

Bachelor's thesis

**HazardEvac: An Evacua-
tion Route Planning Plat-
form with Hazard Simulation**

Pantelis Kanaris

UNIVERSITY OF CYPRUS



DEPARTMENT OF COMPUTER SCIENCE

May 2026

UNIVERSITY OF CYPRUS
Faculty of Pure and Applied Sciences
DEPARTMENT OF COMPUTER SCIENCE

HazardEvac: An Evacuation Route Planning Platform with Hazard Simulation

Pantelis Kanaris

Supervisor

Prof. Demetris Zeinalipour

The Thesis was submitted in partial fulfillment of the requirements for obtaining
the Computer Science degree at the Department of Computer Science
of the University of Cyprus

Declaration

I declare that this dissertation and any accompanying software are my own original work, conducted in full compliance with the University of Cyprus Code of Conduct for Research (<https://www.ucy.ac.cy/legislation/kodikas-deontologias-kai-politikes/>), and with full accountability on my part for the accuracy, integrity, and validity of all content.

May 2026

Acknowledgements

First, I would like to express my gratitude to my supervisor Prof. Demetris Zeinalipour for his continuous guidance throughout the year. The communication between us consistently provided me with valuable feedback, pushing me to achieve great results. I also would like to express my gratitude to my friends and family for providing me with continuous unconditional love and support throughout this journey. Having them by my side, helped me maintain a clear and sharp mind by keeping me motivated. I could not have done this without them.

ABSTRACT

Airborne hazard events, such as toxic industrial releases, radiological accidents, and wildfire smoke episodes, pose a significant threat to public safety and make evacuation decision making particularly challenging. In such scenarios, the hazard does not remain static, but evolves over space and time according to atmospheric transport conditions, wind, and plume dispersion. Although immediate sheltering is often the preferred first response in such emergencies, movement may still become necessary for relocation to a safer shelter, access to emergency reception points, medical need, infrastructure disruption, or family reunification. As a result, the safety of an evacuation route depends not only on the geographic path selected, but also on the time at which each location is reached. Conventional routing approaches typically focus on minimizing travel distance or travel time and therefore provide limited support for situations in which evacuees must move through a dynamically changing contamination field.

To address this problem, this thesis presents, *HazardEvac* a framework for time-dependent evacuation routing under hazardous atmospheric conditions. The proposed framework integrates three main stages. First, a plume generation stage that uses HYSPLIT to produce time-indexed concentration fields for airborne hazards. Second, a hazard integration stage that maps these outputs onto a traversable routing grid while accounting for terrain constraints. Third, a routing stage that applies the Emergency Airborne Release (EAR) algorithm, an exposure-aware routing method that evaluates candidate movements by jointly considering hazard intensity, traversal time, and plume-aware spatial penalties. Through this approach, the system supports single-destination routing, multi-waypoint routing, and multi-destination routing.

To examine the effectiveness of the proposed approach, *HazardEvac* is evaluated under dynamically evolving radiological hazard scenarios using simulation generated plume datasets. The experimental study compares EAR against a Euclidean fastest-path proxy and a Random baseline across multiple transport modes, mobility speeds, plume propagation speeds, and respiratory protection assumptions. The results show that EAR consistently achieves the lowest route-level dose, with average exposure remaining approximately two to three orders of magnitude lower than the corresponding values of the baseline approaches across the tested transport modes. At the same time, the resulting travel times remain close to those of the direct-path baseline, indicating that the reduction in exposure is not achieved through impractically long detours. Lastly additional sensitivity analyses further show that the advantage of EAR is preserved under varying evacuee speeds, plume dynamics, and protection settings, demonstrating both the effectiveness and the robustness of the proposed routing approach.

TABLE OF CONTENTS

ABSTRACT	iv
TABLE OF CONTENTS	v
1 Introduction	1
1.1 Architecture Overview	5
1.2 Thesis Structure	7
2 Background and Related Work	9
2.1 Introduction	9
2.2 Background on Emergency Response Applications	9
2.2.1 Emergency Response and Public Warning Applications	9
2.2.2 Protective Actions in Airborne Hazard Emergencies	12
2.2.3 Concentration and Exposure	13
2.3 Background on Plume Simulation	16
2.3.1 Computational Fluid Dynamics Models	16
2.3.2 Gaussian Plume and Puff Models	17
2.3.3 Lagrangian Dispersion Models	19
2.3.4 Selection of HYSPLIT	21
2.4 Related Work	23
3 HazardEvac: System Overview	26
3.1 Introduction	26
3.2 Hazardous Evacuation Setting	26
3.3 Routing Domain Representation	27
3.3.1 Traversability and Connectivity	28
3.3.2 Path Representation	30
3.4 Time-Dependent Hazard Representation	30
3.4.1 Continuous Hazard Field	30
3.4.2 Grid-Level Hazard Values	31
3.5 Problem Statement	33

3.6	End-to-End Workflow	34
3.7	Main Architectural Components	35
3.8	Data Flow Across the Platform	37
3.9	High-Level View of EAR	39
4	HazardEvac: Plume Generation Layer	40
4.1	Introduction	40
4.2	Scenario Definition and Simulation Inputs	40
4.3	Meteorological Data Acquisition and Preparation	42
4.4	Plume Simulation Workflow	43
4.5	Generated Outputs and Their Role in the System	44
5	HazardEvac: Hazard Integration Layer	46
5.1	Introduction	46
5.2	Overview of the Hazard Integration Layer	47
5.3	Routing Domain Construction	48
5.4	Terrain-Aware Preprocessing	50
5.5	Plume to Grid Mapping	52
6	HazardEvac: Routing and Application Layer	55
6.1	Introduction	55
6.2	Routing Objectives and Cost Model	56
6.3	The Emergency Airborne Release (EAR) Algorithm	61
6.4	Application-Level Route Generation and Result Presentation	66
6.5	Dose-Related Interpretation and Reporting	68
7	Experimental Evaluation	71
7.1	Introduction	71
7.2	Experimental Setup	72
7.2.1	Testbed	72
7.2.2	Compared Approaches	73
7.2.3	Datasets and Scenarios	73
7.2.4	Evaluation Metrics	74
7.2.5	Microbenchmark Factors	75
7.3	Results	76
7.3.1	Radiological Scenario	76
7.3.2	Mobility Speed Variation	78
7.3.3	Plume Propagation Speed Variation	80

7.3.4	Personal Protection Measures	82
7.4	Discussion	83
8	Conclusions and Future Work	84
8.1	Summary of the Thesis	84
8.2	Main Contributions	85
8.3	Limitations	86
8.4	Future Work	87
	BIBLIOGRAPHY	88

1 Introduction

Airborne hazard events, such as toxic industrial releases, radiological accidents, and wildfire smoke episodes, pose a serious threat to public safety and make emergency response particularly challenging. In such incidents, the hazard is not static, but evolves over space and time according to atmospheric transport conditions, wind, turbulence, and the physical characteristics of the released substance. As a result, movement through a contaminated area becomes a spatio-temporal decision problem, where risk depends not only on the selected route, but also on the time at which each location is reached. This issue is especially critical in radiological emergencies, where exposure minimization is essential because ionizing radiation may cause both acute and long-term health effects [1]. Similar concerns arise in chemical releases and smoke-related incidents, where evacuees may be exposed to dangerous airborne contaminants while attempting to move towards safety.



(a) Wildfire smoke plume. Source: LIFO [2].



(b) Industrial fire and smoke release. Source: Focus Taiwan [3].

Figure 1.1: Examples of airborne hazard scenarios involving smoke and industrial releases

In many such scenarios, immediate evacuation is not always the preferred first response. International guidance often recommends initial protective sheltering, especially during the early phase of a hazardous airborne release, where affected populations may be advised to remain indoors and await official instructions [4]. Nevertheless, sheltering does not eliminate the need for later movement. Hazardous conditions may remain operationally relevant for extended periods of time, and individuals may still need to relocate to better shelter, access emergency reception points, seek medical care, reunite with family members, or respond to infrastructure disruption and delayed evacuation orders.

When such movement becomes necessary, the remaining question is no longer simply *whether* to move, but also *how* to move more safely through a contaminated but still traversable environment. This is precisely where route planning becomes important. Classical evacuation routing



(a) Satellite view of smoke dispersion. Source: NASA Scientific Visualization Studio [10].



(b) Emergency response to hazardous release. Source: European Commission Civil Protection Knowledge Network [11].

approaches have mainly focused on objectives such as distance minimization, travel time reduction, road capacity, and congestion handling [5,6]. While these formulations are highly relevant in many emergency management settings, they are not sufficient for hazardous atmospheric dispersion scenarios, where the risk encountered along a route depends on both space and time.

In such settings, a shortest route may pass through regions of high contamination, whereas a geometrically longer route may lead to substantially lower cumulative exposure. Moreover, even the same path may produce different exposure outcomes depending on the departure time, the selected mode of transportation, and the projected plume conditions during traversal. This creates a need for route planning methods that can directly consume time-varying hazard information rather than relying only on static maps or conventional movement costs.

Previous work has also shown that producing such hazard information is itself a complex spatio-temporal problem. Studies on air contaminant dispersion prediction, hazardous gas estimation, and dynamic fire spread modeling demonstrate the importance of spatially and temporally resolved hazard fields for emergency management [7–9]. These works provide valuable support for understanding how hazardous conditions develop, but they do not by themselves determine how an individual should move through the affected area. This gap motivates the integration of hazard prediction with exposure-aware route computation.

Building on this need, more recent work has begun to incorporate hazard or exposure considerations directly into routing and evacuation decisions. Examples include routing in dynamic nuclear dose fields, evacuation under toxic cloud scenarios, wildfire rescue and path planning, and exposure-aware navigation in public-health settings [12–15]. These studies show that route quality in hazardous environments cannot be assessed only through distance or travel time, since the risk accumulated during movement may vary significantly across space and time. However, much of the existing work remains tied to specific scenarios, assumptions, or operational settings, such as indoor evacuation, fleet-based planning, wildfire response, or disease-exposure avoidance. Consequently, there is still limited work on end-to-end evacuation support that

connects atmospheric plume prediction, routing-domain construction, and individualized time-dependent route computation within a unified operational framework.




Figure 1.3: Overview of the *HazardEvac* platform.

Motivated by this gap, this thesis presents *HazardEvac*, an end-to-end platform for evacuation planning under hazardous atmospheric conditions. As illustrated in Fig. 1.3, the platform provides a map-based environment in which the user can configure a hazardous release scenario, inspect the predicted plume over the affected geographic area, define the routing domain, and request an evacuation route through the hazard field. The system first generates plume predictions for a user-defined scenario using the HYSPLIT atmospheric dispersion model. It then transforms the resulting concentration outputs into a routing-ready representation by aligning them with the traversable movement domain. This allows route planning to account not only for where movement is possible, but also for how the hazard is expected to change during traversal. In this way, *HazardEvac* connects plume prediction, spatial visualization, and exposure-aware route computation within a single workflow for practical evacuation support.




Figure 1.4: Overview of the EAR algorithm.

The routing core of the platform is the *Emergency Airborne Release* (EAR) algorithm, illustrated in Fig. 1.4. EAR evaluates candidate movements over the hazard-aware routing domain by considering more than distance or travel time alone. In particular, it accounts for the predicted movement of the plume, the user-designated standoff distance from hazardous regions, the current and forecast concentration values along each candidate movement, and the time required to traverse the terrain. As a result, routing decisions reflect both the changing hazard conditions and the practical movement constraints of the evacuee, allowing the platform to identify routes that remain feasible while reducing exposure to the airborne release.

Overall, *HazardEvac* is designed as a practical decision-support platform for evacuation under airborne hazards whose impact changes over space and time. Rather than focusing on a single incident type or an isolated routing method, it combines plume simulation, hazard-field construction, terrain-aware routing, and exposure-aware path computation within a unified workflow. This enables the platform to support a broader class of emergencies, including radiological releases, toxic airborne incidents, wildfire smoke, and related scenarios where a spatio-temporal hazard field can be generated.

1.1 Architecture Overview

As illustrated in Figure 1.5, *HazardEvac* is organized as a layered system that transforms a hazardous release scenario into evacuation guidance that is aware of the evolving airborne hazard. Rather than treating evacuation as a static shortest-path problem, the architecture is designed to account for the fact that hazard conditions change over space and time and therefore influence route safety throughout movement. To achieve this, the system combines plume generation, hazard-field construction, and route computation within a unified workflow. The overall design is structured into three main layers: the *Plume Generation Layer*, the *Hazard Integration Layer*, and the *Routing and Application Layer*. Together, these layers form the main processing chain of the platform, beginning from scenario definition and ending with the computation and presentation of safer evacuation routes.




Figure 1.5: Overview of the three main layers of HazardEvac: plume generation, hazard integration, and routing/application.

The first layer, namely the *Plume Generation Layer*, is responsible for producing the spatio-temporal hazard field that drives the routing process. Through the web-based interface, the user defines the main properties of a hazardous release scenario, including parameters such as pollutant type, release location, release height, emission rate, release duration, and the spatial configuration of the simulation. For radiological scenarios, the selected pollutant is associated with predefined isotope-related properties used in the corresponding calculations. Based on these inputs, the system retrieves the required meteorological data, stores them locally, and converts them into a format compatible with HYSPLIT. The dispersion workflow is then executed in order to generate a sequence of time-indexed concentration outputs that describe the evolving footprint of the airborne hazard over the area of interest. In this way, the first layer provides the predictive hazard information upon which the rest of the platform depends.

The second layer, namely the *Hazard Integration Layer*, transforms the raw plume outputs into routing-ready structures. In this stage, the geographic area selected by the user is discretized into a regular routing grid that serves as the movement domain for the evacuation algorithms. Terrain-aware preprocessing is then applied to identify which cells are traversable and which should be excluded from movement. In the current implementation, water-covered regions such

as sea and lakes are treated as non-traversable and are therefore removed from the routing domain. Once the traversable grid has been established, the hourly plume outputs are aligned with it by projecting concentration values onto the corresponding routing cells. As a result, each traversable cell is associated with a time-indexed concentration series that captures how the hazard concentration changes over time at that location. Through this process, the continuous atmospheric dispersion output is converted into a hazard-aware routing representation that can be directly consumed by the routing engine.

The third layer, namely the *Routing and Application Layer*, is responsible for computing and presenting the evacuation guidance to the user. This layer takes as input the hazard-aware routing grid produced by the previous stage and applies the *Emergency Airborne Release* (EAR) algorithm in order to identify routes that reduce cumulative exposure while preserving practical movement toward safety. EAR operates on the routing graph in a time-dependent manner, evaluating candidate movements according to the concentration levels encountered during movement, the traversal time of each step, and additional plume-aware spatial penalties. In this way, the routing process explicitly accounts for the fact that the hazard evolves during traversal and that route quality cannot be determined by geometry alone. At the application level, this layer is also responsible for presenting the resulting routes and related outputs through the user interface. In the current system design, the platform supports routing requests involving a single destination, multiple waypoints, and multiple candidate destinations, depending on the operational needs of the user.

The architecture of *HazardEvac* was designed to provide a coherent connection between atmospheric hazard prediction and practical evacuation decision support. By separating plume simulation, hazard-grid construction, and exposure-aware route computation into distinct but connected layers, the system achieves a clear and structured workflow in which each component has a well-defined role. This layered organization also provides the basis for the remainder of the thesis, where each architectural component is examined in greater detail.

1.2 Thesis Structure

The remainder of this thesis is organized as follows.

In Chapter 2, *Background and Related Work*, we present the main background and related work that support this thesis. More specifically, we first examine emergency response and public warning applications in order to place *HazardEvac* within the broader context of digital emergency support systems, and we then introduce the core concepts of concentration, exposure, and radiological dose that are used throughout the remainder of the thesis. Next, we review the principal families of plume simulation approaches for hazardous atmospheric dispersion, comparing CFD-based, Gaussian, and Lagrangian models, and explaining the rationale for selecting HYSPLIT as the plume-simulation engine of this work. Finally, we discuss related emergency management and GIS-based support systems, highlighting how existing approaches contribute to hazard monitoring, decision support, and emergency planning, while clarifying the specific gap addressed by *HazardEvac*.

In Chapter 3, we present the problem setting and the overall system view of *HazardEvac*. More specifically, the chapter first formalizes the hazardous evacuation setting considered in this work, defines the routing domain and its graph representation, and describes the structure of the time-dependent hazard field used throughout the system. It then presents the end-to-end workflow of the platform, beginning with scenario definition and plume generation, continuing through hazard integration and routing domain preparation, and concluding with exposure-aware route computation and result presentation. In addition, the chapter explains the inputs, outputs, and interactions of the main architectural components, thereby establishing both the conceptual foundations and the system-level view upon which the remainder of the thesis is built.

In Chapter 4, we examine the *Plume Generation Layer*. This chapter describes the scenario-definition process, the acquisition and preparation of meteorological data, and the execution of the HYSPLIT-based dispersion workflow that produces the time-indexed concentration outputs used by the platform.

In Chapter 5, we present the *Hazard Integration Layer*. More specifically, we describe how the geographic area of interest is transformed into a traversable routing grid, how terrain-aware preprocessing is applied, and how the plume outputs are mapped onto the routing domain in order to construct the hazard-aware representation consumed by the routing engine.

In Chapter 6, we examine the *Routing and Application Layer*. This chapter presents the *Emergency Airborne Release* (EAR) algorithm in detail, including its cost model, routing logic, and supported routing modes. It also explains how the computed evacuation guidance is produced and presented at the application level. In addition, the chapter includes the methodology used for interpreting routing results in terms of exposure and, where applicable, dose-related reporting.

In Chapter 7, we present the experimental testbed and evaluation methodology used. We then describe the implementation environment, the evaluated scenarios, the baseline approaches, and the metrics used to assess the performance of the proposed system. We then analyze the experimental results in order to evaluate the effectiveness of *HazardEvac*.

Finally, Chapter 8 concludes the thesis by summarizing the main findings of this work, discussing its contributions and limitations, and outlining possible directions for future research and system extensions.

2 Background and Related Work

2.1 Introduction


The objective of this chapter is to present the main background and related work that support the proposed *HazardEvac* framework. Since this thesis brings together emergency response, atmospheric dispersion modeling, and hazard-aware evacuation routing, the chapter first introduces the broader emergency response application context and the core concepts of concentration and exposure. It then reviews the main families of plume-simulation approaches and explains the rationale for selecting HYSPLIT as the hazard-modeling component of this work. Finally, it examines the related literature on emergency management systems, evacuation routing, and hazard-aware path planning, thereby positioning *HazardEvac* with respect to existing approaches.

2.2 Background on Emergency Response Applications


2.2.1 Emergency Response and Public Warning Applications

A number of digital systems are already deployed in practice to support evacuation and emergency response situations, although their primary role is typically to communicate official instructions, improve situational awareness, and coordinate public warnings rather than compute evacuation routes. One important category is that of *authority oriented public warning and crisis-management platforms*. These are systems primarily used by emergency agencies and public authorities to organize, manage, and disseminate warnings at scale. Platforms such as *Genasys Protect* and *Everbridge Public Warning* provide multichannel public warning capabilities that allow authorities to issue location specific alerts, communicate protective instructions, manage warning zones, and disseminate updates through websites, mobile applications, broadcast channels, and related communication pathways [16, 17]. In a similar direction, *Google Crisis Response* operates as a large scale crisis information platform by integrating emergency related information into widely used services such as Google Search and Google Maps, where users may receive crisis alerts, safety information, and route impact warnings associated with wildfires, floods, earthquakes, and other major events [18–20].

A second and more user-centered category is that of *citizen-facing emergency applications*. Unlike authority oriented platforms, which are mainly used to manage and distribute warnings, these applications are designed to support individuals directly during a crisis by delivering alerts



(a) FEMA mobile application.



(b) NERV Disaster Prevention application.

