A Container-Based Software Architecture for Coupling Heterogeneous Agent-Based Evacuation Models using Machine Learning

Aracelly S. Balboa Carvajal[1], Gabriel Astudillo Muñoz[1], Rene Noel[1]

[1]Escuela de Ingeniería Informática – Universidad de Valparaíso
Valparaíso – Chile

aracelly.balboa@alumnos.uv.cl, gabriel.astudillo, rene.noel@uv.cl

Abstract. *Large-scale disaster simulations often require the integration of heterogeneous systems, such as indoor egress models and city-wide traffic simulators. However, coupling these systems presents significant interoperability and scalability challenges due to high computational costs. This work proposes a modular software architecture to bridge this gap. We developed "Oráculo", a containerized microservice system that uses Machine Learning (ML) surrogates to encapsulate the complexity of microscopic Agent-Based Simulations (ABS). The architecture, orchestrated via Docker, allows outdoor simulators to query indoor evacuation rates through a lightweight API, decoupling the internal logic from the external flow. Validation results demonstrate that the ML component achieves an R^2 of 0.87, proving that high-fidelity simulation dynamics can be effectively served as efficient software components.*

1. Introduction

Tsunamis are highly destructive natural phenomena requiring rapid evacuation strategies. To plan for these events, the scientific community relies on simulation software to evaluate scenarios that are impossible to reproduce in real life [Yuksel 2018]. Among these, microscopic Agent-Based Models (ABM) have become the standard for representing individual human behaviors and emergent phenomena like panic [Mls et al. 2023]. However, from a software engineering perspective, simulating a complete city-scale evacuation introduces a critical scalability problem.

Current simulation systems are typically siloed into two categories: "Indoor Environments" (focused on clearing specific facilities) and "Outdoor Environments" (focused on street navigation) [Alvarez and Alonso 2018]. Integrating these two heterogeneous systems to create a holistic simulation is computationally prohibitive; microscopically representing thousands of agents interacting inside every building while simultaneously simulating city traffic consumes excessive processing resources [Muñoz et al. 2022]. This creates an interoperability gap: outdoor simulators often approximate building outflows with static functions, ignoring the complex internal dynamics generated by structural constraints or human behavior [León et al. 2020].

To address this integration and scalability challenge, this work proposes a hybrid software architecture that couples Agent-Based Simulation with Machine Learning. We developed a predictive tool designed as a microservice that encapsulates the complexity of indoor evacuations. The system uses an ABM to generate training data representing various building configurations and panic levels [Zheng et al. 2019]. This data trains a Deep Neural Network (DNN) which is then deployed as an inference engine. This approach allows outdoor simulation systems to integrate realistic building outflow rates via lightweight API calls, bypassing the need for heavy run-time microscopic computation.

The main contribution of this paper is a software artifact—containerized using Docker—that solves the interoperability bottleneck between indoor and outdoor evacuation models. By treating the indoor simulation as a learnable component, we provide a scalable solution for city-wide disaster planning, demonstrating the application of AI Engineering principles to optimize complex simulation systems.

The rest of the article continues as follows. First, Section 2 presents the related work, followed by Section 3 that describes the proposed artifact. Then Section 4 presents the methodology for the validation of the artifact, while Section 5 details the validation results. Section 6 discusses the results, and finally Section 7 presents the main conclusions of the work.

2. Related Work

The simulation of evacuation scenarios has become a critical tool for disaster management, allowing researchers to evaluate strategies that are impossible to test in real-life emergencies. Among the various approaches, Agent-Based Models (ABM) have emerged as the dominant paradigm due to their ability to represent heterogeneous behaviors and complex social interactions [Mls et al. 2023]. This section reviews the state of the art in evacuation modeling, categorized into indoor, outdoor, and mixed environments.

2.1. Indoor Evacuation Models

Indoor evacuation models focus on clearing specific infrastructures, such as buildings or stadiums. [Gutierrez-Milla et al. 2014] developed a distributed crowd evacuation simulator using the Social Force Model, comparing implementations in NetLogo and MPI/C to assess scalability in complex buildings. Similarly, [Kasereka et al. 2018] proposed an ABM in the GAMA platform that integrates the effects of fire and smoke on agent behavior, where the speed and health of individuals decrease as they interact with hazardous environmental factors.

Other studies emphasize the psychological aspects of evacuation. [Rozo et al. 2019] used NetLogo to model evacuations in multi-story buildings, evaluating how signage and route blockage influence congestion. [Kim and Han 2018] introduced an active route choice model in Unity, enabling agents to detect bottlenecks and change their path dynamically based on human characteristics. Furthermore, [Yuksel 2018] applied NeuroEvolution of Augmenting Topologies (NEAT) to allow agents to learn optimal exit routes, while [Şahin et al. 2019] combined fuzzy logic with ABM to simulate human emotions such as panic and stress during egress.

2.2. Outdoor Evacuation Models

Outdoor models generally address large-scale evacuations towards safe zones, particularly in tsunami scenarios. [Hatayama et al. 2018] analyzed evacuation dynamics in a tourist area in Mexico, differentiating the behaviors of residents, tourists, and leaders.

[Aguilar et al. 2019] presented a mass evacuation simulator capable of handling complex agents, including pedestrians and cars, under different environmental conditions.

In the context of the Chilean coast, [León et al. 2020] examined the vulnerability of Viña del Mar using a multi-scale framework that combines least-cost distance and agent-based modeling. Their results highlighted how micro-scale urban vulnerabilities can significantly increase evacuation times. To address computational efficiency in large-scale scenarios, [Muñoz et al. 2022] proposed a model using parallel computing (OpenMP) to simulate coastal cities, validating the approach with real data from Chile and Japan.

