

SIMULATING KNOWLEDGE AND INFORMATION IN PEDESTRIAN EGRESS

Kyle Feuz^{1,2} and Vicki Allan³

¹*Computer Science, Utah State University, Logan, Utah, U.S.A.*

²*Electrical Engineering and Computer Science, Washington State University, Pullman, Washington, U.S.A.*

³*Computer Science, Utah State University, Logan, Utah, U.S.A.*

Keywords: Reinforcement-learning, Pedestrian simulation, Egress assistance, Congestion, Multi-agent systems.

Abstract: Accurate pedestrian simulation is a difficult yet important task. One of the main challenges with pedestrian simulation is providing the simulated pedestrians with appropriate amounts of route knowledge to be used in the route selection algorithm. In this paper, we propose a novel use of reinforcement learning as a means to represent different amounts of route knowledge. Using this techniques we show the impact learning about route distances and average route congestion levels has upon the egress time of pedestrians. We also look at the effect that dynamic congestion information has upon the efficiency of pedestrian egress.

1 INTRODUCTION

In recent years, pedestrian simulation has become an important research topic (Santos and Aguirre, 2004; Pan, 2006; Helbing and Johansson, 2009). Pedestrian simulation models are useful in the design of safe facilities, validation of fire codes, and automatic tracking and surveillance of pedestrians in live video feeds (Antonini et al., 2006). An important area of pedestrian simulation is the route selection algorithm. Most route selection algorithms either assume perfect knowledge of egress routes or they assume no prior knowledge of the egress routes. While common, these methods are not an accurate representation of pedestrian knowledge. Rarely would a pedestrian have complete route knowledge, yet, having no prior route knowledge is also unrealistic for most cases. A few simulators allow route knowledge to be entered manually by the user to simulate different route knowledge for different pedestrians (Gwynne et al., 2001; PTV AG, 2011), which may require a large time commitment by the user to properly set up the environment. We propose a novel application of reinforcement learning to provide pedestrians with individualized knowledge of the building without requiring a large time commitment from the user. Pedestrians can learn about the environment in an initial learning phase, and then the actual simulation is run with different pedestrians having learned various routes.

Another factor which can affect route selection is

congestion. The use of reinforcement learning to supply pedestrian agents with prior knowledge about the building can be extended to include the average congestion levels of the different routes. Using this technique, we can analyze the effect that utilizing congestion knowledge has upon the egress time and efficiency of the simulation. In traffic management, studies conflict as to whether or not providing dynamic information about traffic congestion conditions improves the efficiency of the road network. Some studies indicate that providing such information can lead to road usage oscillation patterns as drivers switch between two alternate routes (Wahle et al., 2002). The question of the effectiveness of providing congestion information has yet to be answered regarding pedestrian egress. The effect of learning typical congestions levels in a building prior to the actual simulation is also unanswered. We fill these gaps by analyzing the effect of incorporating dynamic route congestion information and learned route congestion information into the route selection algorithm.

2 RELATED WORK

Reinforcement learning has been studied extensively for several decades (Kaelbling et al., 1996). Different algorithms and techniques have been developed, each with benefits and drawbacks. In general, re-

inforcement learning algorithms can be divided into two broad categories: model-free learning and model-based learning. The main difference between these two techniques is that in model-based learning, an agent learns about both the transition relationship between states and the reward function, whereas in a model-free technique, an agent only learns about the reward function. For a good survey of reinforcement learning algorithms, consult (Kaelbling et al., 1996).

