



HAL
open science

Data-Driven Crowd Simulation with Generative Adversarial Networks

Javad Amirian, Wouter van Toll, Jean-Bernard Hayet, Julien Pettré

► **To cite this version:**

Javad Amirian, Wouter van Toll, Jean-Bernard Hayet, Julien Pettré. Data-Driven Crowd Simulation with Generative Adversarial Networks. CASA '19 - International Conference on Computer Animation and Social Agents, Jul 2019, Paris, France. pp.1-4, 10.1145/3328756.3328769 . hal-02134282

HAL Id: hal-02134282

<https://inria.hal.science/hal-02134282>

Submitted on 20 May 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Data-Driven Crowd Simulation with Generative Adversarial Networks

Javad Amirian
javad.amirian@inria.fr
Univ Rennes, Inria, CNRS, IRISA
Rennes, France

Jean-Bernard Hayet
jbhayet@cimat.mx
Centro de Investigación en Matemáticas
Guanajuato, Mexico

Wouter van Toll
wouter.van-toll@inria.fr
Univ Rennes, Inria, CNRS, IRISA
Rennes, France

Julien Pettré
julien.pettre@inria.fr
Univ Rennes, Inria, CNRS, IRISA
Rennes, France

ABSTRACT

This paper presents a novel data-driven crowd simulation method that can mimic the observed traffic of pedestrians in a given environment. Given a set of observed trajectories, we use a recent form of neural networks, Generative Adversarial Networks (GANs), to learn the properties of this set and generate new trajectories with similar properties. We define a way for simulated pedestrians (agents) to follow such a trajectory while handling local collision avoidance. As such, the system can generate a crowd that behaves similarly to observations, while still enabling real-time interactions between agents. Via experiments with real-world data, we show that our simulated trajectories preserve the statistical properties of their input. Our method simulates crowds in real time that resemble existing crowds, while also allowing insertion of extra agents, combination with other simulation methods, and user interaction.

CCS CONCEPTS

• **Computing methodologies** → **Intelligent agents; Neural networks; Motion path planning; Real-time simulation.**

KEYWORDS

crowd simulation, content generation, machine learning, intelligent agents, generative adversarial networks

ACM Reference Format:

Javad Amirian, Wouter van Toll, Jean-Bernard Hayet, and Julien Pettré. 2019. Data-Driven Crowd Simulation with Generative Adversarial Networks. In *Computer Animation and Social Agents (CASA '19)*, July 1–3, 2019, PARIS, France. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3328756.3328769>

1 INTRODUCTION

The realistic simulation of human crowd motion is a vast research topic that includes aspects of artificial intelligence, computer animation, motion planning, psychology, and more. Generally, the goal of a crowd simulation algorithm is to populate a virtual scene with

a crowd that exhibits visually convincing behavior. The simulation should run in real time to be usable for interactive applications such as games, training software, and virtual-reality experiences. Many simulations are *agent-based*: they model each pedestrian as a separate intelligent agent with individual properties and goals.

To simulate complex behavior, *data-driven* crowd simulation methods use real-world input data (such as camera footage) to generate matching crowd motion. Usually, these methods cannot easily generate *new* behavior that is not literally part of the input. Also, they are often difficult to use for applications in which agents need to adjust their motion in real-time, e.g. because the user is part of the crowd.

In this paper, we present a new data-driven crowd simulation method that largely avoids these limitations. Our system enables the real-time simulation of agents that behave similarly to observations, while allowing them to deviate from their trajectories when needed. More specifically, our method:

- (1) learns the overall properties of input trajectories, and can generate new trajectories with similar properties;
- (2) embeds these trajectories in a crowd simulation, in which agents follow a trajectory while allowing for local interactions.

For item 1, we use Generative Adversarial Networks (GANs) [4], a novel technique in machine learning for generating new content based on existing data. For item 2, we extend the concept of ‘route following’ [6] to trajectories with a *temporal* aspect, prescribing a speed that may change over time.

Using a real-world dataset as an example, we will show that our method generates new trajectories with matching styles. Our system can (for example) reproduce an existing scenario with additional agents, and it can easily be combined with other crowd simulation methods.

2 RELATED WORK

Agent-based crowd simulation algorithms model pedestrians as individual intelligent agents. In this paradigm, many researchers focus on the *local* interactions between pedestrians (e.g. collision avoidance) using *microscopic* algorithms [2, 5]. In environments with obstacles, these need to be combined with *global* path planning into an overall framework [12]. A growing research topic lies in measuring the ‘realism’ of a simulation, by measuring the similarity between two fragments of (real or simulated) crowd motion [13].

