

Minimizing the evacuation time of a crowd from a complex building using rescue guides

Anton von Schantz^{1,*}, Harri Ehtamo¹

Aalto University, Department of Mathematics and Systems Analysis, P.O. Box 11100, FI-00076 Aalto, Finland

Abstract

In an emergency situation, the evacuation of a large crowd from a complex building can become slow or even dangerous without a working evacuation plan. The use of rescue guides that lead the crowd out of the building can improve the evacuation efficiency. An important issue is how to choose the number, positions, and exit assignments of these guides to minimize the evacuation time of the crowd. Here, we model the evacuating crowd as a multi-agent system with the social force model and simple interaction rules for guides and their followers. We formulate the problem of minimizing the evacuation time using rescue guides as a stochastic control problem. Then, we solve it with a procedure combining numerical simulation and a genetic algorithm (GA). The GA iteratively searches for the optimal evacuation plan, while numerical simulations evaluate the evacuation time of the plans. We apply the procedure on a test case and on an evacuation of a fictional conference building. The procedure is able to solve the number of guides, their initial positions and exit assignments in a single although complicated optimization. The attained results show that the procedure converges to an optimal evacuation plan, which minimizes the evacuation time and mitigates congestion and the effect of random deviations in agents' motion.

Keywords: Multi-agent systems, Numerical simulation, Evacuation, Stochastic optimization, Genetic algorithms

© 2020. This manuscript version is made available under the CC-BY-NC-ND 4.0 license <http://creativecommons.org/licenses/by-nc-nd/4.0/>

*Corresponding author

Email addresses: anton.von.schantz@aalto.fi (Anton von Schantz), harri.ehtamo@aalto.fi (Harri Ehtamo)

1. Introduction

Large complex buildings like airport terminals, high-rise buildings and subway stations need crowd management to ensure safety. In an emergency situation, like a fire or a bomb threat, a well-planned operation is needed to ensure a fast and safe evacuation. It is known that people are willing to follow authorities (Proulx, 2002; Nishinari et al., 2004), and that the strategic use of trained safety personnel, or rescue guides, that lead the crowd out of the building, improves evacuation efficiency (Hou et al., 2014). However, as the crowd size and complexity of the building increases, codes of practice and human intuition are not sufficient to form an optimal evacuation plan. Rather attention needs to be paid to mathematical optimization approaches (Haghani, 2020).

In this paper, we are interested in what will happen in a complex building containing a large crowd, after the sound of a whistle or a fire alarm, in a serious situation, when there is enough yellow-coated rescue personnel, whose members have enough skill and authority to guide the crowd out of the building in minimum time. Even more so, how should these guides be coordinated so that they manage in their task in an optimal way. To our knowledge, this problem has not been solved using a rigorous mathematical optimization approach.

In the research field of crowd and evacuation dynamics, three distinct research streams can be distinguished: experimental research on pedestrian behavior, descriptive mathematical modeling of crowd movement and interactions, and mathematical optimization of evacuation. A good source of references on crowd movement models and optimization approaches is (Vermuyten et al., 2016). The most popular crowd movement model due to its ability to produce realistic physical movement is the microscopic agent-based social force model (Helbing & Molnár, 1995). In it the movement of individual agents are described by Newton's equations of motion.

A typical concern is that focusing mainly on physical movement in crowd evacuation modeling may potentially lead to an underestimation of the evacuation time, since evacuees can engage in a variety of activities that do not move them towards the exits (Gwynne et al., 2016). We assume, that in an evacuation of a large crowd controlled by guides, such activities can be important but rare. From a modeling perspective, these behavioral deviations are small compared to the systemic, deterministic part of motion (Helbing & Johansson, 2013). Deviations from the usual rules of motion can be approximated by adding a Gaussian noise term to the equations of motion (Helbing & Molnár, 1995). So, in this study, our main focus is not to tackle various human behaviors, but rather to consider how the guides should be positioned and what kind of routes they should take in a minimum time evacuation, provided some simple assumptions about the crowd behavior in such a situation are made.

There is a need for mathematical approaches for evacuation optimization (Kurdi et al., 2018; Haghani, 2020). Typically, the optimal guided evacuation has been studied by comparative analysis, where the evacuation is simulated and compared for different guide configurations. The optimal number of guides has been studied in, e.g., (Pelechano & Badler, 2006; Wang et al., 2012; McCormack & Chen, 2014). Also, initial positioning of guides has been studied in, e.g., (Aubé & Shield, 2004; Wang et al., 2012; Cao et al., 2016). It has been suggested

that embedded, peripheral, near exit and uniformly distributed positions improve the evacuation efficiency. The importance of coordination has been noted in (Hou et al., 2014; Cao et al., 2016). Actually, without strategically planned initial positions and exit assignments of guides, their use can worsen the evacuation instead of improving its efficiency.

To our knowledge, there are only two noteworthy studies that use mathematical optimization to study the evacuation of a crowd using guides (Albi et al., 2016; Zhou et al., 2019). The lack of research is probably due to the large state space caused by a moving crowd and the stochastic nature of the problem. In (Albi et al., 2016) the crowd evacuates from a simple building. Its members are assumed to be unfamiliar with the environment. Hence, their movement is a combination of random walk and following nearby agents. Invisible guides are added to the crowd, i.e., the other agents in the crowd do not recognize them as guides. The minimum evacuation time problem is formulated as a stochastic control problem, where the trajectories of the guides are solved. However, the number of guides is not an optimization variable. Albi et al. (2016) do not directly generate an implementable evacuation plan for rescue guides, but rather do a mathematical study on optimal herding of a crowd to an exit. Still, with small changes, their modeling framework can be used for visible rescue guides. Nevertheless, we are skeptical that the solution methods they propose for the problem would be applicable for a more complicated building geometry with the crowd scattered over it.

The study of Zhou et al. (2019) is a simultaneous but independent line of work with ours, and the two works were done without knowledge of each other. Zhou et al. (2019) use a deterministic version of the social force model to describe the crowd movement. At the start they optimize the number and initial positions of the guides using a suitable criterion. After that, they optimize their routes and exits given their starting positions. The two optimization problems are independent of each other, and neither the criterion nor the starting positions of the guides are corrected based on the information got from the second problem. It should be noted that in the first problem any suitable criterion, not only that used by Zhou et al. (2019), can be used to fix the number and starting positions of the guides. All they give a different value for the evacuation time; the minimum value can be obtained only by chance. The model of Zhou et al. (2019) returns an evacuation plan that does not minimize the evacuation time, but does something else.

Furthermore, both (Albi et al., 2016; Zhou et al., 2019) use their modified versions of the social force model, which are in some sense more sophisticated than typically, but they can only be used for optimization and evacuation simulations in simple building geometries. In them, the follower agents might get lost and are unable to follow the guide, if the guide goes into another room, or even behind a corner (Li et al., 2016).

In this paper, we formulate and solve the problem of minimizing the evacuation time of a crowd from a complex building using rescue guides. The problem is formulated as a stochastic control problem, where the objective is to minimize the expected evacuation time. The optimization variables are the number of guides, and their routes represented as origin-destination pairs. The state of the system is modeled with social force equations and simple rules of interaction between guides and exiting agents. The resulting problem solution

space is very large, and it cannot be solved with derivative-based optimization. Thus, we first reformulate the problem as a scenario optimization problem (Calafiore & Campi, 2006), where the expected evacuation time is replaced by its sampled mean. Then we apply a hybrid numerical simulation and genetic algorithm (GA) approach (Goldberg, 1989). In it, the GA iteratively searches for the optimal solution, and numerical simulations evaluate their sample mean evacuation time and steers the randomized search process. We ensure the accuracy and efficiency of the algorithm on a test case, and then apply it on a more complex conference building case. Our three main contributions are that our method applies for complex buildings, it takes into account stochasticity and gives the number of guides, their initial positions and exit assignments needed to minimize the crowd evacuation time in a single optimization.

The paper is structured as follows. In Sec. 2, we discuss the assumptions and mathematical details of the crowd movement model. In Sec. 3, we formulate the optimization problem, and present the GA approach used to solve it. The performance of the GA is analyzed on a test case in Sec. 4, and after that the GA is applied on a more complex conference building case in Sec. 5. Then, in Sec. 6, we revisit and discuss the behavioral assumptions made. Sec. 7 is for implementation details and performance of our algorithm. Finally, Sec. 8 is for conclusion. For more information about parameter values and exact mathematical expressions of the social force model, see Appendix A. And, for a detailed analysis on the effect of stochasticity on our problem, see Appendix B.

2. Evacuation model with guides

The social force model was first presented by Helbing & Molnár (1995). In it the motion of a pedestrian is described by Newton's equations of motion. The model is based on a paradigm from social theory, where an action taken by a pedestrian is understood as a measure of her personal motivation to perform that action. That action is taken according to her psychological processes of assessment of alternatives and utility maximization. In the case of pedestrian behavior, this motivation evokes the physical production of an acceleration force. One can say that a pedestrian acts as if she would be subject to external forces (Helbing & Molnár, 1995; Hoogendoorn & Bovy, 2003; Helbing & Johansson, 2013). Additional physical interaction forces, inspired from driven granular flows, are assumed to come into play, when pedestrians get so close to each other that they have physical contact.

The social force model has been integrated into the Fire Dynamics Simulator with Evacuation (FDS+Evac) of the National Institute of Standards and Technology (NIST, 2020). We will use the version of the social force model found in the FDS+Evac user manual (Korhonen, 2018). But instead of approximating the human body with three overlapping circles (Heliövaara et al., 2012a), we use a single circle. Also, we make a small modification to the social force term. In the original social force model, a distance-dependent term is assumed. However, Karamouzas et al. (2014) have analyzed a large amount of publicly available crowd data, from several outdoor environments and controlled laboratory settings, showing that the social force depends on

the time-to-collision between agents rather than their distance. We will use the social force term presented in (Karamouzas et al., 2014).

2.1. Detailed description

The evacuation model describes the evacuation of a crowd of agents $N \cup G$ from a space $\Omega \subset \mathbb{R}^2$ after the alarm has been given. The space includes obstacles $\mathcal{O} \subset \Omega$ and exits $\mathcal{E} \subset \Omega$. The obstacles are line segments, or walls, with index set W . The crowd consists of exiting agents N and guide agents, or guides, G . Here $N = \{1, \dots, n\}$ and $G = \{n+1, \dots, n+m\}$ are the index sets of the exiting agents and guides, respectively. Exiting agents represent regular people that do not have full knowledge of the exits in the building, and head to their familiar exits by default. The guides represent trained safety personnel that use routes instructed by the evacuation planner. They are assumed to have enough authority to influence the route choices of exiting agents. More specifically, the first time an exiting agent is within the interaction range r^{guide} from a guide, it receives information of the exit the guide is moving towards, and starts also heading there. However, in a threatening situation, we assume that an exiting agent wants to get out of the building as fast as possible. That is why, if an exit is within the visibility range r^{exit} , it starts to head there instead, regardless of where it was previously heading. In all situations, the exiting agents move to the exits using a shortest path (for information on its numerical computation see Sec. 7).

It is assumed that a mixture of socio-psychological and physical forces influence the agents' motion in the crowd. At time t , agent $i \in N \cup G$ with mass m_i and radius r_i likes to move with a certain desired velocity $\mathbf{v}_i^{des}(t)$, where $\mathbf{v}_i^{des}(t) = v_i^{des}(t)\mathbf{e}_i^{des}(t)$. Here, $v_i^{des}(t)$ is its desired speed, and $\mathbf{e}_i^{des}(t)$ is a unit vector that points to the direction that gives the shortest path to the exit it is heading towards. Agent i attempts to change its actual velocity $\mathbf{v}_i(t)$ to $\mathbf{v}_i^{des}(t)$ with a reaction time τ^{reac} .

If agents i and j , $j \in N \cup G, j \neq i$, are about to collide, they experience a repulsive social force $\mathbf{f}_{ij}^{soc}(t)$. When agent i is in contact with another agent j or wall $w \in W$, the physical contact forces $\mathbf{f}_{ij}^c(t)$ or $\mathbf{f}_{iw}^c(t)$, respectively, arise. Additionally, we assume that agent i is affected by a small random fluctuation force $\boldsymbol{\xi}_i(t)$. Typically, the social force model includes a random force term in each agent's equation of motion. This force represents random deviations in behavior; more precisely, situations where two or more behavioral alternatives are equivalent, i.e., whether to pass an obstacle from the left or right hand side. It can also be thought to describe accidental or intentional deviations from the usual rules of motion (Helbing & Molnár, 1995). The force $\boldsymbol{\xi}_i(t)$ is drawn from a truncated bivariate normal distribution with zero $\mathbf{0}$ mean. However, for guides this force equals a zero vector $\mathbf{0}$. For the parameter values used and the exact mathematical expressions of the forces, see Appendix A.

The change of velocity at time t for agent i is then given by the equation of motion:

$$m_i \frac{d\mathbf{v}_i}{dt} = m_i \frac{\mathbf{v}_i^{des} - \mathbf{v}_i}{\tau^{reac}} + \sum_{j(\neq i)} (\mathbf{f}_{ij}^{soc} + \mathbf{f}_{ij}^c) + \sum_w \mathbf{f}_{iw}^c + \boldsymbol{\xi}_i; \quad (1)$$

the change of position vector $\mathbf{x}_i(t)$ is given by the velocity

$$\frac{d\mathbf{x}_i(t)}{dt} = \mathbf{v}_i(t). \quad (2)$$

In addition, we assume the following interactions to take place in the evacuation:

Assumption 1. An exiting agent heads to its familiar exit by default.