The effects of dynamic congestion information in traffic management is a well-studied topic which has not yet received much attention in pedestrian situations. Dia provides a framework for simulating driver behavior with dynamic route information. He leaves as an open question what effect such information will actually have upon route selection behavior (Dia, 2002). Wahle et al. study the effect of dynamic congestion information in traffic scenarios (Wahle et al., 2002). They use simulation models to predicate the effect that different congestion messages will have upon traffic congestion. Their findings indicate that the results are dependent upon the type of information provided, but in general, dynamic information tends to decrease the overall network efficiency as oscillation patterns of road usage develop. Roughgarden shows that selfish routing does not minimize the total latency of a network and provides bounds upon the cost of selfish routing for several different latency functions and network topologies (Roughgarden and Tardos, 2002). However, using game-theory, Helbing et al. discover the emergence of alternating cooperation as a fair and system-optimal road usage behavior in a route choice game (Helbing et al., 2005). They conduct empirical tests using an iterated 2-4 player route choice game. Cooperation tends to emerge when individuals also exhibit exploratory behavior. They do not consider the case of providing dynamic information about the road conditions.

Although dynamic congestion information has not been heavily applied to pedestrian simulation, several researchers have included congestion consideration while modeling pedestrian egress. The work of Hoogendoorn and Bovy includes the cost of congestion when selecting routes and activities to perform (Hoogendoorn and Bovy, 2004). The congestion information can either be derived from the pedestrians current perceptions or it can be based upon future predictions of congestion levels. How this information affects the overall efficiency of the system is not discussed. Banerjee et al consider the complexity issues of dynamically discovering congestion and rerouting agents accordingly (Banerjee et al., 2008). Their model assumes complete route distance information is known to pedestrians and that only pedes-

trians which perceive the congestion will choose new routes. This is in contrast to our model where route information may not be known and where congestion may be known or estimated from previous experience even when the actual congestion cannot be directly perceived.

Pan represents one of the more comprehensive pedestrian behavioral models (Pan, 2006). He includes pedestrian characteristics such as competitive, leader-following, altruistic, queuing, and herding. Using these characteristics, an agent considers visible routes before identifying its currently preferred route. Similarly, Koh, Lin, and Zhou (Koh et al., 2008) define an agent which only considers congestion and obstructions which can be directly perceived by the agent. However, knowledge of the location of the end goals appears to be available to all pedestrians. A common simulation environment, buildingEXODUS, assumes the agents know about a set of user specified routes or all routes in the absence of the specification (Gwynne et al., 2001). VISSIM, a commercially available pedestrian simulator, first processes the building layout to generate perfect route information for the pedestrians (PTV AG, 2011). In VISSIM, the user also has the option of specifying specific routes for specific pedestrian sets (PTV AG, 2011).

3 SIMULATION ENVIRONMENT

Our study is performed using the Pedestrian Leadership and Egress Assistance Simulation Environment (PLEASE) (Feuz, 2011). PLEASE is built upon the multi-agent modeling paradigm where each pedestrian is represented as an individually rational agent capable of perceiving the environment and reacting to it. In PLEASE, pedestrian agents can perceive obstacles, hazards, routes, and other agents. The agents are capable of basic communication to allow for the formation and dissolution of coalitions and the sharing of knowledge. The agents use a two tier navigational module to control their movement within the simulation environment. The high-level tier evaluates available routes and selects a destination goal. The low-level tier, based on the social force model (Helbing and Johansson, 2009), performs basic navigation and collision avoidance.

Typically, reinforcement learning algorithms are used to discover a near-optimal policy. In fact, many reinforcement learning algorithms provably converge to the optimal policy (Kaelbling et al., 1996). One benefit of reinforcement learning to our simulation is the fact that when the search is truncated, a less

than perfect solution is found. These solutions can be used to automatically generate various levels of pedestrian knowledge about the building configuration. These sub-optimal policies do have a unique constraint though as well: they must still be realistic. A learned policy which (when followed) never results in the successful egress of the agent is unacceptable. For this reason, we have implemented the reinforcement learning algorithm using model-based techniques. The details of the implementation follow.