2.3. Mixed Approaches and Gap Analysis

There have been attempts to link indoor and outdoor environments. [Alvarez and Alonso 2018] proposed a methodology for a dam collapse scenario in Spain, connecting sub-models for different evacuation levels; however, the indoor component relied on estimated pre-evacuation times rather than microscopic simulation. Similarly, [Faucher et al. 2020] developed a model for near-field tsunamis in Puerto Rico that initializes pedestrians based on infrastructure data but does not simulate the internal exit dynamics.

Despite these advances, a significant gap remains in the literature: the lack of seamless integration between indoor and outdoor microscopic simulations. Indoor models typically remove agents once they reach the exit, while outdoor models often initialize agents on the streets without considering the actual outflow from buildings [Alvarez and Alonso 2018]. This disconnection limits the ability to analyze how congestion at building exits affects urban escape routes, largely due to the prohibitive computational costs of simulating both environments simultaneously at a microscopic level. Consequently, this work proposes a hybrid approach that integrates a machine learning-based predictor into outdoor simulations to accurately estimate indoor outflows without the computational burden of a fully coupled microscopic model.

3. Proposed Solution

To bridge the gap between indoor and outdoor environments, we designed a modular software architecture named ”Oráculo”. As shown in Figure 1, the system is implemented using Docker containers to ensure portability and modularity. It orchestrates two main computational paradigms: a microscopic simulation engine for data generation and a machine learning module for inference, communicating through a defined pipeline.

This work proposes a hybrid tool designed to bridge the gap between indoor and outdoor evacuation models. As illustrated in Figure 2, the solution workflow consists of configuring simulation scenarios for different buildings, running a microscopic Agent-Based Simulation (ABS) to generate training data, and subsequently training a Machine Learning (ML) model. This resulting predictor allows outdoor simulations to query evacuation rates without the computational overhead of simulating internal agents microscopically.

The predictive model acts as a data-driven approximation of indoor evacuation dynamics. Its function is to estimate, based on the current state of a building, how many

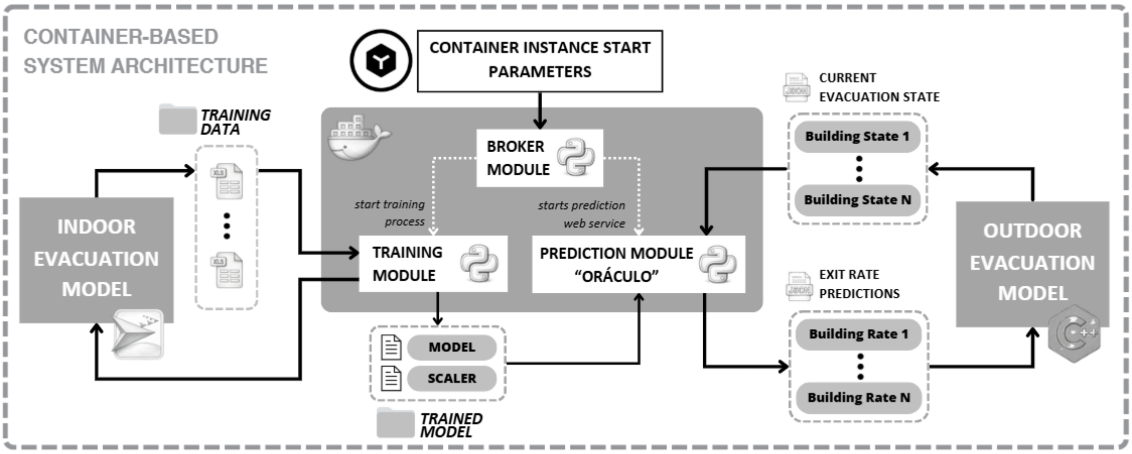

Figure 1. Container-based system architecture illustrating the interaction between the Broker, Training and Prediction module.

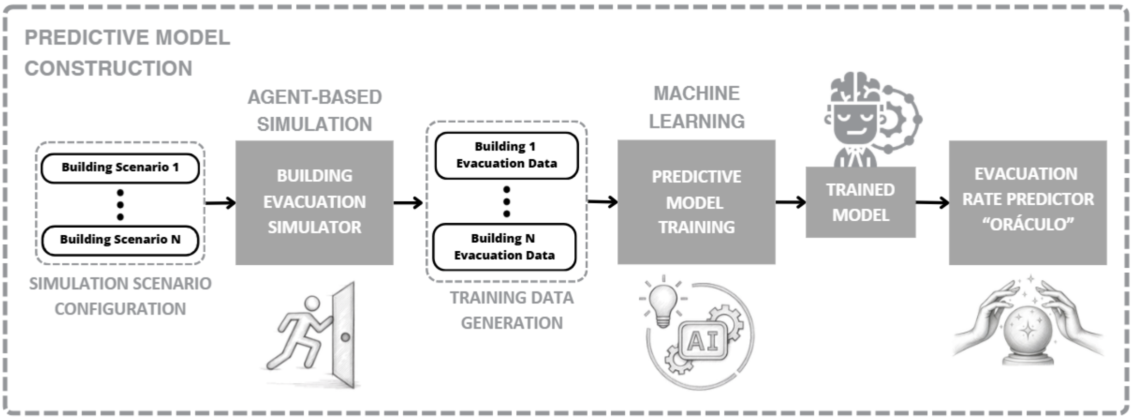

Figure 2. Construction of the predictive model workflow. Adapted from the simulation design.

people will leave the structure during the next simulation interval. To achieve this, it considers variables related to the building configuration, the internal evacuation conditions, and the level of pedestrian congestion near the exits. This estimate is then used by the outdoor simulator to generate agents progressively, enabling a continuous representation of the transition from indoor evacuation to movement in the urban environment, without the need to reproduce in detail the microscopic behavior inside each building.