Assumption 2. The guides' initial positions and target exits are optimization variables. A guide moves from its starting position to its target exit using the shortest path.

Assumption 3. If an exiting agent is moving towards its familiar exit, and a guide comes within the interaction range r^{guide} , it starts to move to the same exit as the guide. If multiple guides are within the r^{guide} range, it starts to follow the closest one.

Assumption 4. If an exiting agent follows a guide, it will not switch to follow another guide.

Assumption 5. If an exit is within the visibility range r^{exit} from an exiting agent, it heads there, whether it previously was following a guide or heading to its familiar exit. If multiple exits are within the r^{exit} range, it heads to the closest one.

3. Optimization model

We are interested in minimizing the evacuation time of the crowd, or equivalently the evacuation time of the last evacuated agent T_{last} . If we denote the evacuation times of the agents in the crowd with t_1, \dots, t_{n+m} , the maximal element of the set is T_{last} . Because of the random fluctuation force term $\xi_i(t)$ in Eq. (1), our problem is stochastic. Thus, as objective function we choose to minimize the expected value of T_{last} with respect to $\xi_i(t), t \in [0, T_{last}], 1 \leq i \leq n$. Recall that for $i \in G$ the random fluctuation force term is equal to $\mathbf{0}$.

We discretize the space Ω into square grid cells ω , so that $\Omega \subset \cup \omega = \bar{\Omega}$. Also, we denote the points associated with an exit by $\varepsilon \subset \mathcal{E}$, so that $\mathcal{E} = \cup \varepsilon$. Each guide $g \in G = \{n+1, \dots, n+m\}$ is associated with a starting grid cell $\omega_g \subset \bar{\Omega}$, and a target exit $\varepsilon_g \subset \mathcal{E}$. The optimization variables are $(\omega_g, \varepsilon_g), g \in G$. We assume that the number m of the possible guides is sufficiently large, so that one or more guides can remain idle in the simulation. The idle guides are mapped to a dummy grid cell outside $\bar{\Omega}$, and their evacuation time is defined to be zero. Our focus here is either a fixed number of active guides used in optimization, or a set of active guides obtained as a result of optimization.

The initial position of a guide g , $\mathbf{x}_g(0) \in \omega_g$, is a prespecified point in its corresponding starting grid cell ω_g . The end position of a guide g , $\mathbf{x}_g(t_g) \in \varepsilon_g$, is any point in its corresponding target exit ε_g . The initial positions of the exiting agents are prespecified $\mathbf{x}_i(0) = \mathbf{x}_{i,0}, i = 1, \dots, n$. The exiting agents can evacuate using any of the available exits; thus it holds for the end positions $\mathbf{x}_i(t_i) \in \mathcal{E}$.

The agents move to the exits according to Eqs. (1), (2) and Assumptions 1-5 defined in Sec. 2. These equations constitute the constraint equations of this problem. Now we can formulate the problem of minimizing the evacuation time of a crowd using rescue guides:

$$\begin{aligned} \min_{(\omega_g, \varepsilon_g)} \quad & \mathbb{E}[T_{last} | \xi_i(t), 1 \leq i \leq n, t \in [0, T_{last}]]; \\ \omega_g \subset \bar{\Omega}, \varepsilon_g \subset \mathcal{E}, g \in G \end{aligned} \tag{3}$$

subject to Eqs. (1), (2); Assumptions 1-5;

$$\mathbf{x}_i(0) = \mathbf{x}_{i,0}, \mathbf{x}_i(t_i) \in \mathcal{E}, i \in N; \mathbf{x}_g(0) \in \omega_g, \mathbf{x}_g(t_g) \in \varepsilon_g, g \in G.$$

To solve problem Eq. (3) we will use scenario optimization (Calafiore & Campi, 2006). The expected evacuation time of the last agent is replaced by its sampled empirical version: M independent identically distributed samples of the random force $\xi_i, 1 \leq i \leq n$, are generated which we denote by $\boldsymbol{\delta}^{(1)}, \dots, \boldsymbol{\delta}^{(M)}$, and then we define the sample mean,

$$\hat{T}_{last} = \frac{1}{M} \sum_{k=1}^M T_{last}^{(k)}. \tag{4}$$

Here, $T_{last}^{(k)}$ is calculated for given $(\omega_g, \varepsilon_g), g \in G$, by numerically simulating the system equations Eqs. (1), (2) with a numerical integration scheme, given the sample vector $\boldsymbol{\delta}^{(k)}$, Assumptions 1-5 and the initial and end conditions (see more implementation details in Sec. 7).

The scenario optimization problem is solved with a genetic algorithm (GA) (Goldberg, 1989). The GA iteratively searches for the optimal solution, while the numerical simulations evaluate the fitness of the found solutions and steer the randomized search process. In our optimization problem, the fitness is the sample mean Eq. (4), and a solution, or chromosome, is an evacuation plan $\{(\omega_{n+1}, \varepsilon_{n+1}), \dots, (\omega_{n+m}, \varepsilon_{n+m})\}$. A single gene in a chromosome contains both the starting grid cell and target exit of a guide. Because a solution has a variable number of active guides, we apply the hidden genes GA (Abdelkhalik, 2013). In it each gene gets a tag, which tells if the gene is active, or if it is hidden or idle. Only the active genes affect the evacuation simulation. However, the genetic information of the idle genes is carried on in the process.

In the GA, a population of a predefined number of solutions is maintained in the consecutive iterations, or generations. At the initial iteration, a population of random solutions is generated. After that, the population goes through three operations, which are selection, crossover and mutation. In selection, the solutions with smallest fitness values are chosen to undergo the next two operations. In this paper, we use elitist selection, where we replace a few of the worst solutions in the current generation with the best solutions of the previous generation.

In the crossover operation, genetic information is combined to create new solutions. In this paper, we use the single-point crossover operator, meaning that the offspring solutions of the two parent solutions contain half

of the genetic material of both parents. The crossover operator is applied with a certain probability, and it is applied on both genes and the tags. If it is not applied, the offspring solutions have exactly the same genetic information as their parents. See Fig. 1 for a depiction of the mechanism; if the gene is colored gray, the gene is hidden, which means that the guide corresponding to the gene is not present in the evacuation, i.e., is idle. In the single-point crossover, the genetic material on the right side of the crossover point of the two parent chromosomes is changing places to create offspring chromosomes.

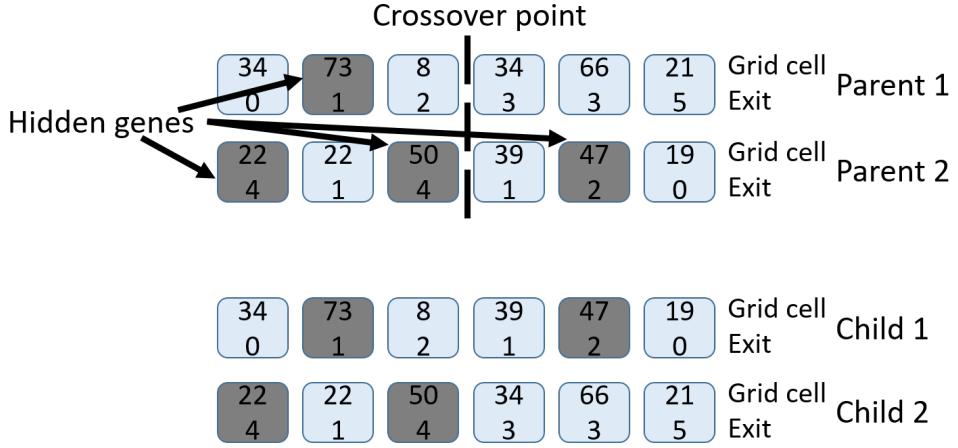


Figure 1: The single-point crossover operation applied on the two parent chromosomes to create offspring solutions.

Then, a mutation operator is applied on the offspring solutions. The mutation operator is applied separately, with a certain probability, on each gene and tag. When applied on a gene, it can either alter the starting grid cell, or the target exit, or both of them. The mutation operator applied on a tag can switch the gene from hidden to active, or vice versa. Finally, the GA goes back to evaluating the fitness values of the solutions of the new generation, i.e., the offspring solutions, and then proceeds to the selection operation, and the iteration continues. The algorithm has converged, when the best solution has not changed in a predefined number of successive iterations.

4. Evacuation of a hexagon-shaped area

Typically, to solve an optimization problem with a GA, the algorithm parameters are tuned manually in a problem-specific manner. Hence, we will first apply the GA on a test case, of which we have some idea what the optimal evacuation plan might be. The algorithm parameters are tuned, so that it converges as efficiently and accurately as possible to the near-optimal solution. The same parameters are also used for the more complex case presented in the next section.

Here, we consider a hexagon-shaped area. We assume it to be an outdoor space with six exits. All the exiting agents have entered through one door, and are thus familiar only with that one. Also, we assume this to be an unhurried situation, where people want to move out of the area, but do not sense such an urgency that

they would by themselves try other exits. Thus, we do not use Assumption 5. The interaction range of a guide is set to $r^{guide} = 5$ m.

For exiting agents $i \in N$, the initial positions \mathbf{x}_i^0 , radii r_i , masses m_i , and desired speeds v_i^{des} are fixed for all simulations. Before fixing the values, the parameters m_i , r_i , and v_i^{des} are drawn from a truncated normal distribution with a cutoff at three times the standard deviation. The mean and standard deviations are 73.5 kg and 8.0 kg, 0.255 m and 0.035 m, and 1.25 m/s and 0.3 m/s, respectively for m_i , r_i and v_i^{des} . On the other hand, for guide agents $g \in G$, we set the values of a typical male: $m_g = 80$ kg, $r_g = 0.27$ m, and $v_g^{des} = 1.15$ m/s. The reaction time is $\tau^{reac} = 0.5$ s. These values are all taken from the FDS+Evac user manual (Korhonen, 2018), and will also be used for the more complex case in the next section.

At the beginning of the evacuation, there are six groups of agents, each of them containing 25 agents, and each of them are located close to an edge of the hexagon; see Fig. 2(a). When the agents start to evacuate, they all head towards the familiar exit (Assumption 1), and form a jam in front of it; see Fig. 2(b).

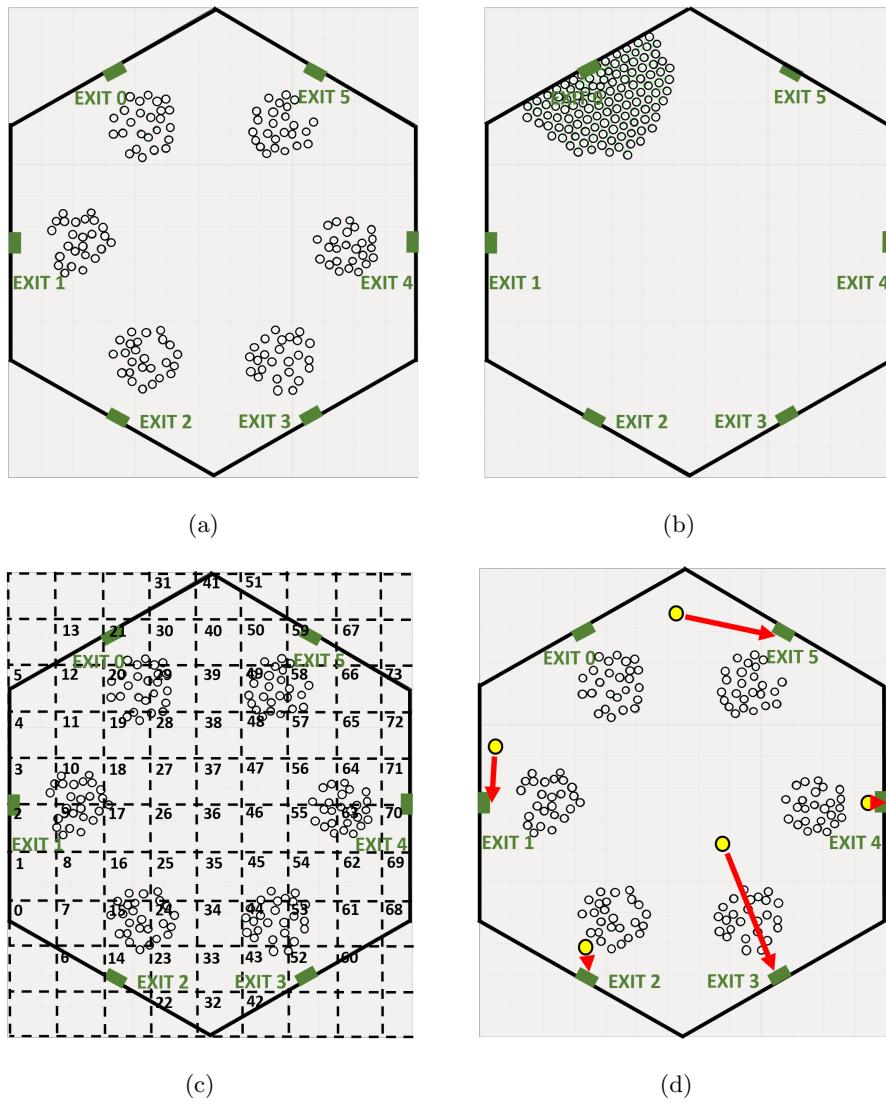


Figure 2: Evacuation of a hexagon-shape area. (a) Initial situation. (b) Without guides, a jam forms after a while. (c) Feasible starting grid cells for guides. (d) The near-optimal evacuation plan.