Each agent builds a model of the building layout and the associated costs of available routes. To do this, the agent abstracts the building layout into a graph-based view. A common abstraction of building layouts is to represent rooms as nodes in the graph and doorways between rooms as edges in the graph. For the purpose of reinforcement learning, however, this abstraction is too coarse-grained. An agent is forced to associate a single cost (the edge weight) between two arbitrary, connected rooms. The true cost actually varies significantly depending upon the agent's location in the room. If time and space considerations are ignored, the building can be discretized into arbitrarily small grid cells, which allows the cost between nodes to be represented more accurately. Of course, this method is too costly in terms of time and space to be practical for buildings of even modest size. PLEASE uses a building representation in between these two extremes. To do this, we introduce the concept of decision points. A decision point is simply a point in the building at which an agent must decide in which direction he will proceed. These points may be placed at any arbitrary location, but in our models, the decision points are always placed at doorways and corridor intersections. We select these locations because they are areas which pedestrian must pass through to move from one area of the building to another. This prevents the systems from forcing a particular path upon an agent. The nodes in the graph represent decision points in the building, and weighted edges between nodes represent the average cost of a path between two decision points. This provides more fine-grained control over the costs learned while still being manageable for larger buildings.

Pseudo-code for the learning algorithm is shown in Figure 1. Initially, the agent's model is empty as the agent has no prior knowledge about the building. Each time an agent passes through a decision point, the agent estimates the cost (based upon distance and/or congestion levels) to all other visible decision points in the room. (See Formula 1-3). Additionally, the agent estimates the cost to other decision points known by the agent to be in the room. The weighted edge between decision points is then

updated to reflect the newly estimated costs. Decision points which are not currently represented in the graph are added as necessary.

Definitions:

model - the adjacency matrix for the building layout representation

d - the decision point whose cost is being updated

dp - decision points in the same room as *d*

alpha - learning parameter of the algorithm, determines the weight applied to new cost estimates

estimateCost - estimates the cost between two decision points. See Formula 1 - 3

model.insert - inserts new rows and columns into the adjacency matrix as needed

```

Begin UpdateCost(DecisionPoint d)
  foreach DecisionPoint dp in room
    if dp isVisible or isKnown
      cost = estimateCost(d, dp)
      if d, dp in model
        tmp = alpha * (cost - model[d][dp])
        model[d][dp] += tmp
        model[dp][d] += tmp
      else
        model.insert(d, dp, cost)
  End

```

Figure 1: Algorithm used by the learning agent to update the estimated cost between decision points.

The agents estimate the cost from one point in the building (*dp1*) to another point in the building (*dp2*) based upon the distance and congestion levels between the two points. This estimate is specified by Formulas 1-3, where *cost* is the estimated cost of moving from *dp1* to *dp2*, w_{cg} is the user-specified weight for congestion costs, w_d is the user-specified weight for distance costs, sp_i is the desired speed of the current pedestrian *i*, sp_j is the current speed of agent *j*, $n_{dp1,dp2}$ is the number of agents along the path from *dp1* to *dp2*, *N* is the total number of agents, s_1 is 1 if $sp_j < sp_i$ and 0 otherwise, $dist(dp1, dp2)$ is the distance between *dp1* and *dp2*, and *maxDistance* is the maximum distance between any two connected decision points which is defined as the length of the diagonal of the building.

Both the distance cost and the congestion cost are weighted by user-specified parameters so that different relative weights can be chosen. Agents in the simulation are able to accurately estimate the distance to visible points within the simulation model as well as being able to estimate the distance to points which they have previously visited. The distance is normalized using the maximum distance between two points on the simulation map. The congestion cost is estimated using the difference in speeds between the

current pedestrian and other pedestrians that exist between the two points in the building. For each pedestrian along the selected route, if its speed is slower than the desired speed of the pedestrian, then a cost is incurred relative to the speed difference. The cost is raised to the square so that smaller speed differences count less than larger differences. Finally, the result is normalized by the worst-case cost (i.e. if every pedestrian in the simulation was along the selected route and was not moving).