Figure 2.1: Examples of deployed citizen-facing emergency support applications. These systems provide alerts, preparedness guidance, map-based situational awareness, and official protective information to the public. Sources: FEMA [21] and NERV [22].

together with practical preparedness and response information. The *FEMA* app, for example, provides real time weather and emergency alerts, preparedness guidance, and shelter related information [23]. The *American Red Cross Emergency* app similarly offers hazard alerts, live weather maps, preparedness content, and shelter or service lookup capabilities [24]. The *Emergency: Severe Weather App* follows the same general approach, focusing on severe-weather alerts, hazard monitoring, storm tracking, preparedness guidance, and shelter information. Outside the United States, the *NERV Disaster Prevention* app provides location-tailored alerts for earthquakes, tsunamis, volcanic eruptions, flooding, landslides, and other related hazards, combining these notifications with a map-centric interface for rapid public awareness [25,26]. Collectively, these applications demonstrate how emergency-support technologies can assist the public not only by notifying users that a hazard exists, but also by helping them interpret the situation and access practical protective resources.

Lastly, a third emerging category is the use of *digital-twin and infrastructure-oriented platforms* for emergency management and resilience analysis. In this context, the *Cyprus Digital Twin (CyDT)* developed by KIOS represents a national, data driven platform intended to support the resilience and management of Cyprus’ critical infrastructure systems. Although such platforms are not designed primarily as public evacuation routing applications, they are important because they illustrate how continuously updated geospatial, infrastructure, and risk related data can be integrated into a unified operational environment for monitoring, analysis, and scenario evaluation [27]. This broader direction is relevant to *HazardEvac*, since it highlights the value of linking dynamic environmental information with decision support rather than treating emergency information as static or isolated data layers.





Figure 2.2: Example of KIOS digital-twin and resilience-oriented platform capabilities, including infrastructure monitoring, environmental risk assessment, and alert-based operational support. Such platforms illustrate how geospatial, infrastructure, and risk-related data can be integrated into a unified environment for analysis and decision support. Source: KIOS [27].

At the same time, the operational emphasis of the above systems remains different from the one addressed in this work. Public warning and crisis information platforms are primarily designed to present alerts, hazard updates, and disruptions that may affect movement, but they do not explicitly compute individualized evacuation routes and do not account for the cumulative exposure that a person may incur while traversing an airborne hazard [18–20, 23–26]. Similarly, digital-twin and infrastructure-oriented platforms emphasize large scale monitoring, data integration, and resilience analysis rather than route computation for evacuees [27]. It is clear that, existing systems help users understand that a hazard exists, where it is located, and what official guidance applies, but they are not designed to provide route guidance through a contaminated environment while explicitly minimizing exposure during movement. In contrast, *HazardEvac* focuses specifically on the movement support problem that arises once hazard awareness has already been established.

2.2.2 Protective Actions in Airborne Hazard Emergencies

Airborne hazard emergencies, such as toxic industrial releases, radiological incidents, and large scale smoke events, require protective actions that aim to reduce the exposure of affected populations as quickly as possible. In such situations, the objective of the authorities is not only to identify the extent of the hazard, but also to determine what form of public action is most appropriate at each stage of the event. Depending on the nature of the release, the meteorological conditions, and the projected development of the plume, authorities may recommend a range of protective measures, including sheltering indoors, controlled evacuation, relocation to safer areas, access restrictions, or exclusion from highly affected zones. These actions are not interchangeable, and the decision regarding which one to apply depends strongly on how the airborne hazard is expected to evolve over time [1,4].



(a) Illustration of shelter-in-place guidance during a radiological emergency. Source: U.S. Environmental Protection Agency (EPA) [28].



(b) Example of first responders operating in protective equipment during a hazardous-material emergency exercise. Source: OPCW [29].

Figure 2.3: Examples of protective action and official emergency response in airborne hazard incidents.

An important point in the context of this thesis is that immediate evacuation is not always the preferred first response. In many airborne contamination events, especially radiological emergencies, guidance often emphasizes the importance of rapid protective sheltering during the initial phase of the release. The rationale is that immediate movement through an active plume may expose individuals to higher contamination levels than remaining indoors for a limited period while the situation is assessed [4]. In practical terms, messages such as *get inside*, *stay inside*, and *stay tuned* reflect the fact that early sheltering can provide an immediate and simple protective action while emergency authorities evaluate the evolving conditions and determine the most suitable follow-up measures.

At the same time, sheltering is not always sufficient as a complete response. Hazardous atmospheric conditions may persist for extended periods of time or they may affect large regions in ways that make later movement necessary. For example, individuals may need to relocate to better shelter, move toward designated reception points, access medical assistance, reunite with family members, or leave areas where contamination remains elevated. In other cases, infrastructure disruption or operational emergency measures may require controlled displacement even after the initial phase of the release has passed. As a result, protective action should not be understood as a simple choice between either sheltering or evacuating. Rather, it is often a staged process in which the recommended action depends on both the current state of the hazard and its expected development over time [1].


This operational reality is particularly important for evacuation planning. If movement occurs while the hazard is still evolving, then the safety of that movement depends not only on the destination, but also on the route followed and the time at which the route is used. A path that is relatively safe at one moment may become significantly more dangerous later as the plume shifts, intensifies, or expands. Conversely, delaying movement for a short period or selecting an alternative route may substantially reduce exposure.

2.2.3 Concentration and Exposure


In order to understand the routing problem addressed in this thesis, it is first necessary to clarify two core concepts that recur throughout the remainder of the document, namely *concentration* and *exposure*. These concepts are central because *HazardEvac* does not treat airborne hazards as simple blocked areas, but as evolving environmental conditions whose intensity changes over space and time. Thus, route quality depends not only on whether a hazard is present, but also on how strong that hazard is and how long an evacuee remains within it. The present subsection therefore introduces the distinction between concentration, exposure, and, in radiological settings, dose, so as to provide the conceptual basis for the hazard representation and route evaluation logic used in the following chapters.

In the context of airborne hazards, *concentration* refers to the amount of hazardous material present in a unit volume of air at a given location and time. In simple terms, it expresses how contaminated the air is. If the concentration is high, the surrounding air contains a larger amount of hazardous substance; if it is low, the substance is more diluted. Concentration is therefore an environmental property of the plume itself.

Exposure, by contrast, refers to the contact of an individual with that hazardous environment over time. Exposure is not only about how contaminated the air is, but also about whether a person is actually within that contaminated area and for how long. A person passing quickly through a region of moderate concentration may receive less exposure than a person who remains



(a) Illustrative example of atmospheric dispersion, airborne concentration distribution, and downwind deposition. Source: Alan Franco [30]; original image credited to bcairquality.ca.



(b) Example of HYSPLIT concentration output showing the spatial distribution of airborne pollutant concentration over a geographic area. Higher contour levels indicate regions of greater airborne concentration near the plume core. Source: NOAA Air Resources Laboratory [31].

Figure 2.4: Illustrative examples of atmospheric dispersion and concentration-field representation. The left figure provides a conceptual view of plume transport and deposition, while the right figure shows a concentration contour map produced by HYSPLIT.

for a longer time in the same area. Thus, concentration describes the hazard in the environment, whereas exposure describes the hazard as experienced by the individual.

This distinction is important for evacuation planning. A plume may exist over a large area, but the actual exposure incurred by an evacuee depends on the specific route followed, the concentration levels encountered along that route, and the duration of traversal.

In radiological scenarios, a further distinction becomes important, namely the distinction between *exposure* and *dose*. While exposure describes contact with radioactive material or radiation in the environment, dose refers more directly to the effect of that exposure on the body. In practical terms, dose is commonly used as the quantity through which radiological impact is interpreted, since it provides a way to relate environmental conditions to possible biological consequences. This is one of the reasons why radiological emergency guidance often expresses protective-action thresholds and reference levels in dose-related units rather than only in terms of airborne concentration [1, 4].

A person may receive a radiological dose through different pathways. One pathway is *external exposure*, where radiation emitted from material in the environment affects the body from outside. Another pathway is *internal exposure*, where radioactive material is inhaled or otherwise enters the body, leading to dose delivery from within. In airborne release scenarios, inhalation is particularly important because evacuees moving through a contaminated plume may breathe in hazardous material while traveling toward safety.




Figure 2.5: Illustration of radiation contamination versus exposure. Source: U.S. Centers for Disease Control and Prevention (CDC) [32].

For the purposes of this thesis, the main routing objective is based on reducing cumulative exposure to the evolving airborne hazard during traversal. Thus the routing engine operates on the hazard intensity represented over the routing domain, rather than directly optimizing biological dose quantities. This design choice keeps the routing process general and allows the same framework to be applied to different airborne hazards whenever a consistent spatio-temporal hazard field is available. At the same time, in radiological scenarios, dose remains an important quantity for interpreting the practical significance of the results. For this reason, dose related measures are used later in the thesis as part of the evaluation and reporting process, where the routing outcomes can be discussed in terms that are meaningful for radiological risk assessment and emergency response.

2.3 Background on Plume Simulation

2.3.1 Computational Fluid Dynamics Models

Computational Fluid Dynamics (CFD) models represent the most detailed class of atmospheric and fluid-flow models considered in evacuation related studies. In general terms, CFD aims to simulate how air moves around objects, structures, and boundaries by resolving the underlying flow behavior over a discretized physical domain. This makes CFD particularly attractive in situations where local geometry strongly influences the movement of hazardous material, such as industrial plants, street canyons, tunnels, building interiors, or other confined and obstacle-rich environments. Because of this level of detail, CFD-based approaches are often regarded as the most physically realistic option when the objective is to understand how a contaminant spreads through a complex local setting [33].




Figure 2.6: Illustrative example of CFD flow behavior in a confined geometry, highlighting the detailed flow structures that can be captured in high-resolution simulations. Source: Technology Gateway [34]

Figure 2.6 illustrates the type of flow behavior that CFD models are capable of capturing. Rather than treating dispersion as a simple large scale plume moving in a single dominant direction, CFD can resolve recirculation zones, vortices, obstacle-induced deflections, and highly localized flow structures. This ability is important in emergency scenarios where the detailed interaction between a hazardous release and the surrounding environment may influence the exposure experienced by nearby individuals. For example, in dense urban settings or industrial sites, small changes in local geometry may alter how contaminants accumulate, disperse, or become trapped in specific regions.

Within the evacuation literature, CFD-based methods are typically used when high spatial realism is more important than rapid computation. They are particularly useful for detailed hazard analysis in local environments, where one wishes to assess how toxic smoke, chemical releases, or other hazardous substances propagate around buildings, machinery, or enclosed spaces. A

representative example is provided by recent work that uses CFD-generated concentration fields to support exposure-aware evacuation planning in hazardous chemical scenarios [35]. In such cases, the fine-grained concentration information produced by the model can be used to evaluate potential evacuation paths with high spatial precision.


Despite these advantages, CFD models also have an important limitation, namely their computational cost. The same high level of detail that makes them physically realistic also makes them expensive to execute, especially when the domain is large, the geometry is complex, or the hazard field must be updated repeatedly over time. For this reason, CFD is often difficult to integrate into real time or large area evacuation-support systems, particularly when route computation must be performed quickly and over regional spatial scales.

2.3.2 Gaussian Plume and Puff Models


Gaussian plume and Gaussian puff models constitute one of the most widely used simplified families of atmospheric dispersion models. In contrast to CFD-based methods, which aim to resolve detailed fluid-flow behavior over complex geometries, Gaussian models rely on idealized assumptions regarding the way contaminants spread in the atmosphere. Their main attraction lies in their computational efficiency. Because they are relatively simple to evaluate, they have been widely used in regulatory analysis, industrial safety studies, air-quality screening, and emergency scenario estimation [36]. For evacuation-related applications, this efficiency is particularly appealing because it allows hazard conditions to be approximated quickly, which is important whenever rapid route or shelter decisions are required.

The basic intuition behind Gaussian models is that, once a hazardous substance is released, it is transported by the mean wind and spreads laterally and vertically due to atmospheric turbulence. Under this formulation, the contaminant cloud is often assumed to follow an approximately Gaussian concentration distribution normal to the main flow direction. Figure 2.7a illustrates the physical idea behind this type of modeling. A source release may initially rise due to buoyancy and momentum, after which it transitions into passive atmospheric dispersion as it moves downwind. In this way, the model seeks to capture the broad shape and spread of the plume without explicitly resolving all of the small-scale fluid-flow structures that would be considered in a CFD simulation.

Within this family, an important distinction exists between *Gaussian plume* and *Gaussian puff* formulations. As illustrated in Figure 2.7b, a Gaussian plume model typically assumes a continuous release under approximately steady meteorological conditions. This makes it suitable when the release is sustained and the atmospheric behavior does not vary too strongly during the time period of interest. A Gaussian puff model, by contrast, represents the release as a sequence of discrete puffs that are transported downwind. This makes puff-based formulations more flex-



(a) Illustrative example of a buoyant source release rising from the source term and transitioning into passive atmospheric dispersion downwind. Source: APSI [37]



(b) Conceptual comparison between Gaussian plume and Gaussian puff representations. The plume model assumes a continuous release, whereas the puff model represents dispersion through discrete released puffs. Source: ResearchGate [38].

Figure 2.7: Illustrative examples of Gaussian-based atmospheric dispersion concepts. The left figure shows the transition from source release to downwind dispersion, while the right figure highlights the difference between continuous plume and discrete puff formulations.

ible in situations where the release is intermittent, finite in duration, or affected by changing meteorological conditions. Although both approaches remain simplified representations of atmospheric transport, the puff model is generally better suited to time-varying releases than the steady-state plume formulation.


The main advantage of Gaussian plume and puff models is therefore their ability to provide rapid approximations of hazard dispersion at low computational cost. This makes them attractive for scenario screening, preliminary emergency analysis, industrial risk studies, and optimization problems in which many candidate solutions must be evaluated quickly. In the evacuation literature, such models have been used to support route planning and shelter-location studies by generating approximate concentration fields that can be integrated with optimization or routing procedures [39]. Their simplicity also makes them appealing when computational resources are limited or when a system must produce approximate results within short time constraints.

At the same time, these models have clear limitations. Their simplifying assumptions reduce their ability to represent complex terrain, obstacle-rich environments, and strongly changing meteorological conditions. They are also less suitable when the hazard evolves in a highly dynamic manner or when detailed local flow behavior significantly influences contaminant transport. As a result, while Gaussian models are effective for fast estimation and broad hazard approximation, they may not provide the level of realism required for evacuation-support systems operating over large, time-varying, and meteorologically diverse geographic regions.


2.3.3 Lagrangian Dispersion Models

Lagrangian dispersion models provide a more flexible representation of atmospheric transport than both CFD-based and Gaussian formulations. Instead of solving the flow field in full detail around every obstacle, or relying on simplified analytical assumptions about plume shape, these models represent the released material through particles or puffs that are transported by the wind field over time. In this way, they are particularly well suited to scenarios in which the hazard evolves continuously under changing meteorological conditions, and where the objective is to model transport over larger geographic regions without incurring the computational cost associated with full CFD simulations [40,41].

The central idea behind this family of models is that the released substance is decomposed into many small elements that move with the atmosphere and spread as time progresses. Their trajectories are influenced by wind, turbulence, and, depending on the model, additional physical processes such as deposition or atmospheric mixing. Figures 2.8a and 2.8b illustrate this principle. Rather than describing the hazard only through a fixed plume outline, the model tracks how released material is transported across space, and the resulting particle or puff distribution can then be aggregated into a concentration field over the region of interest. This makes Lagrangian approaches naturally compatible with time-dependent hazard representation, since the modeled plume evolves explicitly as the underlying atmospheric conditions change.



(a) Illustrative example of a Lagrangian particle-based dispersion representation, where the released material is modeled through transported particles or parcels over the affected urban area. Source: ScienceDirect [42].



(b) Illustrative example of a dispersion footprint represented over a large geographic area, showing how transported particles can be aggregated into a plume-like hazard field. Source: EnviroComp [43].


Figure 2.8: Illustrative examples of Lagrangian dispersion modeling. Such models track the motion of released material through transported particles or puffs, making them suitable for time-varying releases and large scale atmospheric transport.

An important strength of Lagrangian models is that they provide a practical balance between physical realism and computational efficiency. Compared with Gaussian models, they are more suitable for handling time-varying releases and non-steady meteorological conditions. Compared with CFD models, they are much more practical for regional or large area applications, where rapid response and broad spatial coverage are essential. For this reason, Lagrangian dispersion modeling has become particularly attractive in emergency-oriented applications involving radionuclides, chemical releases, wildfire smoke, and other airborne hazards whose effects may extend over substantial spatial scales [44].

Several systems belong to this general modeling family, including HYSPLIT, FLEXPART, NAME, SCIPUFF, and RIMPUFF. Although these models differ in implementation details and target applications, they share the common idea of representing atmospheric transport through moving particles or puffs rather than through purely analytical plume formulas [41, 45]. Their suitability for time-evolving hazards makes them particularly relevant for evacuation support, since the quality of a route depends not only on the spatial extent of the plume, but also on how the hazard changes during traversal. As a result, Lagrangian models are often viewed as especially appropriate for systems that must connect atmospheric hazard prediction with dynamic route planning over large geographic domains.

2.3.4 Selection of HYSPLIT

Among the dispersion-model families discussed above, HYSPLIT was selected in this thesis as the plume simulation engine used to support hazard-aware evacuation planning. HYSPLIT, which stands for *Hybrid Single-Particle Lagrangian Integrated Trajectory*, is a mature atmospheric transport and dispersion modeling system developed by NOAA's Air Resources Laboratory. It has been widely used in studies and operational analyses involving radionuclides, wildfire smoke, volcanic ash, chemical releases, and other airborne hazards [44,46]. This broad adoption is important for the present work, since it indicates that the model is not limited to a single hazard type, but can support a range of emergency-relevant dispersion scenarios.



(a) HYSPLIT, developed by NOAA's Air Resources Laboratory, is a widely used atmospheric transport and dispersion modeling system. Source: NOAA Air Resources Laboratory [47].

(b) Example of HYSPLIT concentration output over a large geographic region, illustrating the type of time-indexed plume footprint that can be used in hazard-aware routing and emergency response analysis. Source: ResearchGate [48].

Figure 2.9: Illustrative examples related to HYSPLIT. The left figure identifies the modeling system selected in this thesis, while the right figure shows the type of concentration output it produces for large scale atmospheric dispersion analysis.

The choice of HYSPLIT is closely related to the overall requirements of *HazardEvac*. The system developed in this thesis requires a model that can generate plume estimates over regional geographic areas, represent how the hazard evolves over time, and do so at a computational cost that remains practical for integration with an evacuation-support framework. In this respect, HYSPLIT provides a suitable balance between physical realism and operational usability. It is more expressive and flexible than simplified Gaussian approaches for time-varying releases and changing meteorological conditions, while also being far more practical than CFD-based simulation for large area routing applications. Since the goal of this thesis is not only to simulate atmospheric transport, but also to connect plume outputs with a routing engine in an end-to-end decision-support workflow, this balance is particularly important.

Figure 2.9b illustrates the type of output that makes HYSPLIT especially useful in this context. The model can generate concentration fields over large domains and across successive time steps, allowing the hazard to be represented as an evolving plume rather than as a single static footprint. Such outputs are well suited to evacuation planning, because they can later be mapped onto a routing grid and used to evaluate how hazard intensity changes across space and time. In other words, HYSPLIT produces exactly the kind of spatio-temporal concentration information required for constructing the time-dependent hazard representation on which the routing process depends.

Another reason for selecting HYSPLIT in this thesis is its practicality from an implementation perspective. Although other Lagrangian systems, such as FLEXPART, also provide strong atmospheric transport capabilities, HYSPLIT was preferred here because it offered a more straightforward workflow for scenario setup, meteorological-data preparation, model execution, and extraction of concentration outputs for integration with the routing framework [45]. This consideration was important in the present work, since the objective was not simply to run a stand-alone plume simulation, but to embed the resulting hazard fields into a larger evacuation-support platform. Ease of use, availability of supporting tools, and a relatively smooth integration path therefore played an important role in the final model choice.