The guides are added to the crowd by applying the hidden genes GA. First, we discretize the space into seventy-four $3\text{ m} \times 3\text{ m}$ grid cells, which are the feasible starting grid cells for the guides; see Fig. 2(c). The six exits in the area are the feasible target exits. Because there are only six exits, and the space is highly regular, we expect that no more than six guides should be needed to obtain the optimal evacuation plan, thus a chromosome in the GA contains at most six active genes. A good guess is that the optimal solution most likely utilizes all six exits evenly. Thus, the evacuation time of the optimal solution should be close to the evacuation that utilizes evenly all six exits. Thus, the algorithm parameters are tuned so as to reach that solution as efficiently and accurately as possible. The crossover probability is set to 0.85, the mutation probability to 0.10, population size to 40, with 30 samples of each chromosome. The two worst individuals are always replaced with two best individuals of the previous generation. The algorithm is considered to have converged to a near-optimal solution, when the best solution has not changed for 15 generations.

The near-optimal solution of the GA is depicted in Fig. 2(d). It includes only five guides (depicted by the yellow circles). The guides head from their initial positions to their prespecified exits along the shortest paths depicted by the red arrows (Assumption 2). The exiting agents start to follow their closest guides, within the r^{guide} range, and do not switch to follow another guide (Assumptions 3 and 4). As it was mentioned earlier, all of the exiting agents have already a single familiar exit, so it is enough to redirect the majority of the exiting agents to the other five exits, thus utilizing all six exits in the evacuation.

One could have expected that the guides would be in symmetric positions close to their target exits. It is not needed, because the evacuation time is affected both by the walking and the queuing time of the agents. Additionally, the flow at the exit is known to be a nonlinear function of the size of the crowd in front of it (Schadschneider et al., 2008). Thus, up to a certain point, the flow can be increased by increasing the size of the crowd in front of the exit. However, after a certain point the flow will start to decrease. Due to these effects, the near-optimal solution is not very sensitive to the starting positions of the guides, as long as the exit utilization is fairly even, and each guide influences about 1/6 of the exiting agents. If the interaction range r^{guide} would be increased, the positions of the guides would have to be chosen more carefully, so that the guides do not also influence exiting agents farther away. In Fig. 3 we see how the hidden genes GA converges to the near-optimal evacuation plan. Starting from the 11th generation, the best solution does not change for 15 generations, and the GA is considered to have converged.

We also try how the solution changes with a fixed number of active guides, i.e., using the GA without hidden genes. We use a larger mutation probability 0.40 for the 30 first generations, and after that we change it to only 0.05. We are able to have a clear exploration phase in the beginning of the algorithm. It is needed for convergence accuracy, when the hidden genes are not used. Otherwise, we use the same settings as with the hidden genes GA. In Fig. 4 we see a comparison of the sample mean evacuation time for the near-optimal evacuation plans with different numbers of guides. We see that adding guides speeds up the evacuation, up to five guides. The slight increase in evacuation time with six guides is a statistical deviation rather than the sixth

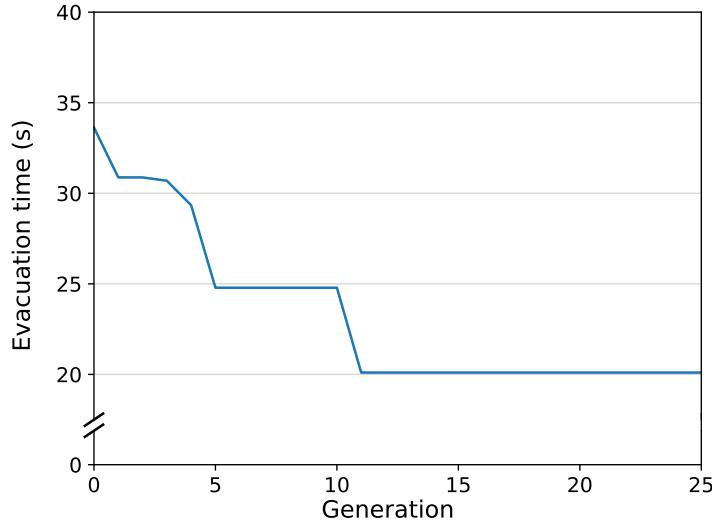


Figure 3: The sample mean evacuation time for the best solution of each generation of the GA, when applied to the test case.

guide slowing the evacuation, as we will soon see from the depictions of the evacuation plan. Notice also that the marginal utility of adding guides is decreasing, i.e., adding the first guide cuts the sample mean evacuation time almost in half, but increasing from four to five guides has only a 5 % benefit.

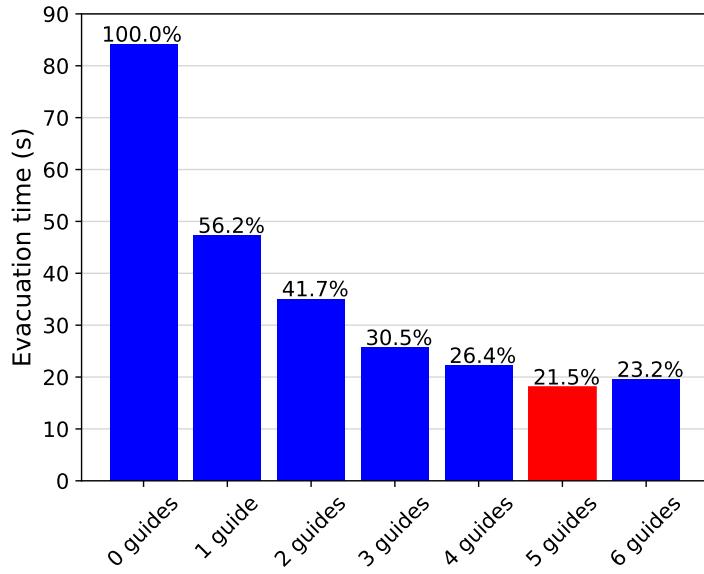


Figure 4: The sample mean evacuation time of the near-optimal evacuation plan for different fixed numbers of guides.

The resulting near-optimal solutions can be seen in Fig. 5. If there is only one guide, it leads half of the exiting agents to the exit on the opposite side of their familiar exit. If there are two guides, the crowd is split into three parts heading towards three different exits. With three guides, the crowd is split into four parts heading towards four different exits. With four guides, the crowd is split into five parts heading towards five different exits. With five guides, we get almost the same solution as with the hidden genes GA, only the starting positions of the guides are slightly altered. With six guides, the sixth guide is put aside to evacuate with the group that is heading to Exit 0. When adding the extra guide, the GA assigns a starting position and target

exit to it so that it does not interfere with the minimum time evacuation plan.

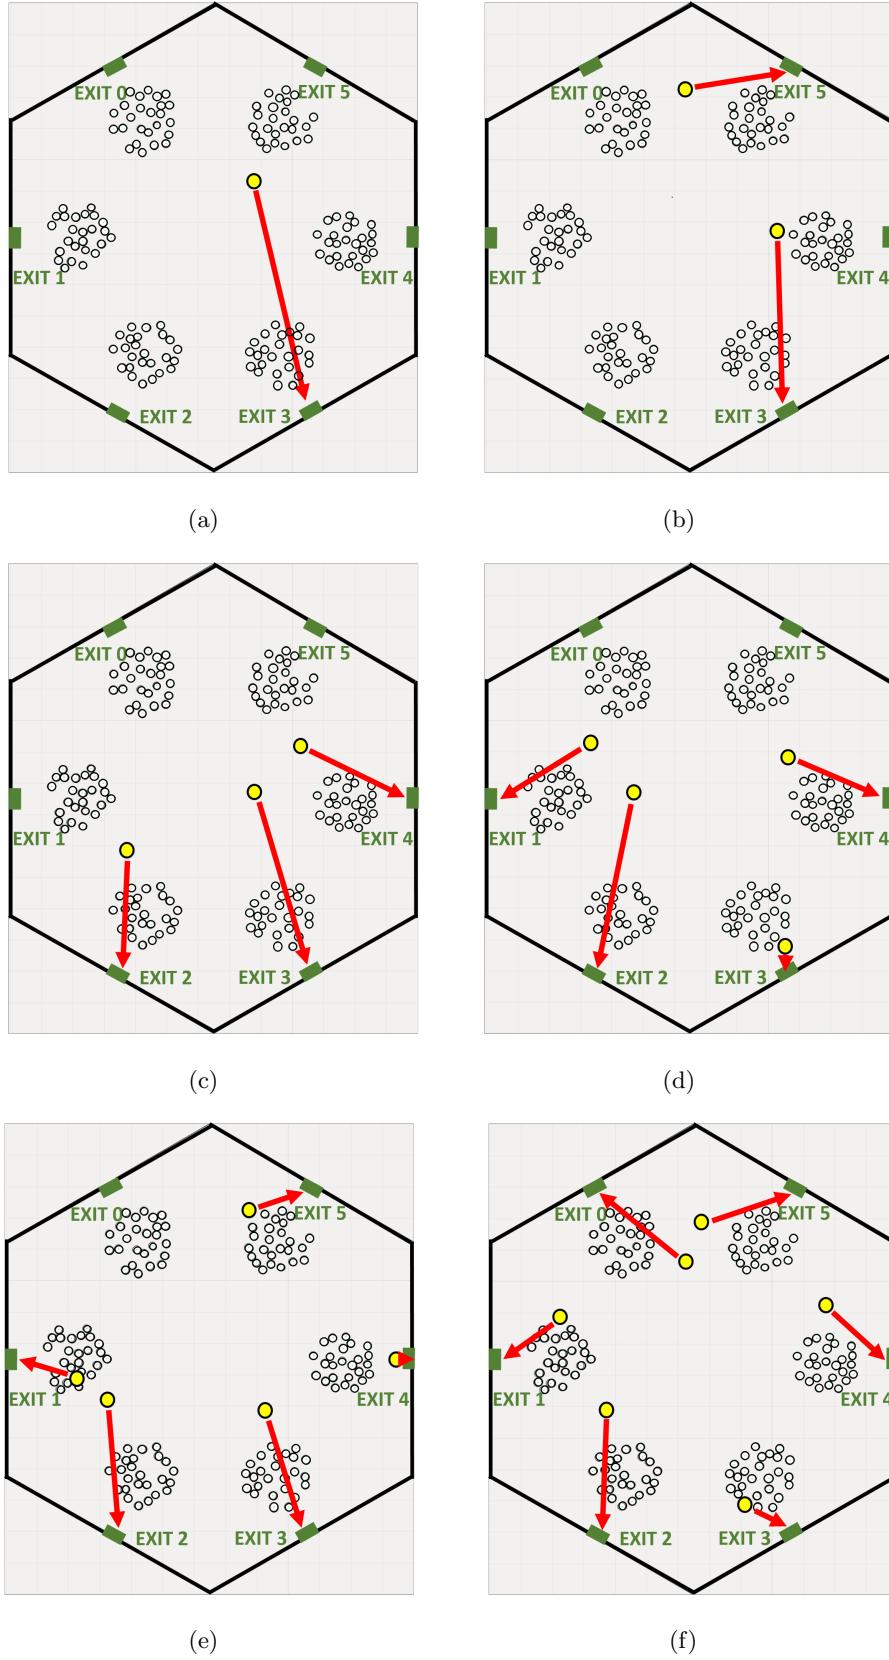


Figure 5: The near-optimal evacuation plans with (a) one guide, (b) two guides, (c) three guides, (d) four guides, (e) five guides and (f) six guides.

5. Evacuation of a conference building

We have now tuned the GA parameters suitable for the test case, and will apply it to the evacuation of a crowd from a conference building in an emergency situation. The initial situation is depicted in Fig. 6.