$$DistCost = \frac{w_d * dist(dp1, dp2)}{(maxDistance)} \quad (1)$$

$$CongCost = w_{cg} * \frac{\sum_{j=0}^{n_{dp1, dp2}} ((sp_i - sp_j) * s_1)^2}{sp_i * N} \quad (2)$$

$$cost = CongCost + DistCost \quad (3)$$

At this point, the agent must select the next route to follow. Pseudo-code for the route selection algorithm is shown in Figure 2. To do this, the agent performs a breath-first search starting from each known decision point (dp) in the current room. If a path is found from the decision point to an end goal (g), the cost of the path is computed as the cost to dp plus the learned cost from dp to g . If no path is found to g , then the cost is computed as the cost of d plus $UNEXPLORED_COST$. $UNEXPLORED_COST$ is a user-specified parameter representing the cost of choosing a route whose destination is not known. With probability p , the agent selects a random decision point to proceed towards, and with a probability of $1 - p$, the agent selects the decision point of least cost. The probability factor represent the probability an agent chooses to explore a different route. When the agent is learning we set this probability to 0.15. This value will reflect the speed with which agents learn a building. When learning congestion cost, this value will also affect the reliability of the learned congestion costs. Agent training happens concurrently for all agents in the simulation. This creates a moving-target problem because congestion levels are constantly fluctuating as agents change their respective policies based upon the congestion levels encountered previously. When the probability of exploring is high, a large number of agents will not be using routes they normally would if they were not exploring which leads to inaccurately learned congestion costs.

Definitions:

dp - decision point in the current room to consideration

$explore$ - normally distributed random value between 0-1

p - probability of exploration

$cost$ - dictionary of costs for decision points considered

$estimateCostTo(dp)$ - similar to $estimateCost$ in Figure 1 but uses the agents position as $dp1$

$BFSCost(dp)$ - the cost found by performing a breadth-first search from dp to the end goal

```

Begin routeSelection()
  foreach DecisionPoint dp in room
    cost[dp] = estimateCostTo(dp) + BFSCost(dp)
  explore = random()
  if explore <= p
    return random DecisionPoint in room
  else
    return arg min cost[dp]
End

```

Figure 2: Algorithm used by the agent to select the desired route of travel.

4 CONGESTION CONSIDERATIONS

As we are interested in the effects of congestion on the egress efficiency of the system, we consider the two cases: 1) ignore current congestion levels and 2) adjust decisions based on directly perceived congestion.

We use the case of ignoring congestion as a base case against which we can compare all other cases. For many situations, we expect that completely ignoring congestion will lead to slower egress times as the building corridors are used inefficiently. However, ignoring congestion is still a feasible pedestrian behavior. Generally, pedestrians prefer to travel along paths they have previously traveled (Ozel, 2001). This may mean that, in spite of congestion, they continue to travel along their preferred route. Congestion might also be ignored if the pedestrian believes that other routes will not decrease their egress time.

Adjusting to directly perceivable congestion is common in many simulation models (Koh et al., 2008; Hoogendoorn and Bovy, 2004). Intuitively, it makes sense that pedestrians adjust their route based upon perceived congestion. From a modeling perspective, this case has the additional benefit of not requiring any additional knowledge about congestion in other areas of the building. The question remaining is, "Does it improve the overall egress times?"

5 KNOWLEDGE CONSIDERATIONS

We are interested not only in the effect of reacting to congestion upon egress times, but also in the effect congestion knowledge has upon egress times. We consider three types of knowledge which pedestrians may have: 1) learned route distance knowledge 2) learned route congestion/distance knowledge, 3) system-provide route congestion/distance knowledge.

The case of route distance knowledge represents pedestrians who have learned route distances but not route congestion levels. These pedestrians primary concern is arriving at the destination rather than the congestion levels along the way. The completeness of the distance knowledge which a pedestrian has is dependent upon the amount of training the agent has.

Knowledge of the average congestion costs is more reflective of reality as pedestrians familiar with a building are also typically familiar with the route usage patterns. This case assumes that pedestrians remember congestion costs from previous experience in the building in addition to the distances between various decision points. The congestion costs are associated with routes between decision points. Each time an agent travels a given route, the expected cost for that route is then updated. The completeness of the distance/congestion knowledge which a pedestrian has is dependent upon the amount of training the agent has.