3.1. Indoor Evacuation Simulator Design

The core component for data generation is a microscopic model developed in NetLogo 3D [Wilensky 1999]. The environment is represented as a three-dimensional grid where each cell corresponds to a physical space of $0.4 \times 0.4 \times 0.4$ meters. As depicted in Figure 3, this discretization allows for the representation of complex high-rise building geometries, including floors, stairs, and exits, using a coordinate system (x, y, z).

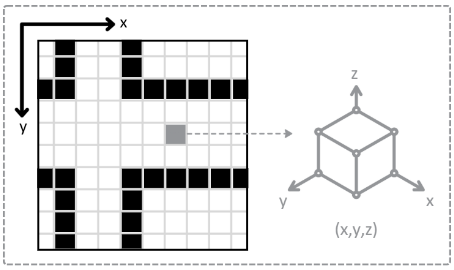

Figure 3. Three-dimensional grid cells representing the simulation environment coordinates (x, y, z).

3.1.1. Agent Behavior and State Machine

Agents in the simulation are modeled using a finite state machine that governs their lifecycle during the emergency. As shown in Figure 4, the agent transitions through five distinct phases:

1. **Indoor Response:** The agent reacts to the alarm.
2. **Indoor Movement:** The agent navigates towards the building exit.
3. **Evacuated:** The agent successfully leaves the building.
4. **Outdoor Response:** A transitional phase representing the decision-making process immediately after exiting.
5. **Outdoor Movement:** The agent moves away from the building towards safe zones until removed from the system.

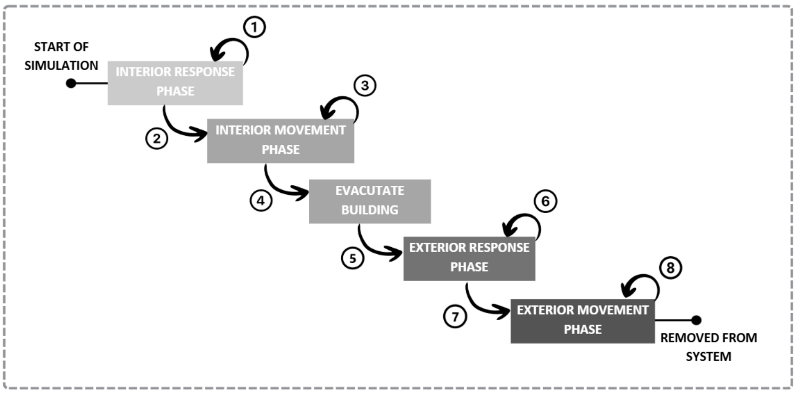

Figure 4. State diagram of a resident agent, from simulation start until removal from the system.

To simulate realistic human behavior under stress, a panic contagion model based on the epidemiological SIR (Susceptible-Infected-Recovered) framework is implemented [Fu et al. 2014]. Figure 5 illustrates this cycle. Agents start as *Susceptible*; if their emotional intensity exceeds a threshold due to interactions, they become *Infected* (Panic), exhibiting random movement. After a specific time, they transition to *Recovered* and resume rational pathfinding, eventually becoming susceptible again.

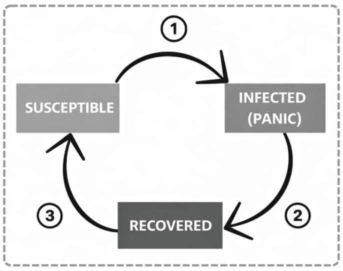

Figure 5. Three-state panic model based on the SIR epidemiological framework.

3.1.2. Perception and Movement

The movement logic is based on local perception and neighborhood rules. Agents utilize the Moore Neighborhood (Figure 6) to evaluate immediate adjacent cells for movement. To avoid collisions and navigate efficient paths, agents also analyze a wider Radial Neighborhood (Figure 7) to assess crowd density.

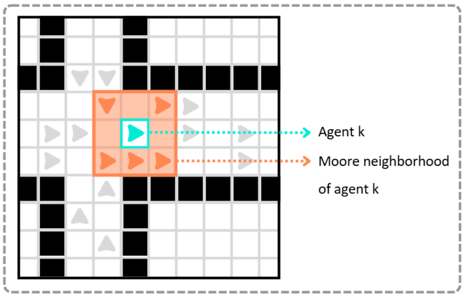

Figure 6. Moore neighborhood (radius 1) of agent k.

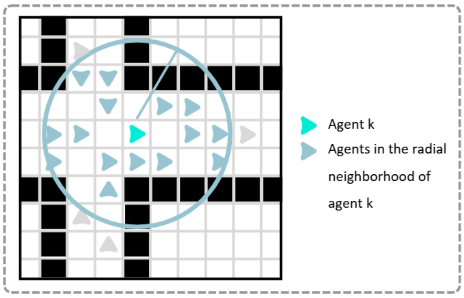

Figure 7. Radial neighborhood of agent k with a 3-cell radius.

Furthermore, the panic propagation and decision-making processes are influenced by the agent's visual perception. As shown in Figure 8, each agent possesses a field of view of 120°. This sensory input allows agents to react to the emotional state of others ahead of them, facilitating the implementation of the emotional contagion dynamics described in the SIR model.