In addition, HYSPLIT has the advantage of being both operationally established and broadly documented. It is supported by a long history of use in atmospheric dispersion studies and emergency-related applications, and its outputs are available in forms that can be processed further for hazard analysis and decision support [46, 49]. For a thesis whose contribution lies in connecting dispersion outputs with exposure-aware evacuation routing, this maturity is particularly valuable. It allows the focus of the work to remain on the integration of hazard modeling with routing and decision support, rather than on the development of a custom dispersion model from scratch.

2.4 Related Work

A substantial body of work has examined how geospatial data, algorithmic analysis, and digital information systems can support emergency management. In most of these works, however, the central objective is not to compute an individualized evacuation route for a person moving through an evolving airborne hazard. Rather, the focus is usually on one or more supporting tasks, such as hazard monitoring, source estimation, route planning for emergency resources, scenario analysis, or public-facing situational awareness. This broader literature is still important for the present thesis, because it shows how spatial information and computational methods can be brought together to support emergency decision making.






Figure 2.10: Representative examples of environmental monitoring, hazard estimation, and inverse-analysis support for airborne or hazardous releases.

One line of research focuses on environmental monitoring and hazard estimation. Zhu *et al.* [52] study the GIS-based design of air-quality monitoring networks for chemical-cluster management. Their goal is to determine where sensing infrastructure should be placed so that hazardous atmospheric conditions can be observed more effectively. In a related direction, Wang *et al.* [7] combine neural-network prediction with the AERMOD system in order to forecast air-contaminant dispersion, while Gao *et al.* [8] examine the spatial and temporal distribution of LPG dispersion concentration in small-scale environments. Other works go one step further toward inverse analysis: Wang *et al.* [53] study hybrid optimization algorithms for estimating hazardous-gas source parameters, and Zhao *et al.* [54] investigate source-seeking for hazardous gas in a chemical cluster under road-network constraints. These studies are valuable because they address how a hazard can be monitored, predicted, or localized. At the same time, they stop before the movement-support problem considered in this thesis, since they do not use those haz-



(a) Example of a GIS-based incident-management interface showing hazard zones and emergency response assets over the affected area. Source: Incident Management Toolkit / HAZUS-style exercise interface [55].



(b) Example of plume information overlaid on geospatial inventory layers for hazard visualization and response planning. Source: FEMA technological hazards material [56].

Figure 2.11: Representative examples of GIS-based emergency-support interfaces for hazardous-event management.

and predictions and measurements to compute evacuation routes for individuals moving through the affected environment.

A second group of works focuses on emergency routing and planning, but mainly from the perspective of logistics, facilities, or large scale coordination. Feng *et al.* [57] propose a GIS-based model for emergency resource scheduling, where the task is to determine suitable delivery routes by considering factors such as road conditions, travel distance, and urgency of demand. Ding *et al.* [58] also formulate emergency logistics scheduling as a multi-objective optimization problem, emphasizing the efficient allocation of limited emergency resources. Similarly, Wang *et al.* [59] study the location-routing problem for emergency medical facilities using deep reinforcement learning, combining facility selection with route planning. These works clearly demonstrate the usefulness of spatial optimization in emergency management, but the unit of analysis is different from that of *HazardEvac*: they focus on distributing resources, selecting facilities, or coordinating operations, rather than on guiding evacuees through a contaminated environment while minimizing the exposure.

There are also works that bring stronger algorithmic support into emergency movement and decision making, but in settings that still differ from the one addressed here. Fu *et al.* [60], for example, combine crowd modeling with reinforcement learning for path optimization in emergency environments, showing how movement strategies can be improved when interactions among agents and obstacles are taken into account. Chondrogiannis *et al.* [6] present a simulation based framework for urban evacuation planning, where evacuation strategies can be assessed through generated simulation models instead of relying only on static network analy-




Figure 2.12: Illustrative example from ASTRO showing spatio-temporal exposure-aware route selection under congestion and outdoor-exposure constraints. Source: Anastasiou et al. [15].

sis. These works are relevant because they move closer to evacuation-path evaluation, but their main concern remains evacuation efficiency and large scale planning rather than exposure-aware guidance under atmospheric contamination.

A particularly relevant study in terms of exposure aware navigation is ASTRO by Anastasiou *et al.* [15]. ASTRO is an A*-based spatio-temporal route optimization framework designed to reduce COVID-19 exposure by predicting congestion in indoor and outdoor spaces and steering users through lower contact areas. This work is especially important because it shows that route planning can be formulated directly around exposure reduction rather than only around distance or travel time. In that sense, it is conceptually closer to *HazardEvac* than traditional evacuation or logistics systems. At the same time, the exposure source in ASTRO is fundamentally different from the one considered in this thesis. ASTRO models exposure in terms of predicted human contact and congestion, whereas *HazardEvac* addresses airborne contamination represented through a time-evolving atmospheric hazard field generated by a plume model. Thus, the two systems share the broader idea of exposure-aware routing, but differ in both the hazard representation and the operational setting.

3 HazardEvac: System Overview

3.1 Introduction

The purpose of this chapter is to present both the problem setting and the system-level view of *HazardEvac*. Since evacuation under airborne hazards depends both on the spatial structure of the environment and on the temporal evolution of the hazard itself, the chapter first introduces the main abstractions used throughout the framework, including the hazardous evacuation setting, the graph-based routing domain, and the time-dependent hazard representation. It then presents *HazardEvac* as an integrated end-to-end platform, describing its overall workflow, the main architectural components, the flow of data across the system, and the role of the *Emergency Airborne Release* (EAR) routing algorithm core within the platform.

3.2 Hazardous Evacuation Setting

In this thesis, we consider that the evacuation will be under airborne hazardous conditions in which an evacuee must move through a traversable geographic region while the hazard evolves over space and time. The scenarios of interest include events such as chemical leaks, radiological releases, and wildfire smoke episodes, where the released material is transported and dispersed by atmospheric processes rather than remaining localized near the source. As a result, the environment through which movement takes place is not static. Instead, the hazard changes continuously as the plume shifts, expands, weakens, or intensifies over the area of interest.

The setting studied in this work assumes that an evacuee begins from a given location within, near, or approaching an affected region and seeks to move toward a desired target. Depending on the situation, this target may correspond to a fixed destination, an intermediate waypoint, or one among several candidate destinations. Movement is assumed to take place over a traversable spatial domain, where some parts of the environment may be inaccessible due to terrain characteristics or other physical constraints. The evacuee may also move under a specific mobility mode, such as walking, bicycling, or driving, which influences the rate at which different parts of the hazardous area are reached. A key characteristic of this setting is that route quality depends not only on the geometric path followed by the evacuee, but also on the time at which each location is reached during traversal. A route that is relatively safe at one moment may become significantly more dangerous later as the plume evolves, while a longer route may lead to lower cumulative exposure if it avoids highly contaminated areas. In this sense, the evacuation problem cannot be treated as a conventional static shortest-path task, where edge costs are assumed




Figure 3.1: Illustrative example of the hazardous evacuation setting considered in this thesis. An evacuee is required to move between selected locations while an airborne hazard plume evolves over the surrounding region, making route safety dependent on both spatial path selection and the time at which movement takes place.

to remain the same throughout the movement process. Instead, it becomes a spatio-temporal decision problem in which both spatial connectivity and the temporal evolution of the hazard must be taken into account.

The objective considered therefore is to support safer movement under dynamically changing airborne hazard conditions. More specifically, is to identify routes that reduce the cumulative exposure incurred during traversal while still preserving practical movement toward safety.

3.3 Routing Domain Representation

In order to support route computation, the geographic area of interest must be transformed into a representation that can be processed algorithmically. The movement domain considered is represented through a grid-based routing abstraction. This representation provides a practical way to describe feasible movement, incorporate terrain constraints, and later associate time-varying hazard information with the corresponding locations of the environment.

Figure 3.2 illustrates this idea. The selected map region is overlaid with a regular grid, so that the area can be treated as a collection of cells rather than as an undivided continuous surface. Each cell is associated with its geographic center, which acts as its representative spatial location.




Figure 3.2: Example of spatial discretization over the area of interest. The selected geographic region is divided into a regular grid of cells, which provides the discrete spatial structure used by the routing domain.

This is useful because hazard values and movement decisions can later be referenced consistently at the level of individual cells. The use of a regular grid also simplifies the definition of neighborhood relations and provides a convenient structure for integrating plume information within the routing space.

An additional benefit of this discretization is that it allows the resolution of the routing domain to be controlled explicitly through the cell size. Smaller cells provide a finer spatial representation and may capture greater local detail, while larger cells reduce computational cost at the expense of spatial precision. The grid-based abstraction therefore offers a flexible balance between representational detail and computational practicality.

3.3.1 Traversability and Connectivity

Once the spatial grid has been constructed, the next step is to determine which cells are actually available for movement. Not all parts of the geographic region are necessarily traversable. Depending on the environment, some cells may correspond to sea, lakes, or other terrain categories through which movement is not possible or should not be considered. For this reason, the routing domain is not defined simply as the full set of grid cells, but rather as the subset of cells that are considered feasible for traversal.

Figure 3.3 illustrates this terrain-aware interpretation of the grid. Cells identified as walkable remain part of the routing domain, whereas cells corresponding to non-traversable terrain are




Figure 3.3: Example of terrain-aware routing-domain construction. Green cells denote traversable locations, while red cells denote non-traversable terrain that is excluded from the routing graph.

excluded. In this way, the grid becomes more than a purely geometric partition of space: it becomes a movement aware representation of the environment. This distinction is important because it ensures that route computation is carried out only over feasible spatial locations, rather than over an idealized surface in which all cells are assumed equally accessible.

Formally, the routing domain is modeled as a graph

$$G = (V, E),$$

where each vertex $g \in V$ corresponds to a traversable grid cell and each edge $(u, v) \in E$ represents the possibility of moving between two neighboring traversable cells. Thus, the set V contains only cells that remain valid after the traversability filtering step, while the set E captures the local adjacency relations among them. In the present work, connectivity is defined using an 8-neighborhood model, meaning that a traversable cell may be connected to its horizontal, vertical, and diagonal neighbors, provided that those neighboring cells are also traversable. Compared with a more restrictive 4-neighborhood representation, this choice provides a more flexible approximation of movement over open geographic space and reduces the directional bias associated with purely orthogonal motion.

3.3.2 Path Representation

Once the routing graph has been defined, an evacuation route can be represented as a path over that graph. Let $s \in V$ denote the start cell and $d \in V$ denote the destination cell. A feasible route is then expressed as an ordered sequence of traversable cells

$$P = \langle g_0, g_1, \dots, g_m \rangle,$$

where $g_0 = s$, $g_m = d$, and each consecutive pair of cells satisfies

$$(g_i, g_{i+1}) \in E \quad \text{for all } i \in \{0, \dots, m-1\}.$$

In other words, a valid path is one in which each step corresponds to a permitted movement between neighboring traversable cells of the routing graph.

3.4 Time-Dependent Hazard Representation

The graph-based routing domain introduced in the previous section provides the spatial structure on which movement can be represented, but it is not sufficient on its own. The reason is that the hazard does not remain fixed over the environment. As the released material is transported and dispersed by atmospheric processes, the affected regions may shift, expand, weaken, or intensify over time. Consequently, the danger associated with a given location cannot be described by a single static value. Instead, it must be represented in a form that captures both its spatial distribution and its temporal evolution. This is expressed through a time-dependent hazard representation, which connects the evolving airborne plume with the discrete routing domain used for evacuation planning.

3.4.1 Continuous Hazard Field

At the continuous level, the airborne hazard can be described through a spatio-temporal concentration field, denoted by

$$C(x, y, t),$$

where (x, y) corresponds to a geographic location and t denotes time. The value of this field expresses the intensity of the hazardous material present in the air at a given place and moment. In this way, the continuous hazard field provides a natural representation of how an airborne release evolves over a geographic region. Rather than treating the hazard as a fixed contaminated zone, it allows the affected area to be described as a time-varying field whose values change as the plume is transported and dispersed by the atmosphere.




Figure 3.4: Example of a continuous airborne hazard field represented through concentration values over a geographic region. Such outputs describe how hazard intensity varies across space at a given time step and form the basis of time-dependent hazard modeling. Source: MeteoInfo [61].

Figure 3.4 illustrates this idea. The plume is represented directly over geographic space, and the concentration values vary across the affected region rather than remaining uniform throughout the footprint. As a result, the hazard encountered by an evacuee depends not only on whether a route intersects the plume, but also on the concentration levels experienced along the way.

3.4.2 Grid-Level Hazard Values

Although the continuous field $C(x, y, t)$ provides a natural description of the evolving plume, the routing process developed in this thesis operates on the discretized graph-based domain introduced earlier. For this reason, the continuous airborne hazard must be translated into a form that is compatible with the routing grid. This is done by associating hazard information with the grid cells that constitute the routing domain. More specifically, each traversable cell g is assigned a time-dependent hazard value

$$c(g, t),$$

which represents the hazard intensity associated with that cell at time t .

Figure 3.5 illustrates this mapping process. The continuous plume footprint is superimposed on the routing grid, and the affected cells can then be identified at a given time step. In practice, this means that the hazard is no longer represented only as a smooth field over geographic space, but also as a sequence of cell-level values defined over the routing domain. As time progresses, both the set of affected cells and their associated hazard intensities may change, producing a




Figure 3.5: Illustrative example of grid-level hazard representation. The plume footprint is superimposed on the discretized routing domain, allowing concentration information to be associated with individual grid cells over time.

time-indexed hazard representation over the graph.

The use of grid-level hazard values is particularly important because it allows route quality to depend on both location and arrival time. A cell that is only weakly affected at one time step may become significantly more hazardous later, while another cell may become safer as the plume moves away or disperses. Therefore, the routing process cannot rely on a single static hazard assignment. Instead, it must reason over a hazard field whose cell-level values evolve over time.

3.5 Problem Statement

We now formalize the routing problem considered. Let

$$G = (V, E)$$

denote the graph-based routing domain, where each vertex $g \in V$ corresponds to a traversable grid cell and each edge $(u, v) \in E$ represents a feasible movement between neighboring traversable cells. Let $c(g, t)$ denote the time-dependent hazard value associated with cell g at time t , and let $s \in V$ denote the initial position of the evacuee. Depending on the operational setting, the movement objective may involve reaching a fixed destination, passing through one or more intermediate waypoints, or selecting one among several candidate destinations. In all cases, movement takes place over the same graph-based and time-dependent hazard domain.

A feasible evacuation route is represented as a path

$$P = \langle g_0, g_1, \dots, g_m \rangle,$$

where $g_0 = s$, the final element corresponds to the chosen target, and each consecutive pair satisfies

$$(g_i, g_{i+1}) \in E \quad \text{for all } i \in \{0, \dots, m-1\}.$$

As the evacuee moves along the path, each visited cell is reached at a particular time that depends on the traversal times of the preceding movements. Since the hazard value of a cell may change over time, the quality of the path depends not only on the sequence of cells it contains, but also on the times at which those cells are encountered. In this sense, the routing problem is inherently time-dependent.

The objective considered in this thesis is to identify a feasible route that reduces the cumulative exposure incurred during traversal while still preserving practical movement toward safety. Informally, this means that the selected path should avoid regions and times of elevated hazard as much as possible, without ignoring the need to make effective progress toward the evacuation goal.

3.6 End-to-End Workflow

HazardEvac operates as an end-to-end workflow that transforms a user-defined hazardous release scenario into route guidance that accounts for the evolving airborne hazard. Rather than treating plume simulation and route planning as isolated tasks, the platform connects them within a unified sequence of processing stages. In this way, the system is able to move from scenario specification, to hazard estimation, to route generation within a single framework.

The workflow of *HazardEvac* is organized into three main phases. The first phase corresponds to *scenario definition*. In this stage, the user specifies the main characteristics of the hazardous release, such as the pollutant type, release location, release duration, and other simulation related parameters. At the same time, the user defines the geographic area of interest, which determines the region over which the plume and the evacuation process will later be analyzed. This phase therefore establishes the initial conditions upon which the remainder of the system depends.

The second phase corresponds to *hazard generation and visualization*. Once the scenario has been defined, the system prepares the meteorological inputs required for dispersion modeling and executes the plume simulation process. The output of this stage is a time-indexed continuous representation of the airborne hazard over the selected region. Rather than producing only a static contamination footprint, the platform generates hazard information that evolves over time, thereby reflecting the dynamic nature of atmospheric releases.

The third phase corresponds to *hazard integration and escape-route generation*. In this stage, the generated hazard outputs are integrated with the user-specified routing domain and the selected evacuation goals, allowing the system to construct a hazard-aware spatial representation over which routes can be computed. Based on this representation, the routing component evaluates feasible movements and produces route guidance that seeks to reduce cumulative exposure while preserving practical progress toward safety. The resulting escape route, together with the associated hazard information, is then presented to the user through the application interface.

An important characteristic of this workflow is that each phase produces outputs that serve as inputs to the next one. The scenario definition phase provides the release and domain parameters needed by the plume generation phase. The plume-generation phase, in turn, provides the hazard information required by the route-generation phase. As a result, the workflow forms a coherent processing chain rather than a collection of disconnected modules. This design is what allows *HazardEvac* to function as a unified evacuation decision-support platform.

3.7 Main Architectural Components

Beyond the end-to-end workflow presented in the previous section, *HazardEvac* is also organized around a set of main architectural components that cooperate to transform a hazardous release scenario into route guidance. At a conceptual level, these components can be grouped according to their role in the system: user interaction, backend orchestration, meteorological-data preparation, plume simulation, intermediate data exchange, and route-generation and visualization. Figure 3.6 summarizes this organization and shows how the main elements of the platform are connected.




Figure 3.6: High-level software architecture of the *HazardEvac* platform. The figure shows the main components involved in user interaction, meteorological-data preparation, plume generation, backend orchestration, local data exchange, and route visualization, together with their principal interactions.

A first major component of the platform is the *user interaction layer*, through which the user defines hazardous release scenarios, selects simulation parameters, specifies the routing domain, and requests route generation. As already shown in the previous chapter, the interface is not limited to displaying final outputs, but also acts as the entry point through which the main scenario parameters of the platform are introduced. In this way, the interface provides the operational bridge between the end user and the internal processing pipeline.

A second component is the *backend orchestration layer*, which coordinates the execution of the main stages of the platform. In the current implementation, this role is performed by the

web-server and application stack shown in Figure 3.6, which receives the user-defined scenario parameters and triggers the corresponding processing tasks. The purpose of this component is to connect the front-end interaction layer with the internal services responsible for meteorological preparation, plume generation, hazard integration, and route computation. Rather than performing only a single isolated task, it acts as the coordinating element of the overall system workflow.

A third important component is the *meteorological-data preparation service*. Since atmospheric dispersion modeling requires meteorological inputs in a compatible format, the system includes a dedicated mechanism for retrieving, validating, caching, and preparing the required weather data for the selected simulation period. This service ensures that the plume-generation stage receives the environmental inputs necessary for executing the atmospheric transport model.

Closely related to this is the *plume-simulation component*, which is centered around HYSPLIT. This component receives the simulation-ready scenario definition together with the prepared meteorological inputs and produces the time-indexed plume outputs that describe the evolving airborne hazard over the area of interest. The output of the plume-simulation stage is subsequently made available to the rest of the system for visualization, hazard integration, and route generation.

A further component of the platform is the *shared data and intermediate storage layer*. As suggested by Figure 3.6, the different stages of the platform exchange information through a common local storage structure. Conceptually, this shared layer allows meteorological files, plume outputs, terrain data, and intermediate routing artifacts to be made available across the different services without requiring each component to regenerate or duplicate the same information independently.

