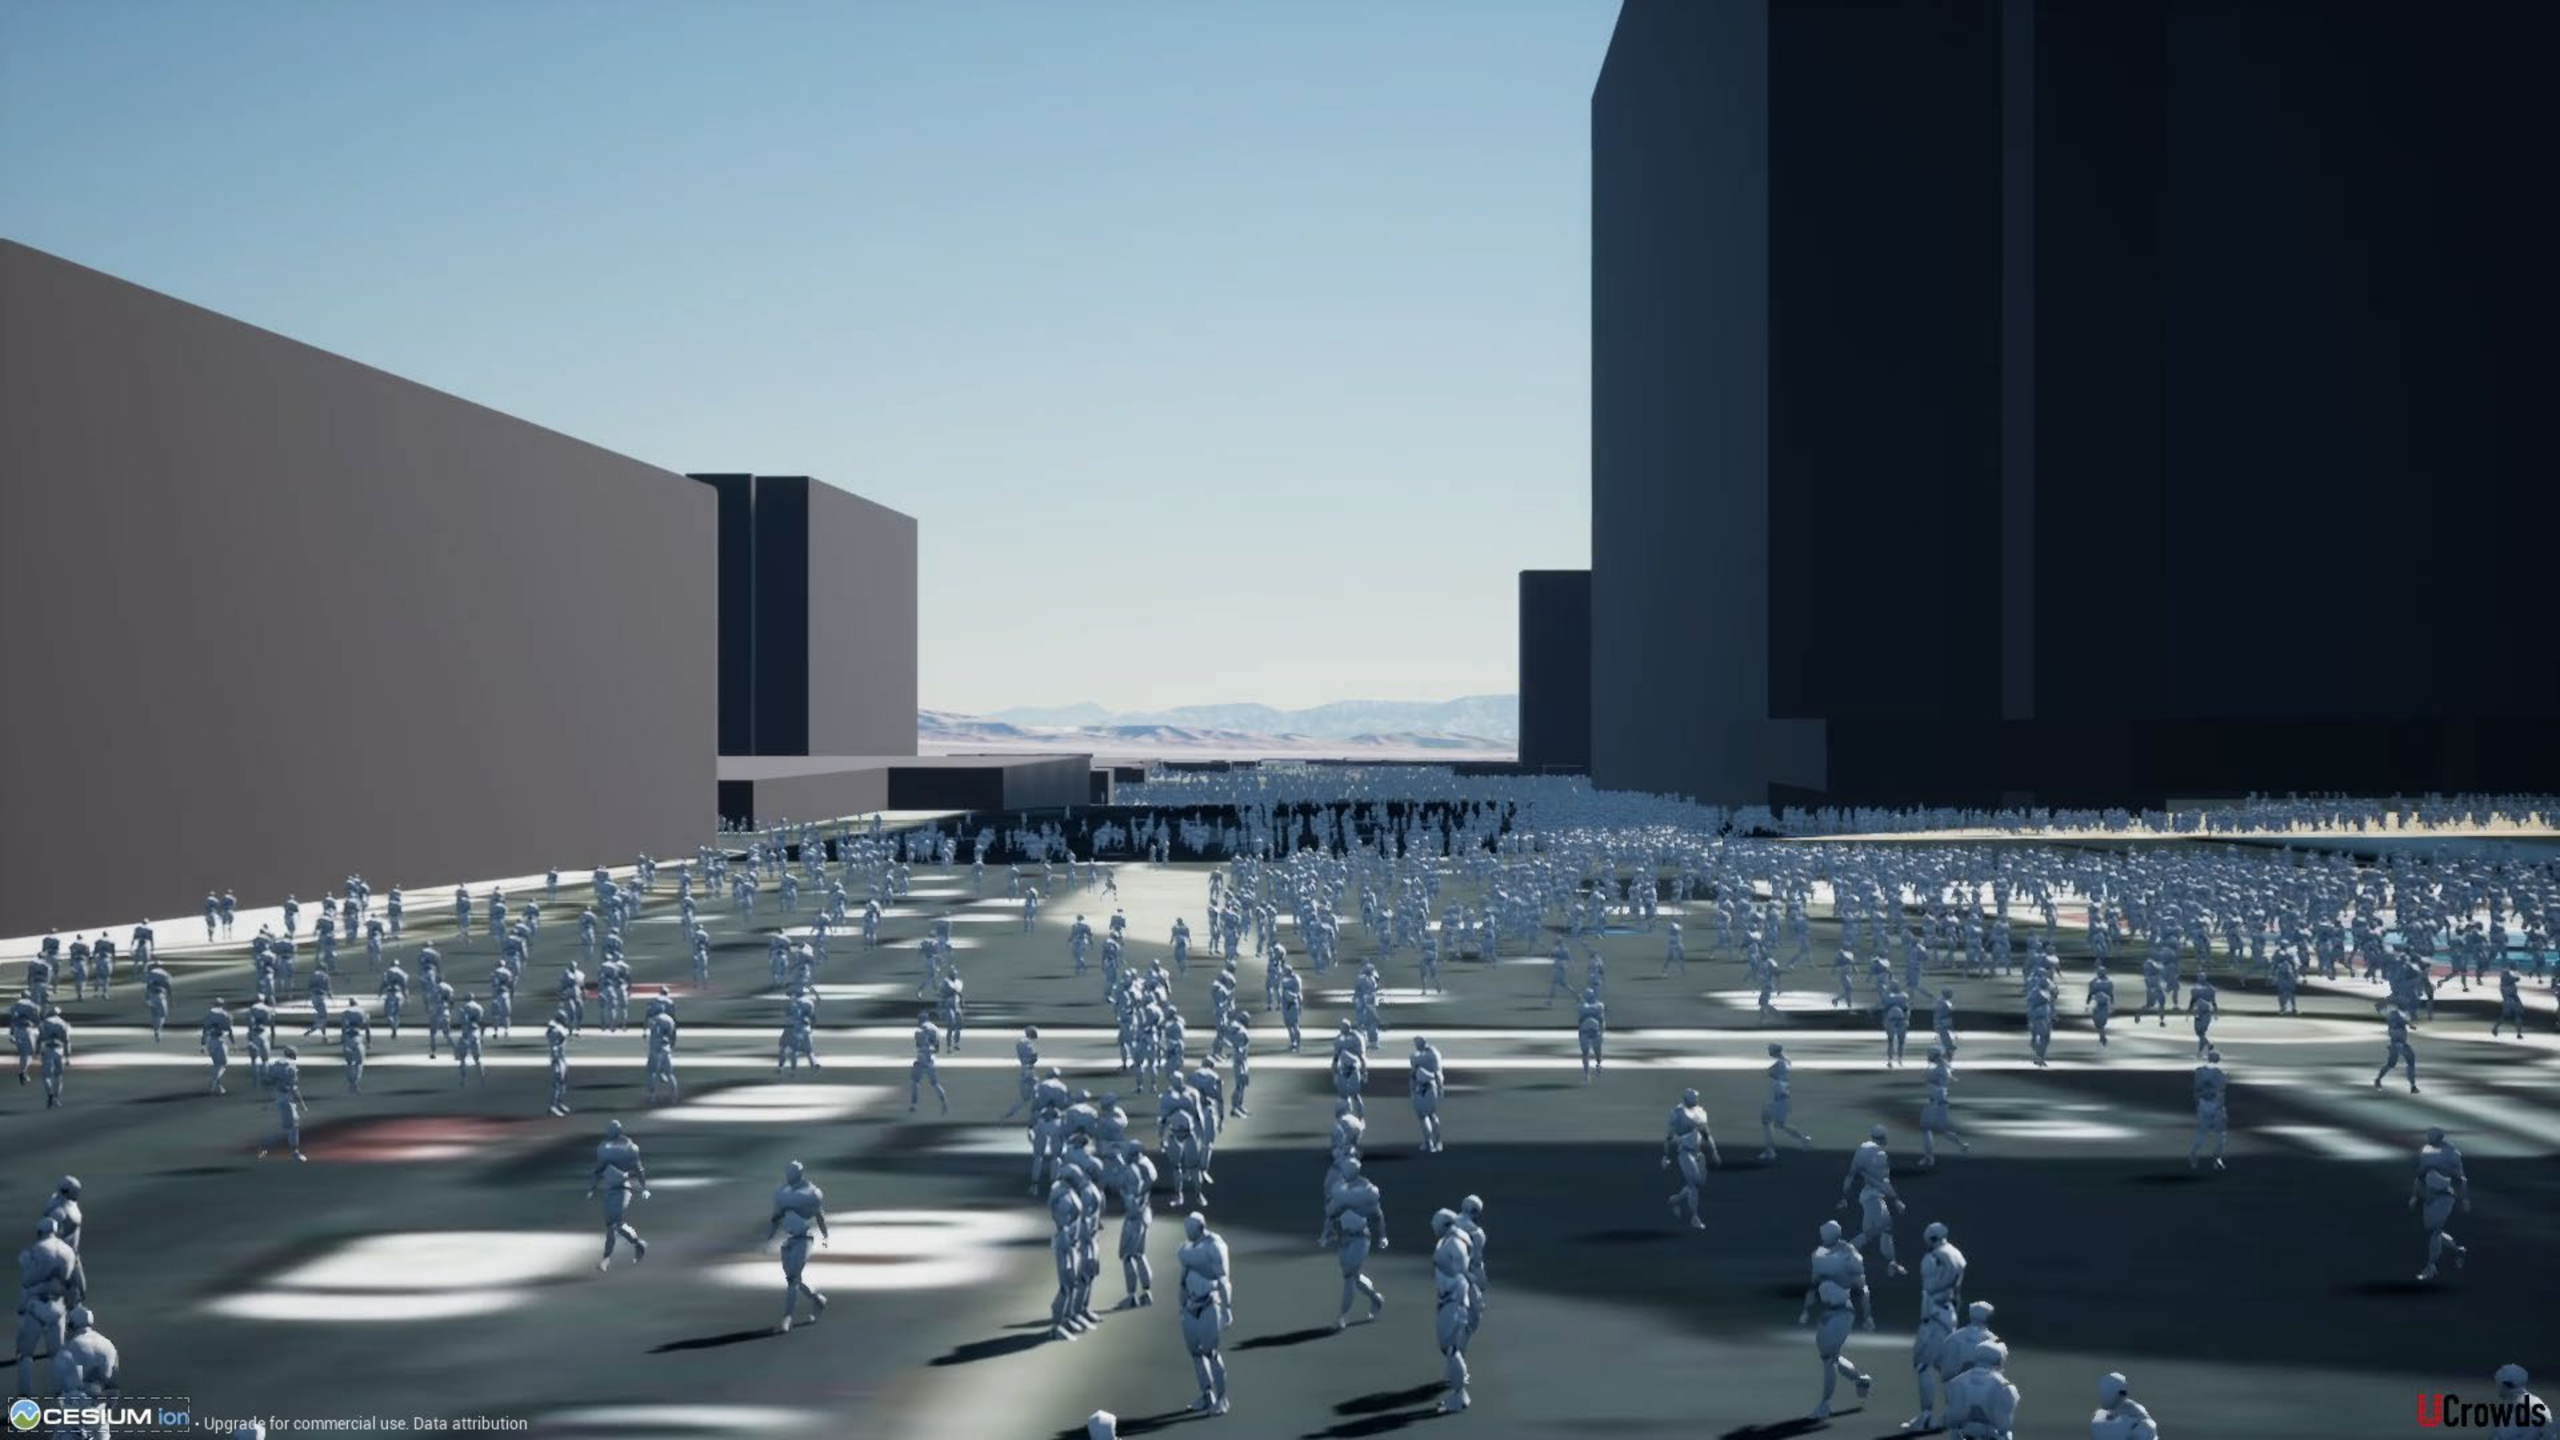
CESIUM ion · Upgrade for commercial use. Data attribution
UCrowds

Contact

I’m looking for collaborations (Dziękuję Ci)

Website	https://webspace.science.uu.nl/~gerae101/
E-mail	R.J.Geraerts@uu.nl
LinkedIn	http://www.linkedin.com/in/rolandgeraerts
Portfolio	www.youtube.com/user/drRolandJan
Address	Princetonplein 5, Utrecht, room 4.08

Universiteit Utrecht