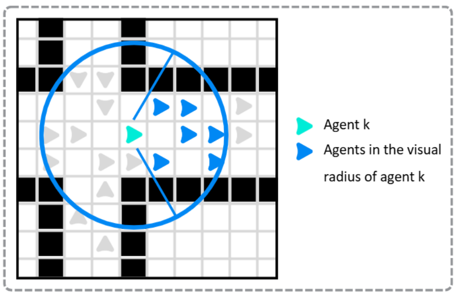

Figure 8. Field of view of agent k, based on a visual radius of 3 cells.

3.2. Evacuation Rate Predictor

The data collected from the simulations described above—capturing the complex interactions of geometry, panic, and movement—is used to train a Dense Neural Network. This predictor takes structural variables (e.g., number of floors, exit width) and dynamic variables (e.g., current occupancy, panic levels) to output the expected evacuation rate for the next time interval.

4. Methodology

The methodology employed in this work follows a two-stage experimental design, as illustrated in Figure 9. The first stage involves the generation of synthetic data through the microscopic simulation of indoor evacuations under varying conditions. The second stage focuses on the training and validation of machine learning models to predict evacuation rates based on the generated data.

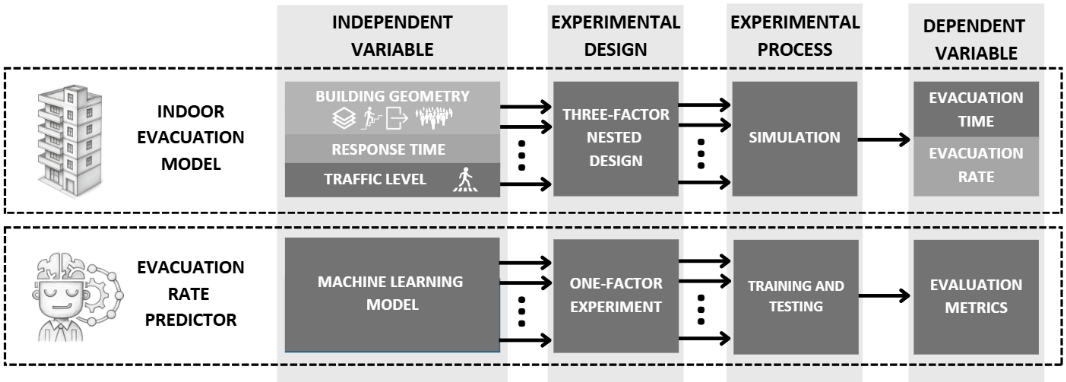

Figure 9. Experimental design showing the manipulation of independent variables (building geometry, response time, outdoor traffic) to obtain dependent variables (evacuation time and rate).

4.1. Case Study: Viña del Mar

To ensure the external validity of the simulation parameters, this study focuses on the city of Viña del Mar, Chile. This coastal city is characterized by high seismic activity and a significant tsunami risk, with a dense urban center located in a flood zone [León et al. 2020].

Population demographics for the agents (age groups and walking speeds) were configured based on the 2017 Census data provided by the National Statistics Institute [Instituto Nacional de Estadísticas 2017]. Two real-world infrastructures were modeled to represent different evacuation complexities:

1. **Residential Building (Acapulco community building):** A 15-story building located in a high-risk zone (Figure 11). It represents a high-density vertical evacuation scenario with limited exit capacity (one stairway).
2. **Educational Facility (República de Colombia School):** A 4-story school building (Figure 12). It represents a scenario with a younger population and multiple exit routes.

4.2. Simulation Experimental Setup

The data generation phase utilized the NetLogo 3D simulator described in the previous section. A nested experimental design was applied to capture the interaction between indoor and outdoor flows.

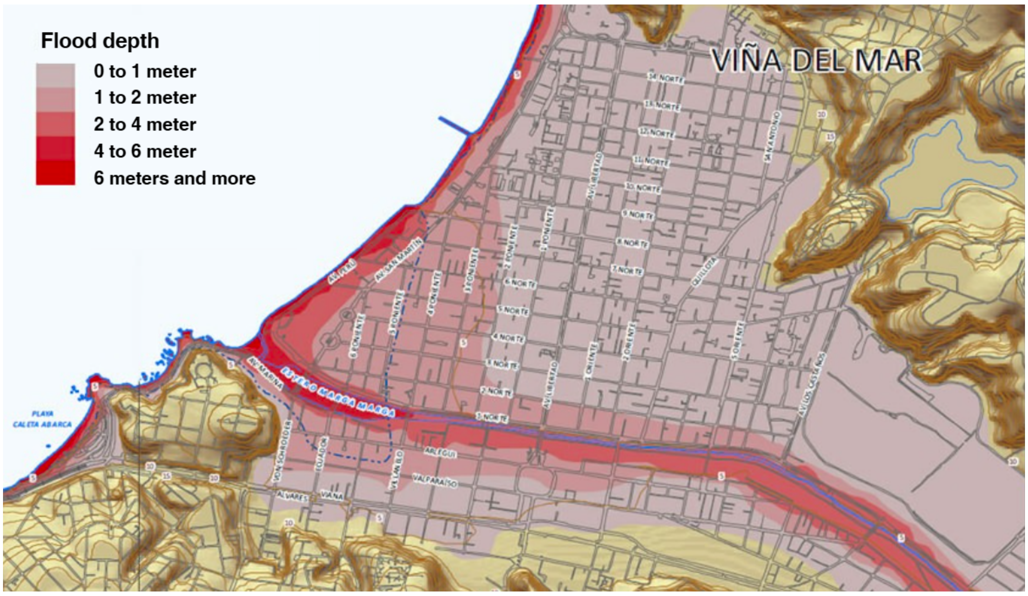

Figure 10. Inundation map of downtown Viña del Mar. Source: SHOA.

Figure 11. Acapulco Building Community.