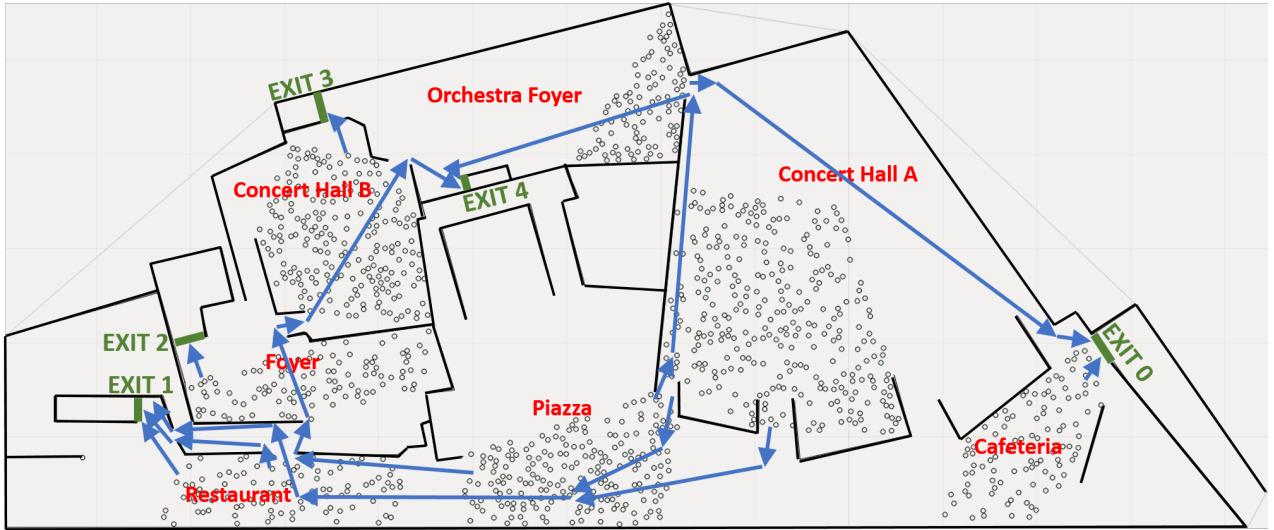
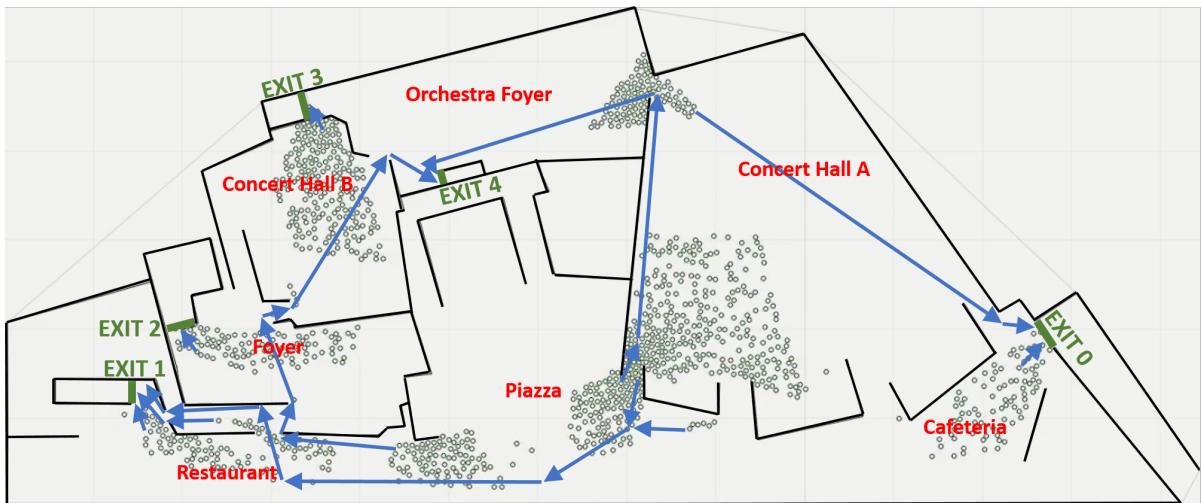


Figure 6: The initial situation of the unguided evacuation of the conference building.

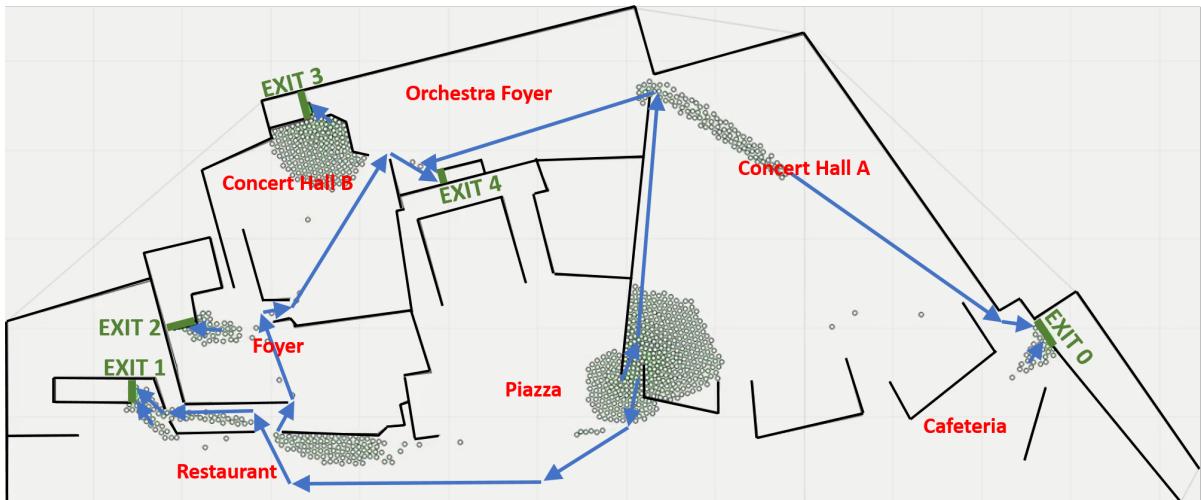
Initially, the crowd of 1100 exiting agents is split into subgroups in seven different rooms: Concert Hall A, Cafeteria, Piazza, Restaurant, Foyer, Concert Hall B and Orchestra Foyer. There are five exits in the building. All agents in a subgroup have the same familiar exit towards which they initially intend to head (depicted by the blue arrows in Fig. 6). In the evacuation model, the agents move to their target exit using their shortest path; note however that even in a small group the agents may take different (shortest) paths to their target exit. This is the case for the group of agents initially in Piazza, which are heading to Exit 4 in Orchestra Foyer. Half of the group heads first to Concert Hall A and move along its left wall to Orchestra Foyer, while the other half takes the route through Restaurant, Foyer and Concert Hall B. The development of one realization, or scenario, of the unguided evacuation is seen in Figs. 7(a) to 8(b).

Almost immediately a large jam emerges at the bottom leftmost entrance of Concert Hall A. Half of the exiting agents from Piazza try to enter Concert Hall A, through the same entrance that the exiting agents in Concert Hall A use to enter Piazza. This jam is cleared out very slowly, and it is the main source of inefficiency for the evacuation.

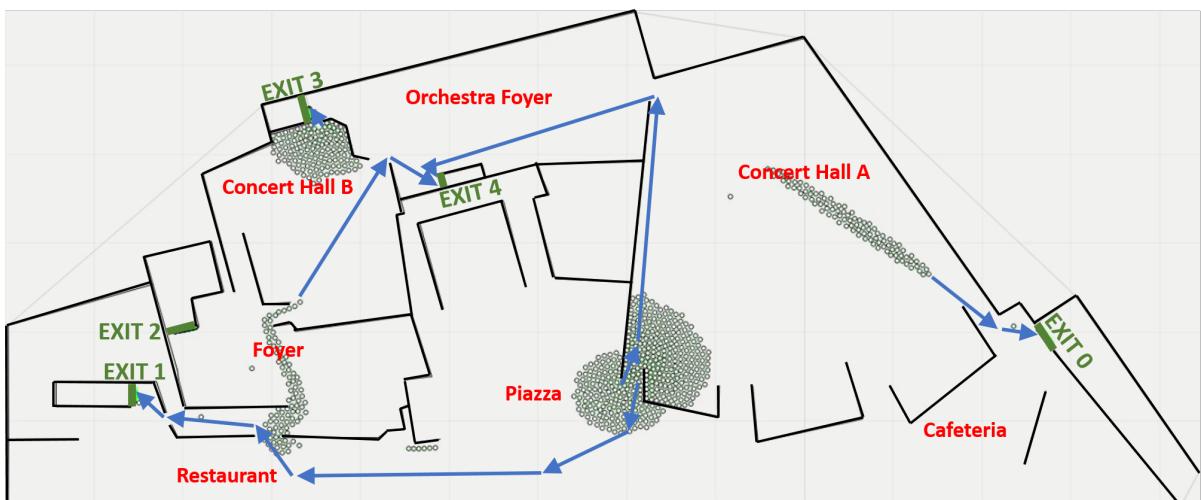
Here, all Assumptions 1-5 are used. We set $r^{guide} = 10$ m. It is larger compared to the 5 m interaction range used for the hexagon-shaped case. We could use the same value in both cases. The parameter seems to have an effect on how close the guide should be positioned to the exiting agents under its influence, and how large these groups of exiting agents are. We here set it larger so that the guide could collect more exiting agents from the large crowd. Perhaps, if it is set smaller here, more guides would be needed to have an effect on the crowd. The exit visibility range is set smaller than the interaction range of the guide, $r^{exit} = 9$ m. As a result of



(a)

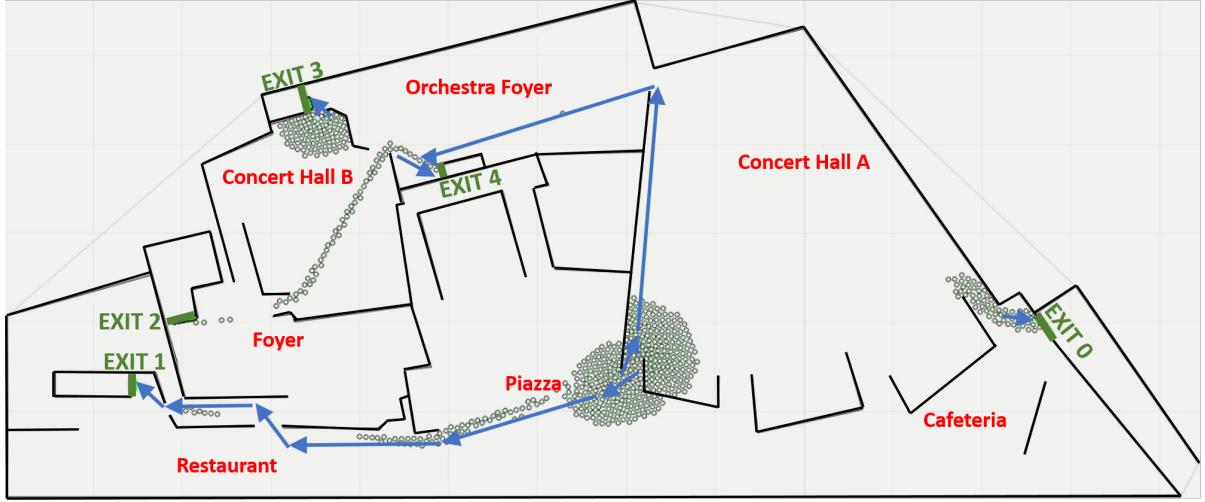


(b)

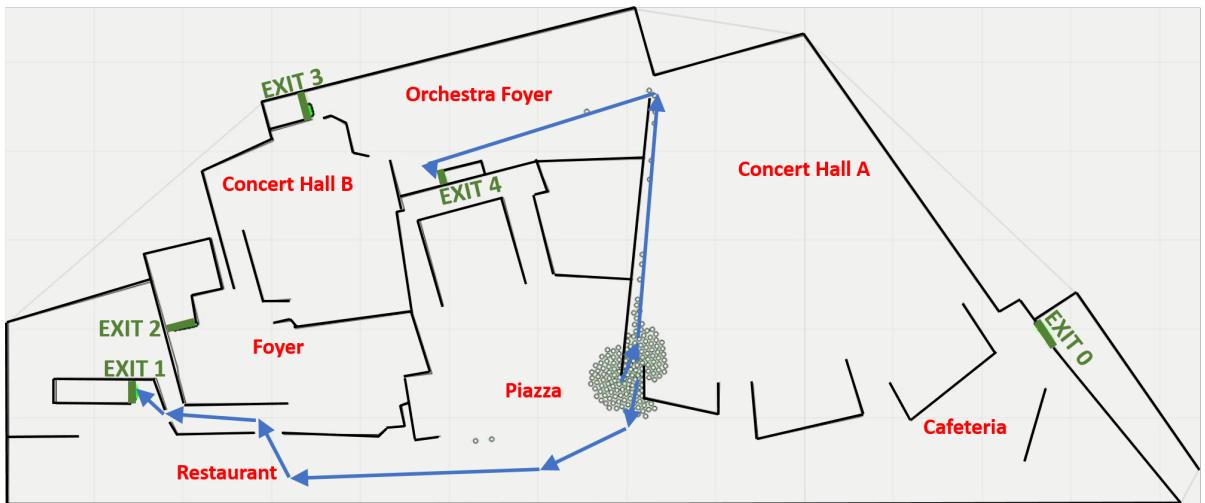


(c)

Figure 7: Snapshots of one realization of the unguided evacuation of the conference building at (a) $t = 5$ s; (b) $t = 15$ s; (c) $t = 30$ s.



(a)



(b)

Figure 8: Snapshots of one realization of the unguided evacuation of the conference building at (a) $t = 50$ s; (b) $t = 205$ s.

this, in some situations, the guides are able to influence exiting agents before they notice a nearby exit. So, the relation between the values of r^{guide} and r^{exit} can have an effect on the exiting agents' exit choice in situations where both a guide and an exit are nearby. Otherwise, the same evacuation model parameter values are used as in the hexagon-shaped case.

5.1. Numerical results

In the GA, we set the number of genes in a chromosome to 10, meaning that the solution can have a maximum of 10 guides. When the GA has converged, all solutions in the final population have less than 10 genes active, which strongly indicates that the optimal solution should have less than 10 guides. The feasible starting grid cells are obtained again by discretizing the building into sixty-two $10\text{ m} \times 10\text{ m}$ grid cells; see Fig. 9. The five exits in the conference building are the feasible target exits.

