IdealCity: a Hybrid Approach to Seismic Evacuation Modeling

(Article begins on next page)

IdealCity: a Hybrid Approach to Seismic Evacuation Modeling

Edoardo Battegazzorre[a], Andrea Bottino[a], Marco Domaneschi[b], Gian Paolo Cimellaro[b]

[a] *Department of Control and Computer Engineering*
[b] *Department of Structural, Geotechnical and Building Engineering, Politecnico di Torino, Corso Duca degli Abruzzi, 24, 10129 Torino, ITALY*

Abstract

In recent years, the analysis of the community response in case of disastrous events has become a research topic of paramount relevance due to the increasing number of calamities like flooding, hurricanes, and earthquakes. In particular, the possibility to use computer simulations to model and study the behavior of thousands of people during an emergency evacuation can provide valuable information to support many processes involved in emergency management. To this end, this work presents IdealCity, a hybrid model for evacuation simulation that couples the representation of the built environment and the transportation network with an agent-based simulation of the urban population. IdealCity can estimate the buildings' damages and debris generated by a seismic event along with their effects on the other model layers (the agents and the roads). Besides that, the simulation takes into consideration as well the emergency response system by modeling shelters, hospitals, and ambulances (each of which has a specific behavior within the environment). The model has been implemented and tested in a challenging test-bed that considers about 900,000 individuals, four different seismic scenarios, and three different times of the day. Results show that IdealCity can be used not only for predicting the population response but also for allowing decision-makers to estimate and intervene on critical response parameters, thus improving the inherent community resilience.

Keywords: Large scale simulations, Agent-based models, Human behavior, Seismic evacuation

---

1. Introduction

In recent decades, the world's population has witnessed and experienced major natural disasters like hurricanes, floods, earthquakes, or landslides that caused irreversible and uncontrollable damages. Large urban areas can amplify their effect due to population density and the interconnections between different infrastructures (such as buildings, transport and distribution networks). In these specific scenarios, the possibility to simulate and study the response of the population in the emergency evacuation can help optimize evacuation plans (Liu et al. [1], Wang et al. [2]), shorten the average evacuation time (Mikulik et al. [3], Meng and Jia [4]) and improve medical care by reducing both the number of people exposed to threats and their exposure time, thus leading to a decrease in injuries and victims. Furthermore, crowd evacuation modeling can offer decision-makers significant insights on the traffic flow and the configuration of resources (such as evacuation vehicles and relief facilities) required in case of large-scale disasters.

The interest in developing effective evacuation modeling (EM) approaches started in the 1980s after the nuclear disasters of Three Mile Island and Chernobyl, which showed that many emergency authorities were unprepared to tackle these events (Urbanik et al. [5], Sheffi et al. [6], Johnson Jr and Zeigler [7]). Then, EM received a subsequent boost from the need to develop preparedness studies for several natural hazards as hurricanes, tropical storms, and wildfires (Hobeika and Jamei [8], Pal et al. [9]).

Researchers presented several approaches to model the characteristics and behavior of a real crowd at different levels of granularity. The *flow-based* methods treat the crowd as a whole and define its behavior in terms of the continuous flow of evacuees from an origin to potential destinations along a graph of predefined possible routes. These approaches allow modeling large-size crowds but involve the adoption of constant physical, demographic, and perceptual attributes for the whole crowd. *Cellular-automata* (CA) models support the representation of each person as a single element. However, they fail to replicate the real and complex crowd behaviors in emergency evacuation processes due to the homogeneous nature of the matrix cells used to discretize the space.

On the contrary, *agent-based models* (ABM) represent the individual evacuees as *agents*, i.e. entities capable of making free movements and taking autonomous decisions based on (i) the data gathered by sensing the environment around them and (ii) the set of rules assigned to each type of agent

[10]. Therefore, ABMs allow fine-grained modeling of the crowd in terms of parameters like age, gender, role, and health conditions. Furthermore, they can model as well the real-world interactions between multiple persons, thus providing a more realistic representation of the evacuation process (Batty [11]).

It is worth noticing a certain overlap between CA and ABM approaches. CA models are usually preferred when the simulation space is represented in the form of a grid, e.g. Geographic Information System (GIS), or when model states and state transition probabilities are known and stable. On the contrary, ABMs are superior when the model basis is a behavioral unit, such as an individual, and the simulation process consists of interactions over time among different types of agents (Clarke [12]). In these situations, ABMs can model complex and different behaviours for each agent (e.g., see the *belief–desire–intention* architectures of Cimellaro et al. [13], Bordini et al. [14]), as well as their fine-grained interactions with other agents and the environment.

In the literature, when dealing with extensive scenarios such as districts or whole cities, ABM-based EM investigations have been extensively applied to assess community resilience and preparedness for natural disasters like tsunamis and flooding. For instance, [15] deals with the innovative concept of "rapid resilience" to tsunamis by modifying the urban morphology of the Chilean city of Talcahuano. Solis and Gazmuri [16] developed and validated an ABM simulation to quantify the tsunami response of the Iquique city's population under different evacuation policies. The evacuation model of Mas et al. [17] integrates a numerical simulation of a tsunami with a casualty estimation, leveraging a multi-agent programming language and GIS data as spatial input. Simulations carried out on the Arahama village (Japan) were assessed through comparison with the recorded outcomes of a real event.

In the context of earthquake EM, most of the works in the literature focuses on the restricted scenario of building evacuation. For instance, Li et al. [18] used an ABM to simulate classroom evacuation in a real event. The ABM parameters were tuned using real-life video data, and the results emphasized how a trained leader can positively affect the evacuation process helping the crowd maintain calm and order. The paper by Xiao et al. [19] simulates the resident evacuation from a building in a real event (the Ludian earthquake in 2014) with the aim of identifying variations to the layout of the inner walls that could reduce the safety escape time. Video footage of the Wenchuan 2018 earthquake in China have been collected by Yang et al.

[20] to analyze differences between real-life escape and simulations. In an attempt to increase the realism in evacuation modelling and show the effect of social attachment on the number of victims, Bañgate et al. [21] developed a multi-agent model of human behavior during a seismic crisis that is based on social attachment theory.

Less studies addressed post-earthquake emergency evacuation at the urban scale. An example is the work of Na and Banerjee [22]. This model focuses on traffic flow and takes into account different types of vehicles. However, it does not consider how the debris generated by the built environment affects the transportation network. The first contribution in this direction was proposed by Torrens [23], which presented a "unified" meta-system to merge physical and human process models within the simulation. Even if this approach includes a detailed representation of the real world, it lacks a structural analysis following civil engineering principles that consider the individual building response to the seismic actions. Moreover, the framework does not consider injured agents and emergency operators and it does not include the effects of emotions into the agent behavior.

Given the relevance of the emotional dimension for the real individuals, several authors started including it into their models. Examples are Lujak and Ossowski [24], which considered the influence of stress on human reactions, Cimellaro et al. [13, 25], which analyzed the impact of human emotions on the emergency evacuation process in several scenarios, and Liu et al. [26], which proposes the inclusion of psychological factors in earthquake behaviour simulation. Different studies focused on the analysis of specific community reactions to earthquakes. For instance, we can find ABM models that include behavioural patterns based on the analysis of real video captures with applications to urban scenarios of different sizes (D'Orazio et al. [27], Bernardini et al. [28], Quagliarini et al. [29], Santos-Reyes and Gouzeva [30], Bernardini et al. [31]). These works report how, in general, after the decision to start the evacuation procedure, pedestrians are essentially attracted by safe areas and, when moving, they try to avoid the area with rubble choosing the path that is wider and clearer of dust and rubble. These actions are influenced by information seeking and sharing activated by people during and after the seismic event. The "herd behaviour" and the presence of ties between individuals are also identified as characteristics of outdoor urban post-earthquake evacuation. Another relevant contribution is the study of Lu et al. [32], which experimentally analyzes how the amount of debris on the travelled path affects the pedestrian walking velocity. Furthermore, safety procedures and recommended behaviours are identified as important additional contributions that can influence the behaviour and outcome of an earthquake evacuation [31]. However, most of these works involve specific case studies and peculiar conditions, often at the small-scale, and not easily transferable to large urban scale simulations.

Since the disadvantage of ABMs is the increase of the computational burden, especially when the simulation involves large crowds, several works addressed the development of real-time applications to direct evacuation or crowd management operations more effectively. Some of these techniques leverage emerging technologies and informative sources to provide a timely response in an efficient way. For instance, Yin et al. [33] developed a knowledge database generated by mobile phone location data to store evacuation plans for typical urban population distributions and accelerate the real-time search of near-optimal evacuation plans in a real emergency. The research works by Kunwar and Johansson [34] and Kunwar et al. [35] combined crowd-sourced spatial databases (i.e., OpenStreetMap) and behavioral models to achieve rapid simulations of large-scale vehicle evacuations. Finally, Li and Zheng [36] put forward a multi-agent-based concept of dynamic calculation and online feedback based on the Internet of Things technologies to improve real-time evacuation strategies of urban traffic congestion.