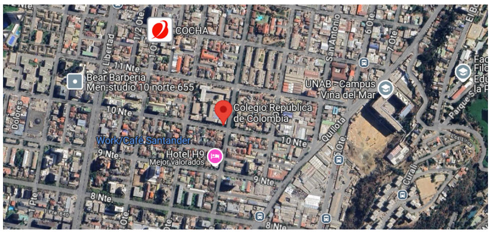

Figure 12. República de Colombia School.

4.2.1. Variables

The independent variables manipulated during the simulations were:

- **Building Type:** Residential vs. Institutional.
- **Response Time (T):** The delay before agents decide to evacuate, set at intervals of $\{30, 60, 90, 120\}$ seconds.
- **Outdoor Traffic Density (D):** The number of transients blocking the exit area, tested at levels of $\{0, 500, 1000, 2000, 3000, 4000\}$ agents.

Controlled variables included corridor widths, stair dimensions, and panic thresholds, which were kept constant to isolate the effect of outdoor congestion. The dependent variable measured was the **Evacuation Rate**, defined as the number of agents exiting the building per 10-second interval.

4.3. Machine Learning Model Validation

The second stage involved training regression models to predict the evacuation rate. The dataset generated from the simulations was split into a training set (80%) and a testing set

(20%). Input features were normalized to ensure scale consistency across variables.

4.3.1. Models and Hyperparameters

Three algorithms were evaluated to determine the most accurate predictor:

1. **Decision Tree Regressor:** Configured with a random state of 42 for reproducibility.
2. **Random Forest Regressor:** An ensemble method configured with 100 estimators ($n\_estimators = 100$) to reduce overfitting [Breiman 2001].
3. **Deep Neural Network (DNN):** A fully connected network implemented in TensorFlow/Keras. The architecture consists of an input layer, three hidden layers (64, 32, and 16 neurons) with ReLU activation, and a linear output layer. The model was trained using the Adam optimizer and Mean Squared Error (MSE) loss function, utilizing Early Stopping to prevent overfitting [Goodfellow et al. 2016].

4.3.2. Performance Metrics

To validate the models internally, the following metrics were used:

- **Mean Absolute Error (MAE):** Measures the average magnitude of errors in the predictions.
- **Mean Squared Error (MSE):** Penalizes larger errors, useful for detecting outliers.
- **Coefficient of Determination (R^2):** Indicates the proportion of the variance in the dependent variable that is predictable from the independent variables.

4.4. Implementation Environment

The experiments were conducted on a workstation equipped with an 11th Gen Intel Core i7 processor (2.80GHz) and 16 GB of RAM. The simulation software used was NetLogo 6.2.2 (3D), and the machine learning pipeline was implemented in Python 3.11 using scikit-learn and TensorFlow libraries.

5. Results

This section presents the findings obtained from the experiments. First, we analyze the impact of outdoor pedestrian traffic on indoor evacuation dynamics using the Agent-Based Simulator. Second, we evaluate the performance of the Machine Learning models trained to predict evacuation rates.

The dataset generated by the simulator and used for training the machine learning models is publicly available in [Balboa 2024b]. In addition, the *docker-oraculo* project, which implements the proposed container-based architecture, is available in [Balboa 2024a].

5.1. Indoor Evacuation Dynamics

The simulation experiments revealed a significant interaction between outdoor traffic density and indoor evacuation efficiency. The results are analyzed below for the two case studies: a residential building and an educational facility.

5.1.1. Case 1: Residential Building (Acapulco Building)

Figure 13 illustrates the evacuation curves for the residential building configured with a single exit and a response time of 60 seconds. The data shows that as outdoor traffic increases, the evacuation efficiency decreases. In the scenario with zero external transients, approximately 80% of occupants evacuated within 600 seconds. However, with a high density of outdoor traffic (3000–4000 agents), this percentage dropped to 70%. The boxplot analysis further confirms this, showing a lower median evacuation rate and higher variability in scenarios with heavy external congestion.

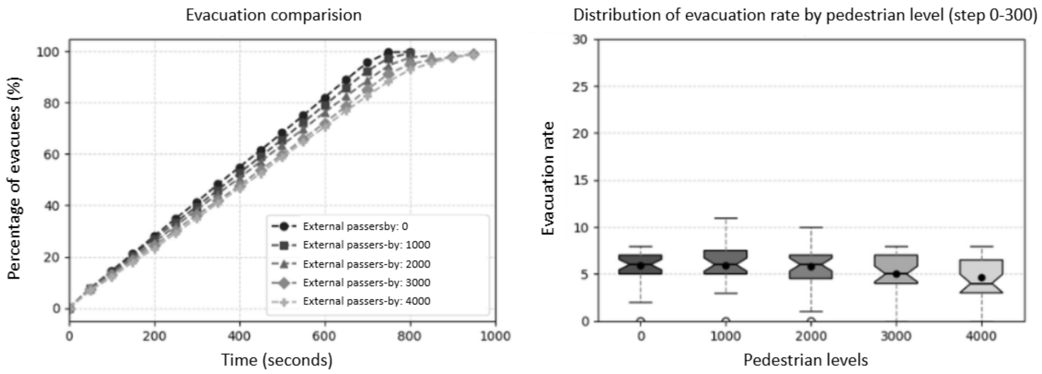

Figure 13. Results for Acapulco Community building with one exit and 60s response time. Left: Evacuation percentage over time. Right: Evacuation rate distribution.

In contrast, when the building was configured with two exits (Figure 14), the negative impact of outdoor traffic was almost entirely mitigated. The evacuation curves remained consistent across all external traffic levels. This suggests that structural improvements, such as increasing exit capacity, can counteract the blocking effect caused by pedestrians on the streets.