Finally, the *routing and visualization component* forms the stage at which the system converts the generated hazard information into evacuation routes. This component consumes the hazard-aware routing representation built from the previous stages and invokes the EAR routing core in order to compute routes that explicitly account for the spatio-temporal evolution of the plume. It also provides the application-level functionality used to display the resulting routes, and present the corresponding guidance to the end user. Although the detailed operation of this layer is deferred to Chapter 6, it is important at this stage to note that it constitutes the point at which all preceding components converge into the final decision-support output of the platform.

3.8 Data Flow Across the Platform

The operation of *HazardEvac* depends on a structured flow of information across its main components. While the previous section identified the principal architectural elements of the platform, the purpose of the present section is to explain how data is passed from one stage of the workflow to the next. At a high level, this flow begins with user-defined scenario parameters, continues through meteorological-data preparation and plume generation, and ends with hazard-aware route computation and result presentation.






Figure 3.7: Scenario-definition interface of *HazardEvac*. Through this interface, the user specifies the main release parameters, timing, pollutant information, source characteristics, and output-grid configuration that initiate the processing pipeline.

The data flow begins at the user interface, where the main scenario parameters are specified. These include the characteristics of the hazardous release, such as pollutant type, release location, release duration, and simulation timing, together with the geographic area of interest and the routing-related selections made by the user. Once submitted, these inputs are passed to the backend orchestration layer, which interprets them as the initial definition of the simulation and routing task. At this point, the system has not yet generated any hazard information, but it has acquired the configuration needed to begin the processing pipeline.

The next stage of the data flow concerns meteorological preparation and plume generation. Based on the user-defined scenario, the platform retrieves and prepares the meteorological inputs required for atmospheric dispersion modeling. These data are then supplied to the plume-



(a) Example of plume-to-grid integration, where the generated airborne hazard field is aligned with the discretized routing domain and associated with grid cells over the area of interest.



(b) Example of terrain-aware preprocessing, where traversable and non-traversable cells are identified before route computation.

Figure 3.8: Intermediate outputs of the hazard-integration stage. The generated plume is first aligned with the routing grid and then combined with terrain-based traversability information in order to construct a hazard-aware routing representation.

simulation component together with the release definition. The result of this stage is a set of time-indexed plume outputs describing how the airborne hazard evolves over the selected region. The system now moves from input parameters to an explicit hazard representation that can be visualized, processed further, and used in later decision-making stages.

Once the plume outputs have been generated, the data flow continues into the hazard-integration stage. Here, the generated hazard field is aligned with the routing domain so that the evolving plume information can be associated with the corresponding cells of the discretized movement space. At the same time, terrain-related information is used to determine which parts of the routing domain are traversable and which are excluded from movement. The result of this integration step is a hazard-aware routing representation in which the routing graph is no longer purely spatial, but is also associated with time-varying hazard values.

The final stage of the data flow takes place in the routing and presentation components. The hazard-aware routing representation, together with the user-selected start, destination, or way-point information, is passed to the EAR routing core. EAR then computes a route that accounts for both feasible movement and the evolving hazard conditions encountered during traversal. The resulting route, along with the associated hazard context and relevant summary information, is returned to the interface for visualization and user interpretation. Thus, the data flow of *HazardEvac* ends in the generation of evacuation guidance that is informed by the entire upstream processing chain, beginning with scenario definition and ending with route presentation.

3.9 High-Level View of EAR

At a high level, the *Emergency Airborne Release* (EAR) component is the routing core of *HazardEvac*. Its role is to compute evacuation paths over a routing domain that is not only spatially structured, but also associated with hazard values that evolve over time. In contrast to conventional route planning methods, which typically optimize only for distance or travel time, EAR is designed to support movement under hazardous atmospheric dispersion, where the risk of a route depends both on the locations traversed and on the time at which those locations are reached. For this reason, EAR treats route computation as a time-dependent and exposure-aware decision problem.




Figure 3.9: Conceptual illustration of the *Emergency Airborne Release* (EAR) approach. Starting from an initial position, the algorithm seeks a route toward the goal while accounting for the predicted plume trajectory and the time-varying spatial evolution of the hazard.

Figure 3.9 illustrates the main intuition behind this process. The evacuee begins from an initial position and seeks to reach a designated goal, while the surrounding hazard evolves over successive time stages. Rather than moving directly through the regions most strongly affected by the plume, EAR aims to identify a safer route that balances forward progress with reduced exposure. The algorithm evaluates candidate movements by considering both the local hazard conditions and the broader expected development of the plume. This means that movement decisions are influenced not only by immediate cell costs, but also by the anticipated direction and structure of the plume as it propagates over the environment.

4 HazardEvac: Plume Generation Layer

4.1 Introduction

Having presented the overall system view of *HazardEvac* in the previous chapter, we now examine the first operational layer of the platform, namely the *Plume Generation Layer*. This layer is responsible for producing the spatio-temporal hazard representation that drives the remainder of the evacuation workflow. Starting from a user-defined hazardous release scenario, it gathers the required simulation inputs, prepares the corresponding meteorological data, and executes the atmospheric dispersion process in order to generate time-indexed concentration outputs over the area of interest.

This chapter describes the operation of the plume generation layer in a step-by-step manner. We first present the process through which the hazardous release scenario and the main simulation inputs are defined. We then describe the acquisition and preparation of the meteorological data required by the dispersion model. Next, we explain the plume simulation workflow itself, and finally discuss the generated outputs and their role in the overall *HazardEvac* pipeline.

4.2 Scenario Definition and Simulation Inputs

Plume generation in *HazardEvac* begins with the definition of a hazardous release scenario through the system interface. At this stage, the user provides the main parameters that characterize the incident to be simulated, thereby establishing the initial conditions of the plume-generation process. These inputs are necessary because atmospheric dispersion modeling requires not only knowledge of the released substance, but also a precise description of when, where, and under what release conditions the incident occurs.

A first group of inputs concerns the *source term* of the release. In particular, the user specifies the pollutant type, the release location, the release height, the emission rate, and the release duration. Together, these parameters define the hazardous material being emitted and the manner in which it enters the atmosphere. They therefore play a central role in determining the subsequent simulated behavior of the plume. In radiological scenarios, the selected pollutant is associated with predefined isotope-specific properties that are later used in the interpretation of the generated hazard information.

A second group of inputs concerns the *temporal configuration* of the simulation. The user specifies the simulation start date and time, as well as the overall duration of the simulated release

Control File Creator

Simulation Name:
my_simulation

Start Date: 04 / 22 / 2026 **Start Time:** 10 : 36 PM

Duration:
5 Hours

Source Points (0):
Click Map to Add Source
No sources added


Release Height (m):
10
☐ Allow multiple sources

Pollutant:
CS37
Emission Rate: 1.0 Emission Hours: 3

Deposition (optional):
☐ Enable deposition settings

Output Grid (0):
Click Map to Set Grid Center
No grids defined
☐ Allow multiple grids

(a) Definition of the main release and simulation parameters.



(b) Definition of the spatial output grid over the area of interest.

Figure 4.1: Scenario-definition interface of *HazardEvac*. The user specifies the main release parameters, simulation timing, pollutant information, and spatial output-grid configuration used to initialize the plume generation workflow.

and plume evolution. This temporal information is critical because the corresponding meteorological inputs must be retrieved for the relevant time window, and because the resulting hazard field is generated as a time-indexed sequence rather than as a single static output. In this sense, scenario definition already anticipates the time-dependent character of the routing problem addressed later.

A third group of inputs concerns the *spatial setup* of the simulation. In addition to defining the hazardous release itself, the user specifies the output grid over the area of interest. This includes the spatial extent of the simulation domain and the grid-related parameters used to sample the resulting concentration field. Through this step, the user determines the geographic region over which the plume will be generated and analyzed. The spatial setup is particularly important because it affects both the coverage of the simulated plume and the form in which the resulting hazard data can later be integrated with the routing domain.

Based on the above information, *HazardEvac* constructs a simulation-ready scenario description that can be passed to the plume-generation backend. This step transforms the user-defined incident configuration into a structured specification suitable for the atmospheric dispersion modeling.

4.3 Meteorological Data Acquisition and Preparation

To execute the dispersion workflow, *HazardEvac* requires meteorological data in a format that is compatible with HYSPLIT. Atmospheric dispersion modeling depends strongly on the surrounding meteorological conditions, since wind, temperature, and other environmental variables directly influence how the plume evolves over time. For this reason, the plume-generation layer includes a dedicated preparation stage that retrieves, stores, and converts the required weather data before simulation can begin.




Figure 4.2: Meteorological data acquisition and preparation workflow in *HazardEvac*. GFS data are retrieved from NOAA, cached locally, and converted into ARL format for HYSPLIT-compatible plume simulation.

As illustrated in Figure 4.2, the process begins with the acquisition of meteorological inputs from the NOAA Global Forecast System (GFS). These data provide the environmental conditions required for the user-specified simulation period and geographic region. In the current demo implementation, the relevant GFS files are downloaded manually in advance and stored locally before plume generation is initiated. This design reflects the practical fact that meteorological-data retrieval may take a non-negligible amount of time and is therefore handled beforehand in order to keep the interactive workflow responsive. In a future operational version of the platform, this step can be automated through a dedicated service that periodically retrieves the latest meteorological data and makes them readily available for simulation.

Once retrieved, the required meteorological files are stored in the local file system. This local caching serves an important practical role, since it allows previously downloaded data to be reused across repeated or overlapping simulation runs. As a result, the system avoids unnecessary repeated downloads and improves the efficiency of the plume-generation workflow, especially when multiple scenarios are explored for similar time periods.

A further preparation step is required because HYSPLIT does not directly consume the downloaded meteorological files in their original form. Instead, the retrieved GFS data must be converted into the ARL format expected by the simulation backend. *HazardEvac* therefore includes an automated conversion stage that transforms the downloaded meteorological inputs into simulation-ready ARL files. These files are then made available to the plume-simulation process described in the next section.

Through these steps, the meteorological preparation stage transforms raw external weather data into reusable, locally available, and HYSPLIT-compatible inputs. In this way, it provides the environmental foundation upon which the subsequent plume simulation process depends.

4.4 Plume Simulation Workflow

Once the hazardous release scenario has been defined and the corresponding meteorological inputs have been prepared, the plume-generation layer proceeds to the execution of the atmospheric dispersion simulation. At this stage, the previously collected information is combined into a simulation-ready representation that can be processed by HYSPLIT.




Figure 4.3: Plume simulation workflow in *HazardEvac*. The user-defined scenario and the prepared meteorological inputs are provided to HYSPLIT, which executes the dispersion simulation and produces time-indexed concentration outputs for successive simulation hours.

The plume simulation workflow begins by forming the complete input set required by the dispersion backend. This includes the user-defined source-term parameters together with the ARL-formatted meteorological files prepared in the previous section. Based on these inputs, *HazardEvac* instantiates the corresponding simulation configuration and passes it to the HYSPLIT execution environment.

Figure 4.3 provides a high-level view of this workflow. The simulation-ready scenario and the prepared meteorological inputs are supplied to HYSPLIT, which then performs the forward dispersion computation over the specified simulation period. Through this process, the system estimates how the released material evolves over the area of interest as a function of the atmospheric conditions and the incident definition.

An important characteristic of this workflow is that the simulation does not produce a single static contamination snapshot. Instead, it generates a time-indexed sequence of plume states that describe the evolving airborne hazard over the selected region. This time-dependent character is essential for the purposes of *HazardEvac*, since the routing process developed later depends on knowing not only where the hazard is present, but also how it changes over time.

4.5 Generated Outputs and Their Role in the System

The final product of the plume-generation layer is a time-indexed description of the evolving airborne hazard over the selected simulation domain. After the simulation has been executed, *HazardEvac* obtains a sequence of concentration outputs that represent the predicted state of the plume at successive time intervals. In the current implementation, these outputs are produced as hourly CDUMP files at the concentration height specified during scenario definition. Each file therefore captures a particular stage in the temporal evolution of the simulated release and, taken together, the full sequence provides a dynamic representation of the hazard over time.

Each generated CDUMP file records the spatial footprint of the plume for the corresponding simulation hour. More specifically, it stores the geographic locations at which non-zero pollutant concentration is present together with the associated concentration values. As a result, the set of generated files does not merely indicate whether the hazard exists in a region, but provides a structured quantitative description of how the concentration field is distributed over space at each time step. Collectively, these outputs form the spatio-temporal hazard representation produced by the plume-generation layer.

The importance of these outputs lies not only in their role as simulation artifacts, but in the fact that they constitute the main information product consumed by the rest of the platform. First, they provide the basis for the visualization of the generated plume over the area of interest, allowing the user to inspect the predicted extent and temporal development of the airborne haz-




Figure 4.4: Illustrative example of a generated CDUMP output file in *HazardEvac*. Each file contains concentration records for a given simulation time step, including the corresponding time reference, geographic coordinates, and pollutant concentration values. These outputs form the main data product of the plume-generation layer and are later consumed by the hazard-integration layer.

ard. Second, and more importantly for the purposes of this thesis, they serve as the input to the hazard-integration layer. In the next stage of the *HazardEvac* pipeline, the generated concentration footprints are aligned with the routing domain so that time-dependent hazard values can be associated with the corresponding routing cells.

From a system perspective, the generated outputs therefore form the bridge between atmospheric dispersion modeling and evacuation routing. Up to this point, the platform has transformed a user-defined hazardous release scenario into predictive environmental information. In the following chapter, these plume outputs are further processed in order to construct the hazard-aware routing representation used by the EAR component. Thus, the generated outputs of the plume-generation layer are not an end in themselves, but the essential intermediate product that enables the remainder of the *HazardEvac* workflow.

5 HazardEvac: Hazard Integration Layer

5.1 Introduction

Having presented the plume-generation layer in the previous chapter, we now turn to the second operational layer of *HazardEvac*, namely the *Hazard Integration Layer*. While the plume-generation layer produces the time-indexed concentration outputs that describe the evolving airborne hazard, these outputs are not yet in a form that can be used directly for evacuation routing. The role of the present layer is therefore to transform the raw plume simulation results into a routing-ready hazard representation that can later be consumed by the EAR component.

More specifically, the hazard-integration layer bridges plume simulation and route computation. Its task is to instantiate the evacuation domain over the user-selected region, determine which parts of that domain are traversable, and align the simulated plume with the resulting routing grid. In this way, the continuous and time-varying hazard information produced by HYSPLIT is converted into a structured representation in which each routing cell can be associated with the corresponding hazard values over time.

This chapter describes the operation of the hazard-integration layer in a step-by-step manner. We first provide an overview of the layer and its role in the overall system pipeline. We then explain how the routing domain is constructed, how terrain-aware preprocessing is used to identify traversable and non-traversable cells, and how the plume outputs are mapped onto the routing grid. Finally, we describe how these elements are combined in order to construct the hazard-aware routing representation that forms the direct input to the routing and application layer presented later in the thesis.

5.2 Overview of the Hazard Integration Layer

The *Hazard Integration Layer* constitutes the intermediate stage between plume simulation and route computation in *HazardEvac*. While the plume-generation layer produces the time-indexed concentration outputs that describe the evolving airborne hazard, the routing layer requires a structured representation in which those hazard values can be associated directly with the movement domain. The role of the hazard-integration layer is therefore to bridge these two stages by transforming the raw dispersion outputs into a routing-ready hazard field.




Figure 5.1: Overview of the hazard-integration workflow in *HazardEvac*. The layer first defines the routing domain over the selected geographic region, then applies terrain-aware preprocessing to identify traversable and non-traversable cells, and finally integrates the generated plume with the routing grid in order to construct the hazard-aware routing representation used by the routing layer.

This layer performs three main tasks. First, it defines the routing domain over the geographic area selected by the user. This step establishes the spatial substrate on which evacuation movement will later be computed. Second, it applies terrain-aware preprocessing in order to distinguish traversable from non-traversable regions, thereby ensuring that the resulting routing space reflects the practical constraints of the environment. Third, it aligns the plume outputs produced by HYSPLIT with the routing grid so that each traversable cell can be associated with the corresponding hazard values over time.

The result of this process is a hazard-aware routing representation in which the evacuation domain is no longer purely spatial, but is also associated with time-dependent hazard information. In this representation, each traversable cell is linked to a time-indexed series of concentration values derived from the plume outputs. This final structure forms the direct input to the routing and application layer, where the EAR algorithm evaluates feasible movements under dynamically evolving hazard conditions.

5.3 Routing Domain Construction

The first step of the hazard-integration layer is to define the spatial routing domain over which evacuation movement will later be computed. In *HazardEvac*, this domain is specified interactively by the user through the map interface. Rather than relying on a fixed preconfigured region, the platform allows the user to select the geographic area of interest directly, thereby adapting the routing domain to the specific hazardous scenario under consideration.




Figure 5.2: Example of routing-domain construction over the whole island of Cyprus

More specifically, the routing domain is defined by drawing a rectangular area on the map. After activating the corresponding control, the user selects two opposite corner points of the desired region. From these two map clicks, the system computes the geographic bounds of the selected area and stores them as $B = [\phi_{\min}, \phi_{\max}] \times [\lambda_{\min}, \lambda_{\max}]$, where ϕ denotes latitude and λ denotes longitude. These bounds determine the spatial extent within which the routing grid will be constructed.




Figure 5.3: Schematic representation of the routing-domain construction. Generated with assistance from OpenAI ChatGPT image generation tools.

Once the rectangular domain has been defined, *HazardEvac* discretizes it into a regular spatial grid. The resolution of this grid is controlled by a user-defined cell size s , expressed in meters. Since the selected domain is represented in geographic coordinates, this cell size must be converted into corresponding angular steps in latitude and longitude. The latitude step is approximated as

$$\Delta\phi = \frac{s/1000}{111.32},$$

where 111.32 denotes the approximate number of kilometers per degree of latitude. Let

$$\phi_c = \frac{\phi_{\min} + \phi_{\max}}{2}$$

be the latitude at the center of the selected domain. The longitude step is then adjusted according to

$$\Delta\lambda = \frac{\Delta\phi}{\cos(\phi_c)}.$$

Using these angular step sizes, the system computes the grid dimensions as

$$W = \left\lceil \frac{\lambda_{\max} - \lambda_{\min}}{\Delta\lambda} \right\rceil, \quad H = \left\lceil \frac{\phi_{\max} - \phi_{\min}}{\Delta\phi} \right\rceil,$$

where W and H denote the number of cells along the longitude and latitude directions, respectively. In this way, the continuous user-selected geographic region is transformed into a discrete

routing grid whose resolution is controlled directly by the selected cell size.

The resulting routing domain can therefore be viewed as a regular set of cells $G = \{g_{x,y} \mid 0 \leq x < W, 0 \leq y < H\}$, where each cell $g_{x,y}$ corresponds to a small geographic region within the selected area. At this stage, the grid represents only the spatial structure of the routing domain itself. Issues of traversability, terrain exclusion, and plume association are addressed in the subsequent steps of the hazard-integration layer.

For visualization purposes, the constructed routing domain is rendered directly on the map as a grid overlay. The platform displays both the outer rectangular boundary of the selected area and the internal vertical and horizontal grid lines, thereby allowing the user to inspect the discretized movement space before proceeding with route setup. In addition, when a new routing domain is drawn, the system clears any previously defined route, start point, and goal markers. This ensures that the routing state remains consistent with the newly selected area and prevents stale route information from being reused across different planning domains.

5.4 Terrain-Aware Preprocessing

After the routing domain has been discretized into a regular grid, *HazardEvac* performs terrain-aware preprocessing in order to enrich the grid with environmental information relevant to route computation. The routing grid defined in the previous section provides only a geometric discretization of the selected area; it does not yet distinguish between terrain types that support efficient movement and those that impose significant mobility penalties. Terrain-aware preprocessing therefore transforms the routing domain from a purely spatial grid into a more realistic movement space.