As in the test case, the crossover probability is set to 0.85, the mutation probability to 0.10, and the

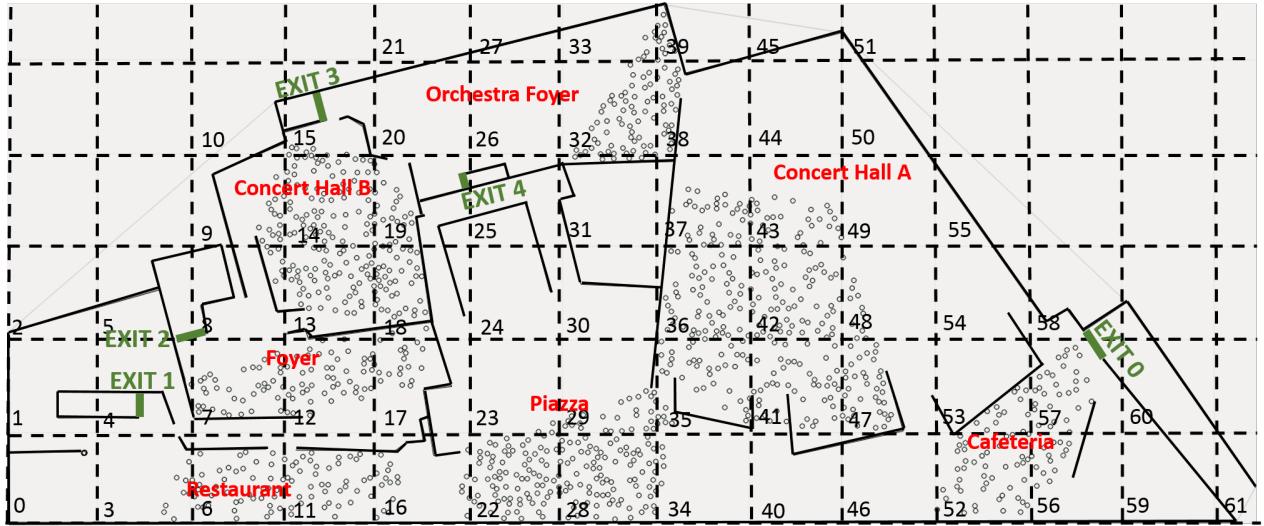


Figure 9: Feasible starting grid cells for the guides in the conference building.

population size to 40, with 30 samples of each chromosome. The two worst individuals are always replaced with the two best individuals of the previous generation. The near-optimal solution given by the hidden genes GA is illustrated in Fig. 10.

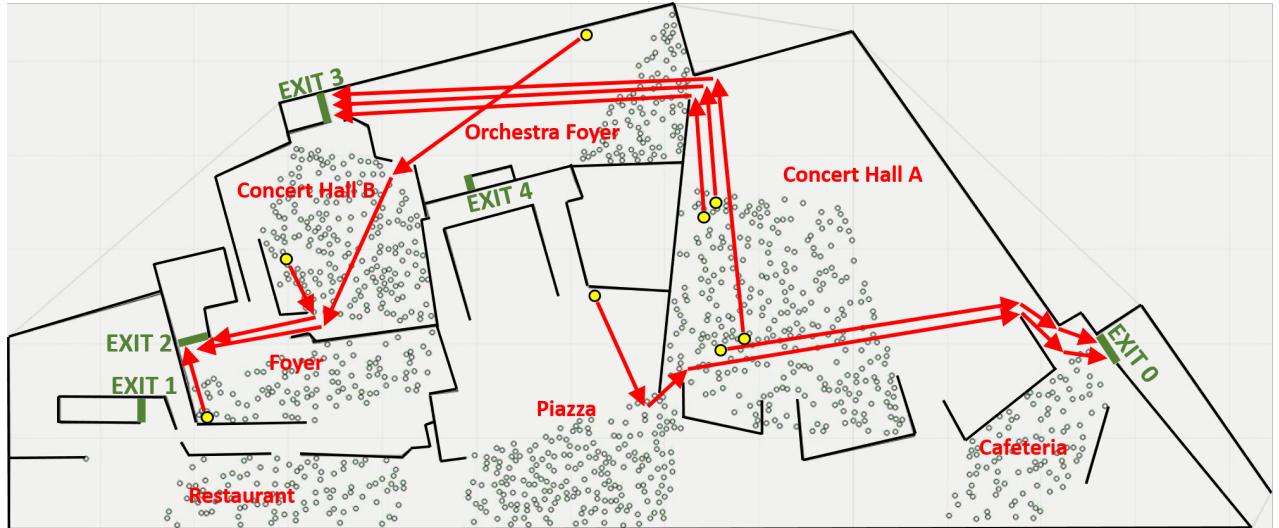


Figure 10: The near-optimal evacuation plan of the conference building.

The total number of active guides is eight (depicted by the yellow circles). They head to their target exits (Assumption 2) along the routes depicted by the red arrows. The exiting agents start to follow their closest guides, within the r^{guide} range, and do not switch to follow another guide (Assumptions 3 and 4). One guide is set in the Orchestra Foyer to head towards Exit 2. Four guides are in Concert Hall A, to herd the large crowd. Out of those four, three are heading towards Exit 3, and one towards Exit 0. In the upper part of Piazza, there is a guide, which leads half of the exiting agents in Piazza to Exit 0. The other half of the exiting agents in Piazza are initially free to move towards Exit 4 using the route through Restaurant, Foyer and Concert Hall B.

The exiting agents located in Cafeteria, Restaurant and Foyer move as they would without guidance. There is also one guide in Foyer and another in Concert Hall B that lead half of the exiting agents from Concert Hall B to Exit 2. The other half of the exiting agents in Concert hall B go to Exit 3 without guidance. Finally, there is one guide located in the Orchestra Foyer, and it is heading towards Exit 2.

The convergence pattern of the GA can be seen in Fig. 11. Starting from the 23rd generation, the best solution does not change for 15 generations, and the GA is considered to have converged.

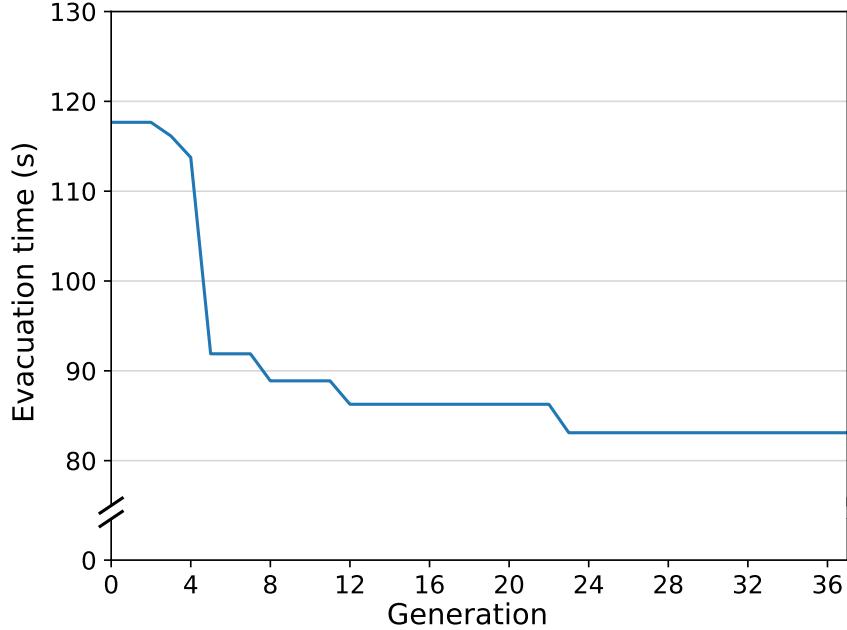


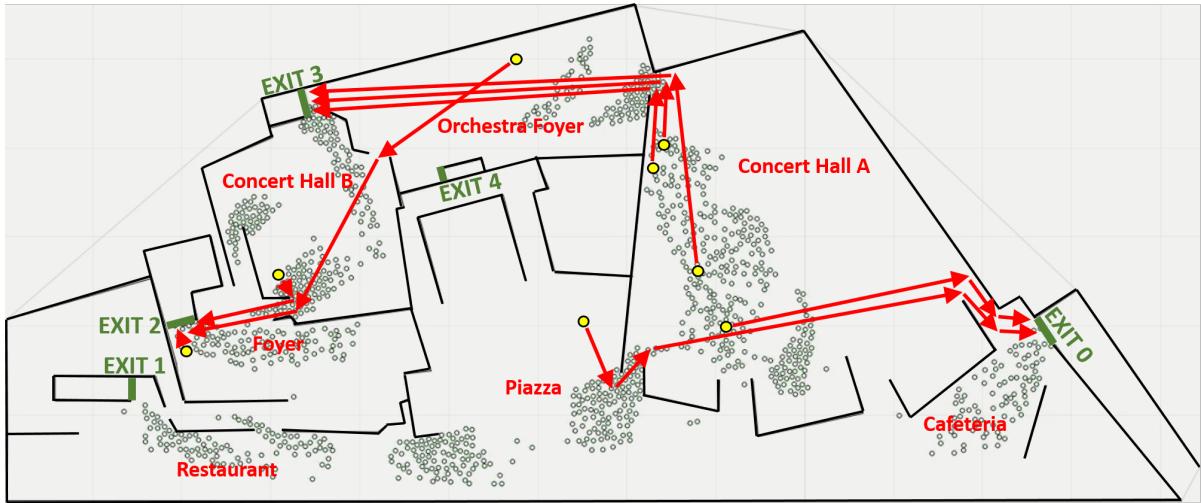
Figure 11: The sample mean evacuation time of the best solution of each generation of the GA, when applied to the conference building case.

In the unguided situation, a large jam forms at the bottom leftmost entrance to Concert Hall A. It is evident that it should be solved to increase efficiency. However, the question is where to redirect this part of the crowd, as not to create a jam somewhere else. In Figs. 12(a) to 13(b) the full evolution of one realization of the near-optimal guided evacuation is illustrated with snapshots.

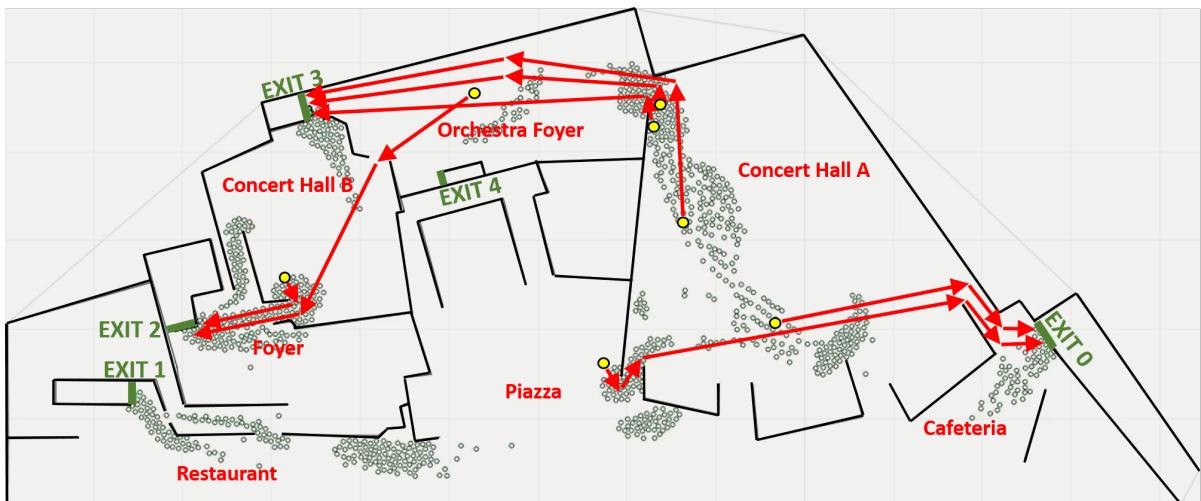
There are many changes to the unguided evacuation. First and foremost, a big jam never emerges at the bottom leftmost entrance of Concert Hall A. This is because the crowd in Concert Hall A is taken to Exit 3 through Orchestra Foyer, and the rest to Exit 0 via the rightmost door of Concert Hall A. Only a couple of exiting agents from Concert Hall A are left out without a guide and move by themselves to their familiar exit, Exit 1.

A large part of the crowd is heading towards Exit 3, and no guide is going towards Exit 4. However, guiding exiting agents to Exit 4 would be unnecessary, as about half of the exiting agents initially heading to Exit 3, once they are in the Orchestra Foyer, detect Exit 4 by themselves, and head there instead of Exit 3 (Assumption 5).