In this paper, we present IdealCity, an agent-based approach aimed at simulating emergency evacuation at the urban level for post-earthquake scenarios. This framework is based on PEOPLES (Cimellaro et al. [37], Kammouh et al. [38]), a hybrid and hierarchical conceptual framework for defining the disaster resilience of communities at various scales. IdealCity couples the models of the transportation network and the built environment with an ABM used to simulate in the urban environment the behavior of individuals assuming different roles. These three elements are defined in individual but interconnected layers that mutually influence each other. Thus, the hybrid characteristics of IdealCity allow, on one side, the estimation of buildings' damage and debris generation caused by the seismic event and, on the other side, the analysis of the consequences of these damages on the other layers of the models. In particular, agents could be killed, injured, or trapped inside damaged buildings. Furthermore, roads could be interrupted, limiting the intervention of emergency operators (e.g., ambulances) and blocking the escape routes for citizens. The proposed hybrid model also implements specific features to reproduce the human behaviour of exposed individuals. In particular, following existing literature works, we considered the effect on pedestrian evacuation speed of various elements such as agent age, emotions, health state, fatigue and the presence of debris ([39, 40, 41, 42, 32, 43]).

The developed model can be tested with different seismic scenarios to investigate the effect of human behavior on the large scale evacuation simulation. Then, including in the model shelters, hospitals, and ambulances (each of which can be associated with specific rules within the ABM environment), we can generate project plans for improving emergency evacuation by studying the community response and the suitability of the involved resources. We underline that the aim of our project was to develop a practical and useful tool for evacuation planning. Therefore, the design of IdealCity strictly involved several specialists, domain experts and stakeholders to identify the input and output variables of the model, define their values and detail the emergency procedures involved in the simulations.

Summarizing, the main features of our work are the following. First, it is a general framework that allows the application of its hybrid model (and the simulation of a variety of seismic events based on established engineering criteria) to different urban communities by simply providing the appropriate input. Second, it takes into account the mutual interaction between different model layers (i.e., the agents, the built environment and the generated debris, and the road network). Third, it includes earthquake behaviours and conditions of exposed individuals, in terms of emotion in human behavior, different levels of agent health (obtained as results of the seismic damage simulation), and the effect of of debris on evacuees' walking speed. Fourth, the simulation of the emergency response network (ambulances, hospitals, shelters) allows decision-makers to estimate critical response parameters at the community level and identify the need of additional resources to manage specific events with objective performance metrics. Finally, it can manage in near real-time large-scale simulations with a high number of individuals (about 900,000 in our tests).

The rest of the paper is organized as follows. In Section 2, we detail the assessment process of the aftermath of seismic events, and, in Section 3, we describe the design of our ABM model. Then, in Section 4, we introduce the implementation details of IdealCity and define all the parameters used in our simulations, whose results are presented and discussed in Section 5. Finally, we draw the conclusions in Section 6.

2. Building damage assessment in the IdealCity model

The ability to predict the expected damage of an earthquake in existing structures is a critical step of paramount relevance not only for evacuation modeling but also for seismic risk assessment, emergency response planning, and risk mitigation. Within IdealCity, we characterize the global seismic behavior of thousands of multi-floor buildings through a *surrogate* model created by simplifying the representation of the real physical system (Domaneschi et al. [44], Marasco et al. [45]). Compared to a refined Finite Element Model, this approach reduces the computational burden by limiting the degrees of freedom to manage, while still predicting the structural seismic response with a reasonable accuracy. The output of our model is the expected per-building damage level (classified, according to [46], into five levels as none, slight, moderate, extensive, or complete damage), the amount and extension of generated debris, and their estimated interaction with the transportation network.

In detail, our model is equivalent to a Single Degree Of Freedom (SDOF) system. For each building, we define a backbone curve representative of the global structural capacity and a hysteretic behavior to account for the shear strength degradation due to cyclic loading. The model adopts as response parameters the roof displacement and the base shear force.

The SDOF's equivalent mass (m_{eq}) is representative of the building one (m), assumed as lumped-mass system. Thus, the equivalent mass is defined as $m_{eq} = \Gamma_i m \phi_i$, where Γ_i and ϕ_i are the modal participation factor of the i-th mode and the correspondent natural vibration mode, respectively. The SDOF's stiffness is the equivalent linear elastic one, followed by post-elastic behavior, where the Rayleigh formulation is used to evaluate the SDOF's equivalent damping.

We define the capacity curve of each reinforced concrete building through a four-linear backbone curve to accurately capture the multi-faced degrading stiffness (from the concrete cracking to the yield point) up to the perfectly plastic response. Conversely, we characterize masonry buildings with more regular post-elastic behavior modeling their capacity with tri-linear backbone curves. A comprehensive description of the surrogate model including validation and calibration details can be found in (Domaneschi et al. [44], Marasco et al. [45]).

The debris generation and extension are estimated through a machine learning (ML) approach [47, 48, 49] that works as follows. First, we created a training data-set of about 200 images, collected from existing reports of seismic evaluation missions (e.g., EERI [50]), which show the debris visibly and measurably (see Figure 1 for an example). For each image, we computed the debris extension applying photogrammetric approaches that use the dimension of known objects as reference. The images are then associated with the normalized debris extension (i.e., the debris extension divided by the building height) and various metadata that characterize the building. These metadata are used by the inference engine to predict the debris extension of unseen buildings for a the given earthquake and include the following features:

- the building materials (i.e., reinforced concrete or masonry), since they result in different behaviors in terms of debris extension;
- building height (which influences the seismic forces and, consequently, the damage level experienced by the building);
- number of floors (taller buildings could produce a larger amount and extension of debris);
- year of construction (which characterizes the differences in construction technologies, methodologies, and regulations);
- seismic magnitude (the stronger the ground motion, the higher the damage probability);
- distance from the epicenter (buildings closer to epicenter suffer stronger ground motions).

Concerning the classifier, we analyzed different approaches assessing their prediction accuracy in terms of R^2 and Mean Absolute Relative Distance (MARD). R^2 is the square of the sample correlation coefficient between the predicted values and the ground truth. MARD is the average distance between each point and the regression line computed by the ML algorithm. Therefore, the lower the MARD, the more accurate the prediction. We trained all the classifiers by K-fold Cross-Validation, with 80% of the images in the training set, and equally dividing the remaining samples between validation and test sets.

The results of our tests showed that Random Forests (RF) and k-Nearest Neighbours (k-NN) are the best performers. In particular, k-NN has the highest MARD score (0.32). However, its low R^2 (0.42) is a possible overfitting indicator. On the contrary, RF achieves the best R^2 (0.62) along with a satisfactory MARD (0.56). Thus, it was selected as the classifier to use in our model after retraining it with the whole set of train samples.

Figure 1: Sample picture from the seismic reports data-set [51]

3. The Agent-Based Model of IdealCity

IdealCity has been conceived as a general framework that allows the simulation of towns with different characteristics. In order to enable this feature, its ABM relies on several input parameters (such as the population age and gender distributions, the number of resources available and their capacities), that can be tailored by the designers to customize their simulated environment.

The ABM designed for our system includes two classes of agents, namely the *individuals* living in IdealCity, and the *ambulances*, which rescue seriously injured persons and transport them to the hospitals. The design is then completed by entities representing the city *buildings*, the emergency *shelters* and the *hospitals* (Figure 2). Buildings belong to different categories (i.e., residential, industrial, and schools) and can contain a number of individuals that vary with the time of the day. Each emergency shelter has a specific and fixed capacity, beyond which it cannot host new individuals. The same happens with the hospitals, except for those that can rapidly deploy a field hospital for disaster victims. In this case, we assume that the hospital has infinite capacity. In our simulations, there is always at least one field hospital available to guarantee that all injured individuals can eventually get medical help. The resources available at hospitals and shelters are managed by the *control room*, where the operators of the emergency team monitor the available information, make decisions about resource management, and communicate with the emergency response team.

The rules and behaviors implemented by our ABM individuals are the following. We divide the population into primary groups (e.g., residential, commercial, and industrial population, whose percentages are a parameter of the model), and we define the indoor and outdoor distributions of each group as a function of the time of the day according to the method defined in Capozzo et al. [52] and HAZUS [46]. As a result of the earthquake, we randomly select the health status of each indoor individual among *healthy*, *lightly injured* (who preserve her/his walking capabilities), *severely injured* (who is not able to walk independently and needs to be rescued by ambulance), or *dead* according to statistical distributions associated with the damage level of the building they are inside. Then, we place outdoor persons in random positions along IdealCity, and their health state depends on their proximity to the debris generated by collapsing buildings.