Publication rights licensed to ACM. ACM acknowledges that this contribution was authored or co-authored by an employee, contractor or affiliate of a national government. As such, the Government retains a nonexclusive, royalty-free right to publish or reproduce this article, or to allow others to do so, for Government purposes only.

CASA '19, July 1–3, 2019, PARIS, France

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-7159-9/19/07...\$15.00

<https://doi.org/10.1145/3328756.3328769>

Complex real-life behavior can hardly be described with simple local rules. This motivates *data-driven* simulation methods, which base the crowd motion directly on real-world trajectories, typically obtained from video footage. One category of such methods stores the input trajectories in a database, and then pastes the best-matching fragments into the simulation at run-time [7, 8]. Another technique is to create pre-computed *patches* with periodically repeating crowd motion, which can be copied and pasted throughout an environment [15]. Such simulations are computationally cheap, but difficult to adapt to interactive situations.

Researchers have also used input trajectories to train the parameters of (microscopic) simulation models [14], so as to adapt the agents' local behavior parameters to match the input data. However, this cannot capture any complex (social) rules that are not part of the used simulation model.

To replicate how agents move through an environment at a higher level, some algorithms subdivide the environment into cells and learn how pedestrians move between them [11, 16]. Our goal is similar (reproducing pedestrian motion at the full trajectory level), but our approach is different: we learn the spatial and temporal properties of complete trajectories, generate new trajectories with similar properties, and let agents flexibly follow these trajectories.

Our work uses Generative Adversarial Networks (GANs) [4], a recent AI development for generating new data. GANs have been successful at generating creative content such as faces [3]. Recently, researchers have started to adopt GANs for short-term prediction of pedestrian motion [1]. To our knowledge, our work is the first to apply GANs in crowd simulation at the full trajectory level.

3 GENERATING TRAJECTORIES

In this section, we describe our GAN-based method for generating trajectories that are similar to the examples in our training data.

As in most crowd-simulation research, we assume a planar environment and we model agents as disks. We define a *trajectory* as a mapping $\pi: [0, T] \rightarrow \mathbb{R}^2$ that describes how an agent moves through an environment during a time period of T seconds. Note that a trajectory encodes *speed* information: our system should capture when agents speed up, slow down, or stand still.

In practice, we will represent a trajectory π by a sequence of n_π points $[\mathbf{p}_0^\pi, \dots, \mathbf{p}_{n_\pi-1}^\pi]$ separated by a fixed time interval Δt ; that is, each \mathbf{p}_i^π has a corresponding timestamp $i \cdot \Delta t$. In our experiments, we use $\Delta t = 0.4s$ because our input data uses this value as well. We will use the notation $\mathbf{p}_{j:k}^\pi$ to denote a sub-trajectory from \mathbf{p}_j^π to \mathbf{p}_k^π .

Given a dataset of trajectories $\Pi = \{\pi_0, \dots, \pi_{N-1}\}$, our generator should learn to produce new trajectories with properties similar to those in Π . We assume that all trajectories start and end on the boundary of a *region of interest* R , which can have any shape and can be different for each environment.

Overview of GANs. A Generative Adversarial Network [4] consists of two components: a *generator* \mathcal{G} that creates new samples and a *discriminator* \mathcal{D} that judges whether a sample is real or generated. The training phase of a GAN is a two-player game in which \mathcal{G} learns to 'fool' \mathcal{D} , until (ideally) the generated samples are so convincing that \mathcal{D} does not outperform blind guessing.

Internally, both \mathcal{G} and \mathcal{D} are artificial neural networks; let θ_G and θ_D be their respective weights. \mathcal{G} acts as a function that converts an m -dimensional noise vector \mathbf{z} to a fake sample $\mathbf{x} = \mathcal{G}(\mathbf{z}; \theta_G)$. \mathcal{D} acts as a function that converts a (real or fake) sample \mathbf{x} to a value $\mathcal{D}(\mathbf{x}; \theta_D) \in [0, 1]$ indicating the probability of \mathbf{x} being real. Training a GAN represents the following optimization problem:

$$\min_{\theta_G} \max_{\theta_D} V(\theta_D, \theta_G), \quad (1)$$

$$V(\theta_D, \theta_G) = \mathbb{E}_{\mathbf{x} \sim p_x} [\log \mathcal{D}(\mathbf{x}; \theta_D)] + \mathbb{E}_{\mathbf{z} \sim p_z} [\log(1 - \mathcal{D}(\mathcal{G}(\mathbf{z}; \theta_G); \theta_D))]$$

where $V(\theta_D, \theta_G)$ is known as the *loss function*. Its first term denotes the expected output of \mathcal{D} for a random real sample. This is higher when \mathcal{D} correctly classifies more input samples as real. Conversely, the second term is higher when \mathcal{D} classifies more generated samples as fake. Here, p_x and p_z are the probability distributions of (respectively) the real data and the noise vectors sent to \mathcal{G} .

