SCHOOL OF COMPUTATION,
INFORMATION AND TECHNOLOGY —
INFORMATICS

TECHNISCHE UNIVERSITÄT MÜNCHEN

Master's Thesis in Informatics

A Risk-Aware Road Cost Framework Driven by Satellite Damage Detection for Point of Distribution (POD) Location–Allocation

Arda Duman

SCHOOL OF COMPUTATION,
INFORMATION AND TECHNOLOGY —
INFORMATICS

TECHNISCHE UNIVERSITÄT MÜNCHEN

Master's Thesis in Informatics

A Risk-Aware Road Cost Framework Driven by Satellite Damage Detection for Point of Distribution (POD) Location–Allocation

Author:	Arda Duman
Examiner:	Prof. Dr. Rainer Kolisch
Supervisor:	Baturhan Bayraktar
Submission Date:	23.02.2026

I confirm that this master's thesis is my own work and I have documented all sources and material used.

Munich, 23.02.2026 Arda Duman

Acknowledgments

This thesis was developed in the context of disaster preparedness and emergency response, with a particular focus on the February 2023 Kahramanmaraş earthquakes. While the study is grounded in technical and scientific methods, it is motivated by the real-world humanitarian challenges observed during this event. First and foremost, I would like to express my sincere gratitude to my thesis supervisor, Baturhan Bayraktar, for his guidance, valuable feedback, and academic insight throughout the development of this research. His support and perspective played a crucial role in shaping this study. I am deeply grateful to my parents, Nilüfer and Oktay Duman, for their unconditional support, patience, and encouragement throughout this process. I would also like to thank my girlfriend, Dila Berfin Baran, for her understanding, emotional support, and continuous encouragement during challenging times. My sincere thanks extend to my friends for their ongoing support and motivation during the course of this work. Beyond the academic scope of this thesis, I respectfully commemorate all those who lost their lives in the February 2023 earthquake. Their loss represents a profound tragedy. I also wish to acknowledge the extraordinary efforts of search and rescue teams, volunteers, medical staff, and citizens who demonstrated remarkable solidarity and dedication in the aftermath of the disaster. This thesis was conceived with the aim of contributing, even in a limited capacity, to disaster preparedness and response planning. It is my hope that the knowledge and methods developed in this study can support future efforts toward building safer and more resilient urban environments. This work is dedicated to the memory of those we lost and to the collective commitment to reduce the human impact of future disasters.

Abstract

This thesis proposes an end-to-end decision support framework that integrates damage information derived from satellite imagery into the problem of locating Point of Distribution (POD)s and assigning demand to them during the early response phase after a major earthquake, when field information is limited. Rather than treating building-scale damage detection as a standalone output, the approach translates detected damage into risk-sensitive road network costs. Impact zones defined around damaged buildings are overlaid with the road network, and the cost of each road segment is increased according to its level of exposure to these zones. In this way, the lowest-cost path reflects not only geometric distance but also the relative difficulty and operational risk of moving through heavily affected areas.

On the resulting network, least-cost path values are computed between supply points and candidate PODs, and between PODs and demand points. These values are normalized to obtain accessibility scores and used as inputs to a mixed-integer location–allocation model that determines which PODs to open and how to assign demand under capacity constraints. Demand is represented through spatial clustering of damaged buildings in order to preserve their distribution pattern, while candidate POD capacities are scaled using an area-based proxy measure.

The framework is evaluated in a case study covering the Onikişubat and Dulkadiroğlu districts following the 2023 Kahramanmaraş earthquakes. The results indicate that incorporating damage-based penalties into road costs can modify certain routes and local assignment decisions. However, the overall accessibility structure and the resulting location–allocation pattern are more strongly influenced by the underlying road network topology and by capacity assumptions. Sensitivity analyses further show that the number of opened PODs reaches a saturation level beyond a given upper bound, and that model feasibility is particularly sensitive to the chosen capacity scaling approach. Overall, the study demonstrates how satellite-based damage detection can be systematically translated into operational planning inputs for risk-aware accessibility analysis and POD planning under time pressure and data constraints.

Contents

1 Introduction

Large-scale earthquakes are spatially highly heterogeneous disasters. In a short time, they can cause severe physical, social, and economic damage across a wide area. The first hours and days after an earthquake form a critical window, often called the "golden hours", when survival outcomes can still be improved. In this period, reducing losses depends not only on search-and-rescue capacity, but also on getting basic needs to people quickly and reliably [Bal+10]. Shelter, water, food, and medical supplies have to reach affected communities in a way that is both safe and repeatable, not as a one-off effort. This is the domain of humanitarian logistics. Unlike commercial logistics, it is not profit-driven, yet it is carried out under intense time pressure, high uncertainty, and tight resource constraints [Hol+13]. For that reason, performance is not judged simply by whether aid arrives in the region, but by who actually receives it, how fast it is delivered, at what cost, and under what operational risks [Van06].

What makes post-earthquake logistics different from a standard facility location or vehicle routing problem is the field itself. Conditions can change quickly, and access can deteriorate sharply, especially around heavily damaged areas. Debris, surface failures, bridge damage, and abandoned vehicles may slow down movement or block it altogether. As a result, routes that look open on a map can become unreliable in practice [GS06]. This is not only a physical constraint; it also creates operational risk. A road may be technically passable, yet still unsafe or unsuitable for repeated use by heavy relief trucks [Lib+14]. Therefore, decision-makers have to act with incomplete field information, often while communication is disrupted and local constraints shift from one hour to the next [GS06].

The 2023 Kahramanmaraş earthquakes provided one of the most striking examples in modern history of these operational challenges and the need for coordination. Immediately after the earthquake, official institutions and a vast network of volunteers mobilized rapidly, and temporary storage and distribution centers were established in different cities. However, the extent of the destruction on the ground and the severe disruption of transportation networks delayed the delivery of aid in some areas, particularly to rural settlements and heavily damaged neighborhoods in the city center. International media and observer reports highlighted that deficiencies in logistical planning and the extent of physical destruction in the initial stages of the earthquake resulted in some areas being without access to basic aid for days [Mic23; Als23]. This case once again underlines the importance of not only capacity management but also a dynamic and risk-sensitive planning mechanism in disaster logistics.

In this context, the need to balance two fundamental objectives in disaster relief distribution becomes clear: accessibility and equity. The accessibility objective reflects an operational efficiency perspective, aiming to satisfy as much demand as possible in the shortest time with limited resources. In contrast, the equity principle emphasizes that areas that are difficult to reach, more vulnerable, or more severely affected by destruction should not be excluded from the system simply because they are far from the center or harder to access. Models that focus only on maximizing overall efficiency may concentrate on "easy-to-reach" areas and reduce the service level in peripheral or highly affected zones. On the other hand, an overly equity-oriented approach may slow down the overall operation and make it harder to use limited capacity effectively. This accessibility–equity trade-off lies at the core of post-disaster facility location and demand allocation problems [NBA16; LA18]. While both objectives are central in the literature, the present study focuses on accessibility maximization. Damage-based demand estimation partially reflects the spatial distribution of need, but equity constraints are not explicitly modeled.

Recent advances in machine learning and remote sensing have enabled rapid, building-level damage assessment from satellite imagery, adding a new dimension to this complex process. However, there is still a clear gap between these technological developments and operational decision-making. Many studies focus mainly on answering the question "what was damaged?" and produce static damage maps. The question of how this information can be translated into changes in network costs, accessibility measures, or strategic POD decisions often remains fragmented. In other words, image-based damage outputs frequently remain as visual products and are not directly integrated as dynamic inputs into logistics optimization models.

The focus of this thesis is to address this gap by proposing an end-to-end framework that transforms satellite-derived damage information into a data-driven decision support input for last-mile distribution. Unlike traditional approaches, demand points are not derived from pre-disaster population data but directly from the spatial concentration of detected building damage. In addition, the transportation network is not treated as a static distance layer; instead, it is modeled as a risk-sensitive cost surface, where road segments are penalized based on their proximity to damaged building impact zones. An important modeling choice should be noted here. Rather than developing a real-time routing policy that updates while vehicles are in motion, this study adopts a deterministic "planning snapshot" based on post-disaster observations. The aim is not to eliminate uncertainty completely, but to reflect the observed spatial footprint of damage indirectly and quickly through a modified cost structure.

The main contributions of this thesis to the literature and practice are based on methodological integration and comprehensive analyses using real-world data. First, the study presents a coherent and reproducible workflow that links satellite-based damage information to risk-sensitive road costs and then to optimized logistics decisions (facility location and demand allocation). Second, the proposed framework is tested using real data layers from the 2023 Kahramanmaraş earthquakes, demonstrating through a concrete case study in the districts of Onikişubat and Dulkadiroğlu how theoretical models can be integrated with complex field conditions. Finally, the study provides systematic insights into how parameters such as capacity constraints, risk penalty coefficients, and the number of facilities affect the solution structure. In doing so, it offers a transparent and data-driven basis for strategic planning in uncertain post-disaster environments.

The remainder of the thesis is organized as follows. Chapter 2 reviews the existing literature on disaster logistics, humanitarian facility location, satellite-based damage assessment, and post-earthquake transportation network analysis, and positions the study within this body of work. Chapter 3 presents the technical definition of the problem, the network components, and the modeling assumptions, and outlines the main research questions and scope of the study. Chapter 4 introduces the methodological framework, from satellite-based damage detection to the generation of risk-sensitive road costs and the mathematical formulation of the optimization model. Chapter 5 reports the computational experiments and case study results based on data from the 2023 Kahramanmaraş earthquakes and discusses the findings from operational and technical perspectives. The final chapter summarizes the main conclusions and outlines directions for future research.

2 Literature Review

This chapter reviews the literature that informs the methodological choices of this thesis and positions the study within prior research on disaster response planning. The review is organized around four interconnected streams: uncertainty and decision challenges in post-disaster logistics, relief facility location and POD planning, remote sensing and deep learning methods for building damage detection, and transportation network disruption and accessibility modeling (including online routing under unknown disruptions). While each stream has developed strong results on its own, the literature often treats them as separate components, leaving a gap between image-based damage assessment and operational decision models. By synthesizing these strands, this chapter motivates the integrated end-to-end workflow proposed in this thesis, which translates detected building damage into risk-aware network costs and leverages them for accessibility analysis and POD location and demand assignment.

2.1 Uncertainty and Decision Challenges in Post-Disaster Logistics

Humanitarian logistics and disaster management involve severe uncertainty, especially during the response phase [ACM16]. Decisions must be made under time pressure, with limited resources and incomplete information. In emergency relief, time is usually the main constraint. The first 72 hours matter most, yet teams still have to act with little clarity about where needs will be highest or how much will be required [Van06].

Early work in the field emphasizes the difficulty of planning when demand, transportation conditions, and available capacity are not known with confidence. For example, Barbarosoğlu and Arda [BA04] propose a two-stage stochastic programming framework to address uncertainty in supply amounts, route capacities, and demand requirements after an earthquake. Similarly, Özdamar et al. [OEK04] combine multi-period and multi-commodity network flow decisions with routing components and note that limited and delayed data can become a major barrier to effective planning. These studies highlight that uncertainty is not a secondary issue in disaster response; it shapes both modeling choices and operational outcomes.

2.2 Facility Location and Distribution Models for Relief Operations

A large body of research studies the location of emergency facilities and PODs as part of post-disaster relief networks [CNP12]. Many models represent demand and network conditions through scenario-based structures, which makes uncertainty explicit and allows the planner to evaluate multiple possible realizations.

Balcik and Beamon [BB08] develop scenario-based variants of maximal covering formulations to support the placement of distribution centers and the planning of inventory under uncertainty. Beyond travel distance, several studies also incorporate service quality, equity, and social impact [GF18]. Loree and Aros-Vera [LA18] integrate deprivation cost ideas into facility decisions, capturing the harm experienced while access to relief is delayed. Holguín-Veras et al. [Hol+13] similarly argue for social-cost objective functions that combine logistics and deprivation costs, defining deprivation cost as the economic valuation of human suffering associated with a lack of access to a good or service. Noyan et al. [NBA16] study last-mile distribution using data from the Van earthquake and propose a stochastic optimization model that balances accessibility and equity objectives. For the Istanbul context, Salman and Yücel [SY15] consider spatial dependencies in link failures and aim to maximize expected covered demand under correlated disruptions.

A common limitation across many location–allocation studies is the way demand and damage are represented. Demand is often derived from pre-disaster proxies such as population density or risk indicators, and network disruptions are commonly imposed through assumed scenarios. This is a practical choice when real-time data are unavailable, but it can disconnect the model inputs from the evolving post-disaster situation. In Istanbul-focused work such as Yenice and Samanlĭoğlu [YS20], scenario-based modeling remains the main mechanism to reflect uncertainty in risk and unmet demand.

2.3 Remote Sensing and Deep Learning for Damage Detection

Recent advances in remote sensing and deep learning have enabled automated damage detection from post-disaster imagery. Convolutional Neural Network (CNN) are widely used to replace manual inspection with scalable, image-driven classification.

Xu et al. [Xu+19] evaluate CNN architectures for building damage assessment using satellite imagery from the Haiti earthquake and show that models comparing pre- and post-event views can achieve strong performance. Several studies propose model designs that improve accuracy and transferability. Alisjahbana et al. [ALZ+24] introduce a two-stage framework (building footprint detection followed by damage classification) and report improved scores on the xView2 benchmark. Bhardwaj et al. [Bha+25] combine ResNet- and U-Net-style components and use attention mechanisms to improve focus on damaged regions. Wang et al. [Wan+25] study adaptation across disaster types and propose semi-supervised fine-tuning strategies. Ekici [Eki21] explores hyperparameter optimization to improve classification performance under class imbalance. Nex et al. [Nex+19] discuss operational constraints and the transferability of CNN models across platforms and geographic contexts.

While these works produce accurate damage maps, their outputs are frequently used for visualization or reporting. In many cases, the detected damage information does not feed directly into a network model or an optimization formulation. This leaves a practical gap between image-based damage detection and decisions such as POD placement and last-mile assignment.

2.4 Transportation Network Risk and Accessibility Modeling

Transportation disruptions after earthquakes strongly affect relief operations, because route availability and travel costs can change rapidly. The literature models these effects using probabilistic methods, vulnerability-based measures, and simulation tools.

Yılmaz et al. [YAE19] estimate segment survivability using analytical relationships that include seismic zone factors and ground motion indicators, and show that optimal routing decisions can differ substantially before and after an event. Toma-Danila [Tom18] uses a Geographic Information System (GIS)-based framework for Bucharest and studies road blockage risk through buffers around vulnerable buildings and Monte Carlo simulation. Costa et al. [Cos+20] model debris effects and evaluate whether road links become physically blocked, then analyze resulting traffic impacts. Kilanitis and Sextos [KS19] incorporate bridge damage and time-varying demand to assess network resilience in a dynamic setting.

A recurring challenge is the translation from remote sensing outputs to network-level costs. Sakuraba et al. [Sak+16] point out that building an operational road graph from imagery-based information often involves manual or semi-automatic steps, which can be slow during the critical early response period. Related discussions also note that real-time road condition information is rarely integrated into planning algorithms in a seamless way [SMG25].

2.5 Research Gap and Positioning of This Study

Taken together, the literature shows strong contributions in three areas: (i) facility location and distribution optimization, (ii) damage detection from imagery, and (iii) transportation network disruption modeling. However, these components are often treated separately. Optimization models typically rely on assumed scenarios for demand and network condition, while damage detection studies often stop at map products. Rather than enumerating disruption scenarios, this thesis uses detected building damage as a direct, data-driven proxy for local access degradation. The resulting proxy is translated into risk-aware edge costs on the road network, providing a consistent mechanism to propagate image-based damage information into shortest-path accessibility measures and location–allocation decisions. Network disruption models may capture physical mechanisms, but do not always connect to POD location decisions through consistent inputs.

This thesis addresses this integration gap with an end-to-end workflow that links damage detection outputs to a risk-aware road cost structure and then to accessibility-based POD location and demand assignment. In the proposed approach, damage detection provides building-level labels, damaged buildings define local impact zones (buffers), these zones translate into risk penalties on road segments, and the resulting risk-aware network costs feed both shortest-path accessibility calculations and the optimization model. By using detected post-disaster conditions as model inputs, the framework reduces reliance on purely hypothetical scenarios and provides a more direct path from imagery to decision support.

3 Problem Definition

This chapter defines the decision problem studied in this thesis and clarifies the modeling scope. The focus is the early response stage after a large earthquake, where rapid planning is required while field information is limited and access conditions can deteriorate unevenly across the road network. The central decision is to select a set of PODs and assign affected demand locations to these PODs in a way that remains feasible under capacity limits and reflects disrupted access through a risk-aware road cost structure.

3.1 Last-Mile Relief Network Structure

The relief network considered in this study consists of three types of nodes: *supply entry points*, *candidate POD sites*, and *demand locations*. Supply points represent locations where aid enters the study area and serves as the origin of inbound flows. Candidate PODs are local distribution sites (e.g., public facilities or open areas) that can be activated to stage and distribute supplies. Demand locations represent aggregated zones of need derived from the spatial concentration of damaged buildings.

From an operational perspective, the planning problem can be summarized by two linked questions: (i) which candidate sites should be activated as PODs, and (ii) once sites are activated, how should demand locations be allocated to them so that service is feasible and access-aware?

3.2 Decision Problem

Let S be the set of supply points, J the set of candidate POD sites, and I the set of demand points. The decision problem is to select which candidate PODs to open, and assign each demand point to exactly one opened POD, subject to capacity limits at opened sites.

Travel between nodes is represented on a road network. Costs are computed as least-cost (shortest) path values, but the edge weights are risk-aware rather than length-only. Each road segment has a base cost derived from its length and a multiplicative penalty that increases with exposure to damaged-building impact zones. In the baseline form, the segment cost can be written as

$$\text{cost} = \text{length}_m\,(1 + \alpha \cdot \text{damage}_n),$$

where $\alpha \geq 0$ controls the contribution of damage exposure and damage_n denotes the damage intensity measure assigned to that segment (derived from buffer-based interactions with damaged buildings). Least-cost paths computed on this weighted network produce two cost matrices: Supply→POD and POD→Demand. These costs are then transformed into normalized accessibility scores and used as inputs to the optimization model.

The output of the decision problem is the opened POD set and a complete demand-to-POD assignment. In later chapters, these outputs are interpreted through accessibility measures and map-based results to illustrate how risk-aware costs influence both routing preferences and location–allocation decisions.

3.3 Research Questions

This thesis is guided by the following research questions:

1. How can damage information extracted from satellite imagery be translated into a road cost structure that reflects post-disaster access conditions?
2. To what extent does integrating satellite-derived damage information into the road cost structure influence accessibility patterns and POD location–allocation decisions under different scenario settings?
3. For the instance sizes considered in this study, can the proposed POD location and demand assignment model be solved to optimality within practical computing limits?

These questions connect the data-driven damage detection output to network analysis and, finally, to the optimization-based site selection decision.