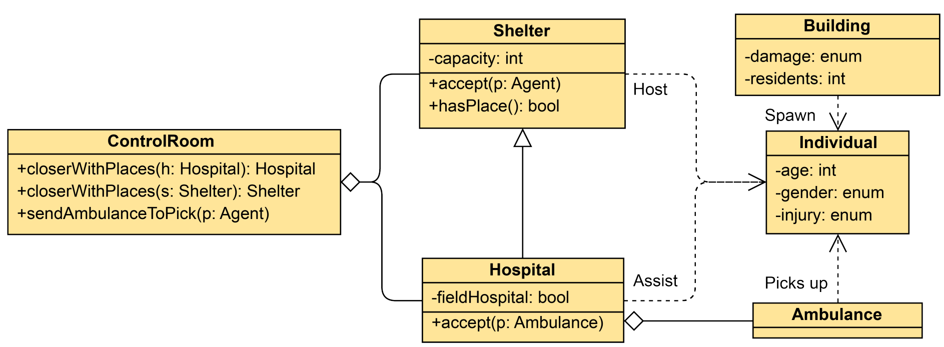

Figure 2: UML class diagram of our ABM model.

Age and gender of each citizen are selected randomly according to the distributions of the population of IdealCity provided as input of the model. The walking speed of each individual takes into account the effects of the percentage of debris coverage on the streets, as proposed in Lu et al. [32]. In details, the initial pedestrian velocity without debris v_{p_0} is selected randomly (with a uniform distribution) in the interval between a comfortable and a maximum gait speed. These two values are taken from the tables reported in the study of Bohannon [39] as a function of age and gender. These random speeds aim at modeling the person's emotional state (the higher the emotion intensity, the higher the speed [40, 41, 42]). If the individual is lightly injured, v_{p_0} is eventually reduced by a variable random percentage in the interval $[40\%, 80\%]$ to simulate injuries that affect in different ways its walking capabilities. During the simulation, the moving speed is decreased as an effect of fatigue until it returns to the comfortable gait speed in a variable time between ten and twenty minutes according to the agent age, injury level, and velocity. The actual walking speed v_p of the agent is then expressed as $v_p = R \cdot v_{p_0}$, where $R \in [0, 1]$ represents the reduction coefficient computed, according to the equations proposed in [32], as a function of the percentage of debris covering the walking path (a value computed by our building damage assessment module). On the contrary, a seriously injured person is unable to walk and needs to be rescued by the emergency medical service, which is activated by the agent itself or by one of its neighbors to request an ambulance. All outdoor pedestrians are spawned at the simulation start. On the contrary, as suggested in [53, 54], indoor persons are spawned from buildings at linear flow rates. In details, the model generates a new exiting individual at random intervals, defined by a Gaussian distribution whose mean is related to the building damage level (the higher the damage, the higher the mean).

The behavior of indoor individuals takes into account the person health state and the damage level of the building and allows to model different subjective responses and other constraints, such as official regulations or recommendations. Healthy individuals can decide to either remain in the close vicinity of the building or reach the nearest emergency shelter according to a random variable that can be parameterized as a function of the building damage level and other simulation choices. This feature aims to provide designers with different options. For example, it allows simulating persons that take individual decisions (e.g., 30% of the healthy residents of slightly damage buildings and 95% of the moderately damaged ones head to a shelter, while the other healthy individuals remain in place) or situations where all residents follow the official recommendation to evacuate. In all cases, individuals remaining in the close vicinity of the building are removed from the simulation. When the person is lightly injured, it walks towards the nearest hospital. As we already stated, seriously injured individuals wait (in place) to be rescued by the emergency service. The behavior of outdoor persons mirrors that of indoor ones. The only exception is represented by healthy individuals, which navigate to the closest emergency shelter irrespective of the state of their homebuilding.

Concerning navigation, we assume that all individuals have a shared knowledge of the locations of nearby hospitals and emergency shelters (in other words, we assume that at least one person knows these pieces of information and shares them with its neighboring peers). When an individual reaches a full target destination, it is directed by the facility personnel to the closest destination of the same type that has yet available places (we suppose that the facility managers are in direct contact with the control room, which knows the overall place availability). If all shelters are full, the individual starts walking toward the closest town exit. Since our simulation includes at least a field hospital, all injured persons are eventually guaranteed to find assistance.

The ambulances are initially distributed among various locations in the city and are in a finite number. The control room is responsible for their management and follows a first-in-first-out practice for processing assistance requests. If an ambulance is available, it is sent to pick a seriously injured inhabitant; otherwise, the request is queued in a waiting list. When the ambulance reaches its target, the paramedical personnel starts loading the patient in the vehicle. This operation requires a random variable time in the interval $[15, 20]$ minutes, after which the ambulance heads to the nearest hospital with available places. If the target hospital gets full during the travel, the ambulance receives from the control room the indication to head towards a new target destination. At arrival, the ambulance has to wait for a random time before it can finish unloading the patient. The waiting time is chosen in the interval $[5, 15]$ minutes with a logarithmic distribution that models the effect of hospital occupancy on service time (the higher the occupancy, the slower the service time). Loading and unloading intervals have been defined after a discussion with medical experts. Finally, when the hospital has accepted the patient, the ambulance becomes available and can serve a new assistance request, if any.

Concluding, the implemented agent model considers both emergency network management and pedestrian evacuation, including some of the typical behaviours of exposed people in outdoor scenarios reported in the literature ([27, 31, 30, 32]). In our model, agents run out of the building and, after the decision to start the evacuation procedure, they are attracted towards safe areas (i.e., shelters and hospitals). In choosing their paths, they seek and exchange information, and in general, they tend to aggregate during the evacuation portraying the typical herd effect. Furthermore, the proposed model allows the interaction of the agents with the generated debris, reproducing interruptions of evacuating paths and the reduction of the evacuation speed where a moderate amount of debris is present. That said, we highlight as well that our model could be improved by taking into account more refined approaches (such as those already tested at smaller scales in [25, 13]) and additional evacuee behaviours (such as the tendency to perform evacuation path selection according to the visible building damage and geometric dimensions of the street, or the attraction for group ties [27]). The reason why we did not consider these features in our work is the heavy computational burden they involve when applied to large-scale simulations, resulting in computation resource issues and an excessive increase of the simulation time. Therefore, we deemed preferable to tackle them as future works.

4. Experimental settings

In this section, we detail the implementation of IdealCity and define all the parameters used in our simulations.

4.1. City model

Concerning the ecological validity of our simulations, the city model used in the experiments reflects the characteristics (i.e., road transportation network, building archetype, year of construction, height classifications, and occupation) of the town of Turin in Italy. Its urban area is about 130 km^2, and the city has more than 900.000 inhabitants. The building stock is representative of a typical Italian one, including a downtown area mainly composed of old masonry buildings and uniform housing zones of reinforced concrete. In the last decade, the residential zones expanded on the grounds of many abandoned factories in the suburbs (included in our model).

The city has 16 hospitals in operation, with an overall capacity in emergencies of about 6,000 places. We selected the San Giovanni Battista university hospital as the only one capable of deploying a field hospital since (i) it is the fourth largest public hospital in Italy, (ii) it is part of the biggest hospital district of Turin, which includes the primary children's hospital, the main gynecological hospital, and the orthopedic trauma center of the city and (iii) it is close to a large urban park (thus, it has room for deploying a large number of tents). The official shelters identified by the National Service of Civil Protection are 31, and they can assist more than 100,000 people.

As for the population, the division in groups (e.g., residential, commercial and industrial population) and their age and gender distributions (see Figure 3) reflects the corresponding distributions of the city population provided by the annual census report released by the Italian National Institute of Statistics [55].

Given the complexity of the model and the number of potentially active agents, we made the following assumptions to simplify our simulations: residents will immediately start moving from the building when the seismic event occurs, all individuals move by walking without using vehicles, there are no disabled people in the buildings, and no altruistic behaviors are considered. In the main simulations, the residents' behaviour models a case where all healthy agents remain in the proximity of their building when it is slightly or not damaged and walks toward a shelter if the damage level is moderate or higher. In Section 5.3, we will discuss how different behavioural models affect the simulations.

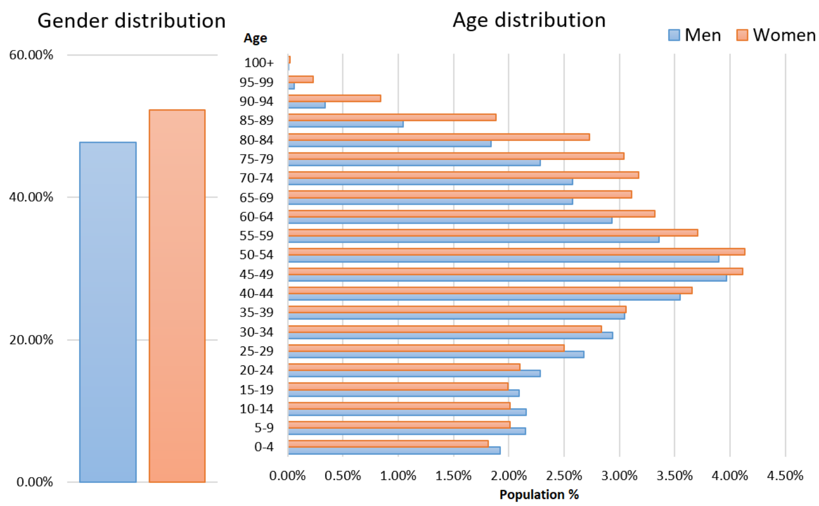

Figure 3: Left: distribution per gender of the whole population of IdealCity used in the experiments. Right: age distributions of the population per gender.