The final case we consider is providing pedestrians with dynamic route congestion information and route distance information. This allows a pedestrian to evaluate all possible routes for distance and congestion, even when those routes are not directly perceivable (i.e. the route cannot be seen). Such information may one day be generally available to pedestrians through personal hand-held devices or public displays (Barnes et al., 2007; Kray et al., 2005; Müller et al., 2008).

6 EXPERIMENTAL SETUP

All the experiments conducted in this paper use four different building layouts (see Figure 3). Building A is designed with specific congestion considerations in mind. To pedestrians in the inner rooms, each room doorway appears to be of equal value. However, the lower doorways lead to a wider corridor and exit and will thus be able to accommodate more pedestrians. Building B is designed to be representative of a general building layout. Buildings C and D are approximations of actual buildings found on the California

State University, Long Beach campus.

6.1 Experiment 1

The purpose of the first set of experiments is to demonstrate the feasibility of using reinforcement learning as a means to represent pedestrian knowledge in a simulation environment. To do this, we show that as the number of learning trials (to which an agent is subjected to) increases, the amount of building knowledge the pedestrian acquires also increases. We show that this increase in building knowledge leads to a corresponding decrease in pedestrian egress times.

For each building, we conduct the test as follows. Five hundred pedestrians are trained in the building for 100 simulation runs during which the agents learn route distance costs. After each simulation run, the agents' current policy is saved to disk so that we can recover the policy learned after any given number of simulations runs.

In order to determine how much knowledge an agent has gained about a particular building, we first need to define some metrics. We consider three key factors affecting route knowledge: 1) the number of known decision points (node knowledge), 2) the number of known paths between decision points (edge knowledge), and 3) the number of decision points known to be direct exits (exit knowledge). Using these metrics, we can then calculate the average amount of knowledge obtained by the agents for each trial run.

Figure 4 shows the average effect of multiple training runs on the total knowledge an agent has. As can be seen from the graphs, the different metrics indicate different knowledge levels, but the values of all metrics show an increase as the number of training runs increases. Agents quickly learn a high percentage of the decision points and paths between decision points, but for the key decision points representing building exits the percentages are lower. This indicates that although the agent learns many internal routes after 100 training runs, they are learning different exits at a slower rate.

As the amount of knowledge pedestrians have increases, so should the efficiency with which agents egress from the building. We measure the egress time of 500 pedestrians randomly distributed throughout the building, averaged over 20 simulations using policies of various training levels. Averaging the results over 20 simulation runs provides relatively small error bars which boost our confidence in the accuracy of the mean egress times obtained for each training level. Figure 5 compares the average egress times ob-




Figure 3: Building layouts used in the congestion experiments.




Figure 4: Percentage of knowledge gained over time using three metrics.

tained when agents have gone through 10, 50, and 100 training runs for each building.

The effect of additional training in building A is minimal. This implies that the additional knowledge gained is not helpful in improving egress times. Building A is fairly simple and therefore the general layout can be learned quickly. Building B and building C both show substantial improvement in egress times as the number of training runs increases indicating that the knowledge gained by the pedestrians is indeed helpful in improving egress times. Building D shows substantial improvement in egress times between 10 and 50 training runs, but then little change occurs between 50 and 100 training runs. This correlates to the previous results in Figure 4, where the amount of knowledge gained between 50 and 100 training runs is much less for building D than it is for the other buildings.

6.2 Experiment 2

The next set of experiments are intended to measure the effectiveness of learning average route congestion costs in addition to route distance costs. The experiments also measure the effectiveness of reacting to currently visible congestion and adjusting the selected route accordingly. Notice the distinction between

learning congestion levels and reacting to current congestion levels. ‘Choose to react’ to congestion or ‘ignore congestion’ does not imply either a knowledge or a lack of knowledge of average congestion levels. It is merely a decision of whether or not to include current congestion levels in the decision-making process. Conversely, having congestions knowledge does not imply that the agent must react to current congestion levels, only that the agent will consider previous learned congestion levels when making the decision. Thus, an agent having no previous congestion knowledge can react to current congestion levels, and an agent having previous congestion knowledge can choose to ignore current congestion levels.