3.4 Scope and Modeling Assumptions

The model represents a planning decision for the early response stage. It focuses on access and allocation rather than detailed operational scheduling. Several assumptions are made to keep the problem tractable and consistent with the available data and GIS layers:

- **Demand representation.** Demand points represent aggregated need locations derived from damaged buildings. Demand magnitude is proportional to the number of damaged buildings within each demand cluster.
- **Capacity proxy.** Candidate POD capacities are derived from site area as a proxy. The capacity scale is aligned with the demand scale through a conversion factor defined in the experimental setup.
- **Risk-aware but static network snapshot.** The road network is assumed to be connected after preprocessing and topology cleaning. However, isolated cases may still occur when point locations fall outside practical routing reach, in which case they are flagged and handled during the experimental phase. Disruption is represented through increased travel costs near damaged buildings rather than through explicit link removals or time-varying updates. The resulting model is a deterministic snapshot conditioned on the available damage signal.
- **Single allocation for interpretability.** Each demand point is assigned to exactly one POD. This produces a clear service partition and supports interpretable, map-based analysis of allocations. This choice trades flexibility for interpretability, which is important in early-response planning.

These assumptions do not claim to resolve all sources of uncertainty present in disaster response. Instead, they reflect what can be reliably supported with the available imagery, derived damage outputs, and road data, while still producing a decision model that can be solved and interpreted within a short planning horizon.

3.5 Chapter Link to the Methodology

The next chapter presents the methodology used to construct the model inputs and solve the decision problem. It describes how damaged buildings are detected from satellite imagery, how damage is mapped into risk penalties on road segments, how shortest-path costs are computed and converted into accessibility scores, and how these quantities are integrated into the POD location and demand assignment formulation.

4 Methodology

This chapter presents a methodological framework for post-disaster aid distribution planning. The proposed approach transforms damage information derived from satellite imagery into a risk-aware cost structure on the road network, computes accessibility measures based on these costs, and subsequently formulates an optimization model to address POD site selection and demand assignment. In this way, the transportation impacts of disaster-induced physical damage are not treated merely as a descriptive or visual output, but are explicitly embedded in the decision-making process.

The framework comprises four main components: (i) Damage Detection, (ii) Road Network Cost Assignment, (iii) Accessibility Computation, and (iv) Optimization Model. Each component produces outputs that serve as inputs to the subsequent stage. Specifically, the damage detection module identifies damaged structures and characterizes their spatial distribution; the road network cost assignment module maps this information onto the transportation network by penalizing road segments according to their exposure to damaged-building influence zones (e.g., buffer-based interactions), thereby producing a risk-aware edge-cost representation; the accessibility computation module then derives shortest-path-based cost and accessibility metrics on the resulting network; and the optimization model leverages these metrics to determine POD locations and assign demand points to the selected PODs.

4.1 Preliminaries and Implementation Context

The workflow links image-based damage detection, GIS-based processing, and road-network analysis within a single pipeline. To clarify the implementation context, two points are stated upfront.

First, the damage detection step uses a deep convolutional neural network with a ResNet18 backbone [He+16]. ResNet is employed as a standard image-recognition backbone to extract visual features from building patches and produce binary damaged/undamaged labels; architectural details and training settings are provided in Section 4.3.

Second, geospatial processing was carried out in a GIS environment using QGIS [QGI25]. QGIS supports layer-based handling of raster and vector data, basic alignment checks, and the export of intermediate products reused in later stages. Similar steps could be implemented with other GIS software or a fully code-based stack; QGIS provides an efficient and reproducible setup for the scope of this study.

4.2 Data Inputs and Preprocessing

This section describes the geospatial datasets used in the proposed framework and the preprocessing steps applied before network analysis and optimization. The workflow uses (i) post-disaster satellite imagery for damage detection, (ii) pre-disaster vector layers for buildings and roads, and (iii) point layers for supply locations and candidate PODs. Demand locations are not provided as an input. Instead, demand points are generated from the spatial distribution of the detected damaged structures (Section 4.2.4). The post-disaster input is a georeferenced optical satellite raster that provides the visual basis for building-level inference. Dataset-specific details (acquisition time, provider, and spatial resolution) are reported in the experimental setup.

All datasets were imported into a single QGIS project and overlaid using their embedded georeferencing information, enabling a direct visual verification that pre-disaster building footprints and road geometries coincide with their expected locations in the post-disaster imagery. In practice, post-disaster imagery is first compiled to provide the visual and spatial basis for the damage detection module, while building footprints and road geometries corresponding to the closest available pre-disaster period are obtained from OpenStreetMap (OSM) and imported as vector layers. In this study, these OSM vector layers were obtained from the Geofabrik distribution of regional extracts and imported into QGIS [Ope25; Geo25]. After confirming the expected overlay in QGIS, these standardized layers form the final inputs for damage detection and provide the consistent spatial foundation for subsequent risk-aware road cost assignment, shortest-path-based accessibility computation, and location–allocation optimization.

4.2.1 Geospatial Data Layers and Key Attributes

The spatial data used in this study were considered in two categories: input layers and layers derived within the study. Input layers provide the basic geometries necessary for damage detection and network analysis.

Derived layers are generated from these inputs and are used in the accessibility calculations and the optimization model. The input layers include: (1) post-disaster satellite imagery (raster), used as the visual and spatial reference for damage detection; (2) pre-disaster building footprints (polygons) and (3) road geometries (lines), obtained from OSM as vector representations of the built environment and the transportation network; (4) supply points, representing the origins of the analysis; and (5) candidate

POD locations, created from OSM; Points of Interest (POI), landuse, and Points of Worship (POFW) classes (the selection and filtering steps are described in Section 4.2.4). Demand points are not available as a standalone dataset. Instead, they are generated from the spatial distribution of the damaged structures identified by the damage detection step (Section 4.2.4). All vector layers include unique identifier fields to support cost-matrix generation in later stages and to index the optimization model. Pre-disaster building footprints and road centerlines are obtained from an open mapping source and used as baseline vector layers. The exact extraction date and filtering rules are documented in the experimental setup.

4.2.2 Coordinate Reference System and Layer Alignment

All spatial data were managed in a single QGIS project. Since each layer included georeferencing information, QGIS overlaid the datasets automatically based on their coordinates. This is important for later spatial operations such as buffer generation, area calculation, and intersection.

Before the analysis, the alignment between the pre-disaster vector layers (building footprints and road geometries) and the post-disaster satellite imagery was checked visually in QGIS. Road geometries generally aligned well with the imagery. For building footprints, small local deviations can occur due to parallax effects. The impact of this issue is discussed in later sections.

To ensure consistent units for distance and area-based operations, all calculations were performed in a projected Coordinate Reference System (CRS) with metric units. When needed, layers were reprojected to the selected CRS and saved to support reproducibility (Project CRS: EPSG:32637).

4.2.3 Road Network Preparation

Road network data were obtained in polyline format and prepared for shortest-path calculations and segment-level cost assignment. Before computing shortest paths, the road layer was inspected to avoid common issues that can break routing results. Inconsistent or repeated segments and occasional gaps were removed or fixed where needed to improve connectivity. This step helps reduce routing and cost-calculation errors in later stages. To ensure routability, the polyline layer was converted into an edge-based network by splitting lines at intersections and removing duplicate/invalid geometries. Small gaps and overshoots that prevent connectivity were corrected using standard GIS topology cleaning operations; dataset-specific thresholds are provided in the experimental setup.

The road layer was treated on a segment (edge) basis. Routing is performed on an undirected graph with edge weights derived from the risk-aware cost field. Directional restrictions and speed profiles are not modeled, as the objective is to represent relative post-disaster accessibility under limited information. For each road segment, the unique identifier field already available in the dataset was used. Since segment length was not available as an attribute in the original layer, it was computed using the Field Calculator ($length) and added as a new field. This length value was used as the initial cost before applying any risk penalty. In the next step (Section 4.4), damage-related risk effects were incorporated into these initial values to obtain the final road costs.

The output of this preprocessing stage is a topologically consistent road network with basic attributes for each segment, including an identifier and length. This dataset is used as input for risk-aware cost assignment (Section 4.4) and shortest-path-based accessibility calculations (Section 4.5).

4.2.4 Point Layers: Supply, Candidate PODs, and Demand

This subsection describes how the point layers were prepared for the analysis. The study uses three main point layers: supply points, candidate PODs, and demand points.

Supply and candidate PODs form the set of decision locations in the optimization model. Demand points are not available as a standalone dataset. Instead, they are derived from the locations of damaged buildings identified by the damage detection module (see Section 4.3).

Supply Points

Supply points were placed to reflect incoming aid, and they were selected along the main access roads into the settlement, at locations suitable for network-based analysis. This choice captures the typical entry directions of incoming aid in a simple but realistic way. Each supply point is assigned a unique identifier and is used as an origin in the subsequent supply-to-POD accessibility calculations. These locations are not intended to represent exact logistics hubs, but rather stylized entry points that capture dominant access directions into the study area.

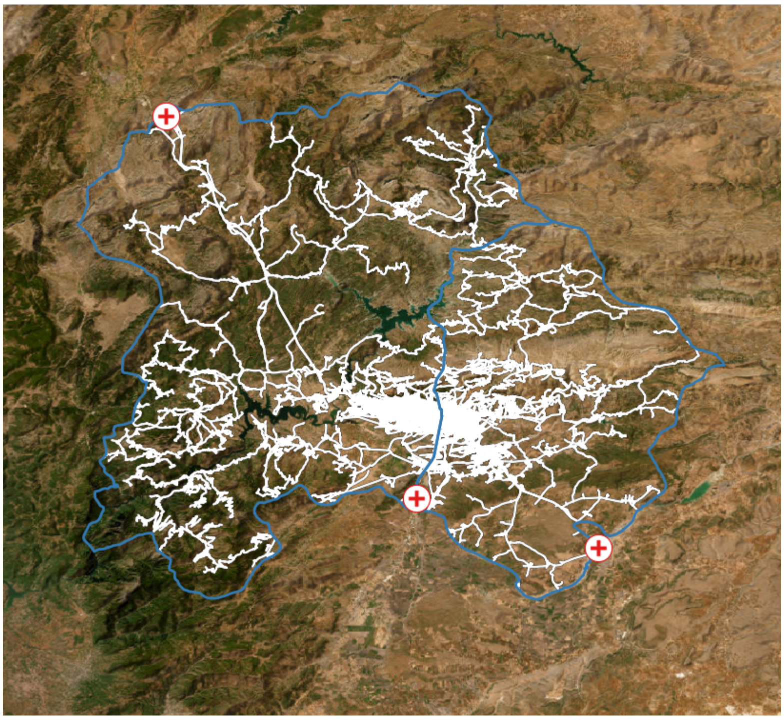

Figure 4.1: Supply points representing stylized entry locations for incoming aid in the road network.

Candidate PODs

The candidate POD set was built from OSM data. Three OSM datasets were used: POI, POFW, and landuse. Each dataset was available as both point and polygon layers. To work with a point-based candidate set, areas were computed from the polygon layers and added as attributes to the corresponding point layers. The point layers were then used in the remaining steps.

From the POI layer, facility types that can serve as PODs were selected: *school*, *hospital*, *clinic*, *university*, *stadium*, *sports_centre*, *market_place*, *shelter*, *playground*, *camp_site*, and *caravan_site*. From the landuse layer, open-space classes were selected: *park*, *grass*, *meadow*, *orchard*, and *farmland*. POFW features were also included.

The area attribute (m^2) was used in two ways. A minimum-area threshold was applied to remove very small sites that are unlikely to operate as PODs; the adopted threshold is reported in the experimental setup. The same area value was also used as a proxy for POD capacity, based on the assumption that larger sites can serve more demand. Finally, candidates from the three sources were merged into a single dataset and duplicates were removed. The resulting candidate POD set was used in the accessibility calculations and in the optimization model.

Demand Points

Demand points were not available as a standalone dataset in this study. Instead, they were derived from the locations of damaged buildings identified by the damage detection module (see Section 4.3). The goal was to represent demand with a smaller set of points, rather than using one point per building.

Damaged building points were clustered using k-means [Mac67]. Since k-means requires the number of clusters **K** to be specified in advance, **K** was selected based on the total number of damaged buildings and the assumed service capacity of PODs. Several nearby **K** values were also tested, and the resulting spatial patterns were visually inspected to avoid unrealistic cluster shapes. Clustering is performed in projected (metric) coordinates. To ensure reproducibility, the k-means procedure is run with a fixed random seed and multiple initializations. The selected number of clusters K follows a capacity-driven granularity criterion, and the final K used in the case study is reported in the experimental setup.

The clustering step is not intended to recover true neighborhoods, but to obtain a computationally manageable and spatially representative demand abstraction that can be consistently integrated into the network-based accessibility analysis and the optimization model [IS15].

One demand point was created for each cluster. Its location was defined as the cluster center, computed as the mean of the coordinates of the damaged buildings belonging to that cluster. Demand magnitude was defined as the number of damaged buildings within the cluster. These demand points were subsequently used in the accessibility calculations and in the optimization model.

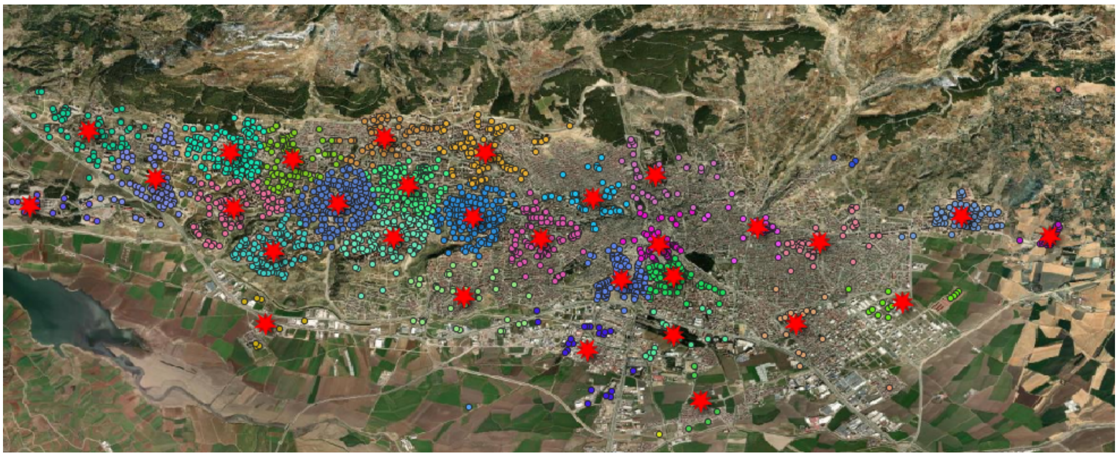

Figure 4.2: Spatial aggregation of damaged buildings into representative demand points using k-means clustering. Colored points indicate damaged buildings assigned to each cluster; red stars denote cluster centers used as demand points in the optimization model.

4.3 Damage Detection

In this thesis, the purpose of the damage detection step is to identify damaged buildings from post-disaster satellite imagery and provide input to two components: deriving demand points from the spatial concentration of damaged buildings, and translating the effects of destruction on the road network into risk penalties at the road-segment level. For this reason, damage detection is formulated as a binary classification task (damaged/undamaged), rather than a detailed damage grading scheme.

The model development followed two stages. First, a ResNet18-based classifier was trained using a dataset of building image patches derived from post-earthquake satellite imagery and labeled as damaged or undamaged. During training and validation, visually similar samples were grouped to avoid placing near-duplicate images in different splits. The data split was then performed at the group level to reduce optimistic performance estimates. Second, the trained model was applied to the study area. The post-disaster image was cropped at each pre-disaster building footprint. Each crop was then classified as damaged or undamaged. The predicted labels were assigned back to the corresponding building polygons and displayed in QGIS as a thematic layer. The set of polygons labeled as damaged was exported as the damaged-building layer and used in the next steps of the workflow. Remaining uncertainties related to image quality and layer alignment are discussed under the limitations of this section.

4.3.1 Training Data and Model Development

The damage detection model was trained on a post-earthquake dataset of cropped building images derived from a recent large-scale earthquake event. As illustrated in Figure 4.3, the dataset includes both undamaged and damaged building patches with varying visual conditions.

Each sample corresponds to a building-area crop extracted from post-disaster satellite imagery.

Figure 4.3: Example image patches from the training dataset. Left: undamaged building samples. Right: damaged building samples.

In this thesis, the task is formulated as a binary classification problem (damaged versus undamaged) to align with the requirements of the decision model.

ResNet18 was used as the backbone architecture and fine-tuned from pretrained weights. All images were resized to a fixed input resolution and normalized prior to training. A compact training setup was adopted, consisting of 8 epochs, a batch size of 8, and a learning rate of 2×10^{-4}.

The dataset may contain visually similar image patches. If near-duplicate images are assigned to different data splits, validation performance may be overestimated. To mitigate this issue, images were grouped based on perceptual similarity and kept together during data splitting. Similarity groups were constructed using perceptual hashing (pHash), and a 5-fold GroupKFold scheme was applied [Ped+11; Buc].

Model performance was evaluated on the validation set using the F1 score and the

Area Under the ROC Curve (AUC). Validation performance remained consistently high across epochs, with F1 scores ranging from approximately 0.982 to 0.988 and AUC values around 0.999. Validation probabilities were calibrated, and the decision threshold was selected to maximize the validation F1 score, resulting in a threshold value of $\tau = 0.464$. Under this threshold, the validation accuracy reached 0.991. The precision–recall trade-off was summarized using a Precision–Recall (PR) curve, which indicates that precision remains high across a wide range of threshold values. While validation scores are high, these values should be interpreted in light of the controlled dataset structure and the binary formulation of the task.

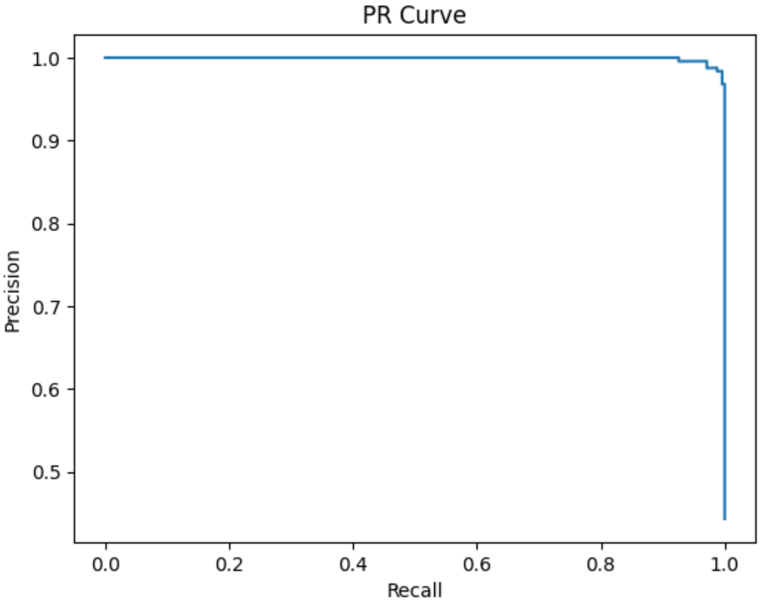

Figure 4.4: PR curve of the damage detection model on the validation set. The curve demonstrates consistently high precision across nearly the full recall range.

4.3.2 Inference on the Study Area

The trained damage detection model was applied to the study area through the pre-disaster building polygons. This step labels each building footprint and creates an output layer for GIS analysis. The building footprints and the post-disaster image share the same coordinate reference system. A Python script crops the post-disaster image using each building polygon. For each footprint, a local image patch is extracted using the footprint extent with a small padding margin, then resized to the network input resolution for classification. This standardizes the inference pipeline across buildings. Each crop was classified as damaged or undamaged. The labels were mapped in QGIS as a thematic layer. Footprints classified as damaged were displayed in red, while undamaged footprints were displayed in green.