4.2. ABM: implementation details

To optimize the execution times, our ABM has been implemented in Unity, a 3D game engine, leveraging the Entity-Component-System (ECS) design pattern and the Unity Job System. Within the ECS paradigm, agents are represented by an *entity*, consisting of one or more *components*. These components are pure data than can be added, removed, or modified by the systems at run-time, and they are used to define one specific aspect of the entity. As an example, a `Transform` component defines the current position of the agent and the `Rotation` component its orientation. Empty components can also be used to tag entities. For instance, in our implementation, the `DestinationReached` component marks any agent that is about to reach its target destination (e.g., the patient for an ambulance, the emergency shelter for an individual). Finally, the *systems* are the elements responsible for implementing specific behaviors or functionalities. Each system operates on a subset of entities defined by specifying the components that entities must possess or excluding entities that have specific components. For instance, we have two different systems to check if an agent has reached its destination: one for the ambulances and one for the pedestrians. The two systems work on entities having a `DestinationReached` component (which is used as a "tag"), but the first manages only entities with the `IndividualComponent` and the second only those having an `AmbulanceComponent`.

The strict separation between data and behaviors in ECS allows exploiting at best the multiple cores available in today's computers thanks to the Unity C# Job System, which is a powerful multithreading API expressly developed for maximizing the execution performances. In this way, it is possible to speed up significantly all the complex tasks that can be executed independently for each agent, achieving huge performance gains. As another advantage, this data-oriented paradigm allows a better usage of memory, which is a critical factor in our simulation, especially for managing pathfinding on a massive scale (i.e., hundreds of thousands of agents).

4.3. Seismic scenarios

In this study, we adopt four different seismic scenarios that are characterized by their epicenter location (the point on the earth's surface vertically above the focus, or hypocenter, of an earthquake), the moment magnitude M_w (directly related to the earthquake energy), and their acceleration time series recorded at the epicenter. We use the Attenuation Ground Motion Prediction Equation (Ambraseys et al. [56]) to estimate the spatial variability of ground motion, i.e., the effect of geometrical attenuation due to the distance between the building location and the epicenter. In our simulations, we also assume that the frequency contents for each time series remains unchanged at any location. Table 1 summarizes the main characteristics of the seismic events selected, showing, for each event, the date, the moment magnitude M_w, the hypocentral depth, and the Peak Ground Acceleration (PGA, which is directly related to the seismic forces).

Event	**Date**	M_w	**Depth [km]**	**PGA [g]**
Northridge	1994-01-17	6.7	11.3	0.84
Kobe	1995-01-17	6.8	17.6	0.82
EL Centro	1940-05-18	6.9	16.0	0.35
Hachinohe	1968-05-16	8.2	26.0	0.23

Table 1: Characteristics of the selected seismic scenarios

The events in Table 1 fall into two groups with different characteristics. The first group includes the Northridge event (recorded at County Hospital parking lot in Sylmar, California) and the Kobe event (recorded at the Japanese Meteorological Agency station, Japan). Both events are representative of the category of near-field earthquakes, i.e., those which occur at a distance close to the fault (about 10 to 60 km). Near-field earthquakes are characterized by higher acceleration and more limited frequencies compared with far-field earthquakes. Furthermore, they may undergo progressive direction effects, containing long-pulse periods of high ranges. On the contrary, the second group covers the far-field seismic scenarios and includes the El Centro (Imperial Valley Irrigation District substation, California) and the Hachinohe (Hachinohe City, Japan) events.

4.4. Methodology

We run three experiments for each of the four scenarios described in the previous section. The input of these simulations is the results of the seismic aftermaths (building damages and debris generated) computed with the methods described in Section 2. As an example, in Figure 4, we show an overview of the Kobe scenario, where we use false colors to identify the different damage levels of each building. Each experiment models a seismic event happening at a specific time of the day (namely, 2 AM, 2 PM, and 5 PM). The event time affects the indoor and outdoor distributions of the residents of IdealCity. As a result, we have different percentages of the dead, seriously/lightly injured, and trapped individuals according to their specific location inside the city when the event happened (see Figure 5 for an example). Some snapshot taken from a running simulation can be seen in Figure 7.

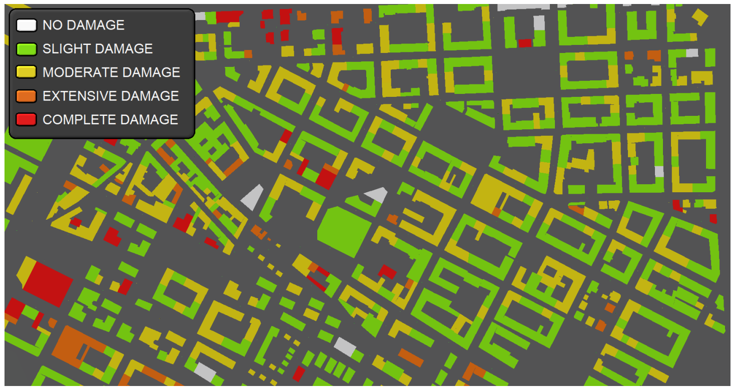

Figure 4: Detail of the IdealCity plan with damage assessment.

For the quantitative assessment of the model, we recorded fine-grained logs for each session (one block of analytic data each minute of simulated time). These data include the average execution time and several other pieces of information, such as the current number of active individuals (i.e., those moving around the city to reach their destination or waiting for rescue) and the secured ones (i.e., those that reached their destination or decided to remain in proximity of their houses when possible), the occupancy of shelters and hospitals, and the number of seriously injured individuals that did not yet receive assistance.

We highlight that, despite the use of various random variables in our model, our approach can be considered as a deterministic one, even running a single simulation for each of the considered scenarios. Indeed, our initial experiments showed, among different runs, negligible statistical differences in the main simulation outputs (e.g., a coefficient of variation lower than 2% for the time to secure injured individuals, and for the occupation and saturation times of shelters and hospitals).

4.5. Input and output data

To conclude this section, we provide an overview of the numerical input data used in the model and of the outputs offered for the emergency assessment, which are summarized in the block diagram in Figure 6. We recall that IdealCity has been envisioned as a general framework and, as such, it can simulate towns with different characteristics and various seismic events by varying its input data and model parameters.

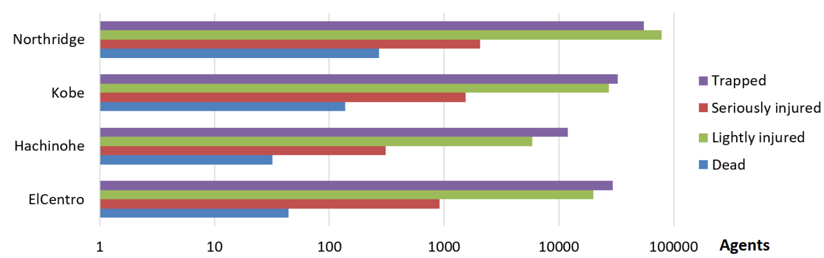

Figure 5: Summary of the simulations at 2 AM for all scenarios (data are represented with a logarithmic scale).

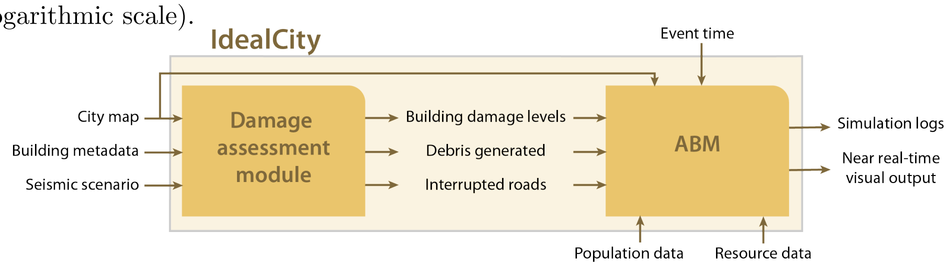

Figure 6: IdealCity: overview of input and output data.

The building damage assessment module receives as input the city map (which defines the transportation network and the built environment), the building metadata defined in Section 2 and the characteristics of the seismic scenario (Section 4.3). This module outputs the per-building damage level, the debris extension and the list of interrupted roads. These data are then fed as input of the ABM module, along with the city map, the population data (i.e., the age, gender, primary group and indoor/outdoor distributions), the resource data (i.e., the number and capacities of shelters and hospitals and the number of available ambulances) and the time of the day when the seismic event happens.