In the current implementation, terrain information is obtained from ESA¹ WorldCover raster data stored locally as GeoTIFF tiles. Since these data are maintained as multiple terrain files rather than as a single global raster loaded in memory, the backend must first determine which files are relevant to the user-selected routing domain. Let the selected routing region be defined by the geographic bounds $B = [\phi_{\min}, \phi_{\max}] \times [\lambda_{\min}, \lambda_{\max}]$, where ϕ denotes latitude and λ denotes longitude. Each available terrain tile T_i is associated with its own spatial extent $E_i = [\phi_i^{\min}, \phi_i^{\max}] \times [\lambda_i^{\min}, \lambda_i^{\max}]$. A tile is considered relevant if its extent intersects the selected routing region, that is, $E_i \cap B \neq \emptyset$. Equivalently, a terrain tile is selected whenever:

$$\phi_i^{\max} \geq \phi_{\min}, \quad \phi_i^{\min} \leq \phi_{\max}, \quad \lambda_i^{\max} \geq \lambda_{\min}, \quad \lambda_i^{\min} \leq \lambda_{\max}.$$

¹European Space Agency

Using this criterion, the backend resolves only the subset of terrain rasters whose spatial coverage overlaps the routing domain, rather than loading all available terrain data globally. This keeps terrain preprocessing focused on the area relevant to the current evacuation task and avoids unnecessary loading of unrelated raster tiles.

Once the relevant terrain rasters have been identified, the system assigns terrain classes to the routing grid on a per-cell basis. More specifically, for each routing cell $g_{x,y}$, the geographic coordinates of its center are computed and used as the representative spatial location of that cell. Let this cell-center mapping be denoted by $\gamma(g_{x,y}) = (\phi_{x,y}, \lambda_{x,y})$. For each such location, the backend samples the corresponding raster value from the selected WorldCover tile covering that point. If the terrain raster is stored in a coordinate reference system different from WGS84, the cell-center coordinates are first transformed into the coordinate system of the raster before sampling is performed.




Figure 5.4: Example of terrain-aware preprocessing in *HazardEvac*. After the routing domain is discretized into a regular grid, terrain information derived from ESA WorldCover data is used to classify cells and distinguish movement conditions across the routing space.

Let the discretized routing grid be denoted by $G = \{g_{x,y} \mid 0 \leq x < W, 0 \leq y < H\}$. After terrain classification, each cell is associated with a terrain label through a mapping $\tau : G \rightarrow \mathcal{T}$, where $\mathcal{T}$ denotes the set of terrain classes used internally by the system. Based on this terrain type, the preprocessing stage derives a terrain-dependent movement modifier $\mu : G \rightarrow \mathbb{R}_{>0}$, which captures the relative ease of traversal over that cell.

In the present implementation, the terrain-speed factors are defined as

$$\mu(g) = \begin{cases} 1.3 & \text{if } \tau(g) = \text{road} \\ 1.2 & \text{if } \tau(g) = \text{urban} \\ 1.0 & \text{if } \tau(g) \in \{\text{unknown, agriculture, grass}\} \\ 0.6 & \text{if } \tau(g) = \text{forest} \end{cases}$$

so that roads and urban cells favor faster traversal, while forests impose strong movement penalties. Lastly, wetlands and water-covered terrain are considered impassable. Consequently, even when two cells are geometrically adjacent, their effective traversal properties may differ significantly depending on the terrain they represent.

The result of this preprocessing step is a terrain-annotated routing grid in which the cells are no longer treated as uniformly traversable. Instead, each cell carries terrain metadata that later influences traversal time and, more generally, route cost. Figure 5.4 illustrates this idea by showing the routing domain after terrain classification, where different terrain conditions are reflected across the discretized movement space.

5.5 Plume to Grid Mapping

After the routing domain has been constructed and enriched with terrain information, the next step is to align the plume outputs with the routing grid. This stage is necessary because the atmospheric dispersion model and the routing engine operate over different spatial representations. On the one hand, HYSPLIT produces hourly concentration outputs over its own simulation lattice, expressed as geographically referenced plume records. On the other hand, the routing layer requires hazard values attached directly to the cells of the routing grid. Plume to grid mapping resolves this mismatch by transforming the hourly plume outputs into cell-level hazard values that can be queried during route computation.

Let a plume record at time t be denoted by $p = (\phi_p, \lambda_p, v_p)$, where ϕ_p and λ_p are the geographic coordinates of the plume output point and v_p is the corresponding concentration value. In the implementation of *HazardEvac*, a plume record is not interpreted merely as a point sample. Instead, each such record is treated as representing a small rectangular spatial footprint over the simulation grid. The dimensions of this footprint are estimated from the spacing of the underlying HYSPLIT source lattice, yielding an approximate latitude and longitude extent $\Delta\phi_p, \Delta\lambda_p$.

Accordingly, the plume record is interpreted as covering the rectangular region

$$R_p = \left[\phi_p - \frac{\Delta\phi_p}{2}, \phi_p + \frac{\Delta\phi_p}{2} \right] \times \left[\lambda_p - \frac{\Delta\lambda_p}{2}, \lambda_p + \frac{\Delta\lambda_p}{2} \right]$$

Let the routing grid be denoted by $G = \{g_{x,y} \mid 0 \leq x < W, 0 \leq y < H\}$, and let the center of each routing cell $g_{x,y}$ be given by $\gamma(g_{x,y}) = (\phi_{x,y}, \lambda_{x,y})$.

The mapping rule used by the system is based on cell-center inclusion. A plume record p is said to contribute to routing cell $g_{x,y}$ if the center of that cell lies inside the rectangular plume footprint R_p , that is, $\gamma(g_{x,y}) \in R_p$.

Equivalently, the contribution condition is

$$\phi_p - \frac{\Delta\phi_p}{2} \leq \phi_{x,y} \leq \phi_p + \frac{\Delta\phi_p}{2}$$

and

$$\lambda_p - \frac{\Delta\lambda_p}{2} \leq \lambda_{x,y} \leq \lambda_p + \frac{\Delta\lambda_p}{2}.$$

In this way, the continuous plume output is projected onto the planner grid using a spatial rule that preserves the local footprint represented by each HYSPLIT output record.

Since multiple plume records may overlap the same routing cell during a given simulation hour, the system must aggregate their contributions into a single cell-level hazard value. Let $\mathcal{P}_{g,t}$ denote the set of plume records at hour t whose footprints contain the center of routing cell g . The concentration value associated with cell g at hour t is then defined as

$$c(g, t) = \begin{cases} \frac{1}{|\mathcal{P}_{g,t}|} \sum_{p \in \mathcal{P}_{g,t}} v_p, & |\mathcal{P}_{g,t}| > 0 \\ 0, & |\mathcal{P}_{g,t}| = 0. \end{cases}$$

Thus, if one or more plume records contribute to a routing cell during a given hour, their concentration values are averaged into a single representative cell-level value. If no plume footprint covers the center of that cell at that hour, the assigned concentration is zero.




Figure 5.5: Illustration of plume-to-grid mapping in *HazardEvac*. Hourly plume outputs are projected onto the routing grid by associating plume footprints with the corresponding routing cells. The resulting cell-level concentration values form the time-dependent hazard field used by the routing layer.

By repeating this process for all routing cells and for all simulation hours, *HazardEvac* constructs a time-indexed hazard field over the routing grid. That is, each cell $g \in G$ becomes associated with a sequence $\{c(g, t_1), c(g, t_2), \dots, c(g, t_n)\}$, where $t_1, \dots, t_n$ denote the discrete simulation hours of the plume outputs. The result of this step is therefore a routing-ready hazard representation in which the spatial discretization used by the routing engine is directly coupled with the temporal evolution of the plume.

Figure 5.5 illustrates how hourly plume outputs are projected onto the routing grid to derive cell-level concentration values for the movement domain. This stage produces the function $c(g, t)$ required by the EAR routing algorithm, allowing *HazardEvac* to evaluate time-varying hazard intensity at the level of individual routing cells and connect atmospheric dispersion outputs with exposure-aware route computation.

6 HazardEvac: Routing and Application Layer

6.1 Introduction

Having described in the previous chapter how the plume outputs are mapped onto the routing grid and combined with terrain information, we now turn to the final operational layer of *HazardEvac*, namely the *Routing and Application Layer*. This layer consumes the hazard-aware routing representation constructed by the hazard-integration layer and uses it to compute the evacuation guidance.

The routing core of this layer is the *Emergency Airborne Release* (EAR) algorithm. EAR operates on the traversable routing graph and evaluates candidate movements by jointly considering multiple factors, including the concentration levels encountered during traversal, the traversal time required for each move, and plume-aware spatial penalties that reflect the predicted structure and evolution of the hazard field. As a result, route computation is no longer based solely on geometric distance or shortest travel time. Instead, the quality of a route depends on both where the evacuee moves and when each part of that route is reached. This makes routing inherently time dependent and requires an algorithm that can explicitly account for the spatio-temporal evolution of the plume during search.

Beyond route computation itself, this layer also includes the application-level functionality that connects EAR to the user-facing platform. Once a route is computed, the system must present it in a form that is meaningful and actionable, together with the associated hazard context and route-level summary information. In addition, since the routing process itself is performed over a concentration-based exposure proxy, this layer is also responsible for the post-processing steps through which route outcomes can be interpreted in dose-related terms for reporting and risk interpretation. This design preserves the generality of the routing engine while still allowing the system to present outputs in units that are operationally meaningful in radiological and similar airborne hazard scenarios.

In this chapter, we first describe the routing objectives and the time-dependent cost model used by EAR. We then present the EAR algorithm itself, first conceptually and then through a small worked example that illustrates how the algorithm evaluates alternative movements. Finally, we explain how computed routes are generated and presented at the application level, and how concentration-based routing outputs are subsequently interpreted and reported in dose-related terms.

6.2 Routing Objectives and Cost Model

The routing objective in *HazardEvac* is to compute evacuation routes that reduce the cumulative hazard encountered during movement while still preserving practical progress toward safety. This objective differs from conventional shortest-path routing, where the route is typically optimized only with respect to distance or travel time. In the present setting, movement takes place inside a hazard field whose intensity evolves over time. Consequently, the quality of a route depends not only on the geometric sequence of cells that it traverses, but also on the time at which each cell is reached. A route that appears acceptable at one moment may become significantly more hazardous later as the plume expands, shifts, or intensifies.




Figure 6.1: Illustration of the EAR cost model. At a given routing state, EAR evaluates the neighboring candidate cells by combining multiple cost components, including the exposure term, traversal time, directional penalty, proximity penalty, and heuristic guidance toward the goal. The resulting combined score determines which candidate movements are favored during route computation.

Let the hazard-aware routing domain be represented as a graph

$$G = (V, E),$$

where each vertex $g \in V$ corresponds to a traversable routing cell and each edge $(u, v) \in E$ connects two neighboring traversable cells. Let $s \in V$ denote the start cell, $d \in V$ the destination cell, and let t_0 denote the departure time. A feasible evacuation path is then represented as

$$P = \langle g_0, g_1, \dots, g_m \rangle,$$

where $g_0 = s$, $g_m = d$, and $(g_i, g_{i+1}) \in E$ for all $i \in \{0, \dots, m-1\}$. The objective of EAR is not merely to find a feasible path from s to d , but to identify one whose cumulative routing cost remains low under the spatio-temporal evolution of the hazard.

A key feature of this setting is that edge costs are time dependent. Consider a transition from cell u to a neighboring cell v beginning at time t . Let $\Delta t(u, v)$ denote the traversal time required for this move. The exposure incurred during the transition is approximated by the average of the concentration at departure and the concentration at arrival, namely

$$\exp(u, v, t) = \frac{c(u, t) + c(v, t + \Delta t(u, v))}{2},$$

where $c(g, t)$ denotes the concentration associated with cell g at time t . This formulation captures the fact that the evacuee is exposed both to the hazard level at the current location and to the hazard level encountered upon reaching the next cell.

EAR augments this exposure term with additional routing components that reflect the operational realities of evacuation. Exposure minimization alone is not sufficient to produce practically useful routes. A route that avoids the highest concentration cells may still be undesirable if it is excessively slow, if it continues in an unfavorable direction relative to the advancing plume, or if it remains too close to the main plume trajectory for an extended portion of the traversal. For this reason, EAR evaluates each candidate movement through four complementary criteria: the exposure incurred during the transition, the traversal time of the movement, a directional penalty relative to the plume trajectory, and a proximity penalty expressing standoff preference from the plume core. These terms are then normalized and combined into a single time-dependent edge cost.

The traversal-time component $\Delta t(u, v)$ discourages unnecessarily slow movements through the hazardous region. This term is important because cumulative exposure is inherently influenced by residence time: even when two candidate routes traverse areas of similar concentration, the one that keeps the evacuee inside the contaminated environment for longer will generally be less

desirable. In the present implementation, traversal time is explicitly terrain-aware. The algorithm does not assume a single fixed movement speed for the entire routing domain. Instead, it begins with a transport-mode base speed v_{mode} , determined by the selected mobility assumption, such as walking, bicycling, or driving, and then modifies this speed according to the terrain type of the traversed cell.

Let $\tau(g)$ denote the terrain class of cell g , and let $\mu(g)$ denote the terrain-speed factor associated with that terrain class. In the current implementation,

$$\mu(g) = \begin{cases} 1.3 & \text{if } \tau(g) = \text{road}, \\ 1.2 & \text{if } \tau(g) = \text{urban}, \\ 1.0 & \text{if } \tau(g) \in \{\text{unknown}, \text{agriculture}, \text{grass}\}, \\ 0.6 & \text{if } \tau(g) = \text{forest}, \end{cases}$$

Thus, the effective speed associated with cell g is

$$v_{\text{eff}}(g) = v_{\text{mode}} \cdot \mu(g).$$

This means that the same transport mode induces different actual traversal speeds depending on the terrain of the cell. Roads and urban cells increase the effective speed relative to the base assumption, while forest cells impose severe mobility penalties.

The implementation further models traversal between adjacent cells as a two-part movement. Specifically, moving from cell u to cell v is decomposed into: (i) travel from the center of u to the edge of u , and (ii) travel from the edge of v to the center of v . If L denotes the side length of a square routing cell in kilometers, then the center-to-edge distance is $L/2$. The corresponding center-to-edge traversal time for a cell g is therefore

$$t_{\text{ce}}(g) = \frac{L/2}{v_{\text{eff}}(g)} \cdot 60,$$

where the factor 60 converts hours to minutes. Using this construction, the total traversal time from u to v is computed as

$$\Delta t(u, v) = t_{\text{ce}}(u) + t_{\text{ce}}(v).$$

Equivalently,

$$\Delta t(u, v) = \frac{L/2}{v_{\text{mode}}\mu(u)} \cdot 60 + \frac{L/2}{v_{\text{mode}}\mu(v)} \cdot 60.$$

The directional component evaluates how the candidate movement is oriented relative to the dominant plume trajectory. The motivation behind this term is predictive rather than purely geometric. If a candidate move continues in the same general direction as the plume propagation, it may keep the evacuee within the advancing hazard field for longer, or even guide the route deeper into the path of the plume. EAR therefore introduces a directional penalty that discourages such movements. Let $\mathbf{m}_{u,v}$ denote the movement vector from cell u to cell v , and let $\mathbf{T}$ denote the dominant plume-trajectory vector. A natural measure of alignment is the cosine similarity

$$\cos \theta(u, v) = \frac{\mathbf{m}_{u,v} \cdot \mathbf{T}}{\|\mathbf{m}_{u,v}\| \|\mathbf{T}\|}.$$

Large positive values indicate stronger alignment with the plume trajectory, whereas lower or negative values indicate divergence from it. In practice, EAR uses a normalized directional penalty,

$$\widetilde{\text{dir}}(u, v),$$

which increases when the candidate move is judged to be unfavorable with respect to plume propagation. The purpose of this term is to bias the search away from moves that risk tracking the plume's expected advance.

The proximity component captures the standoff preference from the plume core. Whereas the directional term evaluates how the movement is oriented relative to the plume trajectory, the proximity term evaluates how close the candidate cell lies to that trajectory. This distinction is important. A move may not be strongly aligned with the plume direction, yet still remain dangerously close to the central region of plume propagation, where concentration is likely to remain elevated. Let $d_{\perp}(v, \mathbf{T})$ denote the perpendicular distance of candidate cell v from the dominant plume trajectory. EAR then defines a normalized proximity penalty,

$$\widetilde{\text{prox}}(v) = f(d_{\perp}(v, \mathbf{T})),$$

where $f(\cdot)$ is a decreasing function of distance. Cells close to the trajectory receive larger penalties, while cells farther away receive smaller ones. The operational purpose of this term is to encourage routes that maintain separation from the most hazardous region of plume propagation, even when local concentrations alone might not be sufficient to enforce that behavior.

To combine these heterogeneous components into a single transition score, EAR normalizes the exposure, traversal-time, directional, and proximity terms, denoted respectively by

$$\widetilde{\text{exp}}(u, v, t), \quad \widetilde{\Delta t}(u, v), \quad \widetilde{\text{dir}}(u, v), \quad \widetilde{\text{prox}}(v).$$

The resulting time-dependent edge cost is defined as

$$w(u, v, t) = w_c \widetilde{\text{exp}}(u, v, t) + w_t \widetilde{\Delta t}(u, v) + w_d \widetilde{\text{dir}}(u, v) + w_p \widetilde{\text{prox}}(v),$$

where w_c , w_t , w_d , and w_p are user-defined weights that control the relative influence of the four cost components. By adjusting these weights, the routing behavior can be made more hazard-averse, more time-sensitive, or more strongly biased toward plume-avoidance geometry.

For a path

$$P = \langle g_0, g_1, \dots, g_m \rangle,$$

the cumulative routing cost is obtained by summing the corresponding edge costs along the traversal sequence:

$$J(P) = \sum_{i=0}^{m-1} w(g_i, g_{i+1}, t_i),$$

where t_i denotes the time at which the transition from g_i to g_{i+1} begins. Since each t_i depends on the traversal times of the preceding moves, the cost of the path is inherently time dependent. Thus, even two geometrically similar paths may incur substantially different cumulative costs if they interact differently with the evolving plume.

In conclusion, the above formulation defines the routing objective used throughout the remainder of this chapter. In essence, EAR seeks routes that balance progress toward safety with reduced cumulative hazard, while explicitly accounting for the fact that both hazard intensity and movement conditions may vary across space and time. The resulting cost model therefore provides the basis for the algorithm presented in the next section, where these components are integrated within a best-first graph-search procedure.

6.3 The Emergency Airborne Release (EAR) Algorithm

Having defined the routing objective and the cost components used by the routing layer, we now describe the *Emergency Airborne Release* (EAR) algorithm itself. EAR is the routing core of *HazardEvac* and operates as a best-first graph-search procedure over the hazard-aware routing domain. Its purpose is to compute feasible evacuation routes that do not merely connect a start cell to a goal cell, but do so in a way that reduces exposure to the evolving hazard while still maintaining practical movement toward safety.

At a high level, EAR starts from the source cell and progressively expands candidate cells according to their current routing priority. For each expanded cell, the algorithm examines its feasible neighbors, evaluates the local movement conditions, estimates the traversal implications of moving into each neighbor, and updates the best known route to that neighbor whenever an improvement is found. The process continues until the destination is reached or until no feasible routing states remain. Once the goal has been found, the final route is reconstructed by following the stored predecessor links backward from the destination to the start.