Overview of Our System. Figure 1 displays an overview of our GAN system. The generator and discriminator both have two tasks: generating or evaluating the *entry points* of a trajectory π (i.e. the first two points \mathbf{p}_0^π and \mathbf{p}_1^π), and generating or evaluating the *continuation* of a trajectory (i.e. the next point \mathbf{p}_{k+1}^π after a sub-trajectory $\mathbf{p}_{0:k}^\pi$). For the continuation tasks, we use concepts from so-called 'conditional GANs' because the generator and discriminator take extra data as input. We will now describe the system in more detail. Parameter settings will be mentioned in Section 5.

Generator. To generate *entry points*, the generator \mathcal{G} feeds a random vector \mathbf{z} to a fully connected (FC) block of neurons. Its output is a 4D vector that contains the coordinates of \mathbf{p}_0^π and \mathbf{p}_1^π .

To generate the *continuation* of a trajectory $\mathbf{p}_{0:k}^\pi$, the generator \mathcal{G} feeds $\mathbf{p}_{0:k}^\pi$ and a noise vector \mathbf{z} to a Long Short Term Memory (LSTM) layer that should encode the relevant trajectory dynamics. LSTMs are common *recurrent* neural networks used for handling sequential data. The output of this LSTM block is sent to an FC block, which finally produces a 2D vector with the coordinates of \mathbf{p}_{k+1}^π . Let $g(\mathbf{z}|\mathbf{p}_{0:k}^\pi; \theta_G)$ denote this result. Ideally, this point will be taken from the (unknown) distribution of likely follow-ups for $\mathbf{p}_{0:k}^\pi$.

The continuation step is repeated iteratively. If the newly generated point \mathbf{p}_{k+1}^π lies outside of the region of interest R , then the trajectory is considered to be finished. Otherwise, the process is repeated with inputs $\mathbf{p}_{0:k+1}^\pi$ and a new noise vector.

Discriminator. The discriminator \mathcal{D} takes an entire (real or fake) trajectory π as input. It splits the discrimination into two tasks with a similar structure as in \mathcal{G} . For the *entry point* part, an FC block evaluates $\mathbf{p}_{0:1}^\pi$ to a scalar in $[0, 1]$, which we denote by $v_e(\mathbf{p}_{0:1}^\pi; \theta_D)$. For the *continuation* part, an LSTM+FC block separately evaluates each point \mathbf{p}_k^π (for $2 \leq k < n_\pi$) given the sub-trajectory $\mathbf{p}_{0:k-1}^\pi$. We denote the result for the k th point by $v_c(\mathbf{p}_{0:k}^\pi; \theta_D)$.

So, for a full trajectory π of n_π points, the discriminator computes $n_\pi - 1$ scalars that together indicate the likelihood of π being real. The training phase uses these numbers in its loss function.

Training. Each training iteration lets \mathcal{G} generate a set Π' of N trajectories for different (sequences of) noise vectors. We then let \mathcal{D} classify all trajectories (both real and fake). The loss function of our GAN is the sum of two components:

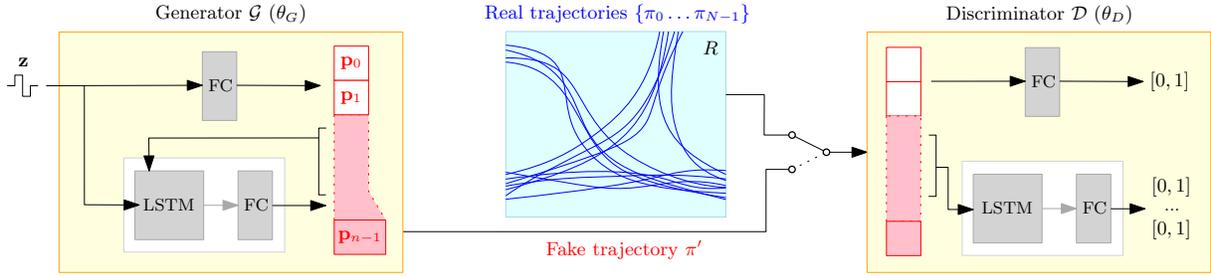


Figure 1: Our GAN architecture for learning and generating pedestrian trajectories.