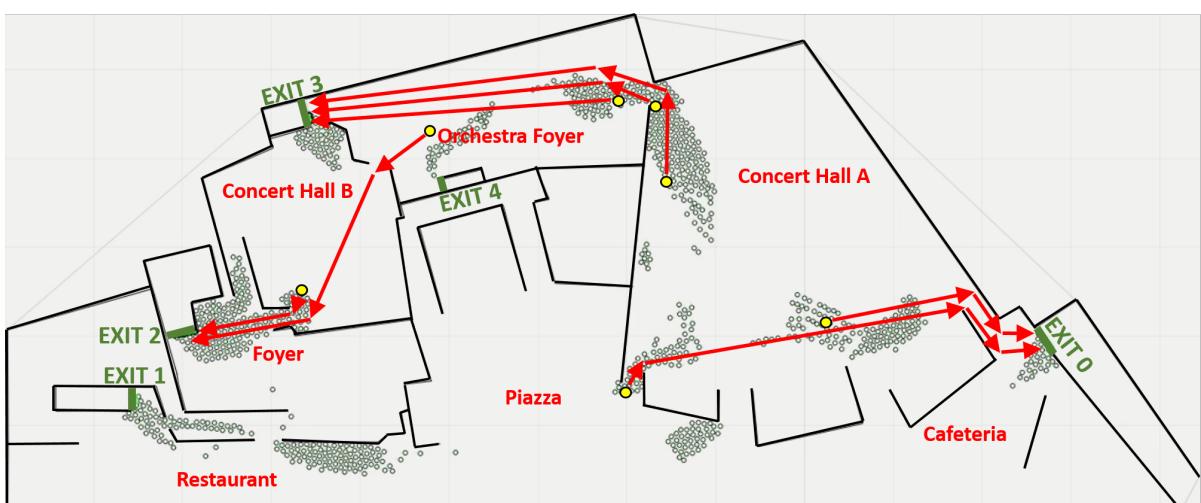
Now without guidance, the exiting agents from Piazza would be heading to Exit 4, which would cause too



(a)

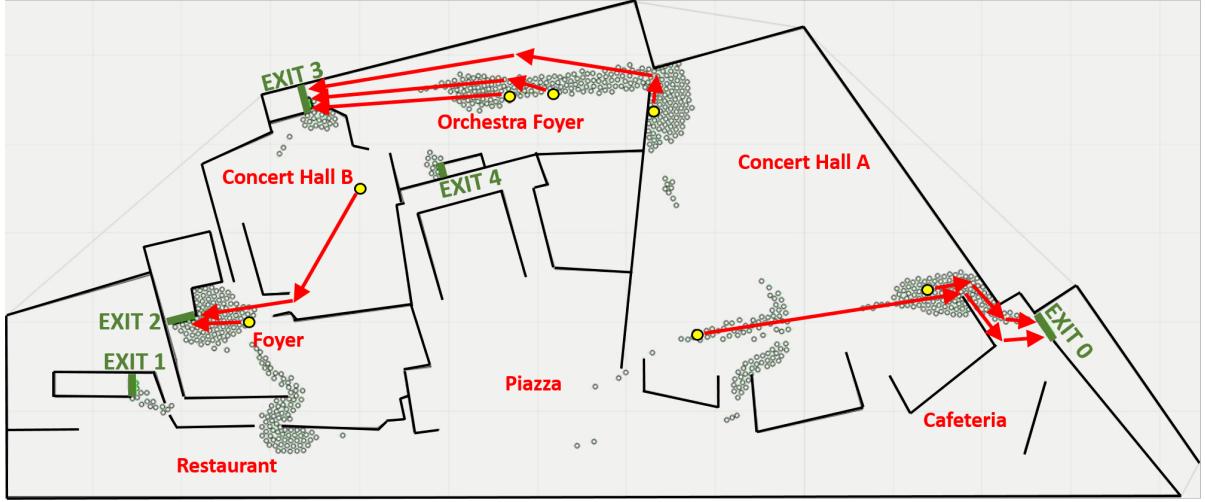


(b)

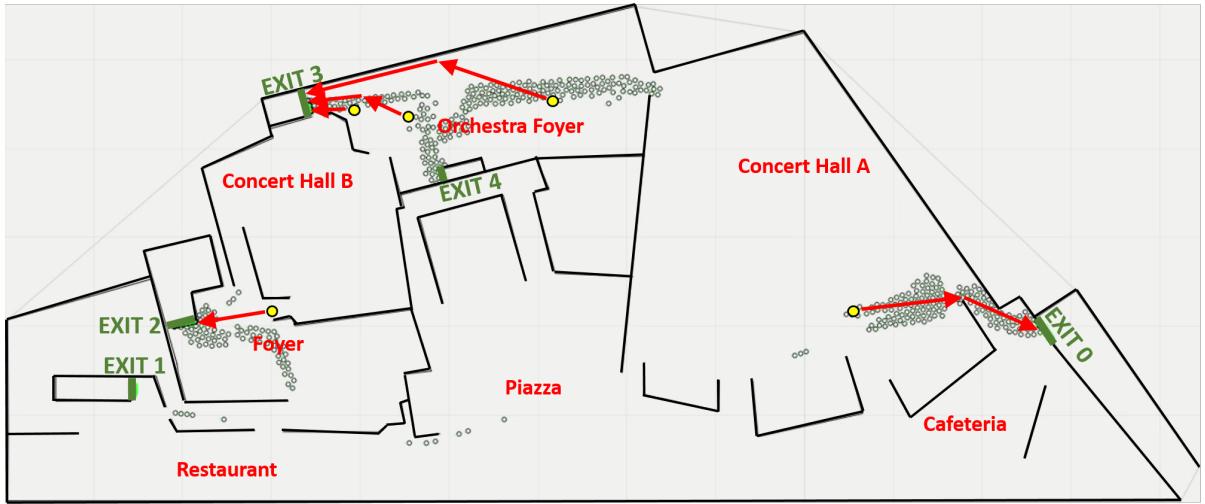


(c)

Figure 12: Snapshots of one realization of the near-optimal guided evacuation of the conference building at (a) $t = 5$ s; (b) $t = 10$ s; (c) $t = 15$ s.



(a)



(b)

Figure 13: Snapshots of one realization of the near-optimal guided evacuation of the conference building at (a) $t = 25$ s; (b) $t = 40$ s.

much traffic there. However, in the near-optimal solution, a guide takes half of the exiting agents from Piazza to Exit 0 through Concert Hall A. The second half of the exiting agents in Piazza that head to Exit 4 through Restaurant, Foyer and Concert Hall B are stopped at the Foyer and redirected to Exit 2.

In a more complicated model, the guide could be set to wait in place for exiting agents crossing paths with it later on. Here, the optimization model circumvents this lack of feature, by setting guides to start moving from farther away, to cross paths with exiting agents just at the right time. This is illustrated by the guide from the upper right corner of Orchestra Foyer, that crosses paths with the exiting agents from Piazza at Foyer, and redirects them to Exit 2. To conclude, the main feature of a good guidance seems to be that all exits are utilized, while none of them are overutilized. Also, jams or counterflows should not either emerge at other parts of the building.

Fig. 14 shows a histogram and a kernel density estimate of the evacuation time of both the unguided

and near-optimal evacuation. For the unguided evacuation, the sample mean is 320.72 s, whereas the sample standard deviation is 20.75 s. For the near-optimal evacuation, the respective values are 83.11 s and 3.23 s. The improvement is tremendous, as the quantities are only about 25 % and 15 %, respectively, of those values of the unguided evacuation. The improvement in standard deviation probably is due to the near-optimal guidance solving the large jam witnessed in the unguided evacuation. The nonlinear physical forces come into play in close contact of agents. When a large crowd is jammed at a bottleneck, the agents in it are in close contact for a prolonged time. Hence, even small movement fluctuations can affect the evacuation time. This has also been experimentally verified in (Helbing et al., 2003; Huang & Guo, 2008). In Appendix B, we have further analyzed the effect of stochasticity on the problem by using metrics from the literature of stochastic programming (Birge, 1982).

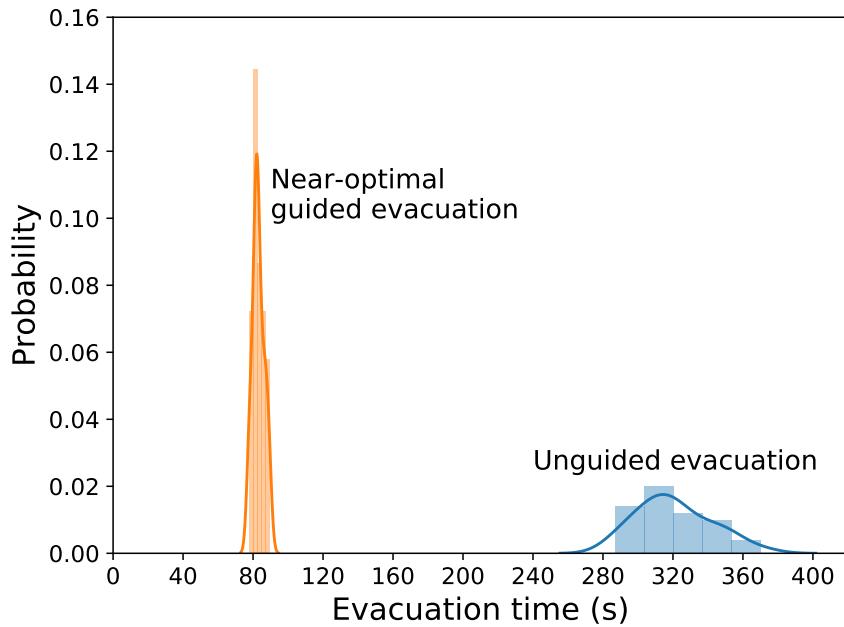


Figure 14: Histograms (staple diagrams) and kernel density estimates (solid lines) of the evacuation times of 30 scenarios.

Also, from another point of view, even though the initial positions of the agents are fixed for all simulations, as time goes on, due to the random force term, the positions will differ between scenarios. Thus, in some sense, the near-optimal evacuation plan is robust to reasonable variations in agents' positions.

6. Behavioral considerations

To emphasize, what we have in mind with our model, or the basic question we ask is this: what will happen in a complex building containing a large crowd, after the sound of a whistle or a fire alarm in a serious situation, when there is enough trained yellow-coated rescue personnel, whose members have enough skill and authority to guide the crowd out of the building as soon as possible. So, our main modeling question here is how to coordinate the guides so that they manage with their task in an optimal way.

A typical concern in modeling crowd evacuation, is that the crude assumption of focusing only on physical movement may potentially lead to an underestimation of the time for the crowd to reach safety, since evacuees are likely to engage in a variety of activities during the operation (Gwynne et al., 2016). We believe that in our situation guided by trained personnel, such evacuation activities, although can be important, are very rare, since people in serious situations are willing to follow authorities (Proulx, 2002; Nishinari et al., 2004).

Even more to the point, it does not matter if some individuals, at some time, would have turned to the right, or stopped for a while, since now guides more or less control their movement. In fact, even for a large group of people, a few guides are enough to control the whole crowd to their target (Dyer et al., 2009; Cao et al., 2016). For a more irregular building geometry, where the crowd is more scattered, more guides are needed. In that case, the coordination between guides becomes very important to facilitate the evacuation. A useful quantitative information that a model should deliver in this case is a suitable macroscopic one, such as the evacuation time of the crowd. It is in fact less prone to globally unnecessary details and fluctuations than microscopic information (Helbing & Johansson, 2013). Our paper solves the number of guides, their initial positions and exit assignments to minimize the crowd evacuation time in a single although complicated optimization problem.

Nevertheless, we acknowledge the importance of taking into account behavioral aspects in crowd evacuation modeling. It has been the focus of our previous research (Heliövaara et al., 2012b, 2013; von Schantz & Ehtamo, 2015, 2019). For example, in (von Schantz & Ehtamo, 2019) we model crowd behavior in an exit congestion. We couple a local decision-making model to the social force model. The decision-making model is based on behavioral assumptions that are verified in our experiment with real humans (Heliövaara et al., 2012b). With our integrated treatment of behavioral and physical aspects, we are able to simulate when, why and how typical phenomena of an evacuation through a bottleneck occur. Most importantly, we attain non-monotonous speed and kinetic pressure patterns near exits in threatening situations. These kind of behavioral phenomena are interesting as such. However, to include them into the optimization would require us to increase the already very heavy computation.

Instead of studying a single agent deviating from its usual movement, it would be interesting to study the possibility of a larger group doing so. For example, it could happen like this: an agent decides to go to another exit and a large 20% part of the crowd decides to follow this agent. However, to add this behavior to a crowd movement model takes careful calibration. It is not a simple task to alter the microscopic dynamics to generate this effect. Moreover, a larger amount of Monte Carlo simulations would probably be needed to estimate the mean of the evacuation time for a given evacuation plan. Perhaps, an easier way than for the model to generate this effect, is to construct representative scenarios. They would describe all drastically different alternative ways the evacuation could develop. In one of the scenarios, 20% of the crowd would go to another exit as opposed to where it goes in the other scenarios. Then, the optimal evacuation plan would be calculated over all these scenarios.

7. Implementation details and performance

Let us first start by discussing the implementation details of the evacuation model. In it, the shortest paths that agents use to move to the exits were calculated using the detailed method presented in depth in (Kretz et al., 2011). The method involves solving the continuous shortest path to an exit. It is solved using the fast marching method (Sethian, 1999), which first discretizes the continuous space into a meshgrid. Then the method works almost like Dijkstra's algorithm for finding shortest paths between nodes in a graph. The solution is a distance map from each point, or grid, in the building to the exit. The direction to which an agent should head towards at each point can be calculated using the gradient direction. The distance maps, i.e., the distances from each grid to every exit are calculated and stored before the optimization simulations.

Interaction forces in the social force equations should be calculated between all agents, but to speed up the computation, we subdivided the simulation domain with a blocklist algorithm with a cut-off distance of 3 m (Yao et al., 2004). This should not affect the crowd dynamics, since the social force term is exponentially decaying.

When using numerical simulations to calculate the sample mean evacuation time, we use a pseudorandom number generator to generate realizations of the random force terms. For each scenario, we store the seed of the pseudorandom number generator. Thus, the realizations of the random force terms are replicable. Then, for a specific seed, the social force equations are solved with the Velocity Verlet numerical integration scheme with the time step $\Delta t = 0.01$ s; see, e.g., appendix of (von Schantz & Ehtamo, 2019) for more details.

The evacuation model is implemented in Python code. Some of the core parts of the code are written as Numba-decorated functions, which translates Python functions to optimized machine code at runtime. Numba-compiled numerical algorithms in Python can approach the speeds of C or FORTRAN (Oliphant et al., 2020). Whereas, the GA is implemented in Bash script that calls the crowd simulation scripts written in Python. For reproducibility, all codes are published (von Schantz, 2020).