Algorithm 1 EAR exposure-aware route computation

Require: Grid $G = (V, E)$ **Require:** start cell s , goal cell g **Require:** start time t_0 (minutes)**Ensure:** Path P and total cost $J(P)$ (or failure)

```
1: Initialize priority queue  $Q \leftarrow \{(f = 0, \text{tie} = 0, s)\}$   $\triangleright$  start from initial cell
2:  $\text{cameFrom} \leftarrow \emptyset$   $\triangleright$  predecessor map
3:  $gScore[s] \leftarrow 0$   $\triangleright$  best known cumulative cost
4:  $\text{minutes}[s] \leftarrow \max(0, t_0)$   $\triangleright$  arrival time at initial cell
5:  $\mathbf{T} \leftarrow \text{InitTrajectoryVector}()$   $\triangleright$  compute predicted plume trajectory
6:  $\text{goalSide} \leftarrow \text{Side}(g, \mathbf{T})$   $\triangleright$  determine on which side of the plume trajectory the goal lies
7: while  $Q \neq \emptyset$  do
8:    $(\_, \_, u) \leftarrow \text{PopMin}(Q)$   $\triangleright$  expand lowest-score state
9:    $t_u \leftarrow \text{minutes}[u]$ 
10:  if  $u = g$  then
11:    output  $\leftarrow \text{ReconstructPath}(\text{cameFrom}, u)$ 
12:    stop
13:   $\mathcal{N}(u) \leftarrow \text{Neighbors}(u)$   $\triangleright$  feasible adjacent cells
14:  for all  $v \in \mathcal{N}(u)$  do
15:     $(\text{exp}, \Delta t) \leftarrow \text{GetMoveCost}(u, v, t_u)$   $\triangleright$  predicted exposure and traversal time
16:    if  $\Delta t = \infty$  or  $\text{exp} = \infty$  then
17:      continue  $\triangleright$  discard infeasible move
18:     $t_v \leftarrow t_u + \Delta t$   $\triangleright$  arrival time at neighbor
19:     $\text{side}_v \leftarrow \text{Side}(v, \mathbf{T})$ 
20:    if  $\text{side}_v \cdot \text{goalSide} < 0$  then
21:       $\text{dir} \leftarrow \text{NormVectorDirCost}(u, v, \mathbf{T})$   $\triangleright$  apply directional penalty
22:    else
23:       $\text{dir} \leftarrow 0$ 
24:     $\text{prox} \leftarrow \text{NormDistToTraj}(v)$   $\triangleright$  evaluate proximity to plume trajectory
25:     $w \leftarrow \text{CostFunction}(\Delta t, \text{exp}, \text{dir}, \text{prox})$   $\triangleright$  local edge cost
26:     $gNew \leftarrow gScore[u] + w$   $\triangleright$  new cumulative cost
27:    if  $gNew < gScore.get(v, \infty)$  then
28:       $\text{cameFrom}[v] \leftarrow u$   $\triangleright$  update predecessor
29:       $gScore[v] \leftarrow gNew$ 
30:       $\text{minutes}[v] \leftarrow t_v$ 
31:       $h \leftarrow w_h \cdot \text{HeuristicNormalized}(v, g)$   $\triangleright$  goal-directed heuristic
32:       $fNew \leftarrow gNew + h$   $\triangleright$  priority score
33:       $\text{tie} \leftarrow \text{tie} + 1$ 
34:       $\text{Update/PushQ}(fNew, \text{tie}, v)$   $\triangleright$  insert or update queue entry
35: output failure,  $\infty$   $\triangleright$  no feasible route found
```

The logic of Algorithm 1 can be understood in a step-by-step manner. In lines 1–6, EAR initializes the main search structures. The priority queue is seeded with the start cell, the predecessor map is prepared for later route reconstruction, the cumulative score of the start state is set to zero, and the initial time is recorded. At the same time, the algorithm computes the dominant plume-trajectory information that is later used to evaluate plume-aware movement penalties.

Lines 7–13 define the main search loop. At each iteration, EAR removes from the queue the routing state with the currently best priority score. If the extracted state corresponds to the goal cell, the search terminates and the route is reconstructed by tracing the predecessor map backward. This is the point at which EAR converts the explored search history into the final evacuation path returned to the application.

Lines 14–18 identify the feasible neighboring cells of the currently expanded state and perform the first stage of local evaluation. For each neighbor, EAR estimates the traversal time and the predicted exposure associated with moving from the current cell into that candidate cell at the current time. If the move is not feasible, for example because the cell is blocked or the transition cannot be evaluated properly, the algorithm simply discards it and continues to the next candidate.

Lines 19–24 incorporate the time and plume aware aspects of the search. EAR first computes the arrival time at the candidate neighbor. It then evaluates whether the movement is at the same or opposite side of the plume trajectory relative to the goal and, if necessary, applies a directional penalty. After that, the algorithm evaluates how close the candidate cell lies to the predicted plume trajectory. In this way, EAR does not treat all feasible moves as equally desirable, but takes into account the broader spatial structure of the hazard.

Lines 25–32 combine the local movement information into a single transition score and decide whether the new route to the candidate neighbor improves upon the best one known so far. If it does, the predecessor of that neighbor is updated, its cumulative routing score is replaced by the improved one, its arrival time is stored, and a new queue priority is computed by combining the cumulative routing cost with a goal-directed heuristic term. The updated neighbor is then inserted back into the priority queue. Finally, if the queue is exhausted without reaching the destination, line 33 returns failure, indicating that no feasible evacuation path could be found under the current routing conditions.

The pseudocode shows that EAR follows the general structure of a best-first search, but augments each routing state with time-dependent hazard information. In particular, the algorithm keeps track not only of which cells have been reached, but also of when they are reached and under what hazard conditions. This is what allows EAR to move beyond purely geometric route selection and to support exposure-aware evacuation guidance.

To make the operation of EAR more concrete, Figure 6.2 presents an illustrative example of a




Figure 6.2: Illustrative numerical example of EAR neighbor evaluation. Starting from the current cell, EAR evaluates the neighboring candidates by considering the expected exposure, traversal effort, and plume-related penalties before selecting the most favorable next move.

local neighbor-evaluation step.

In this example, each candidate move is scored using the weighted combination of predicted exposure, traversal time, directional penalty, proximity penalty, and a goal-directed heuristic term. The inset table in Fig. 6.2 summarizes the values used for the three candidates together with their resulting final scores.

Cell v_1 lies within the plume footprint already at time t_0 and remains within the affected region at the subsequent predicted state. As a result, the move toward v_1 incurs exposure both from the concentration at departure and from the predicted concentration at arrival, yielding the highest exposure term among the three candidates. Cell v_3 , by contrast, is not exposed at the current time step and may therefore appear preferable under a purely static interpretation of the hazard field. However, EAR does not evaluate candidate cells using only present-time conditions. Instead, it evaluates the hazard expected at the time of arrival. In this case, although v_3 is outside the plume at time t_0 , the plume is predicted to advance toward that region at later states, and thus v_3 incurs a non-zero exposure term when the move is evaluated at arrival time.

Cell v_6 differs from both of the above cases. It remains outside the hazardous footprint during the evaluated transition and therefore incurs no exposure cost in this step. In addition, it receives lower proximity and directional penalties than the alternatives, since it remains farther from the predicted plume trajectory and avoids an unfavorable plume-relative movement. Consequently, even though v_6 is not necessarily the most direct geometric move toward the goal, it attains the lowest final score.

This example highlights a key characteristic of EAR: the selected next step is not simply the most direct neighbor toward the destination, but the one that best balances safer traversal with practical progress under the evolving hazard field. In particular, EAR can reject a candidate such as v_3 even when it appears safe *at the current instant*, because the algorithm explicitly accounts for the hazard conditions that will hold when the evacuee actually reaches that cell. This local decision process is repeated throughout the search until the destination is reached or no feasible route remains.

This type of example is important because it highlights the main intuition behind EAR. The algorithm does not simply choose the shortest local move, nor does it always choose the absolutely least contaminated immediate step in isolation. Instead, it evaluates each candidate in a broader operational sense, balancing progress, exposure, movement effort, and plume-aware spatial considerations. The resulting decision is therefore better aligned with the requirements of evacuation under an evolving airborne hazard.

6.4 Application-Level Route Generation and Result Presentation

Once the hazard-aware routing representation has been constructed and the EAR algorithm has been defined, the final step is to expose this functionality through the *HazardEvac* application interface. The role of the application layer is to transform the routing capabilities of EAR into an operational workflow through which the user can define a routing request, invoke route computation, and inspect the resulting evacuation guidance together with its hazard context. In this sense, the application layer is not separate from the routing logic, but rather serves as its operational realization within the platform.

At the user level, route generation begins after the plume outputs have been loaded and the routing domain has been defined over the map. Through the route-planner interface, the user specifies the main routing inputs required for EAR. These include the routing cell size, the start point, the destination or one or more candidate goals, the selected algorithm variant, the transport mode, the personal-protection assumption, and the relative weights of the routing cost components. In addition, the interface allows the optional activation of post-processing based on OpenRouteService, as well as the use of auxiliary visualization options such as terrain preview, route-layer toggles, and walkability overlays. Thus, the application does not merely display the route after computation, but provides the control surface through which the routing request itself is configured.

After the user completes the above configuration and invokes route generation, the frontend sends the corresponding routing request to the backend service. In the present implementation, two execution paths are supported. In the first case, the request is solved directly over an already prepared hazard-aware routing grid. In the second case, the backend first constructs the routing grid from the selected map bounds, the chosen cell size, and the corresponding CDUMP plume outputs, and then solves the routing problem over the resulting discretized domain. In both cases, the backend resolves the selected transport mode and protection settings, applies the user-defined routing weights, instantiates the routing graph, and invokes the selected EAR variant over the hazard-aware grid.

The output returned by the backend contains the information required both for algorithmic interpretation and for application-level presentation. More specifically, the response includes the sequence of path cells, the corresponding geographic coordinates of the computed route, the total routing cost, the estimated total dose in sieverts and millisieverts, the associated WHO risk level, and the walkability information of the routing domain. When optional road-network post-processing is enabled, the response may additionally include an OpenStreetMap/OpenRouteService route geometry together with its estimated distance and duration. In this way, the result




Figure 6.3: Application-level route generation and result presentation in *HazardEvac*. The interface allows the user to configure the routing request, invoke EAR over the hazard-aware routing domain, and visualize the resulting evacuation route together with the plume footprint, routing controls, timeline playback, and route-level summary information.

returned to the frontend is not limited to a single polyline, but constitutes a structured routing product containing both the selected evacuation path and the contextual information necessary to interpret it.

Figure 6.3 illustrates this application-level workflow. On the left side, the interface exposes the route-planner controls through which the user defines the routing problem and adjusts the EAR configuration. On the map itself, the selected routing domain, the evolving plume footprint, the start and goal markers, and the resulting route are displayed together. The interface can present both the algorithmic EAR route and, when enabled, the corresponding post-processed OSM route, allowing the user to inspect the computed guidance from complementary perspectives. In addition, the route summary panel presents application-level information such as total exposure and terrain state, thereby linking the algorithmic output to directly interpretable operational indicators.

An important element of the application layer is the timeline widget, which allows the user to inspect the route together with the temporal evolution of the plume. This component does not recompute the route dynamically during playback. Instead, it serves as a visualization aid through which the already computed route can be examined against successive plume states.

6.5 Dose-Related Interpretation and Reporting

The routing process in *HazardEvac* is performed over a concentration-based exposure proxy rather than directly over effective dose. This is a deliberate design choice. By keeping the routing core independent of isotope-specific dose models, the EAR algorithm remains applicable to a broader class of airborne hazards. Dose estimation is therefore introduced as a post-processing step after a valid route has already been computed. In this way, the platform separates the generic routing problem from the hazard-specific interpretation problem.

In the current implementation, dose estimation is carried out along the computed route on a segment-by-segment basis. Once EAR returns the ordered sequence of cells forming the route, the system examines each consecutive cell transition and determines the concentration encountered during that segment. If no pre-computed segment exposure values are available, the implementation estimates the segment concentration using the average of the concentration at the departure cell and the concentration at the arrival cell, evaluated at the corresponding traversal times. This is consistent with the same time-dependent logic already used during route generation, where the hazard associated with a move depends not only on spatial location but also on the time at which the next cell is reached.

For each route segment, the implementation then estimates the inhaled quantity by combining the encountered concentration with the traversal duration and the physiological and operational assumptions associated with the evacuee. More specifically, the system resolves an inhalation rate and an infiltration or transfer factor from the selected transport mode. Thus, the same external plume concentration does not necessarily imply the same internal exposure under different movement conditions, since both breathing rate and environmental transfer differ across mobility modes. For example, when an evacuee is walking or cycling in open air, the transfer factor is set to 1.0, meaning that the individual is assumed to be exposed to essentially the full outdoor concentration. By contrast, when the evacuee is travelling by car, the transfer factor is reduced to 0.60, reflecting the fact that the vehicle cabin partially filters or reduces the penetration of contaminated outdoor air. At the same time, inhalation rates also differ by movement mode: cycling assumes a higher breathing rate than walking, while driving assumes a lower one. The transport-mode assumptions used during dose-related post-processing are summarized in Table 6.1.

Transport Mode	Speed (km/h)	Inhalation Rate (m ³ /min)	Transfer Factor
Walking	5.0	0.80/60	1.00
Bicycle	25.0	1.20/60	1.00
Car	50.0	0.45/60	0.60

Table 6.1: Transport mode assumptions used during dose related post processing.

In addition, when respiratory protection is selected, the corresponding mask penetration factor is applied so that only a fraction of the airborne particles contributes to the effective inhalation rate. The respiratory-protection assumptions used in the current implementation are summarized in Table 6.2. These values allow the system to estimate how the same route may lead to substantially different internal dose burdens depending on the assumed level of respiratory protection. For example, the NONE setting assumes full penetration, meaning that all inhaled airborne particles contribute to exposure. In contrast, a P100 mask is modeled with a penetration factor of 0.0003, meaning that only a very small fraction of the airborne particles is assumed to pass through the filter. Similarly, OV_AG_P100 still provides strong protection, but with a higher penetration factor than P100, and therefore corresponds to a somewhat higher inhaled burden under the same external conditions.

Mask Configuration	Penetration Factor
NONE	1.0000
P100	0.0003
P3	0.0005
ABEK1P3	0.1000
OV_AG_P100	0.0050

Table 6.2: Respiratory-protection assumptions used during dose-related post-processing.

The conversion from route-level concentration to radiological dose is performed in two stages. First, the encountered mass concentration is converted into activity concentration using the specific activity of the selected radionuclide. Let C_m denote the segment mass concentration in mg/m^3 , and let SA_{iso} denote the specific activity of the selected isotope in Bq/mg . The corresponding activity concentration is computed as

$$C_a = C_m \cdot SA_{\text{iso}},$$

where C_a is expressed in Bq/m^3 . This step converts a mass-based concentration field into an activity-based radiological quantity that can be related to health impact through isotope-specific dose coefficients. Here, *activity* refers to the rate at which radioactive atoms undergo nuclear decay, and is measured in becquerels (Bq), where one becquerel corresponds to one radioactive disintegration per second. In simple terms, this means that the same physical mass concentration does not imply the same radiological significance for all isotopes. A very small mass of a highly active radionuclide may correspond to a much larger activity concentration than the same mass of a less active radionuclide.

Second, the activity concentration is converted into effective dose by combining it with the inhaled air volume and the radionuclide-specific inhalation dose coefficient. Let IR_{eff} denote the effective inhalation rate in m^3/h , after any mask-related adjustment, let Δt denote the segment

traversal time in hours, let DCF_{inh} denote the isotope-specific inhalation dose coefficient in Sv/Bq, and let ϕ denote the relevant infiltration or transfer factor. The segment dose is then estimated as

$$D_{seg} = C_a \cdot IR_{eff} \cdot \Delta t \cdot DCF_{inh} \cdot \phi.$$

This formulation reflects the fact that route-level radiological impact depends not only on the external concentration field, but also on how much contaminated air is actually inhaled under the selected transport and protection assumptions. For example, if two evacuees traverse the same route under the same plume conditions, but one is walking outdoors without a mask while the other is travelling by car with cabin filtering, the first will generally receive a larger inhaled burden because both the transfer factor and the breathing conditions are less favorable. Similarly, if the same person follows the same route but wears a high-efficiency filter mask, the effective inhalation rate is reduced and the resulting dose becomes substantially smaller. In the implementation, the radionuclide database stores the specific activity, half-life, and inhalation dose coefficient of each supported isotope, while the transport and mask profiles provide the operational modifiers used during the conversion.

After all route segments have been processed, the individual segment doses are accumulated to produce the total dose associated with the complete evacuation route. The system reports this value in sieverts, millisieverts, and microsieverts, thereby allowing the route result to be expressed in units that are operationally meaningful for radiological interpretation. Beyond the raw numerical value, the implementation also derives a qualitative risk label by comparing the total route dose, and where relevant the implied dose rate, against predefined reference levels. These reference levels are intended as an interpretive aid and are informed by internationally recognized radiological protection guidance, including WHO and related emergency-radiological guidance, together with the broader public and emergency reference levels commonly used in radiological protection practice. In particular, the implementation distinguishes between levels consistent with the public annual dose limit, the lower part of the emergency reference band, and higher projected-dose conditions associated with protective actions such as evacuation. The purpose of these categories is not to claim that there exists a single universally “safe” route dose threshold for all scenarios. Rather, they provide a practical way to contextualize the computed route result against established radiological reference bands, so that the user or analyst can interpret whether the projected route burden is comparatively low, elevated, high, or critical under the assumed scenario conditions.

7 Experimental Evaluation

7.1 Introduction

This chapter evaluates the performance of *HazardEvac* and, in particular, the effectiveness of the *Emergency Airborne Release* (EAR) algorithm under dynamically evolving airborne hazard conditions. The purpose of this evaluation is to examine whether the routing principles introduced in the previous chapters translate into measurable practical benefits when compared against simpler baseline alternatives. More specifically, the experiments investigate whether a routing strategy that explicitly accounts for time-varying hazard conditions can reduce route-level radiological exposure while preserving operationally acceptable travel characteristics.

The experimental analysis is organized around a radiological evacuation setting in which different routing strategies are compared over the same underlying hazard scenarios. The evaluation focuses on route-level outcomes expressed primarily in terms of average dose, while also considering complementary measures such as travel time and computational behavior. The result figures considered in this chapter include baseline dose and path-time comparisons across transport modes, mobility-speed sensitivity analyses for walking, bicycling, and car movement, plume-speed sensitivity experiments, and an assessment of personal-protection measures under different mask assumptions.

From a methodological perspective, the evaluation is designed to examine not only absolute route quality, but also the sensitivity of the routing behavior to changes in the main operational factors. These include the speed of the evacuee, the propagation speed of the plume, and the use of personal protective equipment.

The remainder of this chapter is organized as follows. Section 7.2 presents the experimental setup, including the testbed, the compared approaches, the datasets and scenarios, the evaluation metrics, and the microbenchmark factors. Section 7.3 then reports the experimental results, beginning with the baseline radiological scenario and continuing with the sensitivity analyses for mobility speed, plume propagation speed, and personal protection measures. Finally, Section 7.4 discusses the main findings, their practical implications, and the limitations.

7.2 Experimental Setup

This section describes the experimental methodology used to evaluate the performance of *HazardEvac* and, in particular, the behavior of EAR under different operating conditions. The objective of the setup is to assess not only whether EAR can reduce route-level radiological exposure when compared with simpler baselines, but also how robust this advantage remains when the main physical and operational factors of the problem are varied. To this end, the evaluation is organized around a common radiological-routing framework in which all compared methods operate over the same underlying plume scenarios and routing domains, while selected parameters such as mobility speed, plume propagation speed, and personal protection assumptions are systematically modified.

7.2.1 Testbed

The experimental evaluation was conducted on a standalone experimental prototype that preserves the core processing pipeline of *HazardEvac*, while simplifying some aspects of the full application-level environment. In particular, the prototype retains the same main sequence of operations used by the platform: plume generation through HYSPLIT, transformation of plume outputs into a routing-ready grid, execution of the routing algorithm over that grid, and post-processing of the resulting routes into dose-related summaries. Thus, although the experiments were not executed through the full deployed application interface, they still evaluate the essential computational workflow underlying *HazardEvac*.

Atmospheric dispersion was generated with HYSPLIT inside a dedicated Docker container. Meteorological inputs were obtained from NOAA GFS data and stored locally according to their timestamps. Since HYSPLIT requires meteorological inputs in ARL format, the downloaded files were converted through a separate preprocessing container that hosted the corresponding conversion utilities. The converted files were then reused across simulations, reducing redundant preparation overhead when multiple experiments shared the same meteorological conditions. In this way, the testbed preserves the same plume-generation workflow used by the main platform.