For each building, we conduct the test as follows. Five hundred pedestrians are trained in the building for 100 simulation runs, learning both route distance costs and average congestion levels. The agents’ current policies are check pointed after 100 training runs so that we can compare the egress times when pedestrians have high levels of knowledge. We then measure the total egress time of 500 pedestrians randomly distributed throughout the rooms, averaged over 20 simulations. There are two parameters that we adjust in these tests: whether the pedestrian reacts to congestion, and what type of knowledge the pedestrian has. Pedestrians can either ignore




Figure 5: Percentage of pedestrians exited over time using three levels of training.




Figure 6: Percentage of pedestrians exited over time using three levels of training.

current congestion levels or react to current congestion levels, and pedestrians can have either learned distance knowledge, learned congestion knowledge (which also includes distance knowledge), or system provided knowledge for both route distances and congestion levels. Therefore we have six cases to consider: 1) ignore current congestion and have learned distance knowledge (Ign-Dist), 2) ignore current congestion and have learned both distance and congestion knowledge (Ign-Cong), 3) ignore current congestion and have perfect distance knowledge provided by the system (Ign-Sys), 4) adjust to congestion and have learned building distance knowledge (Adj-Dist), 5) adjust to congestion and have learned both distance and congestion knowledge (Adj-Cong), and 6) adjust to congestion and have perfect distance and congestion knowledge provided by the system (Adj-Sys).

The results are shown in Figure 6. In every building layout tested, agents which have learned both route distances and congestion levels have faster egress times than agents which have learned only route distances. This indicates that learning average congestion levels and using that knowledge in pedestrian egress is beneficial. However, the same cannot be said about reacting to congestion. In building A, reacting to current congestion always improves performance. This is not surprising because building A

is specifically designed to contain severe congestion problems which are easily mitigated. In Buildings B, C and D, reacting to current congestion yields little change in overall egress time except when pedestrians have only route distance knowledge. In this case, reacting to the current congestion levels actually decreases the overall performance of the agents. This is occurring because the pedestrians are uniformly distributed within the building so the congestion is also well distributed. Thus, when a pedestrian chooses to take an alternate route, they soon discover that it is equally congested. Finally, in building D, reacting to congestion improves performance if the pedestrian has learned previous congestion levels. Although the pedestrians are still uniformly distributed, the routes to the exits are not. Knowing the typical congestion levels allows an agent to make a better decision when reacting to the current congestion levels.

Interesting patterns in the data can also be observed when the egress times of agents with different types of knowledge are compared. In half the buildings (A, and D) utilizing learned congestion levels provides the best egress times, even outperforming system provided information. This is probably due to the oscillation which can occur when dynamic information is provided. As is also seen in traffic management, providing dynamic information can

lead to many pedestrians switching routes simultaneously which decreases the efficiency with which pedestrians are able to evacuate the building. In every building layout tested, when pedestrians have only learned distance information, the performance is the worst of all possibilities considered. Interestingly though, a pedestrian having system information but ignoring current congestion levels and using only distance information is able to egress from most buildings quickly. However, the distance information of such a pedestrian is complete. One would expect that with enough training, pedestrians having learned only distance cost would also be able to egress from buildings with similar efficiency.

7 CONCLUSIONS

Providing agents with perfect knowledge is unrealistic for many pedestrian egress situations. However, manually specifying specific route knowledge can be a difficult and time-consuming task. We have shown that reinforcement learning can be applied to successfully represent different levels of knowledge about a building layout and produces egress times dependent upon the knowledge level of the pedestrians. We have also provided three different metrics for measuring the amount of building knowledge an agent has.