- the success rate for discriminating all entry points:

$$\sum_{\pi \in \Pi} \log v_e(\mathbf{p}_{0:1}^\pi; \theta_D) + \sum_{\pi \in \Pi'} \log(1 - v_e(\mathbf{p}_{0:1}^\pi; \theta_D)),$$

- the success rate for discriminating all other points:

$$\sum_{\pi \in \Pi} \sum_{k=2}^{n_\pi-1} \log v_c(\mathbf{p}_{0:k}^\pi; \theta_D) + \sum_{\pi \in \Pi'} \sum_{k=2}^{n_\pi-1} \log(1 - v_c(\mathbf{p}_{0:k}^\pi; \theta_D)).$$

To let our GAN train faster, we add a third component. For each real trajectory $\pi \in \Pi$, we take all valid sub-trajectories $\mathbf{p}_{k:k+4}^\pi$ of length 5 and let \mathcal{G} generate its own version of $\mathbf{p}_{k:k+4}^\pi$ given $\mathbf{p}_{k:k+3}^\pi$. We add to our loss function:

$$\sum_{\pi \in \Pi} \sum_{k=0}^{n_\pi-5} \|\mathbf{p}_{k+4}^\pi - g(\mathbf{z}|\mathbf{p}_{0:k+3}^\pi; \theta_G)\|$$

i.e. we sum up the Euclidean distances between real and generated points. We observed that this additional component leads to much faster convergence and better back-propagation.

To reduce the chance of ‘mode collapse’ (i.e. convergence to a limited range of samples), we use an ‘unrolled’ GAN [9]. This is an extended GAN where each optimization step for θ_G uses an improved version of the discriminator that is u steps further ahead (where u is a parameter).

4 CROWD SIMULATION

Recall that our goal is to use our trajectories in a real-time interactive crowd simulation, where agents should be free to deviate from their trajectories if needed. This section describes how we combine our trajectory generator with a crowd simulator.

Our approach fits in the paradigm of multi-level crowd simulation [12], in which global planning (i.e. computing trajectories) is detached from the simulation loop. This loop consists of discrete timesteps. In each step, new agents might be added, and each agent tries to follow its trajectory while avoiding collisions.

Adding Agents. To determine when a new agent should be added to the simulation, we use an exponential distribution whose parameter λ denotes the average time between two insertions. This parameter can be obtained from an input dataset (to produce similar crowdedness), but one may also choose another value deliberately. Each added agent follows its own trajectory produced by our GAN.

Trajectory Following. In each frame of the simulation loop, each agent should try to proceed along its trajectory π while avoiding collisions. The main difference with classical ‘route following’ [6]

is that our trajectories have a *temporal* component: they prescribe at what speed an agent should move, and this speed may change over time. Therefore, we present a way to let an agent flexibly follow π while respecting its spatial *and* temporal data. Our algorithm computes a *preferred velocity* v_{pref} that would send the agent farther along π . This v_{pref} can then be fed to any existing collision-avoidance algorithm, to compute a velocity that is close to v_{pref} while avoiding collisions with other agents.

Two parameters define how an agent follows π : the *time window* w and the *maximum speed* s_{max} . An agent always tries to move to a point that lies w seconds ahead along π , taking s_{max} into account. During the simulation, let t be the time that has passed since the agent’s insertion. Ideally, the agent should have reached $\pi(t)$ by now. Our algorithm consists of the following steps:

- (1) Compute the *attraction point* $\mathbf{p}_{\text{att}} = \pi(t_{\text{att}})$, where $t_{\text{att}} = \min(t + w, T)$ and T is the end time of π . Thus, \mathbf{p}_{att} is the point that lies w seconds ahead of $\pi(t)$, clamped to the end of π if needed.
- (2) Compute the *preferred velocity* \mathbf{v}_{pref} as $\frac{\mathbf{p}_{\text{att}} - \mathbf{p}}{t_{\text{att}} - t}$, where \mathbf{p} is the agent’s current position. Thus, \mathbf{v}_{pref} is the velocity that will send the agent to \mathbf{p}_{att} , with a speed based on the difference between t and t_{att} .
- (3) If $\|\mathbf{v}_{\text{pref}}\| > s_{\text{max}}$, scale \mathbf{v}_{pref} so that $\|\mathbf{v}_{\text{pref}}\| = s_{\text{max}}$. This prevents the agent from receiving a very high speed after it has been blocked for a long time.