Polygons classified as damaged were exported as a separate layer. This layer was used in two later steps. First, it supported the derivation of demand points and the clustering procedure (Section 4.2.4). Second, it served as the source layer for assigning risk penalties to road segments in the cost model (Section 4.4).

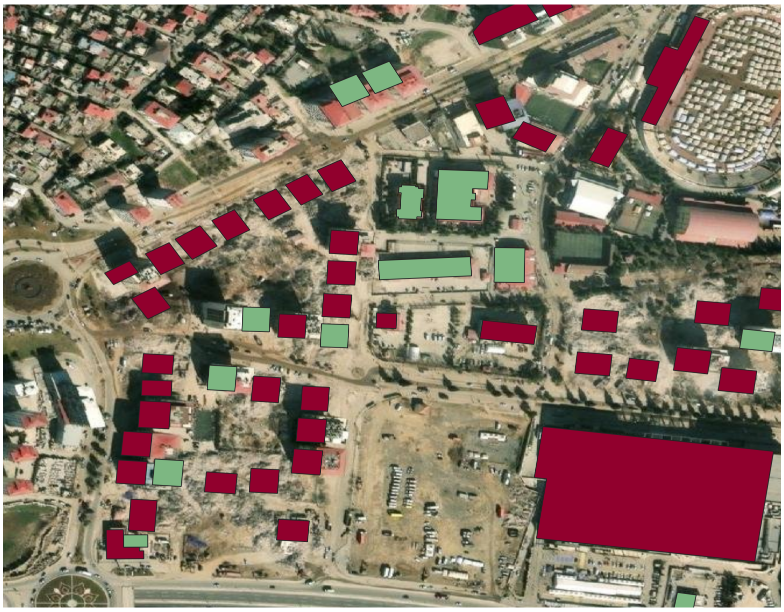

Figure 4.5: Damage detection output mapped on building footprints in QGIS (red: damaged, green: undamaged) over post-disaster satellite imagery.

4.3.3 Limitations

The damage detection step depends on the readability of the post-disaster satellite imagery and on the local overlap between the imagery and the pre-disaster building footprints. Small mismatches can occur in practice. Parallax is a common cause, since the footprint and the visible roof area may not fully align. When this happens, the cropped patch may capture the building only partially, and the model output may change. This can cause some undamaged buildings to be labeled as damaged.

Errors can also increase when the cropped patch contains limited detail. These issues arise at the building level and may carry over to later steps that use the damaged-building layer. Their impact is assessed in the computational experiments and results section.

An important characteristic of the damage detection output is that misclassifications are not symmetric. Due to parallax effects and partial roof visibility, some undamaged buildings may be labeled as damaged, whereas the opposite error (classifying severely damaged buildings as undamaged) was observed much less frequently in qualitative inspection. As a result, the damage signal used in this study can be interpreted as conservative: it may overestimate local damage and exposure, but is less likely to miss heavily damaged structures.

Since the damaged-building layer directly influences the exposure-based road costs, any systematic bias in classification may propagate into the accessibility measures. In particular, false positives (undamaged buildings classified as damaged) can lead to locally inflated road costs and reduced accessibility in those areas. However, because the cost adjustment is proportional and controlled by the parameter α, the network remains connected and routes are penalized rather than removed. This design choice limits structural distortions in the accessibility structure while still discouraging traversal through areas flagged as highly exposed.

4.4 Road Network Cost Assignment

The distribution of damaged buildings obtained from damage detection was used to define a risk component on the road network. The goal is to avoid routes that depend only on geometric length, which can be misleading under post-disaster conditions. Road segments near damaged buildings were assigned an additional penalty, which makes the network costs risk-aware.

For each segment, length was used as the base cost. The damage effect was derived from impact zones around damaged buildings. Final segment cost was then computed by combining length with this damage effect. The resulting risk-aware costs were used to build shortest-path cost and accessibility matrices, and to provide cost inputs for the optimization model.

4.4.1 Input Data and Base Road Cost

Road segments and the set of buildings labeled as damaged are used together in this stage. The road layer is processed segment by segment. Each segment has its own ID, and the cost is defined for that segment. The damaged-building layer consists of building polygons labeled as damaged and provides the risk source. Base road cost is defined as segment length. Length is computed in a metric coordinate system and stored as an attribute in the road layer. This value is the cost before any damage-related penalty. It is updated in the next stage using the damage impact to obtain the final risk-aware cost.

4.4.2 Buffer-Based Impact Zones Around Damaged Buildings

Damaged buildings can affect local road access. Debris and nearby disruptions may make travel more difficult in their immediate surroundings. To reflect this impact on network costs, a fixed-radius impact zone has been defined around each damaged building. In practice, as illustrated in Figure 4.6, a buffer with a fixed radius was created for each of the damaged building polygons, which are the damage detection output. The resulting buffer layer was used as a spatial mask representing the local impact of the destruction on the road network. This mask was then used in the next step to determine whether road segments intersected with these impact zones and to quantify the damage impact on a segment-by-segment basis. The buffer radius is treated as a fixed parameter in the analysis, and its specific value is defined in the next step.

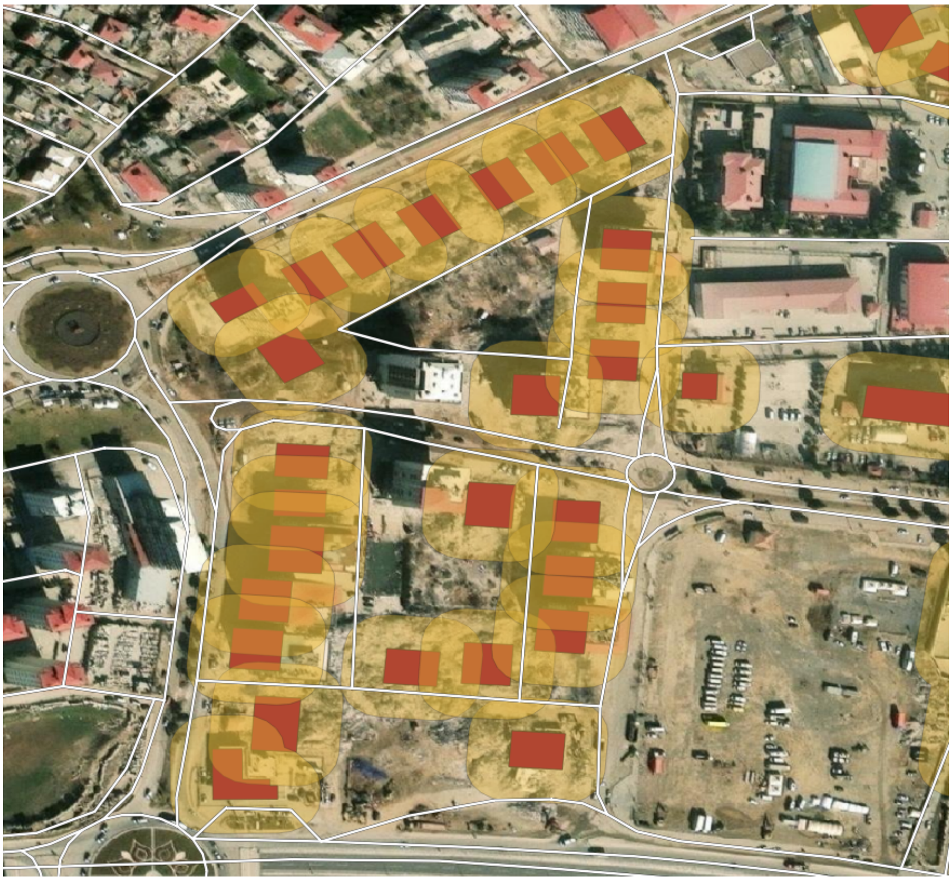

Figure 4.6: Illustration of fixed-radius impact zones around buildings labeled as damaged (buffers) and the road segments intersecting these zones. The number of intersected buffers per segment defines the exposure measure $damage_n$ used in risk-aware edge costing.

To justify the choice of the buffer radius, we rely on a physical interpretation of how collapse-induced debris and response-related constraints can reduce road usability near damaged structures. To capture the localized effect of structural failure on the road network, damage impact zones are created by applying a fixed buffer radius r to each building identified as damaged. In the baseline setting, r is set to 20 meters. This value is specifically chosen to reflect the potential debris footprint of the 3 to 8-story reinforced concrete buildings prevalent in Kahramanmaraş's urban centers [AFA23].

The selection is strongly informed by the empirical correlations proposed by Argyroudis et al. [APA05], which relate the number of storeys (N) to the induced debris width (W_d). According to the average correlation curve ($N = 0.06W_d^2 + 0.35W_d$) provided in that study, an 8-story building ($N = 8$) results in a mean debris spread of approximately 9 meters [APA05]. However, Argyroudis et al. [APA05] emphasize that as building height increases, the dispersion of predicted debris values among different collapse models grows significantly, with "Model C" representing shapes with much wider debris footprints.

Beyond the physical footprint of the wreckage, the logic in [APA05] recognizes that road accessibility is a calculation that factors in the building's set-back distance from the street and the "effective width" required for rescue vehicles to maneuver safely. By adopting a 20-meter radius, we move beyond a simple measurement of debris to create a realistic safety margin that accounts for both unpredictable collapse patterns and the operational space needed for emergency teams.

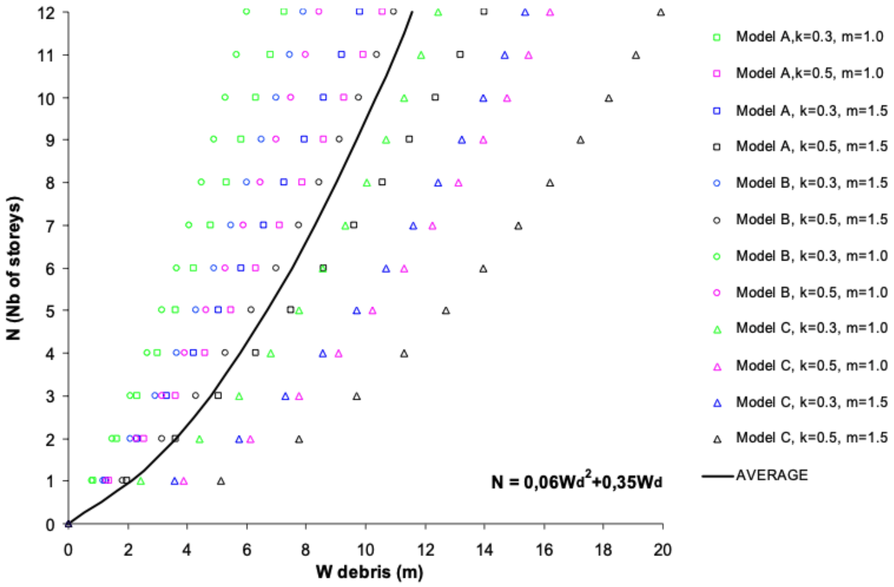

Figure 4.7: Correlation between building height (number of storeys) and induced debris width, showing the mean curve and the variation across different collapse models [APA05]. For an 8-story building, while the mean spread is approximately 9 meters, the model dispersion and operational factors support the use of a more conservative 20-meter buffer radius.

4.4.3 Damage Exposure and Risk-Aware Costing in Road Segments

At this stage, the goal is to measure how damage affects each road segment and to incorporate this effect into segment cost. Road segments were overlaid with the fixed-radius buffer layer defined in Section 4.4.2. For each segment, the number of buffer zones it intersects was computed and denoted by $damage_n$.

Damage exposure is quantified as the *count* of intersected impact zones, rather than the intersection length along the segment. While a length-based exposure measure could provide a finer-grained representation, the count-based measure is preferred to remain consistent with the binary nature of the damage detection outputs and to avoid introducing artificial precision under positional uncertainty (e.g., parallax-related footprint misalignment).

Segments with $damage_n = 0$ are not exposed to nearby destruction, while $damage_n > 0$ indicates segments that pass through one or more impact zones.

Segment costs were then updated using $damage_n$. Length in meters, $length_m$, was used as the base cost. Damage exposure scales this base cost through a multiplier. The cost function is

$$\text{cost} = length_m \left(1 + \alpha \, damage_n\right).$$

Here, α sets the strength of the damage penalty in the road cost. Since the damage signal is derived from satellite imagery and is not a direct measurement of travel time, α is not interpreted as a physically calibrated constant; rather, it acts as a conservative risk-aversion parameter that discourages routing through exposed segments while keeping the network usable. If $damage_n = 0$, the cost reduces to the segment length and no damage effect is applied.

As $damage_n$ increases, the segment cost increases in proportion. The coefficient α is fixed in the computational experiments, and the adopted value is reported in the experimental setup. These risk-aware costs are then used as edge weights in the shortest-path calculations that produce the cost matrices for the optimization model.

4.4.4 Output Layer and Cost Field

At the end of this step, the road layer was updated with the risk-aware costs and saved for later analyses. Two new attributes were added to each road segment: the exposure value $damage_n$ and the corresponding final cost. Each segment therefore retains both its length-based base cost and the updated cost that reflects nearby damage.

Shortest-path calculations in this study use the final cost attribute as the routing weight. The same attribute is used when generating the cost matrices required by the optimization model, which keeps the cost computations consistent across stages. The exported layer represents the risk-aware road network and is used in the accessibility calculations presented in the next section.

4.5 Accessibility Computation

This section computes accessibility using the road network updated with risk-aware costs in the previous step. The goal is to estimate travel costs between the main point sets in the aid distribution system and to produce the matrices required by the optimization model. The costs reflect both route length and the added difficulty near locations identified as damaged through damage detection. Shortest-path costs are computed between supply points, candidate POD locations, and demand points. The results are exported as cost matrices and then transformed into accessibility scores on a common scale.

4.5.1 Network Entities, Preparation and Shortest-Path Cost Computation

Accessibility was computed between three groups of locations: supply points, candidate POD sites, and demand points. Supply points represent entry locations for incoming aid and act as origins in the network analysis. Candidate POD locations form the set of potential distribution sites and are used both in cost computation and in the optimization model. Demand points are derived from the spatial pattern of buildings labeled as damaged through damage detection and are aggregated into a smaller set of representative points via clustering (Section 4.2.4).

Shortest-path analysis requires each point to be connected to the road network. Points were therefore snapped to the nearest road segment before routing. This places each location on the network so it can be used as an origin or destination. The snapping step also reduces routing errors that can arise from small positional offsets. All points are projected onto the nearest road segment prior to routing to guarantee network connectivity. A maximum snapping tolerance δ is used to prevent unrealistic connections; points exceeding δ are flagged for inspection. The value of δ is specified in the experimental setup.

Cost calculations used the risk-aware road network from Section 4.4. Each road segment was weighted by its final cost value. Shortest-path costs were then computed for all pairs in two sets: *supply→candidate POD* and *candidate POD→demand*. For each pair, the result is a single cost value given by the sum of segment costs along the selected route. The minimum-cost paths were obtained with Dijkstra's algorithm [Dij59].

4.5.2 Cost Matrices and Exported Model Inputs

Shortest-path results were prepared in two sets for the optimization model. The first set covers Supply–POD costs, defined as the travel cost from each supply point to each candidate POD. The second set covers POD–Demand costs, defined as the travel cost from each candidate POD to each demand point. Each entry is the minimum-cost path value computed on the risk-aware road network. Using the same cost definition for both sets keeps supply-side access and demand-side service costs comparable.

Costs were exported in a long-table format. In the POD–Demand file, each row contains the fields demand_id, pod_id, and cost. Supply–POD costs follow the same structure. This layout fits indexed parameters in the optimization model and scales well as the number of pairs grows.

The cost values were normalized to obtain accessibility scores that can be compared across origin–destination pairs. Accessibility is denoted by v_{ij} for the POD–Demand pairs and by v_{sj}^{SP} for the Supply–POD pairs. The normalization rule and the exact transformation are defined in Section 4.6.

Additional model inputs were saved as separate CSV tables. The demand table lists each demand point and its magnitude, defined as the number of damaged buildings in the corresponding cluster. The candidate POD table includes the candidate ID and site area; area is used as a proxy for capacity. The supply table lists the supply point IDs. All tables follow a consistent ID scheme so they can be joined reliably in the optimization model. For the study instance, least-cost path values were successfully computed for all required Supply–POD and POD–Demand origin–destination pairs, yielding complete cost tables for the optimization model.

4.6 Optimization Model

This section formulates the POD location and demand assignment problem. The formulation uses the risk-aware network costs and the accessibility scores computed in the previous steps. A limited number of PODs can be opened. Each demand point is assigned to one of the opened PODs, subject to capacity constraints. The objective accounts for accessibility on two links: POD–demand and supply–POD. The next part defines the sets, parameters, and decision variables, and then states the objective function and constraints.

4.6.1 Sets, Parameters, and Decision Variables

Table 4.1 lists the sets, parameters, and decision variables used in the formulation.

Table 4.1: Model Notation

Notation	Description
Sets and Indices	
I	Set of demand points (demand nodes), index i
J	Set of candidate POD locations, index j
S	Set of supply points, index s
Parameters	
d_i	Demand at demand point $i \in I$ (e.g., building count)
K_j	Capacity of POD $j \in J$ (area-proportional capacity, scaled to total demand)
κ	Capacity scaling parameter (total capacity is set to $\kappa \sum_{i \in I} d_i$)
$C \leq \lvert J \rvert$	Maximum number of PODs allowed to open
c_{ij}	Least-cost path cost from POD j to demand point i
c_{sj}^{SP}	Least-cost path cost from supply point s to POD j
$c_{\min}$, $c_{\max}$	Minimum/maximum values of c_{ij} over all (i, j) pairs
$c_{\min}^{SP}$, $c_{\max}^{SP}$	Minimum/maximum values of c_{sj}^{SP} over all (s, j) pairs
ν_{ij}	Normalized POD→Demand accessibility score (from j to i)
ν_{sj}^{SP}	Normalized Supply→POD accessibility score (from s to j)
ν_j^0	Aggregated normalized Supply→POD score for POD j
w_{PD}	Weight of the POD→Demand accessibility term
w_{SP}	Weight of the Supply→POD accessibility term
Decision Variables	
$y_j \in \{0, 1\}$	1 if POD j is opened, 0 otherwise
$x_{ij} \in \{0, 1\}$	1 if demand point i is assigned to POD j, 0 otherwise

The parameters defined in Table 4.1 are combined in the optimization model so that demand allocation and capacity constraints are expressed on a consistent basis. Since demand is represented by aggregated damaged-building counts and POD capacity is proxied by site area, capacities are constructed using an area-proportional allocation rule with a capacity buffer. Let $area_j$ denote the usable site area of candidate POD j and let total demand be $\sum_{i \in I} d_i$. Total available capacity is set to a fixed multiple of total demand, controlled by a buffer factor $\kappa > 0$, and distributed across candidates proportionally to their area:

$$K_j = \kappa \left(\sum_{i \in I} d_i \right) \frac{area_j}{\sum_{k \in J} area_k}.$$

In the implementation, the same parameter κ is denoted as `SCALAR`. This construction preserves monotonicity (larger sites imply larger capacity) and yields $\sum_{j \in J} K_j = \kappa \sum_{i \in I} d_i$. In this sense, κ acts as a capacity buffer that provides slack against proxy and measurement uncertainty. The value of κ is reported in the experimental setup.

The area-based construction of K_j is intended as a structural proxy rather than a precise operational capacity estimate. In the absence of detailed information on site layout, staffing, storage, and handling rates, site area provides a consistent and reproducible scaling rule that preserves relative differences across candidate locations. The buffer factor κ introduces controlled slack to absorb representation uncertainty and to keep the assignment model feasible under proxy-based capacity estimates.