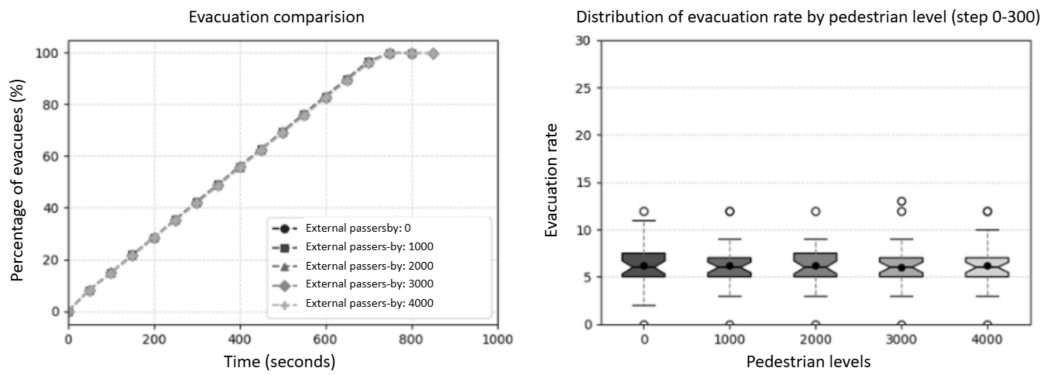

Figure 14. Results for Acapulco Community building with two exits and 60s response time, showing resilience to outdoor traffic.

5.1.2. Case 2: Educational Facility (República de Colombia School)

For the school scenario, we evaluated the combined effect of traffic and response time. With a rapid response time (30s) and a single exit (Figure 15), the impact of outdoor traffic was pronounced; the evacuation completion at 300 seconds dropped from 90% (no traffic) to 70% (3000 transients).

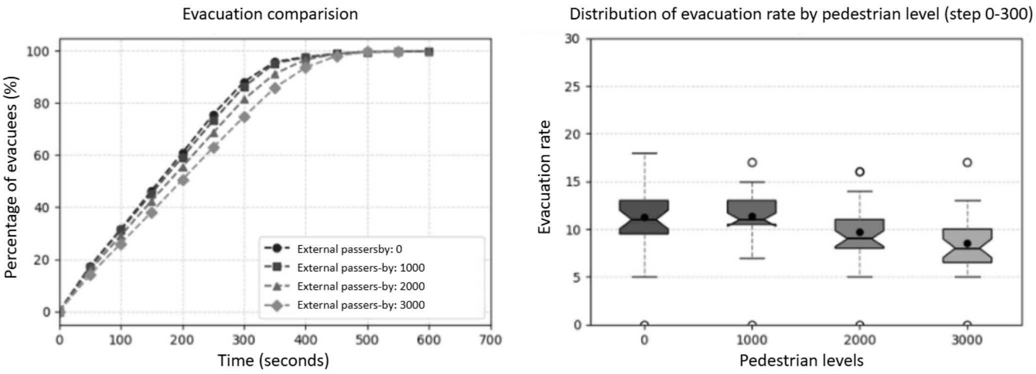

Figure 15. Results for Colegio República de Colombia: 1 exit, 30s response time.

However, when the response time was increased to 90 seconds (Figure 16), the overall efficiency dropped further, and the gap between low and high traffic scenarios widened. Even with two exits enabled (Figure 17), a delayed response time of 90 seconds allowed external congestion to significantly affect the outflow, reducing the evacuated population from 90% to 75% in the high-traffic scenario. This indicates that while multiple exits help, they cannot fully compensate for delayed reaction times combined with heavy outdoor congestion.

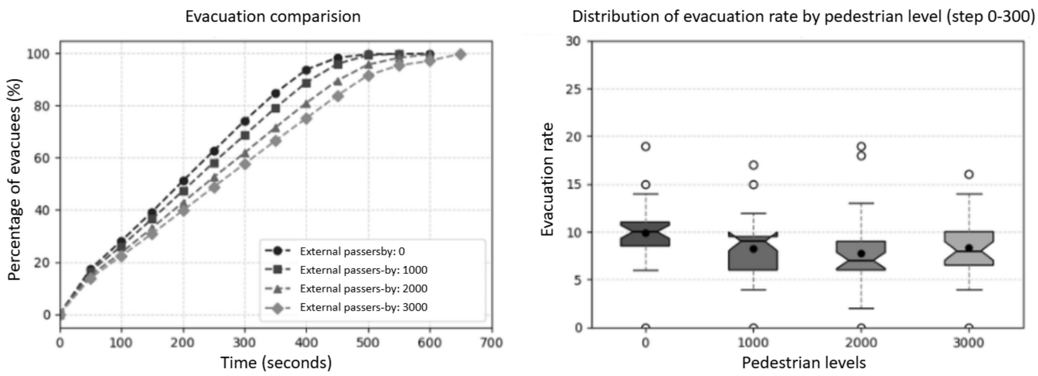

Figure 16. Colegio R. Colombia: 1 exit, 90s response.

5.2. Evacuation Rate Predictor Performance

In the second stage, we evaluated three regression models to predict the evacuation rate (persons per 10s) based on the structural and dynamic features generated by the simulations. Table 1 summarizes the performance metrics: Mean Absolute Error (MAE), Mean Squared Error (MSE), and the Coefficient of Determination (R^2).