Collision Avoidance. The preferred velocity \mathbf{v}_{pref} computed by our algorithm can be used as input for any collision-avoidance routine. In our implementation, we use the popular ORCA method [2]. In preliminary tests, other methods such as social forces [5] proved to be less suitable for our purpose.

5 EXPERIMENTS AND RESULTS

Set-up. We have implemented our GAN using the *PyTorch* library (<https://pytorch.org/>). The input noise vectors are 3-dimensional and drawn from a uniformly random distribution. In both \mathcal{G} and \mathcal{D} , the entry-point FC blocks consist of 3 layers with 128, 64, and 32 hidden neurons, respectively. For the continuation part, the LSTM blocks consist of 62 cells, and the FC blocks contain 2 layers of 64 and 32 hidden neurons. To save time and memory, the LSTM blocks only consider the last 4 samples of a sub-trajectory.

For training the GAN, all FC layers use *Leaky-ReLU* activation functions (with negative slope 0.1), to let the gradient always back-propagate, which avoids gradient-vanishing issues. We train the GAN for 50,000 iterations, using an unrolling parameter $u = 10$.

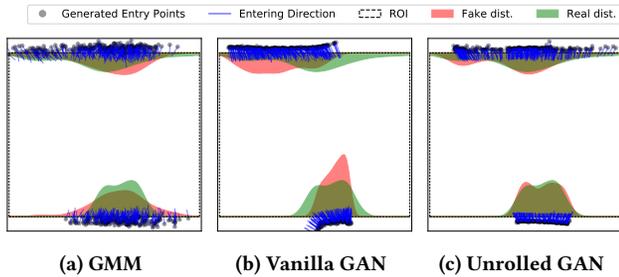


Figure 2: The distribution of entry points created by three different methods.

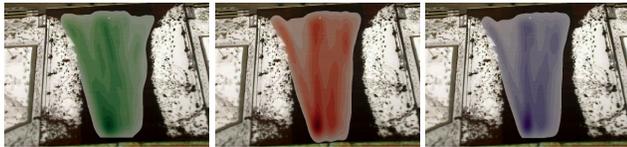


Figure 3: Trajectory heatmaps: the input data, the generated trajectories, and the final simulated agent motion.

In the crowd simulation, we model agents as disks with a radius of 0.25m, and we use a simulation loop with a fixed frame length of 0.1s. In each frame, all agents perform our route-following algorithm (with $w = 5$ and $s_{\max} = 2\text{m/s}$), followed by the ORCA algorithm [2] as implemented by the original authors. We remove an agent when it reaches the end of its trajectory.

We test our method on the *ETH* dataset [10] that contains recorded trajectories around the entrance of a university building. We have defined the region of interest R as an axis-aligned bounding box, and we use only the 241 trajectories that both enter and exit R .

Result 1: Entry Points. To show the performance of our GAN in learning the distribution of entry points, we computed 500 (fake) entry points in the *ETH* scene, and we calculated the distribution of the samples over the boundary of R . We also compared these results against two other generative methods: a Gaussian Mixture Model (GMM) with 3 components, and a ‘vanilla’ GAN variant that does not use the unrolling mechanism. As shown in Fig. 2, the entry points of the unrolled GAN (right) are closer to the real data than those of the other two methods.

Result 2: Trajectories. Next, we used our system to generate 352 new trajectories, and we used them to simulate a crowd. The first two heatmaps in Fig. 3 show that generated trajectories (middle) are similarly distributed over the environment as the real data (left).

The third heatmap shows the final motion of the simulated agents with route following and collision avoidance. In this scenario, agents are well capable of following their given trajectories.

Computation time. We used CUDA to run our GAN on a NVIDIA Quadro M1200 GPU with 4GB of GDDR5 memory. With this set-up, generating a batch of 1024 trajectories (with a maximum length of 40 points) took 152ms, meaning that the average generation time was 0.15ms per trajectory. Thus, after training, the system is sufficiently fast for real-time insertions of trajectories into a crowd.