Then, the framework provides two different types of outputs for the emergency simulation assessment. The first is the set of detailed logs collected during the simulation, which can be further analyzed to extract meaningful information (as we will show in Section 5). The second is the near real-time visualization offered by the Unity rendering engine, which offers the interactive navigation of the city map and zooming features, allowing for a detailed analysis of the simulation evolution.

5. Results and discussion

In this section, we present the results of the assessment of IdealCity. First, we analyze the computational load of the application (Section 5.1). Then, we show how IdealCity can be used to study the behavior of both the population and the emergency response network in different seismic events (Section 5.2). Finally, we illustrate the features that IdealCity offers to decision-makers for the design of efficient resource plans aimed at a prompt and efficient response to emergency events (Section 5.3).

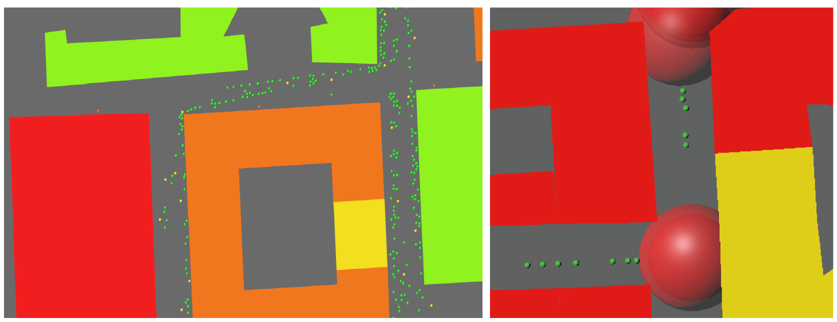

Figure 7: Snapshots of a post-earthquake scenario. Left: overview of a city district with various individuals trying to reach their target destination (green: healthy persons; yellow: lightly injured; red: seriously injured). Right: agents trapped since the debris (the semi-transparent spheres) made all the exit roads from a block impassable.

5.1. Computational load

We run experiments on a Windows workstation with an Intel Xeon Gold 5122 CPU and 192 GB of RAM. The CPU has four cores (and eight threads) running at 3.60 GHz. The amount of available RAM is a vital parameter in our case since, when the number of active agents is very high, the allocated memory quickly reaches tens of GB. Thus, if the size of the RAM is not sufficient, such numbers could result in continuous disk swaps and, consequently, in large performance drops.

The analysis of the behavior of the computational load during the simulation can provide some interesting insights. To perform such analysis, we can plot the ratio between the CPU time and the simulation time, where the latter is controlled by a simulation clock whose tick is generated at each update of the main rendering cycle. In our tests, we arbitrarily created a simulation tick at intervals of 33 ms. However, we observed that the frequency of the simulation clock can be made smaller (up to a factor 10) without causing issues to the stability of the ABM. Since it is clear that the CPU time is also a function of the number of active agents, before drawing the plot, we further normalize the time ratio on a number of 100,000 agents (i.e., we divide it for the ratio between the total number of agents and 100,000).

As an example, we show in Figure 8 the differences between two simulations run in the same scenario (Northridge) at two different daytimes. The first simulation (2 AM) corresponds to the case having the lowest number of outdoor agents at the beginning of the simulation, and the latter (5 PM) is the one with the highest number of outdoor agents. As can be seen, in both tests, the plot stabilizes after about 45 (simulated) minutes on a value ranging between 2.5 and 2.8 for the two cases. Thus, when the simulation clock is set to 33 ms, our simulation can run in real-time when the number of agents is around 40,000 and runs in near-real-time for more substantial values (e.g., for 400,000 agents, one second of simulation requires between 10 and 11.2 CPU seconds) and improvements can be obtained increasing (up to a factor ten) the simulation clock interval. Another thing that can be observed is the significant spike of the time ratio at the beginning of the 5 PM simulation. This effect is due to the computational burden related to the pathfinding routine since, at 5 PM, the initial number of outdoor agents is three orders of magnitudes bigger than that in the 2 AM simulation. Then, when the initial wave of path requests has been served, which happens in a few minutes, the behavior of the two simulations becomes similar.

In Figure 8, we also plot the number of active agents managed by the two simulations, which reaches a maximum of about 600,000 agents and then slowly degrades due to the increased number of secured agents (i.e., those that reached their target destination, a shelter or a hospital, or left the city). We underline that individuals that decide to stay in proximity to their house when it is not damaged are immediately removed from the simulation to reduce the overall computational burden. It should also be noted that the city of Turin can be enclosed in a circumference of about 14 km in diameter. Therefore, crossing the city from side to side at an average walking speed requires less than three hours. As a consequence, the average time required by healthy or lightly injured individuals to reach their destination is relatively small since the spatial location of resources (hospitals and shelters) tries to guarantee uniform access times to the whole population.

We also underline that IdealCity lacks a detailed model of the vehicular network since only the ambulances have been managed. This is a limitation of the model. In a real scenario, people are likely to take their car for reaching a secure place or for leaving the city. The consequence would be a dramatic traffic congestion that has negative consequences for the delay of emergency services and the increase of the travel time of the city dwellers. However, modeling the vehicular traffic along with the other agents in the current experimental settings (a city of 900,000 dwellers) would have resulted in a large increase of the computational burden, thus strongly reducing the near real-time capabilities of our simulator. It is worth underlining that the evacuation models in the literature either model one single class of agents (pedestrian or vehicle, e.g. [22, 28, 33, 57]) or limit the urban area to a much smaller region (i.e. districts, squares or single blocks, e.g. [58, 57]).

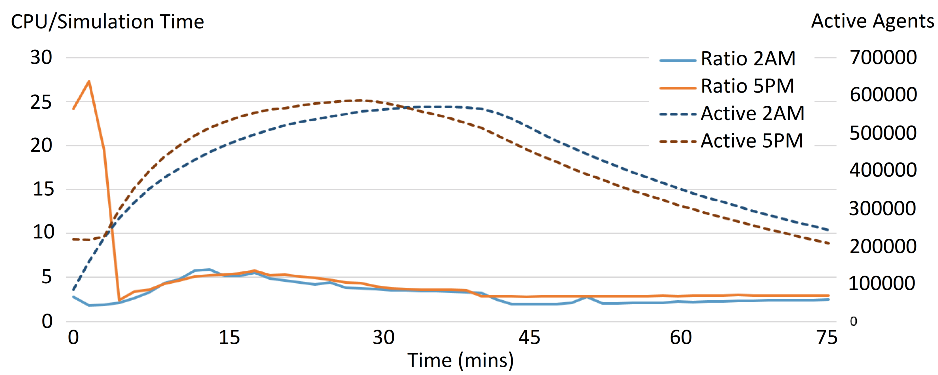

Figure 8: Solid lines: ratio between the CPU time and the simulation time for the same scenario (Northridge) and two different times of the day (2 AM and 5 PM) normalized over 100,000 agents. Dashed lines: number of active agents. The horizontal axis represents the simulation time in minutes.

5.2. Analysis of the population response

A possible way to analyze the population response in the emergency evacuation is to observe the temporal evolution of some of the simulation variables. Examples are the number of secured individuals (i.e., those that reached their final destination) and the number of severely and lightly injured able to reach the hospital to receive treatments. For the sake of brevity, in Figure 9, we show the data collected from the different seismic scenarios and the same time of the day (i.e., 2 AM, the simulations with the lowest number of individuals in the streets when the event happens). The overall number of dead, seriously and lightly injured, and trapped persons of these events is included as a reference in Figure 5.

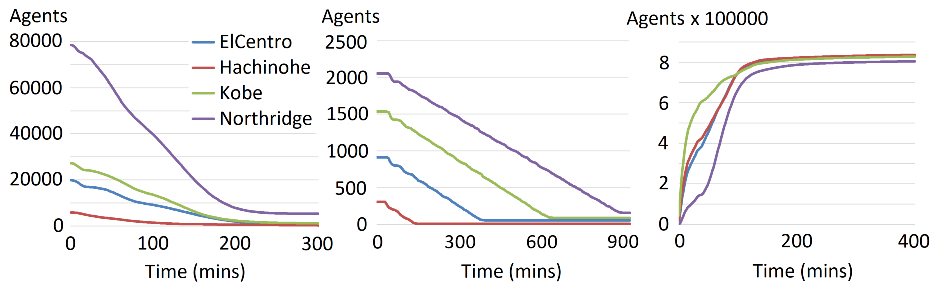

Figure 9: Numbers of lightly injured individuals traveling to hospitals (left), severely injured that are waiting to be rescued (center) and agents that reached a secure place (right) for different seismic events happening at 2 AM.