The Deep Neural Network (DNN) achieved the best performance, with the lowest MAE (1.5301) and the highest R^2 (0.8690), indicating it can explain approximately 87% of the variance in the evacuation rate [8]. The Random Forest model also performed well ($R^2 \approx 0.86$), confirming that non-linear models are suitable for this task. The Decision Tree showed the lowest accuracy, suggesting it failed to capture the complexity of

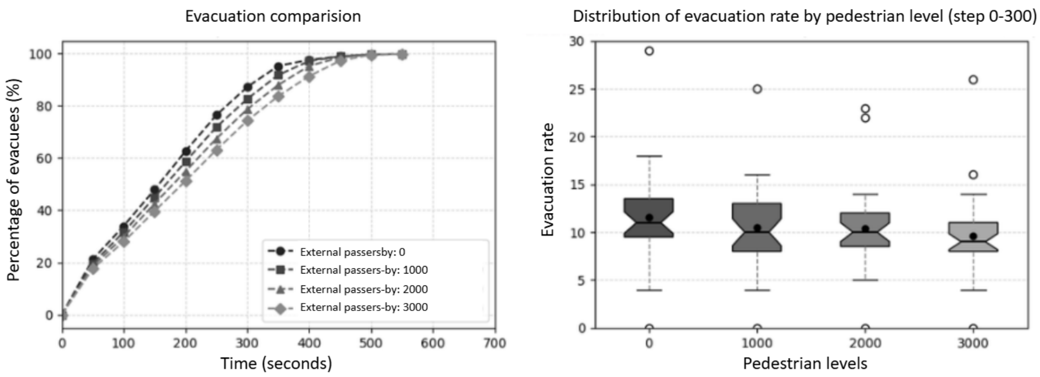

Figure 17. Colegio R. Colombia: 2 exits, 90s response.

Table 1. Performance metrics of the Machine Learning models.

Model	MAE	MSE	R^2 (%)
Decision Tree Regressor	2.0159	8.2674	0.7529
Random Forest	1.5534	4.7353	0.8585
Neural Network (DNN)	**1.5301**	**4.3821**	**0.8690**

the agent interactions. Consequently, the DNN was selected as the core engine for the ”Oráculo” prediction module.

6. Discussion

The central hypothesis of this research posits that the interaction between outdoor pedestrian flow and indoor evacuation processes significantly impacts evacuation times, and that this complex interaction can be effectively modeled using machine learning techniques. The results obtained from the experiments strongly support this hypothesis, revealing critical dependencies between structural configuration, human response time, and urban congestion.

6.1. Impact of Outdoor Flow and Structural Configuration

The simulation results indicate that outdoor traffic acts as a dynamic barrier that can severely compromise evacuation efficiency, particularly in buildings with limited exit capacity. In the case of the residential building (Edificio Acapulco) with a single exit, a high density of outdoor transients (4000 agents) reduced the evacuation completion rate by approximately 10% compared to a scenario with no external traffic. This confirms the existence of a ”blocking effect” where congestion on the sidewalk propagates backward into the building, creating bottlenecks at the discharge point.

However, the introduction of a second exit proved to be a decisive mitigation factor. As observed in the results, structural redundancy almost entirely neutralized the negative impact of outdoor flow. This finding aligns with previous studies emphasizing the importance of multi-exit designs in high-occupancy buildings [Yuksel 2018]. Nevertheless, our results add a nuance to this principle: structural efficiency is not absolute. In the educational facility scenario, when the response time was increased to 90 seconds, the benefit of having two exits diminished in the presence of high outdoor traffic. This suggests that infrastructure improvements must be accompanied by strategies to reduce human reaction times, such as effective alarm systems and regular drills, to prevent the synchronization of indoor outflow with peak outdoor congestion.

6.2. Predictive Capability and Computational Efficiency

The second major finding of this work is the validation of Machine Learning as a viable surrogate for computationally expensive microscopic simulations. The Deep Neural Network (DNN) achieved an R^2 of 0.8690 and a Mean Absolute Error (MAE) of 1.53, outperforming the Decision Tree and Random Forest models. This superior performance can be attributed to the DNN's ability to capture non-linear relationships inherent in crowd dynamics, such as the sudden phase transitions in flow caused by panic contagion (SIR model) and physical obstructions.

This predictive capability addresses the core problem identified in the literature: the disconnection between indoor and outdoor models due to computational costs [Muñoz et al. 2022]. By replacing the complex microscopic simulation of building interiors with a lightweight inference query to the "Oráculo" module, it becomes feasible to simulate city-scale evacuations (like those required for a tsunami in Viña del Mar) while still accounting for the specific exit dynamics of hundreds of individual buildings.

6.3. Limitations

Despite the promising results, this study has limitations. The mobility model employed relies on a Dijkstra-based pathfinding algorithm combined with local avoidance rules. While efficient, it may not capture the full range of physical force interactions observed in high-density crowds as accurately as the Social Force Model [Helbing and Molnar 1998]. Additionally, the current predictor assumes that outdoor congestion is the primary external variable; future iterations should incorporate other environmental factors, such as debris on the streets or the arrival of emergency vehicles.

Furthermore, the integration is currently unidirectional (outdoor simulators query the indoor predictor). A fully coupled bidirectional approach, where the outdoor state dynamically updates the indoor agents' decision-making (e.g., agents deciding to stay inside if they perceive the street is blocked), represents a compelling direction for future research.

7. Conclusions

This work presented a software engineering solution to the challenge of coupling heterogeneous evacuation models. By implementing a hybrid architecture that combines Agent-Based Simulation (ABS) with Machine Learning (ML), we successfully developed a scalable evacuation rate predictor. This tool resolves the computational bottleneck that previously prevented the integration of detailed indoor dynamics into large-scale outdoor simulations [Muñoz et al. 2022].