4.6.2 POD–Demand Accessibility

For each candidate POD $j \in J$ and demand point $i \in I$, a normalized accessibility score ν_{ij} is derived from the corresponding least-cost path value c_{ij}. The score is obtained via inverse min–max normalization to convert travel costs into a standardized accessibility index [Han59]:

$$\nu_{ij} = \frac{c_{\max} - c_{ij}}{c_{\max} - c_{\min}}, \tag{4.1}$$

where $c_{\min}$ and $c_{\max}$ are the minimum and maximum values of c_{ij} over all (i, j) pairs.

4.6.3 Supply–POD Accessibility

For each supply point $s \in S$ and candidate POD $j \in J$, the least-cost path value is denoted by c_{sj}^{SP} and normalized as:

$$\nu_{sj}^{SP} = \frac{c_{\max}^{SP} - c_{sj}^{SP}}{c_{\max}^{SP} - c_{\min}^{SP}}, \tag{4.2}$$

where $c^{SP}_{\min}$ and $c^{SP}_{\max}$ are the minimum and maximum values of c^{SP}_{sj} over all (s, j) pairs.

4.6.4 Aggregation of Supply Accessibility

A single supply accessibility score is computed for each candidate POD j:

$$\nu^0_j = \frac{1}{|S|} \sum_{s \in S} \nu^{SP}_{sj}. \tag{4.3}$$

4.6.5 Model Overview

The model decides which candidate PODs to open and how to assign demand points to the opened sites. Variable y_j indicates whether candidate site j is opened. Variable x_{ij} assigns demand point i to POD j. Each demand point is assigned to exactly one opened POD, and the total assigned demand cannot exceed the capacity of that POD. The number of opened sites is limited by C.

The objective combines two accessibility terms. The first term rewards assignments that provide better POD→Demand accessibility, weighted by the demand magnitude d_i. The second term favors opening sites that are easier to access from the supply side. Weights w_{PD} and w_{SP} control the relative influence of these two components.

4.6.6 Objective Function

The objective maximizes overall accessibility by combining the POD–Demand and Supply–POD accessibility terms.

$$\max \quad w_{PD} \sum_{i \in I} \sum_{j \in J} \nu_{ij}\, d_i\, x_{ij} \;+\; w_{SP} \sum_{j \in J} \nu^0_j\, y_j \tag{4.4}$$

4.6.7 Constraints

The formulation is subject to the following constraints.

1. Maximum number of PODs that can be opened:

$$\sum_{j \in J} y_j \leq C \tag{4.5}$$

2. Assignment of demand points: Each demand point must be assigned to exactly one POD.

$$\sum_{j \in J} x_{ij} = 1 \qquad \forall i \in I \tag{4.6}$$

3. Assignment only to opened PODs: A demand point can only be assigned to an opened POD.

$$x_{ij} \leq y_j \qquad \forall i \in I, \; \forall j \in J \tag{4.7}$$

4. Capacity limitation of PODs: The total demand allocated to a POD cannot exceed its capacity.

$$\sum_{i \in I} d_i \, x_{ij} \leq K_j \, y_j \qquad \forall j \in J \tag{4.8}$$

5. Opened PODs cannot remain empty: This constraint avoids degenerate solutions in which a POD is opened solely to benefit from its supply-side accessibility score, without serving any demand.

$$\sum_{i \in I} x_{ij} \geq y_j \qquad \forall j \in J \tag{4.9}$$

6. Binary decision variables:

$$x_{ij} \in \{0, 1\} \qquad \forall i \in I, \; \forall j \in J \tag{4.10}$$

$$y_j \in \{0, 1\} \qquad \forall j \in J \tag{4.11}$$

5 Computational Experiments

This section reports the computational experiments on the Kahramanmaraş case. The experiments have two aims. First, they examine how mapping damage information to risk-aware road costs affects accessibility and the resulting POD site selection. Second, they demonstrate that the POD location and demand assignment model is solvable at the instance sizes considered in this study. Unless stated otherwise, the study area, input tables, and solver settings are kept fixed, and results are reported through the opened POD set and the demand–POD assignments, supported by map-based outputs.

5.1 Experimental Setup

This section describes the study area and data sources, the instance size and input tables, the baseline parameter setting, and the solver environment.

5.1.1 Study Area and Data Sources

The experiments focus on a study area covering the Onikişubat and Dulkadiroğlu districts of Kahramanmaraş. These districts were among the most heavily affected regions that sustained significant structural damage during the February 2023 earthquake sequence [AFA23]. District boundaries were obtained by filtering a vector layer of Turkey's administrative units, and all analyses were restricted to this extent. Pre-disaster building footprints and road geometries were taken from OSM as vector layers. Post-disaster satellite imagery was added in QGIS through an online map service and overlaid with the pre-disaster layers in the same project. This setup provides a shared spatial reference for cropping building patches in the damage detection step and for the spatial operations used in later stages.

To clarify the temporal meaning of the datasets used in the case study, the following note applies. Pre-event building footprints and road geometries are used as reference geometries to locate where buildings and road segments are expected to be, and the post-event satellite imagery is used to assess whether the corresponding buildings appear damaged or collapsed. Accordingly, the reported dates refer to the snapshot dates of the underlying data sources (i.e., when the vector layers and the imagery represent the situation), rather than to access or download dates.

All OSM-derived vector layers (building footprints, road geometries, and the candidate POD set) correspond to the 01.01.2023 snapshot. Post-disaster imagery was obtained in QGIS via the Esri World Imagery Wayback service [Esr23], using the *World Imagery (Wayback 2023-03-15)* layer (Web Map Tile Service (WMTS)), which provides an imagery snapshot dated 15.03.2023. The imagery layer was used as the visual basis for the damage detection step, while the vector layers provide the spatial index required for patch cropping and subsequent network-based computations.

5.1.2 Instance Size and Input Tables

The experiments use three point sets: supply points, candidate POD sites, and demand points. The instance includes 3 supply points ($|S| = 3$), 562 candidate PODs ($|J| = 562$), and 43 demand points ($|I| = 43$). Demand points are not provided as a standalone dataset. Instead, they are derived from building footprints labeled as damaged by the damage detection step and are aggregated into a smaller set of representative locations.

Demand points were generated by clustering the locations of buildings classified as damaged using k-means. To reduce fragmentation in the resulting demand representation, a small number of very small clusters were merged into nearby clusters based on spatial proximity. The final set used in the case study consists of $|I| = 43$ demand points, where each demand magnitude corresponds to the number of damaged buildings assigned to the cluster.

Model inputs are provided as separate CSV files. `pods.csv` lists candidate POD IDs with their site areas. `demands.csv` lists demand IDs with demand magnitudes. `supplies.csv` contains the supply point IDs. Network costs are stored in two files: `cost_sp.csv` for Supply–POD pairs and `cost_pd.csv` for POD–Demand pairs. Cost tables are saved in long format, where each row records one origin–destination pair and its shortest-path cost.

5.1.3 Key Parameter Settings

This subsection reports the baseline parameter setting used in the Kahramanmaraş case and briefly explains the role of each choice. Unless stated otherwise, experiments use this setting.

Snapping tolerance (δ). All supply points, candidate PODs, and demand points were snapped to the nearest road segment prior to routing. A maximum snapping tolerance δ was enforced to avoid unrealistic connections. In the baseline setting, $\delta = (30)$ m.

Risk parameter (α). The parameter α controls the contribution of damage detection outputs to road segment costs (Section 4.4). In the baseline setting, $\alpha = 0.2$. With this value, each exposure count increases the edge cost by approximately 20% (e.g., $damage_n = 1 \Rightarrow 1.2$, $damage_n = 2 \Rightarrow 1.4$).

Segments near damaged buildings become more costly, so routes tend to avoid these areas when alternatives exist, while connectivity remains intact across the study area. This magnitude is intended to represent a moderate post-disaster slowdown due to debris, congestion, and operational risk, without effectively imposing hard road closures.

Buffer radius (r). To model localized disruption effects, the baseline buffer radius is set to $r = 20$ m. The physical motivation and literature-based rationale for this choice are provided in Section 4.4.2 (see also Figure 4.7). A systematic sensitivity analysis over alternative r values is left for future work.

Maximum number of PODs that can be opened (C). The parameter C caps how many candidate PODs can be opened. In the baseline setting, $C = 100$. Given $|J| = 562$, this value is not intended to be tight. It mainly prevents opening many small sites when capacity becomes binding.

Objective function weights (w_{PD} and w_{SP}). The objective includes POD–Demand and Supply–POD accessibility terms. In the baseline setting, $w_{PD} = 1$ and $w_{SP} = 0.05$. This keeps the focus on serving demand points, while still favoring sites that are not hard to reach from the supply side.

Capacity derivation and buffer factor (κ). Candidate POD capacity is derived from site area and expressed on a building-equivalent scale. Using the area-proportional rule, capacities are set as

$$K_j \; = \; \kappa \left(\sum_{i \in I} d_i \right) \frac{\text{area}_j}{\sum_{k \in J} \text{area}_k}.$$

The baseline setting uses $\kappa = 1.5$.

5.1.4 Solver and Computing Environment

All optimization models were solved using the Gurobi Optimizer (version 13.0.0) [Gur24] through the Python API. Experiments were conducted on an Apple M2 machine with 16 GB of RAM. Gurobi was allowed to use up to 8 solver threads.

Unless otherwise stated, no explicit time limit or user-specified optimality gap tolerance was imposed; Gurobi was allowed to terminate upon proven optimality within its default Mixed-Integer Programming (MIP) tolerances.

5.2 Scenario Design and Parameter Variations

This section defines the scenarios considered in the experiments. The main scenario uses risk-aware road costs and serves as the reference case. Additional scenarios are formulated to isolate the effect of the cost definition and to test sensitivity to key parameters.

5.2.1 Main Scenario (Reference End-to-End Pipeline)

The main scenario in this section serves as the reference case for the Kahramanmaraş study and provides the baseline for all subsequent comparisons. It follows the end-to-end pipeline defined in Section 4 under a fixed dataset and a common set of parameter values. The workflow starts by classifying buildings as damaged or undamaged using the post-disaster satellite imagery. These labels are then mapped onto the road network to construct risk-aware edge costs. Using the resulting network, least-cost paths are computed to obtain Supply–POD and POD–Demand cost tables and their normalized accessibility scores. Finally, these inputs are used to solve the POD location and demand assignment model.

In the main scenario, road costs are not based on geometric length alone. Instead, segment costs are adjusted using an exposure measure derived from fixed-radius impact zones (buffers) around buildings classified as damaged. These risk-aware costs are used as edge weights when computing shortest-path costs for both Supply–POD and POD–Demand pairs. The resulting cost tables are then transformed into normalized accessibility scores and used as inputs to the optimization objective. Unless otherwise stated, the main scenario uses the baseline parameter values reported in Section 5.1.3 (i.e., α, r, δ, C, w_{PD}, w_{SP}, and κ).

Results for the main scenario are summarized using the set of opened PODs, the assignment of demand points to the opened sites, and capacity utilization at each opened POD. To support spatial interpretation, map-based figures are also provided, showing the demand clusters, road segments exposed to damage buffers, the selected POD locations, and representative least-cost routes.

5.2.2 Alternative Scenarios

This section introduces alternative scenarios to clarify which parts of the main setup drive the results and how sensitive the method is to key choices. In all scenarios, the study area, the candidate POD set, the supply points, the solver settings, and the overall workflow are kept the same. Unless otherwise noted, the demand points and baseline parameter values from the main scenario are also used. Each scenario changes only one aspect of the setup (for example, the road-cost definition, the way the damage signal is reflected in costs, the capacity assumption, the objective weights, or the demand representation).

For each alternative scenario, the resulting solution is compared against the main scenario. Differences are examined through the opened POD set and the demand–POD assignments, supported by map-based interpretation where relevant. Capacity utilization is also reviewed to see how pressure shifts across opened sites.

Baseline Cost (Length-Only)

In this scenario, road costs are defined only by geometric length. The damage-based penalty used in the main scenario is turned off, and each road segment is weighted directly by its length attribute, $length_m$. As a result, “least-cost” paths in this setting correspond to distance-minimizing routes on the road network. Using these length-based weights, the Supply–POD and POD–Demand cost tables are recomputed. The same normalization procedure is then applied to obtain accessibility scores, and the optimization model is solved again with these inputs. This scenario serves as a clean reference in which damage information does not enter the road-cost layer, making it easier to interpret whether the differences observed in the main scenario are driven by the risk-aware cost definition.

Maximum Open POD Limit (C) Stress Test

In this scenario, the role of the parameter C, which sets an upper bound on the number of PODs that can be opened, is examined. In the main scenario, C is chosen at a relatively high level so that the solution space is not artificially restricted and the model can open as many PODs as required under the capacity constraints.

To examine the effect of this parameter more clearly, additional runs are performed in which the value of C is gradually reduced while all other components are kept fixed. The study area, the demand points and their magnitudes, the candidate POD set, the supply points, the road cost definition, the capacity derivation rule, the objective function weights, and the solver settings remain unchanged. Only the upper limit on the number of PODs that can be opened is modified.

Lowering C forces the model to serve the total demand using a smaller number of distribution sites. This setup allows the dependence of the solution on the availability of many candidate PODs to be examined in isolation, and makes it possible to observe how demand assignments adjust when the number of open PODs is explicitly constrained.

Safety Factor (`SCALAR`) Sensitivity

In this scenario, the effect of the capacity buffer used in deriving POD capacities was examined. The parameter is denoted by κ in the model, and it appears as `SCALAR` in the implementation. In the main scenario, `SCALAR`=1.5 was used. This setting introduces a practical slack to account for uncertainty in area-based capacity estimates and possible measurement or representation errors.

For the sensitivity analysis, `SCALAR` was varied across several levels, and the model was re-solved each time. This allows us to see how the solution responds as the overall capacity becomes tighter or looser relative to total demand. Apart from `SCALAR`, the rest of the setup was kept fixed: the study area, demand points and demand magnitudes, the candidate POD set, the supply points, the road-cost definition, the objective weights, and the solver settings.

Objective Weight Sensitivity (w_{PD}, w_{SP})

In this scenario, the sensitivity of the solution to changes in the objective function weights is examined. The objective function combines two accessibility-related components: a demand-side term that rewards POD–Demand accessibility weighted by demand magnitude, and a supply-side term that favors POD locations that are easier to reach from the supply entry points (see Section 4.6.6). The relative influence of these two components is controlled by the weight parameters w_{PD} and w_{SP}.

As part of the sensitivity analysis, the model is re-solved under different combinations of these weights. In each run, only the values of w_{PD} and w_{SP} are modified. All other elements of the setup are kept fixed, including the study area, the demand points and their magnitudes, the candidate POD set, the risk-aware road cost definition, the capacity derivation rule and buffer factor (κ), the maximum number of PODs that can be opened (C), and the solver settings.

Varying the weights directly affects which accessibility component plays a more dominant role in the optimization. When the relative importance of w_{SP} is increased, the model places greater emphasis on selecting PODs that are more accessible from the supply side. Conversely, when w_{SP} is kept low, demand-side accessibility and the resulting demand–POD assignment patterns become more influential. To assess how this shift in balance shapes the decision structure, the solutions obtained under different weight settings are compared in terms of the opened POD set and the associated demand–POD assignments.

Real Roads vs Euclidean Distance

The aim of this scenario is to isolate the effect of the road network structure and network-based shortest-path calculations on the solution. In the main scenario, accessibility values are derived from least-cost routes on the real road network. Here, the road network is not used. Instead, accessibility is defined using straight-line (Euclidean) distances between point pairs. This provides a simple baseline to assess how much network topology and corridor structure influence the facility-location and assignment decisions.

Accordingly, the POD–Demand costs (c_{jd}) and the Supply–POD costs (c_{sj}) are computed as Euclidean distances for all relevant point pairs. All point layers are represented in a common projected (metric) coordinate system, and distances are therefore measured in meters. No routing, connectivity checks, snapping to network nodes, or edge-based cost components are used in this scenario.

The Euclidean cost tables are then normalized using the same procedure applied in the main scenario and converted into accessibility scores. These scores are used as inputs to the optimization model. Apart from replacing the cost/accessibility tables, the rest of the setup is kept unchanged: the study area, the demand points and demand magnitudes, the candidate POD set, the capacity derivation rule and buffer factor (κ), the maximum number of PODs that can be opened (C), the objective-function weights, and the solver settings. This way, any differences in the results can be attributed to the use of Euclidean distances in place of network-based costs.

Building-Stock-Based Demand Definition

This scenario tests whether the optimization results depend on how demand is represented. In the main scenario, demand points come directly from the spatial pattern of buildings classified as damaged. The underlying assumption is that, in the early post-disaster phase, demand is concentrated in areas with higher levels of structural damage.

In this alternative, we form demand locations based on the overall settlement pattern. We take centroids of all building footprints in the Kahramanmaraş study area and cluster them using k-means with 43 clusters to remain consistent with the main scenario. After clustering, each group is represented by a single demand point located at the average position of its member building centroids (mean-x, mean-y). With this construction, demand locations follow the overall building distribution rather than the observed damage concentration.

Demand magnitude in this scenario is defined as the total number of buildings in the cluster (count). This treats building density as a practical proxy for potential need rather than tying demand magnitude directly to the damage labels.

The road-cost structure is kept unchanged. Supply–POD and POD–Demand cost tables (`cost_sp.csv` and `cost_pd.csv`) are recomputed on the same risk-aware road network used in the main scenario (see Section 4.4.2 and Section 3.1). Therefore, any differences in the solution can be attributed to the demand locations and demand magnitudes that enter the optimization model. Demand magnitude affects the model through the parameter d_i in both the objective function and the capacity constraints (see Section 4.6).

5.3 Evaluation Metrics

Experimental results were assessed using two types of information: the decisions produced by the model and summary indicators derived from these decisions. The key decision outputs are the set of opened PODs and the assignment of each demand point to one of these opened sites. Both outputs are presented in tabular form and supported with map-based figures.

To describe the solution structure, we report the number of opened PODs and the total demand assigned to each opened site. Assigned demand is presented alongside the corresponding capacity, which allows capacity utilization (load relative to capacity) to be examined and discussed.

For scenario comparisons, we track how much the opened POD sets overlap and what fraction of demand points are assigned to a different POD across scenarios. Together, these indicators show whether changes in the cost definition or parameter values mainly affect the location decisions or the demand assignments.

5.4 Results

This section reports the results for the Kahramanmaraş case. Results under the main scenario are presented first, supported by map-based outputs. Comparisons with alternative scenarios are then reported.

5.4.1 Main Scenario Results

This section reports the results of the reference (main) scenario for the Kahramanmaraş case. To keep the narrative consistent with the pipeline in Chapter 4, results are presented in the same order. We first show the damage detection outputs used to define the damaged-building layer. We then summarize the demand points obtained from clustering and illustrate how damage buffers translate into road exposure. Finally, we report the optimization results in terms of the opened POD set and the resulting demand assignments. Map-based figures are used throughout to support spatial interpretation and to connect intermediate layers to the final location–allocation decisions.

Damage Detection Output and Example Predictions

Figure 4.5 shows the building-level damage labels obtained after applying the classifier to the study area and mapping predictions back onto the building footprints in QGIS. This damaged-building layer is used in two later steps. It provides the input points that are aggregated into demand locations, and it defines the source geometry for the buffer-based exposure layer used in risk-aware road costing.