Using reinforcement learning, we have also shown that learning congestion cost in addition to distance costs leads to quicker egress times. However, reacting to current congestion levels has ambiguous results. This is consistent with similar studies in the traffic management domain. The layout of the building is found to have an impact on the strategy a pedestrian should use to minimize egress time.

REFERENCES

- Antonini, G., Bierlaire, M., and Weber, M. (2006). Discrete choice models of pedestrian walking behavior. *Transportation Research Part B: Methodological*, 40(8):667–687.
- Banerjee, B., Bennett, M., Johnson, M., and Ali, A. (2008). Congestion avoidance in multi-agent-based egress simulation. In *IC-AI*, pages 151–157.
- Barnes, M., Leather, H., and Arvind, D. (2007). Emergency evacuation using wireless sensor networks. *Local Computer Networks, Annual IEEE Conference on*, 0:851–857.
- Dia, H. (2002). An agent-based approach to modelling driver route choice behaviour under the influence of real-time information. *Transportation Research Part C: Emerging Technologies*, 10(5-6):331–349.
- Feuz, K. (2011). Pedestrian leadership and egress assistance simulation environment. Master's thesis, Utah State University.
- Gwynne, S., Galea, E. R., Lawrence, P. J., and Filippidis, L. (2001). Modelling occupant interaction with fire conditions using the buildingexodus evacuation model. *Fire Safety Journal*, 36(4):327–357.
- Helbing, D. and Johansson, A. (2009). Pedestrian, crowd and evacuation dynamics. In *Encyclopedia of Complexity and Systems Science*, pages 6476–6495. Springer.
- Helbing, D., Schönhof, M., Stark, H., and Holyst, J. (2005). How individuals learn to take turns: Emergence of alternating cooperation in a congestion game and the prisoner's dilemma. *Advances in Complex Systems*, 8(1):87–116.
- Hoogendoorn, S. and Bovy, P. (2004). Pedestrian route-choice and activity scheduling theory and models. *Transportation Research Part B: Methodological*, 38(2):169–190.
- Kaelbling, L., Littman, M., and Moore, A. (1996). Reinforcement learning: A survey. *Journal of Artificial Intelligence Research*, 4:237–285.
- Koh, W. L., Lin, L., and Zhou, S. (2008). Modelling and simulation of pedestrian behaviours. In *PADS '08: Proceedings of the 22nd Workshop on Principles of Advanced and Distributed Simulation*, pages 43–50, Washington, DC, USA. IEEE Computer Society.
- Kray, C., Kortuem, G., and Krüger, A. (2005). Adaptive navigation support with public displays. In *Proceedings of the 10th international conference on Intelligent user interfaces, IUI '05*, pages 326–328, New York, NY, USA. ACM.
- Müller, J., Jentsch, M., Kray, C., and Krüger, A. (2008). Exploring factors that influence the combined use of mobile devices and public displays for pedestrian navigation. In *Proceedings of the 5th Nordic conference on Human-computer interaction: building bridges, NordiCHI '08*, pages 308–317, New York, NY, USA. ACM.
- Ozel, F. (2001). Time pressure and stress as a factor during emergency egress. *Safety Science*, 38(2):95–107.
- Pan, X. (2006). *Computational Modeling of Human and Social Behaviors for Emergency Egress Analysis*. PhD thesis, Stanford University, Stanford, California.
- PTV AG (2011). <http://www.ptvag.com/software/transportation-planning-traffic-engineering/software-system-solutions/vissim-pedestrians/>. Accessed on: April 4, 2011.
- Roughgarden, T. and Tardos, É. (2002). How bad is selfish routing? *Journal of the ACM (JACM)*, 49(2):236–259.
- Santos, G. and Aguirre, B. E. (2004). A critical review of emergency evacuation simulation models. In *Building Occupant Movement during Fire Emergencies*.
- Wahle, J., Bazzan, A., Klügl, F., and Schreckenberg, M. (2002). The impact of real-time information in a two-route scenario using agent-based simulation. *Transportation Research Part C: Emerging Technologies*, 10(5-6):399–417.