From an engineering perspective, the system demonstrated high reliability and performance. The deployment of the solution using Docker containers and microservices facilitates its integration into third-party outdoor simulation platforms, adhering to principles of modularity and interoperability. The validation of the Machine Learning component, specifically the Deep Neural Network, yielded a Coefficient of Determination (R^2) of 0.869, proving that the complex non-linear behaviors of agents (derived from the SIR panic model and physical constraints) can be effectively encapsulated and served as a lightweight software component.

The experimental results confirmed that structural factors (e.g., number of exits) and external constraints (e.g., street congestion) significantly impact evacuation efficiency. Our tool captures these nuances without the runtime cost of full microscopic simulation. However, the current mobility logic relies on Dijkstra's algorithm; future iterations will incorporate the Social Force Model to enhance physical realism [Helbing and Molnar 1998].

Future work will focus on establishing a bidirectional coupling interface, allowing the outdoor simulator to dynamically update the state of the indoor predictor in real-time. This would further advance the state of the art in System-of-Systems simulations for disaster management.

References

Aguilar, L., Wijerathne, L., Jacob, S., Hori, M., and Ichimura, T. (2019). Mass evacuation simulation considering detailed models: behavioral profiles, environmental effects, and mixed-mode evacuation. *Asia Pacific Management Review*, 24:114–123.

Alvarez, P. and Alonso, V. (2018). Using microsimulation software to model large-scale evacuation scenarios. the case of sangüesa and the yesa dam collapse. *Safety Science*, 106:10–27.

Balboa, A. (2024a). Casandra docker image - version 01. `https://hub.docker.com/layers/ssstefany/casandra/version01/images/sha256:99d54e65286ff0759468bb6938afc6b274fd2ff15ba2bf12925a65d8b08b934f?uuid=C9ED271D-6AD9-4376-96BD-78AD0DD654F7`. Accessed: 2024-11-27.

Balboa, A. (2024b). Datos de evacuación generados. `https://drive.google.com/drive/folders/1RMdPCQdRH1itZxWpeuMvcl6o1Yqowybz?usp=sharing`. Accessed: 2024-11-27.

Breiman, L. (2001). Random forests.

Faucher, J. E., Dávila, S., and Hernández-Cruz, X. (2020). Modeling pedestrian evacuation for near-field tsunamis fusing alcd and agent-based approaches: A case study of rincón, pr. *International Journal of Disaster Risk Reduction*, 49.

Fu, L., Song, W., Lv, W., and Lo, S. (2014). Simulation of emotional contagion using modified sir model: A cellular automaton approach. *Physica A: Statistical Mechanics and its Applications*, 405:380–391.

Goodfellow, I., Bengio, Y., and Courville, A. (2016). *Deep Learning*. MIT Press. `http://www.deeplearningbook.org`.

Gutierrez-Milla, A., Borges, F., Suppi, R., and Luque, E. (2014). Individual-oriented model crowd evacuations distributed simulation. *Procedia Computer Science*, 29:1600–1609. 2014 International Conference on Computational Science.

Hatayama, M., Kosaka, T., and Hernándeze, A. H. (2018). Analysis on tsunami evacuation options with agent-based simulation in tourist area. In *2018 5th International*

Conference on Information and Communication Technologies for Disaster Management (ICT-DM), pages 1–8. IEEE.

Helbing, D. and Molnar, P. (1998). Social Force Model for Pedestrian Dynamics. *Physical Review E*, 51(5):4282–4286.

Instituto Nacional de Estadísticas (2017). Censo de población y vivienda. Accedido: 2 de junio de 2024.

Kasereka, S., Kasoro, N., Kyamakya, K., Goufo, E. F. D., Chokki, A. P., and Yengo, M. V. (2018). Agent-based modelling and simulation for evacuation of people from a building in case of fire. *Procedia Computer Science*, 130:10–17.

Kim, H. and Han, S. (2018). Crowd evacuation simulation using active route choice model based on human characteristics. *Simulation Modelling Practice and Theory*, 87:369–378.

León, J., Castro, S., Mokrani, C., and Gubler, A. (2020). Tsunami evacuation analysis in the urban built environment: a multi-scale perspective through two modeling approaches in viña del mar, chile. *Coastal Engineering Journal*, 62:389–404.

Mls, K., Kořínek, M., Štekerová, K., Tučník, P., Bureš, V., Čech, P., Husáková, M., Mikulecký, P., Nacházel, T., Ponce, D., Zanker, M., Babič, F., and Triantafyllou, I. (2023). Agent-based models of human response to natural hazards: systematic review of tsunami evacuation.

Muñoz, G. A., Gil-Costa, V., and Marin, M. (2022). Efficient simulation of natural hazard evacuation for seacoast cities. *International Journal of Disaster Risk Reduction*, 81:103300.

Rozo, K. R., Arellana, J., Santander-Mercado, A., and Jubiz-Diaz, M. (2019). Modelling building emergency evacuation plans considering the dynamic behaviour of pedestrians using agent-based simulation. *Safety Science*, 113:276–284.

Şahin, C., Rokne, J., and Alhajj, R. (2019). Human behavior modeling for simulating evacuation of buildings during emergencies. *Physica A: Statistical Mechanics and its Applications*, 528:121432.

Wilensky, U. (1999). Netlogo. Center for Connected Learning and Computer-Based Modeling, Northwestern University, Evanston, IL.

Yuksel, M. E. (2018). Agent-based evacuation modeling with multiple exits using neuroevolution of augmenting topologies. *Advanced Engineering Informatics*, 35:30–55.

Zheng, L., Peng, X., Wang, L., and Sun, D. (2019). Simulation of pedestrian evacuation considering emergency spread and pedestrian panic. *Physica A: Statistical Mechanics and its Applications*, 522:167–181.