To give a qualitative sense of the model output, Figure 5.1 presents representative building patches together with their predicted classes (damaged versus undamaged). The examples are included for illustration and help connect the image-based inference stage to the spatial layers used in the network analysis and optimization model.

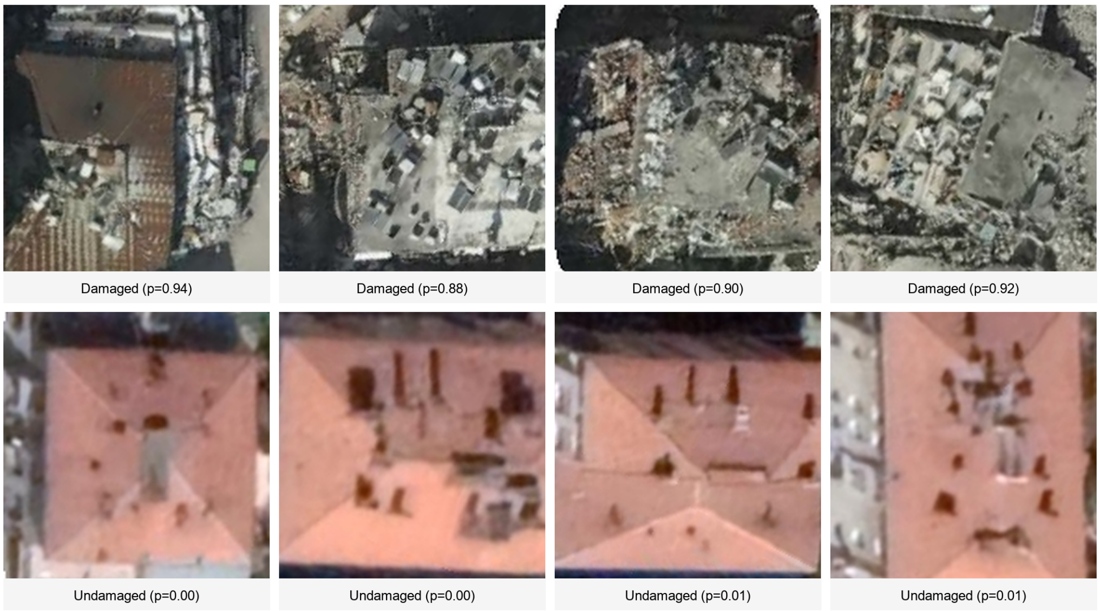

Figure 5.1: Representative building patches classified as damaged and undamaged by the damage detection model. Labels indicate the predicted class and the corresponding model confidence. The examples are shown for qualitative illustration.

Demand Points and Clustering Output

This subsection describes how demand points are constructed from the damage assessment outputs and how they are distributed across the study area in the main scenario. Instead of treating every damaged building as a separate demand location, the damaged-building layer is summarized into a smaller set of representative points. This keeps the network calculations and the optimization model manageable, while preserving the main spatial structure of the observed damage.

Buildings classified as damaged were first clustered using their planar coordinates. The initial k-means step captures the dominant concentrations of damage, but it can also produce very small or isolated clusters. To avoid a fragmented representation, a second-stage post-processing step was applied in which clusters below a minimum size were merged into the nearest neighboring clusters based on spatial proximity. This refinement yields more stable demand zones and reduces sensitivity to minor spatial irregularities in the damaged-building layer.

The two-stage aggregation resulted in a final set of 43 demand points for the study area. Each demand point is placed at the cluster centroid, computed as the mean of the coordinates of the damaged buildings assigned to that cluster. Demand magnitude is defined as the number of damaged buildings in the cluster. In this way, the demand representation captures both where damage is concentrated and how large each demand zone is.

The resulting demand points follow the main spatial concentration patterns observed in the damaged-building layer. These points are then used directly in the accessibility computations and in the POD location and demand assignment decisions in the subsequent stages.

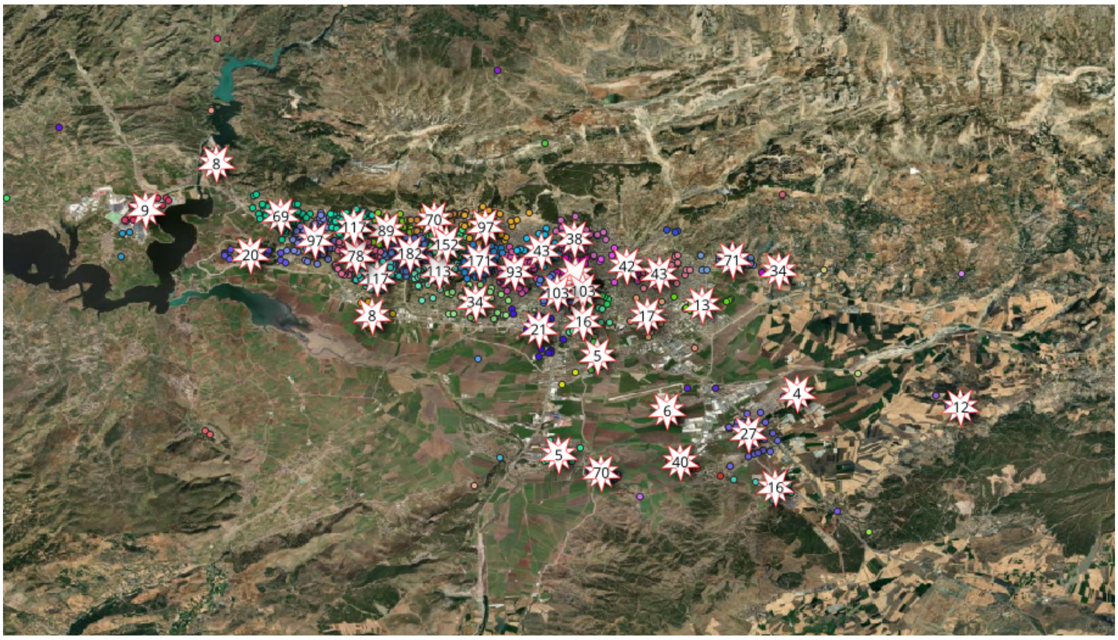

Figure 5.2: Demand points derived from clustered damaged buildings, with labels indicating the number of damaged buildings assigned to each demand point.

Damage Buffers and Road Exposure

This subsection shows how buffer zones defined around damaged structures translate into road exposure and feed into the risk-aware cost definition. Building footprints classified as damaged were converted into impact zones using a fixed-radius buffering step (see Section 4.4.2). The resulting layer provides a simple and consistent proxy for local disruption effects that are hard to observe directly, such as debris, congestion, and short-range access restrictions.

The buffer layer was then overlaid with the road segments. For each segment, the number of distinct impact zones it intersects was computed and recorded as the exposure value $damage_n$. Exposure is measured as the number of buffers a road segment intersects, rather than how long the segment stays inside those buffers. This choice aligns with the binary nature of the damage labels (damaged vs. undamaged) and avoids giving a false sense of accuracy when the imagery and building footprints do not align perfectly, for example, due to parallax. Figure 5.3 illustrates the resulting exposure levels across the study area, where road segments are symbolized according to these computed $damage_n$ values.

Figure 4.6 shows the buffer zones together with the road segments that intersect them. In some places, a segment cuts across several buffers, while in others it stays completely outside the impact areas. Segments with larger $damage_n$ values are assigned higher penalties, which can increase their contribution to total path cost when least-cost paths are computed. In practical terms, routes that traverse areas with dense nearby damage may become more costly under the defined weighting scheme, which means that accessibility measures are not based solely on geometric length.

The exposure layer produced in this step is the direct input to the risk-aware edge costs used in the subsequent analysis. As a result, both Supply–POD and POD–Demand shortest-path costs incorporate damage information through the buffer-based exposure measure.

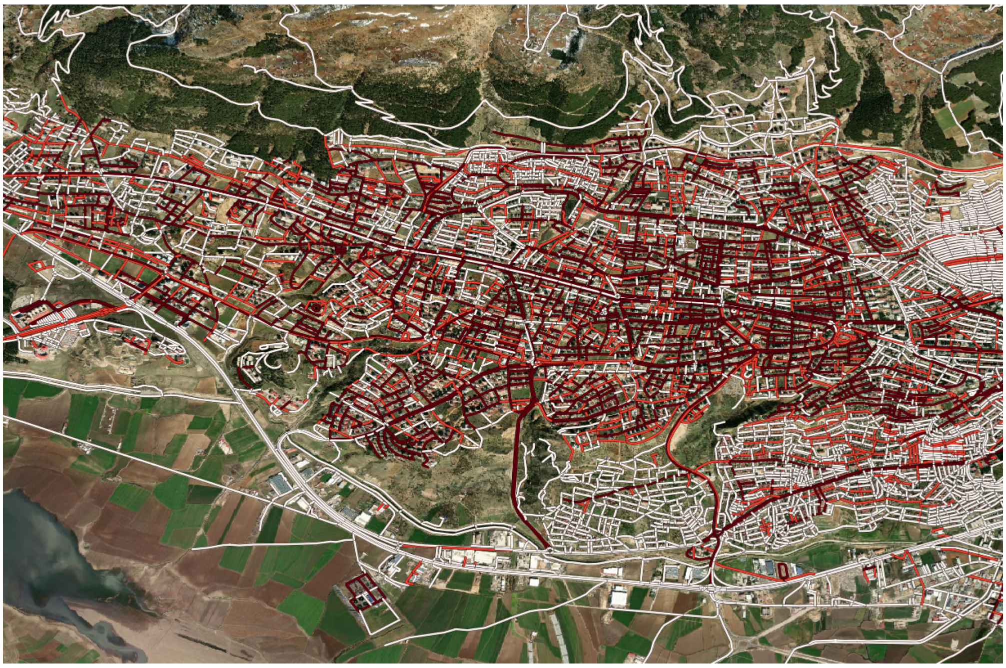

Figure 5.3: Road segments symbolized by their exposure level (*damage$_n$*), defined as the number of damaged-building buffers intersecting each segment. Darker segments indicate higher exposure to nearby damage.

Opened POD Set

This subsection summarizes which PODs are opened by the model under the main scenario and how these sites are distributed across the study area. The solution is obtained under the capacity restrictions and the accessibility inputs computed on the risk-aware road costs. In practice, the model does not open PODs up to the limit C by default; the final number of opened sites follows from the balance between serving demand efficiently and keeping assignments feasible under capacity.

Figure 5.4 shows the opened 23 POD locations together with the demand points. Visually, the selected sites fall close to the main demand concentrations, and there are also locations that sit away from road corridors with higher exposure.

The capacity dimension also matters for the selection. To keep the assignment feasible for all demand points, the model tends to prefer candidates with larger site areas (and hence higher capacity) over small sites with limited capacity. Table 5.1 reports the number of opened PODs, the total demand assigned to each opened site, and the resulting capacity utilization levels.

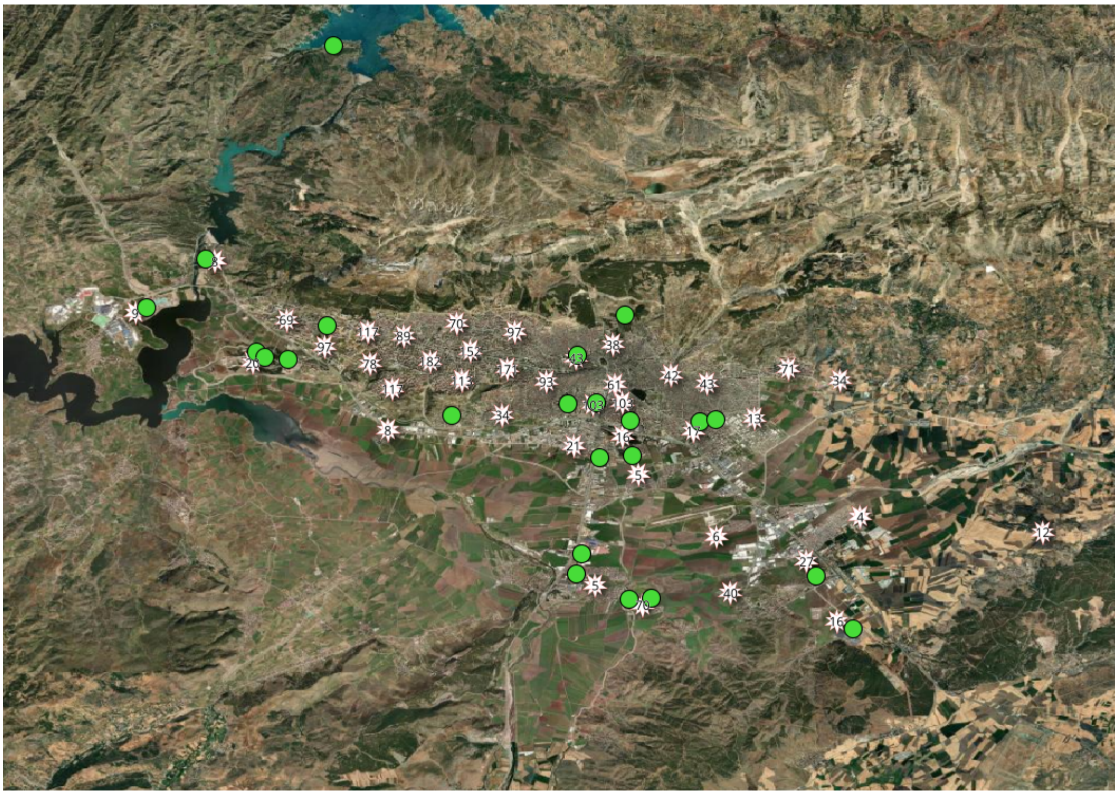

Figure 5.4: Opened POD locations and demand point distribution under the main scenario. Opened PODs are shown as circles, while demand points correspond to cluster centers derived from damaged buildings.

Table 5.1: Capacity utilization of opened PODs under the main scenario.

POD ID	Capacity	Assigned demand	Utilization
5	53.3	43	0.81
31	14.8	13	0.88
32	128.5	127	0.99
39	51.4	43	0.84
61	50.4	48	0.95
67	22.9	21	0.92
164	21.7	16	0.74
213	12.8	9	0.71
260	48.8	40	0.82
262	275.9	272	0.99
307	73.2	69	0.94
317	17.3	16	0.93
374	15.1	13	0.86
376	12.8	12	0.93
378	12.3	8	0.65
379	11.1	5	0.45
386	92.0	88	0.96
389	5.0	5	0.99
411	37.4	34	0.91
413	6.0	4	0.67
487	1395.8	1395	1.00
493	18.2	8	0.44
529	83.9	78	0.93

Demand–POD Assignments

This subsection describes how demand points are assigned to the available PODs under the main scenario. Each demand point is assigned to exactly one POD by the model. Assignments are determined jointly by the accessibility values computed on the risk-aware road network and by the capacity constraints of the candidate sites.

Demand–POD assignments are therefore not driven by geometric proximity alone. The exposure-based penalties embedded in the road costs influence the effective accessibility between locations. As a result, some demand points are assigned to PODs that are not the closest in straight-line distance, but are reached through routes with lower overall risk.

Figure 5.5 illustrates the spatial pattern of the demand–POD assignments for the opened sites. The map shows that many demand clusters are served by a limited subset of PODs, while some sites cover larger service areas due to the combined effect of accessibility and capacity constraints.

The assignment outcomes also display a coherent pattern in terms of capacity usage. PODs located in high-demand areas naturally end up with higher utilization, since they attract and absorb nearby demand clusters. More peripheral PODs, on the other hand, take on part of the load and help prevent the central sites from becoming overly crowded. Overall, the solution spreads demand across multiple sites instead of concentrating it at a single location, which results in a more even service pattern over the network.

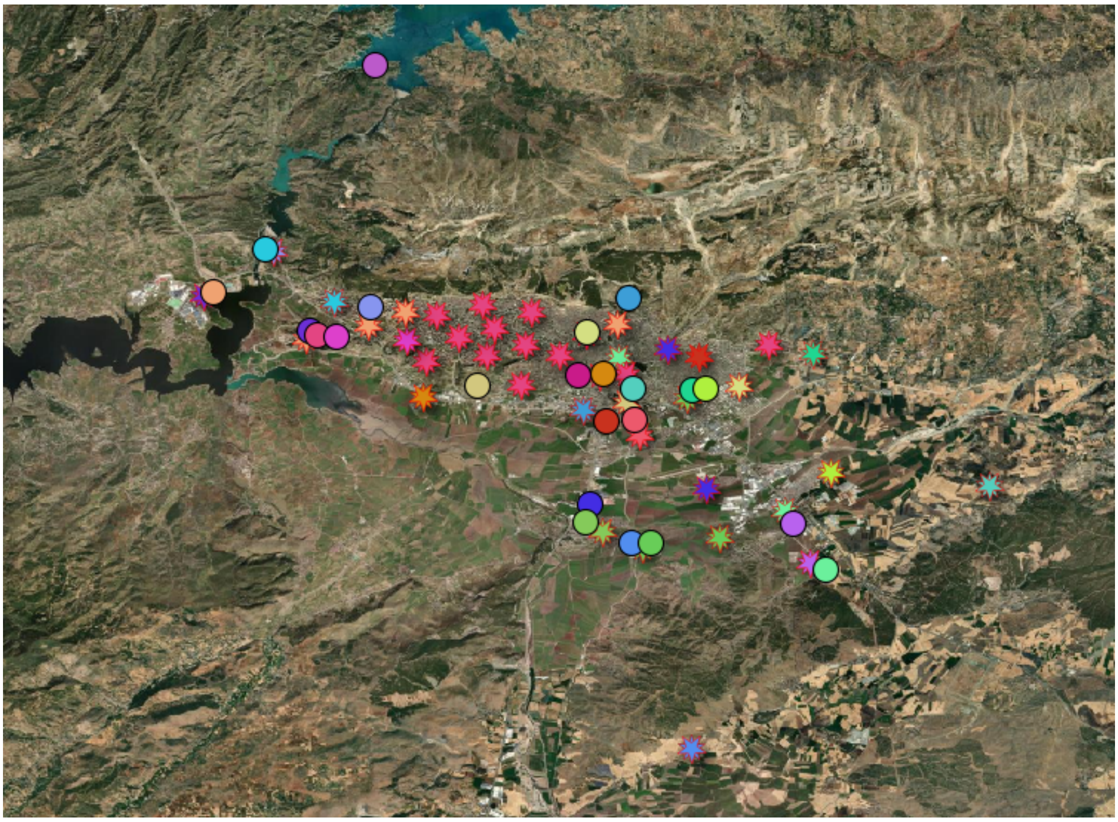

Figure 5.5: Spatial distribution of demand–POD assignments under the main scenario. Colored circles indicate opened POD locations, while star symbols represent demand points. Demand points sharing the same color are assigned to the same POD, illustrating how demand clusters are served by the selected sites.

Shortest Paths Under Risk-Aware Costs

This subsection visualizes the least-cost routes computed on the risk-aware road network used in the main scenario. Rather than reporting only the numerical cost tables, it shows the routes on a map over which these costs are obtained. This makes the link between the risk-aware edge-cost definition on the road network and the accessibility inputs used by the optimization model more tangible.

In the main scenario, least-cost paths are computed with Dijkstra's algorithm using the risk-aware edge weights constructed in Section 4.4. Each edge cost starts from the segment length ($length_m$) and is then scaled up when the segment is exposed to nearby damaged-building buffers. For that reason, “shortest path” in this section means the path with the lowest total weighted cost on this risk-aware network, which may differ from the purely length-minimizing route.