We ran the numerical experiments on the Aalto University high-performance computing cluster Triton. A single generation of the GA was run in parallel on Triton using its computing nodes that are Intel Xeon X5650 2.67 GHz with 48 GB or 96 GB memory, and Xeon E5 2680 v2 2.80 GHz with 64 GB or 256 GB memory. One generation in the GA requires 1200 simulations, because a generation contains 40 solutions, each which have to be run for all 30 scenarios. Due to the user quotas set for Triton users, we only ran 300 simulations in parallel, which means that one generation had to be run in four parts. A single simulation of the evacuation of the hexagon-shaped space could take up to 5 min, while one simulation of the evacuation of the conference could take up to 1 h. Thus, the simulation of one generation took about 20 min and 4 h, respectively. We ran the GA for 26 and 38 generations, for the test case and the conference building case, respectively. Thus, for the test case, the GA converged to the near-optimal solution in 8 h 40 min, whereas, for the conference building case it converged in 152 h, or 6 days 8 h.

8. Conclusion

We have studied the problem of minimizing the evacuation time of a crowd from a complex building using rescue guides. The problem is first formulated as a stochastic control problem, where the objective is the expected evacuation time, and the optimization variables are the number of guides and their routes defined as origin-destination pairs. The system equations are the equations of motion, given by the social force model, and the rules of interaction for exiting agents and guides. The problem is then reformulated as a scenario optimization problem, which we solve with a hybrid numerical simulation and GA approach. With it, we are able to solve the number of guides, their initial positions and exit assignments to minimize the crowd evacuation time in a single optimization.

Typically, to solve an optimization problem with a GA, the algorithm parameters are tuned manually in a problem-specific manner. In our study, we do this by constructing a test case, the hexagon-shaped space in Sec. 4, for which we know beforehand what the expected minimum evacuation time should be. Then we extensively try different parameter configurations to get the GA to converge as efficiently and accurately as possible to the near-optimal solution. This GA parameter configuration is also used for solving the near-optimal evacuation plan for the conference building.

Moreover, it is known that using GAs for problems with large solution space and large number of local minima, introducing noise to the fitness function, and evaluating it by taking multiple samples, improves convergence to the global minimum (Hammel & Bäck, 1994). Our problem is inherently noisy, and we perform multiple Monte Carlo simulations to calculate the expected evacuation time for a given evacuation plan. Thus, this may assure that our GA does not get stuck in a local optimum.

In both cases studied in this paper, the improvement with the near-optimal plan is dramatic, the sample mean evacuation time is only about 25 % of that of the unguided evacuation. Moreover, for the conference building case, we also compared standard deviations, and it was only about 15 % of that of the unguided evacuation. This is probably due to the near-optimal evacuation plan solving all major jams, which decreases the nonlinear physical effects, and hence small deviations in agents' movement does not affect the evacuation time so much.

It is interesting to see how the near-optimal solutions take into account physical phenomena like counterflows, jam formation and flow at bottlenecks being a nonlinear function of crowd size. There are computationally faster evacuation models, like cellular automata and network flow models that are good for planning large-scale evacuations (Løvås, 1998; Abdelghany et al., 2014). However, they cannot model the physics of crowds evacuating from complex confined spaces.

However, it should also be noted that the optimization procedure is computationally very slow. So, as such, this procedure could not be used for online optimization. Faster computation could be achieved by using an implicit integration scheme, where the step size can be set in some cases even 40 times larger than with an explicit

integration scheme (Karamouzas et al., 2017). Alternatively, a position-based dynamics approach could be used, which has been shown to produce real-time simulations (Weiss et al., 2019). Another interesting avenue for future research would be to use neural networks to deal with the large state space (Bertsekas & Tsitsiklis, 1996).

In future research, other objective functions could be used as well, and studied how the optimal evacuation plans differ from that of ours. An evacuation plan should be both fast and safe (Løvås, 1995). If we want to simultaneously take these two objectives into account, we can solve the problem with multi-objective optimization (Saadatseresht et al., 2009). As noted before, one form of risk are rare events that dramatically slow down the evacuation. To account for this, we cannot only minimize the mean evacuation time, but we also have to minimize the variance or some other risk-related measure related to the evacuation time distribution. We could also include physical risks like pressures in the crowd or avoidance of dangerous areas in a building. On the other hand, if we are considering an unhurried evacuation, objectives related to the quality of service can be used, e.g., minimization of the average evacuation time or distance travelled by the evacuees, or time spent in congestions.

Acknowledgements

This study was first funded by a grant from the Foundation for Aalto University Science and Technology, and later with a grant from the Finnish Science Foundation for Technology and Economics. The calculations in this study were performed using computer resources within the Aalto University School of Science "Science-IT" project. We wish to thank our summer assistant Jaan Tollander de Balsch, who considerably helped us develop the simulation codes.

Declaration of interests

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A. Social force model parameters

The social force model parameters have been validated against data, and the collective phenomena observed with the model are robust to reasonable parameter variations (Karamouzas et al., 2014; Korhonen, 2018). In our study, the initial positions \mathbf{x}_i^0 , radii r_i , masses m_i , and desired speeds v_i^{des} , for exiting agents $i \in N$, are fixed for all simulations. Before fixing them, the parameters m_i , r_i , and v_i^{des} are drawn from a truncated normal distribution with a cutoff at three times the standard deviation. The mean and standard deviation are 73.5 kg and 8.0 kg, 0.255 m and 0.035 m, and 1.25 m/s and 0.3 m/s, respectively for m_i , r_i and v_i^{des} . Whereas, for guide agents $g \in G$, we set $m_g = 80$ kg, $r_g = 0.27$ m, and $v_g^{des} = 1.15$ m/s. The reaction time is $\tau^{reac} = 0.5$ s.

Next, we describe the terms of the social force model. Our derivation of the social force term $\mathbf{f}_{ij}^{soc}, i, j \in N \cup G, i \neq j$, follows closely that of (Karamouzas et al., 2014). There, it was inferred from a large public data set that the interaction energy associated with the repulsive social force follows a power law with a sharp truncation at large τ ,

$$E(\tau) = \frac{k_{soc}}{\tau^2} e^{-\tau/\tau_0}. \quad (\text{A.1})$$

Here, k_{soc} is a constant that sets the units of energy. It is not that clearly documented, but by examining the codes provided in the supplemental material of (Karamouzas et al., 2014), we deduce its value to be $k_{soc} = 1.5m_i$. The parameter τ_0 is the interaction time horizon, and it is set to 3 s. The collision time of two agents i and j is denoted with τ . So, Eq. (A.1) defines the interaction energy of a pair of agents, which are on a collision course. This energy is directly related to the social force \mathbf{f}_{ij}^{soc} experienced by agent i due to the interaction with another pedestrian j . In particular,

$$\mathbf{f}_{ij}^{soc} = -\nabla_{\mathbf{x}_{ij}} E(\tau). \quad (\text{A.2})$$

At any given simulation step, we estimate τ by linearly extrapolating the trajectories of the pedestrians i and j based on their current velocities, \mathbf{v}_i and \mathbf{v}_j , and position vectors, \mathbf{x}_i and \mathbf{x}_j . Specifically, a collision is said to occur at some time τ , if the corresponding circles of the pedestrians of radii r_i and r_j intersect. If no such time exists, the interaction force \mathbf{f}_{ij}^{soc} is a zero vector $\mathbf{0}$. Otherwise, $\tau = \frac{b - \sqrt{d}}{a}$, where $a = \|\mathbf{v}_{ij}\|^2$, $b = \mathbf{x}_{ij} \cdot \mathbf{v}_{ij}$, $c = \|\mathbf{x}_{ij}\|^2 - (r_i + r_j)^2$, and $d = b^2 - ac$. Here, $\mathbf{x}_{ij} = \mathbf{x}_i - \mathbf{x}_j$ is their relative displacement, and $\mathbf{v}_{ij} = \mathbf{v}_i - \mathbf{v}_j$ is their relative velocity. By substituting τ into Eq. (A.2), the interaction force can be written explicitly as:

$$\mathbf{f}_{ij}^{soc} = - \left[\frac{k_{soc} e^{\tau/\tau_0}}{\|\mathbf{v}_{ij}\|^2 \tau^2} \left(\frac{2}{\tau} + \frac{1}{\tau_0} \right) \right] \left[\mathbf{v}_{ij} - \frac{\|\mathbf{v}_{ij}\|^2 \mathbf{x}_{ij} - (\mathbf{x}_{ij} \cdot \mathbf{v}_{ij}) \mathbf{v}_{ij}}{\sqrt{(\mathbf{x}_{ij} \cdot \mathbf{v}_{ij})^2 - \|\mathbf{v}_{ij}\|^2 (\|\mathbf{x}_{ij}\|^2 - (r_{ij})^2)}} \right]. \quad (\text{A.3})$$

Here, $r_{ij} = r_i + r_j$. Furthermore, the force \mathbf{f}_{ij}^{soc} is truncated from above to 2000 N.

The original social force model (Helbing & Molnár, 1995) includes also a repulsive social force between agents and walls, \mathbf{f}_{iw}^{soc} . Because, even though the desired velocity \mathbf{v}_i^0 does not point inside walls by construction, the actual velocity \mathbf{v}_i could do so, i.e., an agent could be pushed inside a wall by other agents in the crowd. To avoid this, we use the approach from (Cristiani & Peri, 2017). In it, the desired velocity \mathbf{v}_i^0 is constructed to heavily point away from a wall, when the agent gets close to it.

Physical contact forces come into play, when agents i and j touch each other, $r_{ij} - \|\mathbf{x}_{ij}\| \geq 0$:

$$\mathbf{f}_{ij}^c = k(r_{ij} - \|\mathbf{x}_{ij}\|) \mathbf{n}_{ij} + c_d \Delta v_{ji}^n \mathbf{n}_{ij} + \kappa(r_{ij} - \|\mathbf{x}_{ij}\|) \Delta v_{ji}^t \mathbf{t}_{ij}, \quad (\text{A.4})$$

where $\mathbf{n}_{ij} = (n_{ij}^1, n_{ij}^2) = \mathbf{x}_{ij}/\|\mathbf{x}_{ij}\|$ is the normalized vector pointing from agent j to i , $\mathbf{t}_{ij} = (-n_{ij}^2, n_{ij}^1)$ is the tangential direction, $\Delta v_{ji}^n = -\mathbf{v}_{ij} \cdot \mathbf{n}_{ij}$ is the normal velocity difference, and $\Delta v_{ji}^t = -\mathbf{v}_{ij} \cdot \mathbf{t}_{ij}$ is the tangential velocity difference. The parameters $k = 1.2 \cdot 10^5$ kg/s, $c_d = 500$ kg/s, and $\kappa = 4.4 \cdot 10^4$ kg/(m·s) are force constants. The first term in Eq. (A.4) represents a 'body force' counteracting body compression, the second term a 'damping force', which reflects the fact that the collision of two humans is not an elastic one, and the

third term is a 'sliding friction force' impeding relative tangential motion. The physical agent-wall interaction force \mathbf{f}_{iw}^c is treated similarly and same parameter values are used.

Finally, the components of the random force vector $\boldsymbol{\xi}_i$ follow a truncated normal distribution with mean zero, standard deviation of $0.1m_i$ m/s², and are truncated at three times of the standard deviation.

Appendix B. The value of the stochastic solution

Let us evaluate the effect of stochasticity on our optimization problem using terminology from the literature of stochastic programming (Birge, 1982). First, recall the mathematical definitions from Secs. 2 and 3. Then, let us denote $\boldsymbol{\xi}_i(t), 1 \leq i \leq n, t \in [0, T_{last}]$, with $\boldsymbol{\delta}(t)$. And, because T_{last} depends on $\omega_g, \varepsilon_g, g \in G$ and $\boldsymbol{\delta}$ through the constraint equations, we define:

$$\phi(\omega_g, \varepsilon_g, g \in G; \boldsymbol{\delta}) := T_{last}. \quad (\text{B.1})$$

The problem of Eq. (3) is called the *recourse problem*, and using Eq. (B.1) it can be rewritten:

$$\begin{aligned} \text{RP} := \min_{(\omega_g, \varepsilon_g)} & \mathbb{E}[\phi(\omega_g, \varepsilon_g, g \in G; \boldsymbol{\delta})]; \\ & \omega_g \subset \bar{\Omega}, \varepsilon_g \subset \mathcal{E}, g \in G, \end{aligned} \quad (\text{B.2})$$

subject to the constraints defined in Eq. (3). We obtained the sample mean 83.11 s for RP, in the conference building case, using the detailed solution methodology presented in the paper.

Next, we define the deterministic equivalent problem, or the *expected value problem*:

$$\begin{aligned} \text{EV} := \min_{(\omega_g, \varepsilon_g)} & \phi(\omega_g, \varepsilon_g, g \in G; \mathbb{E}[\boldsymbol{\delta}]); \\ & \omega_g \subset \bar{\Omega}, \varepsilon_g \subset \mathcal{E}, g \in G, \end{aligned} \quad (\text{B.3})$$

subject to the constraints defined in Eq. (3). Recall that the random force terms $\boldsymbol{\xi}_i(t), 1 \leq i \leq n, t \in [0, T_{last}]$, are independent and have zero $\mathbf{0}$ mean. Hence, $\mathbb{E}[\boldsymbol{\delta}] = \mathbf{0}$. Furthermore, in the equation of motion Eq. (1), the random force appears as an additive linear term. Thus, the term disappears and we are left with the deterministic social force equation. The deterministic equivalent problem is also solved with the solution methodology presented in the paper. The EV solution $(\hat{\omega}_g, \hat{\varepsilon}_g), g \in G$ is seemingly similar to the RP solution. Only the initial position of one of the guides is moved from Concert Hall A to Orchestra Foyer.