6 CONCLUSIONS & FUTURE WORK

We have presented a data-driven crowd simulation method that uses GANs to learn the properties of input trajectories and then generate new trajectories with similar properties. Combined with flexible route following that takes temporal information into account, the trajectories can be used in a real-time crowd simulation. Our system can be used, for example, to create variants of a scenario with different densities. It can easily be combined with other simulation methods, and it allows interactive applications.

In the future, we will perform a thorough analysis of the trajectories produced by our system, and compare them to other algorithms. We will also investigate the exact requirements for reliable training. Furthermore, our system generates trajectories for individuals, assuming that agents do not influence each other’s choices. As such, it cannot yet model group behavior, and it performs worse in high-density scenarios where agents cannot act independently. We would like to handle these limitations in future work.

ACKNOWLEDGMENTS

This project was partly funded by EU project CROWDBOT (H2020-ICT-2017-779942).

REFERENCES

- [1] AMIRIAN, J., HAYET, J.-B., AND PETTRÉ, J. Social ways: Learning multi-modal distributions of pedestrian trajectories with GANs. In *CVPR Workshops* (2019).
- [2] VAN DEN BERG, J., GUY, S., LIN, M., AND MANOCHA, D. Reciprocal n-body collision avoidance. In *Proc. Int. Symp. Robotics Research* (2011), pp. 3–19.
- [3] DI, X., AND PATEL, V. Face synthesis from visual attributes via sketch using conditional VAEs and GANs. *CoRR abs/1801.00077* (2018).
- [4] GOODFELLOW, I., POUGET-ABADIE, J., MIRZA, M., XU, B., WARDE-FARLEY, D., OZAIR, S., COURVILLE, A., AND BENGIO, Y. Generative adversarial nets. In *Proc. Int. Conf. Neural Information Processing Systems* (2014), vol. 2, pp. 2672–2680.
- [5] HELBING, D., AND MOLNÁR, P. Social force model for pedestrian dynamics. *Physical Review E* 51, 5 (1995), 4282–4286.
- [6] JAKLIN, N., COOK IV, A., AND GERAERTS, R. Real-time path planning in heterogeneous environments. *Computer Animation and Virtual Worlds* 24, 3 (2013), 285–295.
- [7] LEE, K., CHOI, M., HONG, Q., AND LEE, J. Group behavior from video: A data-driven approach to crowd simulation. In *Proc. ACM SIGGRAPH/Eurographics Symp. Computer Animation* (2007), pp. 109–118.
- [8] LERNER, A., CHRYSANTHOU, Y., AND LISCHINSKI, D. Crowds by example. *Computer Graphics Forum* 26, 3 (2007), 655–664.
- [9] METZ, L., POOLE, B., PFAU, D., AND SOHL-DICKSTEIN, J. Unrolled generative adversarial networks. *CoRR abs/1611.02163* (2017).
- [10] PELLEGRINI, S., ESS, A., SCHINDLER, K., AND VAN GOOL, L. You’ll never walk alone: Modeling social behavior for multi-target tracking. In *Proc. IEEE Int. Conf. Computer Vision* (2009), pp. 261–268.
- [11] PELLEGRINI, S., GALL, J., SIGAL, L., AND VAN GOOL, L. Destination flow for crowd simulation. In *Proc. European Conf. Computer Vision* (2012), pp. 162–171.
- [12] VAN TOLL, W., JAKLIN, N., AND GERAERTS, R. Towards believable crowds: A generic multi-level framework for agent navigation. In *ASCI.OPEN / ICT.OPEN (ASCI track)* (2015).
- [13] WANG, H., ONDŘEJ, J., AND O’SULLIVAN, C. Path patterns: Analyzing and comparing real and simulated crowds. In *Proc. 20th ACM SIGGRAPH Symp. Interactive 3D Graphics and Games* (2016), pp. 49–57.
- [14] WOLINSKI, D., GUY, S., OLIVIER, A.-H., LIN, M., MANOCHA, D., AND PETTRÉ, J. Parameter estimation and comparative evaluation of crowd simulations. *Computer Graphics Forum* 33, 2 (2014), 303–312.
- [15] YERSIN, B., MAÏM, J., PETTRÉ, J., AND THALMANN, D. Crowd patches: Populating large-scale virtual environments for real-time applications. In *Proc. Symp. Interactive 3D Graphics and Games* (2009), pp. 207–214.
- [16] ZHONG, J., CAI, W., LUO, L., AND ZHAO, M. Learning behavior patterns from video for agent-based crowd modeling and simulation. *Autonomous Agents and Multi-Agent Systems* 30, 5 (2016), 990–1019.