To illustrate the routing outputs, representative least-cost paths are mapped for selected POD→Demand pairs appearing in the main-scenario solution. These examples indicate which road segments contribute to the POD–Demand cost table (`cost_pd.csv`) used by the optimization model. They are included to provide a qualitative, spatial view of the cost calculations, not to summarize the full network behavior.

Figure 5.6 presents these representative POD→Demand routes. Segment costs are defined through intersections with damage buffers, consistent with the exposure-based cost assignment in Section 4.3. The figure therefore serves as a visual check that the accessibility inputs entering the optimization model are derived from the same risk-aware cost space defined on the road network.

The purpose of this subsection is to document how the main-scenario road costs translate into the least-cost routes reflected in `cost_pd.csv`. At this point, the main-scenario results are complete; the remainder of this chapter turns to alternative scenarios and comparative findings.

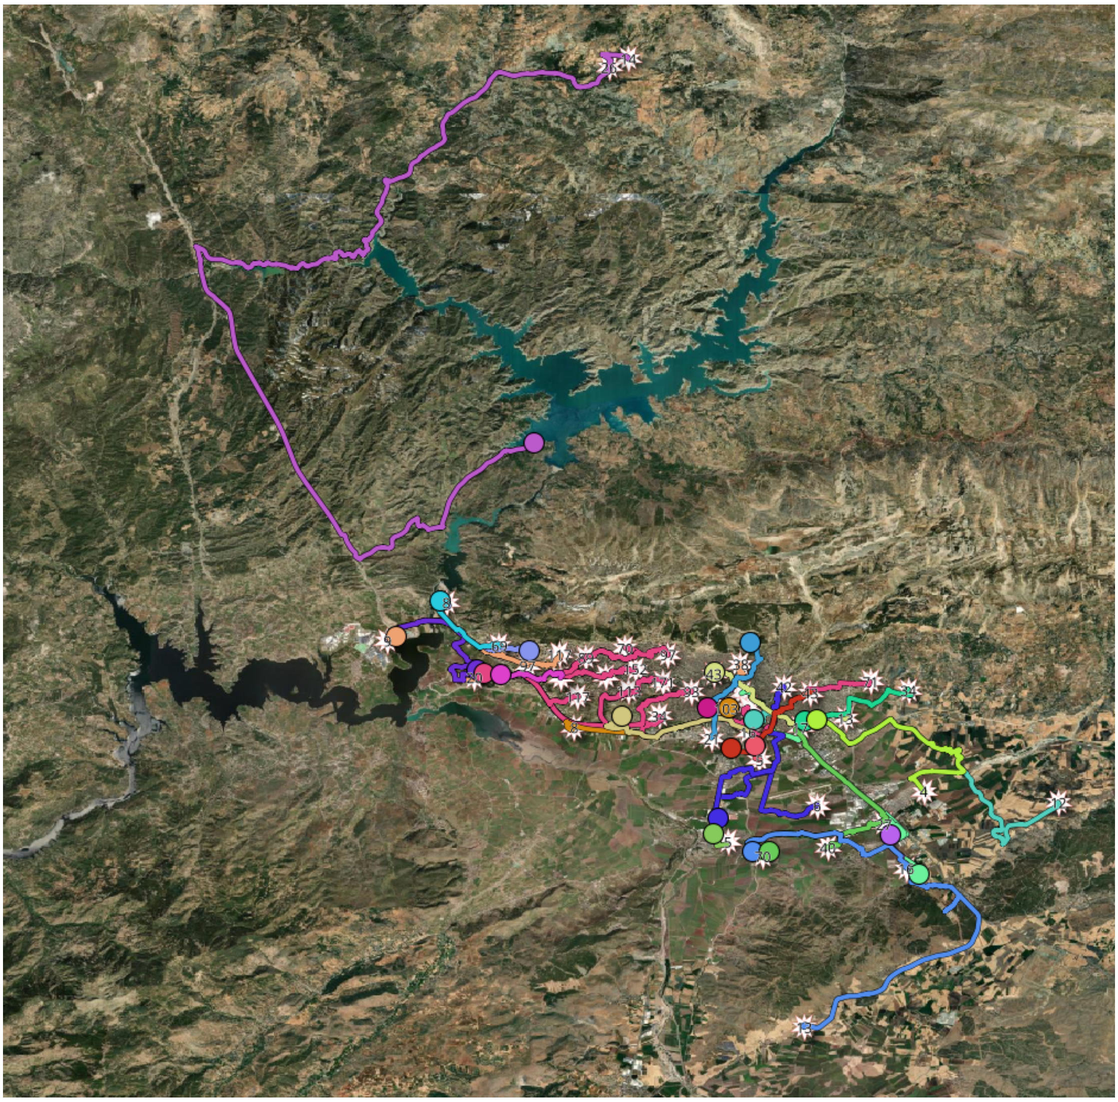

Figure 5.6: Representative least-cost paths between selected POD locations and demand points, computed on the road network weighted by the risk-aware edge-cost field (buffer-based damage penalties). The figure is provided to illustrate the routing outputs used to build the POD→Demand cost table.

5.4.2 Comparative Results Across Scenarios

In this section, results from the alternative scenarios are reported through direct comparisons with the main scenario. The goal is to understand what changes in the solution when only one modeling element is modified (for example, the road-cost definition, the capacity setting, or the objective weights) and to see which parts of the decision structure remain largely unchanged.

The comparison focuses on two aspects. First, changes in the opened POD set are examined, together with their spatial footprint in the study area. Second, we look at how demand–POD assignments shift across scenarios and how these shifts translate into capacity utilization at the opened sites.

Risk-Aware vs Length-Only

In this comparison, only the road cost definition is changed. In the risk-aware scenario, road segments are weighted by an exposure penalty derived from buffers around buildings classified as damaged. In the length-only scenario, this penalty is removed and road costs are defined solely by geometric length. All other elements are held fixed, including the demand points and their magnitudes, the candidate POD set, the supply points, the capacity derivation rule and capacity buffer, the objective weights, and the solver settings.

Both scenarios open 23 PODs each. The two solutions share 22 of these sites, while one site differs between them. Specifically, candidate 378 appears in the opened set under risk-aware costs, whereas candidate 34 appears in the opened set under length-only costs.

Demand assignment patterns follow a similar structure. Out of 43 demand points, 41 are assigned to the same POD in each scenario. The remaining two demand points (demand_id: 28 and 52) are assigned to different PODs across the scenarios.

Although the overall structural similarity between the two solutions is high, the differences indicate that exposure-based penalties influence routing and allocation decisions at the margin. The relatively small number of changed assignments suggests that the road network in the study area remains largely functional, but localized cost adjustments can still shift specific demand–POD matches.

Sensitivity to the Maximum Number of Open PODs (C)

This subsection reports the effect of the parameter C, which specifies the maximum number of PODs that can be opened, on the solution outputs. In all runs, the study area, demand points and demand magnitudes, the candidate POD set, supply points, road cost definition, capacity derivation rule, objective function weights, and solver settings are kept fixed. Only the value of C is varied.

Figure 5.7 illustrates how the number of opened PODs in the solution changes with C. For small values of C, the number of opened sites is directly constrained by this upper bound. As C increases, the number of opened PODs rises and then stabilizes after a certain point. Beyond this point, further increases in C do not lead to changes in the number of opened PODs.

The demand-weighted average POD–Demand accessibility is shown in Figure 5.8a. At low values of C, accessibility remains at relatively low levels. As C increases, a clear improvement in accessibility is observed. Once the number of opened PODs becomes stable, the average accessibility value also remains largely unchanged.

A similar pattern is observed in the total objective function values. Figure 5.8b shows that the objective value is lower for small values of C, increases as C grows, and reaches a plateau after a certain threshold. Beyond this threshold, further increases in C do not result in noticeable changes in the total objective function value.

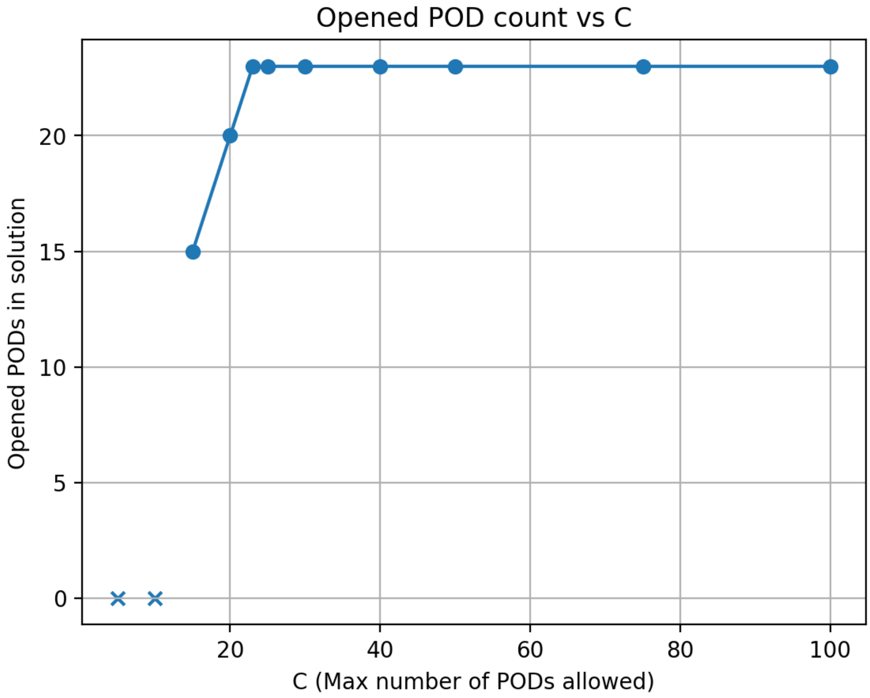

Figure 5.7: Number of opened PODs as a function of the maximum number of sites (C).

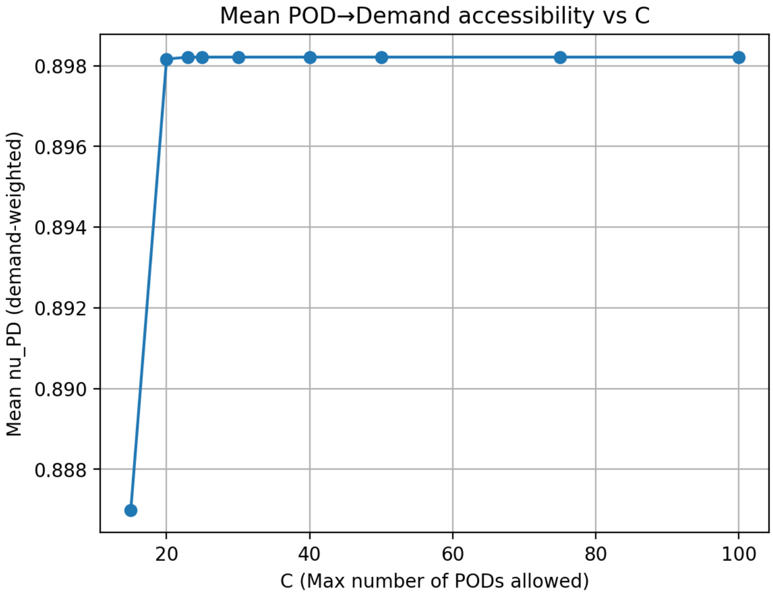

(a) Demand-weighted mean POD–Demand accessibility

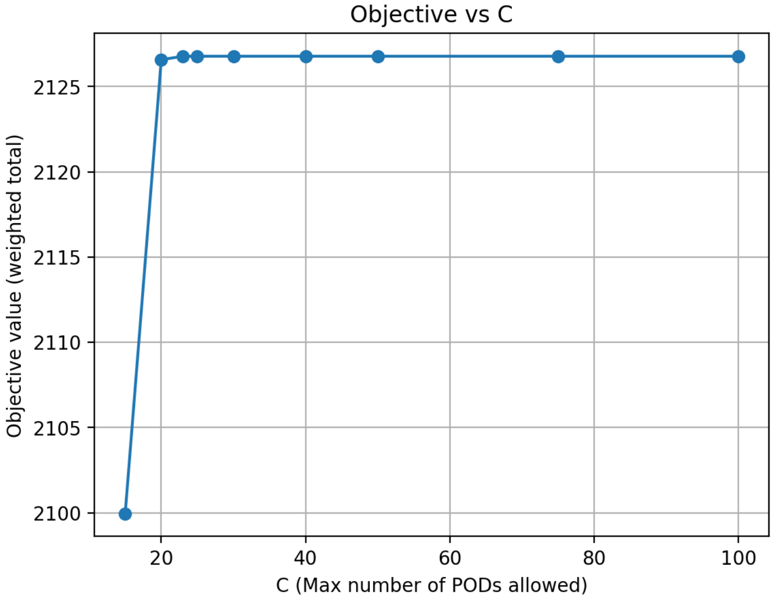

(b) Total objective function value

Figure 5.8: Effect of the maximum number of PODs (C) on demand-weighted mean POD–Demand accessibility and the total objective function value.

Sensitivity to the Capacity Buffer (`SCALAR`)

In this subsection, we compare the solutions obtained under different values of the capacity buffer parameter (`SCALAR`, denoted by κ) used to derive POD capacities. In the main scenario, $\kappa = 1.5$ was used, and the model opened 23 PODs.

When κ was reduced below 1.5, the model did not yield a feasible solution. For `SCALAR` values of 0.8, 1.0, and 1.2, Gurobi reported the problem as *infeasible*. For this reason, the numerical comparisons focus on the range where feasible solutions are obtained (i.e., $\kappa \geq 1.5$).

Within the feasible region, the demand-weighted average availability, $\bar{v}_{PD}$, increases as κ increases (Figure 5.10a). The total objective function value also increases with κ (Figure 5.10b). In terms of facility decisions, the number of opened PODs decreases from 23 in the main scenario to 20 when κ reaches 1.8, and then remains unchanged for larger values of κ (Figure 5.9).

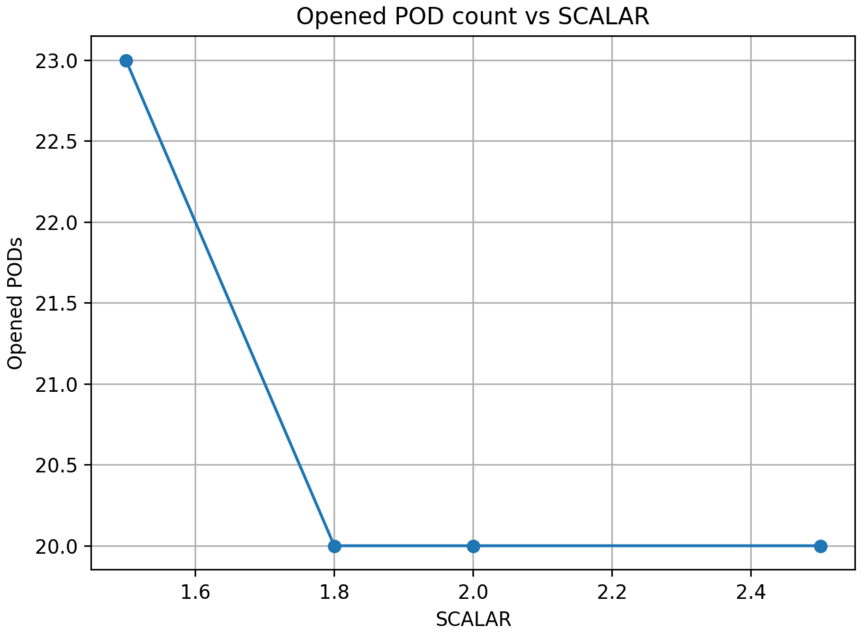

Figure 5.9: Number of opened PODs as a function of the capacity buffer parameter `SCALAR` (κ).

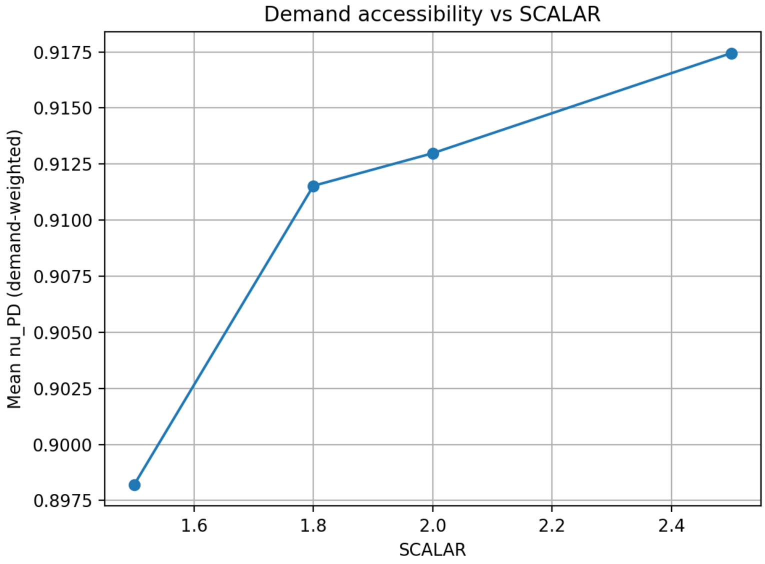

(a) Demand-weighted mean POD–Demand accessibility

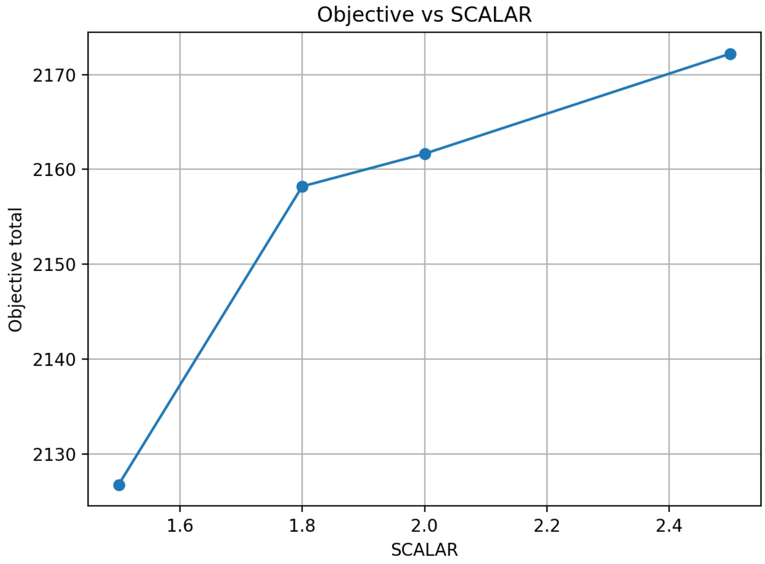

(b) Total objective function value

Figure 5.10: Effect of the capacity buffer parameter `SCALAR` (κ) on demand-weighted mean POD–Demand accessibility and the total objective function value.

Sensitivity to Objective Weights

This subsection examines how the objective-function weights (w_{PD} and w_{SP}) affect the solution. The analysis is based on a two-dimensional logarithmic sweep over (w_{PD}, w_{SP}), complemented by one-dimensional runs in which one weight is varied while the remaining settings are held fixed (Figure 5.12).