The main simplification of the standalone prototype concerns the routing domain. While the full *HazardEvac* platform incorporates terrain-aware preprocessing and distinguishes between different terrain classes and non-traversable cells, the experimental prototype uses a uniform routing surface in which all grid cells are treated as having the same terrain characteristics. Consequently, terrain-dependent mobility variation and terrain-based exclusion of cells are not part of the present experimental evaluation. This means that the reported results isolate the behavior of the hazard-aware routing logic under a simplified movement environment.

All experiments were executed on a workstation equipped with an Intel i7-8700K CPU, 32 GB of RAM, and an NVIDIA RTX 2070 GPU, running a Windows-based operating system. The same machine and software environment were used throughout the evaluation in order to ensure consistency and a fair comparison across the competing routing methods.

7.2.2 Compared Approaches

The experimental evaluation compares EAR against two simpler routing baselines that represent hazard-insensitive movement behaviors of different quality. The first baseline is a *Euclidean* strategy, which is used as a proxy for a fastest-path approach. Since it follows the most direct geometric progression toward the destination, it represents a routing behavior that prioritizes rapid spatial progress but does not explicitly account for the evolving hazard field. The second baseline is a *Random* strategy, which serves as a weaker reference point and provides an indication of route quality in the absence of informed hazard-aware guidance. Accordingly, the main comparisons reported throughout this chapter are between EAR, Euclidean, and Random.

In addition to these reported baselines, a brute-force strategy was also considered during the early stages of experimentation as a conceptual reference for identifying the lowest-exposure route through exhaustive enumeration of candidate paths. However, this approach was not included in the reported evaluation because its computational complexity quickly became prohibitive. In the present routing setting, each step may generate up to eight successor cells, implying exponential growth of the search space on the order of 8^n for a path of depth n . Consequently, exhaustive search became computationally intractable even for modest routing depths and was therefore unsuitable as a practical comparison method for the scenarios considered in this chapter.

7.2.3 Datasets and Scenarios

The experiments are based on radiological dispersion scenarios represented through time-indexed plume outputs. The underlying hazard fields are generated through the plume-generation workflow and then transformed into routing-ready hazard grids using the integration process described in the previous chapters. Each scenario therefore consists of a time-varying radiological concentration field defined over the routing domain, together with a start point, one or more goals, and the operational assumptions required for route generation.

The evaluation is organized around a baseline radiological-routing scenario and a set of controlled microbenchmark variations. In the baseline setting, the compared methods are executed over common plume conditions and transport modes, and their resulting route-level doses are compared. In the sensitivity analyses, one main factor is varied at a time while the others are kept fixed, so that the effect of that factor on route quality can be examined more clearly. The provided result sets indicate that the experiments include 100 routing instances per tested configuration, allowing the reported values to reflect average behavior rather than isolated individual routes.

The scenario family covers several dimensions of operational interest. First, mobility is examined across walking, bicycling, and car travel, both in baseline form and under speed variation. Second, the plume itself is varied through different propagation speeds. Third, the radiological interpretation of the computed routes is examined under different respiratory-protection assumptions.

7.2.4 Evaluation Metrics

The principal evaluation metric used in this chapter is the *average dose* associated with the computed routes, expressed in millisieverts. This metric is central because the main objective of EAR is not merely to reach the destination, but to do so while reducing the hazard exposure accumulated during traversal. The uploaded result figures report this metric directly for the compared methods across transport modes, speed settings, plume conditions, and protection assumptions.

A second important metric is *path time*, expressed in hours. This metric is needed because a low-dose route is not necessarily useful if it becomes operationally unrealistic or excessively slow. Route evaluation must therefore balance hazard reduction against movement practicality. The path-time results reported for the baseline transport-mode comparison provide this complementary view by showing how route duration varies across movement modes and methods.

7.2.5 Microbenchmark Factors

To understand how EAR behaves under different operating conditions, the evaluation varies several controlled microbenchmark factors. The first factor is *mobility speed*. Separate speed sweeps are considered for walking, bicycling, and car movement. For walking, the tested speeds are 5, 10, 15, and 20 km/h. For bicycling, the tested speeds are 10, 20, 30, and 40 km/h. For car travel, the tested speeds are 25, 50, 75, and 100 km/h. These speed ranges allow the evaluation to examine how route-level dose changes as traversal becomes faster within each transport category.

The second factor is *plume propagation speed*. The plume-speed evaluation considers a range of plume speeds, namely approximately 1.7, 3.5, 4.2, 5.2, 7, 10.4, 20.8, and 52.1 knots. This factor is important because the relative difficulty of the routing problem depends strongly on how quickly the hazard field evolves with respect to evacuee movement. A route that is adequate under a slowly moving plume may become significantly less favorable under faster propagation conditions.

The third factor is *personal protection*. The evaluation includes the protection settings NONE, P100, P3, ABEK1P3, and OVAGP100. These assumptions do not change the geometric route-generation problem itself, but they affect the dose-related interpretation of the computed routes by modifying the fraction of inhaled airborne particles. This microbenchmark therefore allows the evaluation to assess how route-level dose outcomes change when different respiratory-protection assumptions are applied to the same underlying hazard-routing setting.

These microbenchmark factors were selected because they capture three distinct but complementary dimensions of the problem: how fast the evacuee can move, how quickly the plume evolves, and how strongly protective measures reduce inhalation burden. Evaluating EAR under all three dimensions provides a more complete understanding of both its effectiveness and its robustness.

7.3 Results

This section presents the experimental results obtained under the radiological-routing setting described in the previous section. We begin with the baseline comparison across transport modes, where the objective is to assess the overall effectiveness of EAR relative to the Euclidean and Random baselines. We then examine a set of sensitivity analyses that isolate the impact of evacuee mobility speed, plume propagation speed, and personal protection assumptions.

7.3.1 Radiological Scenario

We first consider the baseline radiological scenario, in which the three evaluated approaches are compared across the three transportation modes: walking, bicycling, and car movement. Figure 7.1 presents the average dose achieved by each algorithm, while Figure 7.2 reports the corresponding average path time. In both figures, the compared methods are EAR, Euclidean, and Random, and the reported dose values are shown on a logarithmic scale together with indicative public, emergency, and evacuation-related reference thresholds.




Figure 7.1: Baseline radiological dose evaluation across transport modes. Average route dose is reported for EAR, Euclidean, and Random under walking, bicycling, and car movement.




Figure 7.2: Baseline path-time evaluation across transport modes. Average path time is reported for EAR, Euclidean, and Random under walking, bicycling, and car movement.

As shown in Figure 7.1, EAR consistently achieves the lowest average dose across all three transport modes, and its advantage over the two baselines is substantial. More generally, across walking, bicycling, and car movement, the average dose produced by EAR remains approximately two to three orders of magnitude lower than the corresponding values of the Euclidean and Random approaches. This indicates that the benefit of EAR is not confined to a single mobility regime, but persists across the full range of tested transport conditions. Thus, the main trend is consistent throughout the baseline evaluation: faster-path or unguided routing may still reach the destination, but they do not achieve the same reduction in radiological burden as EAR. This result indicates that directly pursuing the destination without accounting for the evolving hazard field is not sufficient in a radiological evacuation setting. EAR performs better because it explicitly incorporates the temporal and spatial structure of the plume during route generation, thereby avoiding transitions that would lead the evacuee into more contaminated regions.

Across all tested transport modes, the average doses associated with EAR remain well below the lower radiological reference levels, whereas the competing methods lie noticeably closer to the lower threshold bands. Although this does not by itself constitute a complete operational safety assessment, it does indicate that hazard-aware routing substantially improves the radiological quality of the selected evacuation routes.

Figure 7.2 complements the dose analysis by showing the corresponding average path times. A key observation is that EAR achieves travel times that remain close to those of the Euclidean baseline, while substantially outperforming the Random strategy in terms of duration. This is an important result, because a routing method that only minimizes dose but produces highly impractical or excessively long detours would be of limited operational value. Instead, EAR manages to maintain travel times that are comparable to the direct-path baseline while still obtaining much lower dose. This indicates that the algorithm achieves a favorable trade-off between safety and evacuation efficiency.

7.3.2 Mobility Speed Variation

We next examine how the routing outcome changes as the speed of the evacuee varies within each transportation category. Figures 7.3, 7.4, and 7.5 report the average dose achieved by the three algorithms for walking, bicycling, and car travel, respectively, under increasing mobility speeds. All three figures are reported on a logarithmic scale.




Figure 7.3: Walking-speed variation. Average route dose is reported for EAR, Euclidean, and Random under walking speeds of 5, 10, 15, and 20 km/h.




Figure 7.4: Bicycle-speed variation. Average route dose is reported for EAR, Euclidean, and Random under bicycle speeds of 10, 20, 30, and 40 km/h.

The first general pattern is that higher mobility speed tends to reduce the accumulated dose for all methods. This is expected, since faster movement reduces residence time inside the contaminated region and therefore decreases total inhalation rate. However, we observe that EAR remains consistently below both baselines across the full tested speed range for each mode. In other words, increasing speed improves all methods, but it does not remove the advantage of hazard-aware routing. EAR continues to provide the lowest average route dose whether the evacuee is moving slowly or quickly.

For walking, shown in Figure 7.3, the relative benefit of EAR is particularly visible at the lower speed range. This is intuitively reasonable: when movement is slow, the evacuee remains inside the contaminated region for longer, and therefore route choice becomes more critical. Under these conditions, simply moving more directly towards the destination is not enough; the algorithm must actively avoid high-concentration regions encountered over time.

A similar pattern appears for bicycling in Figure 7.4. As bicycle speed increases from 10 to 40 km/h, the average dose decreases across all methods, yet EAR consistently remains the lowest-dose approach. This indicates that the benefit of EAR is not limited to one specific mobility regime. Instead, the algorithm continues to exploit the temporal structure of the plume effectively even when the evacuee is able to traverse the domain much more rapidly than in the




Figure 7.5: Car-speed variation. Average route dose is reported for EAR, Euclidean, and Random under car speeds of 25, 50, 75, and 100 km/h.

walking case.

Figure 7.5 shows the same experiment for vehicular movement. Here again, increasing car speed reduces average dose, but EAR remains consistently below the two baselines. The persistence of this advantage is particularly important because one might expect higher vehicular speed to reduce the relevance of routing intelligence by enabling the evacuee to traverse the region quickly enough that the plume structure matters less. The results show that this is not the case: even at the higher speeds, EAR continues to identify routes with lower dose than those produced by the simpler alternatives.

7.3.3 Plume Propagation Speed Variation

We now examine how the algorithms behave when the plume itself evolves at different propagation speeds. Figure 7.6 reports the average dose achieved by EAR, Euclidean, and Random over a range of plume speeds from approximately 1.7 to 52.1 knots, while keeping the vehicular speed fixed at 50 km/h. This experiment shows that the difficulty of the routing problem depends not only on how fast the evacuee can move, but also on how quickly the hazard field changes during traversal.




Figure 7.6: Plume-speed variation. Average route dose is reported for EAR, Euclidean, and Random under increasing plume propagation speeds, with car speed fixed at 50 km/h.

The results show that EAR remains consistently below the two baselines across the full tested plume speed range. Although the absolute dose levels vary less dramatically than in the mobility-speed experiments, the ordering of the methods remains stable: EAR yields the lowest average dose, Euclidean remains higher, and Random performs worst. This indicates that the proposed algorithm is robust not only to changes in evacuee speed but also to changes in the dynamics of the plume itself.

Another observation is that changing how quickly the evacuee moves has a stronger visible influence on accumulated dose than changing how quickly the plume propagates, at least within the tested parameter range. Nevertheless, plume speed is still operationally meaningful, because it governs how quickly lower-dose opportunities may disappear and how rapidly hazardous regions can shift relative to the route. The fact that EAR maintains very low dose values even under faster plume dynamics suggests that the algorithm is able to exploit the time-varying hazard field effectively and identify routes whose traversal remains better aligned with lower-contamination conditions.

From an operational perspective, this is an important robustness result. It indicates that EAR does not depend on a narrow plume-dynamics regime in order to perform well. Instead, it continues to outperform the baselines even when the hazard field changes more rapidly.

7.3.4 Personal Protection Measures

Finally, we examine the influence of respiratory-protection on the reported route-level dose. Figure 7.7 reports the average dose achieved by EAR, Euclidean, and Random under five protection settings: NONE, P100, P3, ABEK1P3, and OVAGP100. In this experiment, the plume speed is approximately 3.47 knots and the bicycle speed is fixed at 25 km/h. Unlike the previous experiments, this microbenchmark does not change the geometry of the routing problem itself. Instead, it changes the dose-related interpretation of the route by modifying the fraction of inhaled airborne particles.




Figure 7.7: Personal-protection evaluation. Average route dose is reported for EAR, Euclidean, and Random under different respiratory-protection assumptions.

The expected overall trend is clearly visible: stronger respiratory protection leads to lower reported dose for all algorithms. The NONE setting, which assumes full inhalation of airborne particles, produces the highest dose values. By contrast, the high-efficiency protection settings, particularly P100 and P3, push the resulting dose to the lowest levels on the plotted scale. Intermediate protection configurations such as ABEK1P3 and OVAGP100 lie between these two extremes, reflecting their weaker but still meaningful particle-filtering assumptions.

At the same time, the relative ordering of the routing methods remains unchanged. EAR consistently achieves the lowest average dose under every protection setting, followed by Euclidean

and then Random. This result shows that protective equipment and hazard-aware routing are complementary rather than interchangeable.

This finding has a practical interpretation. Personal protective equipment can significantly reduce the internal inhalation dose of a radiological plume, but route choice still matters because the external concentration field remains the same across methods. EAR retains its advantage because it continues to steer the evacuee toward lower-contamination regions, and the protection factor then further scales down the inhaled dose along that already improved route. As a result, the algorithmic and protective contributions combine rather than subtract from one another.

7.4 Discussion

The results of this chapter consistently show that EAR achieves lower route-level radiological exposure than the Euclidean and Random baselines across the tested scenarios. This advantage remains visible across different transport modes, mobility speeds, plume propagation speeds, and personal-protection assumptions, indicating that the benefit of hazard-aware routing is both substantial and robust. At the same time, the path-time results suggest that these improvements are not achieved through impractically long routes, but through a better balance between progress and hazard avoidance.

The results also highlight an important practical point: faster movement and stronger personal protection both reduce dose, but neither eliminates the value of better routing. In this sense, mobility, protection, and hazard-aware route selection should be viewed as complementary contributors to evacuation safety rather than as interchangeable alternatives.

Another important observation is that the superiority of EAR is preserved not only in the baseline comparison, but also under the sensitivity analyses. This indicates that the proposed routing logic is not narrowly tuned to a single experimental configuration. Instead, it appears to retain its benefit under changes in both the physical behavior of the plume and the operational characteristics of the evacuee. Such consistency is important for practical deployment, since real emergency conditions are unlikely to match a single fixed parameter setting.

Finally, the reported findings should be interpreted in light of the experimental scope. The evaluation was conducted on a standalone prototype that preserves the core plume to routing to dose processing chain of *HazardEvac*, but does not yet incorporate terrain-dependent heterogeneity or non-traversable cells. Consequently, the present chapter validates the hazard-aware routing logic under a simplified movement model, while future work should extend the evaluation to the full terrain-aware platform. It would also be useful to assess the approach under a broader range of release conditions, geographic settings, and larger-scale evacuation scenarios, so that the generality of the current findings can be examined more fully.

8 Conclusions and Future Work

8.1 Summary of the Thesis

This thesis addressed the problem of evacuation routing under dynamically evolving airborne hazardous dispersion. The main motivation of the work was that conventional shortest-path or purely geometric routing strategies are not well suited to emergency situations in which the hazard field changes over time and where the quality of a route depends not only on where an evacuee moves, but also on when each part of the route is reached. In such settings, direct movement toward a destination may lead an evacuee into areas of intensifying contamination, while a slightly less direct route may lead to substantially lower overall exposure. The central objective of the thesis was therefore to design and evaluate a system capable of supporting time-dependent, hazard-aware evacuation guidance.

To address this problem, the thesis introduced *HazardEvac*, an end-to-end platform that combines plume generation, hazard integration, evacuation routing, and dose-related interpretation within a unified workflow. The proposed system begins from the definition of a hazardous release scenario and the preparation of the corresponding meteorological inputs. It then generates time-varying plume outputs through HYSPLIT, transforms these outputs into a routing-ready hazard field over a discretized movement domain, computes hazard-aware evacuation routes using the proposed *Emergency Airborne Release* (EAR) algorithm, and finally interprets the resulting routes in dose-related terms for reporting and operational understanding.

A substantial part of the thesis focused on the architectural decomposition of this workflow into distinct but connected layers. The plume-generation layer was responsible for producing time-indexed hazard outputs from scenario and meteorological inputs. The hazard-integration layer converted these outputs into a form suitable for routing by constructing the routing domain, associating time-varying hazard values with the grid cells, and preparing the hazard-aware routing representation. The routing and application layer then consumed this representation in order to generate evacuation routes and present them through the platform interface. This layered structure was important because it allowed the thesis to explain the system not merely as a routing algorithm, but as a complete computational pipeline for hazard-aware evacuation support.

Within this broader system, the core algorithmic contribution was EAR. The thesis formulated EAR as a time-dependent routing algorithm that evaluates candidate movements using a multi-criteria cost model combining exposure, traversal time, plume-relative direction, and proximity to the plume trajectory. In this way, EAR moves beyond static shortest-path logic and explicitly accounts for the evolving spatio-temporal structure of the hazard field. Rather than selecting

routes solely on the basis of geometric efficiency, the algorithm attempts to identify routes that reduce the cumulative hazard exposure while preserving practical progress toward safety.

The thesis also addressed the interpretation of route quality in radiological terms. Although routing itself was carried out over a concentration-based exposure proxy, the resulting routes were post-processed to estimate dose-related quantities such as total dose in sieverts, millisieverts, and microsieverts. This conversion incorporated isotope-specific radiological data, transport-mode assumptions, and respiratory-protection factors. As a result, the route produced by EAR could be interpreted not only as a lower-exposure path in relative routing terms, but also as a route associated with a concrete estimated radiological dose under the assumed scenario conditions.

Finally, the thesis evaluated the proposed approach experimentally. The results showed that EAR consistently outperformed simpler routing baselines, namely Euclidean and Random routing, under the tested radiological scenarios. In particular, EAR achieved substantially lower average route-level dose while preserving operationally reasonable path times. This advantage remained visible across different transport modes, mobility speeds, plume propagation speeds, and personal-protection configurations, indicating that the proposed routing logic is both effective and robust under varying conditions.

8.2 Main Contributions

The main contributions of this thesis can be summarized as follows:

1. **An end-to-end hazard-aware evacuation platform.** The thesis designed and implemented *HazardEvac*, a complete workflow that links plume generation, hazard integration, evacuation routing, and application-level result presentation within a unified system.
2. **A hazard-integration methodology for routing under dynamic dispersion.** The work developed the intermediate processing steps required to transform time-varying plume outputs into a routing-ready hazard representation over a discretized movement domain, thereby bridging the gap between atmospheric-dispersion modeling and route computation.
3. **The Emergency Airborne Release(EAR) algorithm.** The thesis proposed EAR as a time-dependent hazard-aware routing algorithm that combines exposure, traversal time, plume-relative direction, and plume-relative proximity in a multi-criteria cost model in order to reduce cumulative route burden while maintaining practical movement toward safety.
4. **A route-level dose interpretation pipeline.** The work introduced a post-processing methodology for converting concentration-based route outcomes into dose-related quan-

tities, allowing the computed routes to be interpreted in operationally meaningful radiological terms.

5. **An experimental evaluation of hazard-aware routing behavior.** The thesis demonstrated through experimental analysis that EAR consistently outperforms simpler baselines across the tested radiological scenarios and sensitivity analyses, thereby validating the value of time-dependent hazard-aware evacuation routing.