These diagrams allow us to gain some interesting insights into the simulations. For instance, if we observe the number of secured individuals (Figure 9, right), we can see that in about three hours and a half after the event, most of the citizens reached a safe place. Given the average distances that individuals have to travel to their destinations, these numbers are reasonable. Indeed, a large part of the pedestrians reach a secure place in about half an hour. The remaining ones continue their journey since their first destination had no places to host them.

The number of available resources affects all emergency evacuation scenarios critically. As we can observe in Figure 10, in all the simulations, all the shelters are filled in less than one hour and a half. Therefore, the remaining healthy individuals have no other choice than leaving the city, which (as stated before) requires at worse about three hours of walk. As for the lightly injured, in all scenarios (exception made for Hachinohe, whose light injured are less than the places available), hospital nominal capacity starts saturating in less than one hour. Thus, wounded people start moving towards the field hospital that eventually guarantees a place for all injured persons that can not find assistance in other facilities. However (as can be seen in Figure 9, left, which shows a drastic decrease of the number of injured accepted in the hospital after about 180 minutes, reaching the field hospital requires longer travel times with possible negative consequences on the agents' health. We should also observe that the time required to rescue all the severely injured can become extremely high, reaching up to 15 hours in the Northridge event (Figure 9, center). Such delays in assistance can lead to an increase in the number of deaths. Summarizing, these observations highlight that the emergency response resources deployed in the simulated city are quite inadequate for a prompt response to extreme events.

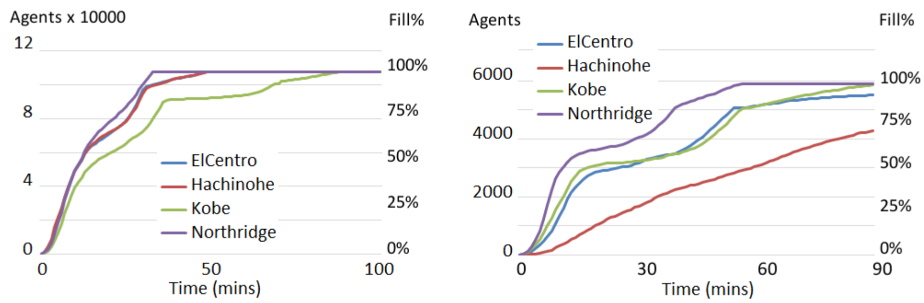

Figure 10: Occupation of Shelters (left) and Hospitals (right) for every scenario, as absolute values and as percentages, relative to the nominal maximum capacity.

Another way to use our simulator is to compare the overall emergency response when the same seismic event happens at different times of the day. As an example, we show the comparison of the response simulation of the Northridge event (i.e., the most damaging one) happening at 2 AM and 5 PM, which are the situations with, respectively, the lowest and highest number of outdoor agents. We recall that different indoor/outdoor distributions result in different distributions of the overall individuals' health levels since indoor persons are more likely to be affected by the consequences of the earthquake.

In Figure 11, we can observe that the shelters' saturation time and the time required to host (almost) all the lightly injured agents are similar in the two simulations. In both cases, the differences between the first parts of the curves can be explained in terms of the different number of indoor agents (which result in different number of injured persons). The healthy street agents can immediately start looking for rescue, while indoor agents must first exit their building, which requires a specific time. As for the lightly injured, their number is higher at 2 AM, which explains the peak in the graph, but the overall response of the medical and sheltering service is similar in the two simulations. Concluding, the major takeaway from these results (which are consistent with that of other combinations of scenario and event time) is again that the primary variable affecting the emergency system response is the number of available resources.

5.3. IdealCity for resource planning

The issues found in the previous sections highlight the need, in these simulations and many other real cases, to estimate and intervene on critical

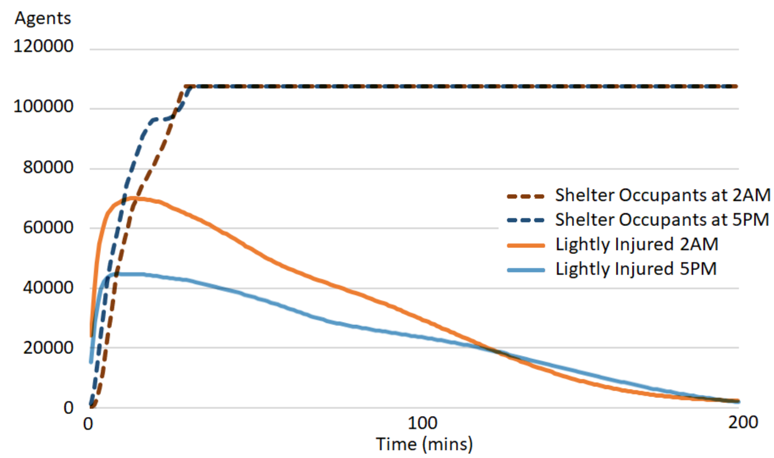

Figure 11: Comparison between the same scenario (Northridge) at 2 AM and 5 PM. The dashed lines represent the number of agents that safely reached a shelter. The solid lines show the number of lightly injured people that are currently traveling to reach an hospital.

parameters such as the number of resources available. IdealCity can provide a contribution to tackle this problem since decision-makers can use the model to prepare resource plans aimed at guaranteeing a certain level of services during the emergency response.

To exemplify the use of IdealCity as a resource planner, we run different simulations on the Northridge event at 2 AM, increasing the number of available resources and then analyzing the response of the simulation variables. In this way, we can estimate the resources required to guarantee a target quality level of the response service.

For instance, Figure 12 shows the time needed to rescue all seriously injured agents as a function of the number of ambulances available. The initial value of about 15 hours for 110 ambulances can be reduced to about two hours when 1,000 ambulances are made available. On the contrary, if we set an arbitrary threshold as a target (e.g., 240 minutes), we can compute the number of ambulances needed (447) to complete patient recovery within that time.

Similar reasoning can be made on other resources. Increasing the overall percentage capacity of shelters and hospitals, we can obtain two pieces of information. The first is the (trivial) indication of the overall number of persons that can be secured. The second (Figure 13) is the time to saturate shelters and (nominal capacity of) hospitals. This value is relevant since it is

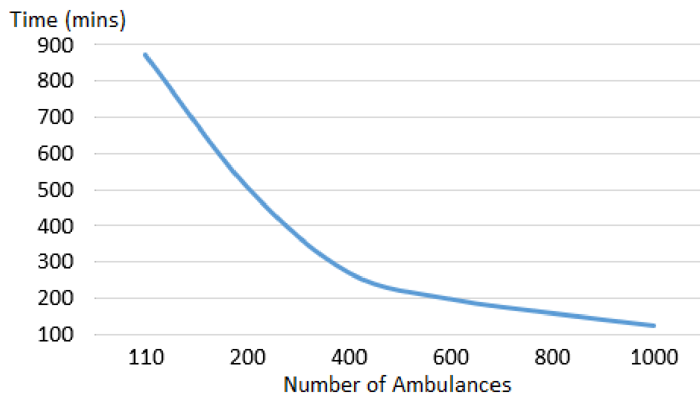

Figure 12: Northridge event at 2 AM: time required by the ambulance service to bring all seriously injured agents to an hospital as a function of the number of available vehicles.

clear that, by definition, an event is a disaster when it exceeds the capacity of the emergency response system. Thus, the saturation times can be used by the planners to define a suitable trade-off between the number of resources deployed and the amount of time available to put in place external resources.

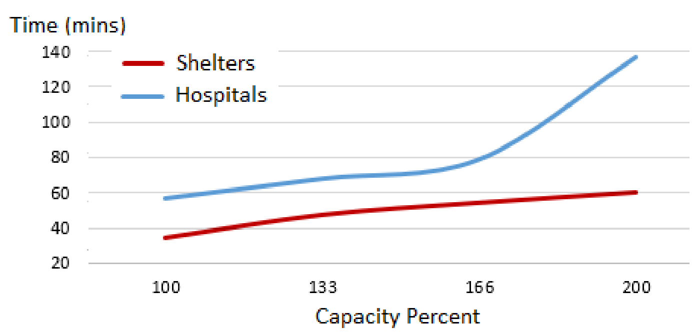

Figure 13: Northridge event at 2 AM: time to reach the nominal capacity of shelters (red line) and hospitals (blue line) when their percentage capacity increases.

Another parameter of IdealCity is the behaviour of indoor individuals. As we stated at the beginning of Section 4.5, the simulations described so far consider a case where healthy indoor agents stay home, and those of damaged buildings walk toward the nearest shelter (this case is referred in the following as B0). In order to analyze how different settings affect the experimental outcomes, we run three other simulations where indoor agents can take different choices according to the following parameters. In simulation B1, all healthy residents follow the official recommendation to head to a shelter. In simulation B2, irrespective of the building damage level, 50% of the residents remain in place, the other head to a shelter. In simulation B3, the average percentage of residents remaining close to the building varies according to the building damage level (namely, 95%, 75%, 50%, 25% and 5% for, respectively, not damaged and slightly, moderately, extensively and completely damaged buildings). We choose as scenario the Hachinohe event at 2 AM since it is of average severity (and, thus, it results on balanced percentages of buildings affected by the different damage levels) and 2 AM is the time of the day with the highest amount of indoor agents.