On the demand side, the outcomes are largely stable. The demand-weighted mean POD–Demand accessibility, $\bar{v}_{PD}$, varies only marginally across most of the tested weight combinations (Figure 5.11b). When w_{PD} clearly dominates, or when w_{PD} and w_{SP} remain within the same order of magnitude, $\bar{v}_{PD}$ stays essentially unchanged (Figure 5.11b). As the ratio w_{SP}/w_{PD} increases, $\bar{v}_{PD}$ declines, but the change is gradual and small in size (Figure 5.11b).

The POD-opening decisions show a clearer response. The number of opened PODs increases as more relative weight is placed on the supply-side term (Figure 5.11a). Under demand-focused settings, the model produces an opened-site structure close to the main scenario, whereas higher emphasis on the supply-side component leads to additional sites being opened (Figure 5.11a). The shift in the opened-site count is not smooth; it occurs through step changes and plateaus, which is consistent with binary opening variables in a mixed-integer formulation (Figure 5.11a).

Overall, the results point to a limited trade-off: demand-side accessibility is largely preserved, while changes in objective priorities are reflected mainly in the facility-opening decisions (Figures 5.11b and 5.11a). The two-dimensional sweep supports the same interpretation, with higher opened POD counts appearing in regions where w_{SP} is relatively larger (Figure 5.12).

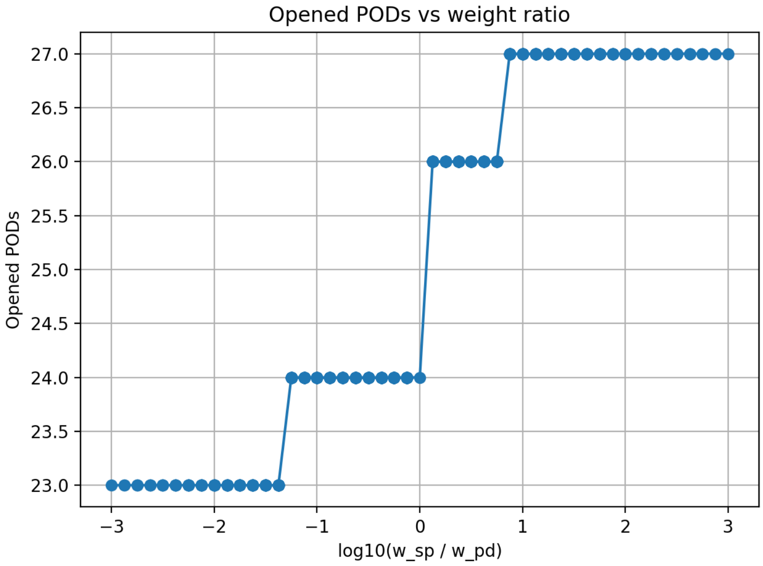

(a) Number of opened PODs vs w_{SP}/w_{PD}

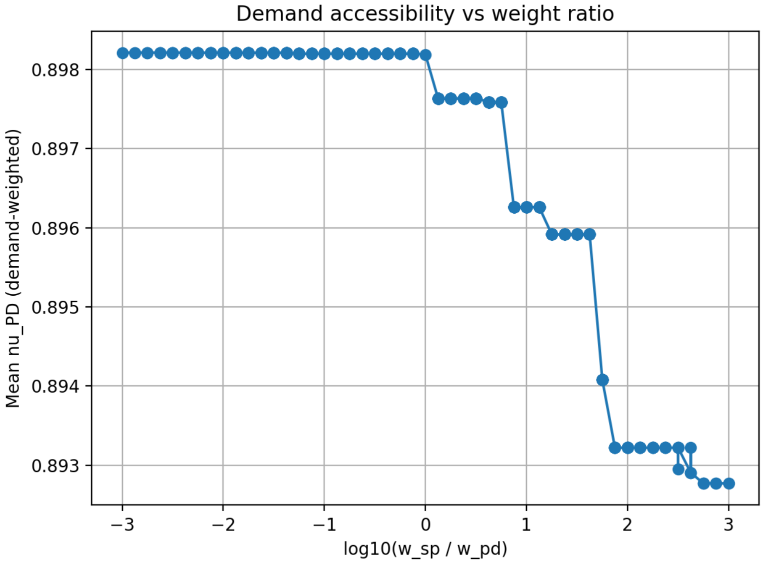

(b) Demand-weighted mean POD–Demand accessibility ($\bar{\nu}_{PD}$)

Figure 5.11: Sensitivity to the weight ratio w_{SP}/w_{PD}. Panel (a) shows the number of opened PODs, and panel (b) shows the demand-weighted mean POD–Demand accessibility $\bar{\nu}_{PD}$. The horizontal axis is reported on a logarithmic scale.

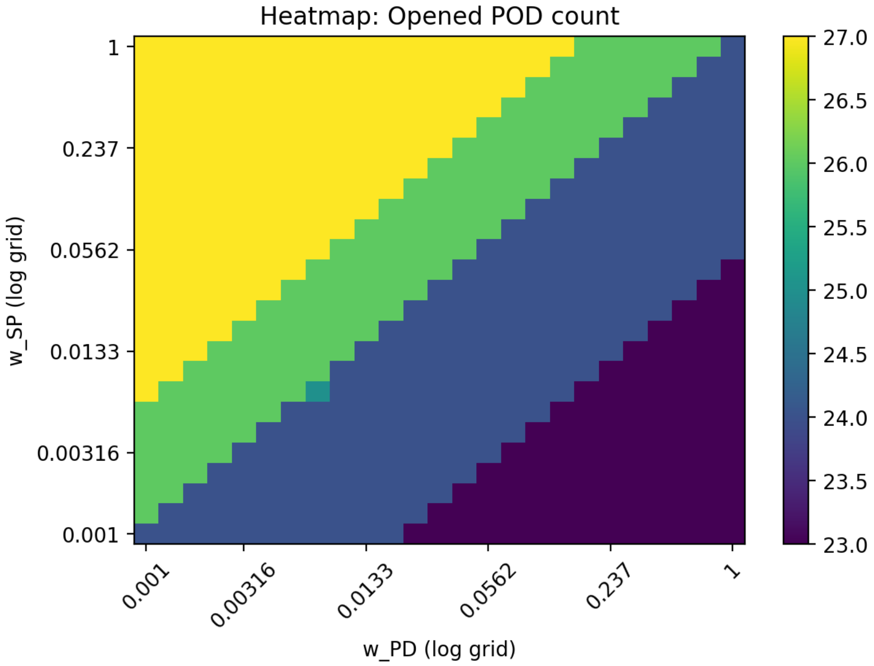

Figure 5.12: Number of opened PODs across the two-dimensional (w_{PD}, w_{SP}) weight grid used in the logarithmic scan. Higher counts appear in regions where the supply-side weight is relatively larger.

Real Roads vs Euclidean Distance

In this subsection, solutions obtained using least-cost route costs on the real road network are compared with solutions obtained under Euclidean distance–based costs. In this comparison, only the `cost_pd.csv` and `cost_sp.csv` tables were regenerated using Euclidean distances. All other input tables, parameter values, and solver settings were kept the same as in the main scenario.

In the Euclidean distance–based scenario, the model opened 22 PODs, whereas the main scenario opened 23. The opened POD sets do not fully overlap: the two solutions share 19 opened sites, and the remaining sites differ between scenarios (4 sites appear only in the main scenario and 3 sites appear only in the Euclidean scenario).

Demand assignments also change noticeably. Out of 43 demand points, 16 are assigned to different PODs compared to the main scenario. Hence, in this scenario, the differences are not limited to the POD-opening decisions but also appear in the demand–POD matching.

This larger divergence is consistent with the fact that Euclidean distance ignores road-network topology and therefore changes the effective accessibility structure more fundamentally than modifying edge weights on the same network.

Table 5.2 reports the derived capacities of the PODs opened under Euclidean distance–based costs, together with the total assigned demand and the corresponding capacity utilization levels. The utilization values fall within a comparable range to those reported for the main scenario: some sites operate close to their capacity, while others remain at lower utilization levels.

Table 5.2: Capacity utilization of opened PODs under the Euclidean distance–based scenario. For each opened site, the table reports the derived capacity, the total assigned demand, and the resulting utilization level.

POD ID	Capacity	Assigned demand	Utilization
5	53.0	42	0.79
31	14.7	13	0.89
32	127.6	127	0.99
34	18.9	16	0.85
39	51.1	43	0.84
61	50.0	48	0.96
67	22.7	21	0.92
164	21.5	20	0.93
189	24.9	20	0.80
213	12.7	8	0.63
260	48.4	40	0.83
262	274.0	272	0.99
307	72.7	70	0.96
317	17.1	8	0.47
374	15.0	13	0.87
376	12.7	12	0.94
379	11.0	6	0.54
386	91.4	88	0.96
411	37.1	34	0.92
420	5.5	5	0.92
487	1386.0	1383	1.00
529	83.3	78	0.94

Building-Stock-Based Demand Representation Results

In this scenario, the model was solved using the settlement-based demand representation. A total of 24 PODs were opened in the solution (upper bound $C = 100$).

Demand points were initially generated from 43 clusters. However, for one demand point, no feasible route could be established on the road network. Since this cluster contained only 6 buildings, it was excluded from the analysis. As a result, the optimization model is reported based on $|I| = 42$ demand points. The total demand entering the model in this scenario equals 4460 (building count).

Figure 5.13 presents the spatial distribution of the settlement-based demand points, while Figure 5.14 shows the locations of the opened PODs. Capacity values, assigned demand, and utilization levels for the opened PODs are reported in Table 5.3.

For this scenario, feasibility was achieved by setting the capacity scaling parameter to $\kappa = 2.0$ (`SCALAR`=2.0). All other components of the framework, including the risk-aware road cost definition, the generation of `cost_sp.csv` and `cost_pd.csv`, and the solver settings, were kept the same as in the main scenario.

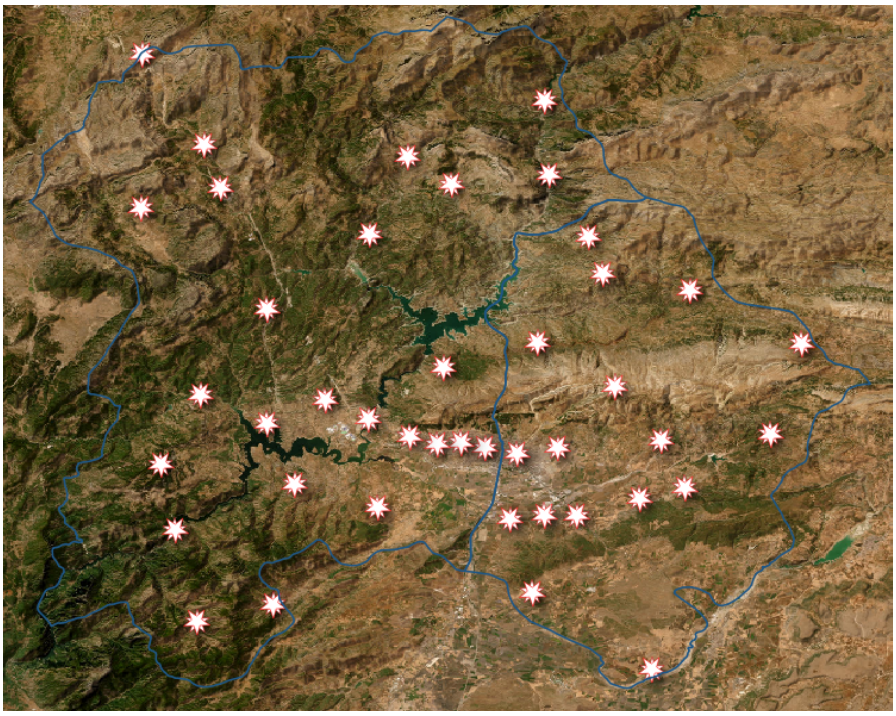

Figure 5.13: Spatial distribution of demand points under the settlement-based demand definition. Star symbols represent demand locations obtained from clustering all building centroids in the study area (cluster centers). The blue boundary outlines the study area covering the Onikişubat and Dulkadiroğlu districts.

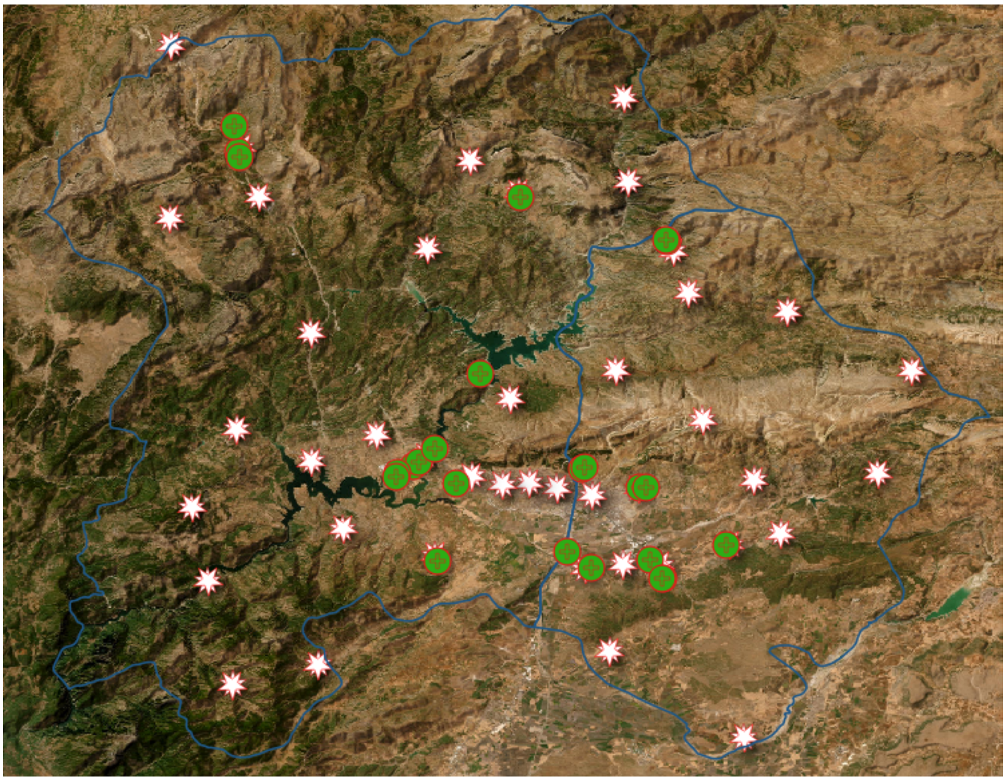

Figure 5.14: Spatial distribution of the opened PODs and settlement-based demand points. Green circles indicate the PODs selected by the optimization model, while star markers represent demand locations derived from clustering all building centroids in the study area. The blue boundary outlines the study area covering the Onikişubat and Dulkadiroğlu districts.

Table 5.3: Spatial distribution of the opened PODs and settlement-based demand points. Green symbols indicate the PODs selected by the optimization model, while star markers represent demand locations derived from clustering all building centroids. The blue boundary outlines the study area.

POD ID	Capacity	Assigned demand	Utilization
10	6.2	6	0.97
11	11.3	11	0.97
32	357.3	326	0.91
61	139.9	115	0.82
67	63.7	1	0.02
68	17.3	5	0.29
164	60.2	52	0.86
205	63.3	10	0.16
208	3.5	2	0.57
209	2.7	2	0.74
260	135.6	25	0.18
262	766.9	244	0.32
286	2.7	1	0.37
307	203.4	66	0.32
357	1.3	1	0.77
373	3.1	2	0.65
386	255.8	193	0.75
404	27.0	1	0.04
436	9.3	1	0.11
471	6.4	4	0.62
487	3879.4	3840	0.99
495	12.0	10	0.84
550	13.1	11	0.84
559	5.6	5	0.90

5.5 Discussion

5.5.1 Summary of Key Findings

This study proposes an end-to-end framework that uses building damage information from satellite imagery to adjust road network costs and connect these costs to accessibility calculations and POD location–allocation decisions. The Kahramanmaraş case study shows that the workflow is technically implementable and that the damage-informed cost structure can affect both site selection and demand assignment.

In the main scenario, the model opened 23 PODs under risk-aware road costs and assigned all 43 demand points while satisfying capacity constraints. When risk-aware costs were replaced with length-only costs, the overall solution pattern remained similar. In contrast, replacing network-based shortest-path costs with Euclidean distances led to larger changes in both the opened POD set and the demand–POD assignments. Sensitivity analyses on C, κ, and the objective weights further illustrate how the solution responds to limits on the number of open sites, capacity scaling, and the balance between demand-side and supply-side accessibility. Finally, the demand representation also matters: changing how demand points are defined can affect both the number of opened PODs and feasibility.

5.5.2 Interpretation of the Main Scenario Results

In the main scenario, the model opens 23 PODs even though the upper limit is $C = 100$. This suggests that the solution is not driven by the upper bound on the number of PODs. Instead, the number of opened sites is determined by the balance between capacity feasibility and the accessibility structure induced by the network costs.

Spatially, the opened PODs concentrate around the main demand clusters, while a smaller set of more peripheral sites helps absorb part of the demand and reduces pressure on central locations. The demand–POD assignments also cannot be explained by geometric proximity alone. Some demand points are not matched to the closest POD in straight-line distance. Instead, they are sent to sites that can be reached through routes that stay away from highly exposed road segments. This suggests that the buffer-based penalty is reflected in the least-cost paths and, as a result, affects the final assignment pattern in localized areas.

A final map-based note is that, for a subset of the selected POD locations, zooming in on satellite imagery reveals tent-like temporary installations. This observation should be interpreted cautiously, since imagery dates and visual interpretation introduce uncertainty. Still, it provides a small qualitative consistency check: some selected candidates appear compatible with locations that may have been used during the response period.

Figure 5.15: Satellite imagery examples of selected open POD locations. When zoomed in, tent-like temporary structures are visible in several sites.

5.5.3 Propagation of Damage-Detection Uncertainty Through the Pipeline

Damage detection is the entry point of the proposed workflow. Minor misalignment between the satellite imagery and the building footprints, especially due to parallax, can cause some intact buildings to be classified as damaged. In our qualitative checks, these mistakes were mostly false positives, so the resulting damage layer can be conservative in parts of the study area.

This uncertainty appears in two places. Demand points are derived from the spatial pattern of buildings labeled as damaged, so small classification errors can slightly change cluster sizes and, in some cases, cluster centers. A more noticeable effect occurs in the road-cost layer. Buffers drawn around predicted damaged buildings can intersect additional road segments, which increases the exposure count $damage_n$ and raises segment costs. In areas where these intersections are frequent, accessibility may therefore be estimated somewhat lower than it would be with perfectly accurate labels.

The model nevertheless treats exposure as a penalty rather than a hard disruption. Road segments are not removed from the network; they simply become more expensive when they are exposed. The strength of this adjustment is controlled by α. In this study, α should be understood as a risk-sensitivity parameter, not as a physically calibrated travel-time delay. This choice keeps the network usable while still discouraging routes that pass through highly exposed areas.

5.5.4 Role of the Road Network: Risk-Aware Costs vs Length-Only vs Euclidean

Risk-aware vs Length-only (Same network, different edge weights)

When risk-aware road costs are compared with the length-only baseline, the opened POD sets are almost the same: 22 of the 23 sites coincide across the two solutions. The demand assignments are similarly stable. Out of 43 demand points, only two are assigned to a different POD.

This pattern suggests that, in this study area, the road network remains broadly usable and the exposure-based penalties do not overturn the main location–allocation structure. The differences that do appear are localized. They tend to occur in parts of the network where exposure is higher, where even a modest penalty can shift least-cost routes and lead to small changes in assignments.