8.3 Limitations

Despite the above contributions, the present work also has a number of limitations that should be acknowledged. First, the experimental evaluation was conducted on a standalone prototype that preserved the core plume to routing to dose processing chain of the system, but did not yet incorporate the full terrain-aware routing functionality of the complete platform. In particular, terrain heterogeneity and non-traversable cells were not included in the reported experiments.

Second, the experimental analysis focused primarily on radiological scenarios and on a bounded set of controlled parameter variations. Although this provided a useful basis for evaluating the effectiveness of EAR, it does not yet capture the full diversity of hazardous atmospheric releases that may arise in practice, such as chemical releases with different physical behaviors or wildfire smoke scenarios with different operational priorities.

Third, the evaluation compared EAR against meaningful but relatively simple baselines, namely Euclidean and Random routing. Although a brute-force strategy was considered conceptually as a possible reference for minimum-exposure routing, it became computationally intractable due to the exponential growth of the search space and could therefore not be included in the reported results. As a result, the evaluation demonstrates EAR's superiority over practical baselines, but not over an exact optimal solver.

Fourth, the dose-related interpretation pipeline depends on scenario assumptions and model parameters, including isotope properties, inhalation rates, transfer factors, and respiratory-protection assumptions. While these choices are necessary for operational interpretation, they also imply that the reported dose values are model-based estimates rather than direct measurements. Consequently, the radiological summaries should be understood as decision-support outputs conditioned on the adopted assumptions.

Finally, the present thesis primarily focused on the design, implementation, and controlled evaluation of the proposed framework. Broader issues such as user studies, operator-centered usability assessment, large-scale field validation, and multi-agent evacuation coordination remain outside the scope of the current work.

8.4 Future Work

The work presented in this thesis opens several directions for future research and development. A first and immediate extension is the incorporation of the full terrain-aware routing functionality into the evaluated pipeline. Future experiments should assess EAR under the complete routing model of *HazardEvac*, including terrain-dependent mobility variation and non-traversable cells, so that the current algorithmic findings can be validated under more operationally realistic movement constraints.

A second important direction is the extension of the evaluation to a broader range of scenarios. This includes testing the framework over different geographic regions, meteorological conditions, release types, and hazard categories. In particular, it would be valuable to examine how the proposed routing logic behaves in chemical-release settings and in wildfire-smoke scenarios, where exposure reduction remains important but the interpretation and protective assumptions may differ from the radiological case.

A third direction concerns real-time and adaptive operation. The current workflow already supports time-dependent plume generation and routing, but future work could investigate tighter integration with live meteorological updates, streaming sensor inputs, and continuously refreshed hazard fields. Such extensions would make the platform more suitable for real-world operational environments in which the hazard estimate evolves as new data become available.

Another promising direction is the treatment of uncertainty. Both plume prediction and evacuation routing are affected by uncertainty in meteorological forecasts, release characteristics, and environmental conditions. Future work could therefore investigate uncertainty-aware or probabilistic routing methods that account not only for expected hazard intensity but also for confidence bounds or alternative hazard realizations.

The routing problem itself could also be extended beyond the single-user setting considered here. Future research may examine multi-agent or population-level evacuation routing, congestion-aware route generation, coordinated evacuation strategies, and the interaction between individual route optimization and system-level evacuation objectives. These issues become especially important when the platform is to support larger-scale evacuation planning rather than only single-route guidance.

Finally, future work should include broader validation activities beyond controlled computational evaluation. This may involve realistic exercises, operator-in-the-loop studies, and usability-driven refinement of the application interface. Such work would help assess not only whether the algorithm computes better routes, but also whether the system presents its outputs in a way that is understandable, actionable, and useful for practical emergency decision support.

BIBLIOGRAPHY

- [1] Radiation and health. [Online]. Available: <https://www.who.int/news-room/questions-and-answers/item/radiation-and-health>
- [2] LIFO. Πώς οι φωτιές κάνουν επικίνδυνο τον αέρα που αναπνέουμε. Accessed: 2026-05-13. [Online]. Available: <https://www.lifo.gr/stiles/optiki-gonia/pos-oi-foties-kanoyne-epikindyno-ton-aera-poy-anapneoyme>
- [3] Focus Taiwan. Nuclear disaster drill / emergency response related coverage. Accessed: 2026-05-13. [Online]. Available: <https://focustaiwan.tw/society/202007150006>
- [4] E. S. A, Nahed. RCCE messages for a radio-nuclear emergency: get inside, stay inside, stay tuned. [Online]. Available: <https://www.emro.who.int/rcce/rn.html>
- [5] S. Kim, B. George, and S. Shekhar, “Evacuation route planning: scalable heuristics,” in *Proceedings of the 15th annual ACM international symposium on Advances in geographic information systems*, ser. GIS ’07. Association for Computing Machinery, pp. 1–8. [Online]. Available: <https://dl.acm.org/doi/10.1145/1341012.1341039>
- [6] T. Chondrogiannis, P. Bouros, and W. Emser, “Simulation-based evacuation planning for urban areas,” in *Proceedings of the 29th International Conference on Advances in Geographic Information Systems*, ser. SIGSPATIAL ’21. Association for Computing Machinery, pp. 297–300. [Online]. Available: <https://dl.acm.org/doi/10.1145/3474717.3483963>
- [7] R. Wang, B. Chen, Y. Wang, Z. Zhu, L. Ma, and X. Qiu, “The air contaminant dispersion prediction by the integration of the neural network and AermodySystem,” in *Proceedings of the 4th ACM SIGSPATIAL International Workshop on Safety and Resilience*, ser. Safety and Resilience’18. Association for Computing Machinery, pp. 1–5. [Online]. Available: <https://dl.acm.org/doi/10.1145/3284103.3284114>
- [8] Y. Gao, H. Zhang, L. Xiao, Y. Fu, and X. Ma, “Spatial and temporal distribution of LPG dispersion concentration in small-scale space,” in *Proceedings of the 5th ACM SIGSPATIAL International Workshop on the Use of GIS in Emergency Management*, ser. EM-GIS ’19. Association for Computing Machinery, pp. 1–8. [Online]. Available: <https://dl.acm.org/doi/10.1145/3356998.3365767>
- [9] Z. Li, F. Wang, X. Zheng, W. Jiang, Q. Meng, and B. Liu, “GIS based dynamic modeling of fire spread with heterogeneous cellular automation model and standardized emergency management protocol,” in *Proceedings of the 3rd ACM SIGSPATIAL International Workshop on the Use of GIS in Emergency Management*, ser. EM-

- GIS '17. Association for Computing Machinery, pp. 1–7. [Online]. Available: <https://dl.acm.org/doi/10.1145/3152465.3152470>
- [10] NASA Scientific Visualization Studio. Global Carbon Monoxide Visualization. Accessed: 2026-05-13. [Online]. Available: <https://svs.gsfc.nasa.gov/11893>
- [11] European Commission – Civil Protection Knowledge Network. Nuclear and Radiological Accidents. Accessed: 2026-05-13. [Online]. Available: <https://civil-protection-knowledge-network.europa.eu/eu-overview-risks/human-induced-risks/nuclear-and-radiological-accidents>
- [12] W. Cheng, Z. Nie, Y. Guan, J. Xing, L. Chen, N. Xue, X. Chen, L. Gao, J. Yan, and W. Deng, “A multistage model for vehicle routing planning in a dynamic nuclear radiation dose field,” vol. 20, no. 6, pp. 8149–8159. [Online]. Available: <https://ieeexplore.ieee.org/document/10461085>
- [13] Modeling and simulation of emergency evacuation from industry buildings under toxic cloud release. [Online]. Available: <https://ieeexplore.ieee.org/document/9084637>
- [14] Y. Wu, K. Lin, W. Li, T. Xiang, Y. Zhang, Y. Yang, Y. Chen, and R. Zhou, “Forest fire emergency rescue and evacuation path planning based on improved APF,” in *Proceedings of the 8th ACM SIGSPATIAL International Workshop on Security Response using GIS*, ser. EM-GIS '23. Association for Computing Machinery, pp. 1–6. [Online]. Available: <https://dl.acm.org/doi/10.1145/3615884.3629421>
- [15] C. Anastasiou, C. Costa, P. K. Chrysanthis, C. Shahabi, and D. Zeinalipour-Yazti, “ASTRO: Reducing COVID-19 exposure through contact prediction and avoidance,” vol. 8, no. 2, pp. 11:1–11:31. [Online]. Available: <https://dl.acm.org/doi/10.1145/3490492>
- [16] “Genasys Protect.” [Online]. Available: <https://protect.genasys.com/search?z=4.5478379079716085&latlon=33.982091306746526%2C35.390252599622954>
- [17] “Safety tips for travelers.” [Online]. Available: <https://www.jnto.go.jp/safety-tips/>
- [18] Google, “Crisis-related alerts in google maps,” <https://support.google.com/maps/answer/9985621?hl=en>, 2026, accessed: 2026-05-11.
- [19] —, “Google crisis response,” <https://crisisresponse.google/>, 2026, accessed: 2026-05-11.
- [20] —, “Public safety and emergency messages,” <https://crisisresponse.google/how-we-help>, 2026, accessed: 2026-05-11.
- [21] Federal Emergency Management Agency, “Mobile Products,” <https://www.fema.gov/about/news-multimedia/mobile-products>, accessed: 2026-05-04.
- [22] NERV, “NERV Disaster Prevention App,” <https://nerv.app/en/>, accessed: 2026-05-04.

- [23] “FEMA Mobile Products | FEMA.gov,” May 2025. [Online]. Available: <https://www.fema.gov/about/news-multimedia/mobile-products>
- [24] “Mobile Apps.” [Online]. Available: <https://www.redcross.org/get-help/how-to-prepare-for-emergencies/mobile-apps.html>
- [25] “NERV Disaster Prevention App.” [Online]. Available: <https://nerv.app/en/>
- [26] American Red Cross, “Emergency: Severe weather app,” <https://apps.apple.com/us/app/emergency-severe-weather-app/id954783878>, 2026, accessed: 2026-05-11.
- [27] KIOS Research and Innovation Center of Excellence, “Cyprus digital twin (cydt),” <https://www.kios.ucy.ac.cy/cyprus-digital-twin/>, 2026, accessed: 2026-05-11.
- [28] United States Environmental Protection Agency, “What should i do in case of a radiological emergency?” Sep. 2019, accessed: 2026-05-13. [Online]. Available: https://19january2021snapshot.epa.gov/radiation/what-should-i-do-case-radiological-emergency_.html
- [29] Organisation for the Prohibition of Chemical Weapons. (2018) First Responders from Asia Enhance Chemical Emergency Response Capabilities. Accessed: 2026-05-13. [Online]. Available: <https://www.opcw.org/media-centre/news/2018/04/first-responders-asia-enhance-chemical-emergency-response-capabilities>
- [30] Alan Franco, “BC Air Dispersion Modelling: An Essential Tool in Air Quality Management,” <https://alanfranco.com/bc-air-dispersion-modelling-essential-tool-air-quality-management/>, accessed: 2026-05-04.
- [31] NOAA Air Resources Laboratory, “HYSPLIT Basic Tutorial: Dispersion Computation Options,” https://www.ready.noaa.gov/documents/Tutorial/view/parm_disp.html, accessed: 2026-05-04.
- [32] Centers for Disease Control and Prevention, “Radiation contamination versus exposure,” Apr. 2024, accessed: 2026-05-13. [Online]. Available: <https://www.cdc.gov/radiation-emergencies/infographic/contamination-versus-exposure.html>
- [33] P. V. Nielsen, L. Rong, and I. Olmedo, “The IEA annex 20 two-dimensional benchmark test for CFD predictions.”
- [34] Technology Gateway, “SEAM Leading the Way in Computational Fluid Dynamics,” <https://technologygateway.ie/articles/seam-leading-the-way-in-computational-fluid-dynamics-cfd/>, accessed: 2026-05-04.
- [35] S. Jo, “A CFD-based decision matrix for evacuation planning: Minimizing exposure to hazardous chemicals,” vol. 12, no. 12, p. 2844. [Online]. Available: <https://www.mdpi.com/2227-9717/12/12/2844>

- [36] A. Venkatram and J. Thé, “Introduction to gaussian plume models.”
- [37] APSI, “Modeling Fire Plumes: Basic Concepts,” https://www.apsi.tech/modeling_fireplumes_basic.html, accessed: 2026-05-04.
- [38] ResearchGate, “Comparison between the Gaussian Plume and Puff Models,” https://www.researchgate.net/figure/Comparison-between-the-Gaussian-plume-and-puff-models_fig14_366356242, accessed: 2026-05-04.
- [39] D. Zhang, Z. Mao, M. Gong, J. Ren, S. Zuo, and X. Chen, “Study on optimization of shelter locations and evacuation routes of gas leakage accidents in chemical industrial park,” vol. 177, pp. 556–567. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0957582023006341>
- [40] S. Du, “A heuristic Lagrangian stochastic particle model of relative diffusion: model formulation and preliminary results,” *Atmospheric Environment*, vol. 35, no. 9, pp. 1597–1607, Mar. 2001. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1352231000004519>
- [41] A. Nanni, G. Tinarelli, C. Solisio, and C. Pozzi, “Comparison between puff and lagrangian particle dispersion models at a complex and coastal site,” vol. 13, no. 4, p. 508. [Online]. Available: <https://www.mdpi.com/2073-4433/13/4/508>
- [42] ScienceDirect, “Lagrangian Atmospheric Dispersion Modelling Figure,” <https://www.sciencedirect.com/science/article/pii/S1352231018305661>, accessed: 2026-05-04.
- [43] Enviroware / EnviroComp, “LAPMOD: Lagrangian Particle Model,” <https://envirocomp.com/capabilities/lapmod.html>, accessed: 2026-05-04.
- [44] T. Chai, X. Ren, F. Ngan, M. Cohen, and A. Crawford, “Estimation of power plant SO₂ emissions using the HYSPLIT dispersion model and airborne observations with plume rise ensemble runs,” vol. 23, no. 19, pp. 12 907–12 933. [Online]. Available: <https://acp.copernicus.org/articles/23/12907/2023/>
- [45] L. Bakels, D. Tatsii, A. Tipka, R. Thompson, M. Dütsch, M. Blaschek, P. Seibert, K. Baier, S. Bucci, M. Cassiani, S. Eckhardt, C. Groot Zwaafink, S. Henne, P. Kaufmann, V. Lechner, C. Maurer, M. D. Mulder, I. Pisso, A. Plach, R. Subramanian, M. Vojta, and A. Stohl, “FLEXPART version 11: improved accuracy, efficiency, and flexibility,” vol. 17, no. 21, pp. 7595–7627. [Online]. Available: <https://gmd.copernicus.org/articles/17/7595/2024/>
- [46] A. F. Stein, R. R. Draxler, G. D. Rolph, B. J. B. Stunder, M. D. Cohen, and F. Ngan, “NOAA’s HYSPLIT atmospheric transport and dispersion modeling system,” vol. 96, no. 12, pp. 2059–2077. [Online]. Available: <https://journals.ametsoc.org/doi/10.1175/BAMS-D-14-00110.1>

- [47] NOAA Air Resources Laboratory, “HYSPLIT Workshop 2018,” <https://www.arl.noaa.gov/news-pubs/news-archive/hysplit-workshop2018/>, accessed: 2026-05-04.
- [48] ResearchGate, “Illustration of Particulate Cesium-137 Concentrations Originated from the Fukushima Accident,” https://www.researchgate.net/figure/Illustration-of-particulate-cesium-137-concentrations-originated-from-the-Fukushima_fig4_276505573, accessed: 2026-05-04.
- [49] S. Villaneuva, K. Cleary, and A. Krupnick, “The societal value of the HYSPLIT air dispersion model.”
- [50] NASA Jet Propulsion Laboratory. (2025) NASA, Aerospace Corporation Study Sharpens Focus on Ammonia Emissions. Accessed: 2026-05-13. [Online]. Available: <https://www.jpl.nasa.gov/news/nasa-aerospace-corporation-study-sharpens-focus-on-ammonia-emissions/>
- [51] HGI. Plume Mapping. Accessed: 2026-05-13. [Online]. Available: <https://www.hgiworld.com/geophysical-services/environmental-services/plume-mapping/>
- [52] Z. Zhu, B. Chen, S. Qiu, R. Wang, L. Ma, and X. Qiu, “The GIS-based approach for optimal design of air quality monitoring network for management of chemical clusters,” in *Proceedings of the 3rd ACM SIGSPATIAL Workshop on Emergency Management using ACM*, pp. 1–8. [Online]. Available: <https://dl.acm.org/doi/10.1145/3152465.3152479>
- [53] Y. Wang, B. Chen, Z. Zhu, and C. AI, “Strategy of hybrid optimization algorithms for source parameters estimation of hazardous gas in field cases,” in *Proceedings of the 5th ACM SIGSPATIAL International Workshop on the Use of GIS in Emergency Management*, ser. EM-GIS ’19. Association for Computing Machinery, pp. 1–6. [Online]. Available: <https://dl.acm.org/doi/10.1145/3356998.3365768>
- [54] Y. Zhao, B. Chen, Z. Zhu, R. Wang, H. Chen, and X. Qiu, “A promising searching method ‘entrotaxis-jump’ for seeking hazardous gas source in a chemical cluster,” in *Proceedings of the 5th ACM SIGSPATIAL International Workshop on the Use of GIS in Emergency Management*, ser. EM-GIS ’19. Association for Computing Machinery, pp. 1–5. [Online]. Available: <https://dl.acm.org/doi/10.1145/3356998.3365766>
- [55] Incident Management Toolkit, “HAZUS-style Exercise Interface,” image/source material cited as: Incident Management Toolkit / HAZUS-style exercise interface. Accessed: 2026-05-13.
- [56] FEMA, “ALOHA Plume Overlaying HAZUS-MH Inventory,” image/source material cited as: ALOHA plume overlaying HAZUS-MH inventory, FEMA technological hazards material. Accessed: 2026-05-13.

- [57] G. Feng, G. Su, and Z. Sun, “Optimal route of emergency resource scheduling based on GIS,” in *Proceedings of the 3rd ACM SIGSPATIAL International Workshop on the Use of GIS in Emergency Management*, ser. EM-GIS ’17. Association for Computing Machinery, pp. 1–5. [Online]. Available: <https://dl.acm.org/doi/10.1145/3152465.3152471>
- [58] Z. Ding, Y. Cao, Y. Chi, X. Li, R. Yao, and L. Guo, “An effective emergency logistics scheduling model based on multi-objective optimization algorithms,” in *Proceedings of the 4th ACM SIGSPATIAL International Workshop on Safety and Resilience*, ser. Safety and Resilience’18. Association for Computing Machinery, pp. 1–6. [Online]. Available: <https://dl.acm.org/doi/10.1145/3284103.3284113>
- [59] S. Wang, J. Zhou, H. Liang, Z. Wang, C. Su, and X. Li, “A new approach for solving location routing problems with deep reinforcement learning of emergency medical facility,” in *Proceedings of the 8th ACM SIGSPATIAL International Workshop on Security Response using GIS*, ser. EM-GIS ’23. Association for Computing Machinery, pp. 50–53. [Online]. Available: <https://dl.acm.org/doi/10.1145/3615884.3629429>
- [60] Y. Fu, W. Shi, H. Zhang, X. Ma, Y. Gao, and D. Guo, “Path optimization of integrating crowd model and reinforcement learning,” in *Proceedings of the 5th ACM SIGSPATIAL International Workshop on the Use of GIS in Emergency Management*, ser. EM-GIS ’19. Association for Computing Machinery, pp. 1–6. [Online]. Available: <https://dl.acm.org/doi/10.1145/3356998.3365765>
- [61] MeteoInfo, “HYSPLIT Output Concentration Plot,” http://meteothink.org/examples/meteoinfolab/trajectory/hyconc_plot.html, accessed: 2026-05-04.