The results of these simulations are summarized in Figure 14, where we show the time to reach the nominal shelter capacity in the four simulations described. We recall that the changes in the behaviour of the healthy indoor individuals do not affect hospitals, which have to manage the same amount of injured in all the simulations. The results clearly show that B1 (i.e., the simulation with the highest number of individual heading towards a shelter) is the most critical condition since shelters saturate their capacity in a short time (and about 20 to 30 minutes less than the other simulations). This result would highlight the scarcity of resources if the politics of recommending all citizen to reach a secure place were indeed applied.

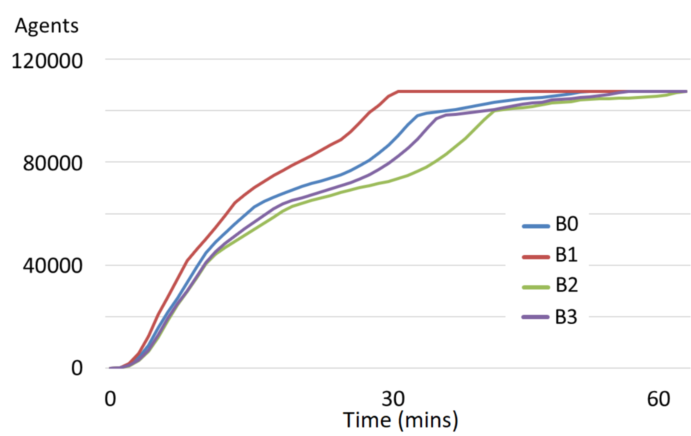

Figure 14: Hachinohe event at 2 AM: effects of different behaviour policies of the healthy indoor individuals on the time to reach the nominal capacity of shelters.

Finally, we underline that IdealCity supports resource planning at the macro and micro levels. While it is relevant to estimate the global effect of the available/planned resources, we can also analyze (at a local scale) where are the most critical locations in terms of emergency response. For instance, in Figure 15, we marked on the IdealCity map the locations of shelters and hospitals using a color-graded scale that is a function of the saturation time of the corresponding facility (red: lower times, green: higher times). Such visualization allows planners to quickly identify the urban district that requires an improvement of the resources needed to face a disaster. For instance, it can be seen from Figure 15 that in the the worst-case scenario (Northridge at 2 AM, i.e., the event with the most significant number of injured individuals) there are hospitals located in densely populated urban areas that are saturated in very short times due to their reduced capacities (from 46 to about 120 places).

Figure 15: Saturation times of the emergency facilities in the Northridge 2 AM simulation (circles: shelters; squares: hospitals).

6. Conclusions

In this work, we have presented IdealCity, a hybrid and multilayered model for simulating and studying the disaster resilience of communities at various scales. The model includes the layers corresponding to the transportation network, the built environment, and an agent-based model capable of simulating human behavior during evacuation. IdealCity allows the estimation of buildings' damage and debris' generation caused by the seismic event as well as their effects on the other layers of the models. Furthermore, it implements the emergency response network in terms of shelters, hospitals, and ambulances.

IdealCity has been tested on a realistic urban environment and with different seismic scenarios. Our experimental results show that the simulations run in near-real time with a large number of agents on a standard PC and that their output can be used to (reliably) analyze the community response. Furthermore, these results show as well that our model can support decision-makers in the design and implementation of specific interventions aimed at improving the inherent community resilience.

Our future works will focus on extending the model validation using different real environments beyond the city of Turin and the four reference seismic events of our simulations. Then, we plan to conjugate the emotional and irrational part of the behavioral model of our agents with different patterns aimed at simulating the interaction among agents (such as altruism [13], the leader-follower phenomenon [59] and pro-social behaviours [31]) as well as specific behaviours of exposed people (such as the path selection described in [31]). Another feature we are planning to introduce is a complete model of vehicular traffic, currently ignored in the proposed ABM model, which will allow increasing the fidelity of our simulations. Finally, we plan to use IdealCity as the core of an immersive Virtual Reality environment for training rescue and civil protection operators in complex and realistic scenarios, without direct risks to individuals and with a reduction of training costs.

7. Acknowledgements

This work was supported by the European Research Council under the Grant Agreement nr. ERC IDEAL RESCUE 637842 of the project IDEAL RESCUE - Integrated Design and Control of Sustainable Communities during Emergencies.

The Disaster Resilience Simulation Laboratory (DRSIL) at Politecnico di Torino DISEG is also gratefully acknowledged.

References

[1] S.-T. Liu, W.-B. Wei, C. Ou, Optimization study on the evacuation mode for the middle train compartment fire of single-track tunnels, Journal of Railway Engineering Society 33 (2016) 91–96.

[2] J.-P. Wang, M.-R. Wang, J.-L. Zhou, Q.-J. Zuo, X.-X. Shi, Simulation based optimal evacuation plan in vertical ship lift: a case study, Engineering Computations (Swansea, Wales) (2020). doi:10.1108/EC-05-2019-0212.

[3] J. Mikulik, W. A. Cempel, S. Kracik, D. Dabal, A simulation model for emergency evacuation time and low-cost improvements of a hospital facility using flexsim healthcare: A case study, in: J. Korbicz, M. Kowal (Eds.), Intelligent Systems in Technical and Medical Diagnostics, Springer Berlin Heidelberg, Berlin, Heidelberg, 2014, pp. 333–342.

[4] Y. Meng, C. Jia, Research and application of metro station evacuation simulation, in: 4th International Conference on Information Science and Control Engineering (ICISCE), 2017, pp. 1123–1125.

[5] T. Urbanik, A. Desrosiers, M. Lindell, C. Schuller, Analysis of techniques for estimating evacuation times for emergency planning zones, Technical Report, Texas Transportation Institute (United States), 1980.

[6] Y. Sheffi, H. Mahmassani, W. B. Powell, A transportation network evacuation model, Transportation research part A: general 16 (1982) 209–218.

[7] J. Johnson Jr, D. J. Zeigler, Modelling evacuation behavior during the three mile island reactor crisis, Socio-Economic Planning Sciences 20 (1986) 165–171.

[8] A. G. Hobeika, B. Jamei, Massvac: A model for calculating evacuation times under natural disasters., Emergency Planning (1985) 23–28.

[9] A. Pal, A. J. Graettinger, M. H. Triche, Emergency evacuation modeling based on geographical information system data, in: CD-ROM Proceedings of the 2003 TRB Annual Meeting, Washington, DC, 2003, pp. 1–16.

[10] E. Bonabeau, Agent-based modeling: Methods and techniques for simulating human systems, Proceedings of the National Academy of Sciences 99 (2002) 7280–7287. doi:10.1073/pnas.082080899.

[11] M. Batty, Cities and complexity: understanding cities with cellular automata, agent-based models, and fractals, The MIT press, 2007.

[12] K. C. Clarke, Cellular Automata and Agent-Based Models, Springer Berlin Heidelberg, Berlin, Heidelberg, 2014, pp. 1217–1233.

[13] G. Cimellaro, S. Mahin, M. Domaneschi, Integrating a human behavior model within an agent-based approach for blasting evacuation, Computer-Aided Civil and Infrastructure Engineering 34 (2019) 3–20.

[14] R. Bordini, J. Hubner, M. Wooldridge, Programming Multi-Agent Systems in AgentSpeak Using Jason, volume 8, 2007. doi:10.1002/9780470061848.

[15] J. Leon, A. March, Urban morphology as a tool for supporting tsunami rapid resilience. a case study of talcahuano, chile, Habitat International 43 (2014) 250–262.

[16] I. Solis, P. Gazmuri, Evaluation of the risk and the evacuation policy in the case of a tsunami in the city of iquique, chile, Natural Hazards 88 (2017) 503–532.

[17] E. Mas, A. Suppasri, F. Imamura, S. Koshimura, Agent-based simulation of the 2011 great east japan earthquake/tsunami evacuation: An integrated model of tsunami inundation and evacuation, Journal of Natural Disaster Science 34 (2012) 41–57. doi:10.2328/jnds.34.41.

[18] M. Li, Y. Zhao, L. He, W. Chen, X. Xu, The parameter calibration and optimization of social force model for the real-life 2013 ya'an earthquake evacuation in china, Safety Science 79 (2015) 243 – 253.

[19] M.-L. Xiao, Y. Chen, M.-J. Yan, L.-Y. Ye, B.-Y. Liu, Simulation of household evacuation in the 2014 ludian earthquake, Bulletin of Earthquake Engineering 14 (2016). doi:10.1007/s10518-016-9887-6.