Network vs Euclidean (Topology matters)

Larger differences appear when the network-based cost setting is compared with the Euclidean-distance baseline. Under Euclidean distances, the model opens a different set of PODs, and 16 of the 43 demand points are assigned to a different site than in the main scenario. This shift highlights how much the road network itself matters for the location–allocation decision.

Euclidean distance treats access as if movement were equally easy in every direction [App+08]. In practice, travel is guided and constrained by the road network [App+08]. Roads form corridors, connections are not equally available everywhere, and detours can be substantial. For this reason, in this case study the presence and structure of the road network has a stronger influence on the solution than simply switching the damage-based penalty on or off.

5.5.5 Sensitivity Results and What They Mean for Planning

Maximum number of PODs (C)

When C is set to low values, it becomes a binding constraint. In this range, the model is forced to operate with too few open PODs, which lowers average accessibility and leads to smaller objective values. As C increases, the number of opened PODs rises and accessibility improves. After a certain point, both measures level off and further increases in C no longer change the solution.

Overall, this pattern suggests that choosing a relatively large C is usually safe in this case study. The model does not open sites just because it is allowed to; it opens additional PODs only when they improve feasibility or accessibility.

Capacity buffer (`SCALAR` / κ)

The model becomes infeasible when κ drops below 1.5. This indicates that feasibility is sensitive to the area-based capacity proxy and to how much slack is built into it. Once the model is in the feasible range, increasing κ improves the demand-weighted accessibility and raises the objective value, while the number of opened PODs tends to decrease. In practical terms, looser capacity allows the model to serve demand with fewer sites and to place more load on the larger or better-positioned PODs.

Objective weights (w_{PD}, w_{SP})

Changing the objective weights leaves the demand-side accessibility almost unchanged across most settings. The clearer effect is on the opening decisions. As the ratio w_{SP}/w_{PD} increases, the model tends to open more PODs. This suggests that the supply-side term mainly influences which sites are opened, while the demand-side term largely preserves the assignment pattern once a feasible set of sites is available.

5.5.6 Demand Representation: Damage-Based vs Building-Stock-Based

Whether demand is defined from damaged buildings or from the entire building stock has a clear effect on the solution structure. Under the damage-based definition, demand points cluster around the city center, where destruction is more concentrated. In this setting, the PODs opened by the model tend to lie close to these core demand areas. A smaller number of sites are opened in outer parts of the district, mainly to balance load and relieve capacity pressure on the central PODs.

When demand is defined from the full building stock, the pattern changes. Demand points are no longer concentrated only in heavily damaged areas; instead, they are distributed more evenly across the district. As demand becomes spatially broader, the solution also becomes more distributed. The model opens more PODs, and the selected sites spread over a wider area rather than clustering around the center. The accessibility structure changes accordingly, since the model now needs to provide reasonable access across a larger geography instead of serving a limited number of central demand clusters efficiently.

In the settlement-based scenario, the need to increase κ to achieve feasibility is consistent with this wider service area. As demand becomes more dispersed, concentrating capacity in only a few large sites becomes harder, and the model requires more flexibility in the capacity scaling. In addition, the fact that one demand point could not be connected to the road network highlights a practical issue. Connectivity and snapping tolerance are not merely technical details; if a route cannot be computed for a demand point, the corresponding cost entries cannot be generated, and the optimization model cannot be fully constructed.

These two demand definitions also reflect two different ways of thinking about the operation. With damage-based demand, the model follows an early-response logic: it assumes that need is more concentrated where destruction is heavier. Using the full building stock shifts the focus to coverage. The aim becomes providing service across the whole district, not only in the hardest-hit zones, which leads to a more citywide planning view.

5.5.7 How the Results Address the Research Questions

The results speak directly to the research questions. For the first question, the study shows how damage information from satellite imagery can be carried into the road-cost layer through buffers and the exposure count $damage_n$, and then used in network-based accessibility calculations. For the second question, the scenario comparisons show that adding risk-aware costs does not change the overall pattern, but it can shift some routes and a few assignments in specific areas. The effect also depends on what cost definition is used. In particular, the road network topology has a stronger impact, as the network versus Euclidean comparison shows. The results also indicate that the solution depends on key assumptions, especially capacity scaling and how demand is defined. It should also be noted that the model does not impose explicit equity constraints; the focus remains on accessibility maximization under different network and cost assumptions. For the third question, the case study confirms that the proposed location–allocation model can be solved to optimality at the instance size considered here and that parameter sweeps are practical. The experiments also show that feasibility is sensitive to the capacity scaling, which makes κ a key practical parameter.

5.5.8 Practical Implications and Limitations That Shape Interpretation

The proposed approach takes a more direct, data-driven route. Instead of building many hypothetical disruption scenarios, it uses the damage that can be observed from imagery and feeds it straight into the road-cost layer. This fits the early response setting, where decisions have to be made fast and detailed field information may not be available yet.

Still, the results need to be interpreted with a bit of care, because the pipeline relies on inputs that are not perfectly certain. The damage labels come from imagery and can include misclassifications, POD capacity is represented through an area-based proxy, and key parameters such as α, r, and κ influence the strength of the cost and capacity effects. The framework also does not aim to update routes in real time. It produces a deterministic planning snapshot for a given moment, based on the available post-disaster observations. For this reason, the outputs are best interpreted as rapid decision support rather than as a fully dynamic operational plan.

5.6 Chapter Summary

In this section, the proposed end-to-end workflow was evaluated computationally on the Kahramanmaraş case study. Damage-assessment outputs were mapped to road costs through a buffer-based exposure measure. Using these costs, accessibility values were computed, and the integrated POD location and demand assignment model was solved. Under the main scenario, all demand points could be assigned while respecting capacity constraints. Scenario comparisons suggest that, on the same road network, the damage-based penalty leads to only limited and mostly local changes in the decisions. By contrast, replacing network-based costs with Euclidean distances, which ignore road-network topology, alters the solution structure more noticeably. Finally, the parameter scans indicate a clear saturation pattern with respect to C, show that feasibility is particularly sensitive to κ, and suggest that the objective weights mainly influence how many PODs are opened.

6 Conclusion

This thesis proposes a framework that uses building damage information from satellite imagery to support POD site selection and demand assignment in the early response period. The idea is to take the spatial damage signal and reflect it in the transportation network by adding a risk-aware cost layer. Accessibility is then computed on this updated network and used as input to the location–allocation model. In this way, the observed damage pattern is not only shown on maps but is also connected to the decision model.

The Kahramanmaraş case study shows that the workflow can be implemented and that the model can produce feasible solutions at the scale considered in this thesis. The results suggest that risk-aware costs do not rewrite the solution everywhere, but they can affect decisions in some local areas. At the same time, the structure of the road network itself plays a stronger role in shaping accessibility and assignments. The parameter analyses also show that capacity assumptions and model settings can change feasibility and the overall solution pattern.

In terms of the research questions, the study demonstrates that satellite-based damage information can be carried into network costs, and that these costs can be used in accessibility calculations. The scenario comparisons also show that adding risk-aware costs can change some routing and assignment decisions, although the overall solution pattern is influenced more strongly by the road network structure and the main modeling assumptions. Finally, the resulting location–allocation model can be solved optimally for the instance considered in this case. Overall, the thesis offers a practical way to use damage seen in satellite imagery for planning when quick decisions are needed and detailed field data is still limited.

6.1 Future Work

This method has several clear directions for future work. One step is to treat damage assessment as uncertain information rather than a single fixed layer. For example, the classifier's confidence scores could be carried into the exposure calculation, or several plausible damage maps could be generated and used to see how sensitive the plan is to classification errors.

A second direction is to study the risk-cost design more systematically. The buffer radius and the risk-penalty strength work as practical knobs in the current setup, but their effects can be explored in a more structured way, and tested under different road layouts and different damage patterns to see when they matter most.

Third, capacity is currently represented through a simple area-based proxy. A more detailed representation could account for operational limits such as staffing, setup constraints, or throughput, or it could treat capacity as uncertain and handle it with robust or stochastic variants of the model.

Another possible extension is to incorporate explicit equity considerations into the model. While this study focuses on accessibility maximization, future work could introduce minimum service guarantees, fairness-based allocation rules, or multi-objective formulations that balance accessibility and equity more directly.

Finally, the pipeline should be tested in other cities to understand how transferable it is. Extending the framework beyond a single planning snapshot is also a natural next step, for instance by adding a time dimension and solving a multi-period version where conditions and inputs can be updated as new information arrives.

Abbreviations

POD Point of Distribution

CRS Coordinate Reference System

POI Points of Interest

POFW Points of Worship

OSM OpenStreetMap

GIS Geographic Information System

AUC Area Under the ROC Curve

ROC Receiver Operating Characteristic

PR Precision–Recall

CNN Convolutional Neural Network

WMTS Web Map Tile Service

EPSG European Petroleum Survey Group

CSV Comma-Separated Values

MIP Mixed-Integer Programming

API Application Programming Interface

List of Figures

List of Tables

Bibliography

[ACM16] D. Alem, A. Clark, and A. Moreno. “Stochastic network models for logistics planning in disaster relief.” In: *European Journal of Operational Research* 255.1 (2016), pp. 187–206.

[AFA23] AFAD. *06 Şubat 2023 Pazarcık-Elbistan (Kahramanmaraş) Mw: 7.7 – Mw: 7.6 Depremleri Raporu*. 140 s. Ankara: T.C. İçişleri Bakanlığı Afet ve Acil Durum Yönetimi Başkanlığı, 2023.

[Als23] L. Alsaafin. *Turkey and Syria face long road to recovery after earthquakes*. Al Jazeera. Feb. 16, 2023. URL: `https://www.aljazeera.com/news/2023/2/16/what-does-the-recovery-process-for-turkey-and-syria-look-like` (visited on 02/15/2024).

[ALZ+24] I. Alisjahbana, J. Li, Y. Zhang, et al. “Deepdamagenet: A two-step deep-learning model for multi-disaster building damage segmentation and classification using satellite imagery.” In: *arXiv preprint arXiv:2405.04800* (2024).

[APA05] S. Argyroudis, K. Pitilakis, and A. Anastasiadis. “Roadway network seismic risk analysis in urban areas: The case of Thessaloniki-Greece.” In: *Proceedings of the Geoline Conference, Lyon, France*. 2005.

[App+08] P. Apparicio, M. Abdelmajid, M. Riva, and R. Shearmur. “Comparing alternative approaches to measuring the geographical accessibility of urban health services: Distance types and aggregation-error issues.” In: *International journal of health geographics* 7.1 (2008), p. 7.

[BA04] G. Barbarosoğlu and Y. Arda. “A two-stage stochastic programming framework for transportation planning in disaster response.” In: *Journal of the operational research society* 55.1 (2004), pp. 43–53.

[Bal+10] B. Balcik, B. M. Beamon, C. C. Krejci, K. M. Muramatsu, and M. Ramirez. “Coordination in humanitarian relief chains: Practices, challenges and opportunities.” In: *International Journal of production economics* 126.1 (2010), pp. 22–34.

[BB08] B. Balcik and B. Beamon. “Facility location in humanitarian relief. International Journal of Logistics-research and Applications.” In: *Int J Logist-Res Appl* (2008).

[Bha+25] D. Bhardwaj, N. Nagabhooshanam, A. Singh, B. Selvalakshmi, S. Angadi, S. Shargunam, T. Guha, G. Singh, and A. Rajaram. "Enhanced satellite imagery analysis for post-disaster building damage assessment using integrated ResNet-U-Net model." In: *Multimedia Tools and Applications* 84.5 (2025), pp. 2689–2714.

[Buc] J. Buchner. *ImageHash: Perceptual image hashing for Python*.

[CNP12] A. M. Caunhye, X. Nie, and S. Pokharel. "Optimization models in emergency logistics: A literature review." In: *Socio-economic planning sciences* 46.1 (2012), pp. 4–13.

[Cos+20] C. Costa, R. Figueiredo, V. Silva, and P. Bazzurro. "Application of open tools and datasets to probabilistic modeling of road traffic disruptions due to earthquake damage." In: *Earthquake Engineering & Structural Dynamics* 49.12 (2020), pp. 1236–1255.

[Dij59] E. W. Dijkstra. "A Note on Two Problems in Connexion with Graphs." In: *Numerische Mathematik* 1.1 (1959), pp. 269–271. DOI: `10.1007/BF01386390`.

[Eki21] B. B. Ekici. "Detecting damaged buildings from satellite imagery." In: *Journal of Applied Remote Sensing* 15.3 (2021), pp. 032004–032004.

[Esr23] Esri. *World Imagery Wayback*. `https://www.arcgis.com/home/item.html?id=dc9c2d7b1c4a4b48b6d61d1d61e9a6c4`. Accessed via QGIS WMTS service. 2023.

[Geo25] Geofabrik GmbH. *Geofabrik Download Server*. Accessed: 2025-09-25. 2025.

[GF18] W. J. Gutjahr and S. Fischer. "Equity and deprivation costs in humanitarian logistics." In: *European Journal of Operational Research* 270.1 (2018), pp. 185–197.

[GS06] A. Goretti and V. Sarli. "Road network and damaged buildings in urban areas: short and long-term interaction." In: *Bulletin of Earthquake Engineering* 4.2 (2006), pp. 159–175.

[Gur24] Gurobi Optimization, LLC. *Gurobi Optimizer Reference Manual*. 2024.

[Han59] W. G. Hansen. "How accessibility shapes land use." In: *Journal of the American Institute of planners* 25.2 (1959), pp. 73–76.

[He+16] K. He, X. Zhang, S. Ren, and J. Sun. "Deep Residual Learning for Image Recognition." In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2016, pp. 770–778. DOI: `10.1109/CVPR.2016.90`.

[Hol+13] J. Holguín-Veras, N. Pérez, M. Jaller, L. N. Van Wassenhove, and F. Aros-Vera. "On the appropriate objective function for post-disaster humanitarian logistics models." In: *Journal of Operations Management* 31.5 (2013), pp. 262–280.

[IS15] C. A. Irawan and S. Salhi. "Solving large p-median problems by a multistage hybrid approach using demand points aggregation and variable neighbourhood search." In: *Journal of Global Optimization* 63.3 (2015), pp. 537–554.

[KS19] I. Kilanitis and A. Sextos. "Impact of earthquake-induced bridge damage and time evolving traffic demand on the road network resilience." In: *Journal of traffic and transportation engineering (English edition)* 6.1 (2019), pp. 35–48.

[LA18] N. Loree and F. Aros-Vera. "Points of distribution location and inventory management model for Post-Disaster Humanitarian Logistics." In: *Transportation Research Part E: Logistics and Transportation Review* 116 (2018), pp. 1–24.

[Lib+14] F. Liberatore, M. T. Ortuño, G. Tirado, B. Vitoriano, and M. P. Scaparra. "A hierarchical compromise model for the joint optimization of recovery operations and distribution of emergency goods in Humanitarian Logistics." In: *Computers & Operations Research* 42 (2014), pp. 3–13.

[Mac67] J. MacQueen. "Some Methods for Classification and Analysis of Multivariate Observations." In: *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability*. Vol. 1. 1967, pp. 281–297.

[Mic23] R. Michaelson. *'It seems too late': hope fading in Turkey's search for earthquake survivors*. The Guardian. Feb. 8, 2023. URL: `https://www.theguardian.com/world/2023/feb/08/it-seems-too-late-hope-fading-in-turkey-search-for-earthquake-survivors` (visited on 02/15/2024).

[NBA16] N. Noyan, B. Balcik, and S. Atakan. "A stochastic optimization model for designing last mile relief networks." In: *Transportation Science* 50.3 (2016), pp. 1092–1113.

[Nex+19] F. Nex, D. Duarte, F. G. Tonolo, and N. Kerle. "Structural building damage detection with deep learning: Assessment of a state-of-the-art CNN in operational conditions." In: *Remote sensing* 11.23 (2019), p. 2765.

[OEK04] L. Ozdamar, E. Ekinci, and B. Kiiciikyazici. "Emergency logistics planning in natural disasters." In: *Annals of operations research* 129.1-4 (2004), pp. 217–245.

[Ope25] OpenStreetMap contributors. *OpenStreetMap*. Accessed: 2025-09-25. 2025.

[Ped+11] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, et al. “Scikit-learn: Machine Learning in Python.” In: *Journal of Machine Learning Research* 12 (2011), pp. 2825–2830.

[QGI25] QGIS Development Team. *QGIS Geographic Information System*. Accessed: 2025-09-25. 2025.

[Sak+16] C. S. Sakuraba, A. C. Santos, C. Prins, L. Bouillot, A. Durand, and B. Allenbach. “Road network emergency accessibility planning after a major earthquake.” In: *EURO Journal on Computational Optimization* 4.3 (2016), pp. 381–402.

[SMG25] J. Shefa, A. B. Milburn, and E. Gralla. “Online path planning for disaster response in the presence of unknown road disruptions: a comparison of research and practice.” In: *Journal of Humanitarian Logistics and Supply Chain Management* (2025).

[SY15] F. S. Salman and E. Yücel. “Emergency facility location under random network damage: Insights from the Istanbul case.” In: *Computers & Operations Research* 62 (2015), pp. 266–281.

[Tom18] D. Toma-Danila. “A GIS framework for evaluating the implications of urban road network failure due to earthquakes: Bucharest (Romania) case study.” In: *Natural Hazards* 93.Suppl 1 (2018), pp. 97–111.

[Van06] L. N. Van Wassenhove. “Humanitarian aid logistics: supply chain management in high gear.” In: *Journal of the Operational research Society* 57.5 (2006), pp. 475–489.

[Wan+25] Z. Wang, C. Wu, F. Zhang, and J. Xia. “Constructing an extensible building damage dataset via semi-supervised fine-tuning across 12 natural disasters.” In: *Journal of Remote Sensing* 5 (2025), p. 0733.

[Xu+19] J. Z. Xu, W. Lu, Z. Li, P. Khaitan, and V. Zaytseva. “Building damage detection in satellite imagery using convolutional neural networks.” In: *arXiv preprint arXiv:1910.06444* (2019).

[YAE19] Z. Yilmaz, A. Aydemir-Karadag, and S. Erol. “Finding optimal depots and routes in sudden-onset disasters: An earthquake case for Erzincan.” In: *Transportation journal* 58.3 (2019), pp. 168–196.

[YS20] Z. D. Yenice and F. Samanlioglu. “A Multi-Objective Stochastic Model for an Earthquake Relief Network.” In: *Journal of Advanced Transportation* 2020.1 (2020), p. 1910632.

Ehrenwörtliche Erklärung

Ich erkläre hiermit ehrenwörtlich, dass ich die vorliegende Arbeit selbständig angefertigt habe. Die aus fremden Quellen direkt und indirekt übernommenen Gedanken sind als solche kenntlich gemacht.

Ich weiß, dass die Arbeit in digitalisierter Form daraufhin überprüft werden kann, ob uner- laubte Hilfsmittel verwendet wurden und ob es sich - insgesamt oder in Teilen - um ein Plagiat handelt. Zum Vergleich meiner Arbeit mit existierenden Quellen darf sie in eine Datenbank eingestellt werden und nach der Überprüfung zum Vergleich mit künftig eingehenden Ar- beiten dort verbleiben. Weitere Vervielfältigungs- und Verwertungsrechte werden dadurch nicht eingeräumt. Die Arbeit wurde weder einer anderen Prüfungsbehörde vorgelegt noch veröffentlicht.

Ort, Datum

Munich, 23.02.2026

Unterschrift