The expected evacuation time for solution $(\hat{\omega}_g, \hat{\varepsilon}_g), g \in G$ is:

$$\text{EEV} := \mathbb{E}[\phi(\hat{\omega}_g, \hat{\varepsilon}_g, g \in G; \boldsymbol{\delta})]. \quad (\text{B.4})$$

To calculate it, we perform Monte Carlo simulations. We obtain the sample mean of 85.59 s for EEV. Thus, the *value of the stochastic solution* is:

$$\text{VSS} = \text{EEV} - \text{RP} = 83.11 \text{ s} - 85.59 \text{ s} = -2.48 \text{ s}. \quad (\text{B.5})$$

So, the sample mean of the evacuation time for the EV solution is close to that of the RP solution. Still, this is not the only side to the matter. We experienced difficulties in getting the GA to converge to the optimum of the deterministic equivalent problem. However, because we had solved the stochastic problem first, we knew what we were seeking. Thus, in the end, we made use of the stochastic problem solution, and performed an exhaustive local search to find the deterministic equivalent problem solution. This convergence issue seems to be a more general concern in solving stochastic problems with their deterministic equivalent using a GA (Hammel & Bäck, 1994). For example, it might be that for a deterministic optimization problem having many local optima, but only one global optimum, when adding stochasticity or averaging over many such deterministic problems makes the problem landscape smoother with only a single optimum.

References

Abdelghany, A., Abdelghany, K., Mahmassani, H., & Alhalabi, W. (2014). Modeling framework for optimal evacuation of large-scale crowded pedestrian facilities. *European Journal of Operational Research*, 237, 1105–1118. doi:<https://doi.org/10.1016/j.ejor.2014.02.054>.

Abdelkhalik, O. (2013). Hidden genes genetic optimization for variable-size design space problems. *Journal of Optimization Theory and Applications*, 156, 450–468. doi:<https://doi.org/10.1007/s10957-012-0122-6>.

Albi, G., Bongini, M., Cristiani, E., & Kalise, D. (2016). Invisible control of self-organizing agents leaving unknown environments. *SIAM Journal on Applied Mathematics*, 76, 1683–1710. doi:<https://doi.org/10.1137/15M1017016>.

Aubé, F., & Shield, R. (2004). Modeling the effect of leadership on crowd flow dynamics. In P. M. A. Sloot, B. Chopard, & A. G. Hoekstra (Eds.), *Cellular Automata* (pp. 601–611). Springer. doi:https://doi.org/10.1007/978-3-540-30479-1_62.

Bertsekas, D. P., & Tsitsiklis, J. N. (1996). *Neuro-dynamic programming* volume 5. Belmont, MA: Athena Scientific.

Birge, J. R. (1982). The value of the stochastic solution in stochastic linear programs with fixed recourse. *Mathematical Programming*, 24, 314–325. doi:<https://doi.org/10.1007/BF01585113>.

Calafiore, G. C., & Campi, M. C. (2006). The scenario approach to robust control design. *IEEE Transactions on Automatic Control*, 51, 742–753. doi:<https://doi.org/10.1109/TAC.2006.875041>.

Cao, S., Song, W., & Lv, W. (2016). Modeling pedestrian evacuation with guiders based on a multi-grid model. *Physics Letters A*, 380, 540–547. doi:<https://doi.org/10.1016/j.physleta.2015.11.028>.

Cristiani, E., & Peri, D. (2017). Handling obstacles in pedestrian simulations: Models and optimization. *Applied Mathematical Modelling*, 45, 285–302. doi:<https://doi.org/10.1016/j.apm.2016.12.020>.

Dyer, J. R. G., Johansson, A., Helbing, D., Couzin, I. D., & Krause, J. (2009). Leadership, consensus decision making and collective behaviour in humans. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 364, 781–789. doi:<https://doi.org/10.1098/rstb.2008.0233>.

Goldberg, D. E. (1989). *Genetic Algorithms in Search, Optimization and Machine Learning*. Boston, MA, USA: Addison-Wesley Longman Publishing Co.

Gwynne, S. M. V., Hulse, L. M., & Kinsey, M. J. (2016). Guidance for the model developer on representing human behavior in egress models. *Fire Technology*, 52, 775–800. doi:<https://doi.org/10.1007/s10694-015-0501-2>.

Haghani, M. (2020). Optimising crowd evacuations: Mathematical, architectural and behavioural approaches. *Safety Science*, 128, 104745. doi:<https://doi.org/10.1016/j.ssci.2020.104745>.

Hammel, U., & Bäck, T. (1994). Evolution strategies on noisy functions: How to improve convergence properties. In Y. Davidor, H.-P. Schwefel, & R. Männer (Eds.), *Parallel Problem Solving from Nature–PPSN III* (pp. 159–168). Springer. doi:https://doi.org/10.1007/3-540-58484-6_260.

Helbing, D., Isobe, M., Nagatani, T., & Takimoto, K. (2003). Lattice gas simulation of experimentally studied evacuation dynamics. *Physical Review E*, 67, 067101. doi:<https://doi.org/10.1103/PhysRevE.67.067101>.

Helbing, D., & Johansson, A. (2013). Pedestrian, crowd, and evacuation dynamics. *arXiv preprint*, . ArXiv:1309.1609.

Helbing, D., & Molnár, P. (1995). Social force model for pedestrian dynamics. *Physical Review E*, 51, 4282. doi:<https://doi.org/10.1103/PhysRevE.51.4282>.

Heliövaara, S., Ehtamo, H., Helbing, D., & Korhonen, T. (2013). Patient and impatient pedestrians in a spatial game for egress congestion. *Physical Review E*, 87, 012802. doi:<https://doi.org/10.1103/PhysRevE.87.012802>.

Heliövaara, S., Korhonen, T., Hostikka, S., & Ehtamo, H. (2012a). Counterflow model for agent-based simulation of crowd dynamics. *Building and Environment*, 48, 89–100. doi:<https://doi.org/10.1016/j.buildenv.2011.08.020>.

Heliövaara, S., Kuusinen, J.-M., Rinne, T., Korhonen, T., & Ehtamo, H. (2012b). Pedestrian behavior and exit selection in evacuation of a corridor—An experimental study. *Safety Science*, 50, 221–227. doi:<https://doi.org/10.1016/j.ssci.2011.08.020>.

Hoogendoorn, S., & Bovy, P. H. L. (2003). Simulation of pedestrian flows by optimal control and differential games. *Optimal Control Applications and Methods*, 24, 153–172. doi:<https://doi.org/10.1002/oca.727>.

Hou, L., Liu, J.-G., Pan, X., & Wang, B.-H. (2014). A social force evacuation model with the leadership effect. *Physica A: Statistical Mechanics and its Applications*, 400, 93–99. doi:<https://doi.org/10.1016/j.physa.2013.12.049>.

Huang, H.-J., & Guo, R.-Y. (2008). Static floor field and exit choice for pedestrian evacuation in rooms with internal obstacles and multiple exits. *Physical Review E*, 78, 021131. doi:<https://doi.org/10.1103/PhysRevE.78.021131>.

Karamouzas, I., Skinner, B., & Guy, S. J. (2014). Universal power law governing pedestrian interactions. *Physical Review Letters*, 113, 238701. doi:<https://doi.org/10.1103/PhysRevLett.113.238701>.

Karamouzas, I., Sohre, N., Narain, R., & Guy, S. J. (2017). Implicit crowds: Optimization integrator for robust crowd simulation. *ACM Transactions on Graphics*, 36, 1–13. doi:<https://doi.org/10.1145/3072959.3073705>.

Korhonen, T. (2018). *Fire dynamics simulator with evacuation: FDS+Evac. Technical Reference and User's Guide*. Technical Report VTT Technical Research Centre of Finland. URL: http://virtual.vtt.fi/virtual/proj6/fdsevac/documents/FDS+EVAC_Guide.pdf.

Kretz, T., Grosse, A., Hengst, S., Kautzsch, L., Pohlmann, A., & Vortisch, P. (2011). Quickest paths in simulations of pedestrians. *Advances in Complex Systems*, 14, 733–759. doi:<https://doi.org/10.1142/S0219525911003281>.

Kurdi, H. A., Al-Megren, S., Althunyan, R., & Almulifi, A. (2018). Effect of exit placement on evacuation plans. *European Journal of Operational Research*, 269, 749–759. doi:<https://doi.org/10.1016/j.ejor.2018.01.050>.

Li, W., Li, Y., Gong, J., & Shen, S. (2016). The Trace Model: A model for simulation of the tracing process during evacuations in complex route environments. *Simulation Modelling Practice and Theory*, 60, 108–121. doi:<https://doi.org/10.1016/j.simpat.2015.09.011>.

Løvås, G. (1995). On performance measures for evacuation systems. *European Journal of Operational Research*, 85, 352–367. doi:[https://doi.org/10.1016/0377-2217\(94\)00054-G](https://doi.org/10.1016/0377-2217(94)00054-G).

Løvås, G. (1998). Models of wayfinding in emergency evacuations. *European Journal of Operational Research*, 105, 371–389. doi:[https://doi.org/10.1016/S0377-2217\(97\)00084-2](https://doi.org/10.1016/S0377-2217(97)00084-2).

McCormack, P., & Chen, T.-Y. (2014). Optimizing leader proportion and behavior for evacuating buildings. In *Proceedings of the 2014 Symposium on Agent Directed Simulation ADS '14* (pp. 13:1–13:6). Society for Computer Simulation International.

Nishinari, K., Kirchner, A., Namazi, A., & Schadschneider, A. (2004). Extended floor field CA model for evacuation dynamics. *IEICE Transactions on Information and Systems*, 87, 726–732.

NIST (2020). Fire dynamics simulator. URL: <https://pages.nist.gov/fds-smv/> accessed: 2020-03-16.

Oliphant, T. et al. (2020). Numba: A high performance python compiler. URL: <https://numba.pydata.org> accessed: 2020-03-16.

Pelechano, N., & Badler, N. I. (2006). Modeling crowd and trained leader behavior during building evacuation. *IEEE Computer Graphics and Applications*, 26, 80–86. doi:<https://doi.org/10.1109/MCG.2006.133>.

Proulx, G. (2002). Movement of people: the evacuation timing. In P. J. DiNenno et al. (Eds.), *SFPE Handbook of Fire Protection Engineering* (pp. 342–366). the National Fire Protection Association.

Saadatseresht, M., Mansourian, A., & Taleai, M. (2009). Evacuation planning using multiobjective evolutionary optimization approach. *European Journal of Operational Research*, 198, 305–314. doi:<https://doi.org/10.1016/j.ejor.2008.07.032>.

Schadschneider, A., Klingsch, W., Klüpfel, H., Kretz, T., Rögsch, C., & Seyfried, A. (2008). Evacuation dynamics: Empirical results, modeling and applications. *arXiv preprint*, . ArXiv:0802.1620.

von Schantz, A. (2020). Minimizing the evacuation time of a crowd from a complex building using rescue guides – code (Version 1.2). Zenodo. doi:[10.5281/zenodo.3831338](https://doi.org/10.5281/zenodo.3831338).

von Schantz, A., & Ehtamo, H. (2015). Spatial game in cellular automaton evacuation model. *Physical Review E*, 92, 052805. doi:<https://doi.org/10.1103/PhysRevE.92.052805>.

von Schantz, A., & Ehtamo, H. (2019). Pushing and overtaking others in a spatial game of exit congestion. *Physica A: Statistical Mechanics and its Applications*, 527, 121151. doi:<https://doi.org/10.1016/j.physa.2019.121151>.

Sethian, J. A. (1999). *Level set methods and fast marching methods: Evolving interfaces in computational geometry, fluid mechanics, computer vision, and materials science* volume 3. Cambridge University Press.

Vermuyten, H., Beliën, J., De Boeck, L., Reniers, G., & Wauters, T. (2016). A review of optimisation models for pedestrian evacuation and design problems. *Safety Science*, 87, 167–178. doi:<https://doi.org/10.1016/j.ssci.2016.04.001>.

Wang, X., Zheng, X., & Cheng, Y. (2012). Evacuation assistants: An extended model for determining effective locations and optimal numbers. *Physica A: Statistical Mechanics and its Applications*, 391, 2245–2260. doi:<https://doi.org/10.1016/j.physa.2011.11.051>.

Weiss, T., Littenecker, A., Jiang, C., & Terzopoulos, D. (2019). Position-based real-time simulation of large crowds. *Computers & Graphics*, 78, 12–22. doi:<https://doi.org/10.1016/j.cag.2018.10.008>.

Yao, Z., Wang, J. S., Liu, G. R., & Cheng, M. (2004). Improved neighbor list algorithm in molecular simulations using cell decomposition and data sorting method. *Computer Physics Communications*, 161, 27–35. doi:<https://doi.org/10.1016/j.cpc.2004.04.004>.

Zhou, M., Dong, H., Zhao, Y., Ioannou, P. A., & Wang, F.-Y. (2019). Optimization of crowd evacuation with leaders in urban rail transit stations. *IEEE Transactions on Intelligent Transportation Systems*, 20, 4476–4487. doi:<https://doi.org/10.1109/TITS.2018.2886415>.