[20] X. Yang, Z. Wu, Y. Li, Difference between real-life escape panic and mimic exercises in simulated situation with implications to the statistical physics models of emergency evacuation: The 2008 wenchuan earthquake, Physica A: Statistical Mechanics and its Applications 390 (2011) 2375 – 2380.

[21] J. Bañgate, J. Dugdale, E. Beck, C. Adam, Solace a multi-agent model of human behaviour driven by social attachment during seismic crisis, in:

2017 4th International Conference on Information and Communication Technologies for Disaster Management (ICT-DM), 2017, pp. 1–9.

[22] H. Na, A. Banerjee, Agent-based discrete-event simulation model for no-notice natural disaster evacuation planning, Computers and Industrial Engineering 129 (2019) 44–55.

[23] P. M. Torrens, High-resolution space–time processes for agents at the built–human interface of urban earthquakes, International Journal of Geographical Information Science 28 (2014) 964–986.

[24] M. Lujak, S. Ossowski, Intelligent people flow coordination in smart spaces, Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics) 9571 (2016) 34–49.

[25] G. Cimellaro, F. Ozzello, A. Vallero, S. Mahin, B. Shao, Simulating earthquake evacuation using human behavior models, Earthquake Engineering and Structural Dynamics 46 (2017) 985–1002.

[26] T. Liu, Z. Liu, M. Ma, T. Chen, C. Liu, Y. Chai, 3d visual simulation of individual and crowd behavior in earthquake evacuation, SIMULATION 95 (2018) 65–81. URL: `https://doi.org/10.1177/0037549717753294`. doi:10.1177/0037549717753294.

[27] M. D'Orazio, L. Spalazzi, E. Quagliarini, G. Bernardini, Agent-based model for earthquake pedestrians' evacuation in urban outdoor scenarios: Behavioural patterns definition and evacuation paths choice, Safety science 62 (2014) 450–465.

[28] G. Bernardini, M. D'orazio, E. Quagliarini, L. Spalazzi, An agent-based model for earthquake pedestrians' evacuation simulation in urban scenarios, Transportation Research Procedia 2 (2014) 255–263.

[29] E. Quagliarini, G. Bernardini, L. Spalazzi, et al., Epes–earthquake pedestrians evacuation simulator: A tool for predicting earthquake pedestrians evacuation in urban outdoor scenarios, International journal of disaster risk reduction 10 (2014) 153–177.

[30] J. Santos-Reyes, T. Gouzeva, Mexico city's residents emotional and behavioural reactions to the 19 september 2017 earthquake, Environmental research (2020) 109482.

[31] G. Bernardini, R. Lovreglio, E. Quagliarini, Proposing behavior-oriented strategies for earthquake emergency evacuation: a behavioral data analysis from new zealand, italy and japan, Safety science 116 (2019) 295–309.

[32] X. Lu, Z. Yang, G. P. Cimellaro, Z. Xu, Pedestrian evacuation simulation under the scenario with earthquake-induced falling debris, Safety science 114 (2019) 61–71.

[33] L. Yin, J. Chen, H. Zhang, Z. Yang, Q. Wan, L. Ning, J. Hu, Q. Yu, Improving emergency evacuation planning with mobile phone location data, Environment and Planning B: Urban Analytics and City Science (2019).

[34] B. Kunwar, A. Johansson, Measuring disaster preparedness of uk cities from open spatial databases, Traffic and Granular Flow, 2013 (2015) 265–271.

[35] B. Kunwar, F. Simini, A. Johansson, Evacuation time estimate for total pedestrian evacuation using a queuing network model and volunteered geographic information, Physical Review E 93 (2016).

[36] F. Li, B. Zheng, Design of the smart city planning system based on the internet of things, International Journal of Smart Home 10 (2016) 207–218.

[37] G. Cimellaro, C. Renschler, A. Reinhorn, L. Arendt, Peoples: A framework for evaluating resilience, Journal of Structural Engineering (United States) 142 (2016).

[38] O. Kammouh, A. Zamani Noori, G. Cimellaro, S. Mahin, Resilience assessment of urban communities, ASCE-ASME Journal of Risk and Uncertainty in Engineering Systems, Part A: Civil Engineering 5 (2019).

[39] R. W. Bohannon, Comfortable and maximum walking speed of adults aged 20—79 years: reference values and determinants, Age and Ageing

26 (1997) 15–19. URL: https://doi.org/10.1093/ageing/26.1.15. doi:10.1093/ageing/26.1.15.

[40] D. Helbing, I. Farkas, P. Molnar, T. Vicsek, Simulation of pedestrian crowds in normal and evacuation situations, volume 21, 2002, pp. 21–58.

[41] M. Haghani, M. Sarvi, Human exit choice in crowded built environments: Investigating underlying behavioural differences between normal egress and emergency evacuations, Fire Safety Journal 85 (2016) 1 – 9.

[42] Y. Mao, X. Fan, Z. Fan, W. He, Modeling group structures with emotion in crowd evacuation, IEEE Access 7 (2019) 140010–140021.

[43] X. Lu, Z. Yang, C. Chea, H. Guan, Experimental study on earthquake-induced falling debris of exterior infill walls and its impact to pedestrian evacuation, International journal of disaster risk reduction 43 (2020) 101372.

[44] M. Domaneschi, S. Marasco, G. Scutiero, A. Noori, O. Kammouh, V. Taurino, G. Cimellaro, Seismic damage assessment of a virtual large-scale city model, Maintenance, Safety, Risk, Management and Life-Cycle Performance of Bridges - Proceedings of the 9th International Conference on Bridge Maintenance, Safety and Management, IABMAS 2018 (2018) 1936–1942.

[45] S. Marasco, A. Zamani Noori, M. Domaneschi, G. Cimellaro, A computational framework for large-scale seismic simulations of residential building stock, Under Review in Engineering structures (Unpublished results), Elsevier (2020).

[46] HAZUS, Fema, multi-hazard loss estimation methodology, earthquake model, user manual, technical manual, Washington, DC, USA: Federal Emergency Management Agency (2015).

[47] I. Goodfellow, Y. Bengio, A. Courville, Deep Learning, MIT Press, 2016. http://www.deeplearningbook.org.

[48] C. M. Bishop, Pattern Recognition and Machine Learning (Information Science and Statistics), Springer-Verlag, Berlin, Heidelberg, 2006.

[49] T. M. Mitchell, Machine Learning, McGraw-Hill, 1997.

[50] EERI, Earthquake clearinghouse, 1999. URL: www.eqclearinghouse.org.

[51] Wikimedia Commons, Schwede66, Earthquake damage on worcester street, corner manchester street. christchurch cathedral is in the background, 4 September 2010. URL: https://commons.wikimedia.org/wiki/, [Online; accessed January 22, 2020]. The file is licensed under the Creative Commons Attribution 3.0 Unported license.

[52] M. Capozzo, A. Rizzi, G. Cimellaro, M. Domaneschi, A. Barbosa, D. Cox, Multi-hazard resilience assessment of a coastal community due to offshore earthquakes, Journal of Earthquake and Tsunami 13 (2019). doi:10.1142/S1793431119500088.

[53] A. Poulos, F. Tocornal, J. C. [de la Llera], J. Mitrani-Reiser, Validation of an agent-based building evacuation model with a school drill, Transportation Research Part C: Emerging Technologies 97 (2018) 82 – 95. URL: http://www.sciencedirect.com/science/article/pii/S0968090X18314670. doi:https://doi.org/10.1016/j.trc.2018.10.010.

[54] Z. Liu, C. C. Jacques, S. Szyniszewski, J. K. Guest, B. W. Schafer, T. Igusa, J. Mitrani-Reiser, Agent-based simulation of building evacuation after an earthquake: Coupling human behavior with structural response, Natural Hazards Review 17 (2016) 04015019. URL: https://doi.org/10.1061/(asce)nh.1527-6996.0000199. doi:10.1061/(asce)nh.1527-6996.0000199.

[55] ISTAT, Population by age, sex and marital status 2019, https://tinyurl.com/turinpop2019, 2019. Accessed: 2019-12-12.

[56] N. Ambraseys, K. Simpson, J. Bommer, Prediction of horizontal response spectra in europe, Earthquake Engineering & Structural Dynamics 25 (1996) 371–400.

[57] V. Karbovskii, D. Voloshin, A. Karsakov, A. Bezgodov, C. Gershenson, Multimodel agent-based simulation environment for mass-gatherings

and pedestrian dynamics, Future Generation Computer Systems 79 (2018) 155–165.

[58] A. Gutierrez-Milla, F. Borges, R. Suppi, E. Luque, Crowd evacuations saas: An abm approach, Procedia Computer Science 51 (2015) 473–482.

[59] F. F. T. Chiang, T. A. Birtch, Understanding Leading, Leader–Follower Relations, and Ethical Leadership in Organizations, Palgrave Macmillan UK, London, 2013, pp. 239–267.